Pillar Axiom

Administrator's Guide

ORACLE

PILLAR AXIOM

Part Number: 4420-00028-1500
Pillar Axiom release 4.1
2011 October

Table of Contents

List of Figures

List of Tables

Preface

Audience

This documentation is intended for field engineers, sales engineers, service technicians, system administrators, and other individuals who need to know about the function and use of the Pillar Axiom Storage Services Manager.

You should have, as appropriate, the necessary skills and experience in administering and working with:

- Computer hardware.
- Graphical user interfaces.
- Configuration of network attached storage (NAS) or storage area network (SAN) environments.
- Various network protocols, such as Common Internet File System (CIFS), Network File System (NFS), Fibre Channel, Internet SCSI (Small Computer System Interface), Network Time Protocol (NTP), and Simple Network Management Protocol (SNMP).
- Network services, such as Domain Name System (DNS), Internet Storage Name Service (iSNS), and Network Information Service (NIS).
- Backup and restore operations.
- Network connectivity trouble resolution tools, such as `ping`, `traceroute`, `arp`, `netstat`, and others.

You should also have a basic understanding and working knowledge of Oracle's Pillar Axiom storage systems.

Before You Read This Guide

Being familiar with certain other technical documentation for Oracle's Pillar Axiom 600 storage system helps you succeed in the use of this guide.

In addition to this guide, review the late-breaking information described in the Pillar Axiom *Customer Release Notes*. That information includes important information that was not available at the time this guide was published, including:

- Errata for technical documents (including this guide).
- Network requirements.
- Known issues.
- Various notations on the operation of the Pillar Axiom storage system.

How This Guide Is Organized

This guide provides procedural and reference information for configuring and managing an Oracle Pillar Axiom 600 storage system.

The guide is divided into nine chapters and an appendix:

- Chapter 1 provides an introduction to the Pillar Axiom Storage Services Manager, which is also referred to as the *graphical user interface (GUI)*.
- Chapter 2 provides instructional information about:
  - The Configuration Wizard.
  - Setting global parameters for the entire system.
  - The capacity and provisioning of logical volumes.
  - Performance profiles.
  - RAID striping.
  - The Pillar Axiom Capacity Planner.
  - Using the Configuration Wizard to create logical volumes (filesystems and LUNs) and File Servers.
  - The Common Internet File System (CIFS) and Network File System (NFS) protocols.
- Chapter 3 provides instructional information about manually managing logical volumes, File Servers, and volume groups.
- Chapter 4 provides instructional information about managing and otherwise working with backups, clones, and snapshots. This chapter also describes the Pillar Axiom VSS Provider plug-in.
- Chapter 5 provides instructional information about managing system alerts and the SNMP feature.

- Chapter 6 provides instructional information about updating the system software and replacing hardware components.
- Chapter 7 provides instructional information about managing administrator accounts.
- Chapter 8 provides instructional information about using the support tools to perform a variety maintenance operations.
- Chapter 9 provides instructional information about viewing the event log and various performance statistics.
- Appendix A provides full reference information for every GUI page you might need to use.

Related Documentation

Table 1 Additional information resources for all systems

Description	Title and part number
The definitions of terms found in Pillar Axiom documentation.	*Pillar Glossary*
An introduction to the hardware and software architecture of a Pillar Axiom storage system.	*Pillar Axiom System Architecture Overview*
A reference to the syntax and use of all `pdscli` commands.	*Pillar Axiom CLI Reference Guide*
A reference to the syntax and use of all `axiomcli` commands.	*Pillar Axiom CLI Guide*
A reference to the SMI-S profiles supported on all Pillar Axiom systems, including how to set the network connection properties.	*Pillar Axiom SMIProvider Reference*
Notations on implementing link aggregation.	*Pillar Axiom Implementation Tips for Link Aggregation*
Instructions for creating formatted statistical data on many facets of system operation.	*Pillar Axiom Statistics Tools User Guide*

Table 1 Additional information resources for all systems (continued)

Description	Title and part number
Instructions for installing hardware components into Pillar and non-Pillar racks and for expanding these systems by adding Bricks and Slammers.	*Pillar Axiom Hardware Installation Guide* for these platforms: Pillar Axiom 300 Pillar Axiom 500 Pillar Axiom 600
Removal and insertion instructions for field replaceable units (FRUs).	*Pillar Axiom Service Guide* for these platforms: Pillar Axiom 300 Pillar Axiom 500 Pillar Axiom 600

Table 2 Additional information resources for SAN systems

Description	Title and part number
Instructions and other supporting information for using the AxiomONE Replication for SAN utility.	*Pillar Axiom Replication User's Guide and Reference for SAN*
Instructions for integrating a SAN Pillar Axiom system into an iSCSI environment (Windows).	*Pillar Axiom iSCSI Integration Guide for Windows Platforms*
Instructions for integrating a SAN Pillar Axiom system into an Oracle 10g RAC environment.	*Pillar Axiom Oracle Integration Guide*

Table 3 Additional information resources for NAS systems

Description	Title and part number
Instructions and other supporting information for using the Pillar Axiom MaxRep Replication for NAS utility.	*Pillar Axiom MaxRep Replication for NAS User's Guide and Reference*
Instructions for integrating a Pillar Axiom system into a Windows environment.	*Pillar Axiom Windows Integration Guide for NAS Systems*
Instructions for integrating a NAS Pillar Axiom system into an NDMP environment.	*Pillar Axiom NDMP Integration Guide for NAS Systems*

Table 3 Additional information resources for NAS systems (continued)

Description	Title and part number
Concepts for multi-protocol file access for UNIX and Windows network clients.	*Pillar Axiom CIFS and NFS Multi-Protocol Planning Guide*
Instructions for listing files in a filesystem that have been changed between two Snap FS snapshots.	*Pillar Axiom SnapDelta FS Reference Guide*

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.

Typographical Conventions

Table 4 Typography to mark certain content

Convention	Meaning
italics	Within normal text, words in italics indicate: • A reference to a book title. • New terms and emphasized words. • Command variables.
`monospace`	Indicates one of the following, depending on the context: • The name of a file or the path to the file. • *Output* displayed by the system on the command line.
`monospace` (bold)	*Input* provided by an administrator on the command line.
>	Indicates a menu item or a navigation path in a graphical user interface (GUI). For example, “Click **Storage** > **Clone LUNs**” means to click the **Clone LUNs** link on the **Storage** page in the graphical user interface (GUI).
...	Used within an expression of a navigation path or within a cascading menu structure. The ellipsis indicates that one or more steps have been omitted from the path or menu structure. For example, in the **Groups** > **Volume Groups** > **Actions** > **...** > **Data Protection** > **Create** menu structure, the **...** implies that one or more menu items have been omitted.

Command Syntax Conventions

Table 5 Typography to mark command syntax

Typographic symbol	Meaning
`[ ]`	Square brackets. Delimits an optional command parameter or a set of optional command parameters.
`{ }`	Curly braces. Delimits a set of command parameters, one of which must be selected.

Table 5 Typography to mark command syntax (continued)

Typographic symbol	Meaning
|	Vertical bar. Separates mutually exclusive parameters.
. . .	Ellipsis. Indicates that the immediately preceding parameter or group of parameters can be repeated.
`monospace`	Indicates the name of a command or the name of a command option (sometimes called a *flag* or *switch*).
italic	Indicates a variable for which you need to supply a value.

Command parameters that are *not* enclosed within square brackets ([ ]) are required.

Important! The above symbols (and font styling) are based on the IEEE Std 1003.1-2004 standard. These symbols are used in the command syntax only to clarify how to use the command parameters. Do not enter these symbols on the command line.

Pillar Contacts

Table 6 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/)

Table 6 Contacts at Pillar Data Systems (continued)

For help with...	Contact...
Professional services and inquiries	globalsolutions@pillardata.com 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 Pillar Axiom Administration

Administrator Accounts and Privileges

Administrator accounts have certain privileges, which depend on the role of the account.

To administer a Pillar Axiom storage system, you must log in from an administrator account. Every account is assigned a specific role that defines system privileges.

Table 7 Administrator privileges by role

Administrator role	Privileges
Primary System Administrator	Performs all configuration, management, and monitoring tasks. This account cannot be deleted or disabled.
Administrator 1	Performs all configuration, management, and monitoring tasks.
Administrator 2	Performs all tasks except: • Create and manage File Servers and administrator accounts. • Modify global, Small Network Management (SNMP), and Network Data Management Protocol (NDMP) settings. • Modify software or hardware configurations. • Shut down the system.
Monitor	Displays system information only; cannot modify the configuration. Can modify own administrator account attributes.
Support	Performs limited customer service-only functions; cannot modify the configuration. **Note:** Only Pillar Data Systems customer service personnel can use this account.

About Administrator Accounts Management

Administrators have specific privileges in the Pillar Axiom storage system based on their account type.

You can create multiple administrator accounts in a Pillar Axiom system. Additional accounts are not necessary, but they are useful if you want to delegate administrator responsibilities. For example, you might choose to create:

- One administrator account. In this way, a designated person can assume responsibility while the Primary system administrator is on vacation. Assign this account to the Administrator 1 role.

  Tip: Pillar strongly recommends that you set up a Type 1 Administrator account when you install the system. Besides the Primary system administrator, only a Type 1 Administrator can modify an account password (including that of the Primary system administrator) without knowing the previous password.

- One or more administrator accounts with read-only privileges. In this way, managers can monitor the system but they cannot change configuration details. Assign these accounts to the Monitor role.

You can create up to 23 administrator accounts.

If you delegated administrative tasks to other administrators, you may need to:

- Modify account attributes (for example, change an administrator's password or disable an account other than the Primary system administrator account).
- Change administrator account security settings.
- Delete obsolete accounts.

At times, you may need to modify the attributes of an administrator account. A Primary system administrator and people who are assigned to the Administrator 1 role can modify their own or another administrator's account.

Some changes take effect immediately. For example, a logged-in administrator's session is terminated when you disable or delete the administrator account.

Other changes affect the administrators the next time that they log in, for example, when you modify the administrator's password or modify the session timeout value.

You can change the security settings for system administrator accounts, including:

- Set the number of failed login attempts that the Pillar Axiom system permits. When the threshold is exceeded, the system disables the account and

writes an entry in the event log. Only a Primary system administrator or Administrator 1 account can re-enable the account, and the system resets the counter upon a successful login. If you do not set this value, there is no limit to the number of unsuccessful login attempts.

- Set the session timeout so that the Pillar Axiom system terminates an administrator's session after a given period of inactivity. If you do not set this value, inactive sessions are terminated after 20 minutes.
- Select Secure Session Only to specify that administrator access to the Pillar Axiom system is over secure HTTP sessions. Upload a secure sockets layer (SSL) certificate to the Pillar Axiom Pilot to authenticate logins.

About Accessing the Oracle Pillar Axiom System

To perform administrator tasks, you must log in to the Pillar Axiom system. At any given time, the following number of administrator sessions can be active:

- 5 active sessions for each administrator account
- 10 total at any given time (2 reserved for Primary System Administrator and Support Administrator)

Note: There is a timeout period after which inactive administrator sessions are terminated by the system.

See also: Account Security Settings Page.

Supported Browsers

The Pillar Axiom system supports a variety of browsers for administering the system.

Table 8 Platform and browser support

Platform	Version	Internet Explorer 6.0 or higher	Netscape Navigator 6.2 or higher	Mozilla 1.7	Mozilla Firefox 2.0 or higher
Redhat Linux	8, 9	N/A	X	X	X
Red Hat Enterprise Linux 3	AS/ES Updates 4 and 5	N/A	X	X	X
Solaris	8, 9, and 10	N/A	X	X	N/A
Windows	NT 4.0, 2000, 2003, XP, and Vista	X	X	X	X

If Microsoft Windows XP is installed on your workstation, optionally set the Desktop appearance to something other than the XP default theme. If you use

the XP default theme, lines for group boxes are invisible and the group-box title is blue in the GUI. Also, disabled options, such as radio buttons, may not be visible. The Windows Classic theme displays visible lines for group boxes. Visible lines indicate related fields more explicitly.

You can download browsers from the following vendors:

Table 9 Browser vendors

Vendor	URL
Microsoft®	www.microsoft.com/
Netscape Communications Corporation	browser.netscape.com/
The Mozilla Organization	www.mozilla.org/

Set the following browser options:

- Ensure that the Preferred Language setting in your browser is set to U.S. English (en-US).
- Set the security level in your browser to no higher than medium-low to accommodate the Pillar Axiom Storage Services Manager's security certificate.

  Note: You cannot log in if your security level is set to a higher level.
- Set the Text Size to Smaller to correct display issues when using the full-screen display.

Optionally enable the following browser features:

- Image support, so that you can see the graphics in the GUI pages. Although you can perform tasks using only the text elements, the graphics provide contextual detail.
- JavaScript, to dim unavailable selections. It is also used in other ways.

Log In to the Pillar Axiom Storage Services Manager GUI

The first time that you use the Primary administrator account to log in to the GUI, use the following default values:

Table 10 Default login values

Field	Default value
Pilot IP address	10.0.0.2
Login Name	administrator
Password	pillar

Tip: If the Pilot IP has not been changed to a customer-specified address, use 10.0.0.2, which was set at the factory. If that IP address is not successful, try 10.0.0.3 or 10.0.0.4, which are the addresses of the Pilot control unit (CU) 0 and CU 1, respectively. If you still cannot log in, ping those addresses and contact the Pillar World Wide Customer Support Center.

1. Start the browser software on your workstation.
2. Specify the IP address of the Pilot management controller or the name of the Pillar Axiom storage system as the address to open.

   The name you enter here is the host name as configured in your site naming services for the Pilot. This name might not be the same as the name entered in the Pillar Axiom GUI.

   Note: System names, by default, are the same as the system serial number. System serial numbers consist of a string of 10 alphanumeric characters. Serial numbers begin with the letter *A* (which designates a Pillar Axiom system), followed by six decimal digits (which comprise a sequentially issued decimal value), and end with three uppercase letters (the last two of which represent a check sum for the preceding eight characters). For example, a system serial number might look like this: **A001539BOW**
3. When prompted, enter your login name and password.

Tip: Because the GUI uses popup windows, configure your browser to allow pop-ups. Many versions of Internet Explorer block pop-ups by default.

The first time that the Primary system administrator account is used to log into the GUI, the following dialog boxes appear:

- Change Password—Prompts you to change the password after the first login before you can proceed.

Tip: If you forget the Primary system administrator password, you can reset it in these ways:

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

- Configuration Wizard—Guides you through the setup and configuration of the system. You can run the Configuration Wizard now or at a later time.

Log Out of the Pillar Axiom Storage Services Manager GUI

When you have completed administrative tasks, log out. If you do not log out:

- An unauthorized user may gain access to the Pillar Axiom system from your workstation.
- One login session is tied up unnecessarily until your session is automatically logged out when the inactivity time limit is reached.

To log out, click **Log Out**, which appears in the top right corner of the GUI page.

The Log In page appears, indicating the log out was successful.

About System Information Provided by the GUI

In addition to context-sensitive help, the Pillar Axiom Storage Services Manager GUI provides key system information in a status bar at the bottom of the window.

Status Bar Information

A lot of information is available at the bottom of every page in the GUI, and additional details are often a click away.

The status bar includes information and buttons that permit you to view items such as system status and administrator accounts. You can also respond to Administrator Actions and cancel background processes when appropriate.

Figure 1 Pillar Axiom Storage Services Manager status bar

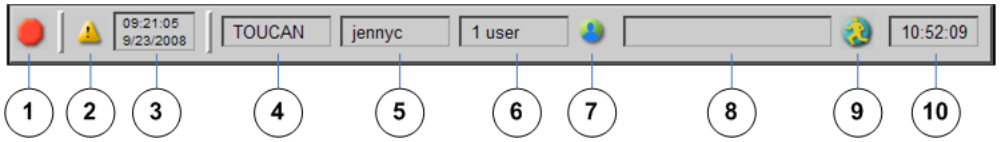

Legend		
	1 System status	6 Number of administrators
	2 Administrator action	7 Display all administrators
	3 Last administrator action	8 Current background process
	4 System name	9 Display background processes
	5 Current administrator	10 System time

Table 11 Status bar details

Icon/field	Description
System Status	Displays the overall system status of the hardware components. A status of Normal requires no action. If, however, the status is Warning or Critical, click the icon to view the Health Summary page to identify the cause of the status.

Table 11 Status bar details (continued)

Icon/field	Description
Administrator Action	Click the Administrator Action icon to open the Administrator Action Items page and respond to any events that require intervention. See About Responding to Administrator Actions.
Last Administrator Action	Displays the date and time of the last administrator action that occurred on the system. This lets you know, especially when there are multiple events, if a new administrator action has occurred since the last time you looked.
System Name	Displays the system name.
Current Administrator	Displays the name of the administrator account currently logged in to the system.
Number of Administrators	Displays the number of administrator accounts currently logged in to the system.
Display All Users	Click the Display All Users icon to open the Logged in Administrators page and view the specific administrator accounts that are currently logged in to the system.
Current Background Process	Displays the current running process or task.
Display Background Processes	Click the Display all Background Process icon to open the All Running Background Processes page. You can select a specific task that is running and, if applicable, cancel it.
System Time	Displays the current system time. To change the time, see Configure the Pillar Axiom System Time.

Red Instructional Text in the GUI

Red text in the GUI (see the image in the following figure) indicates an error and provides instructions on how to resolve the problem. In some instances, the red text is self-explanatory, for example when you make an incorrect selection in the

Actions menu. In other instances, you can mouse over the text and see more explanation as to why a task failed and how to resolve it.

Figure 2 Red instructional text

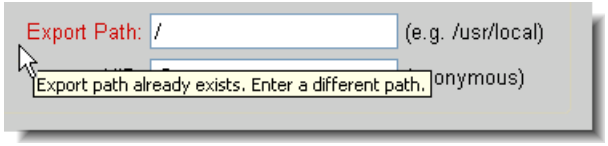

About Licensing Optional Premium Features

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

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

- Pillar Axiom SecureWORMfs - System Perpetual
- Pillar Axiom Storage Domains - System Perpetual
- Pillar Axiom Copy Services Bundle - System Perpetual
- Pillar Axiom MaxRep Replication for NAS - Terabyte Perpetual

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

- Pillar Axiom MaxRep Asynchronous Replication - Terabyte Perpetual
- Pillar Axiom MaxRep Aynchronous Replication with Application Protection - Terabyte Perpetual
- Pillar Axiom MaxRep Synchronous Replication - Terabyte Perpetual
- Pillar Axiom MaxRep Synchronous Replication with Application Protection - Terabyte Perpetual

CHAPTER 2

Configure a New System

About Initial Configuration

Your Pillar Axiom storage system was pre-loaded with management software and an administrative login for you to use. You must configure the system for your environment.

Choose either of the following methods:

- To guide you through the process, Run the Configuration Wizard.
- To manually configure the system using the GUI, begin with the task lists outlined in:
  - About Global Settings Configuration
  - About Storage Resources Creation

About the Configuration Wizard

Use the Configuration Wizard to guide you through the setup and configuration of the Pillar Axiom system. The setup includes the configuration of:

- **Global System Parameters**, such as:
  - System time, using a Network Time Protocol (NTP) server
  - System name
  - Notification and account security
  - Call-Home
- **Storage Allocation and Quality of Service (QoS)**, such as:
  - Performance characteristics
  - Storage capacity and performance
- **Administrator Accounts**

The Configuration Wizard automatically appears after the first login to the Pillar Axiom system. You can, however, run the Configuration Wizard anytime after the first login.

Run the Configuration Wizard

1. Click the **System** icon in the top context pane.
2. Click **Configuration Wizard** at the bottom of the left navigation pane.
3. Click **Next** in the Welcome to the Pillar Axiom Configuration Wizard page.
4. Select the items that you want to configure in the Configuration Tasks page and click **Next**.

   Note: At a minimum, select the following options to configure the Pillar Axiom storage system:

   - **Global System Parameters**
   - **Storage Parameters**

You're now able to configure a variety of system parameters (such as the time) and storage parameters (such as the Quality of Service (QoS)) for a logical volume.

Configure Global Settings

The next set of Configuration Wizard pages guides you through the manual steps to configure global system parameters.

1. In the System Time page, enter the IP addresses of the primary Network Time Protocol (NTP) server and up to two secondary NTP servers (optional).

   Note: If you do not set the system time, the Pillar Axiom system uses the time set on the Pilot.

2. Select the time zone that you want to define for the Pillar Axiom system and click **Next**.

Figure 3 Example NTP server settings

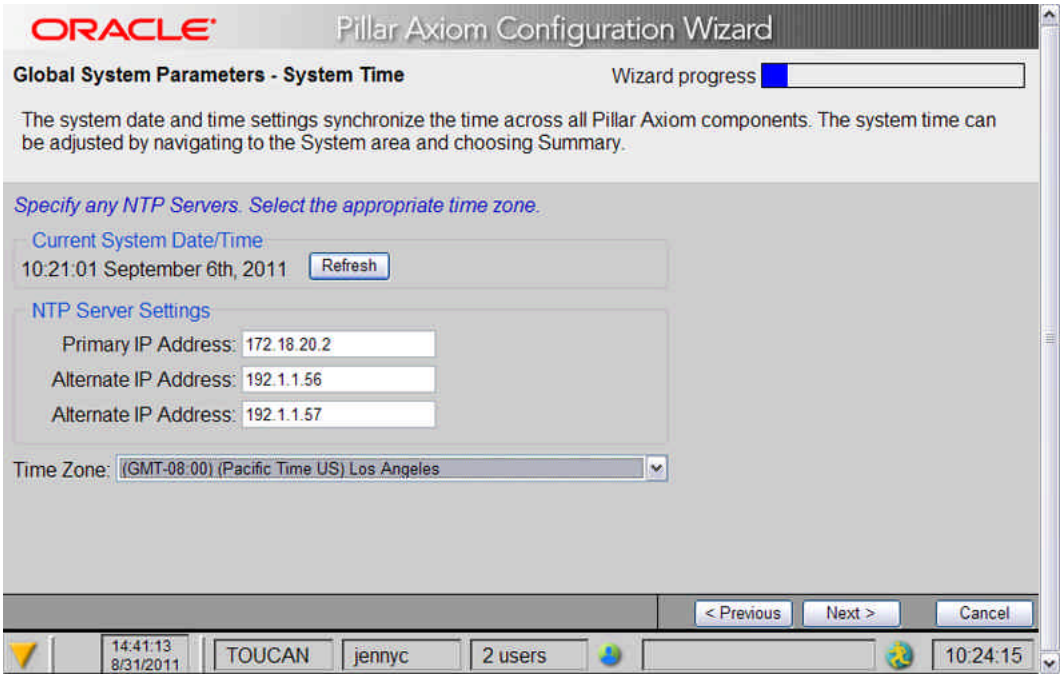

3 In the Management and Data Path Interfaces page, enter the system-wide values and click **Next**.

Figure 4 Example management and data path interfaces values

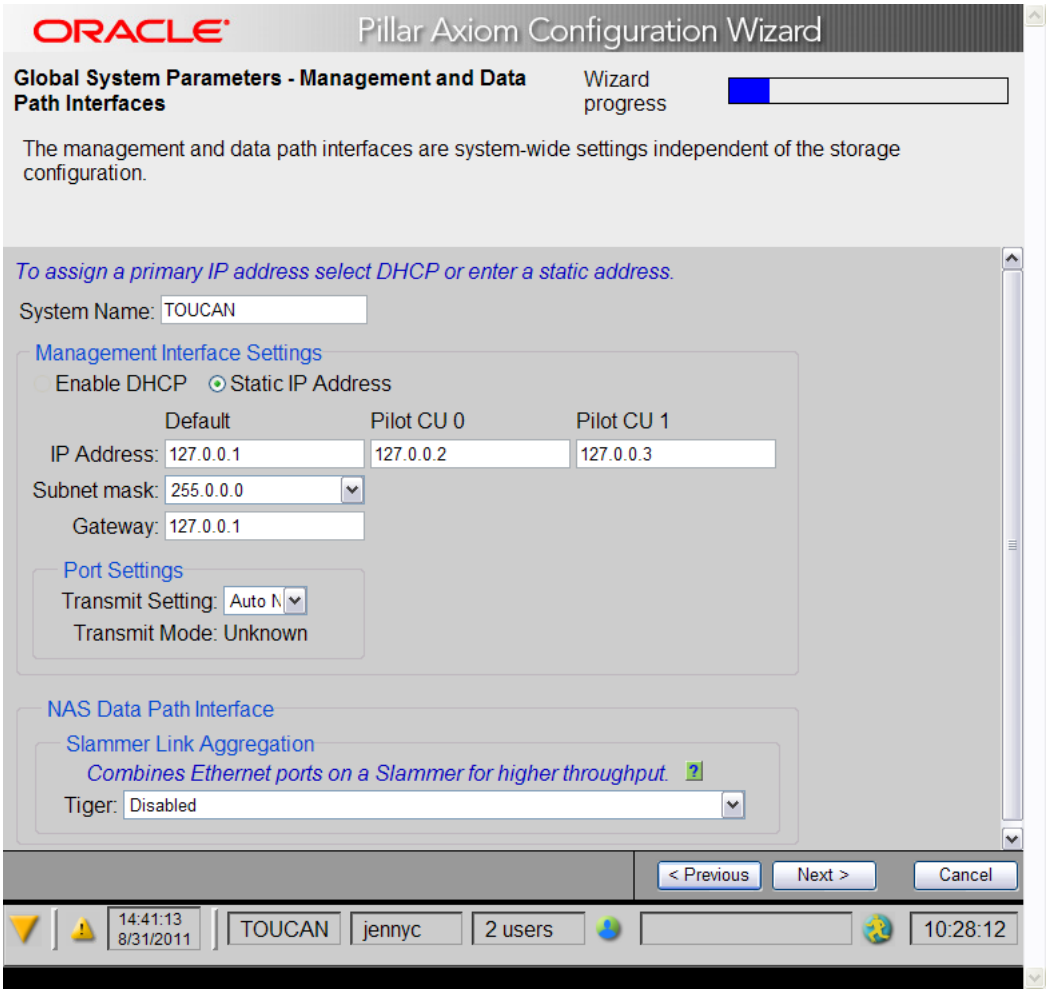

4 Enter the Notification and Account Security values and click **Next**.

Figure 5 Example notification and account security values

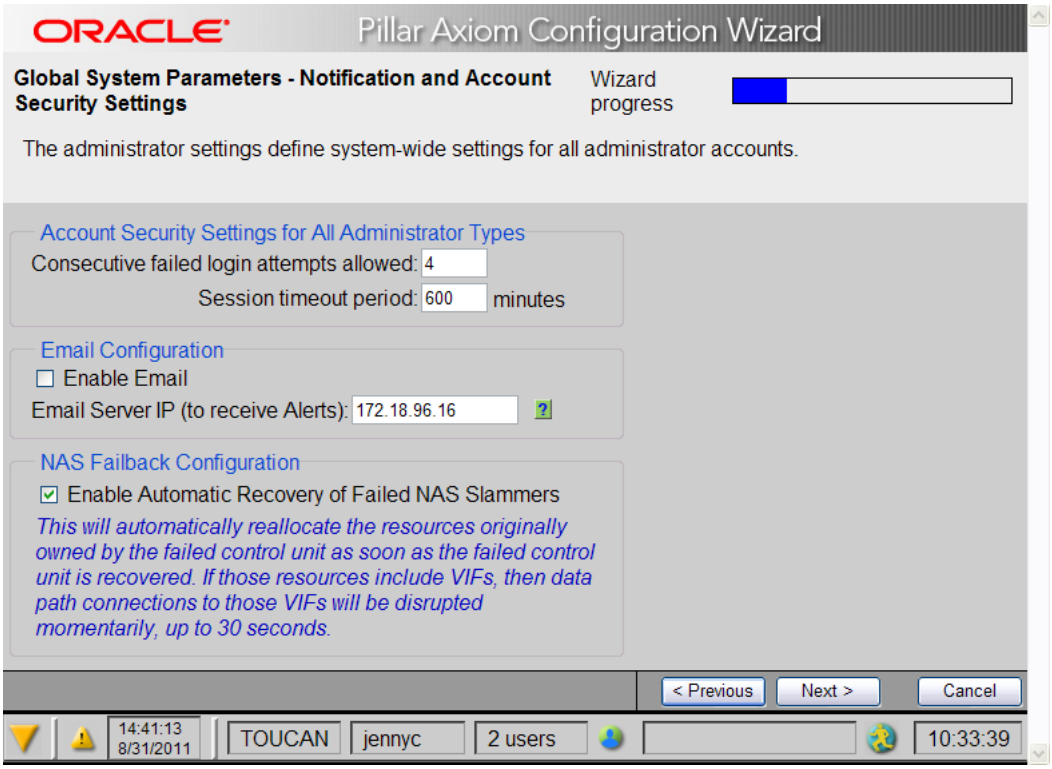

5 Enter values for the Call-Home Settings and click **Next**.

Figure 6 Example Call-Home settings

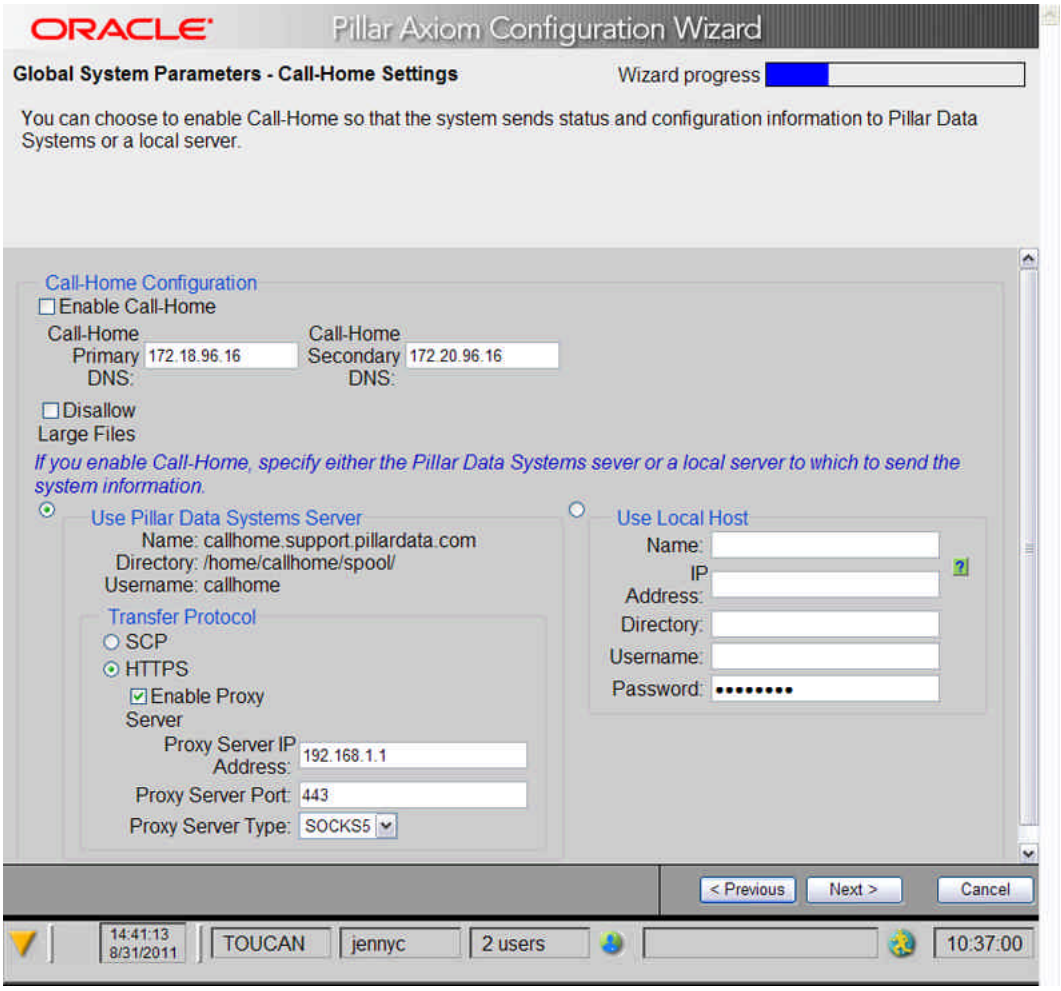

The next set of Configuration Wizard pages guide you through the manual steps to Configure NAS Storage Parameters. Respond to the Wizard's prompts to create a filesystem.

Configure NAS Storage Parameters

1 In the Storage Allocation and Quality of Service (QoS) page, select an option on which to base the QoS settings and click **Next**.

Choose one of:

- **Application Category**: Identifies the generic type of applications that best describes your storage use environment.
- **Performance Profiles**: Identifies the configuration to use for the most optimum throughput.

- **Highest Throughput**
- **Oracle Automatic Storage Management**

  For more information on performance profiles, see About Performance Profiles. For more information about using the Oracle Automatic Storage Management feature, see About Licensing Optional Premium Features.

Note: If others have created custom profiles, you can use those profiles as well. Click **Import Profile(s)** to make them available.

Figure 7 Example NAS storage allocation and QoS values

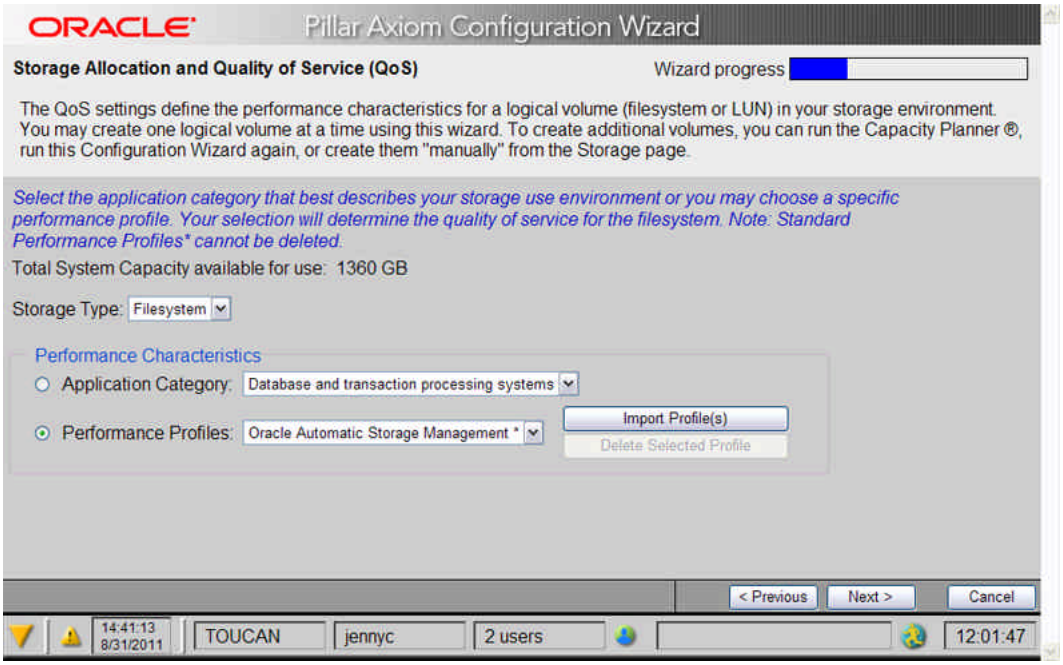

2 Enter the filesystem name as well as the capacity and performance values and click **Next**.

Tip: Some 32-bit operating systems cannot access or manage Network File System (NFS) or combined NFS and Common Internet File System (CIFS) filesystems of two terabytes or greater in size. If in doubt, check the operating system limits for your client and filesystem. For mixed 32-bit and 64-bit clients, limit filesystem sizes such that both clients will be able to work with those filesystems.

Figure 8 Example filesystem capacity and performance values

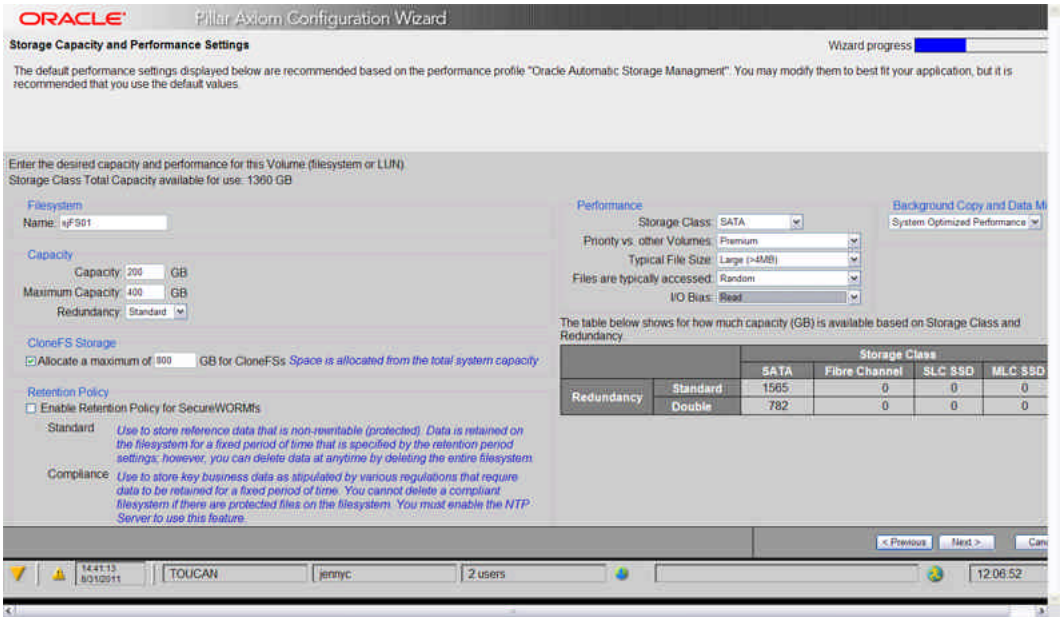

The Configuration Wizard displays recommended performance settings that are based on the application category that you selected in the Storage Allocation and Quality of Service (QoS) page. You can change any of the settings or use the recommended settings.

3 In the File Sharing page, select the type of file sharing that you want for this filesystem and click **Next**.

Choose one or both:

- Windows Share (CIFS): In the Shares page, enter the share name and path, and optionally a share comment; then click **Next**.
- UNIX Export (NFS): In the Exports page, enter the export path, user ID (UID) for anonymous users, and host definitions; then click **Next**.

Your next step is to create a File Server.

Create a File Server

End users connect to a File Server so that they can access the filesystems that are associated with that File Server.

1 In the Choose a File Server page, select **Create a New File Server** and click **Next**.

2 Enter the File Server networking values and click **Next**.

Figure 9 Example File Server values

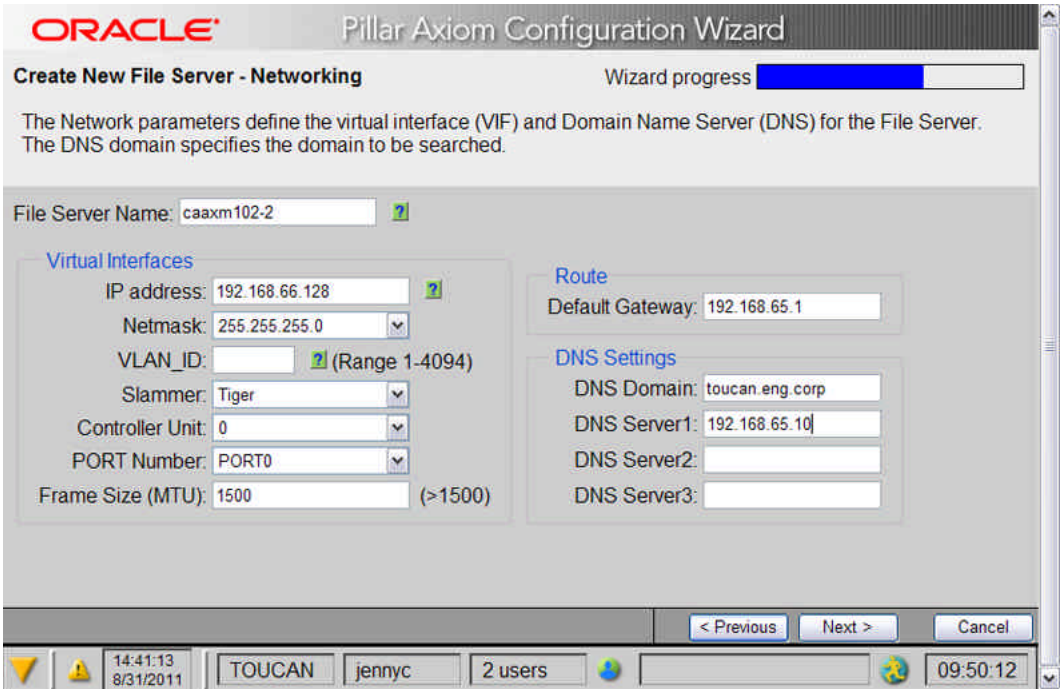

3 Enter the values in the NIS page for clients using the Network Information Service (NIS) and click **Next**.

4 Enter the values in the CIFS page for Windows (CIFS) file sharing and click **Next**.

5 Enter the values in the CIFS Authentication page for Common Internet File System (CIFS) file sharing and click **Next**.

6 Enter the values in the NFS Parameters page for Network File System (NFS) file sharing and click **Next**.

7 Enter the values in the Account Mapping page to map user names between CIFS and NFS domains and click **Next**.

If you are using a single CIFS security domain, enter the NetBIOS name of the CIFS domain that contains definitions for authentication and account mapping.

8 In the Configuration Summary page, review your configuration selections. If you want to change any of the settings, click **Previous**.

9 Click **Finish** to implement the configuration.

Result:
The Configuration Wizard implements your changes and you do not need to restart the system.

To configure additional filesystems, choose one of:

- Create Filesystems (NAS Storage Systems) manually.
- Run the Pillar Axiom Capacity Planner
- Run the Configuration Wizard again for each filesystem that you want to create.

To change the configuration of a File Server, see Modify File Server Attributes.

About Global Settings Configuration

The first time that you log in to configure a Pillar Axiom system, define the following system-wide settings:

- Configure the Pillar Axiom System Time, to set and synchronize the time across all Pillar Axiom components.
- Configure Management and Data Path Interfaces, to enable and configure Dynamic Host Configuration Protocol (DHCP) support, transmission characteristics of the management ports, and link aggregation.
- Configure iSCSI System Settings, to set system-wide iSCSI settings if your configuration requires all iSCSI connections to use CHAP, Access Control, or both. If access control is defined for each initiator, you do not need to configure iSCSI at the system level.
- Configure Notification Settings, to define an electronic mail server in your network that receives Pillar Axiom alerts and forwards them to administrator email accounts.
- Configure Call-Home Settings to enable Call-Home, a feature that notifies Pillar Data Systems about issues in the Pillar Axiom system.
- Configure Security Settings, to define time-out periods and failed login attempts.

The success of other configuration tasks depends on the system-wide settings. For example, if you do not configure the email server, the system cannot send alerts.

Configure the Pillar Axiom System Time

Configure the Pillar Axiom system time so that event and logging timestamps are accurate and time-dependent applications, such as email, work properly.

When you configure the system time, choose one of:

- **Manual**, which allows you to set the system time yourself.
- **Network Time Protocol (NTP) server**, which will synchronize the system time with an external NTP server.

1 Click the **System** icon in the top context pane.

2 Click the **Summary** link in the left navigation pane.

3 Choose **Modify Time Settings** from the **Actions** drop-down list.

4 On the **Details** tab, select the time zone that you want to define for the Pillar Axiom system.

5 Select one of these options:

- **Set Time Manually** and enter the date and time.

  Note: This option is not valid when Pillar Axiom SecureWORMfs filesystems exist on the system. Pillar Axiom SecureWORMfs requires an NTP server.

- **Use NTP Server** and enter the IP addresses of the primary NTP server and up to two alternate NTP servers. If the primary NTP server is not available, the Pillar Axiom system consults the secondary servers in round-robin fashion until a connection is made.

6 To save the system time setting, click **OK**.

For parameter definitions, see the System Time Page.

About Management and Data Path Interfaces

The management interface provides connectivity between the end user data network and the Pillar Axiom Pilot management controller. You can choose which method to use to assign the primary IP addresses to the management interface on the Pilot:

- Dynamic Host Configuration Protocol (DHCP), which assigns a primary IP address dynamically when the Pilot boots.
- Static IP Address assigns a permanent, primary IP address to a control unit (CU) in the Pilot as well as alternate IP addresses for the ports on the partner CU. If the management software cannot access the primary IP address, it accesses an alternate IP address.

The data path interface provides connectivity between the data network and the Pillar Axiom Slammers. If needed, enable and configure link aggregation (trunking) so that you can:

- Use up to four of the Gigabit Ethernet (GbE) ports on a single Slammer control unit (CU) as though those physical links were one logical link, or an enhanced-speed GbE port.

- Connect the Slammer to an appropriately configured GbE switch, and let the switch dynamically balance network traffic.

Link aggregation also provides automatic failover. If one of the ports on the Ethernet switch or Slammer fails or if one of the cables or connectors fails, the Ethernet switch routes all traffic to one of the other aggregated ports. For information on configuring link aggregation, see *Tips on Implementing Link Aggregation*.

Configure Management and Data Path Interfaces

Configure the management interface by setting the IP addressing method (static or dynamic) for the Pilot and the network attached storage (NAS) data path by configuring link aggregation on the Slammers.

1. Click the **System** icon in the top context pane.
2. Click the **Networking** link under the **Global Settings** heading in the left navigation pane.
3. Choose **Modify Network Settings** from the **Actions** drop-down list.
4. Click the **Interfaces** tab.
5. Enter values for the attributes that you want to modify.
6. To save the modified interfaces, click **OK**.

For parameter definitions, see the Interfaces Page.

See also Create Alerts.

About iSCSI System Settings Configuration (Optional)

If you have iSCSI hosts configured to use CHAP, Access Control, or both, you must also set up system-wide iSCSI settings. This configures the authentication and access controls on the Pillar Axiom system in which the host must match to gain access. If you have CHAP and Access Control configured for each initiator, you do not need to configure iSCSI globally.

To configure the iSCSI ports for each Slammer, see Modify iSCSI Port Settings.

Configure iSCSI System Settings

1. Click the **System** icon in the top context pane.
2. Click the **iSCSI** link under the **Global Settings** heading in the left navigation pane.
3. Choose **Modify iSCSI Settings** from the **Actions** drop-down list.
4. Enter values for the attributes that you want to modify.
5. To save the modified interfaces, click **OK**.

For parameter definitions, see the iSCSI Page. For more information about using the iSCSI feature, see About Licensing Optional Premium Features.

About Notification Settings Configuration

You can define an email server to receive alerts from the Pillar Axiom system and send the email messages to designated recipients. If you do not set the email server, the system does not send alerts to administrators when system events occur.

The optional automatic failback feature applies to network attached storage (NAS) systems and defines whether the Pillar Axiom storage system should automatically perform a recovery operation when a previously unavailable NAS Slammer control unit (CU) becomes available. If this feature is not enabled, you must perform the recovery operation manually.

Note: SAN Slammers automatically fail back when a previously unavailable CU becomes available.

Configure Notification Settings

1. Click the **System** icon in the top context pane.
2. Click the **Networking** link under **Global Settings** in the left navigation pane.
3. Choose **Modify Network Settings** from the **Actions** drop-down list.
4. Click the **Notification** tab.

5 Select the **Enable Email** check box.

6 Enter the IP address of the email server.

7 Select the **Enable Automatic Recovery of Failed NAS Slammers** check box (optional).

8 To save the notification settings, click **OK**.

For parameter definitions, see the Notification Page.

See also: Create Alerts.

About Call-Home

The Call-Home feature notifies Pillar Data Systems about issues in the Pillar Axiom system. 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 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 and are listed in the GUI. The system administrator can select a session file and download it to a GUI client machine or send it directly to the currently targeted Pillar Data Systems server.

Call-Home sessions can also be sent to a local Call-Home server. Contact the Pillar World Wide Customer Support Center for details.

Configure Call-Home Settings

1 Click the **System** icon in the top context pane.

2 Click the **Networking** link under **Global Settings** in the left navigation pane.

3 Choose **Modify Network Settings** from the **Actions** drop-down list.

4 Click the **Call-Home** tab.

5 Select the **Enable Call-Home** check box.

6 Enter IP addresses for the Call-Home Primary and Secondary DNS (Domain Name System).

7 Specify the type of server to which to send the Call-Home files:

- Use **Pillar Data Systems Server**
- **Use Local Host**

8 Enter the values in the appropriate configuration fields for the type of server selected in Step 7.

9 To save the Call-Home settings, click **OK**.

For parameter definitions, see the Call-Home Configuration Page.

About Security Settings Configuration

Configure security settings for administrator accounts and browser configuration.

- Set the number of failed login attempts that the Pillar Axiom system permits. When the threshold is exceeded, the system disables the account and writes an entry in the event log. A Primary system administrator or Administrator 1 account can re-enable the account, and the system resets the counter upon a successful login. If you do not set this value, there is no limit to the number of unsuccessful login attempts.
- Set the session time-out so that the Pillar Axiom system terminates an administrator's session after a given period of inactivity. If you do not set this value, inactive sessions are terminated after 20 minutes.
- Select **Secure Session Only** to specify that administrator access to the Pillar Axiom system is over secure HTTP sessions. Upload a secure sockets layer (SSL) certificate to the Pillar Axiom Pilot to authenticate login attempts.

Configure Security Settings

1 Click the **System** icon in the top context pane.

2 Click the **Security** link under the **Global Settings** heading in the left navigation pane.

3 Choose **Modify Security Settings** from the **Actions** drop-down list.

4 Select the **Secure Session Only** check box (optional).

5 If you selected **Secure Session Only**, click **Upload Certificate** and navigate to and select the secure sockets layer (SSL) certificate that you want to use.

6 Enter the values in the failed login attempts and session timeout fields to define the administrator login limits.

7 Enter a login screen message (optional).

8 To save the security settings, click **OK**.

For parameter definitions, see:

- Account Security Settings Page
- *See also:* About Administrator Account Creation

About Storage Resources Creation

Creation of storage resources is the foundation of the initial configuration process. The following sections explain how to:

- Run the Pillar Axiom Capacity Planner to experiment with different allocation models.
- Create LUNs.
- Create Filesystems (NAS Storage Systems).
- Create File Servers (NAS Storage Systems).
- Join File Servers to CIFS Domains (NAS Systems).
- Create Volume Groups.

Before you read the how-to topics, learn about capacity allocations in About Volume Capacity and Provisioning.

About Volume Capacity and Provisioning

- About Over-Committed Volumes
- Free Capacity and Volume Creation
- Provisioning Over-Committed Volumes
- Growth Increments
- Capacity Overhead
- Parity in Reported Capacities
- Reclaiming Capacity

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

About Storage Classes

The Storage Class feature allows you to specify the preferred storage media to use for a logical volume (filesystem or LUN).

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, and solid state drive (SSD). Pillar Axiom 600 systems allow an administrator to explicitly manage volume placement within the overall system storage pool, first by Storage Class and then by relative priority within that Storage Class.

Pillar Axiom storage systems support the following three Storage Classes:

- FC (Fibre Channel drives, 15K RPM)
- SATA (serial ATA drives, 7.2K RPM)
- SLC SSD (single-level cell, solid state drive)

Note: Which Storage Classes are available on a particular Pillar Axiom 600 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 SLC SSD Storage Class, striping for a volume is across all available drives, regardless of the priority setting.

For more information, see Filesystem Page, Quality of Service Tab and LUN, Quality of Service Tab.

About Over-Committed 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 *over commit* a logical volume (LUN or filesystem) by:

- Allocating capacity based on future needs.
- Drawing 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 size 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 LUN or filesystem in increments.

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

A logical volume is over-committed when its maximum capacity is larger than its initial capacity. A logical volume can be overcommitted by any amount.

Storage is provided when write operations to the logical volume require regions that are un-allocated.

Note: The additional space may not be contiguous with previous allocations.

Thin provisioning depends on the relationship between the initial values you set for the following parameters:

- *Initial capacity* that the system makes available to the logical volume. This value is labeled **Total Capacity** after successful volume creation.
- *Maximum capacity* to which the logical volume can grow. This value is labeled **Growth Max** after successful volume creation.

Note: Growth occurs only within the Storage Class on which the volume is based.

The initial capacity can be any value, up to and including the maximum.

Tip: If you do not desire thin provisioning, set the initial capacity of the allocation equal to the maximum capacity.

When a system contains a mix of storage types (Storage Classes), new volume allocation cannot cross storage boundaries. The allocation occurs only within the specified Storage Class.

For more information about this feature, see About Licensing Optional Premium Features.

Free Capacity and Volume Creation

A minimum amount of free space is required to create a new logical volume (filesystem or LUN). 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
- 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 500 GB, for a total physical space allocation of 800 GB.

The amount of capacity required by the volume increases when:

- The QoS priority is set to double redundancy. In this case, the allocation doubles.
- The drive geometry is set to Distributed RAID or RAID 5 with Wide Stripe. In this case, the allocation increases to provide extra capacity for parity and other overhead (see Capacity Overhead and Parity in Reported Capacities).

If a request to create a logical volume fails because of capacity issues, it would be for the following reasons:

- Insufficient capacity remains in the Storage Class specified for the volume.
- Sometimes, the system may need to round your request to a slightly larger size, which then is greater than available capacity.
- You have requested double redundancy, but sufficient capacity is not available on two different Bricks within the specified Storage Class.
- You have requested a Quality of Service (QoS) priority setting of Premium, but sufficient capacity does not exist on that storage band.

Provisioning Over-Committed 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.
- Over-committed (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 over-committed, which defines its thinness, 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 over-committed, which defines its thinness, 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.

Some applications access most or all of the addressable space for a volume. In these cases, the volume transitions from being thinly provisioned to being fully provisioned while the application executes. An example of such an application is the `mkfs` utility, which creates a filesystem on a partition. As `mkfs` executes and formats the filesystem, most or all of the partition is written by the application, causing the underlying volume on the Pillar Axiom system to become fully provisioned. In cases such as these, creating the underlying volume using thin provisioning has little value.

Windows reserves a substantial amount of metadata for a filesystem that has been formatted as an NTFS (New Technology File System) volume. The layout of this metadata causes an early allocation of thinly provisioned space. The primary NTFS metadata consists of the following objects:

- Boot record, which is written to both the beginning and the end of the volume.
- Master File Table (MFT), which is written to both the beginning and the middle of the volume.

To prevent the MFT from becoming fragmented, Windows reserves a buffer around the MFT. The size of this buffer is configurable and can be 12.5%, 25%, 37.5%, or 50% of the drive space. Windows will not create new files in this buffer region until the unused space is consumed. Each time the rest of the drive becomes full, the buffer size is halved. This strategy provides new space for additional write operations.

Pillar does not recommend creating a thinly provisioned LUN that is filled up greater than 90% on the first in-fill, especially with NTFS. NTFS writes all over the LUN causing allocations that do not match the amount of data that is written. A heavily used NTFS filesystem running without much free capacity will eventually use up all the capacity unless the filesystem is de-fragmented periodically. NTFS favors writing into new allocated space instead of reusing previously written space. NTFS works with thin provisioning initially but can quickly use up more allocation than the amount of data the filesystem would show as used.

Because thin provisioning uses Slammer resources and affects performance, a good use of thin provisioning would be for a LUN that has the following characteristics:

- An initial allocation of the amount of existing data plus 10%.

- A maximum growth extent of two times the initial data size plus 10%.

For example, given 420 GB of file data, the administrator should configure the initial LUN size to be approximately 470 GB and a maximum size of approximately 1 TB.

Note: How much capacity NTFS uses depends on many factors, including the size of writes, where the writes are made, and other factors such as the type of storage used in the storage pool.

On Linux platforms, EXT2 and EXT3 filesystems write metadata over the entire range of logical block addresses (LBAs) of the LUN. The drive is organized into block groups and metadata exists at the beginning of each block group. This configuration typically causes the entire LUN to be provisioned when the administrator creates a filesystem. This full provisioning occurs because the metadata write is below the allocation unit used by Pillar Axiom systems. This condition causes the system to expand every allocation extent to the maximum size.

In summary, the success of utilizing thin provisioning depends on the filesystem or the application using the LUN.

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 (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 (filesystem or LUN), the system adds a certain amount of overhead to a request for the capacity of the filesystem or LUN. 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.

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 (LUN or filesystem), the system reconditions the space (by writing a predefined bit pattern) before freeing 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.

About Performance Profiles

Performance profiles determine the quality of service for a logical volume (filesystem or LUN).

When you use the Pillar Axiom Configuration Wizard to configure a logical volume, the Configuration Wizard allows you to select predefined performance characteristics that are based on the category of application that typically accesses the volume or on a more specific performance profile. When using a specific performance profile, you can:

- Select one that you have previously saved.
- Select one of the pre-configured performance profiles that Pillar provides:
    - **Highest Throughput**. This profile stripes data across all available drives in the system. The Highest Throughput profile is intended for use in benchmarking or in environments having a small number of logical volumes to be configured.

Important! Striping data across all available drives can lead to unexpected contention in larger configurations.

- **Oracle Automatic Storage Management (ASM)**. This profile makes use of a wide stripe (which has a strip depth of 1 MB) for the configured logical volume. A wide stripe minimizes the number of seeks required to service data requests in an Oracle ASM environment by matching strip size to the application request size.

A performance profile is defined by means of an XML document. The sample profile illustrated below shows what a valid profile might look like for a filesystem.

```
<Profiles>
  <PerformanceProfile>
    <ProfileName>Exchange Logs</ProfileName>
    <Performance>
      <RelativePriority>Archive</RelativePriority>
      <FileSizeBias>Small</FileSizeBias>
      <AccessBias>Sequential</AccessBias>
      <IOBias>Read</IOBias>
    </Performance>
    <Reliability>
      <Redundancy>Standard</Redundancy>
    </Reliability>
  </PerformanceProfile>
</Profiles>
```

The above XML elements are what the Configuration Wizard supports.

For more information on selecting a performance profile, see Configure NAS Storage Parameters. For information on exporting a performance profile, see Performance Profile Page.

About RAID Array Stripes

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

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) performance 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.

About Enhanced Performance for Oracle ASM

You can enhance the I/O throughput of a logical volume by using a 1 MB strip RAID geometry.

Using the Quality of Service (QoS)-based management tools provided in the Pillar Axiom Storage Services Manager, you can improve the overall performance of logical volumes in an Oracle Automatic Storage Management (ASM) environment by taking advantage of an alternative internal structure of the RAID array. This alternative geometry provides a wide stripe by increasing the stripe width to 1 MB.

The system utilizes this alternative RAID geometry when you select the Oracle ASM performance profile in the Configuration Wizard. Selecting that profile allows the system to optimize the internal structure of the stripe that the RAID array implements so that I/O chunks match those of Oracle 10g systems that utilize ASM. Because Oracle ASM performs 1 MB random I/O operations as its normal access pattern, matching the stripe size utilized in the Brick to the I/O size utilized by Oracle ASM minimizes the I/O required to support a given Oracle ASM workload. LUNs created without the Oracle ASM profile use standard striping.

For more information about this feature, see About Licensing Optional Premium Features.

About Enhanced Performance for Random Write Operations

You can enhance the performance of random write operations on a logical volume under certain circumstances.

Using the Quality of Service (QoS)-based management tools provided in the Pillar Axiom Storage Services Manager, you can improve the overall performance of random-write intensive applications by taking advantage of a nested RAID structure that replaces four I/O operations with a parallel mirrored-write operation.

The system utilizes this RAID geometry when you set the optimization settings for a logical volume to random access with a write I/O bias. In this case, the system allocates space on a Distributed RAID array on which it performs parallel mirrored-write operations (two writes).

Note: This Distributed RAID geometry applies only to newly created volumes and to volumes that are migrated by means of QoS migration that may occur when adding serial ATA (SATA) Bricks to a system containing Fibre Channel Bricks.

In addition, because the data resides on two independent drives, Distributed RAID arrays allow the system to optimize read operations as well. In this case, the system can select the least busy of the two drives.

When the Distributed RAID geometry is not used, the system allocates space on a regular RAID 5 array for the volume and performs four I/O operations (two reads and two writes).

For more information about this feature, see About Licensing Optional Premium Features.

Run the Pillar Axiom Capacity Planner

The Capacity Planner simulates the creation of one or more logical volumes (filesystems or LUNs) based on the predefined performance profile that you select out of a list or on a new simulation. For the predefined profiles, each profile in the list contains a pre-configured collection of Quality of Service (QoS) settings.

In comparison, the Configuration Wizard guides you through the initial setup and actual configuration of your Pillar Axiom system.

For filesystems, the Capacity Planner assumes that a File Server already exists. The Capacity Planner does not include creating shares or exports and does not allow you to locate filesystems or LUNs in a volume group.

At the completion of the Capacity Planner, you can keep the simulated results and use this information to create filesystems and LUNs. You can then create shares and exports for each filesystem.

1 Click the **System** icon in the top context pane.

2 Click **Capacity Planner** at the bottom of the left navigation pane.

3 Choose the type of simulation to run and click **Next**.

Types of simulation:

- **Use Pillar Application Profile configuration.** Click the drop-down list and select a profile.
- **Create a new simulation.**

4 In the Welcome to the Pillar Axiom Capacity Planner page, click **Next**.

5 Enter the values you want to simulate and click **Next** to move through the Wizard pages.

Tip: Some 32-bit operating systems cannot access or manage Network File System (NFS) or combined NFS and Common Internet File System (CIFS) filesystems of two terabytes or greater in size. If in doubt, check the operating system limits for your client and filesystem. For mixed 32-bit and 64-bit clients, limit filesystem sizes such that both clients will be able to work with those filesystems.

6 Review the proposed and predicted Quality of Service (QoS) entries (see the following example).

Figure 10 Example Pillar Axiom Capacity Planner results

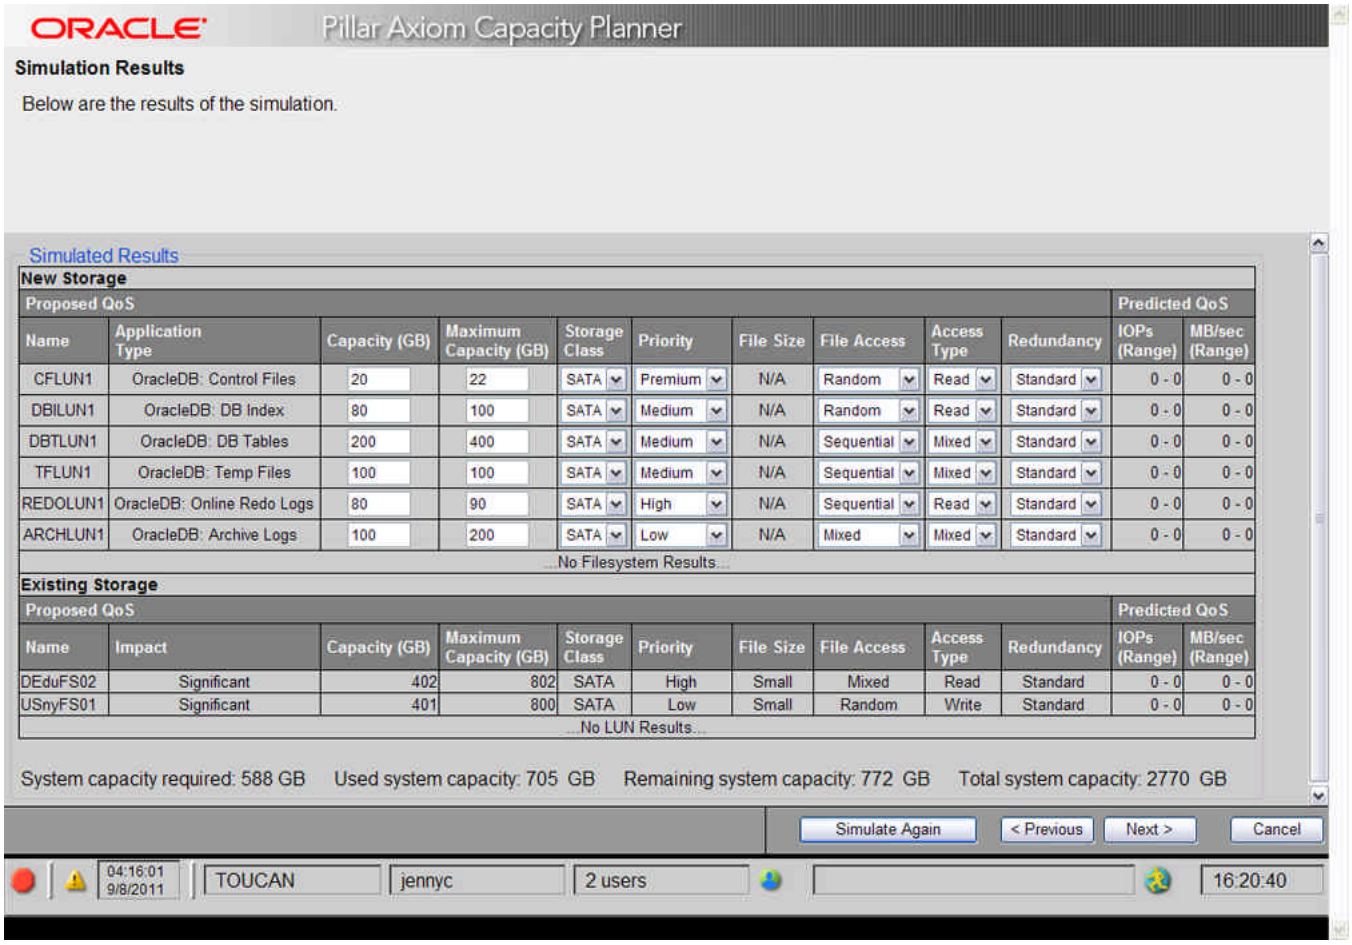

7 When you are satisfied with the simulation results, click **Next**.

8 Determine the disposition of the simulated configuration parameters.

- Save simulated configuration for later use.

  Tip: You can retrieve the saved simulation at a later time. In this way, you do not have to start over again from the beginning.

- Create the simulated configuration.

9 Click **Next**.

10 As needed, provide the requested values and then click **Next**.

11 To end the capacity-planning simulation, click **Finish**.

Result:
Depending on the disposition that you selected for the configuration parameters, either the configuration file is saved for future use or the logical volumes are created.

Note: If any errors are displayed, click **Previous** to go back and correct the errors.

Note: You cannot use the Pillar Axiom Capacity Planner to specify port mapping or masking for LUNs. As such, when you create a LUN using the Capacity Planner, the LUN can be accessed by any SAN host. To configure port mapping or masking after you create the LUN using the Capacity Planner, you can Define LUN Mapping (Optional).

About LUN Creation

A LUN is defined as:

> A logical volume within a storage area network (SAN). Administrators assign storage resources and Quality of Service (QoS) attributes to each logical unit (LUN).

The Pillar Axiom system calculates whether enough storage resources are available to create a new logical volume. Click **Optimizer** on the **Quality of Service** tab to determine how much capacity is available for each combination of Storage Class and Redundancy Quality of Service (QoS) values.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

For information regarding capacity attributes of a logical volume, refer to About Volume Capacity and Provisioning.

Create LUNs

1. Click the **Storage** icon in the top context pane.
2. Click the **LUNs** link in the left navigation pane.
3. Choose **Create LUN** from the **Actions** drop-down list.
4. Enter a LUN name on the **Identify** tab.
5. Select a volume group in which to create the LUN.

   Note: The Volumes option represents the system root.
6. Assign the LUN to a Slammer.

Choose one of:

- Auto-assign the LUN to a Slammer. This allows the Pilot to move resources to the other control unit (CU) in the event of a failover and maximizes performance by balancing the system load with existing LUNs and filesystems.

  Note: If you select this option, you can set up port masking and port mapping after you create the LUN. Select the LUN and choose **Modify LUN** from the **Actions** drop-down list.

- Assign the LUN to a Slammer. Use this option when you want to access the data from a particular port (port masking).
  - Select a Slammer from the drop-down list.
  - Select a control unit (CU) from the drop-down list.

7 Select the host access to the LUN.

Choose one of:

- Only selected hosts to access this LUN. Use this option to map the LUN to one or more specific SAN hosts.
- Allow all hosts to access this LUN and select a LUN number from the drop-down list.

8 Select either Fibre Channel (FC) or iSCSI protocol, or both.

This selection determines the protocol(s) that will be permitted for accessing the LUN.

Important! When you select both FC and iSCSI protocols, the system will use FC optimized and non-optimized paths as a preference over iSCSI paths. Also, the system will not mix load balancing between protocols.

Define LUN Quality of Service Attributes

The Quality of Service (QoS) attributes for a LUN include capacity, performance, and redundancy settings. You can also specify the amount of space you want to allocate on the system for Clone LUNs.

Caution We strongly recommend that you allocate sufficient repository capacity to minimize the chances of running out of this space (which could lead to data inconsistency or loss). To set sufficient capacity, use a value equal to the source volume capacity times the number of replicas times the maximum rate of change. For example, for a 100 GB volume that is projected to have 20 active replicas at a time and no more than a 20% rate of change, use a value of 400 GB for the clone repository.

1 Click the **Quality of Service** tab.

2 Click **Optimizer** and review the available capacity for each possible combination of Storage Class and redundancy values.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

3 Enter values for the LUN that specify its capacity, redundancy, clone space, and performance attributes, including Storage Class.

To accommodate the level of RAID protection required to allocate a newly created logical volume (filesystem or LUN), the system adds a certain amount of overhead to a request for the capacity of the filesystem or LUN. 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.

Caution Under HPUX (and possibly other operating systems), changing the capacity of the LUN without first migrating data from the LUN will result in data loss. Storage Administrators need to be aware of how their operating system will behave when LUN capacity is increased, which may depend on the details of how they have configured their environment.

Note: You can modify a LUN to update these attributes after the LUN has been created.

4 Click **Run Simulation** to see the effect that the settings for this LUN would have on other filesystems and LUNs (optional).

For parameter definitions, see LUN, Quality of Service Tab.

Define LUN Mapping (Optional)

LUN mapping allows you to choose specific SAN hosts that can access the LUN.

Note: This page is available only when LUN mapping is enabled. To enable, click the **Identity** tab and select the **Only selected hosts** option.

1 Click the **Mapping** tab.

2 Select the storage area network (SAN) hosts that you want to allow access to this LUN.

In the drop-down list, the system displays:

- Recognized hosts.
- Unassociated World Wide Names (WWNs) for Fibre Channel (FC) hosts and iSCSI names (iSCSI) for hosts on the SAN network that are not using Pillar Axiom Path Manager.

3 Select the LUN number to assign to the LUN.

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

4 Select the ports on the Slammer control units (CUs) that you want to mask so that the LUN cannot be accessed from the specified ports (optional).

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

5 Click **Create Map**.

6 Click **OK**.

If the iSCSI host that you want to allow access to the LUN is not displayed in the list, you can add it by using Associate Hosts Page option.

For parameter definitions, see the following:

- LUN, Identity Tab

- LUN, Quality of Service Tab
- LUN, Mapping Tab
- LUN, Host Connections Tab

About Filesystem Creation

A filesystem is a logical volume that organizes and catalogs files and assigns resources to a given collection of directories and files in a network attached storage (NAS) system. Administrators can assign different Quality of Service (QoS) attributes to each filesystem. A filesystem must be associated with a File Server.

The Pillar Axiom system calculates whether enough storage resources are available to create a new logical volume. Click **Optimizer** on the **Quality of Service** tab to determine how much capacity is available for each combination of Storage Class and Redundancy Quality of Service (QoS) values.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

For information regarding capacity attributes of a logical volume, refer to About Volume Capacity and Provisioning.

Create Filesystems (NAS Storage Systems)

1. Click the **Storage** icon in the top context pane.
2. Click the **Filesystems** link in the left navigation pane.
3. Choose **Create Filesystem** from the **Actions** drop-down list.
4. Enter a filesystem name on the **Identity** tab.
5. Select a volume group in which to create the filesystem.

   Note: The Volumes option represents the system root.
6. Select **Enable Retention Policy for SecureWORMfs** to create a Pillar Axiom SecureWORMfs filesystem (optional).

Important! You must enable the Network Time Protocol (NTP) server for the system time settings to create a Pillar Axiom SecureWORMfs Compliance filesystem.

- Choose one of:
  - **Standard**—can be deleted any time.
  - **Compliance**—cannot be deleted if it contains files (protected or unprotected).
- On the **Retention Policy** tab, enter the retention period settings.

7 Associate the filesystem with a File Server. Choose one of:

- Select an existing File Server from the **File Server** drop-down list.
- Click **Create New File Server**.
  - On the File Server pages, enter values for network connections, protocol access, and account mapping, and click **OK** to save the new File Server.
  - On the **Filesystem Identify** tab, select the new File Server from the **File Server** drop-down list.

8 Select **Create Automatic Snap FS Schedule** to allow the system to create a snapshot every four hours (optional).

9 Assign the filesystem to a specific Slammer control unit (CU).

Choose one of:

- **Auto-assign Filesystem to a Slammer**

  The system selects the Slammer and CU to which to associate the filesystem.

- **Assign Filesystem to Slammer**

  Identifies the Slammer to which the filesystem is to be assigned. Use the **Assign Control Unit** drop-down list to associate the filesystem with a particular control unit (CU) of the identified Slammer.

For parameter definitions, see:

Filesystem Page, Identity Tab

Filesystem Page, Retention Policy Tab

Define Filesystem QoS Attributes

The Quality of Service (QoS) attributes for a filesystem include capacity, performance, and redundancy settings.

Caution We strongly recommend that you allocate sufficient repository capacity to minimize the chances of running out of this space (which could lead to data inconsistency or loss). To set sufficient capacity, use a value equal to the source volume capacity times the number of replicas times the maximum rate of change. For example, for a 100 GB volume that is projected to have 20 active replicas at a time and no more than a 20% rate of change, use a value of 400 GB for the clone repository.

1 Click the **Quality of Service** tab.

2 Click **Optimizer** and review the available capacity for each possible combination of Storage Class and redundancy values.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

3 Enter values for the filesystem that specify its capacity, redundancy, clone space, and performance attributes.

To accommodate the level of RAID protection required to allocate a newly created logical volume (filesystem or LUN), the system adds a certain amount of overhead to a request for the capacity of the filesystem or LUN. 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.

Tip: Some 32-bit operating systems cannot access or manage Network File System (NFS) or combined NFS and Common Internet File System (CIFS) filesystems of two terabytes or greater in size. If in doubt, check the operating system limits for your client and filesystem. For mixed 32-bit and 64-bit clients, limit filesystem sizes such that both clients will be able to work with those filesystems.

You can modify the filesystem to update these attributes after the filesystem has been created.

4 Click **Run Simulation** to see the effect these settings that this filesystem would have on other filesystems and LUNs (optional).

For parameter definitions, see Filesystem Page, Quality of Service Tab.

About Retention Policy Settings

A Pillar Axiom SecureWORMfs is defined as:

> A type of filesystem used to enforce data retention. Data is stored on a Pillar Axiom SecureWORMfs in a protected (non-rewritable) manner.

The Retention Policy settings specify the period for which files on the Pillar Axiom SecureWORMfs filesystem must be retained.

Note: This page is available only when Retention Policy for Pillar Axiom SecureWORMfs is enabled. To enable, click the Identity tab and select **Enable Retention Policy for SecureWORMfs**.

Note: You must enable the Network Time Protocol (NTP) server for the system time settings to create a Pillar Axiom SecureWORMfs Compliance filesystem.

For more information about this feature, see About Licensing Optional Premium Features.

Define Retention Policy

1 Click the **Retention Policy** tab.

2 Enter values to specify the Master Retention Period for all files on the filesystem.

For parameter definitions, see Filesystem Page, Retention Policy Tab.

Add CIFS Shares to a Filesystem

1. Click the **Shares** tab.
2. Enter a share name and path, and optionally a share comment.
3. To enable the share point immediately, select **Enabled**.
4. Click **Create** to create the new share point.

For parameter definitions, see Filesystem Page, Shares Tab.

Add NFS Exports to a Filesystem

1. Click the **Exports** tab.
2. Enter the export parameters:
   - Export path for the filesystem.
   - The user ID (UID) for anonymous users. Set the UID for anonymous users to zero (0) if you want all data access to be as if from a root user. Set the UID to **nobody** (often -2, but not always) if you want anonymous-user access to be identified as the user `nobody`.
3. Select the host type and enter the associated host value:
   - **All Hosts**: Everyone can access the export point.
   - **Single Host**: Only the specified host can access the export point.
   - **NIS Netgroup**: Everyone within the Network Information Service (NIS) netgroup can access the export point.
   - **Network**: Everyone on the specified subnet can access the export point.
4. To make the export point available to users in read-only mode, select **Read Only**.
5. To allow users as root users on the export point, select **Root Access**.
6. Click **Create**.
7. Click **OK** on the NFS Exports page.
8. From the network file system (NFS) client, mount the filesystem remotely:

- Mount the exported filesystem on an NFS client.
- Change directories to the mount point.
- Create the directories that you want to export to other NFS clients under the root export point.
- Change the directory permissions to the required values (may require root access).
- Unmount the / export point (optional).

9 From the Pillar Axiom Storage Services Manager NAS Exports page, delete the / export point (optional).

10 From the Pillar Axiom Storage Services Manager NFS Export page, export the directories that you created under the root export point.

- Enter the export parameters:
  - Enter a full path name that starts with a slash (/) and does not include the filesystem name.
  - Enter the user ID (UID) for anonymous users. Set the UID for anonymous users to zero (0) if you want all data access to be as if from a root user.
- Select the host type and enter the associated host value:
  - **All Hosts**: Everyone can access the export point.
  - **Single Host**: Only the specified host can access the export point.
  - **NIS Netgroup**: Everyone within the NIS netgroup can access the export point.
  - **Network**: Everyone on the specified subnet can access the export point.

11 Click **Create**.

12 Click **OK** on the NFS Exports page.

For parameter definitions, see Filesystem Page, Exports Tab.

Add NFS Exports to a Filesystem (External File)

1 Click the **Exports** tab.

2 Click **Import from External File**.

3 Enter the path or click **Browse** to locate the external file and click **OK**.

4 On the Import External File page, click **OK**.

5 From the NFS client, mount the filesystem remotely:

- Mount the exported filesystem on an NFS client as root.
- Change directories to the mount point.
- Change the directory permissions to the required values (may require root access).

6 Click **OK** on the NFS Exports page.

For parameter definitions, see Filesystem Page, Exports Tab.

About Filesystem Quotas

You can limit and track the amount of storage used in a shared system through the use of quotas.

Pillar Axiom storage systems allow you to manage the physical space of filesystems by dividing the space among directories. You can accomplish this division through the use of directory quotas. Also, you can further split the storage space assigned to a particular directory by creating nested directories within the parent directory and assigning directory quotas to these subdirectories.

In addition to quota limits at the directory level, you can set quota limits for users and groups, including setting a default limit that applies to all users or all groups that do not have a specific limit.

After you create the directory quota, the system enables usage tracking for users and groups. Usage tracking is enabled independently from default limits or specific quota limits having been set.

Note: Quota setup occurs through use of the Pillar Axiom 600 user interfaces — Pillar Axiom Storage Services Manager or Command Line Interface (CLI), not through commands issued from on a client machine.

When you add quotas to a directory that already has quotas or is empty, the filesystem remains available. However, when you establish the first quota in a directory that already contains files or directories, the filesystem is temporarily unavailable while the system adds the quota.

Similarly, when you delete a quota on a non-empty directory, the system takes the filesystem offline.

When a replica of a filesystem is created, the system transfers the quota settings associated with the original or source filesystem to the replica for the following types of operations:

- Copies
- Snapshots
- Clones
- Pillar Axiom MaxRep Replication for NAS synchronizations
- Block-based NDMP backups (and restores)

Note: File-based NDMP backup operations do not copy quota settings.

About Quotas for Directories

You can limit the amount of physical storage (in megabytes) that a particular directory can consume.

The following figure shows two directory quotas: one for `ny1` and another for `sf2`.

Figure 11 Directory quotas

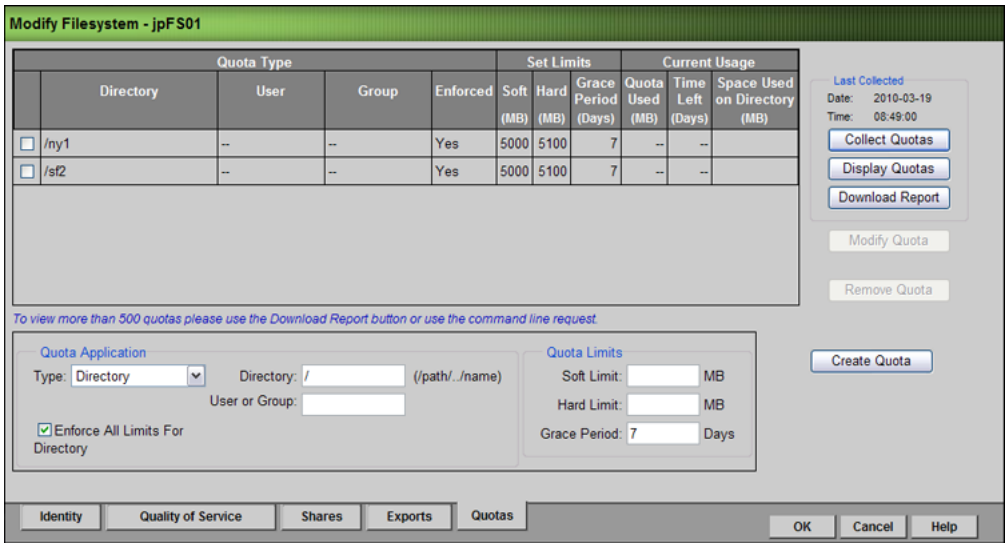

The storage limit for a directory encompasses the individual limits that may be defined for all the filesystem objects (such as files and subdirectories) that you create beneath the directory on which you create the directory quota.

The directory quota is the basis of creating two other quota types, the user and the group quota.

About Specific Quotas for Users and Groups

You can limit the amount of physical storage (in units of megabytes) that a particular user or group can consume.

The following figure shows a specific quota for the user named `jepop` in the directory named `ny1`.

Figure 12 Specific user quotas

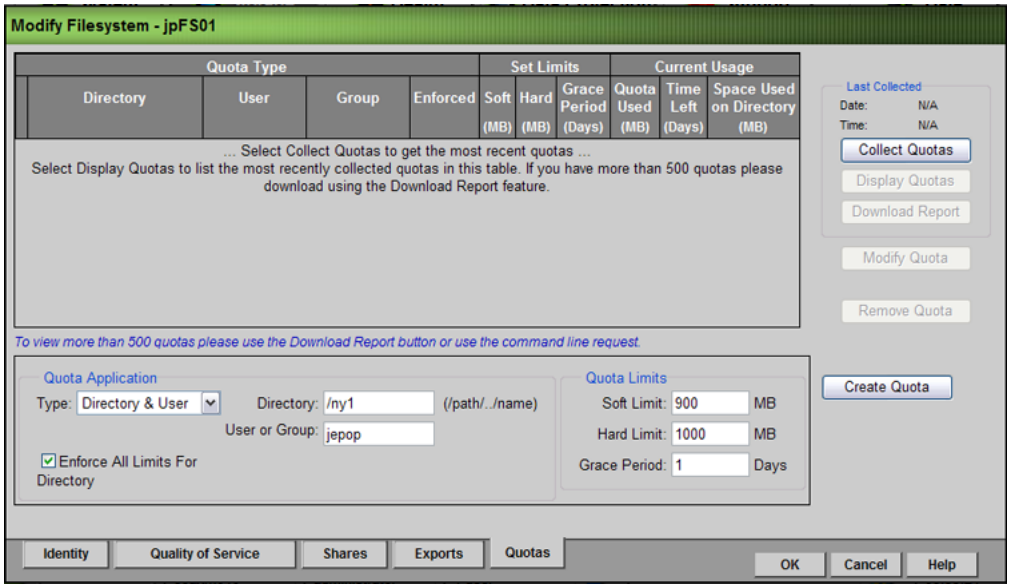

When setting a specific quota limit for a particular *user* within a given directory, you need to set the type of quota to `Directory & User` and then specify the absolute path name of the directory.

Similarly, when setting a specific quota limit for a particular *group* within a given directory, you need to set the type of quota to `Directory & Group` and then specify the absolute path name of the directory.

Important! A specific group quota encompasses the cumulative usage of all the users within a group.

About Default Quotas for Users and Groups

You can use default quotas to allow many users to share the same limit value (hard limit, soft limit, and time limit).

Default limits for groups and users are created by typing an asterisk (*) in the **User or Group** field and setting the desired limit values.

Figure 13 Default quotas

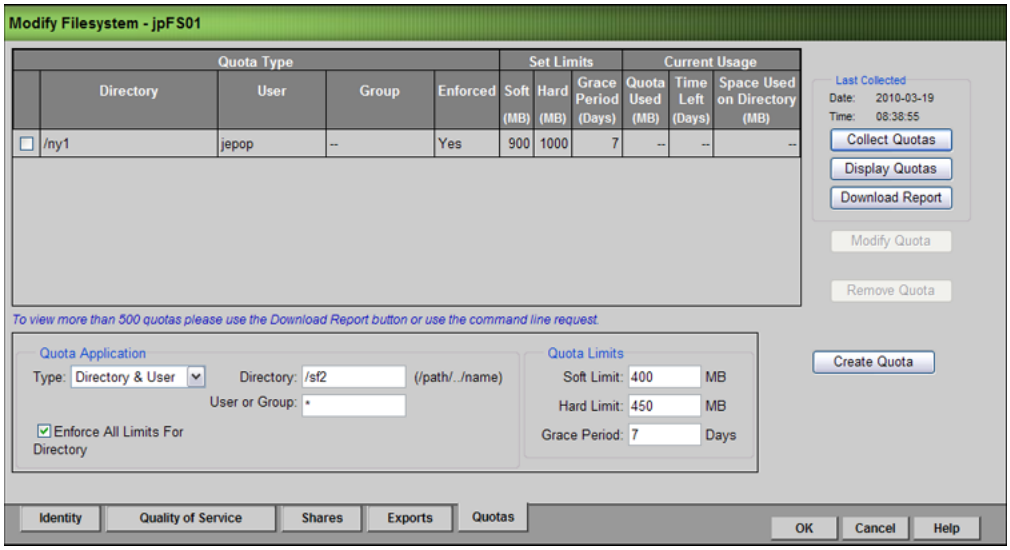

Both user limits and group limits can co-exist on a given folder.

Default limits apply to newly added users and to existing users who currently own files that consume capacity but who have had no quota limits.

You can modify or remove group and user quotas by selecting the quota and then clicking **Modify** or **Remove**, respectively.

You can modify the default limits for users by using an asterisk (*) for user name and changing the appropriate limit values. Changed limit values apply to all users who do not have specific quota limits, existing or new.

When you remove the default limits for users, the users who have no specific limit are bounded only by the directory quota limits, or other nested limits, if they exist.

About Quota Reporting

You can view the list of quota limits for users and groups that are configured on a Pillar Axiom storage system. The limits that are reported include both specific and default settings.

If a quota file already exists, the system lists the time stamp to show when the last collection was made and enables **Download Report**. If no quota file exists from a previous collection request, the system displays no timestamp and disables **Download Report**.

Collect Quotas allows you to generate a report that lists all the quota limits. The system displays the progress of collecting the quotas on the status bar.

After the system collects the quota limits, you can perform the following actions:

- If the number of quota records is 500 or less, use **Display Quotas** to display the records in the GUI.
- If the number of quota records exceeds 500, use **Download Report** to collect the records into a compressed file and to download the file to a location on the client. When the Save As dialog appears in your browser, specify where you want to save the report.

Figure 14 Usage limits report

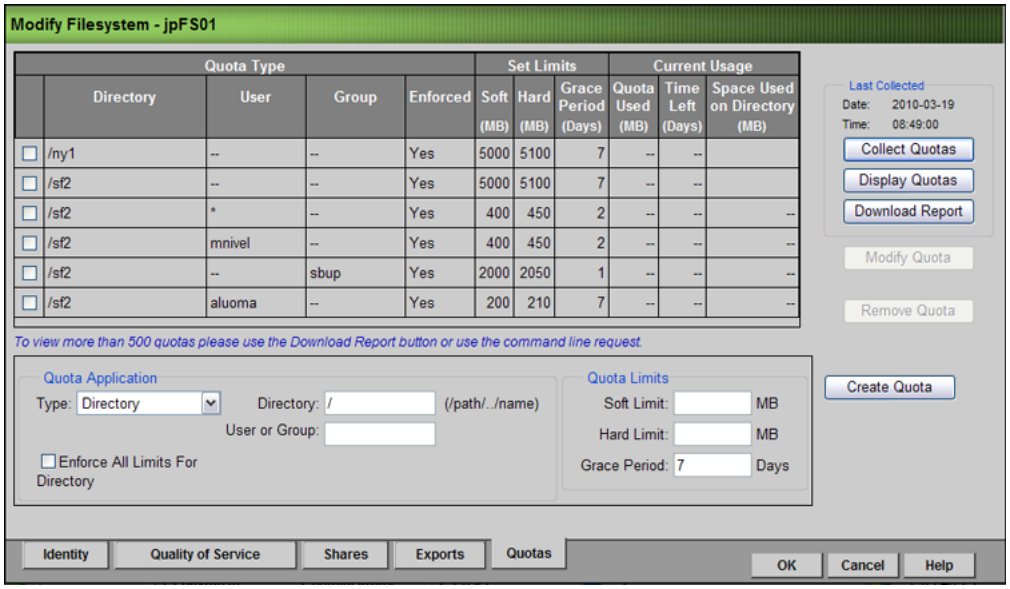

The above illustration shows that:

- The default quota has no usage.
- User `mnivel` has default limits (because no specific limits were set for that user).
- User `aluoma` and group `sbup` have specific limits set.

Add Quotas to a Filesystem

1. Click the **Quotas** tab.
2. Below the **Quota Application** label, select the object that is affected by the quota definition from the **Type** drop-down list:
   - **Directory**, for all users in all directories from the filesystem root.
   - **User**, for a specific user in all directories from the filesystem root.

- **Group**, for a specific group of users in all directories from the filesystem root.
- **Directory & User**, for a specific user in a specific directory.
- **Directory & Group**, for a specific group of users in a specific directory.

Note: You create default limits for groups and users by entering an asterisk (*) in the **User or Group** field and setting the desired soft and hard limit values.

3 Enter values for **Soft Limit** and **Hard Limit** to turn on quota enforcement.

Leave the fields blank to define the values as unlimited. This is functionally equivalent to turning off quota enforcement for the specified limits.

4 Enter an integer that represents a number of days in the **Grace Period** field. A value of 0 (zero) sets an unlimited grace period.

5 Select the **Enforce All Limits For Directory** option (optional).

If you do not select this option, the system does not enforce the quota.

6 Click **Create Quota**.

7 To save the new filesystem, click **OK**.

For parameter definitions, see Filesystem Page, Quotas Tab.

See also:

- About Filesystem Quotas
- Modify Filesystem Attributes

About File Server Creation

A File Server is defined as:

> A network attached storage (NAS) object that is assigned security, network, and protocol access attributes. The attributes apply to all filesystems that are associated with that specific File Server. A Pillar Axiom NAS system requires at least one File Server. Sometimes referred to as a *CIFS server* or *virtual server*.

End users connect to a File Server through a named virtual interface (VIF) so that they can access the filesystems that are associated with that File Server. DNS resolves the named VIF to an IP address, which in turn is linked to one of the VIFs owned by the File Server.

The Primary System Administrator and people who are assigned to the Administrator 1 role can create File Servers. If you are assigned to one of those administrator roles, you might create multiple File Servers because you want:

- Full utilization of all NAS Slammer control units (CUs).

  Tip: We recommend that you assign at least one File Server to each CU, even if they are not used immediately.

- Segmented storage access for different organizations. Assign a different File Server to each organization.

- Separate NFS, CIFS, and multi-protocol filesystems that are associated with dedicated File Servers.

- Different network attributes for filesystems, such as different Domain Name Service (DNS) domains, virtual network interfaces (VIFs), virtual LAN (VLAN) tags, or routes.

  Tip: We recommend that you avoid using VLAN tags. Because you can configure multiple File Servers on the same VLAN that does not use tags, we recommend that you configure your system so that all VIFs are not tagged.

Important! The network switches that are connected to the Pillar Axiom system must support the IEEE 802.1q standard.

In a multiple File Server environment, each File Server need not be on a separate VLAN. Furthermore, File Servers can be on the same VLAN with or without tags. A system administrator can assign VIFs and VLANs to arbitrary File Servers.

If the system administrator enables VLAN tagging, the VLAN ID field must contain a value between 1 and 4094 (inclusive). In this case, in Windows environments, each File Server must exist in a separate security domain, and the VLAN ID identifies a unique security domain.

If you need multiple File Servers, create them before you create filesystems. You can associate a filesystem with a File Server when you create the filesystem; however, you cannot change the association later.

About Network Routes and File Servers

For any File Server, a system administrator can define static and default network routes for communicating with clients on local and remote subnets.

A Slammer control unit (CU) maintains an internal routing table that lists all the MAC addresses of other hosts that are on the same subnet as the CU. How the

Slammer CU routes a network packet depends on whether the specified host can be located in that table.

- When a packet is destined for a host on the same subnet as the Slammer CU, the CU places the MAC address of the target host in the packet header.
- When a packet is destined for a host on a different subnet (a remote host, one that is not in the internal table), the Slammer CU places the MAC address of the default gateway in the packet header. The gateway router will replace that address with the MAC address of the next gateway router along the path to the destination of the packet, and so on.

File Server network routes can be one of the following types:

- **Default route.** A network route to the gateway router through which the Slammer CU can send all non-local traffic (IPs that are not known to the CU). An administrator can define up to eight default routes for a File Server, which provides alternative routes that the File Server's network stack can use if the active route is no longer available. The network stack picks the active default route normally by the order in which the route was defined.

  Note: Due to possible differences in network topologies and delays, there is no route order or preference guarantee. However, when a choice exists, the system will select a static route. When no suitable static route for the target IP is available, the network stack uses the active default route.

  The system configures the active route by choosing the first default route that is usable over one of the active VIFs (virtual interfaces) for that File Server. If the VIF that is used by the active default route fails or is removed, the next default route that matches an active VIF is chosen.

- **Static route.** A network route that is defined for a subnet other than that in which the VIF resides. An administrator can define up to 16 static routes for a File Server. Defining multiple static routes is often done for security or efficiency reasons.

  An administrator can define multiple static routes for the same subnet. In this case, the network stack for the File Server uses the first static route defined. If that route has failed, the network stack uses the next, and so on. If all of the static routes for a given subnet have failed, the network stack uses the default route.

Create File Servers (NAS Storage Systems)

Tip: To utilize all Slammer control units (CUs), Pillar recommends creating at least one File Server for each CU. For example, if you have two NAS Slammers, you should create at least four File Servers.

1. Click the **Storage** icon in the top context pane.
2. Click the **File Servers** link in the left navigation pane.
3. Choose **Create File Server** from the **Actions** drop-down list.
4. Define the network attributes for this File Server.
5. Click **Create** adjacent to the **Additional VIFS** label to create an additional virtual network interface (optional).
   - Enter appropriate values for the virtual interface (VIF).
   - To save the new VIF, click **OK**.
   - Repeat as needed for up to 31 additional interfaces.

   Tip: We recommend that you avoid using VLAN tags. Because you can configure multiple File Servers on the same VLAN that does not use tags, we recommend that you configure your system so that all VIFs are not tagged.
6. Click **Create** adjacent to the **Additional Routes** label to create additional network routes (optional).

   Enter appropriate values for each route. Each route must have a unique set of values for its **Destination** and **Gateway** IP addresses. You can define up to 16 static network routes and up to eight default routes for the File Server.

   Note: Select **Default Route** to identify a given **Gateway** as a default route. The network stack for the File Server uses the active default route when no suitable static route for the target IP is available.
7. Click **OK** to create the File Server.

For parameter definitions, see the following:

- File Server Page, Network Tab
- Virtual Interfaces Page
- Routes Page

Define NIS and NSS Options

1. Click the **NFS** tab.
2. Enter values for the Network Information Service (NIS) domain and up to three NIS servers that you want to associate with this File Server.

3 Click the **Services** tab and select the search order that the Pillar Axiom system uses to resolve hosts and passwords when users try to access the system.

For parameter definitions, see the File Server Page, Services Tab.

Note: If you do not use NIS for host and password resolution, you can upload NIS-equivalent files after the system has successfully created or modified the File Server. However, the Pillar Axiom system must complete File Server initialization before you can upload the files.

See also: Modify File Server Attributes.

Define NFS Options

1 Click the **NFS** tab.

2 Enter values for:

- NFS port number
- Character set
- NFS security settings
- TCP connections (if you support TCP in addition to the default UDP)

For parameter definitions, see the File Server Page, NFS Tab.

Define CIFS Options

1 Click the **CIFS** tab.

2 Enter networking settings for Common Internet File System (CIFS):

- Servers
- Character set
- Opportunistic locking (oplocks)
- CIFS connections over TCP
- CIFS connections with Server Message Block (SMB) Signing

3 Enter values for the CIFS security settings:

- Authentication mode of NTLM (NT LAN Manager) or Active Directory Domain
- Whether Kerberos is to be used for authentication.

4 Enter values for how the File Server accesses the domain controllers:

- Anonymously
- User account

5 Enter values to specify a preferred list of domain controllers (optional).

For parameter definitions, see the File Server Page, CIFS Tab.

Note: After the File Server is successfully created, be sure to join a CIFS domain. Joining a domain registers the File Server with the Windows domain controller. See Join File Servers to CIFS Domains (NAS Systems).

Define Account Mappings

1 Click the **Account Mapping** tab.

2 Indicate whether mapping should be enabled for users that share Common Internet File System (CIFS) and Network File System (NFS) connections (optional).

3 Indicate how mapping should be applied.

Choose one of:

- All security domains.
- A single security domain. Enter the CIFS domain in which the CIFS-to-NFS account mappings are stored.

Note: Account mapping works only when both CIFS and NFS are configured on the File Server.

For parameter definitions, see the File Server Page, Account Mapping Tab.

Review Filesystems Associated With a File Server

1 Click the **Filesystems** tab.

2 Review the filesystem names.

3 To save the new File Server, click **OK**.

For parameter definitions, see the File Server Page, Filesystems Tab.

Join File Servers to CIFS Domains (NAS Systems)

File Servers that support access by Common Internet File System (CIFS) users must join a CIFS domain. The Join CIFS Domain request registers the Pillar Axiom File Server with the Windows domain controller.

When CIFS users try to access filesystems that are associated with the File Server, the Pillar Axiom storage system validates the users against the authentication database that is stored in the Windows domain controller.

1 Click the **Storage** icon in the top context pane.

2 Click the **File Servers** link in the left navigation pane.

3 Select a File Server and choose **Join Windows Domain** from the **Actions** drop-down list.

4 Enter values for the **Administrator** and **Domain Password** fields and click **OK**.

5 Verify that `Yes` appears in the **Joined to CIFS Domain** column after the Pillar Axiom system completes the Join CIFS Domain request.

If `No` appears instead of `Yes`:

- Check the status bar to see if an **Administrator Action Required** icon appears (an exclamation point inside of a yellow triangle).
  - If the icon appears, click it. Perform the suggested corrective action.
  - If the icon does not appear, review entries in the event log and correct the logged errors.
- Run the Join CIFS Domain request again.

For parameter definitions, see the Join Domain Page.

About Volume Groups

Volume groups are organizational units that may contain any grouping of filesystems, LUNs, and nested volume groups.

If you do not create nested volume groups, create all filesystems and LUNs within the default volume group (Volumes) to create a broad, shallow hierarchy.

Figure 15 Default volume group example

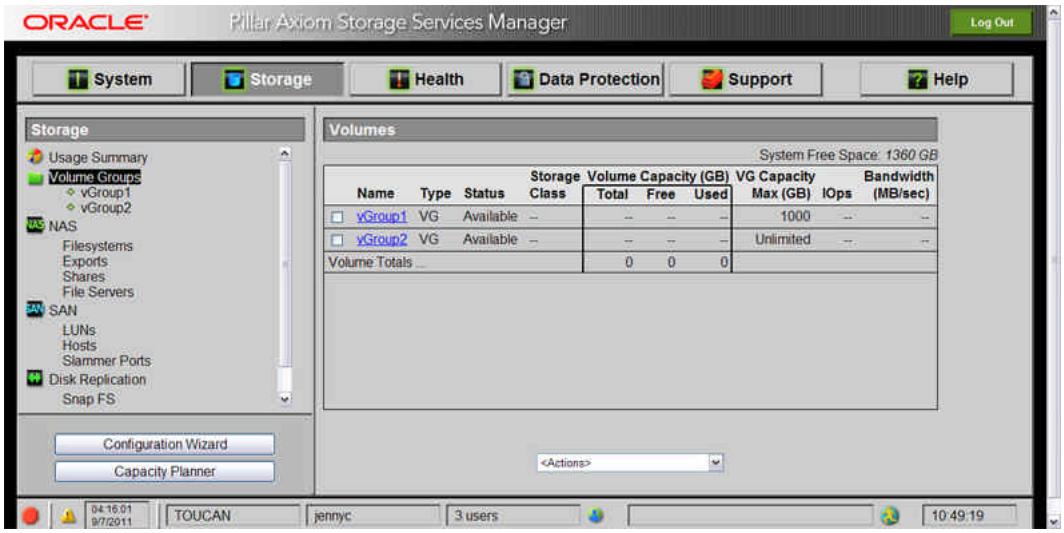

If you create nested volume groups, create filesystems and LUNs within Volumes or a nested volume group to create a narrow, deep hierarchy.

Figure 16 Nested volume groups

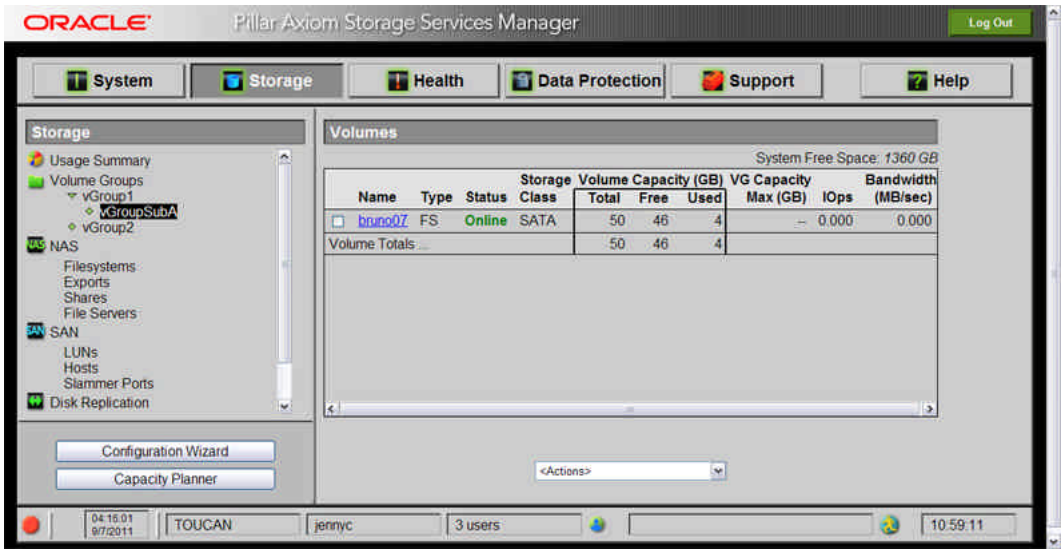

Create Volume Groups

1 Click the **Storage** icon in the top context pane.

2 Click the **Volume Groups** link in the left navigation pane.

3 Choose **Create Volume Group** from the **Actions** drop-down list.

4 Enter a name for the new volume group.

5 Select the volume group location in the drop-down list:

- Select **Volumes** to create a flat structure.
- Select a volume group to create a hierarchical structure of nested volume groups.

6 Select capacity limits for the volume group.

- Select **Unlimited Capacity** to specify that the maximum capacity of associated filesystems, LUNs, and nested volume groups can be increased without constraints.
- Select **Limit Capacity** and enter a value to specify that the maximum capacity of associated objects can be increased up to this volume group constraint.

7 To save the new volume group, click **OK**.

For parameter definitions, see Volume Group Details.

About NFS and CIFS

The ways in which users read and write data files that are stored in a Pillar Axiom storage system are similar to those operations in any networked environment.

These sections explain the operations that are unique to the Pillar Axiom storage environment:

- NFS Protocol Usage
- CIFS Protocol Usage
- Multi-Protocol Usage

Tip: Share this information with the users that you support.

NFS Protocol Usage

A Pillar Axiom system supports network file system (NFS) protocols version 2 and version 3. NFS protocols are explained in detail in these references:

- RFC 1094, NFS: Network File System Protocol Specification, version 2, is located on the Internet Engineering Task Force (IETF) Web site, at www.ietf.org/rfc/rfc1094.txt.
- RFC 1813, NFS Version 3 Protocol Specification, is located online at www.ietf.org/rfc/rfc1813.txt.

Inform NFS users about these points that are specific to Pillar Axiom storage systems:

- Pillar Axiom systems support:
  - NFS versions 2 and 3 over TCP and UDP transports.

    Note: NFS3 using UDP on Solaris is much slower than TCP if the filesystem is mounted read and write.
  - Mount protocol versions 1, 2, and 3 over TCP and UDP.
  - Network Lock Manager (NLM) versions 1, 3, and 4 over TCP and UDP.

  NFS versions 2 and 3 users can share a single filesystem; you do not have to segregate access to a filesystem based on which NFS version is used.

- To get a list of valid mount points, run the following command from an NFS client:

  ```
  showmount -e axiom-IP-addr
  ```

  The command shows the path to a directory with a filesystem-name prefix.

  For example, if a Pillar Axiom administrator creates a filesystem named `fs1`, and a directory named `/dir` under the root of that filesystem, the mount point appears as `/fs1/dir`.

  NFS clients enter ***fileserver-hostname*`:/fs1/dir`** as the NFS mount path.

- Use the *hard* mount option on clients to avoid input and output errors. Use the hard mount option if the network is reliable. For some applications, such as Oracle databases, use the `hard,nointr` mount option.

- NFS file handles remain valid across warm and cold restarts of a NAS Slammer. This means that NFS clients, depending on how the client defines the mount options, generally do not have to remount filesystems after a restart, because their original mounts are still valid.

CIFS Protocol Usage

A Pillar Axiom storage system provides Common Internet File System (CIFS) support for users to access files. Pillar Axiom storage systems follow the Storage Networking Industry Association (SNIA) Common Internet File System (CIFS) Technical Reference (www.snia.org/tech_activities/CIFS/CIFS-TR-1p00_FINAL.pdf).

The CIFS server also follows the standard for NetBIOS over TCP/IP, as defined by these Requests for Comments (RFC):

- RFC 1001, Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods, is located at www.ietf.org/rfc/rfc1001.txt.
- RFC 1002, Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications, is located at www.ietf.org/rfc/rfc1002.txt.

Inform CIFS users about these points that are specific to Pillar Axiom storage systems:

- Where the CIFS Technical Reference document differs from the expectation for Microsoft Windows, the Pillar Axiom implementation of CIFS provides the behavior found on current Microsoft Windows operating systems.

- Pillar Axiom systems support NT access control lists (ACLs), NetBIOS naming service, CIFS browser protocol, and NT LAN Manager (NTLM) 0.12.
- Pillar Axiom systems do not support the following:
  - Windows Distributed File System (DFS)

    Note: A Pillar Axiom system can, however, act as a DFS leaf (the target of a DFS referral).
  - file audit
  - extended attributes

Multi-Protocol Usage

Pillar Axiom network attached storage (NAS) systems provide client access to filesystems through the Network File System (NFS) and Common Internet File System (CIFS) protocols. Administrators configure each File Server to support NFS, CIFS, or both.

When both are configured, the system resolves the NFS and CIFS differences in the multi-protocol filesystems of the File Server.

Table 12 Resolution of multi-protocol differences

Difference in...	Is handled in this way...
Letter case in file names	When a directory contains two or more files with names that differ only in case, the true names of the files are presented to the CIFS client. Typically, the application running on the CIFS client presents these true names to the CIFS user. When performing most CIFS operations, the Pillar Axiom system gives priority to an exact-case match but performs an operation using a file that is not an exact-case match if that is the only file that is available.
Security	Pillar Axiom systems base security for: • CIFS users on a domain controller. • NFS users on a Network Information Service (NIS) database or NIS-alternative files that administrators specify.

Table 12 Resolution of multi-protocol differences (continued)

Difference in...	Is handled in this way...
File locks	During NFS read requests, File Servers ignore CIFS mandatory locks.
	During NFS write requests, File Servers respect CIFS mandatory locks and deny modes.
	Operations that require a shutdown revoke all CIFS opportunity locks (oplocks) first. This strategy allows the completion of any write operations that have been buffered in the CIFS clients.

For complete information regarding multi-protocol implementation in Pillar Axiom systems, refer to the *CIFS and NFS Multi-Protocol Planning Guide*.

CHAPTER 3

Manage Storage Resources

About Storage Management

As a Pillar Axiom administrator, you can modify many attributes of the storage configuration. If more than one administrator modifies the same object at the same time, the simultaneous modifications are not merged. The last changes to be saved take effect.

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

Manage LUNs and SAN Hosts

Display LUN Details

You can display details about all LUNs or you can drill down to the configuration details of a specific LUN.

1. Click the **Storage** icon in the top context pane.
2. Click the **LUNs** link in the left navigation pane.
3. Review the displayed information to ensure that the LUN details are what you expect.

Modify a LUN

You may need to modify the current Quality of Service (QoS) attributes for a LUN, such as increase the capacity or allocate space for Clone LUNs. You can also modify the mapping of a LUN as well as change the Slammer and control unit (CU) to which the LUN is assigned.

If you re-home (move) a LUN from one Slammer to another, the system re-configures the volume at the new location while attempting to maintain the integrity of the data.

⚠ Caution If a client attempts to modify that volume while it is being moved, the client will lose its connection and data may become corrupted or lost. We strongly recommend that, before you re-home a volume, clients unmount the volume to ensure data integrity during the move.

Important! If the LUN is a member of a SAN replication pair, you should isolate the pair before re-homing the LUN to a different Slammer. For more information, see the *Pillar Axiom Replication User's Guide and Reference for SAN*.

When re-homing a LUN from one CU to another on the same Slammer, however, the operation is non-disruptive to client connections and I/O, including operations being performed by Pillar Axiom MaxRep Replication for SAN.

1. Click the **Storage** icon in the top context pane.

2 Click the **LUNs** link in the left navigation pane.

3 Select the LUN that you want to modify.

4 Choose **Modify LUN** from the **Actions** drop-down list.

5 Modify any of the tabbed LUN pages and click **OK**.

For parameter definitions, see the following:

- LUN, Identity Tab
- LUN, Quality of Service Tab
- LUN, Mapping Tab
- LUN, Host Connections Tab

Copy LUNs

You can copy an existing LUN and give the new LUN different Quality of Service (QoS) metrics. This copying allows system resources to be maximized for the task at hand. For example, a copied volume that is used for reporting is assigned a lower performance priority and a higher read-centric access pattern than would the source volume.

You can also copy a Clone LUN. Such a copy will always depend on its source LUN. Also, you cannot choose different QoS attributes for a Clone LUN.

Copy a Clone LUN when you want to test a new application on an exact copy instead of on the original LUN.

Copy a LUN when you need a new LUN with the same starting data as an existing LUN.

Another reason to create a clone or a copy is to preserve a point-in-time view of the data. If you create a clone for this purpose, at a later time you can restore the data to the source LUN.

Unlike a clone, the new blocks for the copy may be on a different set or even a different type of Brick. In other words, a volume copy of a Fibre Channel or solid state drive (SSD) based, premium priority LUN may be created in the low-priority band on SATA Bricks.

1 Click the **Storage** icon in the top context pane.

2 Click the **LUNs** link in the left navigation pane.

3 Select a LUN, and choose **Copy LUN** from the **Actions** drop-down list.

4 Name the LUN copy.

5 Click the **Quality of Service** tab to define the QoS attributes of the new LUN (optional).

6 Click **Optimizer** to see the effect that these settings would have on the amount of capacity (GB) that is available based on the Storage Class and redundancy settings (optional).

7 To save the new LUN, click **OK**.

For parameter definitions, see the LUN, Quality of Service Tab.

Delete Clone LUNs

If you delete a Clone LUN that is the parent or source of other clones, the child clones will not be deleted.

1 Click the **Storage** icon in the top context pane.

2 Click the **Clone LUN** link in the left navigation pane.

3 Select one or more Clone LUNs and choose **Delete Clone LUN** from the **Actions** drop-down list.

4 When prompted to confirm the deletion, click **OK** to delete the selected Clone LUNs.

Delete LUNs

If you need to delete an existing LUN, you can do so if the LUN is not being accessed by users.

Note: When you delete a LUN that is a parent or source for Clone LUNs, all child clones are deleted as well.

1 Click the **Storage** icon in the top context pane.

2 Click the **LUNs** link in the left navigation pane.

3 Select one or more LUNs and choose **Delete LUN** from the **Actions** drop-down list.

4 When prompted to confirm the deletion, click **OK** to delete the LUN.

Download Pillar Axiom Virtual Disk Service Provider

Pillar Axiom Virtual Disk Service (VDS) Provider plug-in allows you to use the Disk Administrator on a Windows 2003 server to configure and manage LUNs on a Pillar Axiom storage system.

Prerequisites:

- System serial number
- Login account: username
- Login account: password

The installer allows the configuration of a single Pillar Axiom system at the time of installation. Additional systems may be configured from a command line tool. A default location for the installation is presented on the install screen, but this can be changed during installation.

1 Click the **Support** icon in the top context pane.

2 Click the **Virtual Disk Service** link under the Utilities Download heading in the left navigation pane.

3 Select **Download VDS for Windows 2003/2008** from the **Actions** drop-down menu.

4 Follow the installation prompts.

5 Restart the Windows server.

Restarting permits the `diskraid` utility to see the VDS providers.

Once the installation completes, you can verify it by running `diskraid.exe` at a command prompt and issuing the command `ListProviders` within `diskraid`.

For the Pillar Axiom VDS Provider to be able to manage a Pillar Axiom system, it most be connected using Fibre Channel. Be sure to register it by using the `registerAxiom.exe` tool available in the `bin` folder in the installation directory.

This registration tool has two functions, add and remove registry entries. Running `registerAxiom.exe` prints the usage directions.

To add a registry entry:

`registerAxiom.exe` *sample-serial user-password*

To remove a registry entry:

`registerAxiom.exe` ***sample-serial***

For parameter definitions, see the Virtual Disk Service (VDS) Page.

Display SAN Host Settings

You can display details about the Pillar Axiom Path Manager SAN host driver settings and LUN connection status as well as configure iSCSI settings.

1. Click the **Storage** icon in the top context pane.
2. Click the **Hosts** link in the left navigation pane.
3. Click the name of a host and select **View Host Settings** from the **Actions** drop-down list.
4. Review the displayed information to ensure that the SAN host details are what you expect.

For parameter definitions, see the following:

- Host Settings, Identity Tab
- Host Information, Configure iSCSI Tab
- Host Information, LUN Connections Tab
- Host Information, Settings Tab

Modify SAN Host Settings

You can change the LUN settings of the SAN host drivers only when the drivers are installed and can communicate with the Pillar Axiom storage system.

Also, if you have SAN hosts that access the LUNs using HP-UX initiator ports and HP HBAs, you can enable the HP-UX option. When this option is enabled, the system determines LUN numbers using the HP-UX addressing scheme, allowing up to 255 LUNs. Also when enabled, the host cannot have a visible LUN using ID 0.

1. Click the **Storage** icon in the top context pane.
2. Click the **Hosts** link in the left navigation pane.
3. Click the name of a host and select **Modify Host Settings** from the **Actions** drop-down list.

4 Make any necessary changes in the tabbed host settings pages and click **OK**.

For parameter definitions, see the following:

- Host Settings, Identity Tab
- Host Information, Configure iSCSI Tab
- Host Information, LUN Connections Tab
- Host Information, Settings Tab

Delete SAN Host Names

If you need to delete an existing SAN host, you can do so only if that host is not connected to the network.

1 Click the **Storage** icon in the top context pane.

2 Click the **Hosts** link in the left navigation pane.

3 Select one or more hosts and choose **Delete Host Name** from the **Actions** drop-down list.

4 When prompted to confirm the deletion, click **OK**.

For parameter definitions, see the Host Settings, Identity Tab.

Associate Hosts

You can create a host-to-HBA association when you do not have the Pillar Axiom Path Manager driver installed on the SAN host. You can do this for hosts that are listed as *unknown* and are referenced by the WWN (for Fibre Channel hosts) or IQN (for iSCSI hosts) of their HBA.

1 Click the **Storage** icon in the top context pane.

2 Click the **Hosts** link in the left navigation pane.

3 Select a SAN host and select **Associate a Host** from the **Actions** drop-down list.

4 Enter a name for the host.

5 Select or add an HBA to use from the WWNs or iSCSI list that is not yet detected by the system.

6 Enter the authentication settings for the specific host (optional).

7 Click **Create** to make the association.

8 Click **OK**.

For parameter definitions, see the Associate Hosts Page.

Modify iSCSI Port Settings

The default for all iSCSI ports is to receive the IP address, subnet mask, and gateway using Dynamic Host Configuration Protocol (DHCP). If you want to manually assign these values, modify each port with the values that you want to use.

1 Click the **Storage** icon in the top context pane.

2 Click the **Slammer Ports** link in the left navigation pane.

3 Click the name of the Slammer and select **Modify iSCSI Port Settings** from the **Actions** drop-down list.

4 Make any necessary changes in the **Settings** tab and click **OK**.

For parameter definitions, see the iSCSI Port Settings Page.

Manage Filesystems

Display Filesystem Details

You can display details about all filesystems or you can drill down to the configuration details of a specific filesystem.

1. Click the **Storage** icon in the top context pane.
2. Click the **Filesystems** link in the left navigation pane.
3. Review the displayed information to ensure that the filesystem details are what you expect.
4. For additional details of a particular filesystem, click the name of the filesystem.
5. Navigate through the tabbed filesystem pages to review the configuration details.

For parameter definitions, see the Filesystem Overview Page.

Display Capacity Usage

At any time, you can display the actual capacity usage of a filesystem and compare that usage to the total system capacity and assigned capacity limits.

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

1. Click **Storage** in the top context pane.
2. Click **Usage Summary** in the left navigation pane.

   In addition to displaying summary usage for the entire system and by type of logical volume, this screen displays a thermometer-style graphic that shows usage by Storage Class.

3 Review the displayed information to ensure that the capacity usage is what you expect.

Figure 17 Usage summary

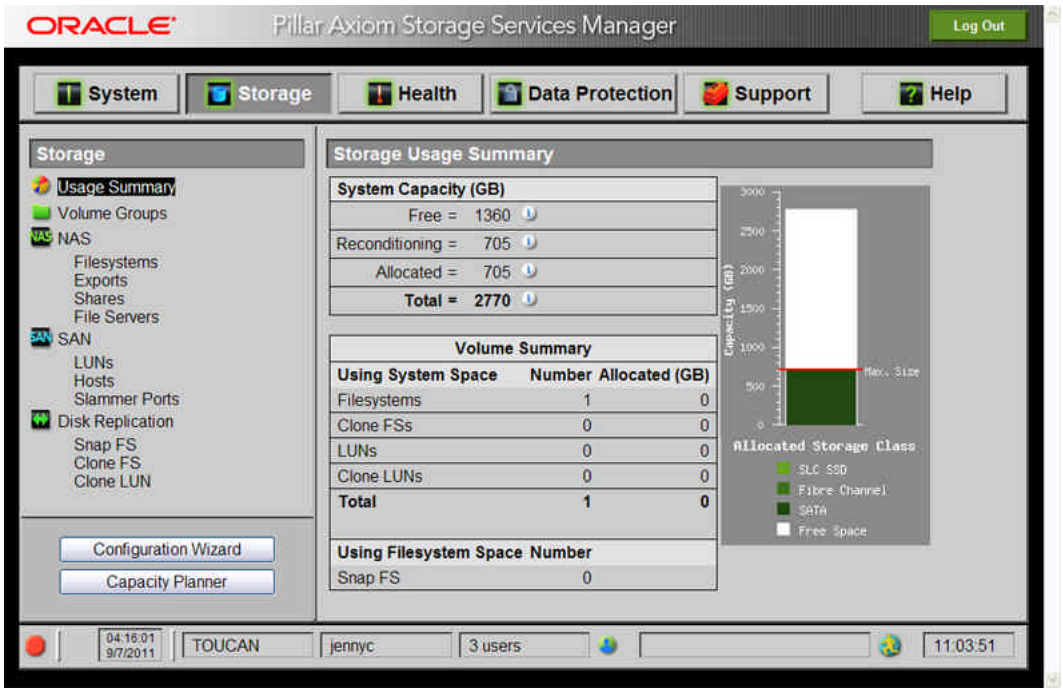

For parameter definitions, see the Storage Usage Summary Page.

Locate a Filesystem on a Slammer

You can physically locate a filesystem on a NAS Slammer based on the File Server that is associated with the filesystem.

Tip: To improve performance, Pillar recommends that you place a filesystem on the same Slammer control unit as the virtual interfaces (VIFs) that will be used to access the filesystem.

Moving a filesystem to a particular Slammer control unit (CU) provides as well the flexibility to balance the load on the system as data resources expand.

1 Click the **Storage** icon in the top context pane.

2 Click the **File Servers** link in the left navigation pane.

3 Click the name of the File Server that is associated with the filesystem.

4 Click the **Filesystem** tab.

5 Review the page to identify the Slammer where the filesystem resides.

For parameter definitions, see the File Server Page, Filesystems Tab.

Modify Filesystem Attributes

Circumstances can arise after you have created a filesystem that require you to change some of its attributes.

For example, you made need to change one or more of these attributes of a filesystem:

- Its name.
- The volume group to which it is assigned.
- The Slammer control unit to which it is assigned.

  If you re-home (move) a filesystem from one Slammer control unit to another, the system re-configures the volume at the new location while attempting to maintain the integrity of the data.

  Caution If a client attempts to modify that volume while it is being moved, the client will lose its connection and data may become corrupted or lost. We strongly recommend that, before you re-home a volume, clients unmount the volume to ensure data integrity during the move.

- Its Quality of Service (QoS) attributes.

  Changing redundancy or a QoS attribute causes data migration for the filesystem, which can take an extended time to complete during which system performance may be affected.

- Its write-once, read-many (WORM) capability.

  An existing WORM filesystem can be upgraded from Standard retention to Compliance.

  Note: Normally, Pillar Axiom SecureWORMfs instances cannot be changed from Compliance to Standard. However, if you need to downgrade a Compliance instance to Standard, contact the Pillar World Wide Customer Support Center for assistance.

Tip: To improve performance, Pillar recommends that you place a filesystem on the same Slammer control unit as the virtual interfaces (VIFs) that will be used to access the filesystem.

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

1 Click the **Storage** icon in the top context pane.

2 Click the **Filesystems** link in the left navigation pane.

3 Select the filesystem that you want to modify.

4 Choose **Modify Filesystem** from the **Actions** drop-down list.

For parameter definitions, see Filesystem Page, Identity Tab.

Modify Quality of Service (QoS) Attributes of a Filesystem

1 Click the **Quality of Service** tab.

2 Enter values for the attributes that you want to modify.

For parameter definitions, see Filesystem Page, Quality of Service Tab.

Modify Filesystem Retention Policy

1 Click the **Retention Policy** tab.

2 Enter values for the master retention period that you want to modify.

Note: The following rules apply when modifying the retention period of a Pillar Axiom SecureWORMfs filesystem:

- **Default**: Can be decreased and increased as long as it falls between the minimum and maximum ranges.
- **Minimum**: Can be decreased but not increased.
- **Maximum**: Can be increased but not decreased.

For parameter definitions, see Filesystem Page, Retention Policy Tab.

Modify CIFS Shares

1 Click the **Shares** tab.

2 Select the share that you want to modify and click **Modify**.

3 Modify the share name and path, or share comment.

4 To enable the share point immediately, select **Enabled**.

5 Click **Update** to use the modified share parameters.

For parameter definitions, see Filesystem Page, Shares Tab.

Modify NFS Exports

You can modify an NFS export by toggling its *root access* parameter, *read only* access parameter, or both. You can also change the collection of hosts that can access the export.

Note: To change the export path, you must delete the export and recreate a new export with the desired path settings.

1 Click the **Exports** tab.

2 Select the export that you want to modify and click **Modify**.

3 To make the export available to users in read-only mode, select **Read Only**.

4 To allow users as root users on the export, select **Root Access**.

5 To provide access to additional NFS clients, click **Additional Hosts**.

Select the type of host access:

- **All Hosts**: Everyone can access the export point.
- **Single Host**: Only the specified host can access the export point.
- **NIS Netgroup**: Everyone within the Network Information Service (NIS) netgroup can access the export point.
- **Network**: Everyone on the specified subnet can access the export point.

Select one or more of the following access modes (optional):

- **Read Only**
- **Root Access**

Select the processing order.

Click **Create** and then **OK**.

6 Click **Update**.

7 Click **OK** on the NFS Exports page.

8 From the Network File System (NFS) client, mount the filesystem remotely:

- Mount the exported filesystem on a NFS client.
- Change directories to the mount point.
- Create the directories that you want to export to other NFS clients under the `root` export point.
- Change the directory permissions to the required values (may require root access).
- Unmount the `/` export point (optional).

9 From the Pillar Axiom Storage Services Manager NAS Exports page, delete the `/` export point (optional).

10 From the Pillar Axiom Storage Services Manager NFS Export page, export the directories that you created under the `root` export point.

- Enter the export parameters:
  - A full path name that starts with a forward slash (/)and does not include the filesystem name.
  - The user ID (UID) for anonymous users. Set the UID for anonymous users to zero (0) if you want all data access to be as if from a root user.
- Select the host type and enter the associated host value:
  - **All Hosts**: Everyone can access the export point.
  - **Single Host**: Only the specified host can access the export point.
  - **NIS Netgroup**: Everyone within the NIS netgroup can access the export point.
  - **Network**: Everyone on the specified subnet can access the export point.

11 Click **OK** on the NFS Exports page.

For parameter definitions, see Filesystem Page, Exports Tab.

Modify a Filesystem Quota

1 Click the **Quotas** tab.

2 Click **Collect Quotas**.

Note: If there are more than 500 quotas, you need to download the list of quotas to the client system to view, using **Download Report**. Once you download the report, unzip the file (using the standard `bunzip2` command) and view the list on the client. If you want to modify a quota, you need to know the fully qualified name (FQN) of the quota and use the `pdscli` utility to modify the quota (see *Pillar Axiom CLI Reference Guide*).

3 Click **Display Quotas**.

4 Select the quota that you want to modify.

5 Select the object that is affected by the quota definition from the **Quota Application** drop-down list:

- **Directory**, for all users in all directories from the filesystem root.
- **User**, for a specific user in all directories from the filesystem root.
- **Group**, for a specific group of users in all directories from the filesystem root.
- **Directory & User**, for a specific user in a specific directory.
- **Directory & Group**, for a specific group of users in a specific directory.

6 Enter values for **Soft Limit** and **Hard Limit** to turn on quota enforcement.

Leave the fields blank to define the values as unlimited. This is functionally equivalent to turning off quota enforcement for the specified limits.

7 Enter an integer that represents a number of days in the **Grace Period** field. A value of 0 (zero) sets an unlimited grace period.

8 Click **Update Quota**.

For parameter definitions, see Filesystem Page, Quotas Tab.

Take Filesystems Offline

1 Click the **Storage** icon in the top context pane.

2 Click the **Filesystems** link in the left navigation pane.

3 Select the filesystem that you want to take offline.

4 Choose **Take Filesystem Offline** from the **Actions** drop-down list.

5 When prompted to confirm the action, click **OK**.

While the filesystem is offline, you cannot access the data.

Put Filesystems Online

You can use this option to force a filesystem back online to do a restore from backup if the filesystem went offline.

1 Click the **Storage** icon in the top context pane.

2 Click the **Filesystems** link in the left navigation pane.

3 Select the filesystem that you want to force online.

4 Choose **Put Filesystem Online** from the **Actions** drop-down list.

5 When prompted to confirm the action, click **OK**.

Copy Filesystems

You can copy an existing filesystem to create a new filesystem that has its own Quality of Service (QoS) metrics. Copying allows system resources to be maximized for the task at hand. For example, a copied volume that is used for reporting is assigned a lower performance priority and a higher read-centric access pattern than would the source volume.

You can also copy a Clone FS. Such a copy will always depend on its source filesystem. Also, you cannot choose different QoS attributes for a Clone FS.

Copy a Clone FS when you want to test a new application on an exact copy instead of on the original filesystem.

Copy a filesystem when you need a new filesystem with the same starting data as an existing filesystem.

Another reason to create a clone or a copy is to preserve a point-in-time view of the data. If you create a clone for this purpose, at a later time you can restore the data to the source filesystem.

Unlike a Clone FS, the new blocks for the copy may be on a different set or even a different type of Brick. In other words, a volume copy of a Fibre Channel or solid state drive (SSD) based, premium priority filesystem may be created in the low-priority band on SATA Bricks.

1. Click the **Storage** icon in the top context pane.
2. Click the **Filesystems** (or **Clone FS**) link in the left navigation pane.
3. Select a filesystem (or Clone FS) and choose **Copy Filesystem** (or **Copy Clone FS**) from the **Actions** drop-down list.
4. Enter a name that is meaningful for the copied filesystem (or Clone FS).
5. Click the **Quality of Service** tab to define the Quality of Service (QoS) attributes of the new filesystem.
6. Click **Optimizer** to see the effect that these settings would have on the amount of capacity (GB) that is available based on the Storage Class and redundancy settings (optional).

   Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.
7. To save the new filesystem, click **OK**.

The NFS exports and CIFS shares associated with the original filesystem are copied. Exports on the new filesystem are mountable from another host. Shares are renamed using the original name plus a number so that they can be on the same File Server.

Delete Clone FSs

If you delete a Clone FS that is the parent or source of other clones, the child clones will not be deleted.

1. Click the **Storage** icon in the top context pane.
2. Click the **Clone FS** link in the left navigation pane.
3. Select one or more Clone FSs and choose **Delete Clone FS** from the **Actions** drop-down list.
4. When prompted to confirm the deletion, click **OK** to delete the selected Clone FSs.

Delete Filesystems

If you need to delete an existing filesystem (including any empty Pillar Axiom SecureWORMfs filesystem), you can do so if the filesystem is empty and is not being accessed by users.

Note: You cannot delete a Pillar Axiom SecureWORMfs Compliance filesystem if it has protected files on it. To delete a non-empty Compliance Pillar Axiom SecureWORMfs filesystem, you must first downgrade it to Standard.

An existing WORM filesystem can be upgraded from Standard retention to Compliance.

Note: Normally, Pillar Axiom SecureWORMfs instances cannot be changed from Compliance to Standard. However, if you need to downgrade a Compliance instance to Standard, contact the Pillar World Wide Customer Support Center for assistance.

When you delete a filesystem that is a parent or source for Clone FSs, all child clones are deleted as well.

1. Click the **Storage** icon in the top context pane.
2. Click the **Filesystems** link in the left navigation pane.
3. Select one or more filesystems and choose **Delete Filesystem** from the **Actions** drop-down list.
4. When prompted to confirm the deletion, click **OK** to delete the filesystem.

Toggle the Pillar Axiom SecureWORMfs File Deletion Setting

This option is enabled by default. When this option is enabled, any unprotected files as well as protected files that have expired can be deleted from the Pillar Axiom SecureWORMfs filesystem.

Disable this option to prohibit the deletion of any files (unprotected or expired). Also, you cannot delete the filesystem when you disable this option. You might want to do this to maintain the files on the filesystem for auditing purposes.

1. Click the **Storage** icon in the top context pane.
2. Click the **Filesystems** link in the left navigation pane.

3 Select the Pillar Axiom SecureWORMfs filesystem that has the files you want to delete or protect.

4 Choose **Enable/Disable File Deletion on a SecureWORMfs** from the **Actions** drop-down list.

Validate File Integrity on a Pillar Axiom SecureWORMfs

You can use this option to verify the data integrity of protected files on a Pillar Axiom SecureWORMfs filesystem.

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

Note: For a given Pillar Axiom SecureWORMfs, the system runs only one validation scan at a time. These scans can be initiated automatically by the system or manually by the administrator. The system queues all subsequent scans.

1 Click the **Storage** icon in the top context pane.

2 Click the **Filesystems** link in the left navigation pane.

3 Select the SecureWORMfs filesystem that you want to verify.

4 Choose **Perform Protected File Integrity Scan** from the **Actions** drop-down list.

Tip: The scan may degrade system performance, so schedule during off-peak hours.

5 Enter the path of the files you want to verify. If you do not enter a path, the system scans all files on the Pillar Axiom SecureWORMfs.

6 When prompted to confirm the action, click **OK**.

Modify Pillar Axiom SecureWORMfs Extended Attributes

About Pillar Axiom SecureWORMfs Extended Attributes

Each directory on a Pillar Axiom SecureWORMfs instance contains a hidden directory named `.attributes`. This directory contains directory entries corresponding to user directory entries.

For every user file on a Pillar Axiom SecureWORMfs instance, a user can access the extended WORM attributes by accessing the contents of the corresponding file in the `.attributes` directory. The file in the `.attributes` directory is the attribute view of the corresponding user file. For example, the `/worm/.attributes/foo.txt.XML` document provides the attribute view for `/worm/foo.txt`.

The attribute view is exported and managed as an XML document. The attribute view contains the following attributes for each file:

- Immutability status.
- Expiration status.
- Data consistency status.
- Creation time (in number of seconds elapsed since *epoch*, which is the time 00:00:00 UTC on January 1, 1970).
- Time when the file is made immutable (in number of seconds elapsed since epoch).
- Expiration time (in number of seconds elapsed since epoch).
- Time of the last successful data integrity check (in number of seconds elapsed since epoch).
- Internal data hash type.
- Internal data hash value.
- Application data hash value.

Pillar Axiom SecureWORMfs uses the Secure Hash Algorithm (SHA)-1 to validate file data. The application hash is a scratch area for use by applications to store application-specific hash of the file data. The Pillar Axiom storage system does not use the application hash in any manner.

Modification of Extended Attributes

Some of the attributes of a Pillar Axiom SecureWORMfs instance can be modified by editing an XML file.

The format of the XML document is shown below:

```
<?xml version="1.0"?>
<ATTRIBUTES>
     <PDSVERSION>2.0</PDSVERSION>
     <Immutable>1</Immutable>
     <Expired>0</Expired>
     <RecordIntegrity>1</RecordIntegrity>
     <CreationTime>1126157232</CreationTime>
     <ImmutableTime>1126157999</ImmutableTime>
     <ExpirationTime>1187157999</ExpirationTime>
     <IntegrityCheckTime>1126198001</IntegrityCheckTime>
     <InternalHashType>SHA-1</InternalHashType>
     <InternalHashValue>
          0x123456789FF98765432112345678 9FF987654321
     </InternalHashValue>
     <ApplicationHashValue>
          0x098765432ABC34567890098765432ABC34567890
     </ApplicationHashValue>
</ATTRIBUTES>
```

Note that before a file is made immutable, values of some attributes are meaningless. The attribute view of a mutable file is shown below with the defaults highlighted or empty:

```
<?xml version="1.0"?>
<ATTRIBUTES>
     <PDSVERSION>2.0</PDSVERSION>
     <Immutable>0</Immutable>
     <Expired>0</Expired>
     <RecordIntegrity>1</RecordIntegrity>
     <CreationTime>1126157232</CreationTime>
     <ImmutableTime></ImmutableTime>
     <ExpirationTime></ExpirationTime>
     <IntegrityCheckTime></IntegrityCheckTime>
     <InternalHashType>None</InternalHashType>
     <InternalHashValue></InternalHashValue>
     <ApplicationHashValue></ApplicationHashValue>
</ATTRIBUTES>
```

Of all the attributes available in the attribute view, only the following may be modified by users or applications with appropriate access rights:

Immutable	This field can only be modified if the file is mutable. Setting the value of this field to 1 makes the file immutable. Once the file is made immutable, this field cannot be modified again.

ExpirationTime This field can be modified at any time. The value is a timestamp that represents a date and time in the future at which point the file expires and becomes mutable. This timestamp must fall within the range specified by the minimum and maximum retention periods defined on the filesystem instance. Before a file is made immutable, expiration time may be decreased, increased, or erased. However, once a file becomes immutable, the expiration time may only be increased.

Note: Expiration time is a 32-bit value. Setting the value to 4294967295 (0xFFFFFFFF) makes a file permanently immutable. In addition, if the expiration time derived from Default Retention Period cannot be represented as a 32-bit value, the value is set to 4294967295.

ApplicationHashValue This field can only be modified as long as the file is mutable. Once the file is made immutable, this field cannot be modified again.

It is possible to write to the un-modifiable fields as long as the values being written are the same as the currently defined values. To modify one or more fields, an application or a user must supply a valid XML document. For example, to update the application hash and the immutability status of a file, a user supplies the following XML document to the Pillar Axiom storage system:

```
<?xml version="1.0"?>
<ATTRIBUTES>
     <PDSVERSION>2.0</PDSVERSION>
     <Immutable>1</Immutable>
     <ApplicationHashValue>
          0x098765432ABC34567890098765432ABC34567890
     </ApplicationHashValue>
</ATTRIBUTES>
```

The fields printed in bold above are required in each XML document supplied to the Pillar Axiom storage system, whereas the fields in italics are optional. It is required that the entire XML document be supplied in a single write.

Manage File Servers

Display File Server Details

You can display details about all File Servers or you can drill down to the configuration details of a specific File Server.

1. Click the **Storage** icon in the top context pane.
2. Click the **File Servers** link in the left navigation pane.
3. Review the displayed information to ensure that the File Server details are what you expect.
4. For additional details of a particular File Server, click the name of the File Server.
5. Navigate through the tabbed File Server pages to review the configuration details.

For parameter definitions, see the File Server Overview Page.

Duplicate File Server

When you need to create a new File Server quickly, you can duplicate an existing one and modify those attributes that are different.

1. Click the **Storage** icon in the top context pane.
2. Click the **File Servers** link in the left navigation pane.
3. Select the File Server that you want to duuplicate.
4. Choose **Duplicate File Server** from the **Actions** drop-down list.
5. Enter the name of the new File Server.
6. Enter values for the attributes that you want to modify on any or all of the tabbed File Server pages.
7. Click **OK** to create the new File Server.

For parameter definitions, see the following:

- File Server Page, Network Tab
- File Server Page, NFS Tab
- File Server Page, CIFS Tab
- File Server Page, Filesystems Tab
- File Server Page, Account Mapping Tab

Modify File Server Attributes

Modify a File Server to change any one of its properties.

For example, you can modify any of the following properties of a File Server:

- Name
- Attributes of its virtual interfaces (VIFs)
- Default and static routes
- Domain Name System (DNS) settings
- Host name resolution search order
- Common Internet File System (CIFS) and Network File System (NFS) configurations
- Account mapping

Tip: On occasion, you may want to rebalance your Pillar Axiom storage system to increase performance.

To do this, you can move File Servers (the virtual network interfaces, or VIFs, and filesystems) from one Slammer control unit to another Slammer control unit. Rebalancing a system might be prudent when you add an additional NAS Slammer to an existing system as well.

1. Click the **Storage** icon in the top context pane.
2. Click the **File Servers** link in the left navigation pane.
3. Click the name of the File Server that you want to modify.
4. Enter values for the attributes that you want to modify on any or all of the tabbed File Server pages.

For parameter definitions, see the following:

- File Server Page, Network Tab
- File Server Page, NFS Tab
- File Server Page, CIFS Tab
- File Server Page, Filesystems Tab
- File Server Page, Account Mapping Tab

Upload NIS alternative Files

You can upload Network Information Service (NIS) alternative files if you do not use NIS lookup to resolve hosts and passwords.

1. Click the **Services** tab.
2. Click **Upload**.
3. Navigate to and select all of the `/etc/passwd`, `/etc/group`, and `/etc/netgroup` files from your local system.
4. To save the modified File Server, click **OK**.

For parameter definitions, see the File Server Page, Services Tab.

Recover NLM Locks

When you suspect that the Network File System (NFS) lock records that are maintained by a Pillar Axiom storage system are incorrect, you can recover the locks maintained by the network lock manager (NLM).

The most common reason these lock records become invalid is when an NFS client has obtained a lock from the Pillar Axiom system and then crashes before this lock could be released. Normally, when an NFS client crashes, it restarts and releases any locks that were held before the crash. If the NFS client is broken in some way, it may not be able to restart. In this case, the Pillar Axiom system will not release the locks.

When the administrator performs NLM lock recovery, the Pillar Axiom system initiates a dialog with each of the NFS clients that it believes might be holding NLM locks. The system asks each of these clients to reacquire the locks it currently holds. At the end of the recovery period (approximately 45 seconds), the system identifies the NFS locks that were held at the start of recovery that have not been reacquired. These NFS locks are then released.

Tip: Under all normal conditions, recovering NLM locks will be successful. An administrator, however, should not perform this operation when the network is having major issues. For example, if a network cable is being replaced, the lock recovery request could potentially result in the Pillar Axiom system being unable to inform an NFS client that the client needs to reacquire the NFS locks it is holding.

1. Click the **Storage** icon in the top context pane.
2. Click the **File Servers** link in the left navigation pane.
3. Select the File Server for which you want to recover the NLM locks.
4. Choose **Recover NLM Locks** from the **Actions** drop-down list.
5. In the confirmation dialog, click **OK**.

Delete File Servers

At times, you may need to delete an existing File Server. You may want to do this when its associated filesystems have been deleted and you no longer have use for the File Server.

If a File Server is associated with any filesystems, you cannot delete the File Server. Those associated filesystems would be unusable, because you cannot change an existing association that a filesystem has with a File Server. You create the association when you create a filesystem, but you cannot modify the association at a later time.

1. Click the **Storage** icon in the top context pane.
2. Click the **File Servers** link in the left navigation pane.
3. Select one or more File Servers and choose **Delete File Server** from the **Actions** drop-down list.
4. When prompted to confirm the deletion, click **OK** to delete the File Server.

Manage Volume Groups

Display Volume Group Details

You can display details about all volume groups or you can drill down to the configuration details of a specific volume group.

1. Click the **Storage** icon in the top context pane.
2. Click the **Volume Groups** link in the left navigation pane.
3. Review the displayed information to ensure that the volume group details are what you expect.

   The content pane displays capacity information for all logical volumes (filesystems and LUNs) and nested groups contained in a volume group.
4. Click the name of a volume group or of a logical volume that is contained in the volume group.
5. Review the configuration details.

For parameter definitions, see the Volume Groups Overview Page.

Modify Volume Group Attributes

At times, you may want to modify certain attributes of a volume group. For example, you may want to nest one volume group within another.

1. Click the **Storage** icon in the top context pane.
2. Click the **Volume Groups** link in the left navigation pane.

   If you want to change the attributes of a nested volume group, click on the *parent* volume group in the navigation pane.
3. In the Volumes content pane, click the name of the volume group that you want to modify.

   Tip: Volume groups have a Type attribute of VG.
4. Enter new attribute values as needed.

5 To save the modified volume group, click **OK**.

For parameter definitions, see Volume Group Details.

See Also: Move a Logical Volume to a Different Volume Group.

About Moving Logical Volumes to Different Volume Groups

You can break the association between a logical volume (filesystem or LUN) and a volume group.

To do so, move the logical volume to a different volume group, which associates the logical volume with the new volume group. You may have to create additional volume groups.

You can:

- Add more volume groups to your current organizational model and move one or more logical volumes into the new volume group.

  For example, if your current organizational model is based on location and your company recently opened a sales office in Tokyo, create a new Japan volume group and move appropriate logical volumes into it.

- Create a new organizational model and move all logical volumes into the new volume groups.

  For example, if your current organizational model is based on location and you want to reorganize based on corporate structure, create new departmental volume groups. Move logical volumes from the volume groups that are named for countries into the volume groups that are named for departments.

Move a Logical Volume to a Different Volume Group

1 Click the **Storage** icon in the top context pane.

2 Click the **Volume Groups** link in the left navigation pane.

3 Select one or more logical volumes (filesystems and LUNs) or nested volume groups from the list of volumes.

4 Choose **Move {Selected}** from the **Actions** drop-down list.

5 Select a single volume group in the Destination box on the **Move Volumes** page.

6 Click **OK** to move the selected items to another volume group.

For parameter definitions, see the Move Volumes Page.

Delete Volume Groups

You can delete a volume group after you have reassigned all of its filesystems and LUNs to different volume groups.

1 Click the **Storage** icon in the top context pane.

2 Click the **Volume Groups** link in the navigation pane.

3 Select one or more volume groups, and choose **Delete** from the **Actions** drop-down list.

If a volume group contains any objects, move or delete those objects before you delete the volume group.

4 When prompted to confirm the deletion, click **OK** to delete the volume groups.

Manage System Tasks

A system task is an action performed by the Pilot in response to a user request or system behavior.

The topics that are covered in this section include:

- Display Task Progress.
- Cancel Tasks.

Display Task Progress

You can review the status of Pillar Axiom system tasks and background processes that are in progress.

1. On any Pillar Axiom Storage Services Manager page, click the **Task in Progress** icon in the bottom right of the task bar.

   The Pillar Axiom Storage Services Manager displays the status of any tasks currently in progress.

2. Click **Close**.

You can also Display the Event Log to review completed system tasks.

Cancel Tasks

1. On any Pillar Axiom Storage Services Manager page, click the **Task in Progress** icon in the bottom right of the task bar.

   Result:
   The Pillar Axiom Storage Services Manager displays the status of any tasks currently in progress.

2. Select the task that you want to cancel and choose **Cancel Task** from the **Actions** drop-down list.

   If a task cannot be cancelled, the Pillar Axiom Storage Services Manager displays a message that informs you that the task cannot be cancelled.

3. Click **Close**.

CHAPTER 4

Manage Backups and Snapshots

About NDMP Backup Management

If you have a data management application (DMA) that uses Network Data Management Protocol (NDMP), you can use the backup and restore functions provided by your DMA instead. Refer to your DMA documentation for instructions.

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 management controller.

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

See the *NDMP Integration Guide for Pillar Axiom NAS Systems* for information on how to configure supported DMAs.

Note: When using NDMP to restore files that have symbolic links, the user and group are not restored.

Create an NDMP User Account

The data management application (DMA) issues commands to the Network Data Management Protocol (NDMP) server on the Pilot management controller to initialize and control backup and restore operations. To issue these commands, the DMA first logs in to the Pillar Axiom storage system. To permit this login, enable NDMP and create an NDMP user account on your Pillar Axiom system.

1. Click the **System** icon in the top context pane.
2. Click the **NDMP Backup Settings** link in the left navigation pane.
3. Choose **Modify NDMP Backup Settings** from the **Actions** drop-down list.
4. Select **Enable NDMP**.
5. Enter values for the NDMP port, user name, and File Server through which backups are performed.

   Note: The NDMP user account has permission only to perform NDMP backup and restore operations.
6. Click **Change Password**, enter between six and eight characters for the NDMP user account password, and click **OK**.
7. To save the new NDMP user account, click **OK**.

For parameter definitions, see the NDMP Configuration Page.

Perform Immediate Data Replication

About Data Replicas and System Capacity

You can create online data replicas in different ways. Each method consumes the capacity in the storage array differently.

The Pillar Axiom storage system ensures that all logical volumes (filesystems and LUNs) that are associated with a particular replica tree reside on[1] the same Slammer control unit (CU).[2] If you change the home of any of these logical volumes, the system changes all of them. This feature applies to all of the following objects:

- Snap FSs
- Clone FSs
- Clone LUNs
- Volume Copies
- Active data migrations because of Quality of Service (QoS) changes

Volume Copies and logical volumes being migrated due to QoS changes are present in the original replica tree until they are detached. Once the Volume Copy or the migrated volume is detached, the volume is removed from the original replica tree and becomes the root of a new replica tree.

However, after you start a Volume Copy operation or the system starts a data migration operation, if you re-home anything in the replica tree, the mechanics are a little different. If the system has not yet detached the copy from its source volume, the copy will be re-homed. If, however, the system has already detached the copy, the copy is no longer in the original replica tree and, so, is not re-homed.

Table 13 Capacity usage by online data replicas

Method	Description	Capacity usage
Clone FS	Creates a readable and writable point-in-time copy of a filesystem.	Consumes system space allocated for clones. Only changes

[1] Sometimes the term *homed on* or *owned by* is used instead of *reside on*.

[2] This discussion of replica trees on a Slammer CU does not apply to replicated objects created by the Pillar Axiom MaxRep Replication for NAS and Pillar Axiom MaxRep Replication for SAN utilities.

Table 13 Capacity usage by online data replicas (continued)

Method	Description	Capacity usage
		to the source or clone are stored.
Snap FS	Creates a read-only copy of a filesystem that users can access through special directories in the filesystem hierarchy.	Consumes part of the capacity of the parent filesystem.
Clone LUN	Creates a readable and writable point-in-time snapshot of a LUN.	Consumes system space allocated for clones. Only changes to the source or clone are stored.
Volume Copy	Creates a block-level, full-image, read-write copy of a logical volume (filesystem or LUN). QoS attributes for a Volume Copy can differ from the QoS attributes of the original.	Consumes free space from system capacity that is equal to the current size of the volume.

The online data replicas identified in the preceding table have the following characteristics:

- They require no prior configuration (other than the initial allocation).
- They are created by explicit one-time operations.
- They are created on the same Pillar Axiom system as the source volume.
- Updates to the source volume are not reflected in the replica. When data changes in the source volume, that change *is not* reflected in the replica.

In comparison, data replicas created by the Pillar Axiom Replication utilities for network attached storage (NAS) and storage area network (SAN) environments are summarized in the following table.

Table 14 Capacity usage by remote data replicas

Method	Description	Capacity usage
NAS-based remote replica	Creates a read-only copy (snapshot) of a source filesystem on a second (or same) Pillar Axiom system. The content of this replica is manually synchronized with the paired source volume. For	Consumes free space from the target filesystem on the target system.

Table 14 Capacity usage by remote data replicas (continued)

Method	Description	Capacity usage
	disaster recovery purposes, a system administrator can make this replica live so applications can continue reading and writing.	
SAN-based remote replica	Creates a clone of a LUN on a second Pillar Axiom system. The content of this replica is automatically synchronized with the paired source volume. A system administrator can use this replica for disaster recovery purposes. The remote replica LUN is normally not accessible. However, when the replica is isolated (replication stopped), a system administrator can map it, which makes it visible to client systems. An isolated replica is readable and writeable.	Consumes system space allocated for Clone LUNs on both systems. Only the changes to the source are stored on the target system.

Remote replicas have the following characteristics:

- SAN-based replicas are configured before they are created.
- NAS and SAN-based replicas are paired with a source volume.
- NAS and SAN-based replicas are created on a remote Pillar Axiom system.
- Updates to the source volume are reflected in the replica. When data changes in the source volume, that change *is* reflected in the replica the next time a synchronization operation takes place.

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

About Immediate Clone FS Creation

You can create an immediate Clone FS at any time. For example, you might want to create an immediate Clone FS right before you make significant changes to the data itself.

The date and timestamp of a filesystem clone virtual directory, as viewed from a client, will be different from the date the clone was created. The time is taken from the source object of the clone, not the creation time of the clone. The actual creation time of the clone is available in the GUI or in pdscli clone query responses.

When you create a Clone FS, you need to make sure you have enough space allocated for the Clone FS. An immediate Clone FS consumes space on the system that was allocated for clones during filesystem creation. The system stores only the changes made to either the source volume or the clone in the allocated clone storage space.

Important! Make sure that the clone space does not fill up, consuming the maximum amount of space allocated. Pillar strongly recommends that you monitor the amount of space available and modify the volume to allocate more clone space as needed.

Create an Immediate Clone FS

Clone FSs create writeable snapshots of a filesystem using partial-block snapshot technology.

1. Click the **Storage** icon in the top context pane.
2. Click the **Filesystems** (or **Clone FS**) link in the left navigation pane.
3. Select the filesystem (or Clone FS) for which you want to create an immediate clone.
4. Choose **Perform Clone FS Now** from the **Actions** drop-down list.
5. Enter a name for the Clone FS and click **OK**.

The name of the new Clone FS appears in the Storage > Clone FS page in the Name column, indented beneath the source filesystem (or Clone FS).

About Immediate Snap FS Creation

You can create an immediate Snap FS at any time, even if you have scheduled recurring Snap FSs. For example, you might want to create an immediate Snap FS right before you make significant changes to the data itself.

The date and timestamp of a filesystem snapshot virtual directory, as viewed from a client, will differ from the date the snapshot was created. The time is taken from the object of the snapshot, not the creation time of the snapshot. The actual creation time of the snapshot is available in the GUI or in pdscli snapshot query responses.

You can create a Snap FS of a SecureWORMfs filesystem for backup purposes, but you cannot restore a SecureWORMfs filesystem from a Snap FS. The system, however, uses a Snap FS of a SecureWORMfs filesystem in the event the SecureWORMfs is recovered using FSCK.

An immediate Snap FS consumes part of the filesystem's capacity.

Create an Immediate Snap FS

1. Click the **Storage** icon in the top context pane.
2. Click the **Filesystems** link in the left navigation pane.
3. Select the filesystem for which you want to create an immediate Snap FS.
4. Choose **Perform Snap FS Now** from the **Actions** drop-down list.
5. Enter a name for the Snap FS and click **OK**.

The name of the new Snap FS appears in the Storage page under Snapshot Name.

Restore a Filesystem from a Snap FS

You can restore a filesystem from a Snap FS at anytime.

1. Click the **Storage** icon in the top context pane.
2. Click the **Snap FS** link in the left navigation pane.
3. Select the Snap FS that you want to restore.
4. Choose **Restore Snap FS** from the **Actions** drop-down list.

Create an Immediate Clone LUN

Clone LUNs create writeable snapshots of a LUN using partial-block snapshot technology.

You can create an immediate Clone LUN at any time. A Clone LUN is defined as:

> A point-in-time, read-write, partial-block snapshot of a LUN that you intend to split from the original LUN for immediate access. A Clone LUN retains the same QoS parameters as the source LUN and consumes storage capacity from the Clone LUN repository that was allocated for the source LUN. A Clone LUN cannot be scheduled from the Pillar Axiom Storage Services Manager; it is an immediate operation. Clone LUNs provide a convenient method to branch from the source data without the need to do a full-block copy operation.

After the clone is created, the system administrator can change the priority setting and location of the clone.

To create a Clone LUN, you need to make sure you have enough space allocated for the Clone LUN. An immediate Clone LUN consumes space on the system that was allocated for clones during LUN creation. The system stores only the changes made to either the source volume or the clone in the allocated clone storage space.

Important! Make sure that the clone space does not fill up, consuming the maximum amount of space allocated. Pillar strongly recommends that you monitor the amount of space available and modify the volume to allocate more clone space as needed. See System Components That Can Be Monitored.

1. Click the **Storage** icon in the top context pane.
2. Click the **LUNs** (or **Clone LUN**) link in the left navigation pane.
3. Select the LUN (or Clone LUN) for which you want to create an immediate clone.
4. Choose **Perform Clone LUN Now** from the **Actions** drop-down list.
5. Enter a name for the Clone LUN.
6. Select the protocol for the Clone LUN.
7. Set the mapping and host connections for the Clone LUN.
8. Click **OK** to create the Clone LUN.

The name of the new Clone LUN will appear on the Storage > Clone LUNs page in the Name column, indented beneath the source LUN (or Clone LUN).

Activate a Clone

You can activate a Clone FS or Clone LUN to make it available to users.

When activating a Clone LUN, assign the volume a LUN number to make it available to all SAN hosts. You can then map the LUN to specific hosts.

1. Click the **Storage** icon in the top context pane.
2. Click the **Clone FS** or **Clone LUN** link under the **Disk Replication** heading in the left navigation pane.
3. Select the clone that you want to activate.
4. Choose the appropriate **Modify Clone** item from the **Actions** drop-down list.
5. On the **Identity** tab, remove the check mark from the **Inactive** option.
6. Click **OK** to activate the clone.

The name of the new volume appears in the Storage page under Filesystems or LUNs.

For more information about this feature, see About Licensing Optional Premium Features.

For parameter definitions, see the Clone Activation Page.

Display Data Replica Details

You can display the name and status of a completed data replica (Clone FS, Snap FS, and Clone LUN). You can display other types of information as well, depending on the type of data replica.

1. Click the **Storage** icon in the top context pane.
2. In the left navigation pane, click the link for one of these types of data replica:
   - **Snap FS**
   - **Clone FS**
   - **Clone LUN**

3 Review the displayed information to ensure that the details for the data replica are what you expect.

For parameter definitions, see the following:

- Clone FS Overview Page
- Snap FS Overview Page
- Clone LUN Overview Page

Manage Snap FS Schedules

A Snap FS schedule defines:

- Intervals at which a Snap FS is created.
- Maximum number of Snap FSs to create.

You can delete Snap FS schedules when they are no longer needed, or if you want to create a new schedule.

Create Snap FS Schedules

You can create replication schedules that in turn create a Snap FS of a filesystem at regular intervals.

1. Click the **Data Protection** icon in the top context pane.
2. Click the **Scheduled Snapshots** link under **Replication Schedules** in the left navigation pane.
3. Choose **Create Snap FS Schedule** from the **Actions** drop-down list.
4. Enter a name for the schedule.
5. Choose the filesystem on which the Snap FS is based from the **FileSystem Name** drop-down list.
6. Select a recurrence interval for the schedule.

   Choose one of:

   - **Run Once**
   - **Hourly**
   - **Daily**
   - **Weekly**
7. Click the **Schedule** tab.
8. Enter the start time and date for this schedule.
9. If needed for the interval that you selected, select a recurrence number. Do so for hourly, daily, or weekly schedules, but not for run-once schedules.

10 Enter the number of Snap FSs to archive. The value differs based on the schedule's recurrence interval.

11 To save the schedule, click **OK**.

For parameter definitions, see the following:

Schedule Configuration Page, Details Tab

Schedule Configuration Page, Schedule Tab

See also: Create an Immediate Snap FS.

Display Snap FS Schedules

You can display details about all Snap FS replication schedules at one time.

Schedule details include:

- Schedule start time and recurrence frequency
- Filesystem on which the Snap FS is based
- Status of each scheduled Snap FS

1 Click the **Data Protection** icon in the top context pane.

2 Click the **Scheduled Snapshots** link under **Replication Schedules** in the left navigation pane.

3 Ensure that the Snap FS schedules that you expect to see are listed.

For parameter definitions, see the Replication Schedule Summary Page.

Delete Snap FS Schedules

You can delete a Snap FS schedule when your data replication requirements change.

1 Click the **Data Protection** icon in the top context pane.

2 Click the **Scheduled Snapshots** link under **Replication Schedules** in the left navigation pane.

3 Select one or more Snap FS names.

4 Choose **Delete Snap FS Schedule** from the **Actions** drop-down list.

5 When prompted to confirm the deletion, click **OK** to delete the schedule.

Restore a Filesystem from a Clone FS

You can restore a particular filesystem back to its state that you previously captured through a Clone FS.

Restoring a filesystem from a Clone FS uses block snapshot technology that allows the filesystem to keep its same identity and to come back on-line in a short amount of time, especially when compared to copying the entire data set back from a tape backup. This restoration process copies only that data that was modified since the snapshot was taken. Furthermore, this process allows access to the data while the background copy is in progress.

Such restoration is often used to restore a filesystem to a known good image in various scenarios, including:

- Some undesirable changes were made.
- An external client application or virus corrupted the filesystem.

1 Click the **Storage** icon in the top context pane.

2 Click the **Clone FS** link in the left navigation pane.

3 Select the Clone FS from which you want to restore the filesystem.

Note: The restoration process resets the creation date of the filesystem to that of the selected Clone FS.

4 Choose **Restore from Clone FS** from the **Actions** drop-down list.

5 In the confirmation dialog, click **OK**.

The system restores the filesystem, during which time system performance may be slightly degraded. The system starts a task in the background to perform the copy operation. When the background task completes, the system writes an event to the event log.

Restore a LUN from a Clone LUN

You can restore a particular LUN back to its state that you previously captured through a Clone LUN.

Restoring a LUN from a Clone LUN uses block snapshot technology that allows the LUN to keep its same identity and to come back on-line in a short amount of time, especially when compared to copying the entire data set back from a tape backup. This restoration process copies only that data that was modified since the snapshot was taken. Furthermore, this process allows access to the data while the background copy is in progress.

Such restoration is often used to restore a LUN to a known good image in various scenarios, including:

- Some undesirable changes were made.
- An external client application or virus corrupted the LUN.

1 Click the **Storage** icon in the top context pane.

2 Click the **Clone LUN** link in the left navigation pane.

3 Select the Clone LUN from which you want to restore the LUN.

Note: The restoration process resets the creation date of the LUN to that of the selected Clone LUN.

4 Choose **Restore from Clone LUN** from the **Actions** drop-down list.

5 In the confirmation dialog, click **OK**.

The system restores the LUN, during which time system performance may be slightly degraded. The system starts a task in the background to perform the copy operation. When the background task completes, the system writes an event to the event log.

About the Pillar Axiom VSS Provider Plug-In

Pillar Data Systems provides a VSS Provider plug-in to facilitate use of VSS-enabled backup applications with your Pillar Axiom storage system.

Microsoft's Volume Shadow Copy Service (VSS) is a SAN-only feature that enables data protection and management services through a standard set of configuration and monitoring capabilities for creating, manipulating and restoring snapshots without shutting down applications or essential services.

For more information about VSS, see:

- Microsoft's Volume Shadow Copy Service Technical Reference (http://technet2.microsoft.com/windowsserver/en/library/b45c7507-b7ad-420e-a981-c273c2012a831033.mspx?mfr=true)
- The Microsoft Developers Network (MSDN) article The VSS Model (Windows) (http://msdn.microsoft.com/en-us/library/aa384625.aspx)

The Pillar Axiom VSS Provider is a VSS hardware provider that allows VSS-enabled applications to make volume shadow copies of data on Pillar Axiom systems without interrupting normal operations. Pillar technology partners FalconStor and InMage offer data replication solutions featuring VSS implementations that do not require the Pillar VSS Provider plug-in.

Refer to your VSS-enabled backup application documentation for instructions on configuring and using VSS with your backup application.

Download and Install the VSS Provider Plug-In

Download the Pillar Axiom VSS Provider plug-in from the Pillar Axiom Storage Services Manager for installation on your SAN host.

Prerequisites:

- The SAN host has TCP/IP Ethernet connectivity to the Pilot management controller.
- For the VSS Provider to create volume shadow copies, the SAN host must have Fibre Channel connectivity to the Slammer storage controller.
- During the installation, you need the system serial number, user name, and password.

1. To download the VSS Provider plug-in, select **Data Protection** and **VSS Provider Download** in the Pillar Axiom Storage Services Manager, and select **Download VSS Provider** from the **Actions** drop-down menu.
2. Follow the instructions to install the VSS Provider plug-in on your SAN host.
3. Verify installation by running the following command at a command prompt:

   ```
   vssadmin List Providers
   ```

   Result:
   This command should return the name of the Pillar VSS Provider, as follows:

   ```
   Provider name: 'PDS VSS HW Provider'
   ```

   If it does, installation was successful and your SAN host can use the VSS Provider to create shadow copies.

The VSS Provider installer allows you to configure a single Pillar Axiom system. To configure additional systems or remove systems, use the `registerAxiom.exe` command line tool:

- To configure additional systems, run this command at a command prompt:

  ```
  registerAxiom.exe serial_number user_name password
  ```

- To remove a configured system, run this command:

  ```
  registerAxiom.exe serial_number
  ```

- Running `registerAxiom.exe` without parameters prints the usage instructions.

For parameter definitions, see the Volume Shadow Copy Service (VSS) Page.

CHAPTER 5

Manage Notifications

About System Notifications

The Pillar Axiom storage system provides various methods of setting up system notifications, including:

- **Alerts**.

  A Simple Mail Transfer Protocol (SMTP) email message that notifies recipients of specified system events, such as informational, warning, or critical events. Alerts are optional and supplement normal event logging and Call-Home notification.

- **Call-Home**.

  A feature of a Pillar Axiom storage system that, when enabled, allows the system to notify Pillar World Wide Customer Support Center of critical issues specific to a Pillar Axiom system. No customer data is transmitted. Call-Home transfers files over the Internet using one of the following user-selected methods:

  - SCP: Uses the secure copy (SCP) method with 1024-bit encryption and secure keys.
  - HTTPS: Uses the Hypertext Transfer Protocol Secure method by sending files directly to Pillar or through a proxy server for security purposes. This method can also be used when the Pillar Axiom system does not have direct access to the Internet.

You must define an email server to receive alerts and send email messages to designated recipients. This email server is also used to send Call-Home notifications to the Pillar World Wide Customer Support Center.

Pillar Axiom systems also support the following protocols for monitoring the configuration of various system components:

- **Storage Management Initiative Specification (SMI-S)**.

  A storage management standard developed by Storage Networking Industry Association (SNIA) that allows multivendor software support of heterogeneous storage devices. Through SMI-S profiles, administrators can

query, for example, device credentials, copy services, masking and mapping of Fibre Channel ports, and so forth.

- **Simple Network Management Protocol (SNMP)**.

  A standard network protocol that is used to monitor Slammers, Bricks, and the drives within the Bricks. Through SNMP traps, administrators can monitor, for example, central processing unit (CPU) temperature and field replaceable unit (FRU) removal and insertion.

System Components That Can Be Monitored

The Simple Network Management Protocol (SNMP) management information base (MIB) is self-documenting and lists Pillar Axiom storage system resources that you can monitor. Download the MIB (a text file) from the **System > SNMP** menu in the Pillar Axiom Storage Services Manager GUI.

Some of the Pillar Axiom resources that a system administrator may wish to monitor are listed below. Some of this information can be used, for example, to graph or otherwise track the trend lines of certain resources, such as that for storage space and its utilization and for I/O operations/sec (IOPS) over certain time periods.

Administrator Actions. These resources are notifications that the Pillar Axiom generates to identify conditions that warrant investigation *and* action. Administrator Actions include, for example:

- Notifications about resources that are not fully operational, indicating a need for maintenance.
- Notifications about storage running low, indicating a need for reallocation or cleanup of resources or possibly the purchase of additional storage. This kind of information is important when an administrator has implemented thin provisioning.

Call-Home or Manual Log Collection. Querying these resources, the administrator can check:

- Time of collection
- Availability status
- Type of information contained in the logs

Running Tasks. Some tasks running in the background are normal management functions such as scheduled snapshots, scheduled upgrades, and so forth, or are the result of some administrative action. Other tasks however may indicate a condition on the Pillar Axiom system worth investigating, such as:

- Pilot restarts
- System restarts
- Topology rediscovery

This category is also useful for seeing when a planned task has completed or may need recovery, such as when replicating a very large filesystem.

Scheduled Tasks. Querying this resource allows the administrator to determine which tasks are scheduled and when they are scheduled. Knowing this

information can be useful in determining whether some traps or events can be expected.

Software Versions. Capturing software versions is useful in a large data center where a single SNMP management utility can keep the administrator from having to log into each Pillar Axiom system individually for the same information to determine which machines need updates or do discover whether a particular software update is complete.

Storage Usage. Monitoring short and long term trends in capacity usage helps the system administrator avoid getting an Administrator Action warning that, for example, Clone LUNs are being deleted to free up capacity. Because you can over allocate logical volumes (filesystems and LUNs) when taking advantage of the thin provisioning feature, such volumes need to be monitored and may require additional physical storage.

System Configuration. Use a central SNMP resource to view the configuration and status of the resources of multiple systems, including:

- LUNs
- File Servers
- Filesystems
- Interfaces
- Exports and shares
- Clones and snapshots
- LUN mapping and masking

Traps. Traps are equivalent to email-based administrator alerts and provide another means of alerting system administrators to unfavorable storage conditions, which may or may not result in an Administrator Action.

About Alert Management

You can create alerts so that you are notified when specific Pillar Axiom system events occur. You may want to display the details of an alert and make changes as needed. You can also test alerts to make sure that the specified email addresses are correct.

Create Alerts

Create alerts so that you are notified when specific events occur in the Pillar Axiom system. You can specify the types of system events that trigger alerts as well as designate the recipients who receive the alerts.

If you do not set up alerts, you can still monitor system events using the event log. Call-Home notifications are also independent of alerts and will be sent to Pillar Data Systems about issues in the Pillar Axiom system.

1. Click **System** in the top context pane.
2. Click **Alerts** in the left navigation pane.
3. Choose **Create Alert** from the **Actions** drop-down list.
4. Enter a name for the alert and a description.
5. Enter the email addresses of alert recipients.
6. Select one or more event types to trigger this alert.
7. Click **Test Email** to make sure that the alert is sent to the correct email addresses and that the SMTP server is properly configured (optional).
8. To save the new alert, click **OK**.

For parameter definitions, see the Alert Details Page.

See also: Configure Notification Settings.

Display Alerts

You can display the details of an alert and determine if any changes are needed.

1. Click the **System** icon in the top context pane.
2. Click the **Alerts** link in the left navigation pane.
3. Click the name of the alert that you want to view.
4. When you are finished, click **OK**.

For parameter definitions, see the Alert Details Page.

Modify Alerts

You can modify the way in which an administrator is notified about Pillar Axiom events. For example, you may want to change the event categories that trigger the alert or you may need to change an email address.

1. Click the **System** icon in the top context pane.
2. Click the **Alerts** link in the left navigation pane.
3. Click the name of the alert that you want to modify.
4. Enter values for the attributes that you want to modify.
5. To save the modified alert, click **OK**.

For parameter definitions, see the Alert Details Page.

See also: Display System Events and Performance Statistics.

Delete Alerts

You can delete an existing alert. For example, you can do this if someone leaves the company and you no longer want event notifications to be sent to an inactive email account.

1. Click the **System** icon in the top context pane.
2. Click the **Alerts** link in the left navigation pane.
3. Select one or more alerts.
4. Choose **Delete Alert** from the **Actions** drop-down list.
5. When prompted to confirm the deletion, click **OK** to delete the alert.

For parameter definitions, see the Alert Details Page.

SNMP Trap Host Management

If you use Simple Network Management Protocol (SNMP) management applications to monitor network devices, you can define SNMP trap hosts to receive Pillar Axiom traps. Any workstation that has an SNMP-based management application installed on it can be a trap host.

Pillar Axiom systems support SNMP version 2c. SET operations from SNMP management applications are not supported.

An *MIB table* is a plain text file that provides the details on all components for which Pillar provides management information.

Note: You can define alerts as an alternative to SNMP.

Create SNMP Trap Hosts

1. Click the **System** icon in the top context pane.
2. Click the **SNMP** link in the left navigation pane.
3. Choose **Modify SNMP Settings** from the **Actions** drop-down list.
4. Enter values in the **Community String** and corresponding **Host IP** and **Port** fields to specify where the traps are directed:
   - For Simple Network Management Protocol (SNMP) queries, enter the client IP address, community string, and port 161.
   - For SNMP traps, enter the client IP address, community string, and port 162.
5. Click **Add Host**.
6. To save the SNMP configuration, click **OK**.

For parameter definitions, see the SNMP Configuration Page.

See also: Create Alerts.

Modify SNMP Trap Hosts

You can modify the hosts that receive Simple Network Management Protocol (SNMP) traps. You may, for example, need to modify IP address of the trap host if you install your SNMP-based management application on a different administrative workstation.

1. Click the **System** icon in the top context pane.
2. Click the **SNMP** link in the left navigation pane.
3. Choose **Modify SNMP Settings** from the **Actions** drop-down list.
4. Select an SNMP host and click **Modify**.
5. Enter values for the attributes that you want to modify and click **Modify Host**.
6. To save the modified SNMP configuration, click **OK**.

For parameter definitions, see the SNMP Configuration Page.

Delete SNMP Trap Hosts

You can delete a trap host from the Simple Network Management Protocol (SNMP) configuration. For example, you might do this after you uninstall an SNMP-based management application from someone's workstation.

1. Click the **System** icon in the top context pane.
2. Click the **SNMP** link in the left navigation pane.
3. Select the trap hosts that you want to delete.
4. Choose **Delete SNMP Hosts** from the **Actions** drop-down list.
5. When prompted to confirm the deletion, click **OK** to delete the trap hosts.

For parameter definitions, see the SNMP Configuration Page.

Notification Settings Management

To manage notification settings, you can:

- Modify Call-Home Settings to enable or disable the Call-Home feature.
- Modify Email Configuration Settings to change the email server settings that are used to receive alerts from the Pillar Axiom system.

About Call-Home Settings Modification

The Call-Home feature notifies Pillar Data Systems about issues in the Pillar Axiom system. 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.

Modify Call-Home Settings

1. Click the **System** icon in the top context pane.
2. Click the **Networking** link in the left navigation pane.
3. Choose **Modify Network Settings** from the **Actions** drop-down list.
4. Click the **Call-Home** tab.
5. Enter the values in the Call-Home configuration fields.
6. To save the Call-Home settings, click **OK**.

For parameter definitions, see the Call-Home Configuration Page.

Test Call-Home

The system sends a Call-Home message to verify that Call-Home is correctly configured.

1. Click the **System** icon in the top context pane.

2 Click the **Networking** link in the left navigation pane.

3 Choose **Test Call-Home** from the **Actions** drop-down list.

Note: Only Primary Administrator or Administrator 1 accounts are allowed to test Call-Home.

4 Confirm that you want to send test Call-Home information to the specificed Call-Home server and click **OK**.

For parameter definitions, see the Notification Page.

Modify Email Configuration Settings

Define an email server to receive alerts from the Pillar Axiom system and send the email messages to designated recipients. If you do not set the email server, the system does not send alerts to administrators of events that have occurred.

1 Click the **System** icon in the top context pane.

2 Click the **Networking** link in the left navigation pane.

3 Choose **Modify Network Settings** from the **Actions** drop-down list.

4 Click the **Notification** tab.

5 Enter the modified values in the **Email Configuration** fields.

6 To save the email configuration settings, click **OK**.

For parameter definitions, see the Notification Page.

CHAPTER 6

Manage Software and Hardware Components

About Software and Hardware Management

Use the Pillar Axiom Storage Services Manager to:

- Manage Software Modules to review software versions and update the software modules on your Pillar Axiom storage system.
- Manage Hardware Components to review the status and configuration of the hardware components on your Pillar Axiom system. You can add Slammers and Bricks to scale performance and capacity of your Pillar Axiom system and replace hardware components as needed.

Manage Software Modules

Use the Pillar Axiom Storage Services Manager to:

- Display Software Versions to review the versions of all software modules on your Pillar Axiom system.
- Download Software Updates to obtain new versions of the Pillar Axiom software.
- Update Pillar Axiom Software to install a new version of software onto your Pillar Axiom system.
- Modify a Software Update Schedule to change the time in which you want to perform a software update.

An authorized representative from Pillar Data Systems will notify you when software updates are available.

Display Software Versions

You can display the versions of all software modules in your Pillar Axiom storage system.

The version information includes:

- Drive and Enclosure Services (ES) firmware in the Brick storage enclosures.
- Software and operating system in the Pilot management controller.
- Software (NAS or SAN) and programmable ROM (PROM) in Slammer storage controllers.

1. Click the **Support** icon in the top context pane.
2. Click the **Software Modules** link in the left navigation pane.
3. Review the displayed information to ensure that the versions are what you expect.

For parameter definitions, see the Software Modules Page.

Download Software Updates

You can download software and firmware updates from the Pillar website and install the software on to the client system.

1. Log in to the Customer Support site:
   - Go to https://support.pillardata.com/login.do.
   - Enter your Email and Password and click **Log In**. If this is your first time, you will need to register.
2. Click **Software Downloads** in the left navigation pane, and select the product for which you want to download software.
3. Expand the categories in the **Release** column of the **View Software** page then click the name of the software to download.
4. Review the details of the update and click the download icon in the **Software Download Details** section.
5. Browse to the location on your local system where you want to save the software update and click **Close** when complete.

See also: About Pillar Axiom Software Updates.

About Pillar Axiom Software Updates

An update installs a new version of software or firmware onto a Pillar Axiom storage system. An update affects one or more of the following components.

- Brick storage enclosures (which requires a system restart):
  - Drive firmware
  - Enclosure Services (ES) firmware
- Pilot management controller:
  - Software
  - Operating system
- Slammer storage controllers:

- Software for network attached storage (NAS) or storage area network (SAN) configurations
- Programmable ROM (PROM) (which requires a system restart)

Update Pillar Axiom Software

Tip: Pillar does not recommend that you upload a software update package to your Pillar Axiom storage system over a slow connection, such as a wide area network (WAN) connection. Instead, use a faster internal network connection (10 Mbit/sec or greater).

1. Click the **Support** icon in the top context pane.
2. Click the **Software Modules** link in the left navigation pane.
3. Choose **Update Software** from the **Actions** drop-down list.
4. Click **Upload Package**.

   Note: Uploading an update package places the package onto the Pilot management controller.
5. Browse to the update package, and click **OK**.
6. Select all software and firmware modules.

   Important! You may need to scroll through the entire list so you can select all modules.
7. Select **Perform Update Now** or **Schedule Update**.

   For a scheduled update, select the date and time on which the Pillar Axiom system should perform the operation. You can schedule updates to occur within 72 hours.

For parameter definitions, see the Software Modules, Details Tab.

Modify a Software Update Schedule

You can schedule software updates to occur at a specified time. For example, you can schedule an update to occur during off-peak hours.

1. Click the **Support** icon in the top context pane.

2 Click the **Scheduled Updates** link in the left navigation pane.

3 Select the scheduled update that you want to modify.

4 Choose **Modify Scheduled Update** from the **Actions** drop-down list.

5 Select a new date and time at which the Pillar Axiom system should perform the operation. You can schedule updates to occur within 72 hours.

6 Click **OK**.

7 Click **OK** in the **Confirm Software Update** dialog box.

About the Effect on QoS in Mixed Brick Configurations

When hard disk drive (HDD) SATA Bricks are added to a Pillar Axiom storage system that contains only Fibre Channel (FC) Bricks, the system automatically changes the relative priority of all existing logical volumes (filesystems and LUNs) to Premium.

When HDD SATA Bricks are added to a Pillar Axiom system that contains only solid state drive (SSD) SATA Bricks, the system does not change the Premium priority of the existing volumes on the SSD Bricks because SSD Bricks have Premium storage only.

If desired, system administrators may then migrate selected logical volumes to the newly added HDD SATA Bricks by changing the relative priority of those volumes.

Manage Hardware Components

Use the Pillar Axiom Storage Services Manager to:

- Display Hardware Component Information to review the status and configuration of the hardware components on your Pillar Axiom storage system.
- Modify System Name to change the name of a Pillar Axiom system.
- Modify Hardware Component Names to change the names of Slammers and Bricks.
- Replace a Hardware Component to remove failed hardware components and add replacement components with the help of the Guided Maintenance feature.
- Identify Hardware Components to locate specific Slammers and Bricks in the Pillar Axiom system.
- Display Tape Storage Devices to identify associated Slammers and device names.

Display Hardware Component Information

At any time, you can review the status and configuration details of hardware components in the system.

1. Click the **Health** icon in the top context pane.
2. Click a component type link in the left navigation pane.

   Choose one of:

   - **Slammers**
   - **Bricks**
   - **Pilot**
   - **UPS**

Display Additional Hardware Component Details

1. Click the name of a component.
2. Navigate through the tabbed component pages to review its configuration details.

   Different details appear based on the component type.

For parameter definitions, see the Hardware Status and Configuration Page.

Modify System Name

You may want to change the name of the system that is displayed in the Pillar Axiom Storage Services Manager.

1. Click the **System** icon in the top context pane.
2. Click the **Networking** link in the navigation pane.
3. Choose **Modify Network Settings** from the **Actions** drop-down list.
4. Click the **Interfaces** tab.
5. Type a new name in the **System Name** field and click **OK**.

Modify Hardware Component Names

You may want to change the display names of the Slammers and Bricks.

The names are displayed in the Pillar Axiom Storage Services Manager.

1. Click the **Health** icon in the top context pane.
2. Click the link for the type of component, Slammer or Brick, in the left navigation pane.
3. Select the hardware component that you want to modify.
4. Choose **Modify Slammer Name** or **Modify Brick Name** from the **Actions** drop-down list.
5. Type a new name and click **OK**.

About Hardware Component Replacement

Your Support Services contract provides guidelines to replace a hardware component in your Pillar Axiom system. Based on the terms of your contract, you can:

- Replace the component yourself.
- Place a service request so that a Service Technician comes to your site to replace the component.

Important! Replacement of hardware components requires Guided Maintenance to:

- Prepare the system for the removal of the old hardware component and the addition of the new component.
- Integrate the replacement component into the system.

We recommend that you check the power supplies annually for accumulated dust. If needed, vacuum the power supplies, even if no component replacement or repair is required. This type of maintenance does not require Guided Maintenance.

Replace a Hardware Component

1. Click the **Support** icon in the top context pane.
2. Click the link in the left navigation pane for the type of component to replace:
   - **Slammers**
   - **Bricks**
3. Select the hardware component to replace.
4. Choose **Repair Slammer Hardware** or **Repair Brick Hardware** from the **Actions** drop-down list.
5. Click the name of the replaceable unit to start Guided Maintenance.

   Note: Although the Repair Hardware page lists the Brick or Slammer chassis as a replaceable unit, Pillar does not currently support chassis replacement.

6 Follow the instructions in the Guided Maintenance feature of Pillar Axiom Storage Services Manager and the appropriate *Pillar Axiom Service Guide* to repair the hardware component.

7 Click **Verify Status** to confirm that the replaceable unit is functioning normally.

8 Click **Finish**.

Identify Hardware Components

Use the identify function to locate specific Slammers and Bricks in the Pillar Axiom storage system. The identify function blinks the light-emitting diodes (LEDs) on the front and back of the target hardware component. You can also choose the **Reverse Identify** function, which blinks the LEDs on all of the Pillar Axiom system hardware components except the target field replaceable unit (FRU).

1 Click the **Health** icon in the top context pane.

2 Click the link in the left navigation pane for the type of component to identify:

- **Slammer**
- **Brick**
- **Pilot**

3 Click the name of any installed hardware component of the selected type.

4 Click **Identify** at the top of the component information page.

5 Follow the instructions in the Guided Maintenance feature of Pillar Axiom Storage Services Manager to identify the hardware component. You can choose one of **Identify** or **Reverse Identify**.

6 Click **Finish** when you are done.

See also the appropriate *Pillar Axiom Service Guide*.

Display Tape Storage Devices

If any tape storage devices are directly attached to the Pillar Axiom storage system, you can identify the:

- Slammer to which the tape storage devices are directly attached.

- Names that are assigned to the tape storage devices.

1. Click the **Data Protection** icon in the top context pane.
2. Click the **Tape Devices** link in the left navigation pane.
3. Review the displayed information about the tape storage devices that are directly attached to Slammers.

For parameter definitions, see the Tape Devices Page.

See also: Create an NDMP User Account.

CHAPTER 7

Manage Administrator Accounts

About Account Management

Administrators have specific privileges in the Pillar Axiom storage system based on their account type. The different types of administrator accounts include:

- Administrator 1:

  A login account that has the authority to perform all administration and configuration tasks.

- Administrator 2:

  A login account that has the authority to perform all administrative and configuration tasks, except:

  - Create, modify, or delete administrator accounts and File Servers.
  - Modify system-wide settings such as Simple Network Management Protocol (SNMP).
  - Modify software or hardware configurations.
  - Shut down the system.

- Monitor:

  A login account that has the authority to perform read-only management tasks in a Pillar Axiom storage system and the ability to modify their own account attributes.

- Primary system administrator:

  A unique login account that has the authority to perform all administration and configuration tasks. This account cannot be deleted or disabled.

About Administrator Account Creation

You can create multiple administrator accounts in a Pillar Axiom system. Additional accounts are not necessary, but they are useful if you want to delegate administrator responsibilities. For example, you might choose to create:

- One administrator account. In this way, a designated person can assume responsibility while the Primary system administrator is on vacation. Assign this account to the Administrator 1 role.

  Tip: Pillar strongly recommends that you set up a Type 1 Administrator account when you install the system. Besides the Primary system administrator, only a Type 1 Administrator can modify an account password (including that of the Primary system administrator) without knowing the previous password.

- One or more administrator accounts with read-only privileges. In this way, managers can monitor the system but they cannot change configuration details. Assign these accounts to the Monitor role.

You can create up to 23 administrator accounts.

Create an Administrator Account

1. Click the **System** icon in the top context pane.
2. Click the **Administrator Accounts** link in the left navigation pane.
3. Choose **Create Account** from the **Actions** drop-down list.
4. Enter values for the account attributes that you need.
5. To save the new account, click **OK**.

For parameter definitions, see the Administrator Configuration Page.

See also: Modify Administrator Account Security Settings.

Display Administrator Accounts

You can display details about all administrator accounts or about a specific administrator account.

1. Click the **System** icon in the top context pane.
2. Click the **Administrator Accounts** link in the left navigation pane.
3. Review the displayed information to ensure that the account details are what you expect.
4. For a specific account, click the login name that is associated with an account.
5. Review the account's configuration details.

For parameter definitions, see the Administrator Accounts Overview Page.

About Administrator Account Modification

If you delegated administrative tasks to other administrators, you may need to:

- Modify account attributes (for example, change an administrator's password or disable an account other than the Primary system administrator account).
- Change administrator account security settings.
- Delete obsolete accounts.

At times, you may need to modify the attributes of an administrator account. A Primary system administrator and people who are assigned to the Administrator 1 role can modify their own or another administrator's account.

Some changes take effect immediately. For example, a logged-in administrator's session is terminated when you disable or delete the administrator account.

Other changes affect the administrators the next time that they log in, for example, when you modify the administrator's password or modify the session timeout value.

Modify Administrator Account Attributes

1. Click the **System** icon in the top context pane.
2. Click the **Administrator Accounts** link in the left navigation pane.
3. Click the login name of the administrator account that you want to modify.
4. Enter values for the attributes that you want to modify.

   Note: You cannot disable the Primary system administrator account.
5. To save the modified account, click **OK**.

For parameter definitions, see the Administrator Configuration Page.

Change Administrator Passwords

You can change administrator passwords if they forget their password and cannot log into the system.

- Primary system administrators and administrators who are assigned to the Administrator 1 role can change the password of any administrator account.

  Tip: If you forget the Primary system administrator password, you can reset it in these ways:

  - Use a Type 1 Administrator account, if one exists, to reset the password. A support administrator cannot reset the Primary system administrator password.
  - Contact the Pillar World Wide Customer Support Center for the encrypted file (for resetting the password). The Pillar World Wide Customer Support Center will send you the encrypted file on a USB key and instruct you on installing the file.

- Administrators who are assigned to the Administrator 2 or Monitor roles can change their own passwords.

1. Click the **System** icon in the top context pane.
2. Click the **Administrator Accounts** link in the left navigation pane.
3. Click the login name of an administrator account.
4. Click **Change Password**.
5. Enter the new password in both fields.
6. To save the modified password, click **OK**.

For parameter definitions, see the Change Password Page.

About Modifying Administrator Account Security Settings

You can change the security settings for system administrator accounts, including:

- Set the number of failed login attempts that the Pillar Axiom system permits. When the threshold is exceeded, the system disables the account and writes an entry in the event log. Only a Primary system administrator or Administrator 1 account can re-enable the account, and the system resets the counter upon a successful login. If you do not set this value, there is no limit to the number of unsuccessful login attempts.
- Set the session timeout so that the Pillar Axiom system terminates an administrator's session after a given period of inactivity. If you do not set this value, inactive sessions are terminated after 20 minutes.

- Select Secure Session Only to specify that administrator access to the Pillar Axiom system is over secure HTTP sessions. Upload a secure sockets layer (SSL) certificate to the Pillar Axiom Pilot to authenticate logins.

Modify Administrator Account Security Settings

1. Click the **System** icon in the top context pane.
2. Click the **Security** link under the **Global Settings** heading in the left navigation pane.
3. Choose **Modify Security Settings** from the **Actions** drop-down list.
4. Select the **Secure Session Only** check box.
5. Click **Upload Certificate** and navigate to and select the secure sockets layer (SSL) certificate that you want to use (optional).
6. Enter the values in the login attempts and session timeout fields to define the administrator login limits.
7. To save the security settings, click **OK**.

For parameter definitions, see the Account Security Settings Page.

See also: About Administrator Account Creation.

Delete Administrator Accounts

You may need to delete an administrator account, for example when someone who has an account leaves the company.

1. Click the **System** icon in the top context pane.
2. Click the **Administrator Accounts** link in the left navigation pane.
3. Select one or more administrator accounts.
4. Choose **Delete Account** from the **Actions** drop-down list.
5. When prompted to confirm the deletion, click **OK** to delete the administrator account.

CHAPTER 8

Perform Maintenance Operations

About Pillar Axiom Support Tools

Someone from the Pillar World Wide Customer Support Center may request that you run one or more of the support tools and send the diagnostic output to Pillar Data Systems.

A Pillar Axiom system is fault tolerant. The system detects anomalies and automatically fails over to a partner component to maintain data availability. No intervention is required, unless a technician is needed to replace a hardware component.

Even fault-tolerant systems with a long mean time between failure (MTBF) rate cannot avoid component failure forever. If a component failure results in system instability, support tools are available to diagnose and fix the issue.

The following tables lists the tools that are provided to help you support a Pillar Axiom system.

Table 15 Pillar Axiom support tools

Tool category	Tools
Collect system information (and download files from the Pilot)	• Collect Debug Logs • Collect Event Logs • Collect Statistics
Verify system operation	• Verify Data Consistency • Verify Storage Redundancy
Resolve system trouble	• Clear System Configuration • Reset System Serial Number
Resolve connectivity trouble	• Resolve Connectivity Trouble

Collect Debug Logs

If a Pillar Axiom hardware component fails, the system writes debug logs so that the issue can be investigated. The Pillar World Wide Customer Support Center may request that you collect the debug logs and send them to Pillar Data Systems for analysis. The logs are not customer-readable.

1. Click the **Support** icon in the top context pane.
2. Click the **Collect System Information** link in the left navigation pane.
3. Select **Collect System Information** from the **Actions** drop-down list.
4. Select the check box to the left of Debug Logs.
5. Click **Options** and select the specific types of debug logs to collect (optional). If you do not click **Options**, all debug logs are selected by default.
6. Click **OK**.
7. Click **OK** again.
8. When prompted that the log-file creation takes time, click **OK** to acknowledge the message.

   The system compresses the debug logs from the selected components into Bzip files and then compresses them into a file named `CHn.tar`.
9. When “Status: Available” appears on the Collect System Information page, select the tar file.
10. Select **Download Information To Client** from the **Actions** drop-down list.
11. When prompted, select a directory available on your workstation as the save-to location.

    Tip: If you are using Internet Explorer on a client system, you may encounter problems downloading large log files from the Pillar Axiom system. Encountering a problem can happen if there are multiple browser windows open at the same time. Try closing one or more of the browser windows.

For parameter definitions, see Collect System Information Page.

Collect Event Logs

If you typically filter the display of event log entries as you work, you may want to collect all logged events from the event log, as well as logs for the management interfaces (GUI and CLI). On occasion, the Pillar World Wide Customer Support Center may request that you collect all event information and send the file to Pillar Data Systems for analysis.

1. Click the **Support** icon in the top context pane.
2. Click the **Collect System Information** link in the left navigation pane.
3. Select **Collect System Information** from the **Actions** drop-down list.
4. Select the check box to the left of Event Log and click **OK**.
5. When prompted that the log-file creation takes time, click **OK** to acknowledge the message.
6. When “Status: Available” appears on the Collect System Information page, select the Event Log.
7. Select **Download Information To Client** from the **Actions** drop-down list.
8. When prompted, select a directory on a local drive as the save-to location.

For parameter definitions, see Collect System Information Page.

Collect Statistics

The Pillar Axiom storage system generates performance statistics for filesystem backups, logical volumes, and network attached storage (NAS) and storage area network (SAN) protocols. The Pillar World Wide Customer Support Center may request that you collect performance statistics and transmit the data to Pillar Data Systems for analysis.

1. Click the **Support** icon in the top context pane.
2. Click the **Collect System Information** link in the left navigation pane.
3. Select **Collect System Information** from the **Actions** drop-down list.
4. Select the check box to the left of Statistics and click **OK**.

5 When prompted that the log-file creation takes time, click **OK** to acknowledge the message.

6 When “Status: Available” appears on the Collect System Information page, select the Statistics Log.

7 Select **Download Information To Client** from the **Actions** drop-down list.

8 When prompted, select a directory on a local drive as the save-to location.

For parameter definitions, see Collect System Information Page.

Verify Data Consistency

Data written to logical volumes (filesystems and LUNs) is stored on drives in certain Bricks. You can use one of the support tools to verify data consistency on the Bricks.

1 Click the **Support** icon in the top context pane.

2 Click the **Verify System Operation** link in the left navigation pane.

3 Select **Data Consistency** from the list of verification tools.

4 Select **Verify and/or View Results** from the **Actions** drop-down list.

5 Select a Brick to verify and click **OK**.

6 Select the priority level to specify how much I/O time to give to this operation.

For example, if you specify High Priority, this operation could impact Brick performance by 30%.

7 When prompted to confirm the verify, click **OK**.

For parameter definitions, see Verify System Operations Page.

Verify Storage Redundancy

Mirror copies are maintained for filesystems along with the original data.

When a drive is replaced, the Pillar Axiom storage system adds that newly discovered drive to appropriate mirror sets to restore full redundancy to the filesystem. A drive is rebuilt when it is added to a mirror set and data is copied to the new drive.

If a second failure occurs before the newly discovered drive is added to mirror sets, the second drive that failed is marked offline. The rebuild process on the new drive is stopped. Use the storage redundancy verification tool if this situation occurs.

1. Click the **Support** icon in the top context pane.
2. Click the **Verify System Operation** link in the left navigation pane.
3. Select **Storage Redundancy** from the list of verification tools.
4. Select **Verify and/or View Results** from the **Actions** drop-down list.
5. Select a filesystem to verify.
6. When prompted to confirm the verify, click **OK**.

For parameter definitions, see Verify System Operations Page.

Clear System Configuration

In extremely rare circumstances, you may need to reset your system configuration.

Prerequisites:

A special encryption file from the Pillar World Wide Customer Support Center.

This special encryption file is time sensitive and must be used by the date that the Pillar World Wide Customer Support Center provided you. This file performs the following actions:

- Deletes all data stored on the Pillar Axiom storage system.
- Resets the configuration to an initial state.
- Resets the system serial number.

Caution Because this action deletes all user data along with the system configuration, the system prompts you to confirm the operation. Be absolutely sure you want to take this action, because all data in your system will be lost.

1. Click the **Support** icon in the top context pane.
2. Click the **Resolve System Trouble** link in the left navigation pane.
3. Click the **Clear System Configuration** link on the Resolve System Trouble page.

4 Read the WARNING text and, when you are ready to proceed, click **Browse**.

5 Navigate to and select the encrypted configuration file that you received from the Pillar World Wide Customer Support Center.

6 Click **OK**.

7 When prompted to confirm the deletion of all data and system configuration, click **OK** to reset the system serial number.

For parameter definitions, see Resolve System Trouble Page.

Reset System Serial Number

In extremely rare circumstances, you may need to reset your system serial number.

Prerequisites:

A special encryption file from the Pillar World Wide Customer Support Center.

This special encryption file is time sensitive and must be used by the date that the Pillar World Wide Customer Support Center provided you. This file performs the following actions:

- Deletes all data stored on the Pillar Axiom storage system.
- Resets the configuration to an initial state.
- Resets the system serial number.

Caution Because this action deletes all user data along with the system configuration, the system prompts you to confirm the operation. Be absolutely sure you want to take this action, because all data in your system will be lost.

1 Click the **Support** icon in the top context pane.

2 Click the **Resolve System Trouble** link in the left navigation pane.

3 Click the **Reset System Serial Number** link on the Resolve System Trouble page.

4 Read the WARNING text, and when you are ready to proceed, click **Browse**.

5 Navigate to and select the encrypted configuration file that you received from the Pillar World Wide Customer Support Center.

6 Click **OK**.

7 When prompted to confirm the deletion of all data and system configuration, click **OK** to reset the system serial number.

For parameter definitions, see Resolve System Trouble Page.

Resolve Connectivity Trouble

Use the Resolve Connectivity Trouble page to identify any communication issues between the Pillar Axiom storage system and the customer network.

1 Click the **Support** icon in the top context pane.

2 Click the **Resolve Connectivity Trouble** link in the left navigation pane.

3 Select the Slammer.

4 Select the control unit.

5 Enter the command in the **Command Line** field.

6 Enter any needed environment variables.

7 Click **Execute**.

For parameter definitions, see Resolve Connectivity Trouble Page.

About Responding to Administrator Actions

Some configuration events in a Pillar Axiom system require administrator intervention to resolve the underlying issue.

The system notifies you of an “administrator action required” event by displaying an exclamation-point (!) icon at the bottom of pages in the GUI. When you click the exclamation-point icon, the GUI displays:

- Information about the event and the time that it occurred
- A recommended action to resolve the issue
- A status field that identifies whether the action has been performed

To resolve the issue, perform the recommended action.

About Filesystem Consistency

The Pillar Axiom filesystem-consistency check is similar to the filesystem check (`fsck`) command that is available on UNIX and Linux operating systems. If filesystem inconsistency is detected, the system performs the following actions:

- Firstly, the system takes the filesystem offline so that no further changes can be made to the data.
- Secondly, the system generates an action-required notification so that you can perform a consistency check.
- Finally, the system checks the filesystem for consistency.
    - If the filesystem is consistent, the check is complete. The system automatically puts the filesystem online.
    - If there are fixable inconsistencies, the Pillar Axiom system attempts to fix the filesystem and reports a completion status.
    - If the filesystem is inconsistent, the system reverts to the last known good Snap FS. The filesystem is not automatically put online, because you might choose to revert to an earlier Snap FS than the latest one.

      Note: If no good Snap FSs exist, restore the filesystem from a backup tape.

Check Filesystem Consistency

Users must remount the filesystem after the consistency check is complete and the filesystem is available again.

1. Click the **Administrator Actions** icon (yellow triangle with an exclamation point) in the status bar.
2. Select an appropriate event from the list on the Administrator Action Items page.
3. Select **Perform Recommended Action** from the **Actions** drop-down list.
4. Identify the time at which to run the verification check:
   - **Run Now** to verify the filesystem immediately.
   - **Schedule** to specify a day and time to verify the filesystem. Choose a start time and day of the week.
5. Click **Close**.

About Clearing Pinned Data

Pinned data can occur when issues arise regarding the Brick storage array. In such a case, data to be written to that array remains in the battery-backed memory of the Slammer storage controller.

Each logical volume (filesystem or LUN) maintains a time-ordered record of committed transactions (set of modified blocks). These records are kept within a dedicated area of the battery-backed memory that belongs to the owning Slammer control unit (CU). The system continuously (but asynchronously) flushes these records to the appropriate Bricks in the background. For NAS filesystems, these records are managed within what are called *journals*. For SAN LUNs, these records are managed within what are called a *write cache*.

Note: The journal and write cache reside on the same Slammer CU as the volume itself. A mirror of the journal or cache is kept on the partner CU. The mirror allows the system to recover from a failure of the owner CU.

An administrator-initiated shutdown request will fail if any user data is still cached and has not yet been written to physical storage. If the Slammers cannot communicate with the Bricks to flush the cached data, the Pillar Axiom storage system retains, or pins, the data in cache.

If you receive a message about pinned data when you initiate a shutdown request, check the **Health** page for details about the Bricks. Resolve any hardware issues that may exist. Hardware issues can prevent communication between Slammers and Bricks and prevents the flushing of the cached data to storage.

Note: If you need additional help, contact the Pillar World Wide Customer Support Center for more information on clearing the pinned data.

Shut Down the Pillar Axiom System

The Pillar Axiom storage system is composed of many hardware components and software processes that have dependencies on other components and processes. To ensure that these dependencies are satisfied and the Pillar Axiom system is shut down in an orderly fashion, use the **Shutdown/Restart** option.

While the system is in a shutdown state, the only actions you can perform are to display system status and to restart the system.

Important! If you need to power off the system for any length of time, remove the batteries.

1. Click the **System** icon in the top context pane.
2. Click the **Shutdown/Restart** link in the left navigation pane.
3. Select one of the following options from the **Actions** drop-down list:
   - **Shutdown Now**
   - **Shutdown in 5 minutes**
   - **Shutdown in 10 minutes**
   - **Restart Now**

   Result:
   Administrators receive an alert notification saying that the Pillar Axiom system has been shut down.

For parameter definitions, see Shutdown/Restart Page.

CHAPTER 9

Display System Events and Performance Statistics

About System Events and Performance Statistics

You can review events that have occurred in the Pillar Axiom storage system. You can also display performance statistics for backups, logical volumes (filesystems and LUNs), or network attached storage (NAS) and storage area network (SAN) protocols.

System Event Severities

The Pillar Axiom system generates events and classifies them by severity.

Table 16 Pillar Axiom event severities

Severity	Explanation
Critical	Access to data is compromised.
Error	Administrator action is required to prevent a hard error.
Warning	Administrator action is required to prevent a soft error from becoming a hard error or critical event.
Informational	A configuration change has been detected or another non-error event has occurred.

Display the Event Log

Review the event log to monitor events that have occurred in the Pillar Axiom system.

1. Click the **Health** icon in the top context pane.
2. Click the **Event Log** link in the left navigation pane.
3. Review the event log details to ensure that the information is what you expect.
4. Choose **Set Filter** from the **Actions** drop-down list (optional).
5. If you have set a filter, enter filter criteria that meets your needs and click **OK**.

For parameter definitions, see the Event Log Page.

See also: Create Alerts.

Filter Event Log Entries

You may not want to see all entries in the event log as you work. Filter the events to limit the type of events that appear.

1. Click the **Health** icon in the top context pane.
2. Click the **Event Log** link in the left navigation pane.
3. Select **Set Filter** from the **Actions** drop-down list.
4. Select one of:
   - **All Events and Processes**
   - **System-Generated Events**
   - **Background Processes**
5. Select one or more event severities to display.
6. Select the **Display Events Occurring On or After** option and enter a date to select events that occurred on this date or later (optional).
7. Click **OK**.
8. Click **Refresh Now** to update the Event Log page with events that match your filter criteria.

For parameter definitions, see the Event Filter Page.

About Performance Statistics

You can display performance statistics for backups, logical volumes (filesystems or LUNs), or network attached storage (NAS) and storage area network (SAN) protocols. Performance statistics are affected by usage patterns and Quality of Service (QoS) settings. For example, if the QoS settings for a filesystem are configured for a large number of operations per second and only a few people are accessing the storage device, the performance statistics show fewer operations per second.

Display Performance Statistics

1. Click the **Health** icon in the top context pane.
2. Click the item below the **Performance** link in the left navigation pane for which you want to display the performance statistics.
3. Review the performance statistics to ensure that the performance of each item is what you expect.

   Note: To update the display, click **Refresh Now**.

For parameter definitions, see the following:

- Performance Backup Page
- Performance Filesystems Page
- Performance NAS Protocols CIFS/NFS Page

APPENDIX A

GUI Field Definitions

About GUI Field Definitions

This section is a reference guide to the pages in the Graphical User Interface (GUI). This reference section supplements the previous task-based sections.

- Use the field definitions to decide what entries to make on the GUI pages.
- Use the Pillar Axiom ranges to learn about quantity, data type, and length ranges for field definitions. See Ranges for Field Definitions.

The GUI pages are listed in alphabetic order. The alphabetic order does not necessarily match the order in which you access the pages to complete administrative tasks.

Tip: On any page in the graphical user interface (GUI), click the **Help** icon in the context pane or **Help**, which is located at the bottom of a page, to display context-sensitive help about that page.

Ranges for Field Definitions

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

Table 17 Quantity ranges

Object	Quantity range
File Servers	Maximum: • 4, for a NAS Slammer on a Pillar Axiom 300 system. • 8, for a NAS Slammer on a Pillar Axiom 500 or Pillar Axiom 600 system. **Note:** In multi-Slammer systems, virtual interfaces (VIFs) that are associated with a File Server can be configured on multiple Slammers. The presence of VIFs is what counts against the limit. Such a File Server is considered to be present on each Slammer on which it has VIFs. **Note:** Virtual local area network (VLAN) tagging does not need to be enabled for more than one File Server. If VLAN tagging is enabled, File Servers do not require a unique VLAN tag.
Virtual interfaces (VIFs) for each File Server	Minimum: 1 Maximum: 32
VIFs for each Slammer port	Maximum: 16 **Note:** A particular virtual interface (VIF) may belong to any File Server
VLANs for each File Server	Minimum: 0 Maximum: 32

Table 17 Quantity ranges (continued)

Object	Quantity range
Network routes for each File Server	Minimum: 0 Maximum (default): 8 Maximum (static): 16
NIS configuration file size	Maximum: 50 MB **Note:** Size limit for each Network Information Service (NIS) file (`/etc/passwd`, `/etc/group`, and `/etc/netgroup`) that is uploaded to the Pilot.
Upload file size	Maximum: 650 MB
Volume groups	Minimum: 1 Maximum: 5000 **Note:** A volume group can contain up to 100 nested groups. Nesting is limited to four levels. Also, the root volume (`/Volumes`) is always available.
Filesystems	Minimum: 1 Maximum (system): 1024 Maximum (NAS Slammer): 1024 **Note:** Clone FSs factor into these limits.
Filesystem size	Minimum: 1 to 2 GB. The exact value depends on these factors • Brick type (Fibre Channel or SATA) • RAID geometry (RAID 5 or Distributed RAID) • Strip size (1 MB or normal) Maximum: System capacity **Note:** All capacity values must be in increments of 1 GB.
Snap FSs	Maximum (for a filesystem): 250 Maximum (for a Pillar Axiom system): 16,000
Pillar Axiom SecureWORMfs retention period	Minimum: 0 days to 1000 years Maximum: 0 days to 1000 years Default: 0 days to 1000 years **Note:** Maximum must be greater than or equal to the minimum.

Table 17 Quantity ranges (continued)

Object	Quantity range
	Note: Default must be greater than or equal to the minimum and less than or equal to the maximum.
NFS exports	Maximum: 1000 for each File Server
NFS host entries	Maximum: 4000 for each File Server
CIFS shares	Maximum: 128 for each File Server
CIFS connections	Maximum for each NAS Slammer (specified memory is the total combined memory of both control units): • 400 for 6 GB memory (Pillar Axiom 300 systems only) • 1200 for 12 GB memory • 6000 for 24 GB memory • 12,000 for 48 GB memory (Pillar Axiom 600 systems only)
CIFS security groups	Maximum: 1024 for each Common Internet File System (CIFS) user
SAN LUNs	Maximum: • 4096 visible for any given SAN Slammer • 4096 visible across all SAN Slammers in a given system (1024 if all LUNs have non-zero clone repositories) • 255 visible for each host **Note:** A visible (active) SAN LUN requires one virtual LUN (VLUN). A clone of a SAN LUN requires a VLUN for the metadata and another for the data repository. If that clone is active, a third VLUN is required, making a total of four VLUNs for the SAN LUN and its clone.
SAN LUN size	Minimum: 1 to 2 GB. The exact value depends on these factors • Brick type (Fibre Channel or SATA) • RAID geometry (RAID 5 or Distributed RAID) • Strip size (1 MB or normal) Maximum: System capacity **Note:** All capacity values must be in increments of 1 GB.

Table 17 Quantity ranges (continued)

Object	Quantity range
Pillar Axiom Path Manager (APM)	Maximum Pillar Axiom systems: 8 for each SAN host
APM data paths	Maximum: 32 to each LUN
APM FC HBA ports	Maximum: 32 for each SAN host
Clone LUNs	Maximum: • Number of available LUNs • 13 active at a time (for a single source)
iSCSI	Maximums for each iSCSI port: • 256 TCP connections • 256 iSCSI initiators • 512 simultaneous commands Maximum for each LUN: 32 persistent reservation registration keys
Administrator accounts	Minimum: 2 Maximum: 23 **Note:** Minimum provides for the Primary system administrator and system administrator
Administrator sessions	Maximum: 10 simultaneous **Note:** Two sessions are reserved for the Primary system administrator and system administrator.
NDMP sessions	Maximum: 10 concurrent

Table 18 Data type and length ranges

Field	Length or Type	Notes
Names for: • Alerts • Brick storage enclosures • File Servers	1 through 16 8-bit Unicode Transformation Format (UTF-8) printable characters.	Embedded spaces are permitted. Invalid characters: • Non-printable characters, including ASCII 0 through 31 • / (slash) and \ (backslash)

Table 18 Data type and length ranges (continued)

Field	Length or Type	Notes
• Filesystems • Pillar Axiom system • Schedules • Slammer storage controllers • Volume groups	UTF-8 is described in RFC 2279, which you can find online with any internet search engine.	• . and .. (dot and dot-dot alone) • Embedded tabs Pillar Axiom processing: • Leading and trailing white space is stripped • Comparison is case sensitive **Tip:** Names of filesystems that you export to NFS users should contain only US-ASCII characters. **Note:** You can have filesystems with the same name if the filesystems are not in the same volume group or File Server.
Names for LUNs	1 through 82 UTF-8 printable characters	Invalid characters: • Nonprintable characters, including ASCII 0 through 31 • / (slash) and \ (backslash) • . and .. (dot and dot-dot alone) • Embedded tabs
Names for SAN hosts	1 through 63 UTF-8 printable characters	
Names for: • DNS domains • NIS domains	1 through 256 UTF-8 printable characters	
Snap FS name	1 through 26 UTF-8 printable characters	Invalid characters: • / (slash) and \ (backslash) • . and .. (dot and dot-dot alone) • Embedded tabs
Snap FS base (mount) name	8 through 33 UTF-8 printable characters	Invalid characters: • / (slash) and \ (backslash) • . and .. (dot and dot-dot alone) • Embedded tabs

Table 18 Data type and length ranges (continued)

Field	Length or Type	Notes
Administrator user name	1 through 16 UTF-8 printable characters	Case-sensitive value Invalid characters: • Embedded spaces • / (slash)
Administrator password	6 through 16 UTF-8 printable characters	• Case-sensitive value • Embedded spaces are permitted.
Administrator login attempts	1 through 20 (integer)	
Optional entries for administrator full names	0 through 40 UTF-8 printable characters	Embedded spaces are permitted.
Optional entries for telephone numbers	0 through 80 UTF-8 printable characters	Embedded spaces are permitted.
Alert descriptions	0 through 80 UTF-8 printable characters	Embedded spaces are permitted.
Email address (emailuser@host)	1 through 64 characters for email user	a-z A-Z 0-9 ! # $ % & ' * + - / = ? ^ _ ` { | } ~ . are permitted, except that . (dot) cannot be the first or last character.
	1 through 255 characters for host	a-z A-Z 0-9 - . are permitted, except that: • 0-9 - . cannot be the first character. • . - cannot be the last character. An IP address cannot be the host part of the email address.
NDMP account user name	1 through 16 UTF-8 printable characters	Case-sensitive value Invalid characters: • Embedded spaces • / (slash)
NDMP account password	6 through 8 ASCII printable characters	Case-sensitive value

Table 18 Data type and length ranges (continued)

Field	Length or Type	Notes
Names for CIFS: • Servers • Domains	1 through 15 ASCII printable characters 33 through 126	
Comments for CIFS servers	1 through 44 ASCII printable characters 32 through 126	Embedded spaces are permitted.
Names for CIFS shares	1 through 80 ASCII printable characters 32 through 126	Embedded spaces are permitted. Invalid characters: • / (slash) and \ (backslash) • : (colon) • control character
Comments for CIFS shares	0 through 256 ASCII printable characters 32 through 126	Embedded spaces are permitted.
CIFS administrator (for domain controller): • User name • Password	0 through 256 UTF-8 characters	Case-sensitive value Invalid characters: • Embedded spaces • / (slash)
Directory paths for CIFS shares	1024 bytes and start with a \ (backslash)	Path includes a filesystem name, which can consist of up to 40 UTF-8 printable characters, plus a NULL terminator.
Directory paths for NFS exports	UTF-8 characters up to 1024 bytes in length; start with a / (slash)	Path includes a filesystem name, which can consist of up to 40 UTF-8 printable characters, plus a NULL terminator.
NFS host name	UTF-8 characters up to 255 bytes in length	Host format: • IP address in dotted-decimal format • Subnet address with both the subnet and mask in dotted-decimal format • Host name

Table 18 Data type and length ranges (continued)

Field	Length or Type	Notes
		• Asterisk (*), to export to all NFS clients (everyone)
IP addresses	0 through 255, in all four parts	IP version 4 (IPv4) dotted-decimal notation (xxx.xxx.xxx.xxx)
Virtual LAN (VLAN) ID (tag)	0 through 4094 (integer)	• 1 through 4094 denote that VLAN Tagging is enabled. • 0 denotes that VLAN Tagging is disabled.
SNMP community string	0 through 255 ASCII printable characters 33 through 126	Invalid characters: • Embedded spaces • Control characters
Chap Secrets	100 UTF-8 characters	Non-character (for example, integer) CHAP secret values are not supported. CHAP secrets should be more than 12 bytes if IPsec is not used on insecure network segments.

Account Security Overview Page

DESCRIPTION	Use the Account Security Overview page to review the security configuration for browsers and administrator accounts.
FIELD DEFINITIONS	**System Name** Identifies the name of the Pillar Axiom system.
BROWSER CONFIGURATION	**Browser Sessions** Identifies whether administrators must manage the Pillar Axiom system in secure GUI sessions. **Certificate Hostname** Identifies the host on which the secure sockets layer (SSL) certificate resides. **Certificate Expiration:** Identifies the date and time when the SSL certificate expires.
ACCOUNT SECURITY	**Consecutive failed login attempts allowed** Identifies the number of times that an administrator can attempt, but fail, to log in to the Pillar Axiom system. **Session timeout period for all administrators** Identifies the inactivity time limit after which an administrator's session is terminated.
SEE ALSO	Modify Administrator Account Security Settings Ranges for Field Definitions

Account Security Settings Page

DESCRIPTION

Use the Account Security Settings page to set the number of login attempts and session timeouts for administrator accounts. When an account is inactive for the session timeout period, the Pillar Axiom system automatically logs the account out of the system. The session timeout period:

- Applies only in property sheets and secondary windows in the Pillar Axiom Storage Services Manager.
- Does not apply in the main window because of activity that occurs to verify the system status and health.

FIELD DEFINITIONS

Secure Session Only

Identifies whether administrators must manage the Pillar Axiom system in secure GUI sessions.

- Enable this option to permit administrators to access the system only in secure sessions.
- Disable this option to permit administrators to access the system in insecure sessions.

Upload Certificate

Permits you to navigate to and select a secure sockets layer (SSL) certificate so that you can:

- Upload the certificate to the Pilot.
- Use the certificate so that administrators can manage the Pillar Axiom system in secure GUI sessions only.

Consecutive Failed Login Attempts Allowed

Identifies the number of times that an administrator can attempt, but fail, to log in to the Pillar Axiom system.

Session Timeout Period for All Administrators

Identifies an inactivity time limit, after which an administrator's session is terminated. In-progress sessions are not affected by changes that you make to the value; current sessions use the old value. Sessions that start after you change the value use the modified session timeout.

The default for the session timeout is 20 minutes.

If the time period contains any values lower than a day, you must include the “T” separator. For example, 2 hours=PT2H; 2 hours and 30 minutes = PT2H30M, where P = period of time; T = time; H = hours, and M = minutes. Other session timeout elements include Y = years and D = days.

Login Screen Message

Specifies a message that is displayed when system administrators log in to the Pillar Axiom storage system. You can enter only 7-bit ACSII characters 32-127

which includes the basic Latin character set but does not include carriage return or other control characters.

SEE ALSO

Modify Administrator Account Security Settings
Ranges for Field Definitions

Additional Hosts Page

DESCRIPTION

Use the Additional Hosts page to create multiple hosts for an NFS export. Each host entry allows you to specify a different access privilege on the NFS export.

FIELD DEFINITIONS

Host Access

Identifies the NFS clients, or hosts, that can mount the NFS export.

- **All Hosts**: Everyone can access the export point.
- **Single Host**: Only the specified host can access the export point.
- **NIS Netgroup**: Everyone within the NIS netgroup can access the export point.
- **Network**: Everyone on the specified subnet can access the export point.
- **Read Only**: Identifies whether the export is made available to users in read-only or in read and write mode.
  - Enable this option so that the export operates in read-only mode.
  - Disable this option so that the export operates in read-write mode.
- **Root Access**: Identifies whether users who are logged in to a client workstation as root are root users on the export.
  - Enable this option so that root users on clients are root users. This is comparable to no_root_squash in Linux root-squashing options.
  - Disable this option so that root users on clients are identified as the user nobody. This is comparable to the root_squash option.
- **Host Processing Order**: Specifies the permissions based on the first matching host entry.

Unix Exports (NFS) List

Displays a list of currently configured NFS exports. Check the box to the left of an export point to select it for modification or deletion.

SEE ALSO

Modify NFS Exports
Add NFS Exports to a Filesystem
Modify NFS Exports
Ranges for Field Definitions

Administrator Accounts Overview Page

DESCRIPTION

Use the Administrator Accounts overview page to review the administrator accounts that are configured on the Pillar Axiom storage system. The **Actions** drop-down list provides options to create, modify, delete, and view administrator accounts.

FIELD DEFINITIONS

Login Name

Lists administrator login, or user names. Click a name to review or modify the administrator account.

Role

Identifies which role is assigned to an administrator account. A role defines which permissions are granted to the administrator.

Account Active

Identifies whether the administrator account is enabled.

- Enabled accounts are active. Administrators whose accounts are enabled can log in to the Pillar Axiom storage system.
- Disabled accounts are inactive. Administrators whose accounts are disabled cannot log in.

Full Name

Identifies the administrator's full name (“first” and “last”).

Email Address

Identifies the recipient's email address. The email server to which the Pillar Axiom system sends alerts must be able to send messages to this email address.

Phone Number

Identifies the recipient's telephone number. The Pillar Axiom system does not verify the validity of this entry.

Note: Enclose telephone numbers with spaces in quotes.

SEE ALSO

Create an Administrator Account
Display Administrator Accounts
Modify Administrator Account Attributes
Delete Administrator Accounts
Ranges for Field Definitions

Administrator Configuration Page

DESCRIPTION

Use the Administrator Configuration page to create, modify, and disable a specific administrator account.

FIELD DEFINITIONS

Login Name

Identifies the administrator's login (user) name.

Role

Identifies the administrator's role. A role defines the administrator's privileges. Choose from:

- Administrator 1, if the person can perform all configuration and administration tasks.
- Administrator 2, if the person can perform all tasks except create, modify, and delete administrator accounts and File Servers; modify global, SNMP, and NDMP settings; modify software or hardware configurations; or shut down the system.
- Monitor, if the person can display information only, and cannot modify the configuration.

For the following predefined roles, you cannot assign administrators to them and you cannot delete them:

- Primary system administrator has the same privileges as the Administrator 1 role.
- support administrator has the same privileges as the Monitor role, as well as privileges to perform support-related tasks.

Password *(Create only)*

Identifies the login password. Passwords are case sensitive, and blank passwords are not permitted.

Retype Password *(Create only)*

Enter the password again.

Full Name

Identifies the administrator's full name (“first” and “last”).

Email

Identifies the recipient's email address. The email server to which the Pillar Axiom storage system sends alerts must be able to send messages to this email address.

Phone

Identifies the recipient's telephone number. The Pillar Axiom system does not verify the validity of this entry.

Note: Enclose telephone numbers with spaces in quotes.

Disable Account

Removes the administrator's ability to access the Pillar Axiom system. This setting takes effect immediately. If the administrator is logged in when you disable the account, the system logs out the administrator immediately.

You cannot disable the Primary system administrator.

Change Password *(Modify only)*

Opens the Change Password Page.

Permit SSH capability *(Support roles only)*

Allows Pillar Data Systems customer service personnel to access the Pillar Axiom system using SSH. Also, the Support Tool and SSH Access license keys must be installed to allow SSH access.

Note: Only Primary and Administrator 1 roles can enable this option. This option can be enabled for Support accounts only.

SEE ALSO

Create an Administrator Account
Modify Administrator Account Attributes
Ranges for Field Definitions

Alert Details Page

DESCRIPTION

Use the Alert Details page to create and modify alerts that are generated when specified events occur.

FIELD DEFINITIONS

Alert Name

Identifies the name of an alert. When an alert is triggered, the Pillar Axiom system sends a notification to the designated recipients.

Description

Describes the alert.

Recipients

Identifies the email addresses of the recipients who are to receive alerts. The email server to which the Pillar Axiom system sends alerts must be able to send messages to these email addresses.

Test Email

Sends a message to the specified email addresses to test recipient email addresses. Recipients should look for a message that is titled "[Axiom-QoS] Test email" in their email in-boxes.

Event Categories to Trigger Alert

Identifies a category of event severities. Choose from:

- **Critical**, which means that access to data is compromised (for example, two drives in a RAID array have failed) or data loss has occurred.
- **Error**, which means that administrator action is required to prevent a "hard" error (for example, a single drive in a RAID array has failed) from becoming a critical event.
- **Warning**, which means that administrator action is required to prevent a "soft" error from becoming an error or critical event.
- **Informational**, which means that a configuration change has been detected or another non-error event has occurred.

For descriptions of events that trigger alerts, see System Event Severities.

Event List

Lists events that are defined in the selected categories.

For details, see System Event Severities.

SEE ALSO

Create Alerts
Modify Alerts
Ranges for Field Definitions

Alerts Overview Page

DESCRIPTION

Use the Alerts overview page to review the alerts that are configured on the Pillar Axiom system. The **Actions** drop-down list provides options to create, modify, delete, and view alerts.

An alert is defined as:

> A Simple Mail Transfer Protocol (SMTP) email message that notifies recipients of specified system events, such as informational, warning, or critical events. Alerts are optional and supplement normal event logging and Call-Home notification.

FIELD DEFINITIONS

Email Notification

Identifies whether email is enabled to receive Pillar Axiom notifications that the system sends when specified alerts are generated.

Email Server IP Address

Identifies the IP address of the email server that receives administrator alert notifications.

Alert Name

Lists the names of configured alerts. Click a name to review or modify the alert settings.

Description

Displays the alert description text.

SEE ALSO

Create Alerts
Modify Alerts
Delete Alerts
Ranges for Field Definitions

Assign Local Groups Page

DESCRIPTION

Use the Assign Local Groups page to assign the Common Internet File System (CIFS) protocol to a local group.

FIELD DEFINITIONS

FileServer

Identifies the File Server.

Local Administration Group

Identifies the local administration group to be used for CIFS protocol for a Pillar Axiom storage system.

Local Backup Group

Identifies the backup group to be used for CIFS.

SEE ALSO

File Server Overview Page
Ranges for Field Definitions

Associate Hosts Page

DESCRIPTION

Use the Associate Hosts page to create a host-to-Fibre Channel (FC) HBA or host-to-iSCSI IQN (or iSCSI name) associations when you do not have the Pillar Axiom Path Manager installed on the host.

FIELD DEFINITIONS

Host Name

Identifies the SAN host that accesses LUNs configured on the Pillar Axiom storage system. If the Pillar Axiom Path Manager is not yet installed, the system displays the World Wide Name (WWN) of the FC HBA or the iSCSI name of the iSCSI device.

Add WWN or iSCSI Name

Allows you to enter an HBA port that the Pillar Axiom system does not yet detect:

- WWN for Fibre Channel (FC)
- iSCSI Name for iSCSI devices

Discovered WWNs or iSCSI Names

Identifies the WWN of the HBA ports that the Pillar Axiom system detects on the network.

Note: Because sometimes a McData switch may return an error for a valid command, the GUI may display a connection to a host port that is not connected.

Create

Creates an association between the specified storage area network (SAN) host and the WWNs of the HBA.

Modify

Changes the selected object's configuration settings.

Remove

Deletes the selected objects.

SEE ALSO

Associate Hosts
Modify SAN Host Settings
Ranges for Field Definitions

Bricks Overview Page

DESCRIPTION

Use the Bricks overview page to review the Bricks that are part of the Pillar Axiom system.

There are three types of Brick, serial ATA (SATA), solid state drive (SSD), and Fibre Channel (FC). FC Bricks come in two flavors: RAID and Expansion. SATA, SSD, and FC RAID Bricks contain two RAID controllers and at least 12 drives. FC Expansion Bricks do not have a RAID controller but instead rely on the controller within the FC RAID Brick.

SATA and SSD Bricks have 13 drives where the 13th drive is used as a spare for automatic failover. FC Bricks do not have a dedicated spare; any drive can be utilized as a spare.

The **Actions** drop-down list provides different options based on the context from which you open the page. When you click the Brick link in the navigation pane from the:

- **Health** context, you can view the status of and change the name that is assigned to a Brick.
- **Support** context, you can repair Brick hardware.

FIELD DEFINITIONS

Brick Name

Lists the names of hardware components. Click a name to display details about that hardware component.

Type

Lists the type of Brick: SATA, SSD, or FC.

Status

Displays the current status of a hardware component. A status of Normal requires no action.

Disk Drives

Displays the current status of a hardware component's drives.

I/O Ports

Displays the current status of a hardware component's input and output (I/O) ports.

Power Supplies and Fans

Displays the current status of a hardware component's power supplies and fans.

Temperature

Displays the current status of a hardware component's temperature. Watch for too high or low temperatures.

SEE ALSO

Manage Hardware Components
Display Hardware Component Information
Modify Hardware Component Names
Identify Hardware Components
Replace a Hardware Component
Ranges for Field Definitions

Call-Home Configuration Page

DESCRIPTION

Use the Call-Home Configuration page to enable and configure the Call-Home settings. Call-Home is a feature that, when enabled, allows the system to send status and configuration information to Pillar World Wide Customer Support Center. No customer data is transmitted.

FIELD DEFINITIONS

Enable Call-Home

Specifies whether Call-Home should be enabled.

- Enable Call-Home support so that the Pillar Axiom system sends status messages to Pillar World Wide Customer Support Center. You can create alerts to send notifications to recipients within your organization.
- Disable Call-Home support so that the Pillar Axiom system does not send the status messages.

Call-Home Primary DNS

Identifies the Domain Name Sever (DNS) that is used to resolve IP addresses. This includes an email server that sends Call-Home messages from the Pillar Axiom storage system to Pillar Data Systems. It can also include the Pillar Data Systems Call-Home server, callhome.support.pillardata.com, that receives Call-Home messages by means of SCP or HTTPS. Call-Home messages contain information about system issues and failures.

Call-Home Secondary DNS

Identifies the IP address of the secondary Domain Name Service (DNS) server in your network that is consulted if the primary server cannot be reached when outward-bound messages are sent.

Disallow Large Files

Identifies whether trace logs and performance statistics are included in Call-Home notifications.

- Disallow large files so that the Pillar Axiom system excludes trace logs and performance statistics from the Call-Home messages that are sent to Pillar World Wide Customer Support Center. You can collect, download, and transmit the trace logs separately if they are needed.
- Allow large files so that trace logs and performance statistics are automatically included in the messages that are sent to Pillar World Wide Customer Support Center.

PILLAR DATA SYSTEMS SERVER

Use Pillar Data Systems Server

Select this option to send Call-Home logs to the Pillar World Wide Customer Support Center.

SCP

Uses secure copy (SCP) with 1024-bit encryption and secure keys to transfer files over the internet.

HTTPS

Sends files either directly to the Pillar World Wide Customer Support Center or through a proxy server for security purposes. It can also be used when the Pillar Axiom system does not have direct access to the Internet. This option uses the internal HTTP, SOCKS_4, or SOCKS_5 access port on the proxy server. The proxy type and port number are dependent on your proxy configuration.

Enable Proxy Server

Select this option to send Call-Home logs through a proxy server for security purposes or when the Pillar Axiom system does not have direct access to the Internet.

Note: For proxy servers, Call-Home supports these characteristics:

- IP address
- Port number
- Type of protocol

Authentication to the proxy server is not supported. However, depending on your security policies, you may have a configuration option or directive that could be set that would not require proxy authentication from a particular set of IP addresses, such as those from the Pillar Axiom Pilot.

Proxy Server IP Address

Identifies the IP address of the proxy server.

Proxy Server Port

Identifies the port that is used by the proxy server to send the Call-Home log files.

Proxy Server Type

Identifies the type of protocol (SOCKS4, SOCKS5, or HTTP) that is used to access the proxy server.

LOCAL HOST

Use Local Host

Select this option to send Call-Home logs to a local server

Name

Identifies the name that you want to use to describe the local host.

IP Address

Identifies the IP address of a server on the local host.

Directory

Identifies the directory path on the target server in which to store the Call-Home log files.

Username

Identifies the name of the user.

Password

Enter the password of the username to access the local host.

SEE ALSO

Configure Notification Settings

Modify Call-Home Settings

Call-Home Logs Page

Ranges for Field Definitions

Call-Home Logs Page

DESCRIPTION

Use the Call-Home Logs page to review the ten most recent Call-Home log files. You can also select and download the Call-Home log files to the client system.

FIELD DEFINITIONS

Name

Identifies the name of the Call-Home file.

Status

Identifies the time status of the Call-Home log file.

Created

Identifies the time the Call-Home log file was created.

Size

Identifies the size of the Call-Home log file.

Type

Identifies the type of the Call-Home log file:

- **Event**: Log files are created when a component operates in degraded or fail mode.
- **Periodic**: Log files are created on a weekly basis.

SEE ALSO

Ranges for Field Definitions

Capacity Planning Wizard Page

DESCRIPTION

Use the Capacity Planning Wizard page to simulate Quality of Service (QoS) settings for particular types of storage use.

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

FIELD DEFINITIONS

Previous

Returns to the previous page so that you can review or modify your entries.

Next

Moves to the next page so that you enter more configuration definitions.

Cancel

Discards your entries. The configuration is unmodified.

SEE ALSO

Run the Pillar Axiom Capacity Planner

Create Filesystems (NAS Storage Systems)

Ranges for Field Definitions

Change Password Page

DESCRIPTION

Use the Change Password page to change your login password. The Primary system administrator and Administrator 1 accounts can also change passwords for:

- Other administrators.
- The NDMP account.

FIELD DEFINITIONS

Password

Identifies the login password. Passwords are case sensitive, and blank passwords are not permitted.

- If you opened this page from the Administrator Configuration Page, this value identifies the Pillar Axiom administrator's password. Administrator account passwords are between 6 and 16 characters long.
- If you opened this page from the NDMP Configuration Page, this value identifies the NDMP account's password. NDMP account passwords are between 6 and 8 characters long.

Retype Password

Enter the password again, exactly as you typed it in the previous field.

SEE ALSO

Change Password Page

Create an NDMP User Account

Ranges for Field Definitions

CIFS Shares Configuration Page

DESCRIPTION

Use the CIFS Shares Configuration page to create and modify Common Internet File System (CIFS) shares for a network attached storage (NAS) filesystem.

FIELD DEFINITIONS

Filesystem

Identifies the name of the filesystem with which the CIFS share is associated. You can define multiple shares for a filesystem.

Name

Identifies the name of a CIFS share. A filesystem must be shared before users can create or access data files. Share names must be unique within a File Server.

Path to Folder

Enter the full path to the CIFS share.

Comment

Describes the CIFS share.

Enabled

Identifies whether the CIFS share is enabled.

- Enabled shares are active. Users can access an enabled share point.
- Disabled shares are inactive. Users cannot access a disabled share point.

Existing Shares List

Displays a list of currently configured CIFS shares.

Apply *(Create only)*

Saves the configuration settings and clears the input fields so that you can define additional settings.

SEE ALSO

Modify CIFS Shares

Ranges for Field Definitions

Clone Activation Page

DESCRIPTION

Use the Clone Activation page to make a clone of a filesystem or LUN available to users.

When activating a Clone LUN, assign the volume a LUN number to make it available to all SAN hosts. You can then map the LUN to specific hosts.

When activating a Clone FS, assign the volume to a File Server to make it accessible to users.

FIELD DEFINITIONS

Volume Backup Name

Identifies the name of the inactive clone.

Type

Identifies the volume type of the clone.

Size

Displays the amount of space allocated to this clone.

Created

Identifies the date and time at which the clone was created.

Volume group location

Identifies the volume group in which the clone is located. Starts at the System (root) level and traverses through parent and child volume groups.

Name for Clone

Specifies the name of the volume.

File Server *(filesystems only)*

Specifies the name of the File Server to which this Clone FS is to be assigned.

LUN Number *(LUNs only)*

Specifies the LUN number to be assigned to the Clone LUN.

Host Access Protocols *(LUNs only)*

Specifies the storage area network (SAN) protocols allowed for accessing this Clone LUN. One or both of:

Fibre Channel
iSCSI

SEE ALSO

Activate a Clone

Create an Immediate Clone FS

Create an Immediate Clone LUN

Clone FS Overview Page

DESCRIPTION

Use the Clone FS Overview page to manage the clones of network attached storage (NAS) filesystems.

FIELD DEFINITIONS

Name

Lists the names of configured filesystems. Click a name to review the filesystem settings.

SecureWORMfs Retention

Identifies the retention settings defined on the filesystem:

- None: There are no retention settings on the filesystem.
- Standard: Files on the Pillar Axiom SecureWORMfs filesystem are protected based on the specified retention settings; however, you can delete protected files on the Pillar Axiom SecureWORMfs by deleting the entire filesystem.
- Compliance Files on the Pillar Axiom SecureWORMfs filesystem are protected based on the specified retention settings. You cannot delete the Pillar Axiom SecureWORMfs if any files are present on the filesystem.

Status

Displays the current status of a clone. The possible states are:

- Inactive: Cannot be accessed from the data path.
- Online: Fully accessible.
- Offline: Not accessible.
- Partial Offline: The actual redundancy level may be different from the redundancy level with which the clone was configured.
- Conservative: Write-back cache has been disabled so journaling has slowed.
- Degraded: All of the copies of a redundant volume are not available. If one copy is missing, it is not fully redundant. This can happen when a write to one copy of the array fails (which may be a 30 second time-out).

File Deletion Enabled

Identifies whether files can be deleted on the Pillar Axiom SecureWORMfs. Protected files cannot be deleted regardless of this setting. When this option is enabled, the files that can be deleted are expired files.

Redundancy

Identifies the redundancy level of the logical volume (filesystem or LUN) as standard or double.

File Server

Identifies the name of a File Server with which the filesystem is associated.

Number of CIFS Shares

Identifies the count of Common Internet File System (CIFS) shares that are defined for the specified filesystem.

Number of NFS Exports

Identifies the count of Network File System (NFS) exports that are defined for the specified filesystem.

Last Protected File Integrity Check

Identifies the date and time at which the system performed the most recent **Protected File Validation** check on the filesystem.

Retention

Identifies the retention period for all files on a Pillar Axiom SecureWORMfs filesystem. Valid values are:

- **Default**: 0 days-1000 years and must be greater than or equal to the minimum and less than or equal to the maximum retention value.
- **Maximum**: 0 days-1000 years and must be greater than or equal to the minimum retention value.
- **Minimum**: 0 days-1000 years.

Creation Time

Identifies the date and time in which the filesystem was created.

Creation Method

Identifies the method that was used to create the Pillar Axiom SecureWORMfs filesystem (generally used for audit purposes):

- **New**: Newly created.
- **Restore**: Restored from a backup image.
- **FSCK**: Recovered using FSCK.
- **Copy**: Created from a copy of the filesystem.
- **Unknown**: The creation method is not known, however, this usually applies to a filesystem that was created with an earlier version of the software.

SEE ALSO

Create an Immediate Clone FS

Ranges for Field Definitions

Clone LUN Overview Page

DESCRIPTION

Use the Clone LUN overview page to manage Clone LUNs. The **Actions** drop-down list provides different options based on the context from which you open the page. When you click the **Clone LUNs** link in the navigation pane from the:

- **Storage** context, you can view, modify, or create an immediate Clone LUN. You can also delete Clone LUNs when they are no longer needed.
- **Data Protection** context, you can create an immediate Clone LUN of the selected LUN.

FIELD DEFINITIONS

Name

Identifies the name that is assigned to a copy of a LUN. LUN copy and LUN names are provided for administration convenience.

Host Access

Identifies if the LUN is mapped to a specific host or if all hosts can access the LUN.

LUN Number

Identifies the unique number that is assigned to a LUN and can be accessed by all storage area network (SAN) hosts if the LUN is not mapped. If the LUN is mapped, the LUN number must be unique to the mapped SAN host.

Status

Displays the current status of a clone. The possible states are:

- **Inactive**: Cannot be accessed from the data path.
- **Online**: Fully accessible.
- **Offline**: Not accessible.
- **Partial Offline**: The actual redundancy level may be different from the redundancy level with which the clone was configured.
- **Conservative**: Write-back cache has been disabled so journaling has slowed.
- **Degraded**: All of the copies of a redundant volume are not available. If one copy is missing, it is not fully redundant. This can happen when a write to one copy of the array fails (which may be a 30 second time-out).

Redundancy

Identifies the redundancy level of the logical volume (filesystem or LUN) as standard or double.

Pinned Data

Identifies whether any user data is pinned in cache and cannot be written to permanent storage.

LUID

Displays the logical unit unique identifier (LUID) of the LUN.

SEE ALSO

Create an Immediate Clone LUN

Ranges for Field Definitions

Collect System Information, Debug Log Details Tab

DESCRIPTION

Use the **Debug Log Details** tab to specify which Slammers and Bricks are of interest, whether SAN hosts should be included, and whether the most recent or all logs should be collected for the Pillar Axiom storage system.

FIELD DEFINITIONS

Slammers and SAN Hosts

Lists the following sources for which system information can be collected:

- Pilot
- Slammers (by name)
- All SAN Hosts
- System Configuration

Bricks

Lists the Bricks by name for which system information can be collected.

Collection Period

Specifies the extent of information coverage for each selected source:

- Most recent log
- All logs

Current Download Bundle

Displays the status of the most recent download bundle.

Unselect All

Automatically removes all items from the selection list. You can then select the items one by one for which you want to collect debug logs.

Select All

Automatically adds all items to the selection list. You can then select the items one by one that you want to remove from the list.

SEE ALSO

Ranges for Field Definitions

Collect System Information Page

DESCRIPTION

Use the Collect System Information page to review collected system information and run the collection tools that are available on a Pillar Axiom storage system.

FIELD DEFINITIONS

Name

Lists the name of the bundle containing the collected system information.

Status

Identifies the status of the download bundle.

Created

Identifies the time and date at which the download bundle was collected from the Pillar Axiom system.

Size

Identifies the size of the download bundle.

Contents

Identifies the types of system information that have been collected and are included in the current download bundle.

Collection Period

Identifies the extent of information coverage for each selected source:

- Most recent log
- All logs

SEE ALSO

Collect Statistics

Collect Debug Logs

Collect Event Logs

Ranges for Field Definitions

Resolve Connectivity Trouble Page

Collect System Information, Summary Tab

DESCRIPTION

Use the **Summary** tab to specify the scope of the system information for the Pillar Axiom storage system that you want to collect.

FIELD DEFINITIONS

Collection Scope

Lists the types of system information to be collected:

- **Debug Logs**, provides a record of unexpected errors, exceptions, and tracing information.
- **Event Log**, provides a record of events that have occurred.
- **Statistics**, provides a collection of performance, capacity usage, and system health information.

Tip: Select Debug Logs by itself or in combination with one or more of the other information types.

Options

Allows you to specify the hardware components, including SAN hosts, that are of interest and the extent of log collection.

Call-Home Option

Allows you to select whether you want the logs to be sent to the Call-Home server as soon as the collection operation completes.

SEE ALSO

Ranges for Field Definitions

Command Line Interface Page

DESCRIPTION

Use the Command Line Interface page to download the Pillar Axiom CLI to your administrative workstation.

The CLI is a client-based application that you can use to run administrative commands from a shell or a script. Through the CLI, you can perform administrative tasks. You can start a session, submit one or more requests to the Pillar Axiom system, and end the session.

On some of the platforms, there are two CLI executables included in the download: `pdscli` and `axiomcli`.

The `pdscli` executable is a full-featured utility that is available on all platforms and supports all commands available in the system. The `pdscli` requires a more complete understanding of the Pillar Axiom system. System administrators can run individual commands, use templates, or integrate the commands into their own scripts. Sample input and output templates as well as representative XML is available for each command.

The `axiomcli` executable is available on some of the platforms and supports a subset of the commands. The `axiomcli` application minimizes the detailed system knowledge system administrators need to complete administrative tasks, such as creating filesystems. The `axiomcli` application is a shell for the `pdscli` and follows more general conventions used by other command line interface (CLI)s.

FIELD DEFINITIONS

Actions List

Displays the operating systems (OSs) on which the CLI runs. Select the OS of the workstation from which you will run the CLI:

- Linux x86
- Linux IA64
- Linux x86_64
- HP-UX PA-RISC
- HP-UX IA64*
- Solaris SPARC*
- Solaris 10 SPARC*
- Solaris 10 x86
- AIX*
- FreeBSD x86
- Windows*
- Mac x86
- Mac PowerPC

* Indicates `axiomcli` supported platforms.

SEE ALSO

Pillar Axiom CLI Reference Guide (for `pdscli`)

Pillar Axiom CLI Guide (for `axiomcli`)

Ranges for Field Definitions

Configuration Wizard Page

DESCRIPTION

Use the Configuration Wizard page to run the Configuration Wizard, which guides you through the initial setup and configuration of the Pillar Axiom system.

FIELD DEFINITIONS

Previous

Returns to the previous page so that you can review or modify your entries.

Next

Moves to the next page so that you enter more configuration definitions.

Cancel

Discards your entries. The configuration is unmodified.

SEE ALSO

Run the Configuration Wizard

Ranges for Field Definitions

Copy Filesystem Page

DESCRIPTION

Use the Copy Filesystem page to make a copy of a filesystem on a Pillar Axiom storage system.

FIELD DEFINITIONS

Name

Lists the names of configured filesystems.

Click a name to review the filesystem settings.

Total Capacity

Identifies the current capacity that is assigned to the specified filesystem.

Free Capacity

Identifies the current unused capacity that is available. Free capacity is total capacity minus allocated capacity.

Used Capacity

Identifies the current capacity usage of the object.

Status

Displays the current status of a filesystem. The possible states are:

- **Online**: The filesystem is online and normal.
- **Offline**: The filesystem is offline.
- **Partial Offline**: The actual redundancy of the filesystem may differ from the redundancy with which the filesystem was configured.
- **Conservative**: Write-back cache on the filesystem has been disabled so journaling has slowed.
- **Degraded**: All of the copies of a redundant filesystem are not available. If one copy is missing, it is not fully redundant. This can happen when a write fails (which may be a 30 sec time-out) to one copy of the array.

File Server

Identifies the name of a File Server with which the filesystem is associated.

Number of Snapshots

Identifies the number of data replicas that have been created for a logical volume (filesystem or LUN).

Snapshots Size (GB)

Identifies the size of the data replica, in GBs.

SEE ALSO

Copy Filesystems
Ranges for Field Definitions

Copy LUN Page

DESCRIPTION

Use the Copy LUN page to make a copy of a LUN on a Pillar Axiom storage system.

FIELD DEFINITIONS

Name

Identifies the name that is assigned to a LUN for administrative purposes. LUN names must be unique across the Pillar Axiom system, not just within its associated volume group.

Total Capacity

Identifies the capacity that is assigned to this LUN.

Free Capacity

Identifies the current unused capacity that is available. Free capacity is total capacity minus allocated capacity.

Used Capacity

Identifies the current capacity usage of the object.

Status

Displays the current status of a LUN. The possible states are:

- **Online**: The LUN is online and normal.
- **Offline**: The LUN is offline.
- **Partial Offline**: The actual redundancy level of the LUN may be different from the redundancy level with which the LUN was configured.
- **Conservative**: Write-back cache on the LUN has been disabled so I/O has slowed.
- **Degraded**: All of the copies of a redundant LUN are not available. If one copy is missing, it is not fully redundant. This can happen when a write fails (which may be a 30 second time-out) to one copy of the array.

SEE ALSO

Copy LUNs
Ranges for Field Definitions

Event Filter Page

DESCRIPTION

Use the Event Filter page to create and modify the event filters that are configured on the Pillar Axiom system.

FIELD DEFINITIONS

Display Events and Processes

Identifies a list of event and process types to display. Choose one of:

- **All Events and Processes**, which displays system-generated events and background processes.
- **System Generated Events**, which displays the system-generated events that are of the selected severity or severities.
- **Background Processes**, which displays background processes. Background processes are long-running configuration requests.

Display Event Severities

Identifies a list of event types. Choose from:

- **Critical events**, which require prompt action to prevent system failures or offline conditions.
- **Error events**, which require event-specific actions to correct the error condition.
- **Warning events**, which are minor conditions that you can address at your convenience.
- **Informational events**, which are informational only and require no action.

Display Events Occurring After

Identifies whether to filter events by occurrence date.

- Enable this option and specify a date so that events that occurred on or after this date and that match the selected filters are displayed.
- Disable this option so that events are filtered by type and severity only.

Calendar

Opens a calendar from which you can select a date for the occurrence date.

SEE ALSO

Filter Event Log Entries
Display the Event Log
System Event Severities
Ranges for Field Definitions

Event Log Page

DESCRIPTION

Use the Event Log page to review entries in the Pillar Axiom event log. The **Actions** drop-down list provides an option to set a filter to display specific types of events.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. This control appears when the log contains no more than one page.

Choose from:

- 1 min
- 5 min
- 10 min

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes. This control appears only when the log contains a single page.

Refresh Now

Updates the page with current data.

Navigation

This control appears only when the log contains more than on page. These links allow you to page through the event log:

- **First**. Displays the first page.
- **Previous**. Displays the page immediately preceding the current page.
- Numbered pages. Displays the first, second, third, and so on, page.
- **Next**. Displays the page immediately succeeding the current page.
- **Last**. Displays the final page.

Severity

Displays the severity level of entries in the Pillar Axiom event log. The **Severity Level** value is one of:

- **Informational**
- **Warning**
- **Error**
- **Critical**

For descriptions of events, see the table of event severities in System Event Severities.

Time Occurred

Identifies the time at which the event occurred.

Type

Displays the type of entries in the Pillar Axiom event log. The Type value is one of:

- **Event**, which indicates a system-generated event.
- **Process**, which indicates a background process.

Description

Briefly describes the event that occurred.

SEE ALSO

Display the Event Log

Filter Event Log Entries

Create Alerts

Ranges for Field Definitions

About System Events and Performance Statistics

File Server Overview Page

DESCRIPTION

Use the File Server overview page to review the NAS File Servers on a Pillar Axiom storage system. The **Actions** drop-down list provides options to create, duplicate, modify, delete, and view File Servers, as well as to associate a File Server with a CIFS domain.

FIELD DEFINITIONS

Name

Lists the names of configured File Servers. Click a name to review or modify the File Server settings.

Number of Filesystems

Displays the number of filesystems that are associated with each of the File Servers.

Status

Displays the current status of each File Server.

NFS

Identifies whether the Network File System (NFS) protocol is configured for this File Server.

CIFS

Identifies whether the Common Internet File System (CIFS) protocol is configured for this File Server.

Joined Domain

Identifies whether the File Server is joined to a CIFS domain. When CIFS users try to access filesystems that are associated with this File Server, the Pillar Axiom system authenticates the users through their user accounts on the Windows domain controller.

- If **Yes**, the Join Domain request succeeded and the File Server is registered with the Windows domain controller. CIFS users can access filesystems.
- If **No**, the Join Domain request failed. Review the configuration of both the File Server and the Windows domain controller. Retry the Join Domain request and ensure that it succeeds.

Local CIFS Groups Assigned

Indicates whether the CIFS protocol has been assigned to any local groups.

VLAN Tag

Identifies the virtual LAN (VLAN) tag that is assigned to the primary virtual network interface of the File Server.

SEE ALSO

Create File Servers (NAS Storage Systems)

Delete File Servers

Duplicate File Server

Modify File Server Attributes

Ranges for Field Definitions

File Server Page, Account Mapping Tab

DESCRIPTION

Use the **Account Mapping** tab to create and modify the account mappings for a File Server. Both Common Internet File System (CIFS) and Network File System (NFS) protocols must be configured on the File Server.

FIELD DEFINITIONS

Enable CIFS to NFS Domain Mapping support

Identifies whether support for mapping CIFS-to-NFS user accounts is enabled.

- Enable CIFS-to-NFS mapping if you have a single authority to assign users with both an NFS account and CIFS account using the same name. Users can use NFS only, CIFS only, or both NFS and CIFS.
- Disable CIFS-to-NFS mapping if you don't have a single authority to assign NFS and CIFS accounts.

You should disable account mapping when:

- Users do not have both NFS and CIFS accounts.
- NFS and CIFS users account names are different.
- The authority to assign account names can be done by different groups.

Note that the user's NFS account is the owner of files and directories that the user creates.

Attempt mapping for all CIFS security domains

Specifies that, when CIFS-to-NFS account mapping is enabled, the mapping is used for all CIFS domains.

Attempt mapping for a single CIFS security domain

Specifies that, when CIFS-to-NFS account mapping is enabled, the mapping is used for a particular CIFS domain.

CIFS Security Domain *(if you specify mapping for a single CIFS security domain)*

Identifies the NetBIOS name of the CIFS domain that contains definitions for authentication and account mapping. The File Server uses this domain to map CIFS accounts to NFS accounts.

Note: Fully qualified DNS names are not supported.

SEE ALSO

Create File Servers (NAS Storage Systems)

Modify File Server Attributes

Ranges for Field Definitions

File Server Page, CIFS Tab

DESCRIPTION

Use the **CIFS** tab to create and modify the Common Internet File System (CIFS) protocol options for a File Server.

FIELD DEFINITIONS

Enable CIFS Support

Identifies whether the File Server supports the Common Internet File System (CIFS) protocol. Enable CIFS if users will access files from client machines that use the CIFS protocol.

NETWORKING

CIFS File Server Name

Identifies the name by which users refer to the CIFS server. Because the CIFS server advertises its services by this name through the NetBIOS Naming Service, this value is sometimes called the *NetBIOS name*.

Important! Avoid using a CIFS File Server name that is already in use by another File Server. Duplicate server names may result in failures, such as an inability to join a Windows Domain.

Server Comment

Describes the NetBIOS name of the CIFS server.

WINS Server 1-3 IP Address

Identifies a Windows Internet Naming Service (WINS) server in your network. The WINS server is the central repository for Windows server names. Broadcasts of NetBIOS names do not cross subnets, so the WINS server can resolve names across multiple subnets.

Note: If you enable Active Directory for authentication, you do not need to enter a WINS server address.

Character Set

Identifies which character set is valid:

- **Standard-ASCII**, which is a 7-bit character set that includes 128 printable and control character values (0-127).
- **IBM-437**, which is an 8-bit character set that includes the Standard-ASCII values plus values from OEM code page 437. Used on most PCs that are sold in the USA.
- **IBM-850**, which is an 8-bit character set that includes the Standard-ASCII values plus values from OEM code page 850. Used on most PCs that are sold in Western European countries.
- **SHIFT-JIS**: An 8-bit, single- and double-byte character set for Japanese (also referred to as SJIS). It is commonly used on non-UNIX platforms.

Opportunistic Locking

Identifies whether opportunistic locks (oplocks) are enabled to cache client data.

- Read caching permits data that is stored on the server to be copied to the client. When a client application requests the data, the read request is satisfied by reading the client's copy of the data. Read caching can last forever. If the data that is stored on the server changes, the client is informed and the read cache is invalidated.
- Write caching permits data changes to be buffered in the client for up to two minutes. When a client application changes the data, the write request is satisfied by changing the client's copy of the data. Client-server connection problems that occur during the two-minute window can result in lost changes. If the client is still running, a popup message alerts the client to the lost changes.

Disable oplocks to require that all client requests must access a data file instead of cache.

Allow CIFS Connections Over TCP Only

Specifies that CIFS connections are permitted directly over TCP connections only. This option restricts CIFS connections using NetBIOS over TCP. This option enables or disables communication over TCP port 445. If this feature is not enabled, communication takes place on TCP ports 137 and 139.

Only allow CIFS connections with SMB Signing

Server Message Block (SMB) signing provides digital signatures in CIFS communications to provide data integrity. Some older CIFS clients do not support SMB signing. In addition, SMB signing can result in a significant performance loss.

The Pillar Axiom system supports both CIFS connections that use SMB signing and CIFS connections that do not use SMB signing. This option disables support for CIFS connections that do not use SMB signing requiring that all CIFS connections use SMB signing.

This option allows you to specify that all CIFS connections use SMB signing even if this means that older CIFS clients will not work and that newer CIFS clients may get poorer performance.

AUTHENTICATION

Require authentication (prevents anonymous access)

Specifies that clients cannot use anonymous access while communicating with the CIFS server. Also, the CIFS server will not use anonymous connection while communicating with the domain controller.

NTLM Domain

When this option is specified, the CIFS server joins the Windows domain in NTLM mode. In this mode, the security protocols used while communicating with the domain controller are weaker than when the **Active Directories Domain** and **Authenticate Using Kerberos** options are specified. The

domain name specifies the Windows domain of which the CIFS server is a member. This option is required when the domain controller is a pre-Windows 2000 server.

Tip: For Windows 2000 and later versions, we recommend that **Active Directory Domain** be selected instead because that option provides stronger security than does the **NTLM Domain** option.

Active Directories Domain

Specifies the Active Directory domain in which to authenticate access to the CIFS server.

When this option is specified, the CIFS server joins the Active Directory domain as a native member server. This enforces stricter security protocols like Kerberos and LDAP while communicating with the domain controller and for authenticating users. This is the preferred method of joining the domain as it provides stronger security.

Authenticate using Kerberos

Specifies that only Kerberos can be used to authenticate an Active Directory domain. Kerberos lets a user request an encrypted ticket from the Key Distribution Center (KDC), which can then be used to request a particular service from a server. The user's password does not have to pass through the network.

When the Pillar Axiom system is a CIFS server, the CIFS client gets the ticket from Kerberos and presents it to the Pillar Axiom system. When the Pillar Axiom system is a CIFS client, the Pillar Axiom system gets the ticket from Kerberos and presents it to a customer-supplied domain controller.

In the past, CIFS authentication was performed using the NTLM protocol. NTLM is considered less secure than Kerberos.

The Pillar Axiom system allows authentication to use either Kerberos or NTLM. In addition, the Pillar Axiom system uses either Kerberos or NTLM when authenticating itself with the CIFS domain controller. This option disables support for NTLM requiring that all authentication be performed using Kerberos.

This option allows you to specify that all authentications use Kerberos even if this means that older CIFS clients will not work or that the Pillar Axiom system may be unable to perform certain operations.

PREFERRED SERVERS

Enable Preferred Servers

In a Windows domain that has multiple Domain Controllers, it may be desirable (for security and latency reasons) to shorten the list of controllers. This option enables a system administrator to identify up to three Domain Controllers that the CIFS File Server will contact when it needs to perform a Windows operation such as joining a domain or authenticating a user.

If you have selected Active Directories Domain for authentication, the File Server uses the Preferred Servers list to communicate with the domain controller to, for example:

- Join a domain.
- Update the Kerberos credentials.
- Update a Kerberos ticket through the Kerberos Key Distribution Center.

If the Enable Preferred Servers option is not selected, the File Server uses the list of Domain Controllers returned by the default settings for DNS (or WINS) in the Pillar Axiom system configuration.

Note: If you have selected NTLM Domain to manage user authentication and service location, the Preferred Servers list will not control the File Serverauthentication process. If you have selected Active Directories Domain, the Preferred Servers list will control the File Serverauthentication process regardless whether the client chooses NTLM or Kerberos authentication.

Use DNS with Preferred Servers

When a Pillar Axiom system joins a Windows domain in Active Directory mode, this option directs the File Server to use the preferred servers first. If the File Server cannot contact any of those Domain Controllers, it will attempt to contact the remaining Domain Controllers from the list that the DNS server returns.

Important! When the Use DNS with Preferred Servers option is not selected, the File Server uses *only* the list of preferred Domain Controllers. If none of the preferred servers are available, the File Server will not attempt to obtain any further Domain Controllers from external sources. Pillar highly recommends that a preferred server list contains addresses of those Domain Controllers of which at least one will be available on the network at any given time.

FILE SERVER ACCESS TO DOMAIN CONTROLLERS

Access Anonymously

Identifies whether the Pillar Axiom File Server communicates with the CIFS domain controller as an anonymous user or as a specific user. The File Server uses the specified account to authenticate users and to perform user-name-to-SID mappings (SID means security identifier). By default:

- Domain Controllers that are installed on pre-Windows 2000 servers or Windows 2000 servers that are members of pre-Windows 2000 domains permit domain controller access by anonymous users. If your network domain controller is configured this way, choose Anonymous.
- Domain Controllers that are installed on Windows Server 2003 machines deny domain controller access by anonymous users. If your network domain controller is configured with the Windows

Server 2003 default or the non-default option for pre-Windows 2000 servers to prohibit anonymous access, choose Specific User Account.

Note: If network traffic is heavy or if the ACL is very large, the system may display the SID number instead of the username. This does not effect the ACL permissions.

Access Using this Account

Enables a Pillar Axiom server to join a Windows domain. This account must have administrative credentials and be a Domain Admin group member.

For complete information on enabling a Pillar Axiom server to join a Windows domain, refer to the *Windows Integration Guide for NAS Systems*.

Username

Identifies a user name for the File Server to access the domain controller when anonymous access is not permitted. This account should be:

- Similar to the Guest account, with the same type of low-security privileges.
- Different from the Administrator account for the domain controller.

Password

Identifies the user password for the Pillar Axiom File Server to access the CIFS domain controller if anonymous access is not permitted.

Log onto (Domain)

This specifies the domain to which the alternate user belongs. The alternate user, when specified, is used to authenticate the CIFS server, while communicating to the domain controller.

SEE ALSO

Create File Servers (NAS Storage Systems)

Modify File Server Attributes

Ranges for Field Definitions

File Server Page, Filesystems Tab

DESCRIPTION

Use the **Filesystems** tab to review thefilesystems that are associated with a File Server.

FIELD DEFINITIONS

Filesystems using this File Server

Displays the names of filesystems that are associated with the File Server as well as the physical location of the filesystems on the Slammer, including the control unit.

SEE ALSO

Create File Servers (NAS Storage Systems)

Modify File Server Attributes

Ranges for Field Definitions

Locate a Filesystem on a Slammer

File Server Page, Network Tab

DESCRIPTION

Use the **Network** tab to create and modify the network definitions for a File Server.

Tip: For increased flexibility, define several File Servers on the same Slammer control unit (CU). By locating the virtual interfaces (VIFs) from different File Servers on the same Slammer CU, you gain the flexibility to rebalance your Pillar Axiom storage system by moving File Servers from one Slammer CU to another CU.

FIELD DEFINITIONS

File Server Name

Identifies the name that is assigned to a File Server. File Server names must be unique on the Pillar Axiom storage system.

VIRTUAL INTERFACES

IP Address

Identifies the IP address associated with the primary virtual interface (VIF) of the File Server.

Netmask

Identifies the subnet mask that is used by the primary VIF.

Slammer

Identifies the NAS Slammer to which the primary VIF is assigned.

Control Unit

Identifies one of the two control units in a NAS Slammer to which the primary VIF is assigned.

Note: A NAS Slammer control unit may have VIFs from different File Servers associated with the CU.

PORT Number

Identifies the Slammer port that is associated with the primary VIF of the File Server. Select a gigabit Ethernet network port on one of the Slammer CUs:

- If Link Aggregation is enabled, select the port having the lowest number in the aggregate.

  Note: You cannot create virtual interfaces on any of the other ports in the aggregate when Link Aggregation is enabled.
- If Link Aggregation is disabled, select any port.

Identify Port

Opens the Identify Hardware Component pages so that you can blink LEDs on the specified hardware component (**Identify**) or on all other components (**Reverse Identify**).

VLAN ID

Identifies the virtual LAN (VLAN) ID that is associated with the primary VIF of the File Server.

Tip: We recommend that you avoid using VLAN tags. Because you can configure multiple File Servers on the same VLAN that does not use tags, we recommend that you configure your system so that all VIFs are not tagged.

- To implement VLAN tagging, assign a value of 1 through 4094, inclusive, if you have connected a VLAN-capable switch to the Pillar Axiom system.
- Leave the field blank to disable VLAN tagging.

Frame Size (MTU)

Identifies the packet frame size. This value defines the maximum transmission unit (MTU).

The frame size (MTU) does not include the Ethernet header portion of the packet. If your network switch has trouble with this, you can set the switch to a larger value or lower the MTU size to correct the problem.

- If your network supports extended Ethernet (jumbo) frames, enter an integer greater than 1500 and less than 16362. Make sure that this Pillar Axiom MTU size matches the network MTU size. If the MTU sizes are mismatched, users may experience I/O “hangs” when the client machines try to process packets that are too large.
- If your network does not support jumbo frames, enter the default frame size of 1500.

Additional VIFS

Depending on mode, the GUI displays **Create**, **Modify**, or **View**, which opens the Virtual Interfaces Page so that you can define, modify, or view secondary virtual network interfaces for the File Server. Secondary virtual network interfaces are optional.

ROUTE

Default Gateway

Identifies the IP address that is assigned to the gateway host. The gateway IP address is used to route messages from this network to other networks.

Additional Routes

Depending on mode, the GUI displays **Create**, **Modify**, or **View**, which opens the Routes Page so that you can access (define, modify, or view) additional routes for the File Server. Additional routes are optional; you do not have to define them.

DNS SETTINGS

DNS Domain

Identifies a Domain Name Service (DNS) domain, which specifies the domain to be searched.

DNS Server

Each field identifies the IP address of one of three Domain Name Service (DNS) servers in your network that translates host names to IP addresses.

If you define more than one DNS server, the Pillar Axiom system accesses them in the order that they are listed. When a match is found, that translation is used.

Note: When a domain controller is unresponsive or its responses are intermittent, the File Server follows the DNS order to contact other domain controllers for join domain operations. This strategy has the appearance of the File Server not following the list order.

SEE ALSO

Create File Servers (NAS Storage Systems)

Modify File Server Attributes

Ranges for Field Definitions

File Server Page, NFS Tab

DESCRIPTION

Use the **NFS** tab to create and modify the Network File System (NFS) protocol options for a File Server.

FIELD DEFINITIONS

Enable NFS Support

Identifies whether the File Server supports the Network File System (NFS) protocol. Enable NFS if users will access files from client machines that use the NFS protocol.

NETWORKING

Port Number

Identifies which port is reserved for NFS requests. If you allow NFS requests only from a secure port, enter a port number that is less than or equal to 1024.

Default: port 2049

Character Set

Identifies which character set is valid:

- Standard-ASCII: A 7-bit character set that includes 128 printable and control character values (0-127).
- UTF8: A Unicode character set that uses an 8-bit transformation format. Defined in RFC 2279.
- ISO8859-1: An 8-bit, single-byte character set that is an extension of Standard-ASCII.
- ISO8859-15: A modification of the commonly used ISO8859-1 character set. Because it includes the Euro symbol, it is commonly used in Europe.
- SHIFT-JIS: An 8-bit, single- and double-byte character set for Japanese (also referred to as SJIS). It is commonly used on non-UNIX platforms in Japan.
- EUC-JP: A multibyte character set for Japanese. It combines elements from the JIS X 0201, JIS X 0208, and JIS X 0212 standards and is commonly used on Unix and Unix-like platforms in Japan.

SECURITY

Non-NFS UID

Identifies the user ID (UID) of the owner of a file that does not have an NFS account, such as a CIFS account.

For example, if you had a special account called `other_protocol` with a UID of `1234`, you would enter **`1234`** in the **Non-NFS UID** field. Doing so allows you to identify a file that is owned by an account that is not an ordinary NFS account.

Non-NFS GID

Identifies the group ID (GID) of a file where the group owner is not an NFS group, such as a Common Internet File System (CIFS) group.

Allow mount from reserved ports only

Identifies whether all NFS mount requests must be received on a secure port that is reserved for NFS requests only.

- Enable the option if NFS mount requests must be received on a secure port that is reserved for NFS requests (including mount requests). The port number must be less than 1024.
- Disable the option if NFS mount requests can be received on any open port.

Allow chown from non-root users

Identifies whether users who do not have root privileges can change the owner of files that are stored in the Pillar Axiom system.

- Enable the option if non-root users can change ownership on only files that they own.
- Disable the option if only root users can change file ownership.

Note: When quotas are enabled, you may set this flag to automatically transfer files to other users if an individual user exceeds a quota limit.

Allow TCP connections

Identifies whether Transmission Control Protocol (TCP) connections are allowed.

- Enable TCP if users can mount filesystems using the Transmission Control Protocol (TCP).
- Disable TCP if users can mount filesystems using only the User Datagram Protocol (UDP).

NIS SERVER

If the NIS Server fields are specified, the status for **NIS entry Service Type** on the **Services** tab shows as Configured. NIS is then one of the services that can be included in the search order for resolving host names and validating user logins.

NIS Domain

Identifies the name of a group of hosts that share a part of their system data (such as user names, passwords, and host names) through NIS for the purpose of maintaining consistent configuration files throughout the network.

NIS Server

Identifies the IP addresses of up to three hosts running an NIS server (from which clients may retrieve host names and login information). These identified hosts service the NIS requests that they get from the Pillar Axiom system.

Note: To minimize the possibility of the client losing access to the Pillar Axiom system, if the system loses communication with the primary NIS Server, the system switches to the secondary NIS Server. If the system loses communication with the secondary NIS Server, the system switches to the tertiary NIS Server.

NETWORK GROUP

/etc/netgroup

Identifies the upload status of the netgroup file:

- **Not Uploaded**: No file has currently been uploaded to the Pilot.
- **Uploaded**: A file has been uploaded to the Pilot.
- **Click OK to Upload**: A filename has been specified on the **Services** tab for the `/etc/netgroup` and is ready to be uploaded.

SEE ALSO

Create File Servers (NAS Storage Systems)

Modify File Server Attributes

Ranges for Field Definitions

File Server Page, Services Tab

DESCRIPTION

Use the **Services** tab to create and modify host name and user login resolution search orders for a File Server.

FIELD DEFINITIONS

Host Name Resolution Search Order

Identifies the search order that the Pillar Axiom system uses to resolve hosts.

User Login Search Order

Identifies the search order that the Pillar Axiom system uses to resolve hosts.

Upload

Appears in Modify mode only. Permits you to:

- Select `/etc/passwd`, `/etc/group`, `/etc/netgroup`, and `/etc/hosts` files that reside on a machine in your network.
- Upload the selected files to the Pillar Axiom storage system.

The system uses the uploaded files if you do not use Network Information Service (NIS).

Note: However, be aware that:

- The ASCII files must use the NFS character set.
- Fields 1 and 3 (user/group name and user/group ID) from `/etc/passwd` and `/etc/group` are used.
- 10 MB is the maximum file size for any of the files.
- The + NIS token is unsupported.
- The Pillar Axiom system does not validate the files.

Download

Appears in Modify mode only. Downloads the previously uploaded `/etc/passwd`, `/etc/group`, `/etc/netgroup`, and `/etc/hosts` files from the Pillar Axiom system to a designated location in your network.

SEE ALSO

Create File Servers (NAS Storage Systems)

Modify File Server Attributes

Ranges for Field Definitions

Filesystem Overview Page

DESCRIPTION

Use the Filesystem overview page to review the filesystems on a Pillar Axiom storage system. The **Actions** drop-down list provides options to create, copy, modify, delete, and view filesystems, as well as to create an immediate Snap FS of a filesystem.

FIELD DEFINITIONS

Name

Lists the names of configured filesystems. Click a name to review the filesystem settings.

SecureWORMfs Retention

Identifies the retention settings defined on the filesystem:

- **None**: There are no retention settings on the filesystem.
- **Standard**: Files on the Pillar Axiom SecureWORMfs filesystem are protected based on the specified retention settings; however, you can delete protected files on the Pillar Axiom SecureWORMfs by deleting the entire filesystem.
- **Compliance** Files on the Pillar Axiom SecureWORMfs filesystem are protected based on the specified retention settings. You cannot delete the Pillar Axiom SecureWORMfs if any files are present on the filesystem.

An existing WORM filesystem can be upgraded from Standard retention to Compliance.

Note: Normally, Pillar Axiom SecureWORMfs instances cannot be changed from Compliance to Standard. However, if you need to downgrade a Compliance instance to Standard, contact the Pillar World Wide Customer Support Center for assistance.

Total Capacity

Identifies the current capacity that is assigned to the specified filesystem.

Free Capacity

Identifies the current unused capacity that is available. Free capacity is total capacity minus allocated capacity.

Used Capacity

Identifies the current capacity usage of the object.

Growth Max

Identifies the maximum capacity limit (in GB) in which the object can expand.

Status

Displays the current status of a filesystem. The possible states are:

- **Online**: The filesystem is online and normal.
- **Online: Replication Target**: The filesystem is online, normal, and in replication target mode. This status applies only to filesystems that are the target in a replication pair that has been defined through use of the Pillar Axiom MaxRep Replication for NAS utility.
- **Offline**: The filesystem is offline.
- **Offline: Replication Target**: The filesystem is offline and in replication target mode. This status applies only to filesystems that are the target in a replication pair that has been defined through use of the Pillar Axiom MaxRep Replication for NAS utility.
- **Partial Offline**: The actual redundancy level of the filesystem may be different from the redundancy level with which the filesystem was configured.
- **Conservative**: Write-back cache on the filesystem has been disabled so journaling has slowed.
- **Degraded**: All of the copies of a redundant filesystem are not available. If one copy is missing, it is not fully redundant. This can happen when a write fails (which may be a 30 second time-out) to one copy of the array.

Storage Class

Identifies the category of physical storage on which the logical volume resides:

- **FC** (Fibre Channel drives)
- **SATA** (Serial ATA drives)
- **SLC SSD** (single level cell, solid state drives)

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

Priority

Identifies the layout of the data for the logical volume:

- **Premium** priority is the highest possible performance. This priority uses the outer 20% of the drive platters.

  Note: I/O is typically faster when data resides on the outer edge of the drive platters.
- **High** priority logical volumes are allocated space in the outer 20%-40% of the drive platters.

  Tip: If sufficient Bricks are present, performance of high priority volumes is enhanced by placing those volumes on a greater number of Bricks.
- **Medium** priority logical volumes are allocated space in the outer 40% to 60% of the drive platters.
- **Low** priority logical volumes are allocated space in the band that is 60% to 80% from the outer diameter of drive platters.
- **Archive** (lowest-priority) logical volumes are allocated space in the inner areas of the drive platters.

Redundancy

Identifies the redundancy level of the logical volume (filesystem or LUN) as standard or double.

File Server

Identifies the name of a File Server with which the filesystem is associated.

File Deletion Enabled

Identifies whether files can be deleted on the Pillar Axiom SecureWORMfs. Protected files cannot be deleted regardless of this setting. When this option is enabled, the files that can be deleted are expired files.

Space for Clone FSs

- **Requested**: The amount of system capacity requested for storing Clone FSs for this filesystem.
- **Allocated**: The amount of system capacity reserved for storing Clone FSs for this filesystem.
- **Used**: The amount of system capacity consumed by Clone FSs for this filesystem.

Number of Snapshots

Identifies the number of data replicas that have been created for a logical volume (filesystem or LUN).

Number of CIFS Shares

Identifies the count of Common Internet File System (CIFS) shares that are defined for the specified filesystem.

Number of NFS Exports

Identifies the count of Network File System (NFS) exports that are defined for the specified filesystem.

IOps

Identifies the current load of the filesystem on performance for input (read) and output (write) operations.

Bandwidth

Identifies the current load of the filesystem on bandwidth in MB/s for read and write operations.

Last Protected File Integrity Check

Identifies the date and time at which the system performed the most recent **Protected File Validation** check on the filesystem.

Retention (d/m/y)

Identifies the retention period for all files on a Pillar Axiom SecureWORMfs filesystem. Valid values are:

- **Default**: 0 days-1000 years and must be greater than or equal to the minimum and less than or equal to the maximum retention value.
- **Maximum**: 0 days-1000 years and must be greater than or equal to the minimum retention value.
- **Minimum**: 0 days-1000 years.

Creation Time

Identifies the date and time in which the filesystem was created.

Creation Method

Identifies the method that was used to create the Pillar Axiom SecureWORMfs filesystem (generally used for audit purposes):

- **New**: Newly created.
- **Restore**: Restored from a backup image.
- **FSCK**: Recovered using FSCK.
- **Copy**: Created from a copy of the filesystem.
- **Unknown**: The creation method is not known, however, this usually applies to a filesystem that was created with an earlier version of the software.

SEE ALSO

Create Filesystems (NAS Storage Systems)

Define Retention Policy

Toggle the Pillar Axiom SecureWORMfs File Deletion Setting

Validate File Integrity on a Pillar Axiom SecureWORMfs

Modify Filesystem Attributes

Copy Filesystem Page

Delete Filesystems

Create an Immediate Snap FS

Take Filesystems Offline

Put Filesystems Online

Performance Profile Page

Ranges for Field Definitions

Filesystem Page, Exports Tab

DESCRIPTION

Use the **Exports** tab to create, modify, and delete the Network File System (NFS) exports for a filesystem.

FIELD DEFINITIONS

Unix Exports (NFS) List

Displays a list of currently configured NFS exports. Check the box to the left of an export point to select it for modification or removal.

Export Path

Identifies the full path to an NFS export point, without the filesystem-name prefix that NFS clients see when they mount this export point. For example, if you export the /dir directory on the fs1 filesystem, the export point appears in the list as `/dir`. NFS clients must mount the `/fs1/dir` mount point.

UID

Identifies the user ID (UID) for anonymous users.

- Set the UID to zero (0) if you want anonymous-user access to be identified as the root user. This is a very insecure setting. Setting anonymous to root allows all users full control.
- Set the UID to the “nobody” value (often -2, but not always) if you want anonymous-user access to be identified as the user “nobody.” In most instances, this is the recommended setting.
- Set the UID to any other user ID that you want to be identified as anonymous-user access.

Create

Creates a new export point based on the export parameters that you entered.

Modify

Modifies the export point that you selected.

Remove

Deletes the selected objects.

Host Access

Identifies the NFS clients, or hosts, that can mount the NFS export.

- **All Hosts**: Everyone can access the export point.
- **Single Host**: Only the specified host can access the export point.
- **NIS Netgroup**: Everyone within the Network Information Service (NIS) netgroup can access the export point.
- **Network**: Everyone on the specified subnet can access the export point. **Netmask** specifies the screen to use to determine which host computers can access this export.
- **Read Only**: The specified host has read-only privileges.

- **Root Access**: The specified host has root permissions.

Additional Hosts

Allows you to create additional hosts.

Import from External File

Allows you to browse to and select a text file that contains export and host definitions. The file must contain entries similar to this example:

```
# This is an example exports file
#
/home/ann pxN(ro,no_root_squash,anonuid=5)
/home/joe @groupN(ro,root_squash,anonuid=150)
/home/sue 192.168.1.7(rw,root_squash,anonuid=50)
/home/ted 192.168.1.9/255.255.255.0
  (rw,root_squash,anonuid=3)
/home/bob pxZ(rw,root_squash,anonuid=22)
  @groupZ(ro,root_squash,anonuid=33)
```

where:

- `/dir/subdir` is a mount point (the `/home/`*name* directories).
- `pxN` (alphanumeric host name), `@groupN` (netgroup), single IP address, or IP address with a subnet mask are valid formats for a machine/host specification.
- `ro` or `rw` is a read-only or read-write specification. If you do not include this specification in the file, `rw` is assumed.
- `no_root_squash` or `root_squash` is a Linux exportfs root-squashing specification. If you do not include this specification, `root_squash` is assumed.
- `anonuid` is the user ID to assign when users access the mount point as the anonymous user (`nobody`). This value applies for the entire export, even though it is defined as a host-entry attribute. If you define multiple host entries for an export, the `anonuid` from the first host entry is used for the anonymous user ID value for the export.

SEE ALSO

Modify NFS Exports

Add NFS Exports to a Filesystem

Ranges for Field Definitions

Filesystem Page, Identity Tab

DESCRIPTION

Use the **Identity** tab to create, modify, and review a filesystem on a Pillar Axiom storage system.

Tip: To improve performance, Pillar recommends that you place a filesystem on the same Slammer control unit as the virtual interfaces (VIFs) that will be used to access the filesystem.

FIELD DEFINITIONS

Filesystem Name *(Create and Modify only)*

Identifies the name that is assigned to a filesystem. Each filesystem name must be unique within its associated File Server and volume group.

You must specify the File Server and the filesystem name, which results in a fully qualified filesystem name of */ fileserver/ filesystem*.

When you modify the name of a filesystem, Clone FS, or Snap FS, the process of changing the name must adhere to the following rules:

- The new name must be unique within the File Server that contains the filesystem.
- Be sure the system completes a rename operation before you start a new rename request.
- Do not start a system software update operation until the rename request completes.
- For Pillar Axiom SecureWORMfs instances:
  - Compliant Pillar Axiom SecureWORMfs instances cannot be renamed.
  - Standard Pillar Axiom SecureWORMfs instances may be renamed.
  - Clones and snapshots of a Compliant Pillar Axiom SecureWORMfs instance may be renamed.

After a filesystem, Clone FS, or Snap FS is created, a name change has the following effects:

- The name change does not affect any existing Common Internet File System (CIFS) shares for the associated filesystem.
- The name change modifies existing Network File System (NFS) exports for the associated filesystem by changing the exported name from the original filesystem to the new name.
- NFS clients that have a filesystem mounted continue to have the filesystem mounted.
- File handles on Snap FS objects continue to be valid, whether the objects are accessed through NFS or CIFS.
- If a filesystem is renamed after a backup has started, the data management application (DMA) may refer to the previous name and cause the backup to fail.

- A DMA may perform a full (rather than an incremental) backup during the next backup process following the name change.

Note: Renaming filesystems does not apply to scheduled snapshots.

Inactive *(Create and Modify clone only)*

Specifies that the cloned filesystem cannot be accessed from the data path.

Created *(View only)*

Identifies the date and time the filesystem was created.

System free space

Identifies the current unused capacity that is available. Free capacity is total capacity minus allocated capacity.

Create Filesystem in *(Create and View only)* or **Volume Group** *(Modify only)*

Identifies the name of a volume group with which the filesystem is associated. Volumes is the label for the system root-volume group.

Retention Policy

Creates a Pillar Axiom SecureWORMfs object. Data is stored on the filesystem in a non-erasable, non-rewritable (protected) manner for a fixed period of time. You can choose one of:

- **Standard**: Allows you to delete the filesystem even if there are protected files. The Protected File Integrity scan occurs during the specified start time.

  Tip: You can convert a standard Pillar Axiom SecureWORMfs instance later to a compliance Pillar Axiom SecureWORMfs instance.
- **Compliance**: You cannot delete the filesystem if there are any files. The Protected File Integrity scan occurs during the specified start time. Also, the Network Time Protocol (NTP) server must be enabled to use the **Compliance** option.

  An existing WORM filesystem can be upgraded from Standard retention to Compliance.

  Note: Normally, Pillar Axiom SecureWORMfs instances cannot be changed from Compliance to Standard. However, if you need to downgrade a Compliance instance to Standard, contact the Pillar World Wide Customer Support Center for assistance.

In **Modify** and **View** modes, **Retention Policy** displays **Disabled** for regular filesystems.

Note: You cannot migrate a Pillar Axiom SecureWORMfs instance from or to a regular filesystem.

When you enable this option, you must define the retention period on the Retention Policy tab.

Filesystem Assignment

This field specifies how to assign filesystems to a Slammer control unit (CU):

- **Auto-assign Filesystem to a Slammer** *(Create only)*

  The system selects the Slammer and CU to which to associate the filesystem.
- **Assign Filesystem to Slammer** *(Create only)*

  Identifies the Slammer to which the filesystem is to be assigned. Use the **Assign Control Unit** drop-down list to associate the filesystem with a particular control unit (CU) of the identified Slammer.
- **Slammer** *(Modify and View only)*

  Identifies the Slammer to which the filesystem is to be assigned.
- **CU** *(Modify and View only)*

  Identifies the particular Slammer CU to which the filesystem is to be assigned. You can optionally allow the system to assign the CU.

File Server *(Create only)* **Associated File Server** *(Modify and View only)*

Identifies the name of a File Server with which the filesystem is associated. You cannot change this association later by performing a Modify action.

View

Opens the File Server Page, Network Tab in read-only mode. Review the File Server definitions to determine which File Server to associate with the filesystem.

Create New File Server *(Create only)*

Opens the File Server Page, Network Tab so that you can create a new File Server.

Background Copy and Data Migration Priority

Identifies the strategy the system should use to control the impact on performance when background tasks need to copy or move data from one location in the storage pool to another:

- **Minimize performance impact**. Restricts the amount of work performed on a loaded system. This option is intended to have a minimal impact on client I/O throughput at the expense of longer copy times.
- **System optimized performance**. Balances the background copy with the incoming client I/O. This option is the default.
- **Maximum performance**. Prioritizes the background copy at the expense of client I/O throughput.

Note: When the system is idle or lightly loaded, the above background task maximizes the amount of work done regardless of the option selected.

The following types of operations are affected by the strategy you select:

- Copy
- Restore
- Quality of Service (QoS) changes in:

- Priority
- Redundancy
- Storage Class

Data transfer operations invoked by the Pillar Axiom MaxRep Replication for NAS utility are not affected by the **Background Copy and Data Migration Priority** setting.

Performance *(Modify and View only)*

Identifies whether the filesystem performance is Normal based on its QoS settings and current usage patterns.

IOPs *(Modify and View only)*

Identifies the current load of the filesystem on performance for input (read) and output (write) operations.

Bandwidth *(Modify and View only)*

Identifies the current load of the filesystem on bandwidth in MB/s for read and write operations.

Pinned Data *(Modify and View only)*

Identifies whether any user data is pinned in cache and cannot be written to permanent storage.

Create Automatic Snap FS Schedule *(Create only)*

Identifies whether to create a schedule so that Snap FSs are created at four-hour intervals.

- Enable this option so that a Snap FS is created every four hours as a scheduled operation.
- Disable this option so that a four-hour schedule is not created. You can create schedules with different recurrence intervals.

SEE ALSO

Create Filesystems (NAS Storage Systems)

Modify Filesystem Attributes

Create Snap FS Schedules

Ranges for Field Definitions

Filesystem Page, Quality of Service Tab

DESCRIPTION

Use the **Quality of Service** tab to create and modify the capacity and performance settings for a filesystem.

When you modify a filesystem, you can increase:

- **Initial Capacity**
- **Maximum Capacity**

FIELD DEFINITIONS

Initial Capacity *(Create only)*

Identifies the initial capacity that is assigned to the object. This value is a soft limit, which means that data can be stored in a logical volume until the maximum capacity is reached.

Current Capacity *(Modify and View only)*

Displays the storage capacity that is currently allocated to the object.

Maximum Capacity

Identifies the maximum capacity to which the logical volume can grow. For a clone, this field identifies how much addressable space will be available.

Redundancy

Identifies how many mirror copies of the original data are stored online.

Important! Pillar highly recommends that you consult with a Pillar Customer Support professional for assistance with sizing your system and creating your logical volumes (filesystems and LUNs).

Redundancy options include:

- **Standard**: stores original data only. Data striping over multiple RAID groups maintains full redundancy, even without mirror copies.
- **Double**: stores original data and one mirror copy, with data striping over multiple RAID groups.

  Note: Double redundancy can only provide true redundancy if your system has enough Bricks to allocate the filesystem or LUN such that no two mirror copies share a RAID group.

A SATA Brick has two RAID groups. A Fibre Channel (FC) Brick has one RAID group.

If the storage pool is becoming depleted, or a large filesystem or LUN is created, it might be necessary to place the filesystem or LUN on more RAID group fragments.

Depending on the ability of the system to allocate sufficient contiguous storage blocks for the size of the logical volume, refer to the following number of RAID groups to configure your volumes for the best performance:

Table 19 Optimum number of RAID groups for best performance

Priority	SATA standard redundancy	SATA double redundancy	FC standard redundancy	FC double redundancy
Archive	4	8	2	4
Low	4	8	2	4
Medium	6	12	3	6
High	8	16	4	8
Premium	8	16	4	8

Note: When the selected Storage Class is **SSD**, all available SSD drives are striped across, regardless of Priority.

For *performance testing purposes only*, create a filesystem or LUN using standard redundancy and the HighThroughput Profile. This is *not* recommended for most applications. Reset your system after you have created a HighThroughput LUN or filesystem before you configure normal filesystems or LUNs for applications.

For examples of how Prioritization and Redundancy are used to create filesystems, see About Filesystem Creation.

Optimizer *(Create and Modify only)*

Displays how much capacity is available for each combination of Storage Class and Redundancy.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

Run Simulation *(Create only)*

Calculates the impact of the QoS settings for the new filesystem on existing logical volumes and displays the results in a table.

Storage Class capacity *(Create and Modify only)*

For the assigned Storage Class, identifies its physical capacity in terms of the following:

- **Free**. Available physical storage that can be allocated.
- **Reconditioning**. Capacity in the process of being released to the free pool.
- **Used**. Physical storage already allocated.
- **Total**. Total physical capacity of the specified Storage Class.

Allocate maximum space for Clone FSs *(Create and Modify only)*

Specifies the maximum amount of space to be reserved for Clone FSs from the total system capacity .

Caution We strongly recommend that you allocate sufficient repository capacity to minimize the chances of running out of this space (which could lead to data inconsistency or loss). To set sufficient capacity, use a value equal to the source volume capacity times the number of replicas times the maximum rate of change. For example, for a 100 GB volume that is projected to have 20 active replicas at a time and no more than a 20% rate of change, use a value of 400 GB for the clone repository.

Storage Class

Identifies the category of physical storage on which the logical volume resides:

- **FC** (Fibre Channel drives)
- **SATA** (Serial ATA drives)
- **SLC SSD** (single level cell, solid state drives)

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

Priority vs. other Volumes

The Priority option determines how much of the system resources are devoted to that volume, including the allocation of Slammer CPU cycles and the allocation of specific portions of the disk platters. The higher the Priority, the greater the allocation of CPU time and the faster the media access time.

The Priority option specifies the layout of data that is stored in the Pillar Axiom system based on one of the following settings (see also the figure below):

- **Premium** priority is the highest possible performance. This priority uses the outer 20% of the drive platters.

  Note: I/O is typically faster when data resides on the outer edge of the drive platters.
- **High** priority logical volumes are allocated space in the outer 20%-40% of the drive platters.

  Tip: If sufficient Bricks are present, performance of high priority volumes is enhanced by placing those volumes on a greater number of Bricks.
- **Medium** priority logical volumes are allocated space in the outer 40% to 60% of the drive platters.
- **Low** priority logical volumes are allocated space in the band that is 60% to 80% from the outer diameter of drive platters.
- **Archive** (lowest-priority) logical volumes are allocated space in the inner areas of the drive platters.

Figure 18 Example of HDD priority bands

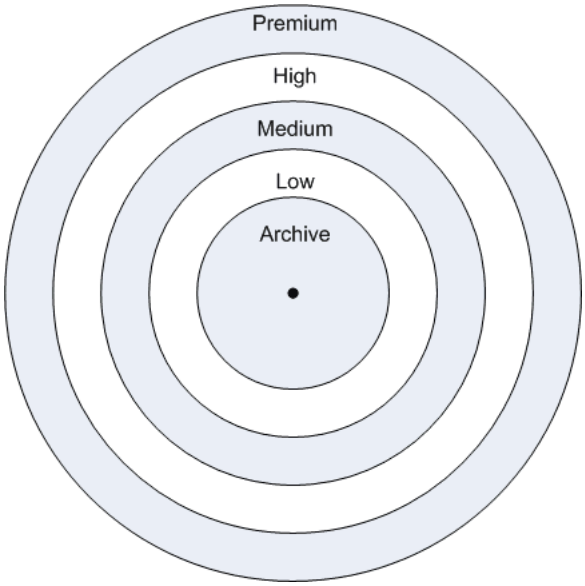

If needed, the Pillar Axiom system allocates space in a higher or lower priority band (within the same Storage Class) than the one that you choose.

The Pillar Axiom system uses the **Prioritization** and **Redundancy** values to calculate whether enough total capacity is available to create the new volume in the specified Storage Class. The calculated results are available in the table that is displayed when you click **Optimizer**.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

Temporary *(Modify only)*

Temporarily changes the priority level of the filesystem or LUN without changing the physical location of the volume. When you want to change the priority back to the original level, modify the QoS of the volume.

One use for temporarily changing the priority level is to allow for a restore operation or data load operation where the original QoS would have been set to optimize for Write Archive or Random Read or Write, respectively.

Important! If the temporary QoS change is not changed back to the original QoS, it can have adverse effects on other system resources. If the change is to be made permanent, the QoS should be changed and the data migrated.

Typical File Size

Identifies the typical size of files that are stored in the specified logical volume. Choose from:

- **Small**, if files are smaller than 20 KB.
- **Medium**, if files are larger than 20 KB and smaller than 4 MB.
- **Large**, if files are larger than 4 MB.

Files are Typically Accessed

Identifies the typical data access method. Choose from the values in the table below.

Table 20 Effects of access and I/O bias

Access bias	I/O bias	Read-ahead in the Slammer	RAID configuration in the Brick
Sequential	Read	Aggressive	RAID 5 Writes data in write-back mode to physical storage in full-stripe extents.
	Write	Conservative	RAID 5 Writes data in write-back mode to physical storage in full-stripe extents.
Mixed and random	Read	None	RAID 5 Writes data in write-back mode to physical storage in full-stripe extents.
Random	Write	None	Distributed RAID

Note: This parameter specifies an optimization bias; it is not a requirement that all data or data operations conform to the specified access method.

I/O Bias

Identifies the typical read-write ratio. Choose one of the following options:

- **Read**, if users or applications read data more often than they write to the data source.
- **Write**, if users or applications write data more often than they read it.

  Important! If you choose **Random** as the access method and **Write** as the I/O Bias, the system creates the filesystem with a Distributed RAID geometry. This geometry provides enhanced write performance but uses twice the capacity. For more information, refer to About Enhanced Performance for Random Write Operations.
- **Mixed**, if the read:write ratio varies.

Note: This parameter specifies an optimization bias; it is not a requirement that all data or data operations conform to the specified I/O bias.

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

One copy is maintained in non-volatile memory on one control unit (CU) of a Slammer. The mirror copy is maintained in one of:

- Battery-backed memory of the partner CU on the Slammer (preferred location). Writes to this copy are equivalent to write-back cache.
- Virtual LUN (VLUN) that is reserved on physical storage for the logical volume if the partner CU is unavailable for the write. Writes to this copy are equivalent to write-through cache.

Writes from the journal to permanent physical storage are equivalent to write-through cache. The system flushes user data and the corresponding metadata as a unit to physical storage.

Quotas will be used *(Create only)*

Identifies whether capacity limits, or quotas, are used in this filesystem. This option enables usage tracking only. Optionally enable quota enforcement explicitly and separately for directories, users, or groups after the filesystem is initialized.

- Enable this option so that the Pillar Axiom system creates a directory quota at the filesystem root (/) and tracks capacity usage in this filesystem.
- Disable this option so that capacity usage is not tracked in this filesystem.

If you plan to use quotas in this filesystem immediately or in the future, enable this option now, before the filesystem contains files and directories. The Pillar Axiom system has to take the filesystem temporarily offline if it contains data when you create quotas and enable usage tracking.

SEE ALSO

Create Filesystems (NAS Storage Systems)

Modify Filesystem Attributes

Ranges for Field Definitions

Filesystem Page, Quotas Tab

DESCRIPTION

Use the **Quotas** tab to create, modify, and delete capacity limits and track usage for individuals or groups who store data in the filesystem.

FIELD DEFINITIONS

Quota List

Displays a list of currently configured quotas. Check the box to the left of quota definition to select it for modification or deletion.

Be aware that if you create a Directory quota and set both the Soft Limit and the Hard Limit to 0 (zero), which means unlimited capacity, the quota does not appear in the list.

Note: You must click **Collect Quotas** and then **Display Quotas** to display the list of quotas configured on the Pillar Axiom system. If there are more than 500 quotas, click **Download Report** to view the list of quotas. See definitions below.

Last Collected

Displays the date and time the quota information was last collected on the system. If the quota list has never been collected on the system, then this field displays N/A.

Collect Quotas

Use **Collect Quotas** to retrieve the list of quotas that are configured on the Pillar Axiom system. Wait for this task to complete before you click **Display Quotas**. Use this option if there are less than 500 quotas configured on the system.

Display Quotas

Use **Display Quotas** to view the list of quotas that are configured on the Pillar Axiom system. You need to use this button each time you view the **Quotas** tab to display the quota list.

Download Report

If there are more than 500 quotas configured on the Pillar Axiom system, you need to download the list of quotas to the client system to view. Once you download the report, unzip (using the standard `bunzip2` command) the file and view the list on the client. If you want to modify a quota, you need to know the fully qualified name (FQN) of the quota and use the `pdscli` to modify it.

Modify Quota

Changes the selected object's configuration settings.

Remove Quota

Deletes the selected objects.

Type

Identifies which directories, users, groups, or combination are covered by the quota. Choose from:

- **Directory**, if the Pillar Axiom system tracks usage for all users and groups and enforces the limits that are defined for the specified directory. Make sure that the directory exists when you create the Directory quota.
- **User**, if the system tracks usage for the specified user in the entire filesystem and enforces the limits that are defined for that user in the filesystem.
- **Group**, if the system tracks usage for the specified group of users in the entire filesystem and enforces the limits that are defined for that group of users in the filesystem.
- **Directory & User**, if the system tracks usage for the specified user and enforces the limits that are defined for that user in the specified directory.
- **Directory & Group**, if the system tracks usage for the specified group of users and enforces the limits that are defined for that group of users in the specified directory.

You can create multiple quotas to track usage and optionally enforce limits at different levels within the directory structure of the filesystem. To create a hierarchy of quotas on a directory tree, start at the lowest level of the hierarchy and work up to the top of the hierarchy. For example, a directory tree resembles the following structure:

Figure 19 Example directory tree structure

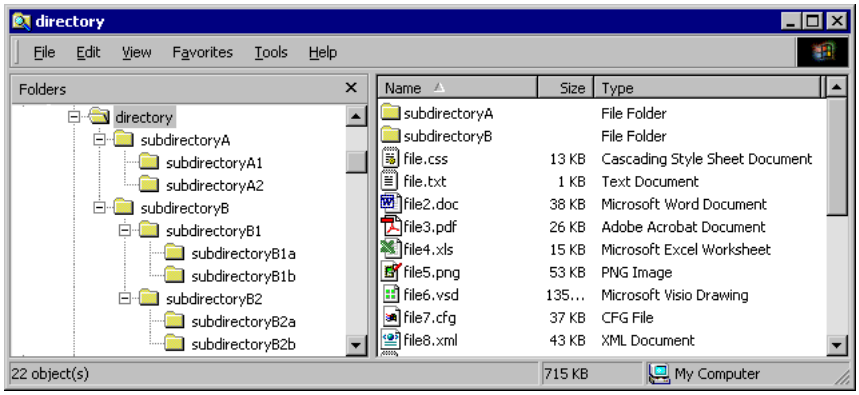

- First, create **Directory**, **Directory & User**, or **Directory & Group** quotas at the `subdirectoryB1*` and `subdirectoryB2*` levels of the hierarchy.
- Second, create those quotas at the `subdirectoryA*` and `subdirectoryB*` levels of the hierarchy.
- Third, create those quotas at the `subdirectoryA` and `subdirectoryB` levels of the hierarchy.
- Finally, create those quotas at the `directory` level of the hierarchy.

By setting quotas at the lower levels first, the system does not have to re-scan lower levels of the hierarchy to create quotas at the higher levels. If you reverse the sequence, however, the system takes longer to create the hierarchy of quotas because each level must be scanned to look for quotas.

Directory

Enter a full path, which starts with a forward slash (/), to the directory that is covered by this quota.

User or Group

Identifies who is covered by a quota.

- If this is a user-level quota, enter a user name (the correct syntax is `username@domain`), user ID (UID), CIFS account name, or NFS account name. The system tracks capacity usage by this user in the specified directory.
- If this is a group-level quota, enter a group name or group ID (GID). The system tracks capacity usage by all group members in the specified directory.

Note: You create default limits for groups and users by entering an asterisk (*) in the **User or Group** field and setting the desired soft and hard limit values.

See Define NIS and NSS Options and File Server Page, Services Tab.

Enforce All Limits For Directory

Identifies whether the quota limits are enforced as well as tracked. All quotas have automatic usage tracking. Enforcement is optional and must be enabled explicitly and separately for each quota.

- Enable this option so that the Pillar Axiom system enforces the capacity limits.
- Disable this option so that the system does not enforce the capacity limits.

Soft Limit

Identifies a soft capacity limit, which means that once the system exceeds this value, the system sends an administrator alert.

Figure 20 Soft limit exceeded alert

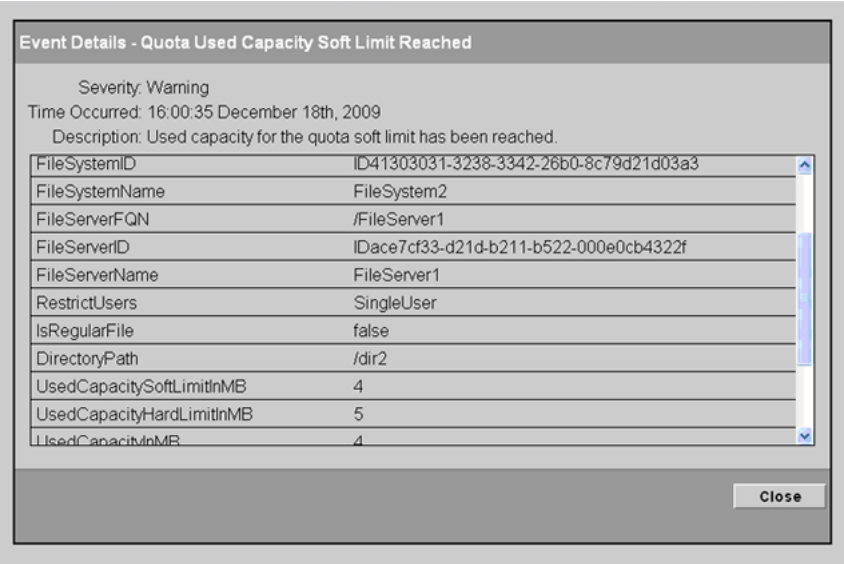

Note: To resolve the situation, either increase the soft limit or delete filesystem objects as needed to drop the used capacity below the limit.

If space is available, the system stores the data until the grace period or the hard limit is reached.

Note: Be aware that:

- If the quota limit of a parent directory is less than that of a child directory, the system enforces the limit of the parent.
- Users sometimes delete directories from the directory structure of a filesystem. If Directory, Directory & User, or Directory & Group quotas are assigned to the deleted directory, the quotas are also deleted.

To define an unlimited value, enter 0 (zero).

Hard Limit

Identifies a hard capacity limit, which means that once the system reaches this value, it rejects write requests and does not store data. Also, a “Disk quota exceeded” event is generated.

To define an unlimited value, enter 0 (zero).

Grace Period

Identifies the number of days that the system can exceed a soft limit. When the time expires, the soft limit becomes a hard limit.

To define an unlimited value, enter 0 (zero).

Create Quota

Adds the new quota definition to the logical volume.

SEE ALSO

About Filesystem Quotas

Modify a Filesystem Quota

Ranges for Field Definitions

Filesystem Page, Retention Policy Tab

DESCRIPTION

Use the **Retention Policy** tab to set the retention periods for a Pillar Axiom SecureWORMfs filesystem. A Pillar Axiom SecureWORMfs filesystem is defined as:

> A type of filesystem used to enforce data retention. Data is stored on a Pillar Axiom SecureWORMfs in a protected (non-rewritable) manner.

Pillar Axiom SecureWORMfs filesystems are intended to be used for archival purposes only. They are not intended to be dynamic production filesystems.

Files written over CIFS connections are protected immediately, not allowing any further modifications.

Files written over NFS connections must be protected explicitly by setting the **Immutable** field for the file in its corresponding attribute file (see Modification of Extended Attributes).

Note: Files under 10 MB are protected in-line. Files 10 MB in size or larger are protected near in-line but without holding up the client.

FIELD DEFINITIONS

Master Retention Period

Identifies the retention period for all files on a Pillar Axiom SecureWORMfs filesystem. Valid values are:

- **Default**: 0 days - 1000 years and must be greater than or equal to the minimum and less than or equal to the maximum retention value.
- **Minimum**: 0 days -1000 years.
- **Maximum**: 0 days - 1000 years and must be greater than or equal to the minimum retention value.

These retention periods can be modified at a later time but must conform to the following restrictions:

- The **Minimum** period can only be made shorter.
- The **Maximum** period can only be made longer.
- The **Default** period must remain within the minimum and maximum values. A change in the default value does not impact files that are already protected.

Protected File Integrity Scan

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

You can manually perform a protected file integrity scan at any time. When multiple scans are scheduled, the system may queue some of those scans depending on the workload. The system shows these queued scans as 0% complete.

See the event log for scan results.

Auditing Mode

This option enables the logging of compliance activities, including:

- The full path to a protected file and the mechanism used to protect the file (manual or automatic).
- Changes to the expiration time of a file. If expiration time is derived from the default retention period, the value of default retention period is also logged in addition to the new and old expiration times and the full path to the file.
- Deletion of an expired file. The log for this event contains the current system time, expiration time of the file, and the full path to the file.
- Detection of data inconsistency in a file. The actual and expected data hash values are logged.

All records in the audit log are in UTF-8 format and have a timestamp and date.

SEE ALSO

Create Filesystems (NAS Storage Systems)

Modify Filesystem Retention Policy

Define Retention Policy

Modification of Extended Attributes

Ranges for Field Definitions

Filesystem Page, Shares Tab

DESCRIPTION

Use the **Shares** tab to create, modify, and delete the Common Internet File System (CIFS) shares for a filesystem.

FIELD DEFINITIONS

Windows Shares (CIFS)

Displays a list of currently configured CIFS shares. Check the box to the left of a share to select it for modification or deletion.

Name

Identifies the name of a CIFS share. A filesystem must be shared before users can create or access data files. Share names must be unique within a File Server.

Enabled

Identifies whether the CIFS share is enabled.

- Enabled shares are active. Users can access an enabled share point.
- Disabled shares are inactive. Users cannot access a disabled share point.

Path to Folder

Identifies the full path to the CIFS share.

Comment

Describes the CIFS share.

Create

Create a new CIFS share based on the share parameters that you entered.

Modify

Allows you to modify the selected share.

Remove

Deletes the selected objects.

SEE ALSO

Modify CIFS Shares

Ranges for Field Definitions

Global Settings Overview Page

DESCRIPTION

Use the Global Settings overview page to select any type of system-wide settings for the Pillar Axiom system. After you select a type, you can review or modify that type of global settings.

FIELD DEFINITIONS

Global Setting Type

Lists types of global settings that you can configure on the Pillar Axiom system. Choose one of:

- **Network**, to review or modify the management and data path interfaces, link aggregation, and notifications on the Interfaces Page.
- **Security**, to review or modify the administrator and browser security settings on the Account Security Settings Page.
- **iSCSI**, to review or modify the iSCSI and iSNS settings on the iSCSI Page.

Global Setting Description

Briefly describes the types of global settings that are configured on Pillar Axiom storage systems.

SEE ALSO

Configure Global Settings

- Configure Management and Data Path Interfaces
- Configure Notification Settings
- Modify Call-Home Settings
- Modify Administrator Account Security Settings
- Modify Email Configuration Settings

Create an Administrator Account

Test Call-Home

Ranges for Field Definitions

Hardware Component, Add Page

DESCRIPTION	Use the Hardware Component, Add page to add Bricks and Slammers to the Pillar Axiom storage system.
FIELD DEFINITIONS	The content pane provides instructional text to add a new hardware component to the Pillar Axiom storage system. Follow the detailed instructions in the *Service Guide* to supplement the instructional text.
SEE ALSO	Replace a Hardware Component *Pillar Axiom Service Guide*

Hardware Component, Identify Page

DESCRIPTION	Use the Hardware Component, Identify page to blink light-emitting diodes (LEDs) on a hardware component that is configured on the Pillar Axiom storage system so that you can locate and service the component.
FIELD DEFINITIONS	The content pane provides instructional text to identify hardware components that are configured on the Pillar Axiom storage system. Follow the detailed instructions in the *Service Guide* to supplement the instructional text.
SEE ALSO	Identify Hardware Components *Pillar Axiom Service Guide*

Hardware Component, Replace Page

DESCRIPTION	Use the Hardware Component, Replace page to remove a Brick or Slammer from the Pillar Axiom storage system and replace the component with a new one.
FIELD DEFINITIONS	The content pane provides instructional text to replace hardware components that are configured on the Pillar Axiom storage system. Follow the detailed instructions in the *Service Guide* to supplement the instructional text.
SEE ALSO	About Hardware Component Replacement *Pillar Axiom Service Guide*

Hardware Component, Summary Page

DESCRIPTION

Use the Component Summary page to review the status of Slammers, Bricks, and the Pilot. If there is a hardware failure, click the failed component. The GUI takes you to component replacement pages.

FIELD DEFINITIONS

Name

Identifies the name that is assigned to a hardware component. Assign unique, meaningful component names to help you more easily locate specific components. The Pillar Axiom system maps the assigned name to the component's serial number and updates the map if you modify the component name.

By default, the Bricks are assigned names such as `/Brick001`, based on a simple sort of the component's internal Fibre Channel WUNAME. This is the logical Brick name and is not necessarily the same as the physical Brick location. You can assign any name to a Brick.

Identify *(Repair only)*

Blinks the LEDs on the specified hardware component so that you can visually identify the component from a group of like hardware components.

Replaceable Unit

Lists hardware components by type. In Repair mode, click a component identifier to start Guided Maintenance on the component.

CU *(Slammers only)*

Identifies a specific control unit (CU) in a Slammer. Each Slammer contains two CUs.

Status

Displays the current status of a hardware component. A status of Normal requires no action.

Part Number

Displays the part number of a hardware component.

Serial Number

Displays the serial number of a hardware component.

Last Read Verify Time *(Brick only)*

Indicates the most recent date and time when a drive was verified that it could be read by the Brick firmware.

SEE ALSO

Display Hardware Component Information

Identify Hardware Components

Replace a Hardware Component and the *Pillar Axiom Service Guide*

Ranges for Field Definitions

Hardware Status and Configuration Page

DESCRIPTION

Use the Hardware Status and Configuration page to select any type of hardware component that is installed on the Pillar Axiom system. After you select a type, you can select and review the status and current configuration of a specific hardware component.

FIELD DEFINITIONS

Component Type

Lists hardware components by type. Click a component identifier to display details about the component.

- **Slammers**, opens the Slammers Overview Page.

  Note: This field is available from both the **Health** and **Support** navigation contexts.
- **Bricks**, opens the Bricks Overview Page.

  Note: This field is available from both the **Health** and **Support** navigation contexts.
- **Pilot**, opens the Pilot Overview Page.

  Note: This field is available only from the **Health** navigation context.

Component Description

Briefly describes the types of hardware components in Pillar Axiom systems.

SEE ALSO

Display Hardware Component Information

Manage Hardware Components

- Display Hardware Component Information
- Modify Hardware Component Names
- Identify Hardware Components
- About Hardware Component Replacement

Ranges for Field Definitions

Health Summary Page

DESCRIPTION

Use the Health Summary page to review the status and health of the Pilot, Slammers, and Bricks that are installed on the Pillar Axiom storage system.

FIELD DEFINITIONS

Name

Identifies the hardware component's name.

Click the component image to the left of the name to open the Hardware Component, Summary Page with detailed information about the hardware component.

Status

Displays the current status of a hardware component. A status of Normal requires no action.

SEE ALSO

About System Events and Performance Statistics

- Display the Event Log
- Filter Event Log Entries

About Alert Management

- Create Alerts
- Modify Alerts

Display Performance Statistics

Manage Hardware Components

- Display Hardware Component Information
- Identify Hardware Components
- Replace a Hardware Component

Ranges for Field Definitions

Host Settings, Identity Tab

DESCRIPTION

Use the Host Settings **Identity** tab to review the storage area network (SAN) host driver information. If you are configuring iSCSI on the host port, you must also configure iSCSI on the Pillar Axiom storage system. See About iSCSI System Settings Configuration (Optional).

FIELD DEFINITIONS HOST INFORMATION

Host Name

Identifies the SAN host that accesses LUNs configured on the Pillar Axiom system. If the Pillar Axiom Path Manager is not yet installed, the system displays the World Wide Name (WWN) of the FC HBA or the iSCSI name of the iSCSI device.

IP Address

Identifies the IP address that is assigned to a SAN host.

Operating System

Identifies the operating system of the Pillar Axiom Path Manager SAN host driver.

Pillar Axiom Path Manager Version

Identifies the version of the Pillar Axiom Path Manager SAN host driver.

Number of LUNs

Identifies the number of LUNs that are mapped to that particular SAN host either because of specific mapping or because the LUN is available to all SAN hosts.

iSCSI HOST PORT INFORMATION

Alias

Identifies the iSCSI alias name and has the following default values:

- Ten-character Pillar Axiom system serial number
- Pillar Axiom model: Axiom600, Axiom500, or Axiom300

iSCSI Name

Identifies the user-assigned iSCSI Qualified Name (IQN) device name. This is the name of the iSCSI initiator for the SAN host, which includes SCSI commands and data requests into iSCSI packets for transfer across the IP network.

IP Addresses

Identifies the IP address that is assigned to a SAN host.

HOST PORT INFORMATION

Port Description

Identifies the name that is assigned to an HBA port (FC only).

Host Port

Identifies the number assigned to an HBA port:

- FC: The World Wide Name (WWN) associated with the port on the HBA
- iSCSI: The IP address of the port.

Type

Identifies the type of protocol, FC or iSCSI.

Speed (Gbps)

Displays the transmission speed, in gigabits per second, of a hardware component.

MFR

Displays the manufacturer of a hardware component.

HBA Model

Displays the model number of a hardware component.

Driver Version

Identifies version of the Pillar Axiom Path Manager SAN host driver.

Firmware Version

Displays the firmware version that is installed on a hardware component.

SEE ALSO

Modify SAN Host Settings

Ranges for Field Definitions

Host Information, Configure iSCSI Tab

DESCRIPTION

Use the Host Information **Configure iSCSI** tab to configure the storage area network (SAN) hosts to communicate with the Pillar Axiom storage system if the system requires authentication.

If you access this page from the **Hosts > Configure iSCSI Hosts** link, you can configure multiple hosts at the same time.

If you access this page from the **Hosts > Modify Host Settings** link, you can configure iSCSI for the selected host.

FIELD DEFINITIONS

Host Name *(Configure iSCSI Hosts only)*

Identifies the SAN host that accesses LUNs configured on the Pillar Axiom system. If the Pillar Axiom Path Manager is not yet installed, the system displays the World Wide Name (WWN) of the Fibre Channel (FC) HBA or the iSCSI name of the iSCSI device.

Enable Authentication

Specifies the authentication of the host (initiator) during login.

Note: If the initiator on the SAN host has been configured to require Challenge-Handshake Authentication Protocol (CHAP) authentication, login will fail unless either the Pillar Axiom system has been configured to authenticate to All Initiators or it has been set to Per-Initiator and the Enable Authentication option has been selected. In either case, you must specify the **CHAP Name** and **CHAP Secret** for the initiator.

Chap Name

CHAP authentication requires an exchange of the iSCSI initiator name (the CHAP name) and secrets (encrypted CHAP passwords) between two devices.

CHAP Secret

Identifies the encrypted CHAP authentication password (secrets) used in the exchange of user names and secrets between two devices. Both devices must support Point-to-Point Protocol (PPP) authentication.

Note: The Pillar Axiom system supports up to 100 UTF-8 non-integer characters. However, when connecting to Windows servers, you must limit the Secret to a value between 12 and 16 characters in length.

Retype CHAP Secret

Re-enter the encrypted CHAP authentication password.

Enable Access Control

Specifies that the Pillar Axiom system must reject iSCSI login attempts from initiators that have not explicitly been granted permission by the user through the Pillar Axiom Storage Services Manager interface.

Modify *(Configure iSCSI Hosts only)*

Allows the administrator to change the host name, iSCSI authentication, and access control for a specific iSCSI host.

Update *(Configure iSCSI Hosts only)*

Implements the new settings for the iSCSI host.

SEE ALSO

Modify SAN Host Settings

Ranges for Field Definitions

Host Information, LUN Connections Tab

DESCRIPTION

Use the Host Information **LUN Connections** tab to review the storage area network (SAN) host connections information.

FIELD DEFINITIONS

LUN Name

Identifies the name that is assigned to a LUN for administrative purposes. LUN names must be unique across the Pillar Axiom system, not just within its associated volume group.

LUN Number

Identifies the unique number that is assigned to a LUN and can be accessed by all SAN hosts if the LUN is not mapped. If the LUN is mapped, the LUN number must be unique to the mapped SAN host.

LUN Type

Identifies whether the displayed LUN is one of:

- **Source**, an original volume
- **Clone**, a replica (Clone LUN)

LUN Name on Host

Identifies the name that the LUN is known by this SAN host.

Host Port

Identifies the name that is assigned to an HBA port (FC only).

No. of Paths Optimized/Non Optimized

Identifies the number of paths to the LUN that are optimized (fastest path available) or non-optimized.

Slammer

Identifies the name of the Slammer on which the LUN connection resides.

CU PORTs

Identifies the status, connected or not connected, of each port on the control units (CUs).

SEE ALSO

Modify SAN Host Settings

Ranges for Field Definitions

Host Information, Settings Tab

DESCRIPTION

Use the Host Information **Settings** tab to review and modify the load balancing settings of LUNs.

FIELD DEFINITIONS

Load Balancing

Identifies the type of load balancing that the storage area network (SAN) hosts should perform to access LUNs that are configured on the Pillar Axiom storage system. Choose from:

- **Static** load balancing across multiple paths to the configured LUNs.

  The software selects the best available path, and all commands are sent over that path until the path is no longer operational, in which case the failed path a fails over to another appropriate path.

- **Round-robin** load balancing across multiple paths to the configured LUNs.

  Commands are sent by turn over the best available paths, which ensures that LUN commands are evenly distributed over any path that is available to access the LUNs.

Update Settings

Changes the load balancing settings of the selected LUNs.

PATH MANAGER SETTINGS

LUN #

Identifies the unique number that is assigned to a LUN and can be accessed by all SAN hosts if the LUN is not mapped. If the LUN is mapped, the LUN number must be unique to the mapped SAN host.

LUN Name

Identifies the name that is assigned to a LUN for administrative purposes. LUN names must be unique across the Pillar Axiom system, not just within its associated volume group.

Local Host Name

Identifies the name that the LUN is known by this SAN host.

Select All

Selects all objects.

HOST SETTINGS

HP-UX Compatibility Mode

Use this option when the SAN hosts that access the LUNs have HP-UX initiator ports and HP HBAs. When this option is enabled, the system determines LUN numbers using the HP-UX addressing scheme, allowing up to 255 LUNs. Also when enabled, the host cannot have a visible LUN using ID 0. You can verify the current host mappings in the **LUN Connections** tab.

SEE ALSO

Modify SAN Host Settings

Ranges for Field Definitions

Host Information, LUN Connections Tab

Interfaces Page

DESCRIPTION

Use the Interfaces page to create and modify the Pillar Axiom management and data-path interfaces.

FIELD DEFINITIONS

System Name

Identifies the name that is assigned to the Pillar Axiom storage system. Assign a meaningful system name. This value appears in the status bar and in multiple locations in the user interfaces (GUI and CLI).

MANAGEMENT INTERFACE

Enable DHCP

Identifies whether Dynamic Host Configuration Protocol (DHCP) is enabled.

- Enable DHCP if you have a DHCP server that automatically assigns IP addresses to network clients. This setting makes the Pillar Axiom system known to the DHCP software.
- Disable DHCP if you do not use it in your network.

Static IP Address

Identifies whether a permanent IP address will be assigned to the Pilot in a Pillar Axiom system. Choose this option if you do not use DHCP.

Public Interface

Enter an IP address that is permanently assigned to the primary control unit (CU) in the Pilot.

Pilot CU 0 and **Pilot CU 1**

Enter the IP addresses that are permanently assigned to the ports on the CUs in the Pilot. You can use these static IP addresses as an alternate method to access the active Pilot.

Subnet Mask

Enter a subnet mask for the public IP address that is permanently assigned to the Pillar Axiom system.

Gateway

Enter the public IP address of the gateway server in the subnet of which the Pillar Axiom system is a member.

NAS DATA PATH INTERFACE

Slammer Link Aggregation

Identifies whether link aggregation is enabled and, if so, over which ports. Pillar Axiom systems support the IEEE 802.3ad standard.

- **Disabled**: Turns off link aggregation.
- **PORT0, PORT1**: Enables link aggregation to treat physical links to PORT 0 and PORT 1 as one logical link. Ports 2 and 3 are separate links. Configure virtual network interfaces on port 0.

- **PORT2, PORT3**: Enables link aggregation to treat physical links to PORT 2 and PORT 3 as one logical link. Ports 0 and 1 are separate links. Configure virtual network interfaces on port 2.
- **PORT0, PORT1,PORT2**: Enables link aggregation to treat physical links to PORT 0, PORT 1, and PORT 2 as one logical link. Port 3 is a separate link. Configure virtual network interfaces on port 0.
- **PORT0, PORT1,PORT2,PORT3**: Enables link aggregation to treat physical links to PORT 0, PORT 1, PORT 2, and PORT 3 as one logical link. Configure virtual network interfaces on port 0.
- **PORT0, PORT1** and **PORT2, PORT3**: Enables link aggregation to treat physical links to PORT 0 and PORT 1 as one logical link and to PORT 2 and PORT 3 as another logical link. Configure virtual network interfaces on port 0 for the first set and on port 2 for the second set.

Note: The specified port settings above apply to *both* control units of the Slammer. Also, each link aggregation group spans a single control unit only.

PORT SETTINGS

Identifies the network port settings for the Pillar Axiom management interface.

Important! Use care when setting the transmit speed and duplex mode. A management interface setting that is not supported by the external network could result in loss of access to the Pilot. If access is lost, contact the Pillar World Wide Customer Support Center for assistance.

Transmit Setting

Select from the drop-down list the speed and duplex mode that you want the Pilot management interface to use:

- Auto negotiate
- 10 Mb/sec, half duplex
- 10 Mb/sec, full duplex
- 100 Mb/sec, half duplex
- 100 Mb/sec, full duplex
- 1000 Mb/sec, half duplex
- 1000 Mb/sec, full duplex

Note: Auto negotiate is the default transmit setting. We recommend the default setting for all but special circumstances.

Transmit Mode

Displays the actual speed and duplex mode being used by the management interface.

SEE ALSO

Configure Management and Data Path Interfaces

Modify System Name

Ranges for Field Definitions

I/O Port Details Page

DESCRIPTION

Use the I/O Port Details page to review the status of the Fibre Channel interfaces of Bricks and Slammers. You can also review the Fabric Switch interfaces, as well as the data path and management interfaces of Slammers.

FIELD DEFINITIONS

Port

Lists by type the Fibre Channel ports on a Brick or the ports on a Slammer:

- For Bricks:
  - **RAID Controller Module** (0 or 1)
  - **FC0** through **FC3** or **Cascade** (for FC Bricks only)
  -
- For Slammers:
  - **Private Interconnect Module** (PIM)
    - **FC0** through **FC2**
    - **FS0** through **FS9**
    - **ETH0** through **ETH2**
  - **Network Interface Module** (NIM)
    - 2–port NIM: **PORT0** and **PORT1**
    - 4–port NIM: **PORT0** through **PORT3**

CU *(Slammers only)*

Identifies a control unit (CU) of the Slammer.

Type *(Slammers only)*

Identifies the types of network ports for data path traffic between the customer network switches and the Pillar Axiom Slammers:

- **Copper**, which identifies RJ-45 copper interfaces.
- **Optical-L**, which identifies longwave optical small form-factor pluggable (SFP) transceiver interfaces.
- **Optical-S**, which identifies shortwave optical SFP transceiver interfaces.

Connected

Identifies the connection status of the port.

The Pillar Axiom user interfaces (the GUI and CLI) show that host Fibre Channel (FC) HBA ports are either **Connected** or **Not Connected** to the Slammer ports. The meaning of **Connected** is that the HBA port on the SAN host has logged in to the port on the Slammer using the FC protocol. In most operating systems, host ports log in to the Slammer ports immediately after the two are physically connected and enabled and remain logged in until the physical connection is broken. So, **Connected** effectively means that there is an enabled physical connection between the ports.

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

Speed

Displays the transmission speed, in megabits per second, of the port.

Topology *(Slammers only)*

Identifies the Fibre Channel (FC) transport topology in use by the ports in the network interface module (NIM) to connect to the storage area network (SAN) employed by the customer:

- **Fabric**, which means that the port is an N_Port in a switched fabric (FC-SW).
- **Loop**, which means that the port is an NL_Port in an arbitrated loop (FC-AL).
- **Point-to-Point**, which means that the port is an N_Port that is connected to another N_Port, back to back (FC-P2P).
- **Public Loop**, which means that the port is an NL_Port that is connected to a loop in which one port in the loop is an FL_Port in the fabric (FC-FLA).

Note: The topology used by Storage System Fabric (SSF) between the Slammer PIMs and the Brick RAID controllers is private and therefore not reported.

SFP Status

Displays the status of the small form-factor pluggable (SFP) transceiver.

- **Bypassed**
- **Bypassed-No SFP**
- **Bypassed-Incorrect Speed**
- **Bypassed-Read Error**
- **Bypassed-Incorrect Type**
- **Bypassed-Lost Sync**

If the interface module itself should fail, the SFP status shows **Hardware Failure**.

Note: The **SFP Status** and **SFP Vendor Info** fields display information only when version 2 private interconnect modules (PIMs) are connected to version 2 SATA controllers using optical SFPs. In all other cases, these two fields are blank.

SFP Vendor Info

Displays the vendor's part number for the SFP. If that information is not available, the system displays **Unknown**. See also the preceding note.

SEE ALSO

Display Hardware Component Information

Ranges for Field Definitions

iSCSI Page

DESCRIPTION

Use the iSCSI page to configure system-wide iSCSI settings if you have iSCSI hosts configured to use Challenge-Handshake Authentication Protocol (CHAP), Access Control, Internet Storage Name Service (iSNS), or some combination of these parameters. This configures the authentication and access controls on the Pillar Axiom storage system in which the host must match to gain access. If you have CHAP and Access Control configured on a per-initiator basis, then you do not need to configure iSCSI globally.

FIELD DEFINITIONS

System Name

Identifies the name that is assigned to the Pillar Axiom system.

ACCESS CONTROL

Choose one of these methods:

None

Axiom Access Control

Specifies that the Pillar Axiom system rejects iSCSI login attempts from initiators that have not explicitly been granted permission by the user through the Pillar Axiom Storage Services Manager interface.

iSNS Access Control

If you enable iSNS-based access control, only those initiators that are members of Discovery Domains of which the Pillar Axiom system is a member will be allowed to login to the Pillar Axiom system.

iSCSI CONFIGURATION

iSCSI Device Name

Identifies the user-assigned iSCSI Qualified Name (IQN) device name. This is the name of the iSCSI initiator for the SAN host, which includes SCSI commands and data requests into iSCSI packets for transfer across the IP network.

Alias

Identifies the iSCSI alias name and has the following default values:

- Ten-character Pillar Axiom system serial number
- Pillar Axiom model, Axiom500 or Axiom300

Authentication

Identifies the authentication of the host (initiator) during login. Select one of:

- **Per-Initiator**: Specifies that CHAP authentication is required only for those iSCSI connections for which it is configured for each host.

- **All Initiators**: Specifies that CHAP authentication is required for all iSCSI connections to the Pillar Axiom system, regardless of what is configured for each host.

Note: If the initiator on the SAN host has been configured to require CHAP authentication, then login will fail unless either the Pillar Axiom system has been configured to authenticate to All Initiators or it has been set to Per-Initiator and the Enable Authentication option has been selected. In either case, you must specify the CHAP Name and CHAP Secret for the initiator.

Enable Header Digest

When an iSCSI initiator logs in to the Pillar Axiom system, they negotiate the parameters for the iSCSI session. If the initiator does not give the system a choice regarding the use of iSCSI header digests, the system complies with what the initiator wants. If the initiator gives the system a choice and if Enable Header Digest is enabled, the system will choose to use header digests, regardless of the initiator's preference.

Note: When selected, this parameter provides additional error checking for the header portion of the iSCSI packet.

Enable Data Digest

When an iSCSI initiator logs in to the Pillar Axiom system, they negotiate the parameters for the iSCSI session. If the initiator does not give the system a choice regarding the use of iSCSI data digests, the system complies with what the initiator wants. If the initiator gives the system a choice and if Enable Data Digest is enabled, the system will choose to use data digests, regardless of the initiator's preference.

Note: When selected, this parameter provides additional error checking for the data portion of the iSCSI packet.

Enable Bi-Directional CHAP

Enables the CHAP protocol to be used for both requests for data (from the iSCSI initiator) and responses to requests (from the iSCSI target). If bi-directional CHAP support is disabled for the Pillar Axiom system, bi-directional CHAP must be disabled for all initiators, or initiator login will fail.

Axiom CHAP Secret

Identifies the encrypted CHAP authentication password (secrets) used in the exchange of user names and secrets between two devices. Both devices must support Point-to-Point Protocol (PPP) authentication.

Note: The Pillar Axiom system supports up to 100 UTF-8 non-integer characters. However, when connecting to Windows servers, you must limit the Secret to a value between 12 and 16 characters in length.

Retype Axiom CHAP Secret

Re-enter the encrypted CHAP authentication password used.

iSNS CONFIGURATION

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.

Important! If an iSCSI port has no access or loses access to the iSNS server, the Pillar Axiom system reports iSNS error events but continues to operate normally. For iSNS Access Control to function correctly, at least one Pillar Axiom iSCSI port must have access to the iSNS server during a restart; otherwise, all iSCSI logins are rejected.

For information on configuring the Microsoft iSNS Server, refer to the *Pillar Axiom iSCSI Integration Guide for Windows Platforms*.

Enable iSNS Server Registration

Choosing this option allows Pillar Axiom iSCSI targets to be registered in the iSNS server.

For discovery of the iSNS server IP address, specify either DHCP or static addressing:

- **DHCP**. For this option, you must use a Microsoft DHCP server that has been configured by the Microsoft iSNS Server installer to return the server IP address using DHCP option 43 (vendor-specific) or DHCP option 83 (iSNS).

  Note: Microsoft does not support DHCP option 83 until the Windows Server 2008 release.
- **Static**. For this option, specify:
  - **Server IP Address**
  - **Server TCP port**

SEE ALSO

Configure iSCSI System Settings

Modify SAN Host Settings

Ranges for Field Definitions

iSCSI Port Settings Page

DESCRIPTION

Use the iSCSI Port settings page to configure the settings for the Slammer ports dedicated to Internet SCSI (Small Computer System Interface), or iSCSI, protocol.

FIELD DEFINITIONS

TCP Port Number

Identifies the Transmission Control Protocol (TCP) port number that is configured on the iSCSI port.

MTU (bytes)

Identifies the maximum transmission unit (MTU) value.

The frame size (MTU) does not include the Ethernet header portion of the packet. If your network switch has trouble with this, you can set the switch to a larger value or lower the MTU size to correct the problem.

If your network supports extended Ethernet (jumbo) frames, enter an integer greater than 1500 and less than 9001. Make sure that this Pillar Axiom MTU size matches the network MTU size. If the MTU sizes are mismatched, performance may be severely degraded.

DHCP

Select **DHCP** if you have a DHCP server that automatically assigns IP addresses to network clients. This setting makes the Pillar Axiom system known to the DHCP software.

Manual

Select **Manual** if you do not use DHCP in your network.

Primary IP Address

Enter an IP address that is permanently assigned to the primary control unit (CU) in the Pilot.

Netmask

Enter a netmask for the IP address that is permanently assigned to the Pillar Axiom system.

Gateway IP Address

Enter a gateway IP address that is permanently assigned to the Pillar Axiom system.

CU

Identifies a control unit (CU) in a hardware component.

Modify

Allows you to edit the settings for the selected iSCSI port.

Update

Applies the changes you made to the settings for the selected iSCSI port.

Cancel

Cancels the edit mode and retains the original settings for the selected iSCSI port.

SEE ALSO

Modify iSCSI Port Settings

Modify SAN Host Settings

Ranges for Field Definitions

Join Domain Page

DESCRIPTION

Use the Join Domain page to establish a connection between a File Server enabled for Common Internet File System (CIFS) operations and the Windows domain controller that contains valid CIFS user accounts. The Pillar Axiom storage system uses the connection to authenticate CIFS users.

FIELD DEFINITIONS

Administrator

Identifies the administrator user name for the domain controller that is associated with the Pillar Axiom system.

Domain Password

Identifies the administrator password for the CIFS domain controller.

SEE ALSO

Join File Servers to CIFS Domains (NAS Systems)

Ranges for Field Definitions

LUN, Host Connections Tab

DESCRIPTION

Use the LUN **Host Connections** tab to review the connections.

FIELD DEFINITIONS

Refresh Table

Updates the page with current data.

LUN Name

Identifies the name that is assigned to a LUN for administrative purposes. LUN names must be unique across the Pillar Axiom storage system, not just within its associated volume group.

Host

Identifies the storage area network (SAN) host that accesses LUNs configured on the Pillar Axiom system. If the Pillar Axiom Path Manager is not yet installed, the system displays the World Wide Name (WWN) of the Fibre Channel (FC) HBA or the name of the Internet SCSI (Small Computer System Interface), or iSCSI, device.

LUN Number

Identifies the unique number that is assigned to a LUN and can be accessed by all SAN hosts if the LUN is not mapped. If the LUN is mapped, the LUN number must be unique to the mapped SAN host.

Host Port

Identifies the name that is assigned to an HBA port (FC only).

Status of FC Ports by Slammer Port

Identifies the status, connected or not connected, of each Fibre Channel port on the control units (CUs).

Status of iSCSI Ports by Slammer Port

Identifies the status, connected or not connected, of each iSCSI port on the CUs.

SEE ALSO

Create LUNs

Modify a LUN

Ranges for Field Definitions

LUN, Identity Tab

DESCRIPTION

Use the LUN **Identity** tab to review details of the selected LUN and to assign LUNs to storage area network (SAN) hosts.

FIELD DEFINITIONS

LUN Name

Identifies the name that is assigned to a LUN for administrative purposes. LUN names must be unique across the Pillar Axiom system, not just within its associated volume group.

Inactive *(Create and Modify clone only)*

Specifies that the cloned LUN cannot be accessed from the data path.

Create LUN in *(Create only)* or **Volume Group** *(Display and Copy LUN only)*

Lists the name of the volume group where the volume (filesystem or LUN) is located.

LUN ASSIGNMENT

Auto-Assign LUN to a Slammer *(Create only)*

The Pillar Axiom system assigns the LUN on the Slammer.

Assign LUN to Slammer *(Create only)*

Specifies a specific Slammer in which to assign the LUN. Use this option when you know from where you want to access the data.

Assign Control Unit *(Create only)* or **Default Configured Slammer** *(Modify only)*

Specifies a specific control unit (CU) on the Slammer in which to assign the LUN.

LUN ACCESS

Only selected hosts

Specifies only designated SAN hosts to access this LUN using the LUN number.

All hosts may access this LUN using LUN Number

Specifies all SAN hosts to access this LUN using the LUN number.

Fibre Channel

Specifies the protocol in which the hosts can access this LUN.

Tip: Fibre Channel paths will always be used at a preference over iSCSI paths. Also, load balancing will not be mixed between these two protocols.

iSCSI

Specifies the protocol in which the hosts can access this LUN.

BACKGROUND COPY AND DATA

Identifies the strategy the system should use to control the impact on performance when background tasks need to copy or move data from one location in the storage pool to another:

MIGRATION PRIORITY

- **Minimize performance impact**. Restricts the amount of work performed on a loaded system. This option is intended to have a minimal impact on client I/O throughput at the expense of longer copy times.
- **System optimized performance**. Balances the background copy with the incoming client I/O. This option is the default.
- **Maximum performance**. Prioritizes the background copy at the expense of client I/O throughput.

Note: When the system is idle or lightly loaded, the above background task maximizes the amount of work done regardless of the option selected.

The following types of operations are affected by the strategy you select:

- Copy
- Restore
- Quality of Service (QoS) changes in:
  - Priority
  - Redundancy
  - Storage Class

Data transfer operations invoked by the Pillar Axiom MaxRep Replication for SAN utility are not affected by the **Background Copy and Data Migration Priority** setting.

SEE ALSO

Create LUNs

Modify a LUN

Ranges for Field Definitions

LUN, Mapping Tab

DESCRIPTION

Use the LUN **Mapping** tab to create and modify the LUN-to-host mapping settings for a LUN. This tab is only available when the LUN mapping option is enabled.

FIELD DEFINITIONS

Host

Identifies the SAN host that accesses LUNs configured on the Pillar Axiom storage system. If the Pillar Axiom Path Manager is not yet installed, the system displays the World Wide Name (WWN) of the Fibre Channel (FC) HBA or the iSCSI name of the iSCSI device.

Available LUN Number

Displays the available numbers that you can assign to the LUN. All SAN hosts can access the LUN by this unique identifier.

Assigned Slammer

Displays available Slammers in which to assign the LUN.

Select Port Mask

Displays the physical Slammer ports to exclude (mask) so they cannot access the LUN.

Create Map

Creates the LUN-to-host mapping based on your selections.

Modify Map

Changes the LUN-to-host mapping based on your selections.

LUN MAPPING

Host Name

Identifies the SAN host that accesses LUNs configured on the Pillar Axiom storage system. If the Pillar Axiom Path Manager is not yet installed, the system displays the WWN of the FC HBA or the iSCSI name of the iSCSI device.

LUN Number

Identifies the unique number that is assigned to a LUN and can be accessed by all SAN hosts if the LUN is not mapped. If the LUN is mapped, the LUN number must be unique to the mapped SAN host.

Port Mask Set

Identifies whether or not the physical port is masked so that the selected LUN cannot be accessed from this port.

Remove Map

Removes the mapping so that all SAN hosts can access the LUN using the LUN number.

SEE ALSO

Create LUNs

Modify a LUN

Define LUN Mapping (Optional)

Ranges for Field Definitions

LUN Performance Page

DESCRIPTION

Use the LUN Performance page to review performance statistics for LUNs.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. Choose from:

- 1 min
- 5 min
- 10 min

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes.

Refresh Now

Updates the page with current data.

Name

Identifies the name that is assigned to a LUN for administrative purposes. LUN names must be unique across the Pillar Axiom system, not just within its associated volume group.

Capacity

Identifies the maximum capacity limit that is assigned to the object.

Priority

Identifies the priority of the specified LUN in relation to other LUNs.

IOPs

Identifies the current performance for input (read) and output (write) operations for the LUN.

MB/s

Identifies the data transfer rate of the specified LUN.

Collection Period

Identifies the time at which information was last collected from the Pillar Axiom system.

SEE ALSO

Display Performance Statistics

Ranges for Field Definitions

LUN, Quality of Service Tab

DESCRIPTION

Use the **Quality of Service** tab to create and modify the capacity and performance settings for a LUN.

FIELD DEFINITIONS

Initial Capacity *(Create only)*

Identifies the initial capacity that is assigned to the logical volume. This value is a soft limit, which means that data can be stored in a logical volume until the maximum capacity is reached.

Maximum Capacity

Identifies the maximum capacity to which the logical volume can grow. For a clone, this field identifies how much addressable space will be available.

Redundancy

Identifies how many mirror copies of the original data are stored online.

Important! Pillar highly recommends that you consult with a Pillar Customer Support professional for assistance with sizing your system and creating your logical volumes (filesystems and LUNs).

Redundancy options include:

- **Standard**, to store original data only. Data striping over multiple RAID-5 groups maintains full redundancy, even without mirror copies.

  Note: Standard does not maintain redundancy at the LUN level; however, it does provide sufficient data protection for most purposes.
- **Double**, to store original data and one mirror copy, with data striping over multiple RAID-5 groups.

  Note: Double Redundancy can only provide true redundancy if your system has enough Bricks to allocate the filesystem or LUN such that no two mirror copies share a RAID group.

A SATA Brick has two RAID groups. A Fibre Channel (FC) Brick has one RAID group.

If the storage pool is becoming depleted, or a large filesystem or LUN is created, it might be necessary to place the filesystem or LUN on more RAID group fragments.

Depending on the ability of the system to allocate sufficient contiguous storage blocks for the size of the logical volume, refer to the following number of RAID groups to configure your volumes for the best performance:

Table 21 Optimum number of RAID groups for best performance

Priority	SATA standard redundancy	SATA double redundancy	FC standard redundancy	FC double redundancy
Archive	4	8	2	4
Low	4	8	2	4
Medium	6	12	3	6
High	8	16	4	8
Premium	8	16	4	8

Note: When the selected Storage Class is **SSD**, all available SSD drives are striped across, regardless of Priority.

For *performance testing purposes only*, create a filesystem or LUN using standard redundancy and the HighThroughput Profile. This is *not* recommended for most applications. Reset your system after you have created a HighThroughput LUN or filesystem before you configure normal filesystems or LUNs for applications.

Optimizer

Displays how much capacity is available for each combination of Storage Class and Redundancy.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

Run Simulation *(Create only)*

Shows how the QoS settings for this LUN will affect the performance of other LUNs. Values that are less than the currently configured settings are highlighted.

Storage Class capacity *(Create and Modify only)*

For the assigned Storage Class, identifies its physical capacity in terms of the following:

- **Free**. Available physical storage that can be allocated.
- **Reconditioning**. Capacity in the process of being released to the free pool.
- **Used**. Physical storage already allocated.
- **Total**. Total physical capacity of the specified Storage Class.

Clone LUN Space

Specifies the maximum amount of space to make available on the Pillar Axiom system for Clone LUNs.

Caution We strongly recommend that you allocate sufficient repository capacity to minimize the chances of running out of this space (which could lead to data inconsistency or loss). To set sufficient capacity, use a value equal to the source volume capacity times the number of replicas times the maximum rate of change. For example, for a 100 GB volume that is projected to have 20 active replicas at a time and no more than a 20% rate of change, use a value of 400 GB for the clone repository.

Storage Class

Identifies the category of physical storage on which the logical volume resides:

- **FC** (Fibre Channel drives)
- **SATA** (Serial ATA drives)
- **SLC SSD** (single level cell, solid state drives)

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

Priority vs. other Volumes

The Priority option determines how much of the system resources are devoted to that volume, including the allocation of Slammer CPU cycles and the allocation of specific portions of the disk platters. The higher the Priority, the greater the allocation of CPU time and the faster the media access time.

The Priority option specifies the layout of data that is stored in the Pillar Axiom system based on one of the following settings (see also the figure below):

- **Premium** priority is the highest possible performance. This priority uses the outer 20% of the drive platters.

  Note: I/O is typically faster when data resides on the outer edge of the drive platters.
- **High** priority logical volumes are allocated space in the outer 20%-40% of the drive platters.

  Tip: If sufficient Bricks are present, performance of high priority volumes is enhanced by placing those volumes on a greater number of Bricks.
- **Medium** priority logical volumes are allocated space in the outer 40% to 60% of the drive platters.
- **Low** priority logical volumes are allocated space in the band that is 60% to 80% from the outer diameter of drive platters.
- **Archive** (lowest-priority) logical volumes are allocated space in the inner areas of the drive platters.

Figure 21 Example of HDD priority bands

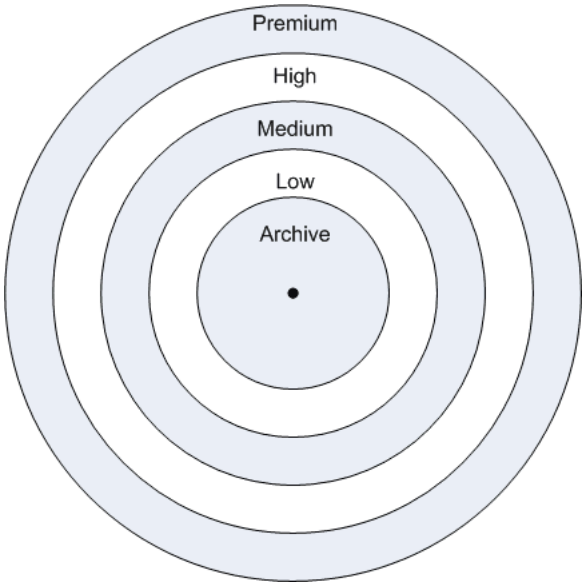

If needed, the Pillar Axiom system allocates space in a higher or lower priority band (within the same Storage Class) than the one that you choose.

The Pillar Axiom system uses the **Prioritization** and **Redundancy** values to calculate whether enough total capacity is available to create the new volume in the specified Storage Class. The calculated results are available in the table that is displayed when you click **Optimizer**.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

Data is typically accessed

Identifies the typical data access method. Choose from the values in the table below.

Table 22 Effects of access and I/O bias

Access bias	I/O bias	Read-ahead in the Slammer	RAID configuration in the Brick
Sequential	Read	Aggressive	RAID 5 Writes data in write-back mode to physical storage in full-stripe extents.
	Write	Conservative	RAID 5

Table 22 Effects of access and I/O bias (continued)

Access bias	I/O bias	Read-ahead in the Slammer	RAID configuration in the Brick
			Writes data in write-back mode to physical storage in full-stripe extents.
Mixed and random	Read	None	RAID 5 Writes data in write-back mode to physical storage in full-stripe extents.
Random	Write	None	Distributed RAID

Note: This is an optimization bias, not a requirement that all data or data operations conform to the value.

I/O Bias

Identifies the typical read-write ratio. Choose from:

- **Read**, if users or applications read data more often than they write to the data source.
- **Write**, if users or applications write data more often than they read it.

  Important! If you choose Random as the access method and Write as the I/O Bias, the system creates the LUN with a Distributed RAID geometry. This geometry provides enhanced write performance but uses twice the capacity. For more information, refer to About Enhanced Performance for Random Write Operations.
- **Mixed**, if the read-write ratio varies.

Note: This is an optimization bias, not a requirement that all data or data operations conform to the value.

The system stores all writes of user data and system metadata in mirrored copies of the journal. Although the filesystem or LUN is journaled, for some of the profiles, write caching is disabled since it tends to have very low hit rates.

One copy is maintained in non-volatile memory on one control unit (CU) of a Slammer. The mirror copy is maintained in one of:

- Battery-backed memory of the partner CU on the Slammer (preferred location). Writes to this copy are equivalent to write-back cache.
- Virtual LUN (VLUN) that is reserved on disk for the logical volume if the partner CU is unavailable for the write. Writes to this copy are equivalent to write-through cache.

Writes from the journal to permanent disk storage are equivalent to write-through cache. The system flushes user data and the corresponding metadata as a unit to disk.

SEE ALSO

Create LUNs

Modify a LUN

Ranges for Field Definitions

Move Volumes Page

DESCRIPTION

Use the Move Volumes page to move logical volumes (filesystems or LUNs) or nested volume groups to a different volume group.

FIELD DEFINITIONS

Move Selected Item

Displays a list of filesystems, LUNs, and volume groups. Select one or more objects to move them to a new location.

Destination

Displays a list of volume groups and nested volume groups. Select one destination into which the selected filesystems, LUNs, and volume groups will be moved.

SEE ALSO

Move a Logical Volume to a Different Volume Group

Ranges for Field Definitions

NAS Exports Overview Page

DESCRIPTION

Use the NAS Exports Overview page to review the Network File System (NFS) exports that are configured on the Pillar Axiom storage system. The **Actions** drop-down list provides options to create, modify, delete, and view NFS exports.

FIELD DEFINITIONS

Export

Lists the names of configured NFS exports. Click a name to review or modify the export settings.

File Server

Identifies the name that is assigned to a File Server.

Filesystem

Identifies the name that is assigned to a filesystem.

Host(s)

Identifies the NFS clients, or hosts, that can mount the NFS export.

- **All Hosts**: Everyone can access the export point.
- **Single Host**: Only the specified host can access the export point.
- **NIS Netgroup**: Everyone within the Network Information Service (NIS) netgroup can access the export point.
- **Network**: Everyone on the specified subnet can access the export point.
- **Read Only**: The specified host has read-only privileges.
- **Root Access**: The specified host has root permissions.

SEE ALSO

Modify NFS Exports

Add NFS Exports to a Filesystem (External File)

Ranges for Field Definitions

NAS Protocols Page

DESCRIPTION

Use the Performance NAS Protocols page to select any type of performance statistics for network attached storage (NAS) protocols. After you select a type, you can select and review those statistics.

FIELD DEFINITIONS

CIFS/NFS

Review how many connections are in use by Common Internet File System (CIFS) and Network File System (NFS) clients for the specified time period.

TCP/IP

Review the current TCP/IP performance statistics for output (Send) and input (Receive) packets.

SEE ALSO

Display Performance Statistics

Performance NAS Protocols CIFS/NFS Page

NAS Protocols Performance, TCP/IP Page

Ranges for Field Definitions

NAS Protocols Performance, TCP/IP Page

DESCRIPTION

Use the NAS Protocols Performance, TCP/IP page to review the performance statistics for output (Send) and input (Receive) packets for the specified time period.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. Choose from:

- 1 min
- 5 min
- 10 min

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes.

Refresh Now

Updates the page with current data.

Slammer

Identifies the name of the Slammer that contains TCP/IP statistics.

Control Unit

Identifies the control unit (CU) of the Slammer that contains the statistics.

Send (MB/s)

Displays the statistics for output (Send) packets.

Receive (MB/s)

Displays the statistics for input (Receive) packets.

Collection Period

Identifies the time at which information was last collected from the Pillar Axiom system.

SEE ALSO

Display Performance Statistics

Ranges for Field Definitions

NAS Shares Overview Page

DESCRIPTION

Use the NAS Shares Overview page to review the Common Internet File System (CIFS) shares that are configured on the Pillar Axiom storage system. The **Actions** drop-down list provides options to create, modify, delete, and view CIFS shares.

FIELD DEFINITIONS

Name

Lists the names of configured CIFS shares. Click a name to review or modify the share settings.

File Server

Identifies the name of a File Server with which the filesystem is associated.

Filesystem

Identifies the name of a filesystem with which the share is associated.

Path

Identifies the full path to the CIFS share.

Enabled

Identifies whether the CIFS share is enabled.

- Enabled shares are active. Users can access an enabled share point.
- Disabled shares are inactive. Users cannot access a disabled share point.

SEE ALSO

Modify CIFS Shares

Ranges for Field Definitions

NAS Storage Overview Page

DESCRIPTION

Use the NAS Storage Overview page to select any type of network attached storage (NAS) object that is configured on the Pillar Axiom storage system. After you select a storage object type, you can select and manage a filesystem, export, share, or File Server.

FIELD DEFINITIONS

NAS Storage Object Type

Lists types of NAS objects stored in the Pillar Axiom system. Click a type link to display a list of the selected type of objects.

- **Filesystems** opens the Filesystem Overview Page.
- **Exports** opens the Filesystem Page, Exports Tab.
- **Shares** opens the Filesystem Page, Shares Tab.
- **File Servers** opens the File Server Overview Page.

NAS Storage Object Description

Briefly describes the types of NAS storage objects that can be configured on Pillar Axiom systems.

SEE ALSO

Create Filesystems (NAS Storage Systems)

Modify Filesystem Attributes

Add NFS Exports to a Filesystem (External File)

Add NFS Exports to a Filesystem (External File)

Modify NFS Exports

Define CIFS Options

Modify CIFS Shares

Create File Servers (NAS Storage Systems)

Modify File Server Attributes

Ranges for Field Definitions

NDMP Configuration Page

DESCRIPTION

Use the NDMP Configuration page to enable the Network Data Management Protocol (NDMP) and to modify the NDMP settings.

The Pillar Axiom storage system supports:

- CommVault Galaxy 6.1.
- Symantec Veritas NetBackup 6.0.
- EMC NetWorker 7.3.
- BakBone NetVault: BackUp 7.
- Oracle Secure Backup 10.1.0.3.

FIELD DEFINITIONS

Enable NDMP

Identifies whether NDMP support is enabled.

- Enable NDMP to use NDMP backup applications to back up data.
- Disable NDMP to deactivate support for NDMP backup applications.

Port

Identifies a TCP port that is assigned as the listen port for the NDMP daemon. The default NDMP port number is 10000.

Username

Identifies the NDMP user's login name.

Change Password

Opens the Change Password Page.

File Server

Identifies the File Server with which the NDMP attributes are associated. Although you enable or disable the NDMP feature system-wide, all NDMP backup and restore operations must go through this one File Server.

View

Opens the File Server Page, Network Tab in read-only mode so that you can review the File Server definitions.

SEE ALSO

Create an NDMP User Account

Ranges for Field Definitions

Networking Overview Page

DESCRIPTION

Use the Networking Overview page to review the settings for Pillar Axiom networking properties.

FIELD DEFINITIONS

System Name

Identifies the name that is assigned to the Pillar Axiom storage system. Assign a meaningful system name. This value appears in the status bar and in multiple locations in the user interfaces (GUI and CLI).

MANAGEMENT INTERFACE

DHCP

Identifies whether Dynamic Host Configuration Protocol (DHCP) is enabled.

IP Address

Identifies the IP address that is assigned to the primary control unit (CU) in the Pilot.

Subnet Mask

Identifies the subnet mask for the IP address that is permanently assigned to the Pillar Axiom system.

Gateway

Identifies the IP address of the gateway server in the subnet of which the Pillar Axiom system is a member.

Transmit Setting

Identifies the requested port speed and duplex setting for the management interface on the Pillar Axiom system.

Transmit Mode

Identifies the actual port speed and duplex setting that the management interface is running.

NAS DATA PATH INTERFACE

Slammer Link Aggregation

Identifies whether link aggregation (IEEE 802.3ad standard) is enabled for each Slammer.

NOTIFICATION

Email

Identifies whether email notification is enabled.

Email Server IP

Identifies an email server that receives alerts.

Automatic Failback

Identifies whether the Pillar Axiom storage system should perform an automatic recovery operation when a previously unavailable Slammer control unit (CU) becomes available.

CALL-HOME

Call-Home

Indicates whether Call-Home is enabled.

- **Enabled**: The Pillar Axiom system sends status messages and alerts to Pillar Data Systems.
- **Disabled**: The Pillar Axiom system does not send status messages and alerts.

Call-Home Server

Identifies whether the Pillar Data Systems sever or a local server receives the system information.

Name

Identifies the name of the Call-Home server.

Transfer Protocol

Identifies the method used to send Call-Home logs to the Call-Home server.

Proxy Server

Identifies whether a proxy server is enabled.

Proxy Server IP Address *(if Proxy Server is enabled)*

Identifies the IP address of the proxy server.

Proxy Server Port *(if Proxy Server is enabled)*

Identifies the port number used to access the proxy server.

Proxy Server Type*(if Proxy Server is enabled)*

Identifies the protocol the Pillar Axiom system uses when accessing the proxy server:

- HTTP
- SOCKS4
- SOCKS5

Large Files

Indicates whether large files are sent to the Call-Home server.

SEE ALSO

Configure Management and Data Path Interfaces

Modify System Name

Ranges for Field Definitions

NFS Exports Configuration Page

DESCRIPTION

Use the NFS Exports Configuration page to create, modify, view, and delete Network File System (NFS) exports for a filesystem.

When you modify an NFS export, you can modify only the access permissions and the hosts that have access to the export.

FIELD DEFINITIONS

Filesystem

Identifies the name of the filesystem with which the NFS export is associated. You can define multiple exports for a filesystem, and you can define multiple hosts for each export.

Export Path

Enter a full path to the NFS export. To create an export for the filesystem, the path must start with a forward slash (/). Each entry must export a different path name. To assign different permissions to different hosts on a specific export, create multiple host entries for the same export.

For example, if you want to create an export for the directory `users` in the filesystem `test`, you must first export `test` to a client with root privileges. This export will use a forward slash (/) to export the entire directory.

From the client with root privileges, do the following:

1. Mount the filesystem `test` to a suitable mount point on the client (for example, `/mnt`).
2. Create the directory users from this client as `/mnt/users` and set the permission on users from the client.
3. Export the directory users to other clients, as required.

This export will be for `/users`, not `/test/users`.

UID

Identifies the user ID (UID) for anonymous users.

- Set the UID to zero (0) if you want anonymous-user access to be identified as the root user. This is a very insecure setting. Setting anonymous to root allows all users full control.
- Set the UID to the “nobody” value (often -2, but not always) if you want anonymous-user access to be identified as the user “nobody.” In most instances, this is the recommended setting.
- Set the UID to any other user ID that you want to be identified as anonymous-user access.

Host Access

Identifies the NFS clients, or hosts, that can mount the NFS export.

- **All Hosts**: Everyone can access the export point.

- **Single Host**: Only the specified host can access the export point.
- **NIS Netgroup**: Everyone within the Network Information Service (NIS) netgroup can access the export point.
- **Network**: Everyone on the specified subnet can access the export point. **Netmask** specifies the screen to use to determine which host computers can access this export.
- **Read Only**: The specified host has read-only privileges.
- **Root Access**: The specified host has root permissions.

Create *(Create only)*

Creates a new export having the specified export parameters and host definitions.

Modify *(Create only)*

Allows the access permissions for a selected export to be changed by copying the attributes of the selected host into the editable fields so that you can modify the current settings. To apply the new permissions, click **Update**.

Update *(Create only)*

Updates the access permissions for the selected export. This button is available if **Modify** was previously clicked.

Cancel *(Create only)*

Cancels the update operation. This button is available if **Modify** was previously clicked.

Remove *(Create only)*

Deletes the selected exports from the list.

Additional Hosts

Allows you to create additional hosts.

Unix Exports (NFS) *(Create only)*

Displays a list of currently configured NFS exports. Check the box to the left of an export point to select it for modification or removal.

SEE ALSO

Modify NFS Exports

Add NFS Exports to a Filesystem (External File)

Ranges for Field Definitions

Notification Page

DESCRIPTION

Use the Notification page to configure the electronic mail (email) server that receives alerts from the Pillar Axiom storage system and sends the email messages to the designated recipients.

FIELD DEFINITIONS

System Name

Identifies the name that is assigned to the Pillar Axiom system.

EMAIL

Enable Email

Identifies whether email is enabled.

- Enable email if you intend to define alerts to send email notifications.
- Disable email if you do not want to send email notifications from the Pillar Axiom system.

Email Server IP

Specifies an email server that receives alerts. Enter the IP address of an email server in which the recipients have email accounts.

Enable Automatic Failback of NAS Control Units

Specifies whether the Pillar Axiom system should perform an automatic recovery operation when a previously unavailable Slammer control unit (CU) becomes available.

Note: The Failback process disrupts network attached storage (NAS) service for all filesystems and virtual interfaces that have been failed over to the surviving CU for approximately 30 sec.

- Enable this option if interruptions of up to 30 sec in data-path connections are permitted while the automatic recovery takes place. If enabled, a NAS Slammer CU that is in Failed Over status will transition to Failback status and if that Failback succeeds, the CU will be placed in Normal status and back in service.
- Disable this option if you want to be prompted to perform the operation manually when the Slammer CU is ready for recovery.

SAN Slammer failover is enabled automatically and cannot be disabled.

SEE ALSO

Collect Debug Logs

Ranges for Field Definitions

Performance Backup Page

DESCRIPTION

Use the Performance Backup page to review performance statistics for backups.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. Choose from:

- 1 min
- 5 min
- 10 min

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes.

Refresh Now

Updates the page with current data.

Filesystem

Opens the Filesystem Page, Identity Tab in read-only mode so that you can review the filesystem settings.

Slammer

Identifies the name of the Slammer on which a backup or restore operation is running.

Control Unit

Identifies the control unit (CU) of the Slammer on which a backup or restore operation is running.

Start Time

Identifies the time at which the Pillar Axiom system started a backup or restore operation.

Duration

Identifies how much time has been spent on a backup or restore operation. Duration is measured from the start time to the time at which the system received the request to gather backup statistics.

Backup Average (GB/hr)

Identifies the average rate at which the Pillar Axiom system performs a backup operation. The system calculates the average by dividing the amount of data that was transferred (in GB) by the duration.

Restore Average (GB/hr)

Identifies the average rate at which the Pillar Axiom system performs a restore operation. The system calculates the average by dividing the amount of data that was transferred (in GB) by the duration.

Number of Files

Identifies the number of files that the Pillar Axiom system transferred during backup and restore operations on a filesystem.

Status

Displays the status of a backup or restore operation.

SEE ALSO

Display Performance Statistics

Ranges for Field Definitions

Performance Filesystems Page

DESCRIPTION

Use the Performance Filesystems page to review performance statistics for filesystems.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. Choose from:

- 1 min
- 5 min
- 10 min

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes.

Refresh Now

Updates the page with current data.

Name

Lists the names of configured filesystems. Click a name to review the filesystem settings.

Capacity

Identifies the maximum capacity limit that is assigned to the object.

Priority

Identifies the priority of the specified filesystem in relation to other filesystems.

IOPs

Identifies the current load of the filesystem on performance for input (read) and output (write) operations.

MB/s

Identifies the data transfer rate of the specified filesystem.

Collection Period

Identifies the time at which information was last collected from the Pillar Axiom system.

SEE ALSO

Display Performance Statistics

Ranges for Field Definitions

Performance NAS Protocols CIFS/NFS Page

DESCRIPTION

Use the Performance NAS Protocols CIFS/NFS page to review how many connections are in use by clients using Common Internet File System (CIFS) or Network File System (NFS) protocols.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. Choose from:

- 1 min
- 5 min
- 10 min

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes.

Refresh Now

Updates the page with current data.

File Server

Lists the names of configured File Servers. Click a name to review the CIFS and NFS statistics for that File Server.

Connections

Identifies how many connections are in use by NFS and CIFS clients.

Collection Period

Identifies the time at which information was last collected from the Pillar Axiom system.

SEE ALSO

Display Performance Statistics

Ranges for Field Definitions

Performance Overview Page

DESCRIPTION

Use the Performance Overview page to select any type of performance statistics that are available on the Pillar Axiom system. After you select a type, you can select and review those performance statistics.

FIELD DEFINITIONS

Performance Type

Lists types of performance statistics in Pillar Axiom systems. Click a type link to display performance statistics of that type.

- **Backup** opens the Performance Backup Page.
- **Filesystems** opens the Performance Filesystems Page.
- **LUNS** opens the LUN Performance Page.
- **NAS Protocols** opens the Performance NAS Protocols CIFS/NFS Page.
- **SAN Protocols** opens the Performance SAN Protocols Overview Page.

Performance Description

Briefly describes the types of performance statistics in Pillar Axiom systems.

SEE ALSO

Display Performance Statistics

Ranges for Field Definitions

Performance Profile Page

DESCRIPTION

Use the Performance Profile page to create and save a set of Quality of Service (QoS) settings for future use. You can import the profiles into the Configuration Wizard to create additional filesystems or LUNs based on these QoS settings.

FIELD DEFINITIONS

Performance Profile Name

Enter the name that you want to call the performance profile.

Filesystem/LUN Name

Identifies the name that is assigned to a filesystem or LUN.

Priority vs. Other Volumes

The Priority option determines how much of the system resources are devoted to that volume, including the allocation of Slammer CPU cycles and the allocation of specific portions of the disk platters. The higher the Priority, the greater the allocation of CPU time and the faster the media access time.

The Priority option specifies the layout of data that is stored in the Pillar Axiom system based on one of the following settings (see also the figure below):

- **Premium** priority is the highest possible performance. This priority uses the outer 20% of the drive platters.

  Note: I/O is typically faster when data resides on the outer edge of the drive platters.
- **High** priority logical volumes are allocated space in the outer 20%-40% of the drive platters.

  Tip: If sufficient Bricks are present, performance of high priority volumes is enhanced by placing those volumes on a greater number of Bricks.
- **Medium** priority logical volumes are allocated space in the outer 40% to 60% of the drive platters.
- **Low** priority logical volumes are allocated space in the band that is 60% to 80% from the outer diameter of drive platters.
- **Archive** (lowest-priority) logical volumes are allocated space in the inner areas of the drive platters.

Figure 22 Example of HDD priority bands

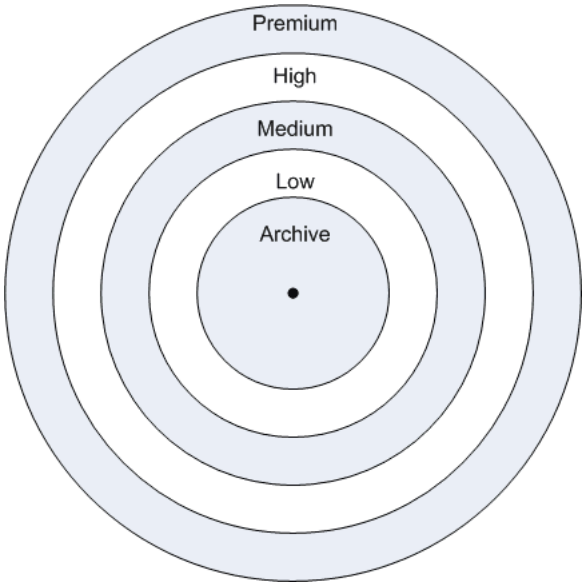

If needed, the Pillar Axiom system allocates space in a higher or lower priority band (within the same Storage Class) than the one that you choose.

The Pillar Axiom system uses the **Prioritization** and **Redundancy** values to calculate whether enough total capacity is available to create the new volume in the specified Storage Class. The calculated results are available in the table that is displayed when you click **Optimizer**.

Note: The capacity values displayed in the Optimizer table represent the sizes of the largest volume you can create in a particular Storage Class, given one of two performance configurations.

File Size

Identifies the typical size of files that are stored in the specified logical volume. Choose from:

- **Small**, if files are smaller than 20 KB.
- **Medium**, if files are larger than 20 KB and smaller than 4 MB.
- **Large**, if files are larger than 4 MB.

Files Access Bias

Identifies the typical data access method. Choose from:

Table 23 Effects of access and I/O bias

Access bias	I/O bias	Read-ahead in the Slammer	RAID configuration in the Brick
Sequential	Read	Aggressive	RAID 5

Table 23 Effects of access and I/O bias (continued)

Access bias	I/O bias	Read-ahead in the Slammer	RAID configuration in the Brick
			Writes data in write-back mode to physical storage in full-stripe extents.
	Write	Conservative	RAID 5 Writes data in write-back mode to physical storage in full-stripe extents.
Mixed and random	Read	None	RAID 5 Writes data in write-back mode to physical storage in full-stripe extents.
Random	Write	None	Distributed RAID

Note: This is an optimization bias, not a requirement that all data or data operations conform to the value.

I/O Bias

Identifies the typical read-write ratio. Choose from:

- **Read**, if users or applications read data more often than they write to the data source.
- **Write**, if users or applications write data more often than they read it.
- **Mixed**, if the read-write ratio varies.

Note: This is an optimization bias, not a requirement that all data or data operations conform to the value.

The system stores all writes of user data and system metadata in mirrored copies of the journal. Although the filesystem or LUN is journaled, for some of the profiles, write caching is disabled since it tends to have very low hit rates.

One copy is maintained in non-volatile memory on one control unit (CU) of a Slammer. The mirror copy is maintained in one of:

- Battery-backed memory of the partner CU on the Slammer (preferred location). Writes to this copy are equivalent to write-back cache.
- Virtual LUN (VLUN) that is reserved on disk for the logical volume if the partner CU is unavailable for the write. Writes to this copy are equivalent to write-through cache.

Writes from the journal to permanent disk storage are equivalent to write-through cache. The system flushes user data and the corresponding metadata as a unit to disk.

Redundancy

Identifies how many mirror copies of the original data are stored online.

Important! Pillar highly recommends that you consult with a Pillar customer support professional for assistance with sizing your system and creating your LUNs.

Redundancy options include the following:

- **Standard**, to store original data only. Data striping over multiple RAID-5 groups maintains full redundancy, even without mirror copies.

  Note: Standard does not maintain redundancy at the LUN level; however, it does provide sufficient data protection for most purposes.
- **Double**, to store original data and one mirror copy, with data striping over multiple RAID 5 groups.

Double redundancy can provide true redundancy only if your system has enough Bricks to allocate the filesystem or LUN such that no two mirror copies share a RAID group.

Note: SATA Bricks have two RAID groups for each Brick, FC Bricks have one RAID group for each Brick.

If the storage pool is becoming depleted, or a large filesystem or LUN is created, it might be necessary to place the filesystem or LUN on more RAID group fragments.

Depending on the system's ability to allocate sufficient contiguous storage blocks for the size of the LUN, use the following criteria to configure your LUNs.

- A small to medium sized filesystem or LUN will require two SATA Bricks or four FC Bricks to allow placing the four default extents on four separate RAID groups.
- A small to medium sized double redundant LUN or filesystem will require four SATA Bricks or eight FC Bricks to allow placing the two mirror copies on eight separate RAID groups.

For *performance testing purposes only*, create a filesystem or LUN using standard redundancy and the HighThroughput Profile. This is *not* recommended for most applications. Reset your system after you have created a HighThroughput LUN or filesystem before you configure normal filesystems or LUNs for applications.

Export a copy to your management client

Select this option if you want to save a copy of the performance profile to your local client.

SEE ALSO

Display Performance Statistics

Run the Configuration Wizard

Ranges for Field Definitions

Performance SAN Protocols Overview Page

DESCRIPTION

Use the Performance SAN Protocols Overview page to review performance statistics for storage area network (SAN) protocols.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. Choose from:

- 1 min
- 5 min
- 10 min

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes.

Refresh Now

Updates the page with current data.

Slammer

Identifies the name of the Slammer that contains TCP/IP statistics.

Control Unit

Identifies the control unit (CU) of the Slammer that contains the statistics.

Network Interface

Identifies the physical port on the control unit (CU).

Type

Identifies the Slammer port connection type, Fibre Channel (FC) or Internet SCSI (Small Computer System Interface), or iSCSI.

Send (MB/s)

Displays the statistics for output (Send) data.

Receive (MB/s)

Displays the statistics for input (Receive) data.

Collection Period

Identifies the time at which information was last collected from the Pillar Axiom system.

SEE ALSO

Display Performance Statistics

Ranges for Field Definitions

Pilot Overview Page

DESCRIPTION

Use the Pilot Overview page to review the status of the Pilot management controller that is installed on the Pillar Axiom storage system.

FIELD DEFINITIONS

Identify *(Pilot control units)*

Blinks the light-emitting diodes (LEDs) on the specified hardware component so that you can visually identify the component from a group of like hardware components.

Unit Serial Number

Identifies the serial number that is assigned to the hardware component.

Mode

Displays the current operational mode of the two control units (CUs) within the Pilot.

- **Active** indicates which CU performs all configuration tasks that administrators request.
- **Standby** indicates which CU acts as a secondary device and does nothing unless the active CU fails over to this standby control unit.

Status

Displays the current status of a control unit (CU) within the Pilot. A status of Normal requires no action.

SEE ALSO

Manage Hardware Components

Identify Hardware Components

Ranges for Field Definitions

Remote Replication Overview Page

DESCRIPTION

Use this page to set the management IP address of the web server where the remote data management application is installed. This allows you to use the separate file or volume replication management interface to manage remote data replication.

FIELD DEFINITIONS

File Replication

Allows you to enter the management IP address for file-based remote data replication.

Volume Replication

Allows you to enter the management IP address for volume-based remote data replication.

Replication Command Line Interface

Allows you to download the Pillar Axiom Replication utility bundle for NAS and SAN systems.

SEE ALSO

Ranges for Field Definitions

Remote Replication Settings Page

DESCRIPTION

Use the Remote Replication Settings page to define the management IP address to use file and volume replication applications.

FIELD DEFINITIONS

Management Interface IP Address

Identifies the IP address of the web server where the remote data management application is installed.

SEE ALSO

Ranges for Field Definitions

Replication Command Line Utility Page

DESCRIPTION

Use the Replication Command Line Utility page to download the Pillar Axiom replication utility package. This package is supported on the following platforms:

- Linux (NAS and SAN systems)
- Solaris (NAS systems only)
- Windows (NAS and SAN systems)

Note: The download package bundles the Pillar Axiom MaxRep Replication for NAS and the Pillar Axiom MaxRep Replication for SAN utilities. After downloading the package, you can install either utility, or both, as needed.

FIELD DEFINITIONS

<Actions>

Displays the operating systems on which the command line utilities run.

Select the download package that is appropriate for your client platform and save the package to a temporary location on that client.

SEE ALSO

Pillar Axiom MaxRep Replication for NAS User's Guide and Reference

Pillar Axiom Replication User's Guide and Reference for SAN

Replication Overview Page

DESCRIPTION

Use the Replication Overview page to select any type of data replicas that can be created on a Pillar Axiom storage system. After you select a type, you can select and manage a Snap FS, Clone FS, or a Clone LUN.

FIELD DEFINITIONS

Replica Type

Lists types of data replicas in Pillar Axiom systems. Click a type link to display details about data replicas or their source logical volumes (filesystems orLUNs).

- **Snap FS** opens the Snap FS Overview Page
- **Clone FS** opens the Clone FS Overview Page.
- **Clone LUN** opens the Clone LUN Overview Page.

Replica Description

Briefly describes the types of data replicas in Pillar Axiom systems.

SEE ALSO

For Snap FSs:

Create an Immediate Snap FS
Display Data Replica Details
Create Snap FS Schedules
Delete Snap FS Schedules

For Clone FSs:

Activate a Clone
Create an Immediate Clone FS
Display Data Replica Details

For Clone LUNs:

Activate a Clone
Create an Immediate Clone LUN
Display Data Replica Details

Ranges for Field Definitions

Replication Schedule, Run Daily Page

DESCRIPTION

Use the Replication Schedule, Run Daily page to create a replication schedule that runs every day.

FIELD DEFINITIONS

Start Time

Identifies the time at which the Pillar Axiom system starts a scheduled operation.

Start Date

Identifies the date on which the Pillar Axiom system performs a scheduled operation.

Recurrence *(if* Schedule Intervals *is not* ***Run Once****)*

Identifies how often the Pillar Axiom system should perform the scheduled operation. Valid values vary based on the schedule's recurrence interval or frequency. Choose from:

- **Hourly**: 1 through 24, inclusive
- **Daily**: 1 through 7, inclusive
- **Weekly**: 1 through 4, inclusive

Limit Number of Snapshots To

Identifies the maximum number of scheduled snapshots to create. Valid values are based on the schedule's recurrence interval or frequency.

SEE ALSO

Create Snap FS Schedules

Delete Snap FS Schedules

Ranges for Field Definitions

Replication Schedule, Run Hourly Page

DESCRIPTION

Use the Replication Schedule, Run Hourly page to create a replication schedule that runs every hour.

FIELD DEFINITIONS

Start Time

Identifies the time at which the Pillar Axiom system starts a scheduled operation.

Start Date

Identifies the date on which the Pillar Axiom system performs a scheduled operation.

Recurrence *(if* Schedule Intervals *is not* ***Run Once****)*

Identifies how often the Pillar Axiom system should perform the scheduled operation. Valid values vary based on the schedule's recurrence interval or frequency. Choose from:

- **Hourly**: 1 through 24, inclusive
- **Daily**: 1 through 7, inclusive
- **Weekly**: 1 through 4, inclusive

Limit Number of Snapshots To

Identifies the maximum number of scheduled snapshots to create. Valid values are based on the schedule's recurrence interval or frequency.

SEE ALSO

Create Snap FS Schedules

Delete Snap FS Schedules

Ranges for Field Definitions

Replication Schedule, Run Weekly Page

DESCRIPTION

Use the Replication Schedule, Run Weekly page to create a replication schedule that runs every week.

FIELD DEFINITIONS

Start Time

Identifies the time at which the Pillar Axiom system starts a scheduled operation.

Start Date

Identifies the date on which the Pillar Axiom system performs a scheduled operation.

Recurrence *(if* Schedule Intervals *is not* ***Run Once****)*

Identifies how often the Pillar Axiom system should perform the scheduled operation. Valid values vary based on the schedule's recurrence interval or frequency. Choose from:

- **Hourly**: 1 through 24, inclusive
- **Daily**: 1 through 7, inclusive
- **Weekly**: 1 through 4, inclusive

Limit Number of Snapshots To

Identifies the maximum number of scheduled snapshots to create. Valid values are based on the schedule's recurrence interval or frequency.

SEE ALSO

Create Snap FS Schedules

Delete Snap FS Schedules

Ranges for Field Definitions

Replication Schedule Summary Page

DESCRIPTION

Use the Replication Schedule Summary page to schedule snapshots to create a point-in-time copy of user data.

FIELD DEFINITIONS

Replication Schedules

Lists types of replication schedules in Pillar Axiom systems. Click a type link to display replication schedules of that type.

Click **Scheduled Snapshots** to schedule the automatic creation of Snap FSs. Create replication schedules on:

- Schedule Configuration Page, Details Tab
- Schedule Configuration Page, Schedule Tab

SEE ALSO

- Create Snap FS Schedules
- Delete Snap FS Schedules
- Ranges for Field Definitions

Resolve Connectivity Trouble Page

DESCRIPTION

Use the Resolve Connectivity Trouble page to help resolve a connectivity issue between the specified Slammer and the customer network.

FIELD DEFINITIONS

Slammer

Identifies the Slammer for which you want to resolve connectivity issues.

Control Unit

Identifies a specific control unit (CU) in a Slammer and the physical network port in the network interface module on that Slammer CU.

Command Line

Identifies a command to perform in order to resolve a connectivity issue between the specified Pillar Axiom Slammer and the customer network. Enter the command and command-line arguments:

- *required arguments*
- *[optional arguments]*

A File Server configuration contains virtual interfaces (VIFs). The name of a File Server is internally mapped to a socket identifier (SOCK) and a virtual server identifier (VSID). The environment variables SOCK and VSID provide values to various Slammer commands which then use those values to identify the target File Server.

The syntax of the SOCK and VSID environment variables is defined as follows:

SOCK=/vserver*N*	Provides a prefix that the system uses to address an internal structure associated with a particular File Server. When this variable is not set, the system addresses a more generic structure within the internal network that interconnects the Pilot and the Slammers (this network is called the *private management interface*, or *PMI*). Most commonly, this variable is set to a particular value in order to run a command against the data interface.
VSID=*virtualServerID*	Identifies by number (*virtualServerID*) a particular File Server.

Note: For a given File Server, the value for *N* and the value for *virtualServerID* are the same value.

Tip: If you have one (and have had only one) File Server, the value of SOCK would be **`/vserver0`** and the value for VSID would be **`0`**. Otherwise, you can execute `ifconfig` repetitively using various settings for SOCK and checking each returned IP address against the address configured for the control unit to determine the correct value for SOCK. Although the maximum number of supported File Servers is eight at any given time, the maximum value for SOCK is **`/vserver255`** and that for VSID is **`255`**.

Table 24 Slammer commands

Command	Syntax	Description
`ping`	`ping -c` *count ipaddr* **`-c`** ***count*** Replace `count` with an unsigned number, which is the number of packets to send. ***ipaddr*** IP address of a host machine. Environment variables (optional): `SOCK`	Sends an echo request to a specified host. Default sends the echo request from the private management interface (PMI), unless the SOCK environment variable is set. If used, this variable must be set to the virtual interface (VIF) from where you want the echo request to be sent. Example: The following command requests an echo from the host whose IP is 198.168.1.2 by sending five packets to that address. **`ping -c 5 198.168.1.2`**
`route`	`route show` Environment variables (optional): `SOCK`	Displays the route table. Default displays the route table of the PMI, unless the SOCK environment variable is set. If used, this variable must be set to the virtual interface that contains the route table that you want to display. The Pilot and the Slammer communicate to one another over the PMI. Example: The following command displays the route table associated with virtual server 0.

Table 24 Slammer commands (continued)

Command	Syntax	Description
		`route show SOCK=/vserver0`
`traceroute`	`traceroute` [`-s` *ipaddr*] [`-r`] [`-v`] *host* **`-s`** ***ipaddr*** Source address (*ipaddr*) of the `traceroute` probe packet. **`-r`** Bypass the normal routing tables and send directly to the attached host. **`-v`** Verbose mode. ***host*** The target host. Environment variables (optional): SOCK	Traces the route that an IP packet takes to reach a specified host. Default sends the probe packet from the PMI interface, unless the SOCK environment variable is set. If used, this variable must be set to the virtual interface from where you want the `traceroute` probe to come. Example: The following command displays all route information between the PMI and the host whose IP address is 10.20.5.29. **`traceroute -v 10.20.5.29`**
`arp`	`arp -a` **`-a`** Display current ARP entries. Environment variables (optional): SOCK	Displays the IP-to-physical address translation tables used by the Address Resolution Protocol (ARP). Default displays the ARP table of the PMI interface, unless the SOCK environment variable is set. Example: The following command displays the ARP table located in the Pilot. **`arp -a`**
`ifconfig`	`ifconfig` Environment variables (optional): SOCK	Displays interfaces and related information. Default displays interface information for the PMI, unless the SOCK environment variable is set. Example: **`ifconfig`**

Table 24 Slammer commands (continued)

Command	Syntax	Description
`wbinfo`	`wbinfo` [`-t`] [`-m`] [`-I` *ipaddr*] [`-N` *name*] **`-t`** Verifies that the workstation trust account that was created when the CIFS server was added to the Windows domain is working and that the shared secret is good. **`-m`** Lists trusted domains. **`-I`*ipaddr*** Converts an IP address to a NetBIOS name (WINS). **`-N` *name*** Converts a NetBIOS name to an IP (WINS). Environment variables (required): `SOCK`, `VSID`	Queries information from the Common Internet File System (CIFS) `windbind` daemon. By default, this command displays the help information. **Note:** Only the parameters listed in the Syntax column are supported for this command. Example: The following command requests the CIFS server to confirm that the status of the joined domain is Yes. **`wbinfo -t`** **`VSID=0 SOCK=/vserver0`**
`netstat`	`netstat` [`-s`] [`-r`] [`-pip`] [`-pudp`] [`-ptcp`] **`-s`** Request statistics instead of connections. **`-r`** Displays the routing table. **`-pip`** Displays IP level statistics. Good for finding issues with fragments for NFS over UDP. **`-pudp`** Displays UDP level statistics. **`-ptcp`** Displays TCP level statistics. Environment variables (required): `SOCK`	Shows networking statistics for a NAS Slammer control unit. If no options are specified, `netstat` returns information about the listening sockets and established connections. Example: The following command requests the IP packet statistics that are associated with virtual server 0. **`netstat -pip`** **`SOCK=/vserver0`**
`nmblookup`	`nmblookup` [`-U` *server*] [`-R`] *hostName* **`-U`*server*** Specifies the IP address or name of the WINS server. Pillar recommends using IP address to be consistent with the value	Looks up NetBIOS names on a WINS server and maps them to IP addresses. This command requires that both the SOCK and VSID environment variables be set. Example 1:

Table 24 Slammer commands (continued)

Command	Syntax	Description
	configured in the File Server definition. **-R** Causes the WINS server to return names stored on the server. *hostName* Name of a NetBIOS host. Environment variables (required): SOCK, VSID	The following command requests the WINS server at 172.10.1.20 to resolve the cifs01 CIFS server to an IP address. **nmblookup -U 172.10.1.20 -R cifs01 VSID=0 SOCK=/vserver0** Example 2: The following command requests the WINS server at 172.10.1.20 to provide the IP addresses of the domain controllers serving the domain01 NTLM domain. **nmblookup -U 172.10.1.20 -R domain01#1c VSID=0 SOCK=/vserver0** **Note:** The two hexadecimal digits (**1c**) following the # character specify the NetBIOS name type as *Domain Controllers*, which conforms to the NetBIOS suffix convention.
nslookup	nslookup *hostName* [*serverName*] *hostName* Display information about the specified host. *serverName* Specifies the name of the DNS server. Environment variables (optional): SOCK	Queries domain name server (DNS). By default, sends the domain name request over the PMI interface, unless the SOCK environment variable is set. If used, this variable must be set to the virtual server from where you want the query to be sent. Example: The following command looks up the IP address of ca-lab.eng using the DNS at 10.20.0.10. In this example, the request is sent from the NAS Slammer port that is associated with virtual server 0.

Table 24 Slammer commands (continued)

Command	Syntax	Description
		`nslookup ca-lab.eng 10.20.0.10`
`perf`	`perf [-c] [-l numThreads]` **`-c`** Reset the counters after returning the statistics. **`-l`*numThreads*** Displays the statistics for the specified number of threads that consume the most CPU cycles. Environment variables: None	Checks CPU utilization of the specified Slammer. This command returns the CPU statistics for the following categories: • The idle process • The kernel time • All process IDs that are running When you use the `-c` option, the next invocation of `perf` will return data accumulated from that point. Example: The following command displays the utilization statistics for all threads running on the selected Slammer and then clears the counters for that Slammer. **`perf -c`**

Environment Variables

Identifies the space-delimited pairs of environment variables and values to use while executing the command.

Execute

Performs the specified command.

Command Output

Displays the results of a command that was run to resolve a connectivity issue.

SEE ALSO

The set of tools listed in About Pillar Axiom Support Tools

Ranges for Field Definitions

Resolve System Trouble Page

DESCRIPTION

Use the Resolve System Trouble page to run the resolution tools that are available on a Pillar Axiom storage system.

FIELD DEFINITIONS

Resolution Tool Type

Lists types of problem-resolution tools that are provided to support Pillar Axiom systems. Click a type link to run the selected resolution tool.

For details about how to resolve different types of system problems, see:

- Clear System Configuration
- Reset System Serial Number

Resolution Tool Description

Briefly describes the tools that you can use to resolve issues on the Pillar Axiom system.

SEE ALSO

The support tools listed in About Pillar Axiom Support Tools

The actions listed in About Responding to Administrator Actions

Ranges for Field Definitions

Routes Page

DESCRIPTION

Use the Routes page to create and delete additional network routes for a File Server. Create the primary route on the File Server Page, Network Tab.

You can define up to eight default network routes. Only one of these default routes is active at any given time. The system configures the active route by choosing the first default route that is usable over one of the active VIFs (virtual interfaces) for that File Server. If the VIF that is used by the active default route fails or is removed, the next default route that matches an active VIF is activated.

You can define up to 16 static network routes for a File Server.

FIELD DEFINITIONS

Default Route

Identifies the network route as a default route. This option allows input for the gateway but disables the IP address and Netmask fields.

Destination

Identifies a destination subnet for this network route. All computers in the subnet are reached through the specified gateway.

Netmask

Identifies a subnet mask for the range of IP addresses at the route's destination.

Gateway

Identifies the IP address that is assigned to the gateway host. The gateway IP address is used to route messages from this network to other networks.

Create

Adds the new configuration to the selected object.

Routes

Displays a list of currently configured network routes. Check the box to the left of a route to select it for deletion.

Remove

Deletes the selected objects.

SEE ALSO

Create a File Server

Modify File Server Attributes

Ranges for Field Definitions

SAN Hosts Overview Page

DESCRIPTION

Use the SAN Hosts Overview page to review the storage area network (SAN) hosts defined on a Pillar Axiom storage system. The **Actions** drop-down list provides options to modify and view host settings, as well as to delete host names.

FIELD DEFINITIONS

Host Name

Identifies the names of the SAN hosts that are available the Pillar Axiom system.

Host Port

Identifies the initiator, the World Wide Name (WWN) of a Fibre Channel (FC) host bus adapter (HBA) or the iSCSI IP address on the SAN hosts.

Type

Identifies the type of host interface, FC or iSCSI.

Pillar Axiom Path Manager

Identifies whether or not the Pillar Axiom Path Manager driver is communicating, or, if it is not registered. If it is not registered, you need to install and run the driver.

Number of LUNs

Identifies the number of LUNs that are mapped to that particular SAN host either because of specific mapping or because the LUN is available to all SAN hosts.

Host Port Status

Identifies the status, connected or not connected, of the HBA.

SEE ALSO

Display SAN Host Settings

Modify SAN Host Settings

Delete SAN Host Names

Associate Hosts

Define LUN Mapping (Optional)

Ranges for Field Definitions

SAN LUNs Overview Page

DESCRIPTION

Use the LUNs Overview page to review the SAN LUNs that have been defined on a Pillar Axiom storage system. The **Actions** drop-down list provides options to create, copy, modify, delete, and view those LUNs, as well as to create an immediate Clone LUN.

FIELD DEFINITIONS

Name

Identifies the name that is assigned to a LUN for administrative purposes. LUN names must be unique across the Pillar Axiom system, not just within its associated volume group.

Total Capacity (GB)

Identifies the capacity that is assigned to this LUN.

Growth Max (GB)

Identifies the maximum capacity limit (in GB) in which the object can expand.

Host Access

Identifies the SAN hosts that can access the LUN.

LUN Number

Identifies the unique number that is assigned to a LUN and can be accessed by all SAN hosts if the LUN is not mapped. If the LUN is mapped, the LUN number must be unique to the mapped SAN host.

Status

Identifies the current status of each LUN.

Storage Class

Identifies the category of physical storage on which the logical volume resides:

- **FC** (Fibre Channel drives)
- **SATA** (Serial ATA drives)
- **SLC SSD** (single level cell, solid state drives)

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

Priority

Identifies the layout of the data for the logical volume:

- **Premium** priority is the highest possible performance. This priority uses the outer 20% of the drive platters.

  Note: I/O is typically faster when data resides on the outer edge of the drive platters.

- **High** priority logical volumes are allocated space in the outer 20%-40% of the drive platters.

  Tip: If sufficient Bricks are present, performance of high priority volumes is enhanced by placing those volumes on a greater number of Bricks.
- **Medium** priority logical volumes are allocated space in the outer 40% to 60% of the drive platters.
- **Low** priority logical volumes are allocated space in the band that is 60% to 80% from the outer diameter of drive platters.
- **Archive** (lowest-priority) logical volumes are allocated space in the inner areas of the drive platters.

Redundancy

Identifies the redundancy level of a LUN as standard or double.

Pinned Data

Identifies whether any user data is pinned in cache and cannot be written to permanent storage.

LUID

Identifies which LUNs have been configured on the Pillar Axiom storage system. Enter the logical unit unique identifiers (LUIDs) of the configured LUNs.

Space for Clone LUNs (GB)

Specifies the amount of space to make available on the Pillar Axiom system for Clone LUNs.

IOps

Identifies the current performance for input (read) and output (write) operations for a LUN.

Bandwidth

Identifies the data transfer rate of the specified LUN.

SEE ALSO

Create LUNs

Display LUN Details

Modify a LUN

Delete LUNs

Copy LUNs

Create an Immediate Clone LUN

Ranges for Field Definitions

SAN Slammer Ports Page

DESCRIPTION

Use the SAN Slammer Ports page to review the topology of the network ports on each of the SAN Slammer control units (CUs).

FIELD DEFINITIONS

Slammer

Identifies the name of the SAN Slammer.

Control Unit

Identifies a SAN Slammer CU.

Network Interface

Identifies the physical port on the CU.

Type

Identifies the type of host interface, Fibre Channel (FC) or Internet SCSI (Small Computer System Interface) (iSCSI).

WWN (FC) or MAC Address (iSCSI)

Identifies the unique identifiers of the host bus adapter (HBA) ports that the Pillar Axiom system detects on the network. For FC networks, this identifier is the World Wide Name (WWN). For iSCSI networks, this identifier is the Media Access Control (MAC) address.

IP Address (iSCSI)

Identifies the IP address of the iSCSI SAN Slammer ports that the Pillar Axiom system detects on the network.

iSNS Server Status (iSCSI)

Displays the status of the SAN Slammer ports that support the Internet Storage Name Service (iSNS) protocol:

- **Connected** (able to communicate with the iSNS server)
- **Connection Lost** (no longer able to communicate with the iSNS server)
- **Not Connected** (not able to communicate with the iSNS server)

When no status shows, the port is not an iSCSI port.

Target Portal Group Tag (iSCSI)

Identifies the group of iSCSI target ports through which connections for a single session can be made. This allows an iSCSI target to designate multiple ports that can have connections within a single session.

Topology

Identifies the Fibre Channel (FC) transport topology in use by the ports in the network interface module (NIM) to connect to the storage area network (SAN) employed by the customer:

- **Fabric**, which means that the port is an N_Port in a switched fabric (FC-SW).
- **Loop**, which means that the port is an NL_Port in an arbitrated loop (FC-AL).
- **Point-to-Point**, which means that the port is an N_Port that is connected to another N_Port, back to back (FC-P2P).
- **Public Loop**, which means that the port is an NL_Port that is connected to a loop in which one port in the loop is an FL_Port in the fabric (FC-FLA).

Note: The topology used by Storage System Fabric (SSF) between the Slammer PIMs and the Brick RAID controllers is private and therefore not reported.

Negotiated Link Speed

Displays the transmission speed in gigabits per second for the port.

Medium Type

Identifies the types of network ports for data path traffic between the customer network switches and the Pillar Axiom SAN Slammers:

- **Copper**, which identifies RJ-45 copper interfaces.
- **Optical-L**, which identifies longwave optical small form-factor pluggable (SFP) transceiver interfaces.
- **Optical-S**, which identifies shortwave optical SFP transceiver interfaces.

SEE ALSO

Create LUNs

Ranges for Field Definitions

SAN Storage Overview Page

DESCRIPTION

Use the SAN Storage Overview page to review the logical units (LUNs) and hosts that are configured on the Pillar Axiom storage system. The **Actions** drop-down list provides options to create, modify, delete, and view LUNs and hosts.

FIELD DEFINITIONS

LUNs

A logical volume within a storage area network (SAN). Administrators assign storage resources and Quality of Service (QoS) attributes to each LUN.

Hosts

Clients that expose the block-based storage provided by LUNs as filesystems and shares.

Slammer Ports

A topology overview of the network ports on the Slammer control units.

SEE ALSO

Create LUNs

Display LUN Details

Modify a LUN

Copy LUNs

Delete LUNs

Schedule Configuration Page, Details Tab

DESCRIPTION

Use the Schedule Configuration page, **Details** tab to create and modify a data replication schedule.

FIELD DEFINITIONS

Schedule Name

Identifies the name of a scheduled operation, which is an action to be performed at the specified time or at regular intervals.

Filesystem Name

Identifies the name of a filesystem from which a scheduled Snap FS will be created.

Schedule Intervals

Identifies the intervals at which the Pillar Axiom system performs a scheduled operation. Choose from:

- **Run Once**
- **Hourly**
- **Daily**
- **Weekly**

Enable Schedule

Identifies whether the schedule is enabled.

- Enable the schedule so that the operation is performed at the scheduled time.
- Disable the schedule so that operations are not performed. This permits you to define a schedule before the source logical volume (filesystem or LUN) has been made available to users.

SEE ALSO

Create Snap FS Schedules

Delete Snap FS Schedules

Ranges for Field Definitions

Schedule Configuration Page, Schedule Tab

DESCRIPTION

Use the Schedule Configuration page, **Schedule** tab to create and modify a data replication schedule.

FIELD DEFINITIONS

Start Time

Identifies the time at which the Pillar Axiom system starts a scheduled operation.

Start Date

Identifies the date on which the Pillar Axiom system performs a scheduled operation.

Recurrence *(if **Schedule Intervals** is not **Run Once**)*

Identifies how often the Pillar Axiom system should perform the scheduled operation. Valid values vary based on the schedule's recurrence interval or frequency. Choose from a value listed in the following table.

Table 25 Schedule recurrence values

Recurrence interval	Valid values
Hourly	1 through 24, inclusive
Daily	1 through 7, inclusive
Weekly	1 through 4, inclusive

Archive the most recent {n} snapshots *(if **Schedule Intervals** is not **Run Once**)*

Identifies the maximum number of scheduled snapshots to create. Valid values are based on the recurrence interval or frequency of the schedule.

SEE ALSO

Create Snap FS Schedules

Delete Snap FS Schedules

Ranges for Field Definitions

Scheduled Operations Page

DESCRIPTION

Use the Scheduled Operations page to review and modify operations that are scheduled on the Pillar Axiom system.

FIELD DEFINITIONS

Start Time

Identifies the time at which the Pillar Axiom system starts a scheduled operation.

Operation

Identifies the name of a scheduled operation.

Filesystem

Identifies the name of the filesystem on which the scheduled operation will be performed.

SEE ALSO

Modify a Filesystem Quota

Create Snap FS Schedules

Update Pillar Axiom Software

Verify Storage Redundancy

Verify Data Consistency

Check Filesystem Consistency

Ranges for Field Definitions

Scheduled Updates Page

DESCRIPTION

Use the Scheduled Updates page to review the scheduled software and firmware updates on the Pillar Axiom storage system. The **Actions** drop-down list provides options to modify and delete the schedules.

FIELD DEFINITIONS

Scheduled Update

Identifies the name of an update package file. An update package contains one or more firmware or software modules. If more than one module is in a package, you can install them all together or separately.

Start Time

Identifies the time at which the Pillar Axiom system starts a scheduled operation.

System Restart Required

Identifies whether a shutdown and restart is required after a firmware or software module is updated or a hardware component is installed.

SEE ALSO

Update Pillar Axiom Software

Display Software Versions

Modify a Software Update Schedule

Ranges for Field Definitions

Select Update Page

DESCRIPTION

Use the Select Update page to review available software updates for the Pillar Axiom storage system.

FIELD DEFINITIONS

Package

Identifies the name of an update package file. An update package contains one or more firmware or software modules. If more than one module is in a package, you can install them all at once or separately.

Remove

Deletes the selected objects.

Add

Adds the new configuration to the selected object.

SEE ALSO

Download Software Updates

Update Pillar Axiom Software

Ranges for Field Definitions

Shutdown/Restart Page

DESCRIPTION

Use the Shutdown page to stop software processes and shut down the Pillar Axiom hardware components in an orderly fashion. Users cannot access data while the system is down.

Tip: If you need to power off the system for more than 48 hours, remove the batteries.

FIELD DEFINITIONS

Actions List

Identifies when to gracefully shut down the Slammers and Bricks in the Pillar Axiom system. Choose from:

- **Shutdown Now**
- **Shutdown in 5 minutes**
- **Shutdown in 10 minutes**
- **Restart Now**

Restart Now is the only option that restarts the system automatically after the shutdown is complete. If you choose any of the other shutdown options, manually restart the system.

SEE ALSO

Shut Down the Pillar Axiom System

Ranges for Field Definitions

Slammers Overview Page

DESCRIPTION

Use the Slammers Overview page to review the Slammers that are part of the Pillar Axiom system. The **Actions** drop-down list provides different options based on the context from which you open the page. When you click the Slammers link in the navigation pane from the:

- **Health** context, you can view the status of and change the name that is assigned to a Slammer.
- **Support** context, you can repair Slammer hardware.

FIELD DEFINITIONS

Slammer Name

Lists the names of hardware components. Click a name to display details about that hardware component.

Type

Lists the type of Slammer, NAS or SAN.

Control Unit

Identifies a control unit (CU) in a hardware component.

Status

Displays the current status of a hardware component. A status of Normal requires no action.

I/O Ports

Displays the current status of a hardware component's input and output ports.

Power Supplies

Displays the current status of the power supplies of a hardware component.

Fans

Displays the current status of the fans within a hardware component.

Batteries

Displays the current status of the batteries within a hardware component.

Temperature

Displays the current temperature within a hardware component.

Important! Watch for temperatures that are too high or too low.

Slammer Diagnostics

Runs diagnostics on a Slammer control unit (CU). During the diagnostics, the CU is taken off line for several minutes.

If you run diagnostics on a SAN Slammer, you need to disconnect the CU from the public network and attach a loop back connector between the ports on the Slammer CU.

SEE ALSO

Manage Hardware Components

- Display Hardware Component Information
- Modify Hardware Component Names
- Identify Hardware Components
- Replace a Hardware Component

Ranges for Field Definitions

SnapDelta FS Download Page

DESCRIPTION

The Pillar Axiom SnapDelta FS utility 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 FSs.

External applications sometimes need to scan the contents of a Pillar Axiom 600 filesystem to extract information. After an initial scan, these applications need to perform periodic re-scans 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

FIELD DEFINITIONS

SnapDelta FS

Select the option in the **Actions** drop-down list to download the SnapDelta FS utility for use on a Linux platform.

SEE ALSO

Pillar Axiom SnapDelta FS Reference Guide

Snap FS Overview Page

DESCRIPTION

Use the Snap FS Overview page to manage Snap FSs of NAS filesystems. The **Actions** drop-down list provides different options based on the context from which you open the page. When you click the Snap FS link in the navigation pane from the:

- **Storage** context, you can restore a filesystem from a Snap FS and delete Snap FSs when they are no longer needed.
- **Data Protection** context, you can create an immediate Snap FS of the selected filesystem.

Note: You cannot restore a Pillar Axiom SecureWORMfs filesystem from a Snap FS.

FIELD DEFINITIONS (STORAGE CONTEXT)

Source Filesystem

Identifies the name of the filesystem from which a Snap FS was created.

Snapshot Name

Identifies the name of a data replica.

Created

Identifies the date and time at which the data replica was created.

Size (GB)

Identifies the size of the data replica, in gigabytes.

FIELD DEFINITIONS (DATA PROTECTION CONTEXT)

Filesystem Name

Lists the names of configured filesystems. Click a name to review the filesystem settings.

Total Capacity

Identifies the current capacity that is assigned to the specified filesystem.

Free Capacity

Identifies the current unused capacity that is available. Free capacity is total capacity minus allocated capacity.

Used Capacity

Identifies the current capacity usage of the object.

Status

Displays the current status of a filesystem. The possible states are:

- **Online**. The filesystem is online and normal.
- **Offline**. The filesystem is offline.

- **Partial Offline**. The actual redundancy level of the filesystem may be different from the redundancy level with which the filesystem was configured.
- **Conservative**. Write-back cache on the filesystem has been disabled so journaling has slowed.
- **Degraded**. All of the copies of a redundant filesystem are not available. If one copy is missing, it is not fully redundant. This can happen when a write fails (which may be a 30 second time-out) to one copy of the array.

File Server

Identifies the name of a File Server with which the filesystem is associated.

Number of Snapshots

Identifies the number of Snap FSs that have been created from the specified filesystem.

Snapshots Size (GB)

Identifies the size of the data replica, in gigabytes.

SEE ALSO

Create an Immediate Snap FS

Create Snap FS Schedules

Ranges for Field Definitions

Snap FS Schedule Summary Page

DESCRIPTION

Use the Snap FS Schedule Summary page to review completed filesystem snapshots. The **Actions** drop-down list provides options to create, modify, delete, and view Snap FS schedules, as well as to create an immediate Snap FS.

FIELD DEFINITIONS

Snap FS Schedule Name

Identifies the name of a schedule. Click a name to review or modify the schedule settings.

Start Time

Identifies the time and date on which the Pillar Axiom system started a schedule recurrence.

Frequency

Identifies the interval at which the Pillar Axiom system starts a recurrent schedule.

Filesystem

Identifies the name of the filesystem from which a Snap FS was created.

Status

Displays the current status of a schedule.

SEE ALSO

Display Data Replica Details

Create Snap FS Schedules

Delete Snap FS Schedules

Create an Immediate Snap FS

Ranges for Field Definitions

SNMP Configuration Page

DESCRIPTION

Use the SNMP Configuration page to create, modify, and delete Simple Network Management Protocol (SNMP) trap hosts.

FIELD DEFINITIONS

SNMP Hosts

Displays a list of trap hosts from which you can select one or more to be the source of the operation.

Modify button

Changes the selected object's configuration settings.

Remove button

Deletes the selected objects.

Authorized Host IP

Identifies the IP address of a workstation with an SNMP-based management application installed. This workstation receives the SNMP traps that the Pillar Axiom system generates.

Community String

Identifies a community for which a specific host should receive traps that the Pillar Axiom system generates. You can specify different community strings for each UPS device that the Pillar Axiom system monitors so that multiple administrators can receive specific types of SNMP traps.

Note: When an SNMP administrator does not specify a community string for read-only access, the typical value used by SNMP servers and clients is `public`.

Port

Identifies the port on which an SNMP host listens for the SNMP traps that the Pillar Axiom system generates.

Add Host button

Adds the specified trap host to the SNMP configuration.

SEE ALSO

Create SNMP Trap Hosts

Modify SNMP Trap Hosts

Delete SNMP Trap Hosts

Ranges for Field Definitions

SNMP Settings Page

DESCRIPTION

Use the SNMP Settings page to determine whether the Simple Network Management Protocol (SNMP) software feature is enabled and to review and delete the SNMP hosts when the feature is enabled. The **Actions** drop-down list provides options to modify the SNMP settings and to delete SNMP hosts.

FIELD DEFINITIONS

Authorized Host IP

Identifies the IP address of a workstation that receives the Pillar Axiom SNMP information.

Community String

Identifies the community string that is defined for the SNMP host that the Pillar Axiom system generates.

Port

Identifies the port on which the SNMP host listens.

SEE ALSO

Create SNMP Trap Hosts

Modify SNMP Trap Hosts

Delete SNMP Trap Hosts

Ranges for Field Definitions

Software Configuration Page

DESCRIPTION

Use the Software Configuration page to select any of the types of software configuration operations that are available on a Pillar Axiom storage system. After you select a type, you can perform the specified software configuration operation.

FIELD DEFINITIONS

Software Configuration Type

Lists types of software configuration operations that are available on Pillar Axiom systems. Click a type link to review or configure that type of operation.

- **Software Modules** opens the Software Modules Page.
- **Scheduled Updates** opens the Scheduled Updates Page.

SEE ALSO

Download Software Updates

Update Pillar Axiom Software

Display Software Versions

Modify a Software Update Schedule

Ranges for Field Definitions

Software Modules, Details Tab

DESCRIPTION

Use the Software Modules **Details** tab to install new versions of firmware and software on the Pillar Axiom storage system or to schedule an update.

FIELD DEFINITIONS

Last Software Update Staged

Identifies by file name the most recent update package that is available for installation. An update package contains one or more firmware or software modules. If more than one module is in a package, you can install them all together or separately.

Upload Package

Permits you to navigate to and select a software update package so that you can:

- Copy the package from its distribution media to the staged packages that are available for installation.
- Install the package components on the Pillar Axiom system.

Software Module

Identifies the name of a software or firmware module that is installed on the Pillar Axiom system.

Note: When using an administrator account during a software upgrade process and the package version and the installed version of an update component are the same, the check boxes for those software modules will be disabled (greyed out) and a notation of `Requires Support administrator account role` is displayed adjacent to the module name.

Table 26 Software and firmware modules

Name	Description
Pilot OS	Operating system for the Pilot.
Pilot Software	Software that runs on the Pilot, such as the GUI interface and web server, online help, Small Network Management Protocol (SNMP), and Network Data Management Protocol (NDMP).
Slammer PROM Slammer PROM AX600	Programmable ROM (PROM), which includes BIOS and netboot code, for network attached storage (NAS) and storage area network (SAN) Slammers. Pillar Axiom 600 systems display the *AX600* suffix.

Table 26 Software and firmware modules (continued)

Name	Description
Slammer Software Slammer Software AX600	NAS or SAN software that runs on Slammers. Pillar Axiom 600 systems display the *AX600* suffix.
Brick Firmware Brick SATA2 Firmware	RAID firmware for serial ATA (SATA) Bricks. Pillar Axiom systems that contain version 2 SATA controllers display the *SATA2 Firmware* suffix. **Note:** Version 2 SATA RAID controllers have 16 ports to support the Storage System Fabric (SSF). Version 1 SATA controllers have 13 such ports.
Brick FC Firmware	RAID firmware for Fibre Channel (FC) Bricks.
Brick Disk Drive Firmware	Drive firmware for Bricks.

Installed Version

Identifies the version number of a software module.

Package Version

Identifies the version number of a software module after the software update is complete.

System Restart Required

Identifies whether a shutdown is required after a firmware or software module is updated or a hardware component is installed.

Restart System (optional)

Select this option to force a restart of the system, even though it is not required by a particular software module update. This applies to Brick and Slammer related updates.

Software Component Dependencies

Opens the Software Update, Dependencies Page.

SEE ALSO

Download Software Updates

Update Pillar Axiom Software

Download Software Updates

Modify a Software Update Schedule

Ranges for Field Definitions

Software Modules Page

DESCRIPTION

Use the Software Modules page to review the versions of software and firmware modules that are currently installed on the Pillar Axiom system. The **Actions** drop-down list provides an option to update software.

FIELD DEFINITIONS

Software Module

Identifies the name of a software or firmware module that is installed on the Pillar Axiom system.

Table 27 Software and firmware modules

Name	Description
Pilot OS	Operating system for the Pilot.
Pilot Software	Software that runs on the Pilot, such as the GUI interface and web server, online help, Small Network Management Protocol (SNMP), and Network Data Management Protocol (NDMP).
Slammer PROM Slammer PROM AX600	Programmable ROM (PROM), which includes BIOS and netboot code, for network attached storage (NAS) and storage area network (SAN) Slammers. Pillar Axiom 600 systems display the *AX600* suffix.
Slammer Software Slammer Software AX600	NAS or SAN software that runs on Slammers. Pillar Axiom 600 systems display the *AX600* suffix.
Brick Firmware Brick SATA2 Firmware	RAID firmware for serial ATA (SATA) Bricks. Pillar Axiom systems that contain version 2 SATA controllers display the *SATA2 Firmware* suffix. **Note:** Version 2 SATA RAID controllers have 16 ports to support the Storage System Fabric (SSF). Version 1 SATA controllers have 13 such ports.
Brick FC Firmware	RAID firmware for Fibre Channel (FC) Bricks.
Brick Disk Drive Firmware	Drive firmware for Bricks.

Installed Version

Identifies the version number of a particular software module.

Version

Identifies the release number associated with the compatibility matrix for this Pillar Axiom system. This matrix identifies what software components will work with which specific hardware component versions in the system.

SEE ALSO

Download Software Updates

Update Pillar Axiom Software

Ranges for Field Definitions

Software Modules, Schedule Tab

DESCRIPTION

Use the Software Modules **Schedule** tab to schedule a time at which to update software and firmware modules on a Pillar Axiom storage system.

Tip: You may want to schedule the update to occur during off-peak hours.

FIELD DEFINITIONS

Perform Update Now

Performs the update of the uploaded package immediately.

Schedule Update

Specifies the date and time to schedule the update. You can schedule one future update at a time. If the time of any scheduled update is still in the future, you cannot schedule another update.

Start Time

Identifies the time at which the Pillar Axiom system starts a scheduled operation.

SEE ALSO

Download Software Updates

Update Pillar Axiom Software

Ranges for Field Definitions

Software Update, Dependencies Page

DESCRIPTION

Use the Software Update, Dependencies page to view information about dependencies among the software modules that are installed on the Pillar Axiom system.

FIELD DEFINITIONS

Software Module

Identifies the name of a software or firmware module that is installed on the Pillar Axiom system.

Table 28 Software and firmware modules

Name	Description
Brick Disk Drive Firmware	Disk drive firmware for Bricks.
Brick Firmware	RAID firmware for serial ATA (SATA) Bricks.
Brick FC Firmware	RAID firmware for FC Bricks.
Pilot OS	Operating system for the Pilot.
Pilot Software	Software that runs on the Pilot, such as the GUI interface and web server, online help, Small Network Management Protocol (SNMP), and Network Data Management Protocol (NDMP).
Slammer PROM	Programmable ROM (PROM), that includes BIOS and netboot code, for NAS and SAN Slammers.
Slammer Software	NAS or SAN software that runs on Slammers.

Dependencies within the Staged Update

Identifies that the specified software module has dependencies on other software modules. Install or update prerequisite software first so that this module's dependencies are satisfied.

System Dependencies

Identifies what other software components are needed to install a particular package or software component

Close

Closes the read-only page of information.

SEE ALSO

Download Software Updates

Update Pillar Axiom Software

Display Software Versions

Modify a Software Update Schedule

Ranges for Field Definitions

Software Update, Select Package Page

DESCRIPTION

Use the Software Update, Select Package page to select an update package that contains software updates.

FIELD DEFINITIONS

File Name

Identifies the name of an update package. An update package contains one or more firmware or software modules. If more than one module is in a package, you can install them all at once or separately.

Browse

Opens a browse dialog box so that you can navigate to and select a package.

SEE ALSO

Download Software Updates

Update Pillar Axiom Software

Ranges for Field Definitions

Statistics Tools Page

DESCRIPTION

The Pillar Axiom Statistics Tools allows you to collect Pillar Axiom system statistics and to parse the information into comma-separated values (CSV) files for use in report generators.

FIELD DEFINITIONS

Statistics Tools

Select an option in the **Actions** drop-down list to download the Statistics Tools for use on the following platforms:

- Linux
- Windows

SEE ALSO

Pillar Axiom Statistics Tools User Guide

Storage Usage Summary Page

DESCRIPTION

Use the Storage Usage Summary page to review storage capacity details for filesystems, LUNs, and data replicas.

Sizes reported in capacity usage summaries and total, used, and free system capacities are *raw* physical capacities.

Note: When you request a certain capacity for a newly created logical volume, the system adds some overhead to your request. The capacity consumed and reported for RAID 5 logical volumes includes that overhead. This overhead varies depending on the type of Brick on which the volume is located:

- For SATA Bricks, the overhead is 20%.
- For FC Bricks, the overhead is 10%.

The capacity consumed and reported for Distributed RAID logical volumes is twice that requested, regardless of Brick type.

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

FIELD DEFINITIONS

SYSTEM CAPACITY (GB)

Free

Identifies the current unused storage capacity that is available on the Pillar Axiom system. Free capacity is total capacity minus allocated capacity.

Reconditioning

Identifies the amount of storage capacity that is in the process of being released back to the free pool.

Allocated

Identifies the total amount of storage capacity that is reserved on the system.

Total

Identifies the total storage capacity, free plus allocated capacity, of the system.

VOLUME SUMMARY

Filesystems

Identifies the number of filesystems that have been created on the Pillar Axiom system and the amount of storage capacity allocated for them.

Clone FSs

Identifies the number of Clone FSs that have been created on the Pillar Axiom system and the amount of storage capacity allocated for them.

LUNs

Identifies the number of LUNs that have been created on the Pillar Axiom system and the amount of storage capacity allocated for them.

Clone LUNs

Identifies the number of Clone LUNs that have been created on the Pillar Axiom system and the amount of storage capacity allocated for them.

Total

Identifies the total number of logical volumes (filesystems and LUNs and replicas that have been created on the Pillar Axiom system.

Snap FS

Identifies the number of Snap FSs that have been created on the Pillar Axiom system.

Usage Graph

Displays how much storage capacity is used in each of the relative priority levels and how much capacity is free.

SEE ALSO

About Volume Capacity and Provisioning

Display Capacity Usage

Display Filesystem Details

Display Data Replica Details

Ranges for Field Definitions

System Summary Page

DESCRIPTION

Use the System Summary page to review system status and configuration for the Pillar Axiom storage system.

FIELD DEFINITIONS

System Status

Displays the overall system status of the hardware components. A status of **Normal** requires no action. Investigate any other status:

- Look for notifications in the status bar that indicate that some action is required and take the suggested action. See About Responding to Administrator Actions.
- Check the Health Summary Page.

Management IP

Identifies the IP address of the Pillar Axiom system.

System Name

Identifies the name that is assigned to the Pillar Axiom system. The system name also appears in the status bar.

Model

Displays the model number of a hardware component.

System Serial Number

Identifies the system serial number (SSN) that is assigned to the Pillar Axiom system.

Date and Time

Identifies the current Pillar Axiom date and time.

Time Zone

Identifies the current time zone that is defined for the Pillar Axiom system.

NTP

Identifies whether the Pillar Axiom system synchronizes its clocks with:

- Network Time Protocol (NTP) servers
- A manual time setting that you specify

Status Bar Icons and **Messages**

- Click the **System Status** icon (one of the following images) to open the Health Summary Page.

   Green circle: **Normal** status.

   Yellow inverted triangle: **Warning** status.

   Red octagon: **Critical** status.

For more information, see Status Bar Information.

- Click the **Administrator Actions** icon, a yellow triangle containing an exclamation point (!), to determine what corrective actions should be performed.
- Text in the status bar identifies the name of the system, your login name, and the number of logged-in administrators (users). Click the **Administrators** icon (a person's head) to display details about additional logged-in administrators.
- Text in the status bar identifies the currently running background task. Click the **Running Tasks** icon (a running person) to display details about additional background tasks.

SEE ALSO

Configure Global Settings

Configure the Pillar Axiom System Time

Create an Administrator Account

Create Alerts

Create an NDMP User Account

Create SNMP Trap Hosts

Manage System Tasks

Modify System Name

About Responding to Administrator Actions

System Time Page

DESCRIPTION

Use the System Time page to synchronize the Pillar Axiom clock time with a Network Time Protocol (NTP) server or to set the date and time manually.

FIELD DEFINITIONS

System Name

Identifies the name that is assigned to the Pillar Axiom storage system.

Time Zone

Identifies the time zone that is used for the system's local time. Select the time zone of the site at which the Pillar Axiom system is installed. Select a different time zone if you want to display local time as something other than the installed site's time.

Set Time Manually

Identifies that the Pillar Axiom system synchronizes its clocks with a date and time that you set. Pillar Axiom clocks are synchronized with each other, and their time may differ from other clocks in your network.

Date

If you set the Pillar Axiom date and time manually, select the current month, day, and year.

Or, click the date selector icon to choose a date from the representation of calendar pages.

Time

If you set the Pillar Axiom date and time manually, select the current hours and seconds.

Use NTP Server

Identifies that the Pillar Axiom storage system synchronizes its clocks with Network Time Protocol (NTP) servers.

- Enable NTP if you want to specify a primary NTP server and up to two alternate NTP servers with which the system synchronizes its clocks.
- Disable NTP if you want to manually set the date and time for the Pillar Axiom system.

Primary NTP Server

Identifies the IP address that is assigned to an NTP server with which the Pillar Axiom system synchronizes its clocks.

Alternate NTP Server

Identifies the IP address of alternate NTP servers. If the primary NTP server is unavailable, the Pillar Axiom storage system consults the alternate servers in

round-robin fashion until the system connects to an available NTP server. Enter IP addresses for up to two alternate NTP servers.

SEE ALSO

Configure the Pillar Axiom System Time

Configure Notification Settings

Ranges for Field Definitions

Tape Devices Page

DESCRIPTION

Use the Tape Devices page to review the tape devices that are attached to the Pillar Axiom system. The **Actions** drop-down list provides an option to check for attached tape devices.

FIELD DEFINITIONS

Refresh Interval

Identifies the interval at which the page is updated with current data. Choose from:

- **1 min**
- **5 min**
- **10 min**

Set

Saves the specified refresh interval. The data on the page is next updated after that number of minutes.

Refresh Now

Updates the page with current data.

Device Name

Identifies the name and other identification details of a tape device that is attached to the Pillar Axiom system. For example, if the device name is `Tape-20-1`, the components of that name mean the following:

- `Tape` is the device type such as robot or tape.
- `2` is the target ID (for SCSI devices, this the SCSI identifier).

  Note: SCSI tape devices are not supported on Pillar Axiom 600 systems.
- `0` is the LUN ID.
- `1` is a unique identifier to prevent duplicate names.

Slammer

Identifies the name of the Slammer to which the tape storage device is attached.

Control Unit

Identifies the specific Slammer control unit (CU) to which this tape device is attached.

Inquiry Data

Displays information about the tape device. Different types of tape devices report different inquiry data.

Device Type

Identifies the type of tape device that is attached to the Pillar Axiom system:

- **Tape**

- **Robot**

SEE ALSO Replace a Hardware Component

Ranges for Field Definitions

Tools Overview Page

DESCRIPTION

Use the Tools Overview page to select any of the types of support tools that are available on a Pillar Axiom storage system. After you select a type, you can select and perform the specified support operation to help diagnose and resolve system issues.

FIELD DEFINITIONS

Support Tool Type

Lists types of tools that are provided to support Pillar Axiom systems. Click a type link to display the support tools of that type.

- **Verify System Operations** opens the Verify System Operations Page.
- **Collect System Information** opens the Collect System Information Page.
- **Resolve System Trouble** opens the Resolve System Trouble Page.
- **Resolve Connectivity Trouble** opens the Resolve Connectivity Trouble.

SEE ALSO

The set of tools listed in About Pillar Axiom Support Tools

Ranges for Field Definitions

Upload SSL Certificate Page

Description

Use the Upload SSL Certificate page to select a secure sockets layer (SLL) certificate so that you can upload the certificate to the Pilot. Use the certificate so that administrators can mange the Pillar Axiom system in secure GUI sessions only. The SSL certificate requires both the certificate and the key files.

FIELD DEFINITIONS

Certificate File

The certificate file should contain the public key of the certificate along with the certificate details. This file should be a certificate in X.509 format, encoded according to Distinguished Encoding Rules (DER), and then encoded according to Privacy Enhanced Mail (PEM).

Key File

The key file should be the PEM encoded private key that corresponds to the certificate. This file must not be password protected.

SEE ALSO

Ranges for Field Definitions

UPS Page

DESCRIPTION

Use the UPS page to review the current status of the uninterruptible power supply (UPS) battery and power.

Note: The Pillar Axiom system supports PowerNet Management Information Base (MIB) version 3.4.4.

FIELD DEFINITIONS

IP Address

Identifies the IP address that is assigned to the UPS.

SNMP Community String

Enter the community string of the specified UPS. The community string is required to allow Small Network Management Protocol (SNMP) queries to the specified IP address.

Model

Displays the model number of a hardware component.

Firmware Revision

Displays the firmware version that is installed on a hardware component.

Serial Number

Displays the serial number of a hardware component.

Power Source

Displays the source of the UPS power. The value is one of:

- AC (alternating current)
- Battery

Battery Status

Displays the current status of the UPS batteries.

SEE ALSO

Display Hardware Component Information

Ranges for Field Definitions

Utilities Download Page

DESCRIPTION

Use the Utilities Download page to download the following Pillar Axiom utilities:

- SnapDelta FS
- Command Line Interface (CLI)
- Statistics Tools
- Virtual Disk Service

FIELD DEFINITIONS

Command Line Interface

Downloads a client-based application that allows you to pass administrative commands to the Pillar Axiom system from a shell or script.

SnapDelta FS

Downloads a client-based application that allows you to determine changes within a Pillar Axiom filesystem during the interval defined between two Snap FSs.

Statistics Tools

Downloads a client-based application that allows you to work with statistics derived from the Pillar Axiom system.

Virtual Disk Service

Downloads a Windows-based plug-in that allows you to use Windows to manage Pillar Axiom storage devices.

SEE ALSO

Command Line Interface Page

Download Pillar Axiom Virtual Disk Service Provider

SnapDelta FS Download Page

Statistics Tools Page

Virtual Disk Service (VDS) Page

Verify System Operations, Data Consistency Tab

DESCRIPTION

Use the Verify System Operations **Data Consistency** tab to verify the integrity of the data stored on a particular Brick.

Important! Verifying data consistency takes a significant amount of time and impacts system performance. Pillar recommends performing this operation at times when the system is not being fully utilized.

FIELD DEFINITIONS

Verify Data Integrity on Brick

Lists the names of the Bricks installed in the Pillar Axiom system. Select the Brick to be verified.

High Priority

Specifies that data verification on the selected Brick should be performed as a *high* priority task.

Important! A priority level of *high* could impact performance of the selected Brick by up to 30%.

Low Priority

Specifies that data verification on the selected Brick should be performed as a *low* priority task.

Important! A priority level of *low* could impact performance of the selected Brick by up to 10%.

Last Results

Lists the date and the failed or passed status of a Brick the last time its data integrity was verified.

SEE ALSO

Verify System Operations, Storage Redundancy Tab

Ranges for Field Definitions

Verify System Operations Page

DESCRIPTION

Use the Verify System Operations page to run the verification tools that are available on a Pillar Axiom storage system and to review verified system information.

FIELD DEFINITIONS

Verification

Lists types of verification tools that are provided to support Pillar Axiom systems. Click a type link to run the selected verification tool.

For details about how to verify different types of system operations, see:

- Verify Data Consistency
- Verify Storage Redundancy

Result

Displays the result of the verification test.

Date Last Run

Displays the last date in which the verification test was run.

Target

Identifies the hardware components from which system information was collected.

SEE ALSO

The set of tools listed in About Pillar Axiom Support Tools

Ranges for Field Definitions

Verify System Operations, Storage Redundancy Tab

DESCRIPTION

Use the Verify System Operations **Storage Redundancy** tab to verify the redundancy of a filesystem.

Important! Verifying storage redundancy takes a significant amount of time and impacts system performance. Pillar recommends performing this operation at times when the system is not being fully utilized.

FIELD DEFINITIONS

Check Redundancy on Filesystem

Specifies the filesystem for which you want to verify the storage redundancy.

Last Results

Lists the date and the failed or passed status of each filesystem the last time their storage redundancies were verified.

SEE ALSO

Verify System Operations, Data Consistency Tab

Ranges for Field Definitions

Verify System Operations, Summary Tab

DESCRIPTION

Use the Verify System Operations **Summary** tab to choose the areas of the system to check.

Important! Verifying data consistency and storage redundancy takes a significant amount of time and impacts system performance. Pillar recommends performing these operations at times when the system is not being fully utilized.

FIELD DEFINITIONS

Data Consistency

Select this option to check the integrity of the data on a specific Brick.

Storage Redundancy

Select this option to check the redundancy of a specific filesystem.

SEE ALSO

Verify System Operations, Data Consistency Tab

Verify System Operations, Storage Redundancy Tab

Ranges for Field Definitions

Virtual Disk Service (VDS) Page

DESCRIPTION

Use the Virtual Disk Service (VDS) page to download the Pillar Axiom VDS Provider plug-in to your administrative workstation. The VDS API provides you with a method to manage storage devices and allows you to:

- Create, delete, and extend LUNs
- Mask and unmask LUNs
- Obtain status of storage devices (Slammers, Bricks, disk drives, and LUNs)

VDS runs on the Windows 2003 platform.

SEE ALSO

Ranges for Field Definitions

Download Pillar Axiom Virtual Disk Service Provider

The Microsoft *Virtual Service Technical Reference*

Virtual Interfaces Page

DESCRIPTION

Use the Virtual Interfaces page to create and delete secondary virtual network interfaces for a File Server. Create the primary virtual network interface on the File Server Page, Network Tab.

Tip: With regard to the virtual interfaces (VIFs) that are associated with a particular File Server, Pillar recommends the following:

- Locate all VIFs for a File Server on the same Slammer control unit (CU).

  Locating the VIFs on multiple CUs can impact performance, because the VIF selected using the Domain Name System (DNS) may not be on the Slammer CU on which the filesystem resides. When you cannot locate all the VIFs for a File Server on the same CU but must use several CUs, avoid using the DNS name to mount or map the VIFs. Instead, use the IP address.

  For example, consider a File Server with the following VIF IPs:
  - 10.1.2.30 on Slammer CU 0
  - 10.1.2.31 on Slammer CU 1

  For filesystems located on CU 0, you should use 10.1.2.30 for mapping or mounting requests.
- Register all of these VIFs with DNS as belonging to this File Server, even when all the VIFs cannot be located on the same CU.
- Avoid using virtual LAN (VLAN) tags.

FIELD DEFINITIONS

Virtual Interfaces

Displays a list of currently configured VIFs. Check the box to the left of an interface definition to select it for modification or deletion.

Remove

Deletes the selected objects.

IP Address

Identifies the IP address that is assigned to the VIF.

Netmask

Identifies the subnet mask that is assigned to the VIF.

Slammer

Identifies the Slammer that is assigned to the VIF.

Control Unit

Identifies the Slammer CU that is assigned to the VIF.

PORT Number

Identifies the port on the Slammer CU that is assigned to the VIF. Select a gigabit Ethernet network port on one of the Slammer CUs.

- If link aggregation (IEEE 802.1AX-2008) is enabled, select port 0 (zero). You cannot create virtual interfaces on alternate port 1 when link aggregation is enabled.
- If link aggregation is disabled, select either port.

Identify Port

Opens the Identify Hardware Component pages so that you can blink the light-emitting diodes (LEDs) on the specified hardware component (**Identify**) or on all other components (**Reverse Identify**).

VLAN ID

Identifies the virtual LAN (VLAN) ID that is assigned to the VIF.

Tip: We recommend that you avoid using VLAN tags. Because you can configure multiple File Servers on the same VLAN that does not use tags, we recommend that you configure your system so that all VIFs are not tagged.

To implement VLAN tagging, assign a value of 1 through 4094, inclusive, if you have connected a VLAN-capable switch to the Pillar Axiom system. Leave the field blank to disable VLAN tagging.

Frame Size

Identifies the packet frame size. This value defines the maximum transmission unit (MTU).

The frame size (MTU) does not include the Ethernet header portion of the packet. If your network switch has trouble with this, you can set the switch to a larger value or lower the MTU size to correct the problem.

- If your network supports jumbo frames (extended Ethernet), enter an integer greater than 1500 and less than 16362. Make sure that this Pillar Axiom MTU size matches the network MTU size. If the MTU sizes are mismatched, users may experience I/O hangs when the client machines try to process packets that are too large.
- If your network does not support jumbo frames, enter the default frame size of 1500.

Create

Adds the secondary virtual network interface to the File Server configuration.

Remove

Removes the selected VIF from the File Server.

SEE ALSO

Create a File Server

Modify File Server Attributes

Ranges for Field Definitions

Volume Group Details

DESCRIPTION

Use the Volume Group Details page to create and modify the organizational units that group any number of logical volumes (filesystems and LUNs).

FIELD DEFINITIONS

Name

Identifies the name that is assigned to a volume group. Each volume group name must be unique within its parent volume group.

Create Volume Group In *(Create)* or **Volume Group Location** *(Modify)*

Identifies the location of the volume group. Starts at the System (root) level and traverses through parent and child volume groups.

CAPACITY LIMITS

Unlimited Capacity or Limit Capacity

Identifies whether the volume group is configured with unlimited or limited capacity. Choose from:

- **Unlimited Capacity** if the volume group should not constrain the maximum capacity of associated logical volumes and nested volume groups.
- **Limit Capacity** if the maximum capacity of associated logical volumes and nested volume groups cannot exceed the value that you specify.

Maximum Capacity

Identifies the maximum capacity for the volume group. The maximum capacity of the logical volumes and nested volume groups that are associated with the volume group cannot exceed this value.

A value of 0 (zero) identifies that the volume group is configured with unlimited capacity. You can increase the maximum capacity of associated logical volumes and nested volume groups without constraints.

SEE ALSO

Create Volume Groups

Modify Volume Group Attributes

Move a Logical Volume to a Different Volume Group

Delete Volume Groups

Ranges for Field Definitions

Volume Groups Overview Page

DESCRIPTION

Use the Volume Groups Overview page to review the volume groups and logical volumes (filesystems and LUNs) that are configured on the Pillar Axiom system. The **Actions** drop-down list provides options to create, modify, delete, move, and view volume groups and logical volumes, as well as to create an immediate Clone FS or Clone LUN.

When dashes appear in a field, it means that the field is not applicable to a specific object type.

FIELD DEFINITIONS

Name

Lists the names of configured volume groups. Click a name to review or modify the volume group settings.

Type

Identifies the object type:

- **FS** is a NAS filesystem.
- **LUN** is a SAN LUN.
- **VG** is a volume group.

Status

Identifies the current status of each logical volume and volume group.

Storage Class

Identifies the category of physical storage on which the logical volume resides:

- **FC** (Fibre Channel drives)
- **SATA** (Serial ATA drives)
- **SLC SSD** (single level cell, solid state drives)

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

Total

Identifies the total volume capacity (free plus used capacity) of the Pillar Axiom system.

Free

Identifies the current unused volume capacity that is available. Free capacity is total capacity minus used capacity.

Used

Identifies the current volume capacity usage of the object.

VG Capacity Max (GB)

Identifies the maximum capacity limit for the object.

IOps

Identifies the current load of the filesystem on performance for read (input) and write (output) operations.

Bandwidth (MB/sec)

Identifies the current load of the filesystem on bandwidth in MB/sec for read and write operations.

Volume Totals

Identifies the total allocated capacity, both free and used, for all volumes on the Pillar Axiom system.

SEE ALSO

Create Volume Groups

Modify Volume Group Attributes

Ranges for Field Definitions

Volume Shadow Copy Service (VSS) Page

DESCRIPTION

Use the Volume Shadow Copy Service (VSS) page to download the Pillar Axiom VSS Provider plug-in to your administrative workstation. VSS allows you to create and maintain shadow copies of volumes and files, including open files.

During backups:

- Applications continue to write data.
- Open files are included in the backup.
- Users are not locked out.

VSS runs on the Windows 2003 platform.

SEE ALSO

Download and Install the VSS Provider Plug-In

Microsoft's *Volume Shadow Copy Service Technical Reference*

Ranges for Field Definitions

Index

W