FlashBlade//EXA Administration Guide

Version 4.6.7

PURESTORAGE®

Copyright Statement

Pure Storage, Inc. 2555 Augustine Drive, Santa Clara, CA 95054 USA *http://www.purestorage.com*

Direct comments to DocumentFeedback@purestorage.com.

Table of Contents

Chapter 1
About this Guide

The Pure Storage FlashBlade® Administration Guide provides information about FlashBlade//EXA, its features, and how to use them. This guide is written for array administrators who view and manage the Pure Storage FlashBlade//EXA storage systems. FlashBlade//EXA arrays are administered through the Purity for FlashBlade (Purity//FB) graphical user interface (GUI) or command line interface (CLI). Users should be familiar with system, storage, and networking concepts, and have a working knowledge of Windows or UNIX.

For full information about the CLI commands referenced in this document, refer to the *FlashBlade CLI Reference*.

Organization of this Guide

This guide is organized as follows:

Purity//FB Overview - Defines FlashBlade array naming and numbering conventions.

Purity//FB GUI Overview and (subsequent topics) - Describes the use of the browser-based Purity//FB graphical user interface (GUI) to administer a FlashBlade array.

Choosing an Active Directory Account or Directory Services Configuration - Provides an overview of Active Directory and Directory Services configurations for array and file management and accessibility, as well as configuration instructions for each.

A Note on Format and Content

FlashBlade technology is evolving rapidly. As with all advanced information technologies, once basic architecture is in place, implementation develops at different rates in its different facets, each preceded by feature-by-feature detailed design.

This edition of the guide describes the properties and behavior of arrays that run the latest Purity//FB release. It may include information about planned capabilities whose external form has been specified at the time of the release. Such material is included to provide users of this release with information for design planning purposes, and is subject to change as new functionality is implemented. Material relating to not-yet-implemented functionality is identified as such in the text.

Documentation and Support

All related guides are or will soon be available in the Pure Storage Knowledge Base, which is located at *http://support.purestorage.com*.

Contact Us

We welcome your feedback about Pure Storage documentation and encourage you to send your questions and comments to *documentfeedback@purestorage.com*.

Product Support

If you are a registered Pure Storage user, log in to *http://support.purestorage.com* to browse our knowledge base, view the status of your open support cases, and view the details of past support cases.

You can also contact Pure Storage Technical Services at *support@purestorage.com*.

Chapter 2
Purity//FB Overview

Purity for FlashBlade (Purity//FB™) is the operating environment that queries and manages the FlashBlade hardware, networking, and storage components. The Purity//FB software is distributed with the FlashBlade array.

Purity//FB provides three ways to administer the FlashBlade array: through the browser-based graphical user interface (Purity//FB GUI), the command line interface (Purity//FB CLI), and FlashBlade REST API.

Purity//FB Conventions

Purity//FB follows certain conventions.

Object Names

Valid characters are letters (A-Z and a-z), digits (0-9), and the hyphen (-) character. File system names can also include the underscore (_) character. The first and last characters of a name must be alphanumeric, and a name must contain at least one letter.

Most objects in Purity//FB that can be named, including file systems, data vips, and subnets can be 1-63 characters in length.

Array names can be 1-56 characters in length. The array name length is limited to 56 characters so that the names of the individual controllers, which are assigned by Purity//FB based on the array name, do not exceed the maximum allowed by DNS.

Purity allows the use of lowercase and uppercase in names and preserves capitalization in command output. However, duplicate naming with different capitalization is not allowed. Additionally, search ignores capitalization differences.

Array and File System Sizes

Array and file system sizes are specified as integers followed by one of the suffix letters `K`, `M`, `G`, `T`, `P`, representing `KiB`, `MiB`, `GiB`, `TiB`, and `PiB`, respectively, where " `Ki` " denotes 2^10, " `Mi` " denotes 2^20, and so on.

Note that the raw capacity value displayed on the **Hardware > Health** page displays gibibyte (`G`), gigabyte (`GB`), tibibyte (`T`), terabyte (`TB`), pebibyte (`P`), and petabyte (`PB`) units.

Time Zones

Purity//FB GUI dates and times are displayed in the user's local time zone, which is determined by the browser's local time. The user's local time zone appears at the bottom of the navigation pane of the Purity//FB GUI.

Purity//FB CLI dates and times are displayed in the time zone of the array, which is set during FlashBlade installation. The array time zone appears and can be edited in the **Time** panel of the **System > Settings** page.

IP Addresses

FlashBlade supports two versions of the Internet Protocol: IP Version 4 (IPv4) and IP Version 6 (IPv6). IPv4 and IPv6 addresses follow the addressing architecture set by the Internet Engineering Task Force.

An IPv4 address consists of 32 bits and is entered in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from `0` to `255` representing a group of 8 bits. Here are some examples:

```
puredns setattr --domain mydomain.com --nameservers 192.0.2.10
purenetwork setattr --address 192.0.2.0 DataVip01
puresubnet create --prefix 192.0.2.0/24 --vlan 100 Sub01
```

An IPv6 address consists of 128 bits and is written in the form `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Colons separate each 16-bit field. Leading zeros can be omitted in each field. Furthermore, consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`) . For example, IPv6 address `2001:0db8:85a3:0000:0000:8a2e:0370:7334` becomes `2001:db8:85a3::8a2e:370:7334`.

To use an IPv6 address in a URL or URI, enclose the entire address in square brackets (`[]`) . When specifying a URL or URI with a port number, append the port number after the end of the entire address. Here are some examples:

```
puredns setattr --domain mydomain.com --nameservers 2001:db8:85a3::8a2e:370:7334
purenetwork setattr --address 2001:db8:85a3::8a2e:370:7334 DataVip01
puresubnet create --prefix 2001:db8:85a3::/64 --vlan 100 Sub01
```

SMB Port Assignments

This section provides the minimum ports required on each VLAN for SMB to function successfully. This information should be provided to your network team who configures the firewall.

Management Network Ports

The following table provides the ports required on the management network. Without these ports opened, SMB and LDAP will not be able to authenticate. These ports must be open between the FlashBlade management VIP and the LDAP/Active Directory (AD) server(s).

Port	TCP	UDP	Purpose
53	Yes	Yes	DNS used to locate the AD hosts and services.
123	Yes	Yes	NTP to avoid time skews, which needs to be limited for authentication.
389	Yes	Yes	Used to manage array integration with a directory service using LDAP. Active Directory requires both TCP and UDP. LDAP only requires TCP.
636	Yes	Yes	Used to manage array integration with a directory service using LDAPS (LDAP over TLS/SSL).
445	Yes	No	Used by SMB join/unjoin AD and perform all authentication/authorization via the management network.

Firewall Ports

The following table lists the required firewall ports for FlashBlade access. These ports are outgoing traffic from the FlashBlade to the public network.

Port	TCP	UDP	Purpose
443	Yes	No	Used for both HTTPS and SSH over HTTPS for IPv4. Use the following IP block and hostnames: • IP block 52.40.255.224/27 • Hostname: cloud-support.purestorage.com Static IPs are not supported for IPv6. Use the hostname: .cloud-support.purestorage.com Add all items above to your whitelist.

Data Network Ports

The following table describes the required ports on the data network. These ports must be open between the FlashBlade data VIP and SMB hosts.

Port	TCP	UDP	Purpose
53	Yes	Yes	Share name resolution.
445	Yes	No	SMB data traffic. SMB 2.1 is supported; therefore port 139 is not needed.

About NFSv4.1 Support

In the current implementation of the NFSv4.1 protocol, FlashBlade supports the following subset of core NFSv4.1 protocol features, which are compliant with Linux clients:

- Native byte-range locking semantics required by Linux clients
- All NFS traffic through single secure port (2049)
- Top level root mountable with file systems as directories
- Exportable to NFSv3 and NFSv4.1 clients simultaneously
- Client qualification: CentOS/RedHat OS

- Management support using the FlashBlade GUI, CLI, and REST API
- Kerberos

The following features of the NFSv4.1 protocol are currently unavailable:

- pNFS
- Delegation
- Referrals
- Shared reservation

⚠ **Important!** NFSv4.0 is not supported.

Top Level Mounting

In the current implementation of NFSv4.1, hosts can mount at the top level root (`/`) with file systems as directories. File systems can be exported to different clients with different export rules and permissions. Note that issuing `ls` on `/` displays the file system that is mountable to that host, and issuing `df` on `/` displays space information for `/` and any mountable, actively used (accessed within the last 60 seconds) file systems.

Figure 2.1 Top Level Mounting

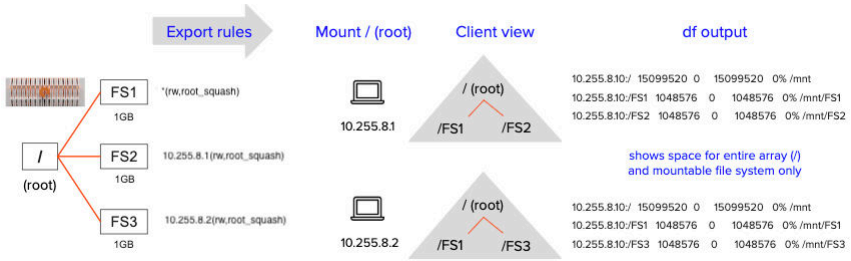

Troubleshooting

Help! Who do I contact about problems with the FlashBlade Array?

Contact a member of your Pure Storage® account team, visit us at *http://support.purestorage.com*, or email Pure Technical Services at *support@purestorage.com*.

If you are contacting Pure Technical Services about a technical issue, they may ask you to open an RA session so that they can help you diagnose the issue.

Error Messages

The following error messages may appear when you perform operations through the Purity//FB GUI or CLI:

Service Temporarily Unavailable

An internal service is currently unavailable. This is a temporary issue. Wait a few minutes and then try the Purity//FB GUI or CLI operation again. If you still get the error message, contact Pure Technical Services at *support@purestorage.com*.

Unexpected Error

The Purity//FB GUI or CLI operation that you performed generated an unexpected error. Contact Pure Technical Services at *support@purestorage.com*.

Chapter 3
Purity//FB GUI Overview

The Purity//FB graphical user interface (GUI) is a web application used to view and administer the FlashBlade array.

The screens of the Purity//FB GUI differ slightly depending upon if your FlashBlade array is a single-chassis or is a multi-chassis configuration. The following figure displays the GUI for a single-chassis configuration.

Figure 3.1 Purity//FB GUI

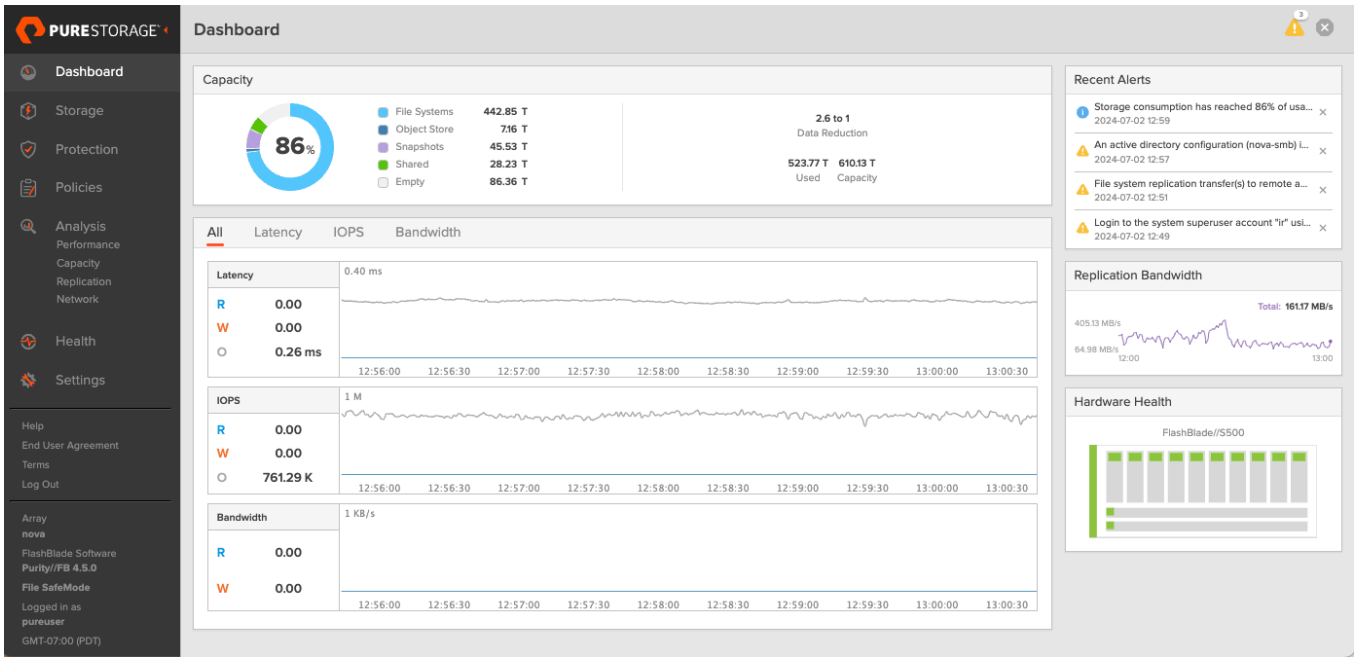

The Purity//FB GUI contains the following pages:

Dashboard

Represents a graphical overview of the array, including hardware status, recent alerts, storage capacity, replication bandwidth, and input/output (I/O) performance metrics.

Storage

Displays summary array details including connection and replication information. In addition, quick links are provided to access file systems, snapshots, object store user, account, bucket, and object information, and export rules. Additionally, a **File Systems** tab provides a list of file systems on the array and its attributes, including provisioned size, capacity usage details, and export rules. A **Servers** tab displays information for all servers on the array.

Protection

Displays backup and replication information. A list of file system snapshots, their source array, source file system, replication policy, and creation date is displayed. Tabs for **Snapshots**, **File Replica Links**, and **Object Replica Links** are provided to view details about file system snapshots and file and object replication link information.

Policies

Displays policy information. Policies and their rules can be viewed and managed from this page.

Analysis

Displays historical array information, including space and I/O performance metrics, from various viewpoints.

Health

Displays array health including alerts, hardware status, and parity.

Settings

Displays array-wide system and network settings. Manage array-wide components, including network interfaces, system time, connectivity and connection configurations, directory services, support settings, SSL certificates, and alert settings.

Purity//FB GUI Navigation

The dark gray navigation pane that appears along the left side of the Purity//FB GUI contains links to the Purity//FB GUI pages.

Figure 3.2 Purity//FB GUI - Navigation Pane

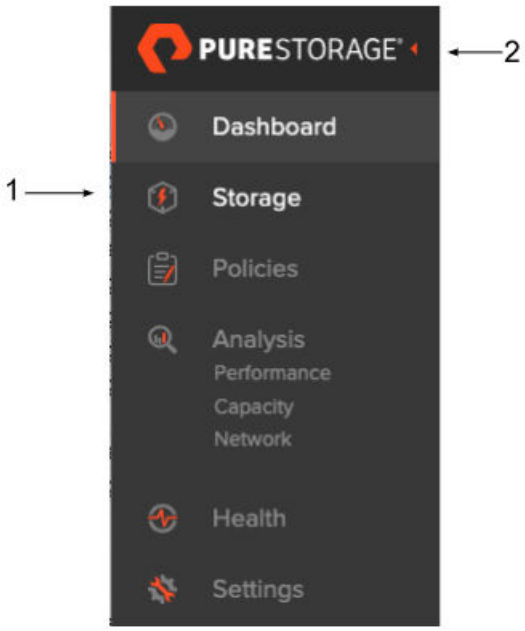

Table 3.1 Navigation Pane Description

Callout	Description
1	Purity//FB GUI links
2	Show/Hide button

Click a link in the navigation pane to analyze or configure the information that appears in the page to its right. For example, click the **Storage** link to view information about the FlashBlade array, file systems, and object store..

Click the **show/hide** button to toggle between the expanded and collapsed views of the navigation pane. The show/hide button appears to the right of the Pure Storage logo.

The navigation pane also includes links to the following external sites:

Help

Launches the FlashBlade user guide, FlashBlade CLI Reference, REST API guide, community and support websites.

End User Agreement

Launches the Pure Storage End User Agreement pop-up window.

Terms

Launches the Pure Storage Legal Terms, Programs, and Product Information web page.

Log Out

Logs the current user out of the Purity//FB GUI.

Note: If no security update has been installed, the last installed Purity//FB security update will not appear at the bottom of the navigation pane. Additionally, if Purity//FB was recently updated, and the update includes the security update, the security update will not appear.

The bottom of the navigation pane displays FlashBlade array name and Purity//FB version information, the version of the last installed Purity//FB security update (if it still applies), the name of the user who is currently logged into the Purity//FB GUI, and the user's local time zone. The dates and times that appear in the Purity//FB GUI are based on the user's local time zone, which is determined by the browser's local time. Addiitionally, if enabled, File SafeMode and/or Object SafeMode are displayed.

The page to the right of the navigation pane displays the information and configuration options for the selected Purity//FB GUI link. The information on each page is organized into panels, charts, and lists.

Figure 3.3 Purity//FB GUI Page

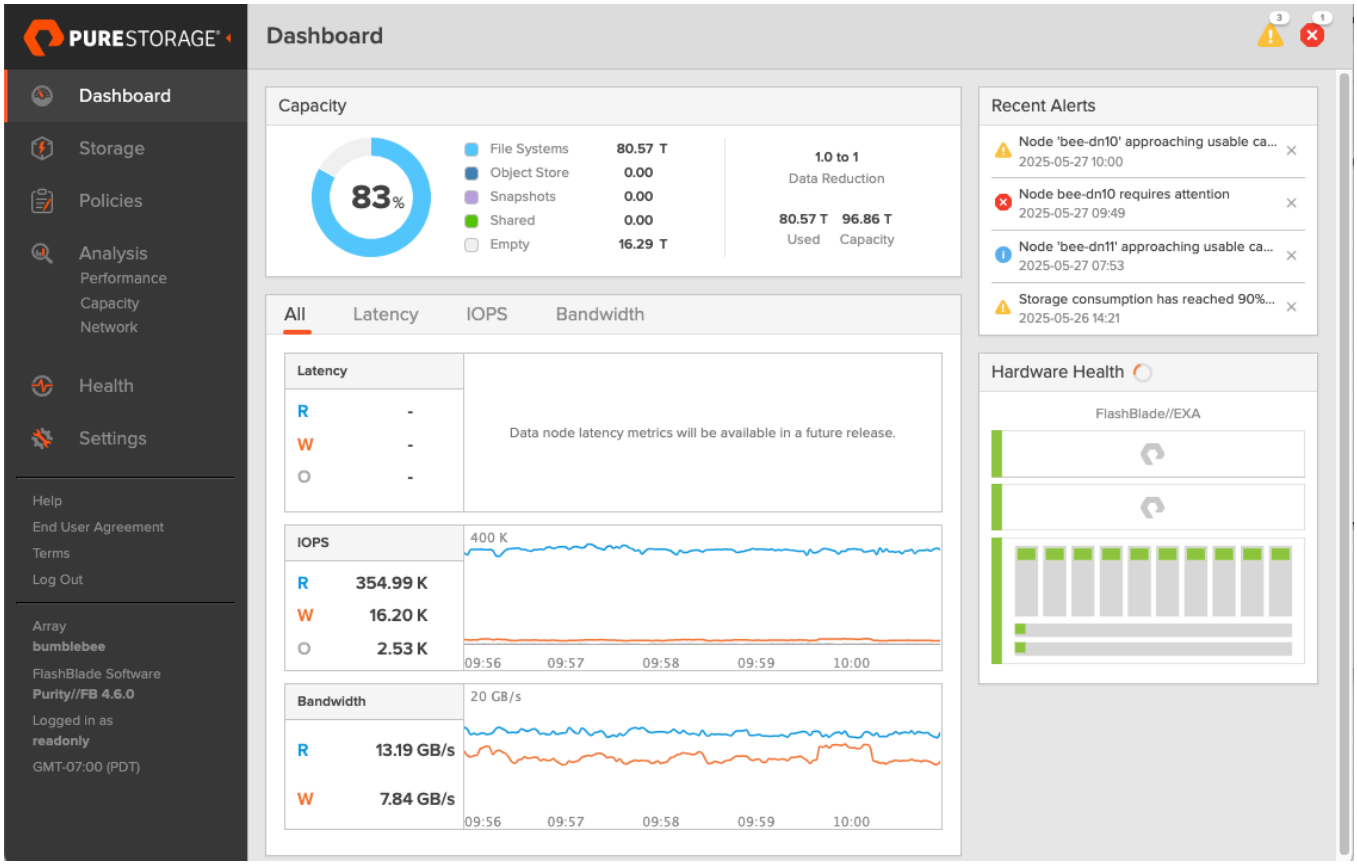

The alert icons that appear to the far right of the title bar indicate the number of open **Warning** and **Critical** alerts. Alerts in the `open` and `closing` states are included in the alert count displayed by the icons.

When an alert is in the `closing` state, it is still considered open. If after 24 hours there is no recurrence or increase in severity, the alert will close. However, if there is a recurrence or increase in severity, the alert will move from the `closing` state back to the `open` state. To analyze alerts in more detail, click the **Health** link.

Various panels, such as **Storage > File Systems** and **Health > Alerts**, contain lists of information. The total number of rows in a list output is displayed in the upper-right corner of the list. Some lists can be very large, extending beyond hundreds of rows.

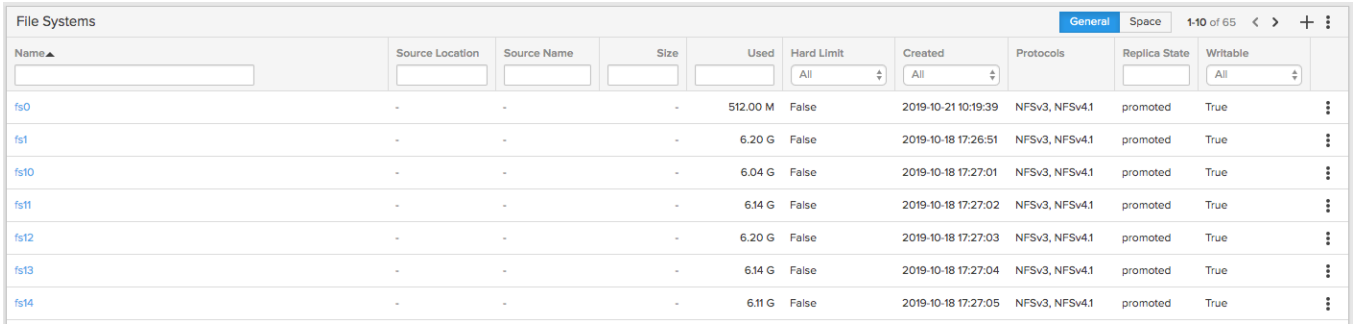

Name	Source Location	Source Name	Size	Used	Hard Limit	Created	Protocols	Replica State	Writable
					All	All			All
fs0	-	-	-	512.00 M	False	2019-10-21 10:19:39	NFSv3, NFSv4.1	promoted	True
fs1	-	-	-	6.20 G	False	2019-10-18 17:26:51	NFSv3, NFSv4.1	promoted	True
fs10	-	-	-	6.04 G	False	2019-10-18 17:27:01	NFSv3, NFSv4.1	promoted	True
fs11	-	-	-	6.14 G	False	2019-10-18 17:27:02	NFSv3, NFSv4.1	promoted	True
fs12	-	-	-	6.20 G	False	2019-10-18 17:27:03	NFSv3, NFSv4.1	promoted	True
fs13	-	-	-	6.14 G	False	2019-10-18 17:27:04	NFSv3, NFSv4.1	promoted	True
fs14	-	-	-	6.11 G	False	2019-10-18 17:27:05	NFSv3, NFSv4.1	promoted	True

Pagination divides a large list output into discrete pages. Pagination is in effect if the number of lines in the list output exceeds 25 rows for tabs displaying a single table, or 10 rows for tabs displaying multiple tables. To move through a paginated list, click **<** to go to the previous page, or click **>** to go to the next page.

Purity//FB GUI Login

Logging in to the Purity//FB GUI requires the management virtual IP address or fully-qualified domain name (FQDN) and an Purity//FB login username and password; this information is provided during the FlashBlade installation.

FlashBlade is installed with one administrative account with the username `pureuser`.

As a best practice, for security purposes Pure Storage recommends that the password for the account be changed immediately upon first log in using the `pureadmin` CLI command. Refer to the `pureadmin` command in the *FlashBlade CLI Reference* for details.

Note: Upon initial log in, a password alert will be generated asking you to change the `pureuser` account password. This alert will appear as a weekly reminder until the password is changed. Once the password is changed, you cannot change it back to the default value.

Pure Storage tests the Purity//FB GUI with the most recent version of the following web browsers:

- Apple Safari
- Google Chrome
- Mozilla Firefox

To launch the Purity//FB GUI login screen, open a browser and type the virtual IP address or FQDN into the address bar, and press **Enter**.

Figure 3.4 Purity//FB GUI - Login Screen

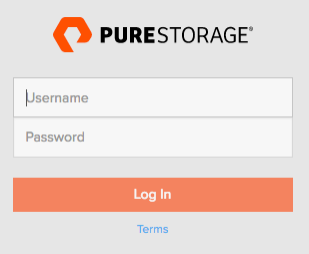

Logging in to the Purity//FB GUI

Log in to the Purity//FB GUI to view and administer the FlashBlade array.

To log in to the Purity//FB GUI:

1. Open a web browser.
2. Type the virtual IP address or fully-qualified domain name of the FlashBlade in the address bar and press **Enter**. The Purity//FB GUI login screen appears.
3. In the **Username** field, type the FlashBlade user name. For example, `pureuser`.
4. In the **Password** field, type the password for the FlashBlade user. For example, `pureuser`.
5. Click **Log In** to log in to the Purity//FB GUI.

Viewing and Accepting the End User Agreement

Once the FlashBlade array has been installed, you can read and accept the terms and conditions of the End User Agreement (EUA). Note that the EUA is only presented upon the first login.

Note: Only users with **array_admin** permissions can accept and update the EUA. Users without **array_admin** permissions can only view the EUA.

Upon initial login of the GUI, a pop-up window appears displaying the EUA. Carefully review the EUA before proceeding. To accept the EUA, click **Accept**. Once the EUA has been accepted, the next time that the EUA is accessed, the **Accept** button is replaced with a **Close** button along with the acceptance date and time displayed.

You can view the EUA at any time by clicking **End User Agreement** from the navigation pane to open the EUA pop-up window.

If you do not wish to accept the EUA at this time, or have negotiated a separate EUA with Pure Storage, click **Close**. The **Close End User Agreement** pop-up window appears with information about the continued usage of Pure Storage products. Click **Close** to close the pop-up window and proceed with your GUI tasks.

Chapter 4
The Purity//FB GUI Dashboard Page

The **Dashboard** page displays a running graphical overview of the array's storage capacity, performance, and hardware status.

Figure 4.1 Dashboard

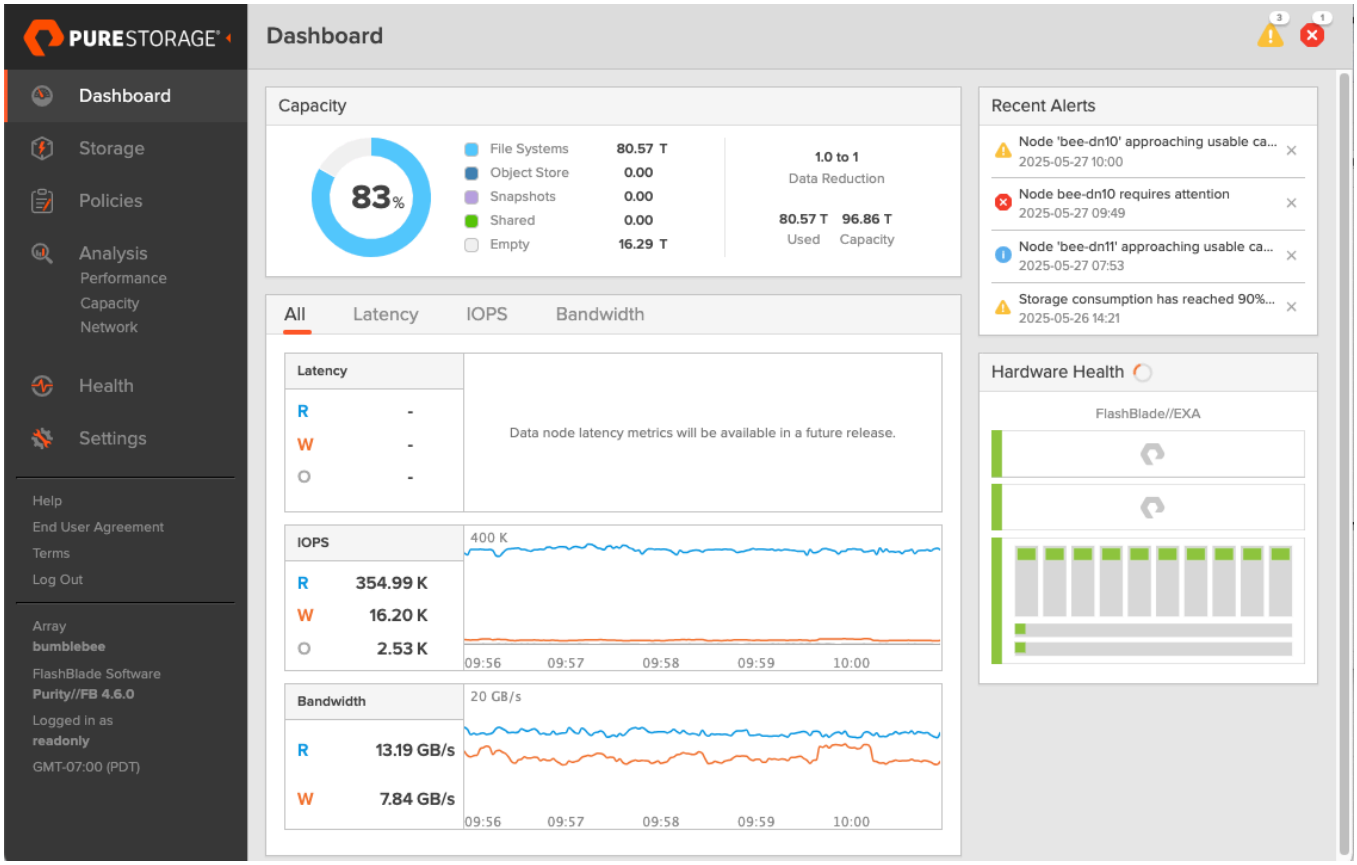

The **Dashboard** page contains the following panels and charts:

- **Capacity**
- **Performance Charts**
- **Recent Alerts**
- **Hardware Health**

Capacity

Capacity is displayed in the Dashboard in the **Capacity** panel.

The capacity wheel displays the percentage of array space occupied by file system, object store, snapshots, and shared data. The percentage value in the center of the wheel is calculated as **Used** capacity/**Total** capacity in tebibytes (T).

The capacity data is broken down into the following components:

- **File Systems**: Amount of space that the written data occupies on the array 's file systems after reduction via data compression.
- **Object Store**: Amount of space that the written data occupies on the array 's buckets after reduction via data compression.
- **Snapshots**: Amount of space that the array 's snapshots occupy after data compression.
- **Shared**: Physical space occupied by deduplicated data, meaning that the space is shared with other filesystems/buckets and snapshots as a result of data deduplication.
- **Empty**: Unused space.
- **Data Reduction**: Ratio of the size of the written data (i.e. Used) versus the amount of space the data occupies after data compression (i.e. Physical) and/or deduplication.
- **Used**: Amount of space that the written data occupies on the array after reduction via data compression.
- **Capacity**: Total usable capacity of the array.

Purity DeepReduce™

Purity DeepReduce™ is a data reduction feature running Global Nearline Similarity based algorithms to generate higher data reduction rates in the given dataset on FlashBlade//E systems. Built to handle today's surge of pre-reduced, immutable, and application-compressed datasets, this advanced similarity-based reduction engine helps unlock meaningful efficiency gains across challenging workloads. Engineered for reduction at scale, **DeepReduce**™ seamlessly augments as a background process and works in conjunction with FlashBlade//E's existing inline compression technologies to maximize effective capacity. With **significantly higher data reduction than conventional deduplication**, customers can now achieve greater efficiency, lower TCO, and more predictable storage economics at scale. Progress

and effectiveness can be monitored, with metrics such as percentage of data blocks scanned by DeepReduce™ and achieved data reduction ratios reported.

Performance

The performance panel displays IOPS and bandwidth values in real time.

Figure 4.2 Dashboard - Performance Charts

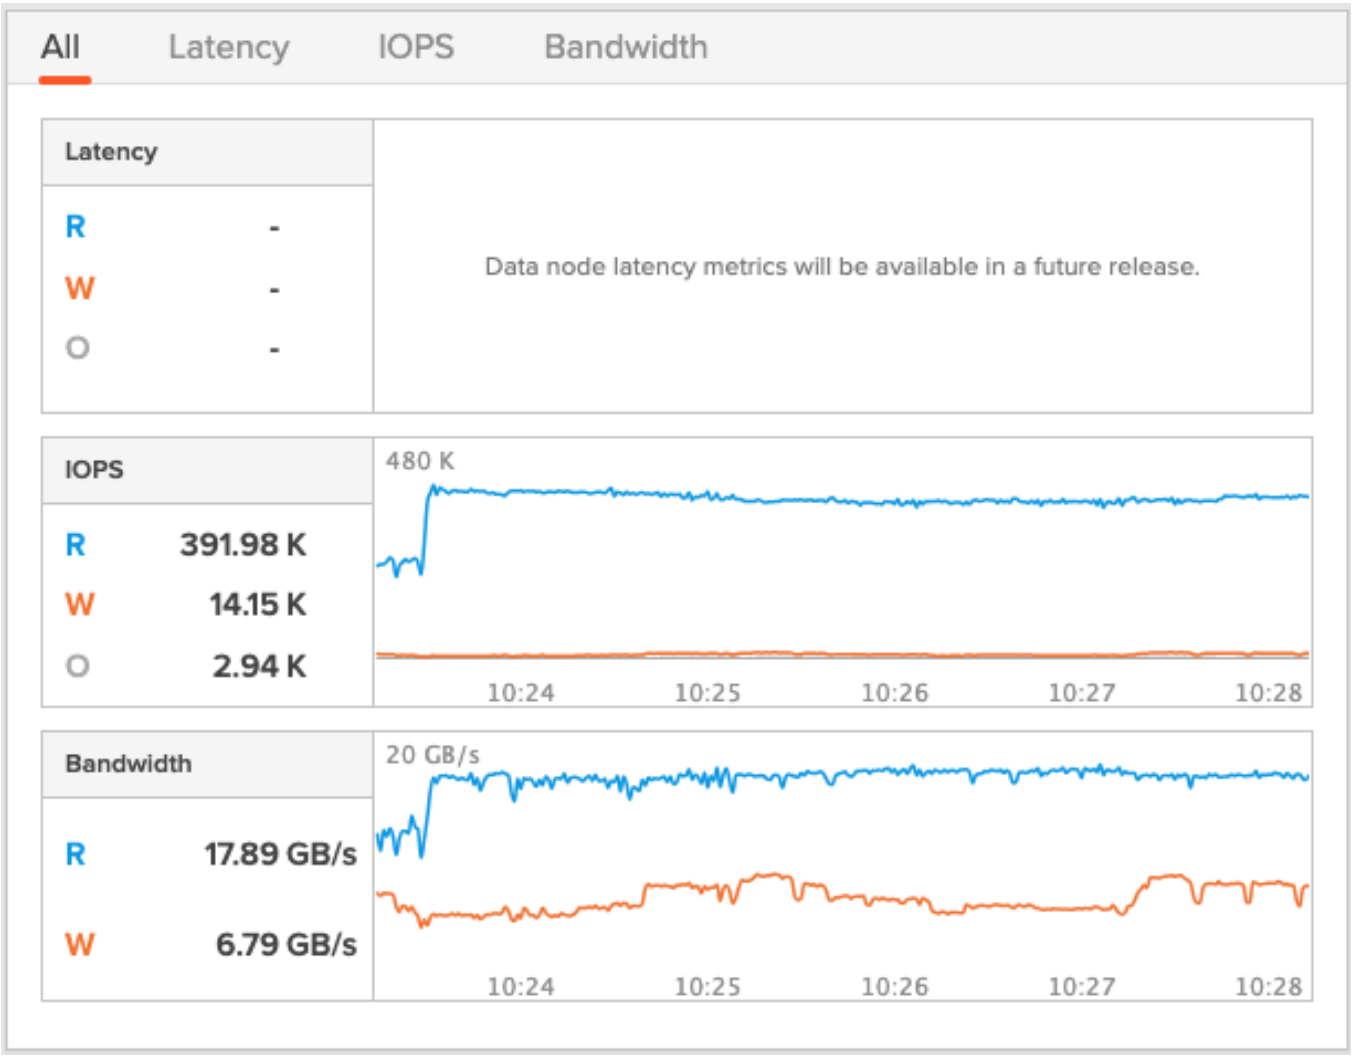

The performance metrics are displayed along a scrolling graph; incoming data appears along the right side of each graph every second as older data drops off the left side after 5 minutes. Each performance chart includes R, W, and O (if applicable) values, representing the most recent data samples.

Hover over any of the charts to display metrics for a specific point in time. The values that appear in the point-in-time tooltips are rounded to two decimal places.

The performance panel includes Latency, IOPS, and Bandwidth charts. The **All** chart displays all three performance charts in one view.

Note: The Latency chart will be available for data node latency metrics in a future release.

IOPS

The **IOPS** (Input/output Operations Per Second) chart displays I/O requests processed per second by the array. This metric counts requests per second, regardless of how much or how little data is transferred in each.

- **Read IOPS (R)** - Number of read requests processed per second.
- **Write IOPS (W)** - Number of write requests processed per second.
- **Other IOPS (O)** - Number of metadata operations processed per second.
- **Average IO Size** - Average I/O size per request processed. Requests include reads and writes.

Bandwidth

The **Bandwidth** chart displays the number of bytes transferred per second to and from the array (both file systems and buckets) . The data is counted in its expanded form rather than the reduced form stored in the array to truly reflect what is transferred over the storage network. Metadata bandwidth is not included in these numbers.

- **Read Bandwidth (R)** - Number of bytes read per second.
- **Write Bandwidth (W)** - Number of bytes written per second.

By default, the performance charts display performance metrics for the past 5 minutes. To display more than 5 minutes of historical data, select **Analysis > Performance**.

Recent Alerts

The **Recent Alerts** panel displays a list of alerts that Purity//FB saw that is both flagged and not in the closed state. The list contains recent alerts of all severity levels. If no alerts are logged, the panel displays **No recent alerts**.

Figure 4.3 Dashboard - Recent Alerts Panel

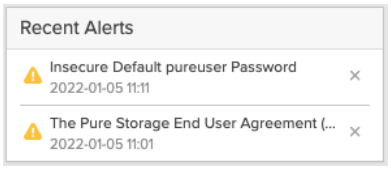

To view the details of an alert, click the alert message.

To remove an alert from the **Recent Alerts** panel, click the clear flag (**X**) button. The alert will no longer be displayed in the **Recent Alerts** panel, but will still appear on the **Health** page.

To view a list of all alerts, including ones that are in no longer open, go to the **Health** page.

Hardware Health

The **Hardware Health** panel displays the operational state of the FlashBlade array chassis, blades, and fabric modules.

Figure 4.4 Hardware Health Panel - FlashBlade//EXA

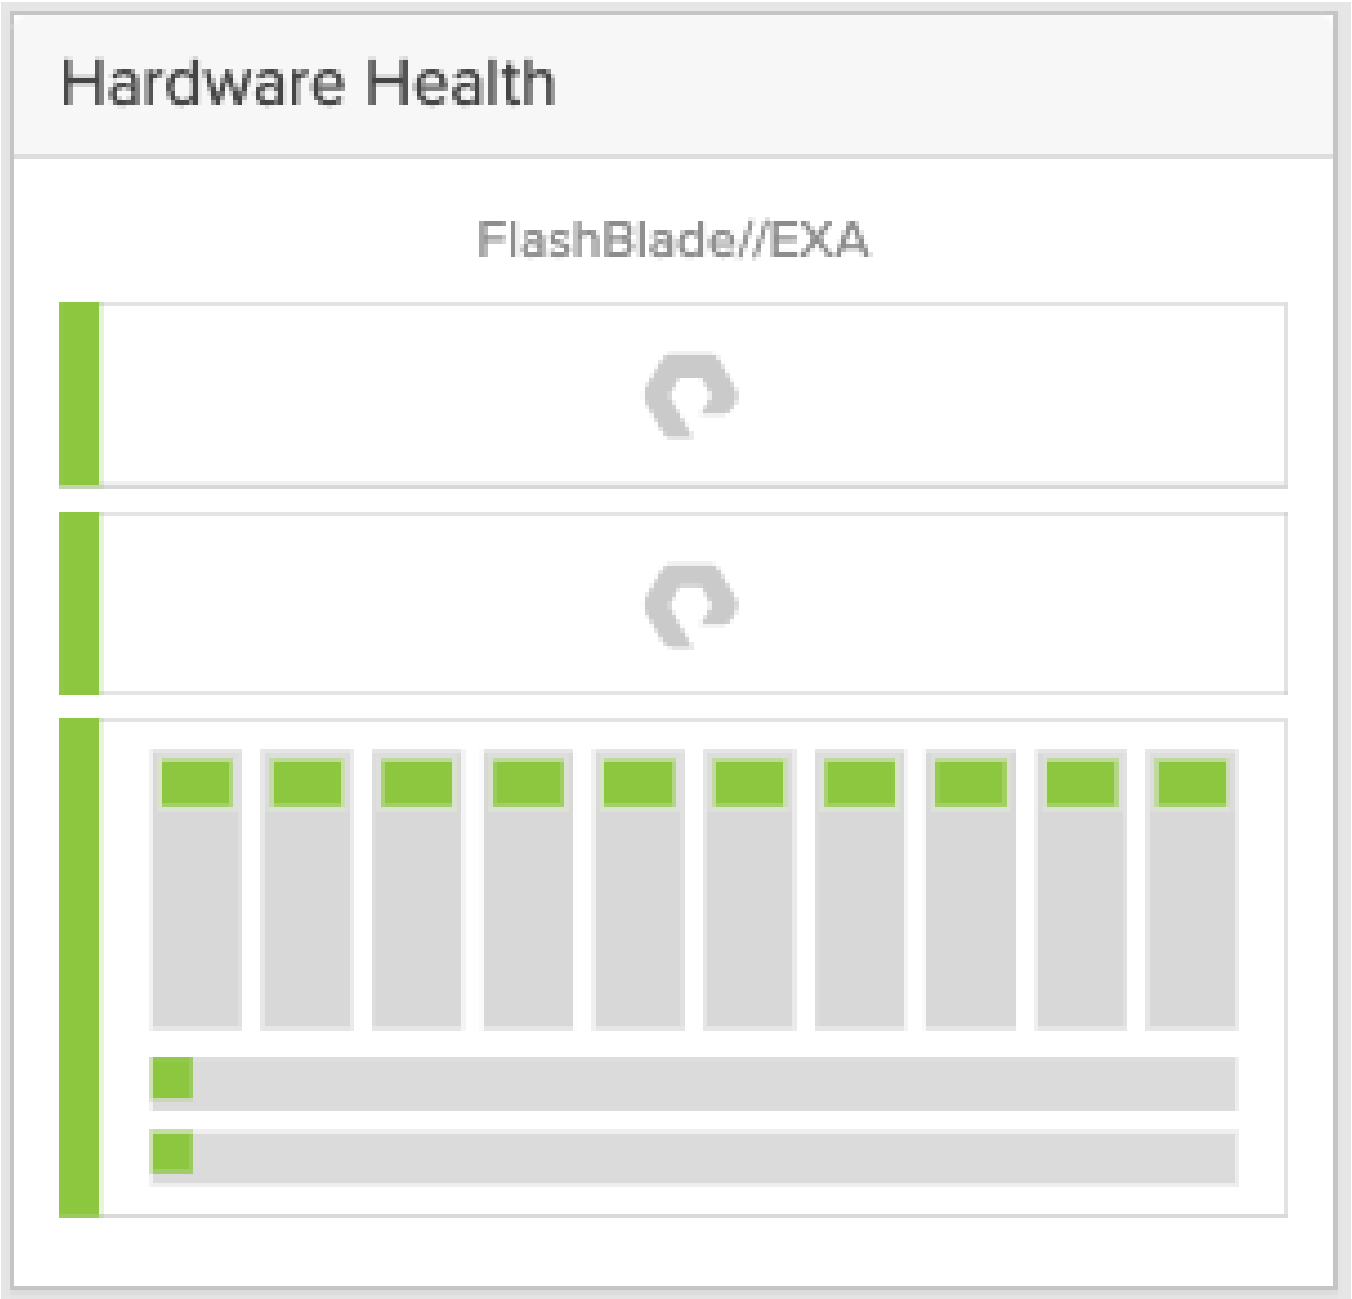

To analyze the hardware components in more detail, click the **Health** link.

Chapter 5
The Purity//FB GUI Storage Page

The **Storage** page displays and manages the array 's servers and file systems.

Accessing a file system requires at least one valid data virtual IP address (data vip). Data vips are configured through the **Settings** > **Network** page.

Figure 5.1 Purity//FB Storage Page

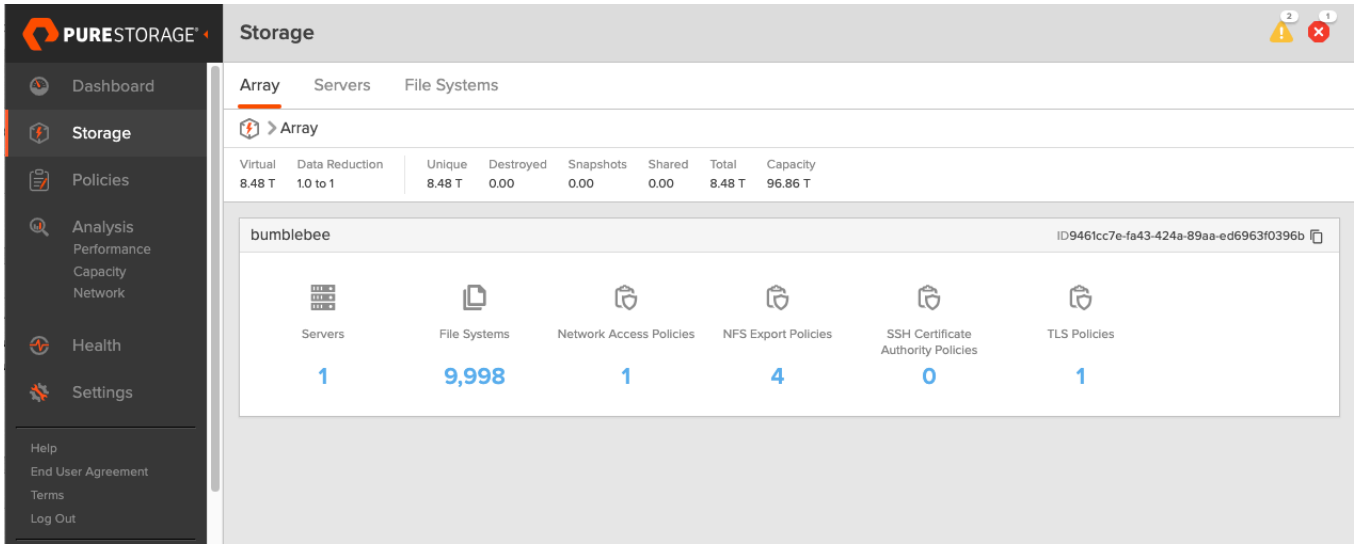

Array

By default the **Array** tab is displayed when navigating to the **Storage** page.

The **Array** tab displays an array summary panel.

Figure 5.2 Array Tab

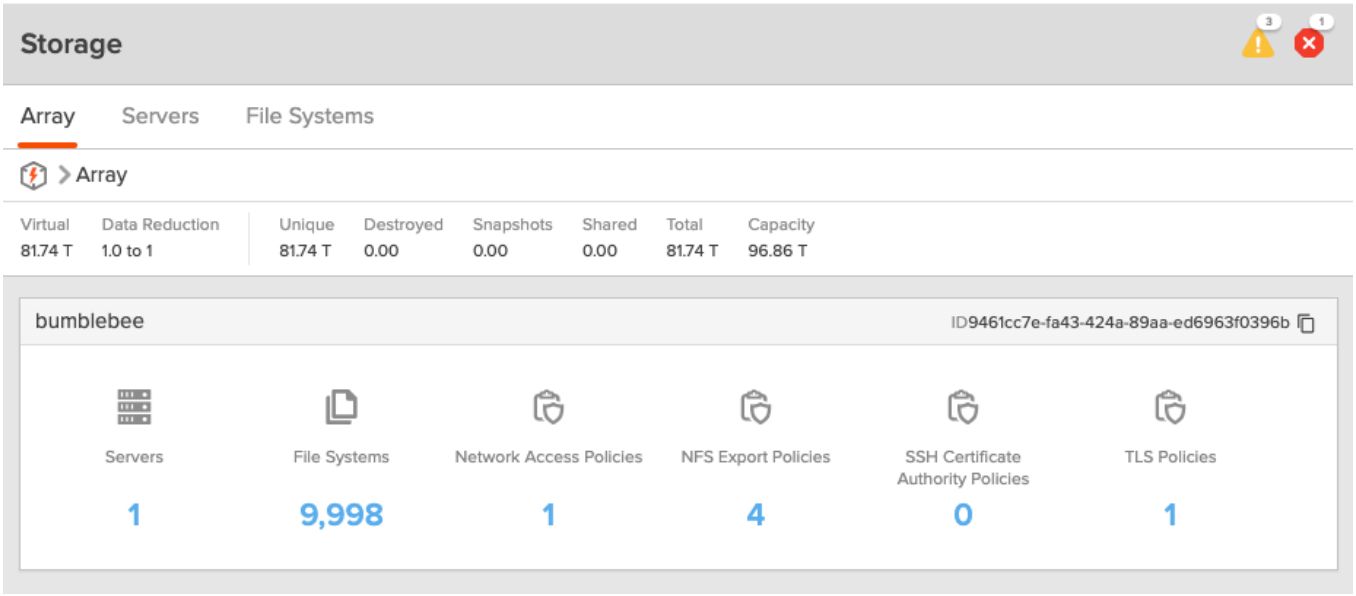

The array summary panel displays counts of various types of metadata objects. Click a count to jump to the page containing the details for associated objects. For example, clicking the **File Systems** count opens the **Storage** > **File Systems** tab in the GUI.

Servers

The **Servers** tab displays information for all servers on the array.

Note: IPv6 is not supported with this release for multi-tenancy file access and file systems with NFS and SMB access enabled need to have access to UID/GID sources via LDAP or AD.

Note: Egress traffic for DNS/AD/LDAP/Kerberos/SMB-Audit-Syslog-Server goes over the management network with an exception where the external server IP and FlashBlade datavip is in the same subnet.

A server is deployed with each FlashBlade. They contain network and identity information used to access data on the FlashBlade. Contained within each server is a network interface, NFS Directory Services configuration, up to one Active Directory account, a DNS configuration, and list of exported file systems accessible via the server. There is a single default server on all FlashBlades named "_array_server". Additional servers can be added to the FlashBlade to implement multi-tenancy. Multi-tenancy allows users, or tenants, to share a FlashBlade. Each tenant will have access to their own private data with their dedicated server.

The **Servers** panel displays configured servers used for mult-tenancy. The **_array_server** is a default server on the array.

Note: Server fully qualified names needs to be unique on the array. This is unlikely to happen often as most AD and LDAP domains are unique to that group and will not conflict unless two Servers are using a common name and domain name like **test.local**

Servers Panel

By default general information about each server is displayed, which includes the following information:

- **Name**: The server name. Note that **_array_server** is the array's default server.
- **DNS**: The DNS service that will be used the Server.

Clicking any server directs you to that server's details page.

In addition to displaying servers, the **Servers** panel allows you to create and configure servers.

Server Details

A server's details page is accessed by clicking a server from the **Servers** panel.

The server detail page provides the following panels:

- **File System Exports**: Lists all exported file systems.
- **Interfaces**: The IP(s) and subnet(s) that will be used for the clients that will be accessing this server.
- **DNS Settings**: DNS setting to be used by this Server.
- **Directory Service**: The server's Directory Service. Configuring a Directory Service is optional for a server.
- **Active Directory Account**: Optional Active Directory Account. Required for any SMB client access.

Creating a Server

To create a server for multi-tenancy, perform the following:

Note: IPv6 is not supported with this release for multi-tenancy file access and file systems with NFS and SMB access enabled need to have access to UID/GID sources via LDAP or AD.

Note: Egress traffic for DNS/AD/LDAP/Kerberos/SMB-Audit-Syslog-Server goes over the management network with an exception where the external server IP and FlashBlade datavip is in the same subnet.

1. From the **Storage** page navigation sidebar, click **Servers** tab.
2. In the **Servers** panel, click the **Add** interface button (+). The **Create Server** pop-up window appears.
3. In the **Name** field, enter a name for the server.
4. Click **Create**.

Exporting File Systems

⚠ **Important!** Beginning with Purity//FB 4.5.2, case sensitivity is enforced for all NFS file systems.

- All NFS mounts require a case-sensitive client mount command that precisely matches the file system's share name as displayed in the FlashBlade's **`purefs list`** output for NFS shares.
- Before upgrading to Purity//FB 4.5.2, audit and update all client-side configurations, including entries in /etc/fstab, automount maps, and custom deployment scripts, to ensure exact case-matching of the NFS share name to prevent disruption.

To export file systems for multi-tenancy, perform the following:

⚠ **Important!** If you require the same data to be accessible by NFS and SMB, you must create an export for each protocol.

1. From the **Storage** page navigation sidebar, click **Servers** tab.
2. Select the server of which you are exporting the file system to and the details page of that server appears.
3. In the **File Systems Exports** panel, click the **Add** interface button (+). The **Create File System Export** pop-up window appears.
4. In the **File System** field, enter the name of the file system you wish to export.
5. In the **Export Name** field, enter a name for the exported file system, which is the name the clients will use to connect to the file system.
6. In the **Type** field, select **NFS** or **SMB**, which is the type of policy the file system is attached.
7. (When the NFS type is selected) In the **NFS Export Policy** list, select the policy.
8. (When the SMB type is selected) In the **SMB Client Policy** list, select the policy
9. (When the SMB type is selected) In the **SMB Share Policy** list, select the policy.

10 Click **Create**.

Editing a File System Export

To edit a file system export, perform the following:

1 Select **Storage >Servers**.

2 Click the server used for export.

3 Click the **Edit** button in the row of the file system export you want to edit. The **Edit File System Export** pop-up window appears.

4 Enter a new export name.

5 Select a policy to attach.

6 Click **Save**.

Deleting a File System Export

To delete a file system export, perform the following:

1 Select **Storage** > **Servers**.

2 Click the server used for export.

3 To delete a single file system export:

- **a** Click the **Delete** button in the row of the file system export you want to delete. The **Delete File System Export Configuration** pop-up window appears.
- **b** Click **Delete**.

4 If you wish to delete multiple file system exports:

- **a** Click the **More Options** button in the header of the **File Systems Exports** panel and select **Delete**. The **Delete File System Exports** pop-up window appears.
- **b** From the **Exsting File System Exports** list, select the file systems you wish to delete. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all file systems by clicking the checkbox next to the search box. All selected file systems appear in the **Selected File System Exports** list.
- **c** Click **Delete**.

File Systems

The **File Systems** tab displays information for all file systems on the array.

The **File Systems** panel displays configured file system information while the **Destroyed File Systems** panel displays destroyed file systems that are pending eradication.

File Systems Panel

By default **General** information about each file system is displayed, which includes the following information:

- **Name**: File system name. Clicking a file system name displays any snapshots belonging to that file system, and also provides the option of creating snapshots.
- **Source Location:** The location where the file system was created.
- **Source Name:** If created on the target array, the name of the source file system.
- **Size**: Provisioned size of the file system. If not set, the provisioned size is unlimited.
- **Virtual**: Total amount of pre-compressed data written to the file system.
- **Hard Limit**: Whether the file system's provisioned size has a hard limit (true) or not (false).
- **Created**: The file system creation date.
- **Protocols**: One or more protocols configured for the file system. A dash (-) means the file system has no protocols assigned.
- **Promotion Status**: Whether the source or replicated file system is promoted or demoted.
- **Writable**: Whether the file system is writeable (true) or not (false). When the file system is changed from writable disabled to enabled, it is recommended to remount the file system to avoid stale client cache issues.

The size, rules, and protocols of each file system can be configured.

Click **Space** to view file system space information. The columns in the **File Systems** panel display the following information:

- **Name**: File system name. Clicking a file system name displays any snapshots belonging to that file system, and also provides the option of creating snapshots.
- **Size**: Provisioned size of the file system. If not set, the provisioned size is unlimited.
- **Virtual**: Total amount of pre-compressed data written to the file system.
- **Available**: The available space of the file system.
- **% Available**: The percentage of available space.
- **Data Reduction**: RRatio of the size of the written data (i.e. Used) versus the amount of space the data occupies after data compression (i.e. Physical) and/or deduplication.
- **Unique**: The unique physical space occupied by customer data.
- **Snapshots**: Amount of space occupied by snapshots after reduction from data compression.
- **Total**: The total used space (physical space and snapshots).

You can search the file systems list by entering search criteria in the search box under each column.

Clicking any file system directs you to that file system's details page.

In addition to displaying file systems, the **File Systems** panel allows you to create and maintain file systems.

File System Details

A file system's details page is accessed by clicking a file system from the **File Systems** panel.

Figure 5.3 File System Details Page

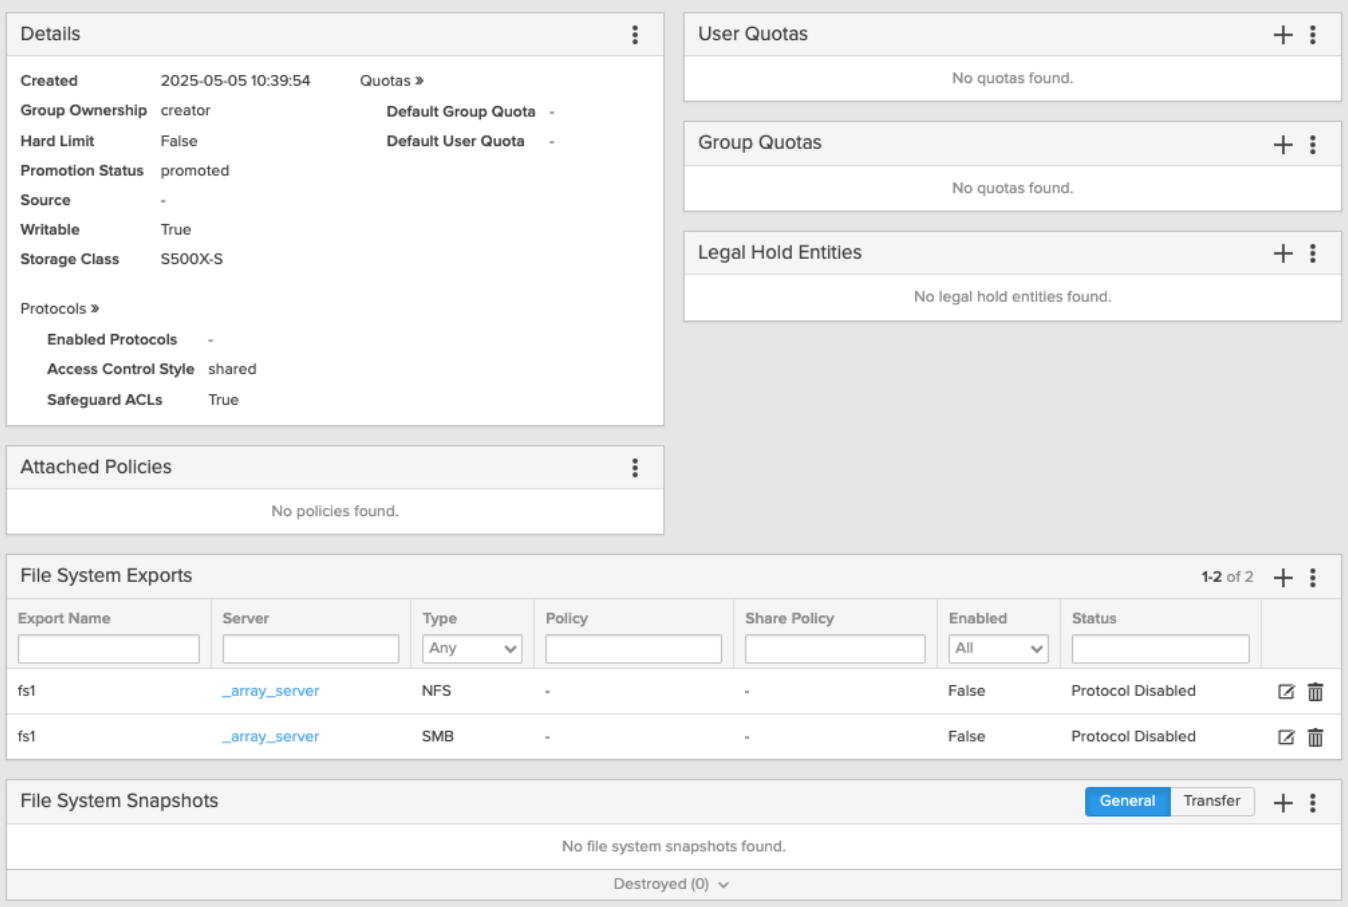

The file systems detail page provides the following panels:

- **Details**: Provides an overview of the file system's configured parameters.
- **File System Exports**: Lists all file system exports.

File Systems Exports Panel

The **File Systems Exports** panel displays the following information:

- **Export Name**: Exported file system name.
- **Server**: The server used for multi-tenancy. The default server on the array is **_array_server**.
- **File System**: The file system name.
- **Type**: The directory service type, NFS or SMB.

- **Policy**: The NFS export policy. See *NFS Export Policies* for additional information.
- **Share Policy**: The SMB share policy. See *SMB Share Policies* for additional information.
- **Enabled**: Whether the export is enabled (**True**) or not (**False**).
- **Status**: The current export status.

File Locks Panel

Refer to *Viewing and Managing File Locks*.

Destroyed File Systems Panel

The **Destroyed File Systems** panel is located directly below the **File Systems** panel. Clicking the arrow beside the **Destroyed** heading expands the panel.

Once a file system has been destroyed, it is moved to the **Destroyed File System** panel making it inaccessible to clients. This is also known as an eradication pending period. During this period the file system and snapshot is left intact for 24 hours. If the destroyed file systems are not recovered within the 24-hour period, they will be eradicated and their space will be reclaimed by the array. If you need to reclaim space before the 24 hour expiration, file systems can be manually eradicated.

- **Name**: File system name.
- **Size**: Provisioned size of the file system.
- **Virtual**: Total amount of pre-compressed data written to the file system.
- **Hard Limit**: Whether the file system's provisioned size has a hard limit (true) or not (false).
- **Unique**: The unique physical space occupied by customer data.
- **Time Remaining**: Indicates the file systems expiration period in the eradication pending period. Time is represented in hours (**h**) and minutes (**m**) . When time expires the file systems will be eradicated and their space will be reclaimed by the array.
- **Created**: The create date and time of the file system.

If a file system is destroyed, its snapshots are also destroyed and will appear in the **Destroyed Snapshots** panel. Those snapshots cannot be recovered unless the file system is recovered.

Pure File SafeMode

Pure File SafeMode is configured during installation and is designed to protect the system from accidental changes. File SafeMode prevents users from performing some operations such as file system snapshot restore and file system snapshot eradication. For more information, contact Pure Storage Technical Services.

About Creating and Exporting File Systems

Warning: Samba adapter was deprecated in Purity//FB 4.3.0. If file systems were created using Samba adapter, ACLs contain UID/GID instead of SIDs, meaning domain information is dropped. Migrating the FlashBlade system to another primary domain may lead to unauthorized access being granted to users with the same RID (Relative Identifier) from that domain, and is not recommended.

Creating and exporting a file system involves creating the file system and enabling protocol support for the file system. FlashBladefile systems support the Hypertext Transfer Protocol (HTTP), the Network File system (NFS) and Server Message Block (SMB) protocols.

Exporting a file system requires at least one data virtual IP address (data vip). The data vip can be created before or after creating the file system. Data vips are managed through the **Settings** > **Network** > **Subnets** panel. For more information about data vips, refer to *Subnets and Interfaces*.

For more information about using directory services with file system protocols, refer to *Directory Services Overview*.

When creating a file system, consider which protocol(s) you want to configure:

- **NFS**. File sharing over NFS requires that you configure the NFS export rules and then enable the protocol. Per NFS limitations, each user can belong to a maximum of 16 groups. If users belong to more than 16 groups and a directory service contains the group information, configure file sharing over NFS with a directory service.

  Follow these steps to set up file sharing over NFS without a directory service:

  1. Create the data vip. Data vips are managed through the **Settings** > **Network** > **Subnets** panel.
  2. Create the file system. Optionally set a provisioned size for the file system.
  3. Define the NFS export rules.
  4. Enable the NFS protocol.

- **NFS with a directory service**. For file sharing over NFS where a user might belong to more than 16 groups, configure the FlashBlade directory service to integrate with a directory server or files will be accessed via SMB and NFS at the same time. This assumes that you already have a directory service, such as OpenLDAP or Active Directory set up.

  Follow these steps to set up file sharing over NFS with a directory service:

  1. Configure the directory service for protocol support. The directory service is managed through the **Settings** > **Security** > **Directory Service** panel.
  2. Create the data vip. Data vips are managed through the **Settings** > **Network** > **Subnets** panel.
  3. Create the file system. Optionally set a provisioned size for the file system.
  4. Define the NFS export rules.
  5. Enable the NFS protocol.

- **SMB with Active Directory**. FlashBlade offers SMB Native authentication mode. It is recommended for customers requiring Kerberos authentication.

  The default SMB authentication mode is SMB Native. Follow these steps to set up file sharing over SMB in Native mode:

  1. Prepare Active Directory for multi-protocol support.

     Prepare Active Directory for multi-protocol support by setting the `gidNumber` and `uidNumber` attributes for each domain user and group that will be accessing the SMB shares. The `gidNumber` and `uidNumber` attributes can be set through the Microsoft Identity Management for UNIX service or PowerShell.

  2. Configure the directory service for protocol support.

     The FlashBlade array must be integrated with an Active Directory service. The directory service is managed through the **Settings** > **Security** > **Directory Service** page. For more information about the directory service, refer to *Directory Services Overview*.

  3. Create the data vip. Data vips are managed through the **Settings** > **Network** > **Subnets** panel.
  4. Create the file system. Optionally set a provisioned size for the file system.

- **HTTP access**. Enabling the HTTP protocol turns on HTTP and HTTPS access to a file system.

  Before you enable the HTTP protocol, consider the following access control limitations:

- All HTTP APIs by default are executed on behalf of a root user.
- Certain HTTP APIs provide the ability to specify a user by specifying the UID as part of the request.
- Access checking in HTTP is based solely on the client's crafted request. Any user can read, modify, and delete any files on a file system when HTTP is enabled.
- The use of HTTP connections, as opposed to HTTPS connections, potentially gives unauthenticated users access to the file system. Anyone with network access to the data vip has access to the file system.

Follow these steps to set up file system for HTTP access:

1. Create the data vip. Data vips are managed through the **Settings** > **Network** > **Subnets** panel.
2. Create the file system. Optionally set a provisioned size for the file system.
3. Enable the HTTP protocol.

File System Size

File system size represents the provisioned size allocated to a file system. The size is a quota of space that is set to help gauge the fullness of a file system.

When the amount of data written to the file system reaches a certain percentage of its quota, alerts are generated notifying administrators of the threshold reached. For example, Purity//FB generates a `WARNING` alert when the amount of data written to the file system reaches 90% and then 100% of its provisioned size. This behavior differs slightly if a hard limit is set on the file system. When the file system usage reaches 90% of its provisioned size, the system generates a `WARNING` alert. When the file system usage reaches 100% of its provisioned size, the system generates a `CRITICAL` alert. See **File System Usage Limits** for more information. To view alert details, select **Health** > **Alerts**.

The file system size can be blank or set to any value between `0` and `8191`P. If the field is left blank or set to `0`, the provisioned size will default to an unlimited size.

Purity//FB does not enforce any restrictions nor does it take any further action when space consumption reaches or exceeds the provisioned size. To address the issue, decrease the percentage of storage space used by either deleting files or increasing the size of the file system. In contrast, if a hard limit is set for the file system, all writes are halted once the provisioned size is reached. See **File System Usage Limits** for details.

File system sizes can be changed at any time.

File System Usage Limits

⚠ **Important!** File system hard limits cannot be enabled for file systems with unlimited sizes.

To help manage resource usage, a hard limit can be enabled on the provisioned size of a file system. Enabling a hard limit effectively limits writes of the file system to its provisioned size. The hard limit can be enabled when creating or editing a file system.

When the system detects the file system has exceeded its provisioned size, all future writes are halted until usage decreases below the provisioned size. This occurs within minutes of the provisioned size being reached. Note that the provisioned size can be exceeded during this detection period when a hard limit is enabled. For example, a 5G write is in-process for a file system whose provisioned size is 10G and current usage is 9G. When the system halts the write, the file system's provisioned size will be slightly exceeded. Data writes only resume after data has been managed (deleted or truncated) to below the file system's provisioned size.

A hard limit cannot be enabled if the file system usage has already exceeded its provisioned size unless forced. You can force the hard limit by using the **Edit File System** dialog box from the **File Systems** screen

Additionally, if a hard limit is enabled, the file system's size cannot be reduced to a value less than its current space usage. You can force shrink the file system's size by using the **Edit File System** dialog box from the **File Systems** screen.

You can disable a hard limit from the **Edit File System** dialog box on the **File Systems** screen.

User and Group Usage Limits

Similar to file system hard limits, hard limits (quotas) can be set for users and groups on a per- file system basis to help manage file system resource usage.

⚠ **Important!** User and group quotas are not supported for SMB. If quotas are set on the file system, all writes over SMB are charged against a single NOBODY user and NOBODY group. For additional information, contact Pure Technical Services.

When the system detects that a user or group has exceeded its quota, all future writes to the file system are halted until usage decreases below the set quota size via deletion or truncation of that user or group's data from the file system. Note that a quota can be exceeded during this detection period. For

example, if a user is writing a large file that causes them to exceed the quota, their write will fail, but their total usage may have slightly exceed their set quota.

During file system creation, a default group quota and/or user quota can be set, which is applied to all users and/or groups in the file system. Additionally, once a file system has been created, quotas can be set per individual user ID or group ID (quotas can be set explicitly) for that file system. Note that explicit quotas take priority over default quotas. For example, if a file system was created with default user or group quota and then an explicit quota was set for a specific group or user of that file system, the explicit quota is used. If that explicit quota is later deleted, the default quota then takes effect.

User and group quotas are displayed on the **Details**, **User Quotas**, and **Group Quotas** panels from a file system 's detail screen. To access the detail screen, from the **File Systems** panel, click the file system whose details you wish to view.

File System Group Ownership

Group ownership can be configured when creating and editing file systems.

By default, files and directories in a file system will follow the file system's creator's group, where the owning group of all newly created files and folders is the same as the creator's primary group. File systems can be optionally configured to inherit the parent directory's group ID (GID). This can be useful for supporting BSD clients on FlashBlade since new files and directories in BSD systems are given the group of the directory, which contains them by default. This can also be useful in cases where you wish to migrate data from existing storage that uses BSD-style GID inheritance to FlashBlade and retain the GID interactions.

When creating and editing file systems, to configure group ownership you are provided with the following options:

- **Creator**: The owning group of new files and directories is the primary group of the user who creates them. This is the default setting.
- **Parent Directory**: The owning group of new files and directories is the primary group of the parent directory. New files and directories inherit the parent directory's group ID.

The **Details** panel of a file system's detail screen reflects the file system's group ownership.

Figure 5.4 Viewing a File System's Details Panel

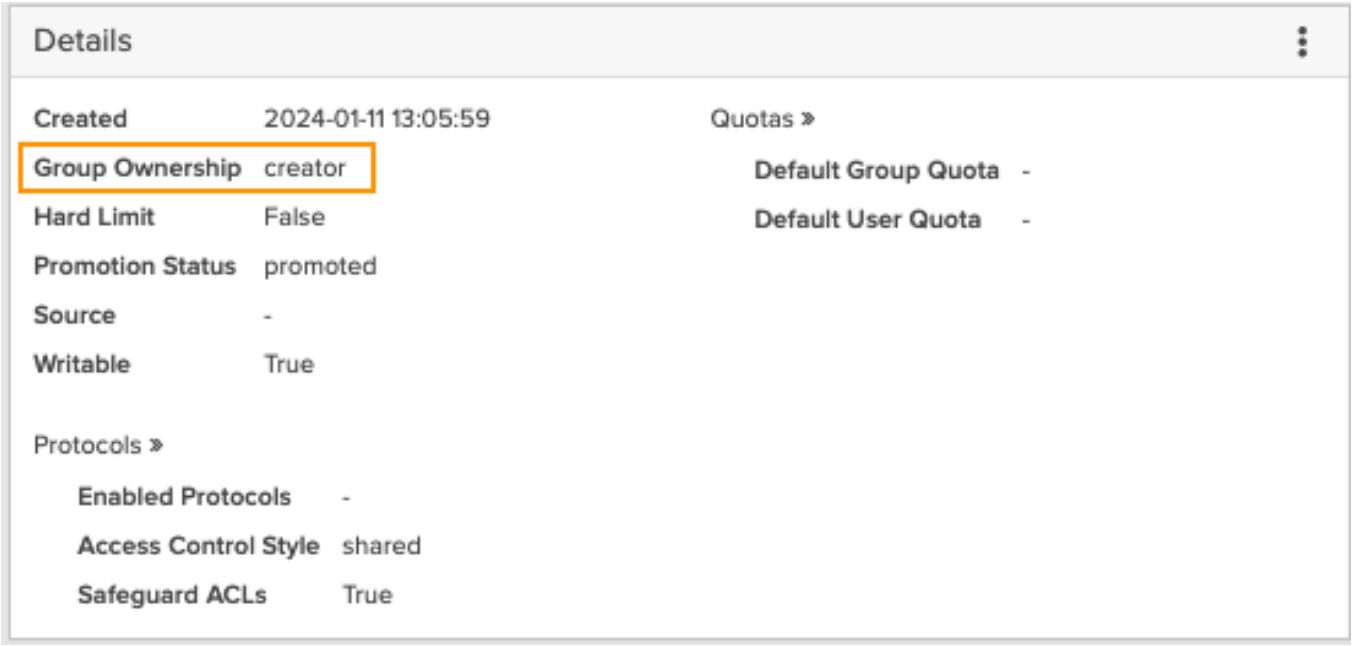

ACLs

An Access Control List (ACL) is a list of security permissions associated with a directory or file. Specific rules can be added to the ACL to define who can access the file and what actions they can perform. When an ACL is applied to a directory or file, then that directory or file will reference its list of rules to determine whether to grant or deny access, including which actions are allowable, such as; read, write, and execute.

ACEs

An ACL consists of up to 1,820 Access Control Entries (ACEs), plus three OGE (owner, group, everyone) entries. Note that certain clients may have lower limits from what the server supports.

Each ACE contains detailed permissions information. ACEs consist of four fields separated by colons — type, flag, principal, and permissions.

- **type** - if the defined permissions are allowed (A) or denied (D).
- **flag** - defines both child object inheritance and if the ACE is defining group permissions.
- **principal** - defines to whom the permissions apply.

- **permissions** - defines the permissions that the specified principle will be allowed or denied.

For example, the ACE `A:d:user@purestorage.com:R` means that the user `user@purestorage.com` is allowed (`A`) read permissions (`R`) and that new sub-directories will inherit this ACE (`d`).

Access Control Styles

For implementations where more than one file system export protocol is enabled, Purity//FB 3.1.1 and later allows you to set the file system's access control style to dictate how the protocols will interact.

The following access control styles are available:

- **Shared** - NFS and SMB clients are both able to set permissions on, and access files and directories over any protocol.
- **NFS** - Only NFS clients can set permissions on files and directories; however, files and directories can be accessed over any protocol.
- **SMB** - Only SMB clients can set permissions on files and directories; however, files and directories can be accessed over any protocol.
- **Independent** - NFS and SMB clients are both able to set permissions on files and directories and can access files and directories created over any protocol. Permissions set by SMB clients will not affect NFS clients and vice versa. NFS clients will be restricted to only using mode bits to set permissions.
- **Mode Bits** - NFS and SMB clients are both able to set permissions on files and directories, but permissions are enforced in mode bits-only format. When SMB clients set an ACL it will be converted to mode bits.

Note the following behaviors and restrictions:

- If you set the access control style to **SMB** and then disable the SMB protocol, the access control style is automatically reset to **Shared**.
- If you set the access control style to **NFS** and then disable both NFSv3 and NFSv4.1, the access control style is automatically reset to **Shared**.
- If neither the NFSv3 or NFSv4.1 protocols are enabled, the access control style cannot be set to **NFS**.
- If the SMB protocol is not enabled, the access control style cannot be set to **SMB** or **Independent**.

File system attributes that were set prior to Purity//FB 3.1.1 will be transitioned to the following access control styles as described in the following table with Purity//FB 3.1.1 and later:

Table 5.1 Access Control Style Transitions

Pre Purity//FB 3.1.1 File System Attribute	Purity//FB 3.1.1 and Later Access Control Style
n/a	SMB
n/a	NFS
n/a	Shared
SMB native, NFSv3, NFSv4.1	Independent
SMB shared	Mode Bits

ACL Safeguarding

If using the **Shared** or **NFS** access control styles, consider enabling ACL safeguarding.

When using the SMB access control style, ACL safeguarding is automatically enabled.

ACL safeguarding controls how a mode bit change will impact the ACL on the file or directory. When disabled, modification of mode bits will result in the removal of any existing ACL. When enabled, modification of mode bits will result in the modification of the existing ACL to merge the permissions from the mode bits with the permissions in the ACL.

Note: ACL safeguarding is not available for **Independent** or **Mode Bits** access control styles.

Setting the access control style and enabling ACL safeguarding are performed when creating a file system and can be modified at a later date.

Group Policy Objects

A group policy allows administrators to define security policies for users. A Group Policy Object (GPO) is a collection of group policy settings. When joining an Active Directory (AD) domain, FlashBlade applies two specific group policies as defined by the GPO within the AD domain: backup operator and the privilege to set system access control lists (SACLs). GPOs, group policies and their privileges are configured on the client. FlashBlade resynchronizes with AD and updates its GPO in 30 minute intervals.

- **backup operator**: Allows select users full read access to file systems in order to create backup copies of them. This grants access to a file regardless of if the user has the appropriate Access

Control Element (ACE) set in the Access Control List (ACL). Full read/write access for this policy requires the SeBackupPrivilege and SeRestorePrivilege privileges. Note that this provides identical permissions to the `fs.smb_backup_operators` tunable (tunable enabled by Pure Technical Services).

- **privilege to set up SACLs**: Allows select users to perform security filtering based on SACLs attached to file system folders and files that are served by FlashBlade. Access is granted with the SeSecurityPrivilege privilege.

Multi-Domain Trust

Multi-Domain Overview

A multi-domain environment consists of two or more domains that communicate with each other. These domains are made up of trees and forests. A tree is a collection of one or more domains that share a common namespace. For example, *engineering.flashblade.com* and *it.flashblade.com* share the same domain name of *flashblade.com*, therefore are two domains of the same tree. A forest is a group of trees, with each tree having a different domain names. For example, *engineering.flashblade.com* and *it.flashblade.com* are one tree and engineering.fb.com and it.fb.com are in a different tree, therefore are two trees in a forest domain consisting of *flashblade.com* and *fb.com*.

Multi-domains allow customers to effectively manage more than one internet or intranet domain with different namespaces. FlashBlade supports direct communication with domains it has joined, called “joined domains” and direct connections with “trusted domains”, which are domains with different domain names but can be authenticated. Trusted domains are authenticated to communicate with joined domains. This means that trusted domains (domains outside of the primary domain) can authenticate users and groups for access across multiple domains via Active Directory and Kerberos (KDC Servers) which manage cross-domain and cross-forest authentication. A cross-forest trust allows users from other forests to access the FlashBlade if their forest has bidirectional trust.

In a multi-domain trust environment where Active Directory and Kerberos now manage authentication, the Global Catalog, which is isolated within a single domain, is only used for optimizing search performance.

Multi-Domain Trust Setup

Multi-Domain Trust is enabled by default and does not require additional network configuration. Note that it only supports an Active Directory configuration but not a Directory Service configuration.

Note: Mutli-domain is only supported using Active Directory and SMB clients, it does not support other Directory Services or NFS clients.

Requirements:

- Use the `puread` command to join the Active Directory server to a domain.
- Bidirectional trusts.
- DNS configuration:
  - Be able to resolve LDAP/KDC SRV queries for all trusted domains by a configured DNS server.
  - Every forest root domain should be able to resolve Global Catalog SRV queries (on which domain is Global Catalog available).
  - The above configuration applies for all clients accessing the FlashBlade.
- Each trusted domain controller should be searchable by a configured FlashBlade account, including the Global Catalog.
- Site for DNS queries is discovered on Fabric Modules (FM). It is recommended that FMs and Blades are in the IP segment. Note that if not on the same site, LDAP queries from blades may not be sent to LDAP/GC servers in the same site.
- All trusted domains must support LDAPS (LDAP over SSL).
- At least one Global Catalog available per forest with a reachable port 3269.
- Universal, Global, and Domain local groups are supported.

Not Supported:

- One-way trusts (incoming or outgoing directions).
- Dedicated LDAP server for all trusted domains.

NFS Behavior in a Multi-Domain Trust

NFS and SMB interoperability is not supported in the current multi-domain trust for all users whose home domain is not in the joined domain. Support of RFC2037 is not manageable in multi-forest environments due to the challenges for the system administrator to ensure the user/group name mappings to uidNumber/gidNumber are unique within the whole Active Directory topology across forests. This may cause security risks if misconfigured.

NFS functionality of *nobody/nogroup* within the multi-domain trust behaves differently than the traditional use in UNIX and Linux. In a UNIX and Linux environment, users with the *nobody/nogroup* should not own

any files and their UID/GID is 0 (zero). It is primarily used to squash a user's privilege by setting their UID/GID to *nobody*.

In the multi-domain trust environment, Active Directory reads the nobody/nogroup from a user in trusted domains. The FlashBlade evaluates user permissions by ACLs and not by the owner/group. Clients from a trusted domain can create those files only via the SMB protocol and their SID (Secure ID) is stored in the ACL (Access Control List) as ACE (Access Control Entry). Therefore, the use of *nobody/nogroup* in FlashBlade means that it will not find the UID/GID of clients in joined domains, additionally it does not map clients from trusted domains to a UID/GID.

NFS and File Locking

The Network Lock Manager (NLM) protocol works with the NFS protocol to ensure file locks are visible across all NFS clients for I/O coordination and to help clients coordinate access to files. The NLM protocol allows all NFS clients mounting the same NFS shared file system to see file locks set by other clients.

NLM locks are considered advisory locks in that an NFS client application must check for the existence of a lock in order to coordinate access; the NFS service itself does not automatically prevent a client from accessing a file if it has the security permission to do so, and the service does not check for the existence of or try to obtain a lock on a file ahead of time.

The NLM service is enabled by default with the NFS protocol service and cannot be disabled.

Viewing and Managing File Locks

The **File Locks** panel lists all active file locks for NFS and SMB file systems.

The **File Locks** panel is displayed directly below the **File Systems** panel on the **Storage** > **File Systems** page.

You can select between showing general and detailed information. By default, the following general information about file locks is displayed:

- **Client IP**: The client IP address.
- **File System Name**: The file sytem name.
- **Path**: The directory path to the file system with the file lock.

- **Access Type**: The type of accessibility under a specific protocol lock. The protocol locks for NFSv3, NFSv4.1, or SMB can have an access type of exclusive or shared. The protocol lock for DOS can have the access type of read-only, write-only, read-write, or no access. The access types can be exlusive or shared, which are NFSv3, NFSv4.1, or SMB locks.
- **Created**: The date and type the file lock was created.
- **Protocol**: The protocol type of the file lock. It can be DOS, NFSv3, NFSv4.1, or SMB.

You can alternately click **Details** to view the following additional detail information:

- **Inode**: The Inode ID.
- **Lock Offset**: The offset of the lock relative to the file in bytes.
- **Lock Length**: The length in bytes displayed.

File locks are updated on demand. To retrieve the most current file lock information, click **Reload Cache**.

File locks can be deleted. The deletion of locks is to invalidate locks from disconnected clients. NFSv4.1 and SMB file locks can be invalidated before their locks timeout and cannot be renewed intentionally. To delete file locks, select **More Options > Delete Locks** from the **File Locks** panel header. Select the locks you wish you to delete and click **Delete**.

All blocked and active NLM locks can be cleared by selecting **More Options > NLM Reclamation** from the **File Locks** panel header, and selecting the confirmation and clicking **Initiate** from the **NLM Reclamation** pop-up window. All existing NFS file locks will be deleted and monitored clients will be notified that their locks have been removed. A 45 second grace period is provided to allow applications to restore their locks.

Fast Remove

When a directory and its contents are no longer required, they can be removed to reclaim space.

Removing the contents of a directory recursively by running the `rm -r` command causes the client to delete files one at a time. This can be a time consuming and expensive operation.

The fast remove feature allows you to quickly remove large directories by offloading this work onto the server. When the fast remove feature is enabled, a special pseudo-directory named `.fast-remove` is created in the root directory of the NFS mount. To remove a directory and its contents, run the `mv` command to move the directory into the `.fast-remove` directory.

In the following example, from the root directory of the NFS mount, a directory called `big_dir` is being moved into the `.fast-remove` directory.

```
mv big_dir/.fast-remove/
```

Once a directory has been moved into the `.fast-remove` directory, it is immediately destroyed, and through background operations, the space it occupied will be freed.

The fast remove feature can be enabled or disabled at any time. Fast remove is disabled by default.

After you enable the fast remove feature, view the contents of the root directory of the NFS mount to verify that the `.fast-remove` directory has been created.

For example,

```
//Run on the NFS mount to view the content of the root directory.
ls -al
drwxrwxrwx   1  root  root      0 Oct 30 10:46 .
drwxrwxrwx  25  root  root  12288 Jun 30 16:44 ..
drwxr-xr-x   1  root  root      0 Oct 30 11:27 .fast-remove
…
```

The `.fast-remove` directory cannot be renamed, deleted, or moved into other directories.
The `.fast-remove` directory name is a reserved name, so you cannot create a regular directory named `.fast-remove`, regardless of the fast remove enabled/disabled status.

You cannot move individual files to the `.fast-remove` directory. To remove a file, either use the `rm` command or move the directory enclosing the file to the `.fast-remove` directory.

The `.fast-remove` directory can only be removed by disabling the fast remove feature.

The fast remove feature is powerful and comes with the following cautionary notes:

- Once a directory has been moved into the `.fast-remove` directory, it is no longer recoverable.
- With fast remove, users can remove any files that are contained in a directory that they can rename, even if they do not have read, write, or execute permissions on the child directories. This breaks from the standard UNIX permissions model. To restrict access to the `.fast-remove` directory, set the directory permissions.

By default, the `.fast-remove` directory permissions are set to `drwxr-xr-x`. These permissions can be changed to grant or restrict access to the `.fast-remove` directory. Disabling and then re-enabling the fast remove feature will reset the settings to the default permissions.

Creating and Exporting a File System

Creating and exporting a file system involves creating the file system and configuring protocol support for the file system.

Before you create and export a file system, verify that a data vip has been created. Clients connect to the file system via the data vip. Data vips are managed through the **Settings** > **Network** > **Subnets** panel.

If you are integrating the FlashBlade array with a directory service, also verify that the directory service has been properly configured and enabled. The directory service is managed through the **Settings** > **Security** > **Directory Service** panel.

To create and export a file system:

1. From the **Storage** > **File Systems** page, click the add (+) button in the heading of the **File Systems** list. The **Create File System** pop-up window appears.
2. In the **Name** field, type the name of the directory to be exported.
3. In the **Data Node Group** field, select a data node group.
4. In the **Special Directories** section, to enable the ability to quickly remove large directories using the fast remove feature, select **Fast Remove**. For more information about the fast remove feature, refer to *Fast Remove*.
5. Under **Protocols**, select the **NFSv4.1** checkbox to enable. Select if you wish to not create exports, create rules-based exports, or create policy-based exports.
   a. If you select to create rules-based exports, configure the NFS export rules to define the access rights and privileges a client has to the file system:
      - For IPv4 addresses, `host` represents `ddd.ddd.ddd.ddd` for IP address, `ddd.ddd.ddd.ddd/dd` for netmask, or `*(options)` for all clients.
      - The format is `host(options)`.
      - For IPv6 addresses, `host` represents `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx` for IP address, `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx` for netmask, or `*(options)` for all clients.
      - The `options` within parenthesis represents a comma-separated list of NFS export options. For example, `10.20.225.2(ro,root_squash)`. Multiple Kerberos security values are specified with colons. For example, `10.20.255.2 (ro,root_squash,sec=krb5:krb5i:krb5p)`. Valid export options are listed in *NFS Export Rules*. Note that Kerberos security authentication is supported for NFSv3 and NFSv4.1.

b If you select the create policy-based exports, select or create an export policy.

6 If you wish to set group ownership parameters, click **Additional Settings**.

a Under **Group Ownership**, select **Creator** (default) or **Parent Directory**.

- **Creator**: The owning group of new files and directories is the primary group of the user who creates them.
- **Parent Directory**: The owning group of new files and directories is the primary group of the parent directory. New files and directories inherit the parent directory's group ID.

7 Click **Create**.

If you wish to edit a file system at a later date, perform the following:

1 Locate the file system you wish to edit in the **File Systems** list.

2 Click the **More Options** button in the same row as the file system and select **Edit**. The **Edit File System** pop-up window appears. All configuration parameters set during creation can be edited except the file system's name.

3 Click **Save** when you have completed editing.

Changing the NFS Export Rules or Policy of a File System

NFS export rule changes will be synchronously applied to all currently mounted clients. If you apply an NFS export policy, any previously configured NFS export rules for the file system will be overwritten.

To change the NFS export rules or policy of a file system:

1 From the **Storage > File Systems** page, click the **More Options** button in the row of the file system you want to edit and select **Edit**. The **Edit File System** pop-up window appears.

2 If you wish to change rules-based exports, configure the NFS export rules to define the access rights and privileges a client has to the file system:

- For IPv4 addresses, `host` represents `ddd.ddd.ddd.ddd` for IP address, `ddd.ddd.ddd.ddd/dd` for netmask, or `*(options)` for all clients.
- The format is `host(options)`.
- For IPv6 addresses, `host` represents `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx` for IP address, `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx` for netmask, or `*(options)` for all clients.
- The `options` within parenthesis represents a comma-separated list of NFS export options. For example, `10.20.225.2(ro,root_squash)`. Multiple Kerberos security values are specified with colons. For example, `10.20.255.2 (ro,root_squash,sec=krb5:krb5i:krb5p)`.

Valid export options are listed in *NFS Export Rules*. Note that Kerberos security authentication is supported for NFSv3 and NFSv4.1.

3 If you wish to change the policy-based export, select or create a new export policy.

4 Click **Save**.

Renaming a File System

To rename a file system:

1 Select **Storage** > **File Systems**.

2 From the **File Systems** panel, locate the file system you wish to rename and click the **More Options** button at the end of the row where the file system appears and select **Rename**. The **Rename File System** pop-up window appears.

Alternatively, click the file system you wish to rename to open its details page. Click the **More Options** button in the **Details** panel and select **Rename**. The **Rename File System** pop-window appears.

3 In the **New Name** field, enter the file system's new name.

4 Click **Rename**.

Enabling and Disabling Fast Remove

⚠ **Important!** The fast remove feature is available for removing directories only, not individual files.

Enable the fast remove feature to quickly remove large directories.

To enable or disable fast remove:

1 From the **Storage** > **File Systems** page, click the **More Options** button in the row of the file system you want to edit and select **Edit**. The **Edit File System** pop-up window appears.

2 Click the **Fast Remove** toggle button to switch between enabled (blue) and disabled (gray) status.

3 Click **Save**.

4 In the root directory of the NFS mount, confirm that a pseudo-directory named `.fast-remove` has been created.

By default, the `.fast-remove` directory permissions are set to `drwxr-xr-x`. Optionally use `chown` or `chmod` to set the desired access permissions.

To remove a directory, run the `mv` command to move the directory into the `.fast-remove` directory.

⚠ **Important!** Once a directory has been moved into the `.fast-remove` directory, it is immediately destroyed and unrecoverable.

Promoting and Demoting File Systems

⚠ **CAUTION:** If Pure File SafeMode is enabled, contact Pure Storage Technical Services to demote a file system.

To promote and demote file systems, perform the following:

1. Fom the **Storage** > **File Systems** page, locate the file system that you wish to promote or demote from the **File Systems** panel.
2. Click **More Options** > **Promote** or **More Options** >**Demote** at the end of the row in which the remote file system appears.
3. Click **Promote** or **Demote** when prompted for confirmation.

Enabling and Disabling File System Writes

From both the **File Systems** table and from the **Details** panel for a specific file system, you enable and disable file system writes.

1. To enable or disable file system writes for a single file system:
   a. Select **Storage** > **File Systems**.
   b. In the File Systems table, locate the file system for which you wish to enable or disable writes and click **More Options** > **Enable Writes** or **More Options** > **Disable Writes**.
   c. Click **Enable** or **Disable**.
2. To enable or disable file system writes for multiple file systems:

   Note: By default, the file systems displayed reflect filtering that was applied at the parent table. Unchecking **Apply external table filter** removes any previous filtering and all items as shown by default in the parent table are displayed. If no filter was applied at the parent table, the **Apply external table filter** checkbox will not appear.

   a. Select **Storage** > **File Systems**.

b In the File Systems heading, click **More Options** > **Enable Writes** or **More Options** > **Disable Writes**.

c If you selected **Enable Writes**, the **Enable Writes** pop-up window appears. From the **Readonly File Systems** panel, select the file systems you wish to make writeable. The selected file systems will appear in the **Selected File Systems** panel. Click **Enable**.

d If you selected **Disable Writes**, the **Disable Writes** pop-up window appears. From the **Writeable File Systems** panel, select the file systems you wish to make read only. The selected file systems will appear in the **Selected File Systems** panel. Click **Disable**.

Viewing File System Performance

From both the **File Systems** table and from the **Details** panel for a specific file system, you can view the following:

- Historical file system performance
- Group performance
- User performance

To view performance statistcs, peform the following:

1. Select **Storage** > **File Systems**.
2. Click the **More Options** button in the row of the file system that you wish to view performance information and select one of the following:
   - **View Historical Performance**
   - **View Performance by Groups**
   - **View Performance by Users**

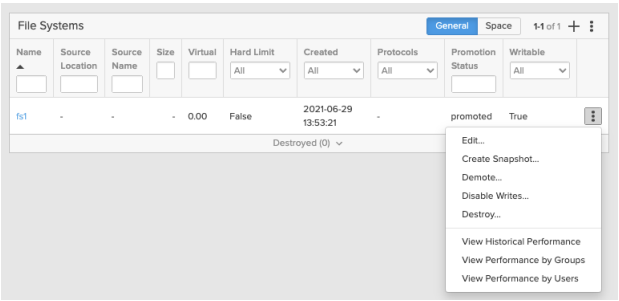

You can alternately select a file system from the **File Systems** table. From the selected file system's **Details** panel, click the **More Options** button to select one of the options described above.

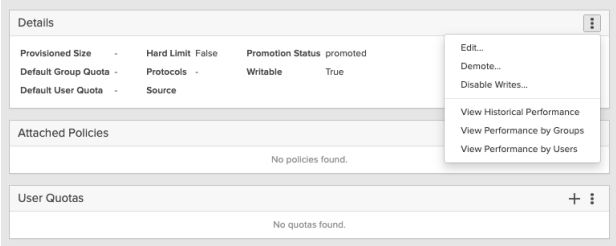

- Selecting **View Historical Performance** opens the peformance charts for that selected file system (see *The Performance Chart* for details).
- Selecting **View Performance by Groups** or **View Performance by Users** opens the **Groups** or **Users** performance table for that selected file system (see *Displaying Client, Group, and User Performance Statistics* for details).

Destroying a File System

Destroying a file system will also destroy all related snapshots. All protocols of a file system must be disabled before it can be destroyed; disabling protocols terminates all client connections to the file system. Any enabled protocols will be disabled by Purity//FB when destroying a file system.

To destroy a file system, perform the following:

1. Select **Storage** > **File Systems**.
2. To destroy a single file system:
   a. Click the **More Options** button in the row of the file system you want to destroy and select **Destroy**. The **Destroy File System** pop-up window appears. A message displaying the number of protocols that will be disabled and number of snapshots that will be destroyed is displayed.

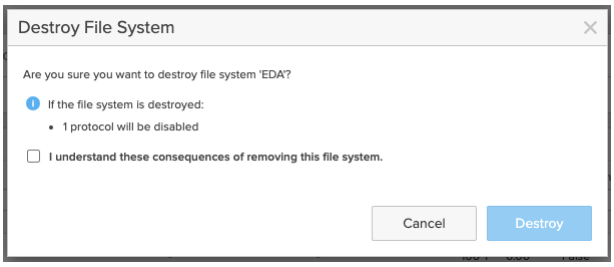

   b. Select the confirmation checkbox.
   c. Click **Destroy**.
3. If you wish to destroy multiple file systems:

Note: By default, the file systems displayed reflect filtering that was applied at the parent table. Unchecking **Apply external table filter** removes any previous filtering and all items as shown by default in the parent table are displayed. If no filter was applied at the parent table, the **Apply external table filter** checkbox will not appear.

a Click the **More Options** button in the header of the **File Systems** panel and select **Destroy**. The **Destroy File Systems** pop-up window appears.

b From the **Destroyable File Systems** list, select the file systems you wish to destroy. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all file systems by clicking the checkbox next to the search box. All selected file systems appear in the **Selected File Systems** list. If protocols are enabled on a file system, that file system will not appear in the **Destroyable File Systems** list.

c Click **Destroy**.

The destroyed file system appears in the **Destroyed File Systems** panel and begins its 24-hour eradication pending period.

During the eradication pending period, you can recover the file system to bring it back from its previous state, or manually eradicate the destroyed file system to reclaim the physical storage space occupied by the file system.

When the 24-hour eradication pending period has lapsed, Purity//FB starts reclaiming the physical storage occupied by the file system. Once reclamation starts, either because you have manually eradicated the destroyed file system, or because the eradication pending period has lapsed, the destroyed file system can no longer be recovered.

Recovering a File System

Destroyed file systems have 24 hours to be recovered. When a file system is recovered, all related snapshots will be also be recovered unless the snapshot was explicitly destroyed. When the 24 hour eradication pending state expires, the file system and its implicitly destroyed snapshots will be eradicated and you will no longer be able to retrieve the data.

To recover a destroyed file system:

1 Select **Storage** > **File Systems**.

2 In the **Destroyed File System** panel, click the clock icon in the row of the file system you want to recover. The **Restore File System** pop-up window appears.

Note: By default, the file systems displayed reflect filtering that was applied at the parent table. Unchecking **Apply external table filter** removes any previous filtering and all items as shown by

default in the parent table are displayed. If no filter was applied at the parent table, the **Apply external table filter** checkbox will not appear.

If you wish to recover multiple file systems, click **More Options** > **Recover** in the **Destroyed File Systems** panel under the **File Systems** panel. The **Recover File Systems** pop-up window appears. From the **Destroyed File Systems** list, select the file systems you wish to recover. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all file systems by clicking the checkbox next to the search box. All selected file systems appear in the **Selected File Systems** list.

3 Click **Recover**. The recovered file system appears in the **File Systems** panel.

Eradicating a File System

During the eradication pending period, you can manually eradicate the destroyed file system to reclaim physical storage space occupied by the destroyed file system. Eradicating a file system will also eradicate all of its related snapshots. Once reclamation starts, the destroyed file system and snapshots can no longer be recovered.

CAUTION: If Pure File SafeMode is enabled, contact Pure Storage Technical Services to perform this action.

To eradicate a destroyed file system:

1 Select **Storage** > **File Systems**.

2 In the **Destroyed File Systems** panel, click the trash icon at the end of the row of the file system you want to eradicate. The **Eradicate File System** pop-up window appears.

Note: By default, the file systems displayed reflect filtering that was applied at the parent table. Unchecking **Apply external table filter** removes any previous filtering and all items as shown by default in the parent table are displayed. If no filter was applied at the parent table, the **Apply external table filter** checkbox will not appear.

If you wish to eradicate multiple file systems, click **More Options** > **Eradicate** in the **Destroyed File Systems** panel under the **File Systems** panel. The **Eradicate File Systems** pop-up window appears. From the **Destroyed File Systems** list, select the file systems you wish to eradicate. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all file systems by clicking the checkbox next to the search box. All selected file systems appear in the **Selected File Systems** list.

3 Click **Eradicate**. The array will begin reclaiming the eradicated space of the file system.

Open Files and Sessions Management with the Microsoft Management Console

The Microsoft Management Console (MMC) can be used with FlashBlade to manage open files and sessions.

Open files and sessions can be managed through MMC using the `fsmgmt.msc` snap-in (**Share Folders**), which allows administrators to manageopen files and sessions directly from the console on a Windows client.

In order to use MMC on the Windows client to manage open files and sessions, SMB protocol must be enabled on the file systems and an Active Directory account must be set up on the FlashBlade. Refer to *Creating and Exporting a File System* and *Creating an Active Directory Account* for additional information.

The following actions can be executed for open files using MMC:

- View open files.
- Close open files.

The following actions can be executed for sessions using MMC:

- View active SMB sessions.
- Close active SMB sessions.

For more information about MMC, refer to Microsoft MMC documentation.

Setting Up the Microsoft Management Console

Open files and sessions can be managed directly from a Windows client through the Microsoft Management Console (MMC) using the Share Folders snap-in. For more information about MMC, refer to Microsoft MMC documentation.

To set up the **Shared Folders** Microsoft Management Console (MMC) snap-in, perform the following on the Window's client:

1. Click **Start** > **Run**. The **Run** pop-up window appears.
2. Enter **mmc** and click **OK**. The Console appears.
3. Click **File** > **Add/Remove Snap-in**. The **Add or Remove Snap-ins** pop-up window appears.

4 Select **Shared Folders** from the **Available snap-ins** panel and then click **Add**. The **Shared Folders** pop-up window appears.

5 Select **Another computer** and enter the name or IP address of the AD account computer name of the FlashBlade.

6 Under **View**, select **All** and then click **Finish**. The **Shared Folders** pop-up window closes.

7 In the **Add or Remove Snap-ins** pop-up window, click **OK**.

The Console now displays the **Shared Folders** snap-in.

Managing Open Files with the Microsoft Management Console

When clients connected over SMB (Server Message Block) leave files open, it can prevent other users from accessing those files. Typically, these locks are not listed in the **File Locks** panel of the GUI because it does not include OpLocks (opportunistic locks).

Once the **Shared Folders** Microsoft Management Console (MMC) snap-in has been set up, the **Open Files** folder displays all open files of the corresponding SMB shares.

Perform the following to manage open files:

1 Open MMC on the Windows client.

2 Select **Open files** under **Shared Folders**.

3 To close one or more files, right-click the file (or selection of files) and click **Close File**. When prompted for confirmation, click **Yes**.

The file(s) is closed and all locks are released.

Chapter 6
The Purity//FB GUI Policies Page

The **Policies** page displays and manages the array's NFS and SMB export policies and rules, and object and file system snapshot policies.

The left side of the **Policies** page provides a navigation pane for quick access to the following:

- **Data Eviction**: Displays and allows you to manage data eviction policies. See *Data Eviction Policies* for additional information.
- **NFS Export**: Displays and allows you to manage NFS export policies and NFS export rules. See *NFS Export Policies* for additional information.
- **Object Store Access Policies**: Displays and allows you to manage object store access policies. See *Object Store Access Policies* for additional information.
- **Management Access Policies**: Displays and allows you to manage the array's management access policies. See *Management Access Policies* for additional information.
- **Password**: Displays and allows you to manage the array's password policy. See *Password Policies* for additional information.
- **SSH Certificate Authority**: Displays and allows you to manage the SSH login rules for users, including local and remote users on the array. See *SSH Certificate Authority Policies* for additional information.
- **Network Access**: Displays and allows you to manage network access policies and network access rules. See *Network Access Policies* for additional information.
- **TLS**: Displays and allows you to manage TLS policies on the system. See *TLS Policies* for additional information.

NFS Export Policies

NFS export policies are made up NFS export rules that govern how files are shared over NFS. NFS export rules define the access rights and privileges that an NFS client has to the file system.

FlashBlade allows multiple file systems to use the same policy, thereby facilitating scaling of file systems by allowing a single change in an NFS export policy to affect many file systems.

NFS Export settings allows you to view, create, and manage NFS export policies and NFS export rules.

NFS Export settings contains the **NFS Export Policies** and **NFS Export Rules** panels.

The **NFS Export Policies** panel manages and displays the following information for all configured NFS policies:

- **Name**: The policy name.
- **Type**: The NFS policy type.
- **Enabled**: If the policy is enabled (**True**) or disabled (**False**).

The **NFS Export Rules** panel manages and displays the following information for all NFS Export Rules:

- **Policy**: The policy to which a rule is attached.
- **Client**: The IP address, host name, subnet in CIDR notation, netgroup, or anonymous (*).
  - **IP address**: An IPv6 or IPv4 IP address.
  - **host name**: A host name in short form (without dots) or in FQDN form, that can include the wildcards * or **?**.
    - *: Matches any number of characters.
    - **?**: Matches a single character in a specific position.
  - **subnet in CIDR notation**: An IPv6 or IPv4 IP address and the length of the netmask separated with a slash (/).
  - **netgroup**: An Active Directory, LDAP, or NIS netgroup.
  - *: The asterisk (*) character specifies no filtering is applied.
- **Permission**: If the file system is read-only (**ro**) or read-write (**rw**).
- **Security**: The rule's security protocol. Values can be **krb5**, **krb5i**, **krb5p**, or **sys**. Note that multiple security protocols can be specified during rule creation and will be displayed.
- **Access**: The NFS export access rule. Values can be **no-squash**, **root-squash**, or **all-squash**.

- **Anon UID**: Anonymous user IDs.
- **Anon GID**: Anonymous group IDs.
- **Transport Security**: The transport security type. Values can be **none**, **tls**, or **mutual-tls**.
- **Secure**: If client ports are restricted to secure ports (**True**) or not (**False**).
- **fileid-32bit**: If 32-bit inodes are used (**True**) or 64-bit inodes are used (**False**).
- **atime**: If the timestamp is updated when a file is accessed (**True**) or not (**False**).
- **Index**: The index of the rule as ordered in its policy.

NFS export policies are comprised of an ordered list of NFS export rules. An NFS export rule has a single client specification of IP address (IPv6 or IPv4), subnet in CIDR notation, netgroup (from AD, LDAP, or NIS), or anonymous (*), and consists of a single set of access permissions and one or more security settings (for NSFv4). Clients must be unique within an NFS export policy.

Clicking a policy in the **NFS Export Policies** panel or in the **NFS Export Rules** panel will open the details page for rules associated with that policy. Information about each of the policy's rules is displayed in the detail page's **NFS Export Rules** panel (same information as listed above, except sorted for each client). IP address rules are in the first group, then host names, subnet in CIDR notation, or netgroup, followed by anonymous (*) in the last group. The **Index** column displays the order in which each group's rules will be executed. The rules can be reordered within each group.

The **Details** panel displays the policy name and if it is currently enabled. It also allows you to disable and rename the policy.

The **File System Exports** panel displays all file system exports governed by the policy. It also allows you to create and delete file system exports from the policy.

For additional information about export rules, see *NFS Export Rules*.

NFS Export Rules

The NFS export rules define the access rights and privileges that an NFS client has to the file system. An NFS client's access rights and privileges include rules for users and groups. Any NFS export rule changes will be synchronously applied to all currently mounted clients. NFS rules apply to both NFSv3 and NFSv4.1. The NFS export rules of a file system can be changed at any time.

Note: Kerberos security authentication is supported for NFSv3 and NFSv4.1.

User and Group IDs that are zero (0) have root privilege and IDs that are non-zero do not have root privilege.

Valid export rules are listed below:

Table 6.1 Export Rules

Export Rules	Rules Description
`ro`	Read-Only grants an NFS client and their users or groups the privilege to read the file system.
`rw`	Read-Write grants an NFS client and their users or groups the privilege to read and write to the file system.
`sec=sys`	No cryptographic protection. This is the default.
`sec=krb5`	Enable Kerberos security authentication.
`sec=krb5i`	Enables Kerberos security authentication and integrity checking of NFS operations.
`sec=krb5p`	Enables Kerberos security authentication, integrity checking of NFS operations, and NFS traffic encryption.
`fileid_32bit`	Allows 32-bit inode support for clients.
`nofileid_32bit`	Disables 32-bit inode support for clients.
`root_squash`	Prevents client users and groups with root privilege from mapping their root privilege to a file system. All users with UID 0 will have their UID mapped to `anonuid`. All users with GID 0 will have their GID mapped to `anongid`.
`no_root_squash`	Allows root users and groups to access the file system with root privilege.
`all_squash`	Maps all UIDs (including root) to anonuid and all GIDs (including root) to anongid.

Table 6.1 Export Rules (continued)

Export Rules	Rules Description
`no_all_squash`	Prevents the remapping of user and group IDs to `anonuid 65534` or `anongid 65534`. All users and groups will retain their IDs unless `root_squash` is also specified.
`anonuid`	Any user whose UID is affected by `root_squash` or `all_squash` will have their UID mapped to `anonuid`. The default `anonuid` is `65534`.
`anongid`	Any user whose GID is affected by `root_squash` or `all_squash` will have their GID mapped to `anongid`. The default `anongid` is `65534`.
`atime`	After a read operation has occurred, the inode access time is updated only if any of the following conditions is true: the previous access time is less than the inode modify time, the previous access time is less than the inode change time, or the previous access time is more than 24 hours ago.
`noatime`	Disables the update of inode access times after read operations.
`secure`	Prevents NFS access to client connections coming from non-reserved ports. Applies to NFSv3, NFSv4.1, and auxiliary protocols MOUNT and NLM.
`insecure`	Allows NFS access to client connections coming from non-reserved ports. Applies to NFSv3, NFSv4.1, and auxiliary protocols MOUNT and NLM.

Note: NFS export rules sort is stable, and rules prioritization is as follows: IP address > host name >wildcard > netmask > netgroup > star (*). IP address > Host Name > Wildcard > Netmask > Netgroup > *. NFS export policies follow a different prioritization: IP address > Host Name | Wildcard | Netmask | Negroup > *. The user interfaces (GUI and CLI) enforce the prioritization so that you cannot, for example, place a Netmask rule in front of an IP address rule within a policy. The order of Netmask and Netgroup,

however, are interchangeable within a policy. The rules within a policy are then evaluated in the exact order as they are listed.

NFS export rules can be added to NFS export policies, or applied during file system creation as an NFS export rule string.

During NFS export policy creation, you can optionally add NFS export rules (see step 4 under *Creating NFS Export Policies and Rules*). Multiple NFS export rules added to an NFS policy are ordered (indexed) and are evaluated and applied by Purity//FB in numerical order. Available rules are listed in the table *Table 6.1* .

When creating a file system you can optionally apply an NFS export policy, or specify an NFS export rule string (see step 8 under *Creating and Exporting a File System*). Available rules are listed in the table *Table 6.1* .

The following information provides more details and guidelines about creating NFS export rule strings.

The following default export rules will be applied to NFS clients if none are specified:

- `rw`
- `no_root_squash`

Rules are applied in the form of `host(options)`, where `host` is an IP address or netmask and `options` are export rules enclosed in parenthesis.

The `RULES` argument represents the export rules that apply to the file system. Options can be applied to a single or multiple IP addresses, Netmasks, or LDAP group. See the syntax examples:

- Options applied to a single client: `'host(options)'`
- Options applied to multiple clients: `'-options host1 host2…'`
- Represent all clients with an asterisk: `*`

The `host` parameter must belong to one of the following categories:

- IPv4 IP address: `ddd.ddd.ddd.ddd`

  IPv4 Netmask: `ddd.ddd.ddd.ddd/dd`
- IPv6 IP address: `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`

  IPv6 Netmask: `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/ddd`

Multiple rules may be specified by separating each rule with a space. For example,

```
10.60.50.83/32(rw,root_squash) 2620::/16(ro)
```

The following guidelines apply when creating NFS export rules:

- If rules are not defined before enabling the NFS protocol, then all clients will have `rw` access and `no_root_squash` privileges to the file system. This is equivalent to specifying export rule `*(rw,no_root_squash)`.
- Omitting the `(options)` portion of a rule is equivalent to specifying `(ro,root_squash)`. For example, an export rule that is defined as `*` will have `read-only, root_squash` access. This is equivalent to specifying export rule `*(ro,root_squash)`.
- If the rule is defined but one or more export options are not specified, then the undefined option(s) will default to `ro` access and `root_squash` privileges. For example, an export rule that is defined as `*(rw)` will have `root_squash` privileges. Likewise, an export rule that is defined as `*(no_root_squash)` will have `ro` access.
- A rule cannot include conflicting options. For example, `ro` and `rw` cannot be included in the same rule. Likewise, `root_squash` and `no_root_squash` cannot be included in the same rule.
- If a client IP address matches multiple rules, Purity//FB uses the first rule it encounters in priority based on rule category. IP address rules have the highest priority. Specifically, the priority of matching rules is as follows: IP address > Subnet > Netgroup > Wildcard.
- If an address matches multiple rules in the same category, then the first (leftmost) rule applies.
- By default, the FlashBlade file systems use 64-bit inode numbers. If your environment contains clients that can only support 32-bit inode numbers, add a third export option named `fileid_32bit` to the NFS export rule of the file system that requires 32-bit inode support. For example, `*(rw,no_root_squash,fileid_32bit)`.
- Multiple security types are specified with a colon separator. For example, `sec=krb5:krb5i:krb5p:sys`.

Here is an example of a file system list output taken from the CLI with various NFS export rules set for each file system:

```
purefs list --protocol-specific nfs
Name    Protocols Rules
File1   nfs       *
File2   nfs       *()
File3   nfs       -
File4   nfs       10.20.225.2(rw)
File5   nfs       2620::/16(no_root_squash)
File6   nfs       *(rw,no_root_squash)
File7   nfs       *(rw,root_squash,fileid_32bit)
File8   nfs       10.20.30.40(root_squash,anonuid=7777,anongid=8888)
```

In the list output above, the NFS export rules are described as follows:

Export Rules	Rules Description
`*`	All clients have default `ro` access and `root_squash` privileges to the file system.
`*()`	All clients have default `ro` access and `root_squash` privileges to the file system.
`-`	Export rules have been deleted, rendering the file system inaccessible.
`10.20.225.2(rw)`	Client `10.20.225.2` has `rw` access and `root_squash` (default) privileges.
`2620::/16(no_root_squash)`	Client `2620::/16` has `ro` (default) access and `no_root_squash` privileges.
`*(rw,no_root_squash)`	All clients have `rw` access and `no_root_squash` privileges. This is the default rule used when creating a file system.
`*(rw,root_squash,fileid_32bit)`	All clients have `rw` access and `root_squash` privileges. The file system also supports clients that require 32-bit inode support.
`10.20.30.40 (root_squash,anonuid=7777,anongid=8888)`	Client 10.20.30.40 has `rw` access. Any access attempts with UID 0 will be mapped to UID 7777. Any access attempts with GID 0 will be mapped to GID 8888.

Netgroups

Purity//FB supports netgroups with AD, LDAP, and NIS.

Netgroups allows a set of hosts to be referenced as a single entity, or a *netgroup*.

Netgroups are used in export rules to grant NFS mount attributes to an entire set of hosts within a netgroup. This allows you to change the netgroup membership and affect every reference to that netgroup.

Netgroups are stored in LDAP as objects with the class nisNetgroup. Each netgroup has a name and may contain an unlimited number (or none) of member netgroups (sub-netgroups), as well as an unlimited number (or none) of tuples.

On a network using NIS, the ***netgroup*** input file on the master NIS server is used for generating the ***netgroup*** map. The ***netgroup*** map contains the basic information in the ***netgroup*** input file, which includes information in a format that speeds lookups of netgroup information among hosts.

FlashBlade queries directory service for a given netgroup, following any member netgroups and reading every tuple until the client attempting to mount the NFS file system is encountered within the netgroup, or until all referenced hosts are exhausted.

Referencing a netgroup with an export group must be expressed in the format `@<netgroup_name>(<export_rule>)`. The `@` indicates the specified `netgroup_name` is a netgroup. `(<export_rule>)` specifies the export rule to which the netgroup is referenced. Multiple export rules can be specified in a comma-separated list. Parentheses are required.

For example, `@netgroup1(rw,no_root_squash)` indicates that every host mentioned in the netgroup (`netgroup1`), may mount the given file system with read-write (`rw`) and root privilege (`no_root_squash`) attributes for that file system.

Creating NFS Export Policies and Rules

To create an NFS export policy and its rules, perform the following:

1. From the **Policies** > **NFS Export** settings page, click the **Add** (+) button in the heading of the **NFS Export Policies** panel. The **Step 1: Create NFS Export Policy** pop-up window appears.
2. Enter the policy name.
3. (Optional) By default, the policy will be enabled upon creation. To disable the policy, click the **Enabled** checkbox.

4 Click **Create**. The **Step 2: Add NFS Export Rule** pop-up window appears. You can continue to configure the export rule or click **Skip** to configure the rule at later time.

5 In the **Client** field, enter the client as an IP address (IPv6 or IPv4), host name in short or FQDN form (including wildcards), subnet in CIDR notation, netgroup (Active Directory, LDAP, or NIS), or anonymous (*).

6 Set the access permission to **Read-Only** or **Read-Write**.

7 Select one or more security protocol. The security options krb5, krb5i, and krb5p are not supported by NFSv3.

8 Select one of the following NFS export access rules: **no-squash**, **root-squash**, or **all-squash**.

9 (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous UID** field enter one or more UIDs.

10 (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous GID** field, enter one or more GIDs.

11 Set the data transport security to **none**, **tls**, or **mutual-tls**.

12 (Optional) Select to restrict client ports to secure ports only.

13 (Optional) Select to use 32-bit inodes instead of 64-bit inodes.

14 (Optional) Select to update the timestamp when accessed.

15 Click **Save**.

Once created, the policy and rule will appear in the **NFS Export Policies** and **NFS Export Rules** panels.

Creating File System Exports

A file system export can be associated with one policy. Multiple file system exports can use the same policy.

To create a file system export, perform the following:

1 From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy to which you wish to create a file system export. The detail page for that policy appears.

2 Click the Add (+) button in the heading of the **File System Exports** panel. The **Create File System Export** pop-up window appears.

3 From the **Server** field, select the server to which which you are exporting the file system.

4 In the **File System** field, select or enter the name of the file system you wish to export.

5 In the **Export Name** field, enter a name for the exported file system, which is the name the clients will use to connect to the file system.

6 In the **Type** field, **NFS** is pre-selected (cannot be changed), which is the type of policy the file system is attached.

7 The **NFS Export Policy** field is pre-populated with the name of the NFS export policy which you selectedd to create the file system export .

8 Click **Create**.

Editing File System Exports

When editing a file system export, only the file system export's name can be edited.

To edit a file system export, perform the following:

1 From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy wkth the file system export you wish to edit. The detail page for that policy appears.

2 In the **File System Exports** panel, click the Edit button in the same row as the file system export you wish to edit. The **Edit File System Export** pop-up window appears.

3 In the **Export Name** field, enter a new name for the exported file system, which is the name the clients will use to connect to the file system.

4 Click **Save**.

Deleting File System Exports

Note: When removing a file system export from an NFS export policy, some clients may lose access to the removed file system export.

1 To delete a single file system export from an NFS export policy:

   a From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy from which you wish to delete a file system export. The detail page for that policy appears.

   b In the **File System Exports** panel, click the Delete button on the same row as the file system export you with to delete. The **Delete File System Export Configuration** pop-up window appears.

   c Click **Delete**.

2 To delete multiple file system exports from an NFS export policy:

a From the **Policies** > **NFS** > **NFS Export Policies** panel, click the NFS export policy from which you wish to delete multiple file systems. The detail page for that policy appears.

b Click **More Options** > **Delete** in the heading of the **File Sysetm Exports** panel. The **Delete File System Exports** pop-up window appears.

Note: By default, the file systems displayed reflect filtering that was applied at the parent table. De-selecting **Apply external table filter** removes any previous filtering and all items as shown by default in the parent table are displayed.

c From the **Existing File System Exports** panel, select the file system exports you wish to remove from the policy. If you wish to remove all file systems, click the checkbox next to the search box to select all file systems. Each selected file system appears in the **Selected File System Exports** panel.

d Click **Delete**.

Adding Additional NFS Export Rules to NFS Export Policies

To add an NFS export rule to an existing NFS export policy, perform the following:

1. From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy to which you wish to add a new rule. The detail page for that policy appears.
2. Click the **Add** (+) button in the heading of the **NFS Export Rules** panel. The **Add NFS Export Rule** pop-up window appears.
3. In the **Client** field, enter the client IP address (IPv6 or IPv4),host name in short or FQDN form (including wildcards), subnet in CIDR notation, netgroup (Active Directory, LDAP, or NIS), or anonymous (*).
4. Set the access permission to **Read-Only** or **Read-Write**.
5. Select one or more security protocol. The security options krb5, krb5i, and krb5p are not supported by NFSv3.
6. Select one of the following NFS export rules: **no-squash**, **root-squash**, or **all-squash**.
7. (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous UID** field enter one or more UIDs.
8. (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous GID** field, enter one or more GIDs.
9. Set the data transport security to **none**, **tls**, or **mutual-tls**.
10. (Optional) Select to restrict client ports to secure ports only.
11. (Optional) Select to use 32-bit inodes instead of 64-bit inodes.

12 (Optional) Select to update the timestamp when accessed.

13 Select if the new rule will be applied after the existing rule(s) (**Default**) or **Before** the existing rule(s). If you select **Before**, select which existing rule the new rule will be moved before.

14 Click **Add**.

Editing NFS Export Rule Clients

NFS export rules are created with a single client specification. You can bulk-edit the rules' clients to a new client.

To edit NFS export rule clients, perform the following:

1 From the **Policies** > **NFS Export** > **NFS Export Rules** panel, click **More Options** > **Edit** in the heading of the panel. The **Edit Rules** pop-up window appears.

2 In the **Existing Rules** panel, select the rules for which you wish to change the client. Each selected rule appears in the **Selected Rules** panel.

3 In the **Client** field, enter the new client as an IP address (IPv6 or IPv4), host name in short or FQDN form (including wildcards), subnet in CIDR notation, netgroup (Active Directory, LDAP, or NIS), or anonymous (*).

4 Click **Save**.

Editing NFS Export Rules

To edit an existing NFS export rule, perform the following:

1 From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy whose rule you wish to edit. The detail page for that policy appears.

2 In the **NFS Export Rules** panel, click **More Options** > **Edit** at the end of the row in which the rule you wish to edit appears. The **Edit NFS Export Rule** pop-up window appears.

3 In the **Client** field, enter the client as an IP address (IPv6 or IPv4), host name in short or FQDN form (including wildcards), subnet in CIDR notation, netgroup (Active Directory, LDAP, or NIS), or anonymous (*).

4 Set the access permission to **Read-Only** or **Read-Write**.

5 Select one or more security protocol. The security options krb5, krb5i, and krb5p are not supported by NFSv3.

6 Select one of the following NFS export access rules: **no-squash**, **root-squash**, or **all-squash**.

7 (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous UID** field enter one or more UIDs.

8 (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous GID** field, enter one or more GIDs.

9 Set the data transport security to **none**, **tls**, or **mutual-tls**.

10 (Optional) Select to restrict client ports to secure ports only.

11 (Optional) Select to use 32-bit inodes instead of 64-bit inodes.

12 (Optional) Select to bypass updating the timestamp when accessed.

13 Click **Add**.

Reordering NFS Export Rules

The order that NFS export rules are applied is displayed in the **Index** column of the **NFS Export Rules** panel. NFS export rules can be reordered per client type.

Reordering may be required if a client is in two netgroups to determine which rule takes precedent. For example `rule1` with the client `@group1` is indexed at 4 and `rule2` with the client `@group2` is indexed at 5. `ClientX` belongs to both `@group1` and `@group2`. With the current rule ordering, only `rule1` will be applied to `ClientX`. If `rule2` is moved before `rule1`, then only `rule2` will be applied to `ClientX`.

Note that NFS export rules sort is stable, and rules prioritization is as follows: IP address > host name >wildcard > netmask > netgroup > star (*).

To move an NFS export rule, perform the following:

1 From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy whose rule you wish to move. The detail page for that policy appears.

2 In the **NFS Export Rules** panel, click **More Options** > **Move** at the end of the row in which the rule you wish to move appears. The **Move NFS Export Rule** pop-up window appears.

3 In the **Before** list, select the new order of the rule.

4 Click **Move**.

Once moved, the **NFS Export Rules** panel is updated reflecting the new rule order.

Removing NFS Export Rules

1 To remove a single NFS export rule:

a From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the policy from which you wish to remove a rule. The detail page for that policy appears.

b In the **NFS Export Rules** panel, click **More Options** > **Remove** on the row of the rule you wish to remove. The **Remove Rule** dialog box appears.

c Click **Remove**.

2 To remove multiple rules from an object store access policy:

a From the **Policies** > **NFS** > **NFS Export Policies** panel, click the policy from which you wish to remove multiple rules. The detail page for that policy appears.

b Click **More Options** > **Remove** from the heading of the **NFS Export Rules** panel. The **Remove NFS Export Rules** pop-up window appears.

Note: By default, the rules displayed reflect filtering that was applied at the parent table. Selecting **Show all items** removes any previous filtering and all items as shown by default in the parent table are displayed.

c From the **Available Rules** panel, select one or more rules from the list. The selected rule(s) appears in the **Selected Rules** panel. To select all rules, select the checkbox directly below the **Available Rules** heading.

d Click **Remove**.

Renaming NFS Export Policies

NFS export policies can be renamed at any time after their creation.

To rename an NFS export policy, perform the following:

1 From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy you wish to rename. The detail page for that policy appears.

2 Click **More Options** > **Rename** in the heading of the **Details** panel. The **Rename NFS Export Policy** pop-up window appears.

3 Enter the policy's new name and click **Rename**.

All panels displaying the policy are updated with the new name.

Enabling and Disabling NFS Export Policies

Unless explicitly set to disabled, NFS export policies are enabled by default when they are created.

NFS export policies can be disabled or enabled at any time after their creation.

Note: When disabling a policy, all client access to all attached file systems will be restricted.

To enable or disable an NFS export policy and its rules, perform the following:

1. From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click the NFS export policy you wish to enable or disable. The detail page for that policy appears.
2. To enable the policy:
   a. Click **More Options** > **Enable** in the heading of the **Details** panel. The **Enable Policy** pop-up window appears.
   b. Click **Enable**.
3. To disable the policy:
   a. Click **More Options** > **Disable** in the **Details** panel. The **Disable Policy** pop-up window appears.
   b. Click **Disable**.

Deleting NFS Export Policies

1. To delete a single NFS export policy:
   a. From the **Policies** > **NFS Export** > **NFS Export Policies** panel, click delete button at the end of the row in which the policy you wish to delete appears. The **Delete NFS Export Policy** pop-up window appears.
   b. Click **Delete**.
2. To remove multiple rules from an object store access policy:
   a. From the **Policies** > **NFS** > **NFS Export Policies** panel, click **More Options** > **Delete** in the heading of the **NFS Export Policies** panel. The **Delete Policies** pop-up window appears.

      Note: By default, the policies displayed reflect filtering that was applied at the parent table. Selecting **Show all items** removes any previous filtering and all items as shown by default in the parent table are displayed.

   b. From the **Available Policies** panel, select one or more policies from the list. The selected policies appears in the **Selected Policies** panel. To select all policies, select the checkbox directly below the **Available Policies** heading.
   c. Click **Delete**.

Password Policies

The FlashBlade array provides a single password policy, **management**, that governs rules such as password length, attempts before lockout, length of lockout, etc. The password policy can be edited, allowing you the flexibility to fine-tune password requirements to your organization's needs.

To view your array's password policy, from the **Policies** page, click **Password**. The **Details** panel alows you to view the current rules of the array's password policy.

The **Details** panel displays the following:

- **Name**: The name of the password policy.
- **Enabled**: Whether the policy is enabled (**True**) or disabled (**False**).
- **Minimum Password Length**: The minimum allowable password length for the array, from 1 to 100 characters.
- **Maximum Login Attempts**: The maximum number of login attempts before lockout, from 1 to 100. If 0 is displayed, this rule is disabled.
- **Lockout Duration**: The lockout duration after the maximum number of login attempts, from 1 second to 90 days in seconds, minutes, hours, or days. If 0 is displayed, this rule is disabled.
- **Password History**: The number of passwords tracked to prevent password reuse. If 0 is displayed, this rule is disabled.
- **Minimum Password Age**: The minimum duration before the password can be changed from its last change.
- **Maximum Password Age**: The maximum duration before the password must be changed.
- **Minimum Character Groups**: The minimum number of character groups required in the password. There are four character groups:
  - Group **1**: a-z
  - Group **2**: A-Z
  - Group **3**: 0-9
  - Group **4**: Special characters

- **Minimum Characters Per Group**: The minimum number of characters per group to count the group as present.
- **Enforce Username Check**: Checks whether the password contains the username is enabled (**True**) or disabled (**False**).
- **Enforce Dictionary Check**: Checks whether the password is from a dictionary of known compromised passwords is enabled (**True**) or disabled (**False**).

Editing Password Policies

To edit a password policy, perform the following:

1. From the **Policies** > **Details** panel, click **More Options** from the panel header and then select **Edit Policy**. The **Edit Password Policy** pop-up window appears.
2. In the **Name** field, optionally enter a new policy name. By default, the policy name is **management**.
3. Select the **Enable** checkbox to ensure that the policy is enabled once you are finished editing.
4. In the **Minimum Password** field, enter the minimum character length for the password, from 1 to 100.
5. In the **Maximum Login Attempts** field, enter the maiximum allowable number of login attempts before lockout occurs, from 1 to 100. If 0 is entered, the maximum login attempts rule is disabled.
6. In the **Lockout Duration** field, enter the lockout time once the maximum number of login attempts is reached, from 1 second to 90 days. Values can be entered in seconds, minutes, hours, and days. If 0 is entered, the lockout duration rule is disabled.
7. In the **Password History** field, enter the number of passwords tracked to prevent reuse of old passowrds, from 1 to 64. If 0 is entered, the password history rule is disabled.
8. In the **Minimum Password Age** field, enter the minimum duration before a password can be changed from its last change, from 0 hours to 1 week. Values can be entered in hours and days. For example, 1d.
9. In the **Maximum Password Age** field, enter the maximum age of a password before it must be changed, from 1 day to 99,999 days. Values can be entered in hours, days, week, and years. For example, 3w.
10. In the **Minimum Character Groups** field, enter the minimum number of character groups required to be present in the password. There are four character groups:
    - 1: a-z character group. Entering **1** indicates the password must contain one or more members of the a-z character group.
    - 2: A-Z character group. Entering **2** indicates the password must contain one ore more members of the A-Z and a-z character groups.

- 3: 0-9 character group. Entering **3** indicates the password must contain one or more members of the 0-9, A-Z, and a-z character groups.
- 4: special character group. Entering **4** indicates the password must contain one or more membesr of the special characters, 0-9, A-Z, and a-z character groups.

11 In the **Minimum Characters Per Group** field, enter the minimum number of characters allowed per character group to count the group as present. For example, if **1** is specified in the **Minimum Character Groups** field, entering **2** in the **Minimum Characters Per Group** field indicates that two characters from character group 1 (a-z) must be used in the password.

12 Select the **Enforce Username Check** checkbox to check if the password contains the username.

13 Select the **Enforce Dictionary Check** checkbox to check if the password appears in a dictionary of known compromised passwords.

14 Click **Save**.

The password policy is updated and the **Details** panel is updated reflecting the changes.

Disabling Password Policies

Password policies can be disabled at any time.

To disable a password policy, perform the following:

1 From the **Policies** > **Details** panel, click **More Options** from the panel header and then select **Disable Policy**. The **Disable Policy** pop-up window appears.

2 Click **Disable**.

The password policy is disabled and the **Details** panel is updated reflecting the policy status.

Enabling Password Policies

To enable a password policy, perform the following:

1 From the **Policies** > **Details** panel, click **More Options** from the panel header and then select **Enable Policy**. The **Enable Policy** pop-up window appears.

2 Click **Enable**.

The password policy is enabled and the **Details** panel is updated reflecting the policy status.

SSH Certificate Authority Policies

An SSH Certificate Authority Policy manages the SSH login rules for users, including local and remote users on the array. FlashBlade supports X.509 certificate or public key format signing authority to verify SSH logins with an SSH certificate. The signing authority is specified per policy and the policy can be attached to either array or per admin user. This is a more secure method of login compared to each user uploading their public key for SSH. The central authority can use these policies to assign CA certificates and public keys to policies then add users and arrays to those policies.

When a user is added to a policy, they can log in to an array using SSH plus their SSH certificate.

Customer now has the option to set an expiration time for the issued SSH certificate for enhanced SSH login management of the FlashBlade users. Each SSH certificate can (and generally should) have an expiration time, after which point that particular SSH certificate becomes invalid. If the SSH certificate authority policy is configured with an X.509 certificate, that may have an expiration time, after which point we will not accept any SSH certificate signed by the key corresponding to that X.509 certificate.

User login activities can be monitored in the session log, indicating whether the user logged in using an SSH certificate or a public key. The details of the certificate and public key are also displayed in the session log.

Adding a Public Key

To add a public key, perform the following:

1. From the **Settings** > **Public Key** tab, click **Add** (+). The **Add Public Key** pop-up window appears.
2. In the **Name** field enter the name for the public key.
3. In the **Key** field enter your key value by copying and pasting the public key.
4. Click **Add**.

Creating an SSH Certificate Authority Policy

To create a policy, perform the following:

1. From the **Policies** > **SSH Certificate Authority** tab, click the **Add** (+) button in the heading of the **SSH Certificate Authority** list. The **Create SSH Certificate Authority Policy** pop-up window appears.
2. In the **Name** field enter the name of the SSH certificate policy.

3 The **Enabled** check box is enabled by default. Uncheck the box to create a disabled policy. You can enable it at a later time.

4 The **Static Authorized Principals** field is optional. If a name is entered then any time the user logs in with the certificate, the principal in the certificate must also match the entered name. If no static authorized principals are set, then successful login will require that the SSH certificate has the username being logged in to as one of its principals. Note that you can add multiple name/principals in the static authorized principal list.

5 In the **CA Certificate** field select the CA certificate you wish to use for the policy or import a new CA certificate by clicking **Import CA Certificate** beneath the field. Leave this field blank if this policy is intended to be used with a CA public key

6 In the **CA Public Key** field select the CA public key you wish to use for the policy or import a new CA certificate by clicking **Add Public Key** beneath the field. Leave this field blank if this policy is intended to be used with a CA certificate.

7 Click **Create**.

Adding Users to an SSH Certificate Authority Policy

To add a user to a policy, perform the following:

1 From the **Policies** > **SSH Certificate Authority** tab, click the policy which you want the user added to.

2 In the **Members** panel, click **Add User...** from the menu tree. The **Add User** pop-up window will appear.

3 In the **Name** field enter a username and click **Add.** The username will now be listed under the **Members** panel.

Adding an Array to an SSH Certificate Authority Policy

To add an array to the policy, perform the following:

Note: Adding an array to a SSH Certificate Authority Policy will mean that admin users who do not explicitly have an SSH policy attached will be governed by this SSH policy. Additionally, enabling the SSH policy for an admin will auto disable the SSH public key access to the FlashBlade, if configured. Alternatively, enable SSH certificate access by attaching a policy directly to the user or attaching a policy to the array.

1 From the **Policies** > **SSH Certificate Authority** tab, click the policy which you want the array added to.

2 In the **Members** panel, click **Add Array...** from the menu tree. The **Add Array** pop-up window will appear.

3 In the **Name** field enter a username and click **Add**. The array will now be listed under the **Members** panel. The current name of the array is listed by default.

Network Access Policies

Network access policies are made up network access rules that govern client access to the FlashBlade interfaces (CLI, GUI, console, REST API,and SNMP).

Note: For network access policies, console access configuration is related to configuring/restricting superuser console acess. Network access policies do not manage console access for users with pureuser credentials.

The **Network Access Policies** panel alows you to view and manage network access policies and network access rules.

The **Network Access Policies** panel displays the following information for all configured network access policies:

- **Name**: The policy name.
- **Type**: The policy type.
- **Enabled**: If the policy is enabled (**True**) or disabled (**False**).

Network access policies are comprised of an ordered list of network access rules. A network access rule has a single client specification of IP address (IPv6 or IPv4), subnet in CIDR notation, or anonymous (*), and grants or denies access to FlashBlade interfaces.

Clicking a policy in the **Network Access Policies** panel will open the details page for rules associated with that policy.

Figure 6.1 Viewing Network Access Policy Rules

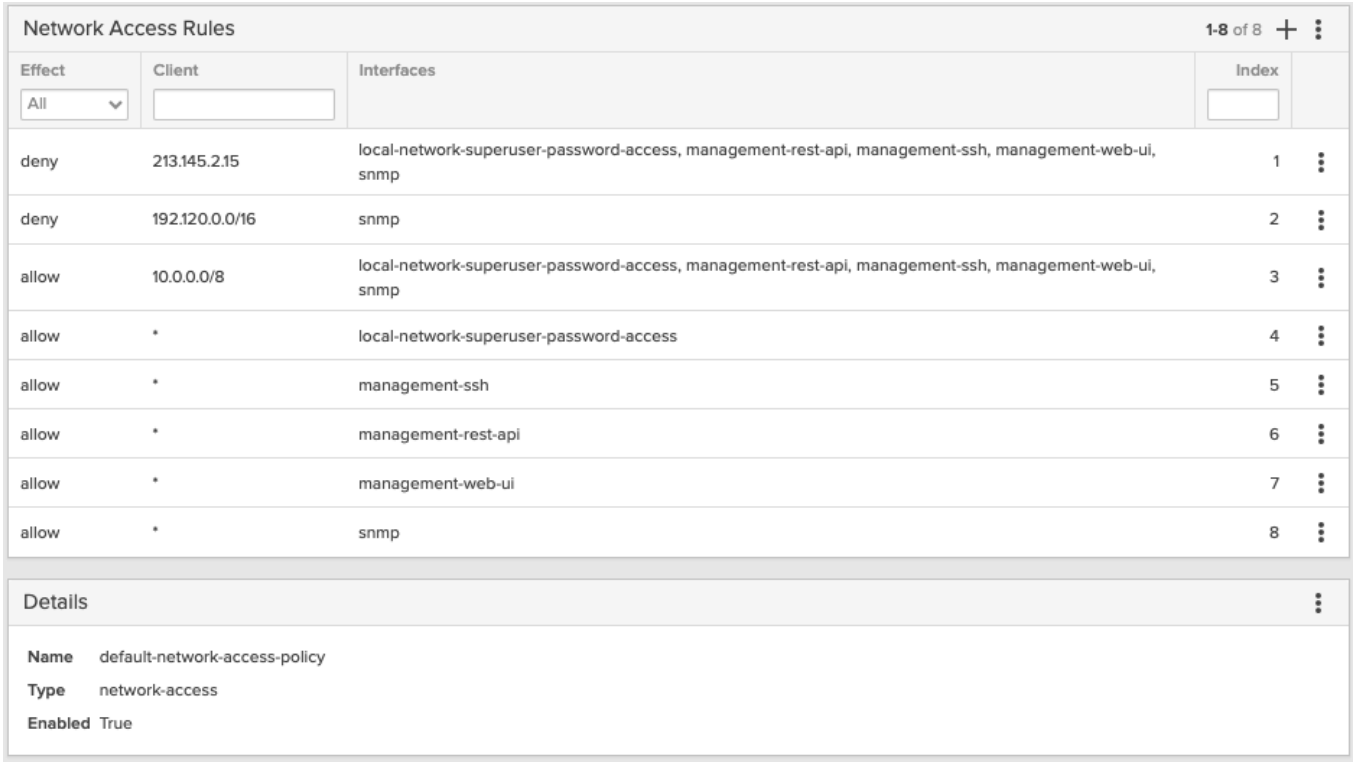

Information about each of the policy's rules is displayed in the **Network Access Rules** panel.

- **Effect**: Whether the client has access (**allow**) or not (**deny**) to an interface.
- **Client**: The client IP address or subnet (in CIDR notation). If anonymous, then * is displayed.
- **Interfaces**: The interface to which the rule is applied. Values can be as follows:

  Note: The **local-network-super-password-access** interface requires FlashBlade superuser credentials. Contact Pure Technical Services for additional information.

  - **local-network-superuser-password-access**: CLI, GUI, console
  - **management-rest-api**: REST API
  - **management-ssh**: CLI
  - **management-web-ui**: GUI
  - **snmp**: SNMP
- **Index**: The order in which the rule is executed, adhering to the following rule ordering restrictions:

- Deny rules are ordered before allow rules as explicit deny rules take precedent over allow rules.
- Rules with explicit IP addresses must be placed before anonymous (*).

The **Details** panel displays the policy name, policy type, and if it is currently enabled. It also allows you to rename the policy.

Note: If the system is configured to allow superuser access, five network access rules are included by default with the built-in policy, **default-network-access-policy**. Contact Pure Technical Services for additional information.

FlashBlade provides the built-in policy, **default-network-access-policy**. By default, four network access rules are included (five if the system is configured to allow superuser access), which allow access to all FlashBlade interfaces. This policy can be renamed and its rules can be removed and modified. Additional rules can also be added to the policy.

Renaming a Network Access Policy

To rename a network access policy, perform the following:

1. From the **Network Access Policies** panel, click **More Options** > **Rename** at the end of the row in which the policy you wish to rename appears in the **Network Access Policies** panel. The **Rename Network Access Policy** pop-up window appears.
2. Enter the new policy name.
3. Click **Rename**.

The **Network Access Policies** panel is updated with the policy's new name.

Adding a Network Access Rule to a Network Access Policy

To add a rule to a network access policy, perform the following:

1. Select the policy to which you wish to add a new rule in the **Network Access Policies** panel. The details page for that policy appears.
2. Click the **Add** (+) button from the header of the **Network Access Rules** panel. The **Add Network Access Rule to Policy** pop-up window appears.
3. In the **Client** field, enter the client as an IP address (IPv6 or IPv4), subnet in CIDR notation, or anonymous (*).
4. Select to **Allow** or **Deny** access to the specified client.

5 At the **Interfaces** field, click the **Add** (+) button. The **Add Interfaces** pop-up window appears.

   a Select the interface(s) to which access is being allowed or denied from the **Available Interfaces** column. Each selected interface appears in the **Selected Interface** column.

   b Click **Add** to add the selected interface(s) and return to the **Add Network Access Rule to Policy** pop-up window.

6 Click **Add**.

Editing a Network Access Rule

To edit a network access rule, perform the following:

1 Select the policy whose rules you wish to edit in the **Network Access Policies** panel. The details page for that policy appears.

2 From the **Network Access Rules** panel, locate the rule you wish to modify and select **More Options** > **Edit** at the end of row in which the rule appears. The **Edit Network Access Rule** pop-up window appears.

3 In the **Client** field, enter the client as an IP address (IPv6 or IPv4), subnet in CIDR notation, or anonymous (*). Note that ordering of network access rules is such that deny rules will be ordered before allow rules and rules with explicit IP addresses will be ordered before anonymous (*)

4 Select to **Allow** or **Deny** access to the specified client. Note that deny rules are ordered before allow rules as explicit deny rules take precedent over allow rules. For example, if there are two allow rules (rule 1 and rule 2), rule 2 cannot be changed to a deny rule.

5 At the **Interfaces** field, perform the following to remove or add interfaces:

   a Remove any interfaces by clicking the **Remove** (**x**) button next to the interface name.

   b Add any new interfaces by clicking the **Add** (+) button and then selecct the interface(s) to which access is being allowed or denied from the **Available Interfaces** column. Click **Add**.

6 Click **Save**.

Reordering Network Access Rules

Network access policies are comprised of an ordered list of network access rules.

Note the following network access rule ordering restrictions:

- Deny rules must precede allow rules.
- Rules with explicit IP addresses must be placed before anonymous (*).

Network access rules can be reordered after adding or editing rules in order to adhere to rule ordering restrictionsd, while ensuring continued system access.

To reorder a network access rule, perform the following:

1. Select the policy whose rules you wish to reorder in the **Network Access Policies** panel. The details page for that policy appears.
2. From the **Network Access Rules** panel, locate the rule you wish to move and select **More Options** > **Move** at the end of row in which the rule appears. The **Move Network Access Rule** pop-up window appears.
3. In the **Before** field, select which rule will execute AFTER the rule being currently modified.
4. Click **Move**.

The index column of the **Network Access Rules** panel is updated displaying the new order of the network access policy's rules.

Removing a Network Access Rule from a Network Access Policy

To remove a network access rule from a network access policy, perform the following:

1. Select the policy from which you wish to remove one or more network access rules in the **Network Access Policies** panel. The details page for that policy appears.
2. To remove a single network access rule:
   a. Select **More Options** > **Remove** at the end of row in which the rule appears. The **Remove Rule** pop-up window appears.
   b. Click **Remove**.
3. To remove multiple network access rules:
   a. From the **Network Access Rules** panel header, select **More Options** > **Remove**. The **Remove Network Access Rules** pop-up window appears.
   b. From the **Available Rules** column, select the rules you wish to remove. Each selected rule will appear in the **Selected Rules** column.
   c. Click **Remove**.

The **Network Access Rules** panel is updated displaying the remaining rules.

TLS Policies

Transport Layer Security (TLS) is a cryptographic protocol that encrypts data in transit on FlashBlade systems. TLS Policies can govern the properties and requirements of any TLS negotiation that they apply to. But TLS Policies do not force clients to use TLS in protocols where it is not supported, such as SMB, or not mandatory by default, such as NFS. For protocols such as NFS, please refer to the relevant protocol policies (NFS policies or SMB policies) for controls that allow you to make TLS mandatory. TLS can be used for a variety of inbound traffic, including, but not limited to:

- Manageability via web browser or REST API
- S3 data traffic
- S3 IAM/STS traffic
- Replication
- NFS over HTTP

TLS versions and ciphers for inbound traffic to FlashBlade can be controlled for your environment through TLS policies. This includes the ability to disable old TLS versions and ciphers deemed insecure, re-enable them if needed, and select identity certificates that the array will present to clients during TLS negotiation. For applications that require stronger security through bidirectional certificate verification, FlashBlade supports mutual TLS (mTLS), ensuring both the client and server authenticate each other during the TLS handshake.

A default TLS policy (**default-tls-policy**) is included with FlashBlade.

The default TLS policy uses special “default” values for enabled ciphers and the minimum TLS version. You may also use “default” as a cipher to enable, or minimum TLS version, when creating and modifying policies.

The value of “default” is determined for any given software version by Pure Storage, and will be updated with software upgrades in the future to comply with latest security practices. If it is used in a policy, it allows you to pick up new security improvements automatically upon upgrade, such as removal of insecure ciphers. The “default” enabled cipher can be used in combination with other enabled ciphers, or in combination with disabled cipher specifications, in order to add to or remove from the default set of ciphers.

The exact meaning of “default” on a given software version can be viewed by navigating to **Policies** > **Network Security** > **TLS** > **TLS Defaults** panel.

Summary details and membership information about TLS policies can be viewed by clicking a policy name in the **TLS Policies** panel.

Figure 6.2 Viewing TLS Policy Summary Details

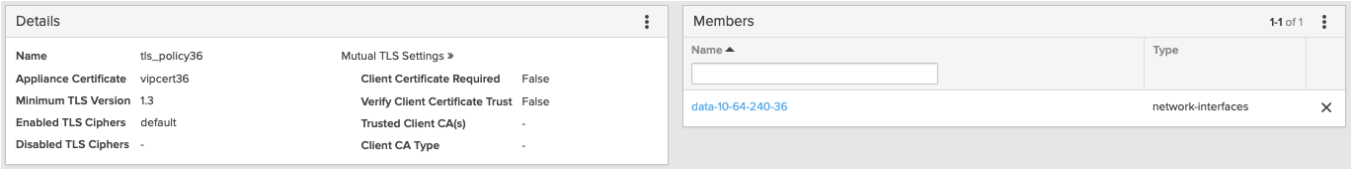

A TLS policy's effective values can be viewed in the **Details** panel (see *Viewing a TLS Policy's Effective Values* for instructions). Additionally, TLS policies can be attached to, and removed from, network interfaces from the **Members** panel (see *Attaching TLS Policies to Network Interfaces* and *Removing TLS Policies from Network Interfaces* for instructions).

Custom TLS policies for inbound traffic can be created and managed on FlashBlade, allowing you control over the minimum TLS version and ciphers required for your environment's services. Mutual TLS can also be enabled with custom TLS policies to provide additional security. Use the **TLS Policies** panel to create and manage custom TLS policies. See *Creating a TLS Policy*.

Creating a TLS Policy

Custom TLS policies can be created for inbound traffic. Different requirements can be set for different services on the system.

The following restrictions apply:

- Ciphers and cipher suites must be specified in OpenSSL format.
- 64 individual ciphers can be specified.
- Ciphers containing PSK and SRP are not allowed.
- Appliance certificates must be "array"-type certificates.

Additionally, custom TLS policies can be created with mutual TLS authentication to provide enhanced security. Note that your browser must be configured with a proper client certificate before attaching a TLS policy with mutual TLS to the management vip. If not configured, access to the GUI will be lost until the browser is configured.

To create a custom TLS Policy, perform the following:

1. Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click Add (+). The **Create TLS Policy** pop-up window appears.
2. In the **Name** field, enter the policy name.
3. In the **Appliance Certificate** field, select the appliance certificate.
4. In the **Minimum TLS Version** field, select the minimum TLS version. If not selected, the default minimum TLS version is “default”
5. In the **Enabled Ciphers** field, click Add (+). Enter a comma-separated list of ciphers/cipher suites you wish to be enabled with this policy and click **OK**. If not specified, the default enabled ciphers will be “default”.
6. In the **Disabled Ciphers** field, click Add (+). Enter a comma-separated list of ciphers/cipher suites you wish to be disabled with this policy and click **OK**.
7. (Optional) Click the down arrow next to **Mutual TLS Settings** to display mutual TLS options.
   a. Select the **Client Certificates Required** checkbox to require client certificates during TLS negotiation.
   b. Select the **Verify Client Certificate Trust** checkbox to require the validation of client certificates against the trusted CA.
   c. Perform one of the following:
      - From the **CA Certificate** list, select the CA certificate used to verify client authenticity.
      - From the **CA Certificate Group** list, select the CA certificate group used to verify client authenticity.
8. Click **Create**.

Editing a TLS Policy

Note: Pure’s default TLS policy, **default-tls-policy**, can be combined with other specified ciphers in the **Enabled Ciphers** field.

The default TLS policy, as well as custom TLS policies, can be edited as your environment and security requirements evolve over time.

The following restrictions apply:

- Ciphers and cipher suites must be specified in OpenSSL format.
- 64 individual ciphers can be specified.

- Ciphers containing PSK and SRP are not allowed.
- Appliance certificates must be “array”-type certificates.

To edit a TLS Policy, perform the following:

1. Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click the Edit button in the same row as the policy you wish to edit. The **Edit TLS Policy** pop-up window appears.
2. Modify values as needed.
3. If you wish to enable or disable mutual TLS, click the down arrow next to **Mutual TLS Settings** to display the mutual TLS options and modify as needed.
4. Click **Save**.

Deleting a TLS Policy

Note: The array's default TLS policy, **default-tls-policy**, cannot be deleted.

1. To delete a single TLS policy:
   a. Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click Delete button in the same row as the policy you wish to delete. The **Delete TLS Policy** pop-up window appears.
   b. Click **Delete**.
2. To delete multiple TLS policies:
   a. Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click **More Options** > **Delete Polices** in the header of the panel. The **Delete TLS Policies** pop-up window appears.
   b. Select the TLS policies you with to delete from the **Available Policies** panel.
   c. Click **Delete**.

Viewing a TLS Policy's Effective Values

To view the effective values for a TLS Policy, perform the following:

1. Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click policy whose effective values you wish to view. The details screen for that policy appears.
2. In the header of the **Details** panel, click **More Options** > **List effective ciphers**. The content of the **Details** panel updates to display all effective ciphers configured for the TLS policy.

Setting the Array's Default TLS Policy

Any TLS policy can be set as the array default policy and applied to all network interfaces that do not have a TLS policy attached to them. Once a TLS policy is set as the array's default TLS policy, it cannot be disabled or deleted unless you set another policy to replace it as the default policy.

Before configuring a TLS policy as the array's default policy, review the TLS policy's effective values to verify that the policy is the correct policy to apply as the array's default (see *Viewing a TLS Policy's Effective Value*).

To set an array's default TLS policy, perform the following:

1. Navigate to **Settings** > **System** > **Attached Policies**.
2. In the **Attached Policies** panel header, click **More Options** > **Set Default TLS Policy**. The **Set Default TLS Policy** pop-up window appears.
3. Select the TLS policy you wish to set as the array's default from the **New TLS Policy** field.
4. Click **Save**.

Attaching TLS Policies to Network Interfaces

To attach a TLS policy to one or more network interfaces, perform the following:

1. Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click policy you wish to attach to a network interface. The details screen for that policy appears.
2. In the header of the **Members** panel, click **More Options** > **Add Network Interfaces**. The **Add Network Interfaces** pop-up window appears.
3. Select the network interfaces to which you wish to attach the TLS policy from the **Available Network Interfaces** panel.
4. Click **Add**.

Removing TLS Policies from Network Interfaces

1. To remove a TLS policy from a single network interface:
   a. Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click TLS policy you wish to remove from a network interface. The details screen for that TLS policy appears.
   b. In the **Members** panel, click the delete button in the same row of the interface from which you wish to remove the TLS policy. The **Remove Member from TLS Policy** pop-up window appears.

c Click **Remove**.

2 To remove a TLS policy from multiple network interfaces:

a Navigate to **Policies** > **Network Security** > **TLS**. In the **TLS Policies** panel, click TLS policy you wish to remove from multiple network interfaces. The details screen for that TLS policy appears.

b In the header of the **Members** panel, click **More Options** > **Remove**. The **Remove Members from TLS Policy** pop-up window appears.

c Select the interfaces from which the TLS policy will be removed from the **Current Members** panel.

d Click **Remove**.

Chapter 7
The Purity//FB GUI Analysis Page

The **Analysis** page displays historical array data, such as I/O performance trends and storage capacity and consumption data. It also provides network port statistics. Real time and historical performance data is presented in charts.

At the top right corner of each chart is a legend depicting the type of metrics dispalyed in the chart.

Clicking a metric from a chart legend filters out the graph for that metric. The filtered metric legend item will appear dimmed and its corresponding graph will be hidden. Multiple metrics can be filtered out at any given time.

Performance

The **Analysis** > **Performance** page displays a series of rolling charts consisting of real-time performance metrics for the array.

The curves in each chart are comprised of a series of individual data points. Hover over any part of a chart to display values for a specific point in time. The values that appear in the point-in-time tooltips are rounded to two decimal places.

Figure 7.1 Viewing the Performance Page

The Performance Chart

The following example displays the IOPS statistics on the array at precisely `5:00:00` on `March 12`.

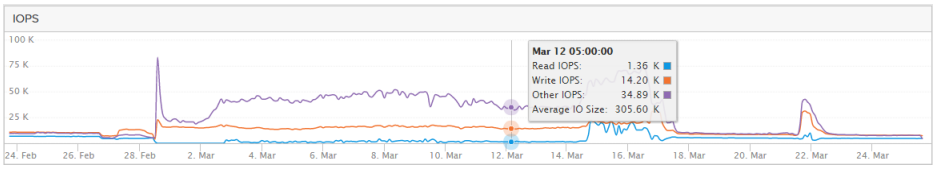

The Performance panel includes the Latency, IOPS, and Bandwidth charts.

Latency

The Latency chart displays the average latency times for various operations.

- **Read Latency (R)** - Average arrival-to-completion time, measured in milliseconds, for a read operation.

- **Write Latency (W)** - Average arrival-to-completion time, measured in milliseconds, for a write operation.
- **Other Latency (O)** - Average arrival-to-completion time, measured in milliseconds, for all other metadata operations.

IOPS

The IOPS (Input/output Operations Per Second) chart displays I/O requests processed per second by the array. This metric counts requests per second, regardless of how much or how little data is transferred in each.

- **Read IOPS (R)** - Number of read requests processed per second.
- **Write IOPS (W)** - Number of write requests processed per second.
- **Other IOPS (O)** - Number of metadata operations processed per second.
- **Average IO Size** - Average I/O size per request processed. Requests include reads and writes.

Bandwidth

The Bandwidth chart displays the number of bytes transferred per second to and from all file systems. The data is counted in its expanded form rather than the reduced form stored in the array to truly reflect what is transferred over the storage network. Metadata bandwidth is not included in these numbers.

- **Read Bandwidth (R)** - Number of bytes read per second.
- **Write Bandwidth (W)** - Number of bytes written per second.

By default, the performance charts, as described above, are displayed for the array. Click the **File Systems** or **Object Store** tabs to view file system or object store-specific latency, IOPs, and bandwidth performance charts.

The **File Systems** and **Object Store** tabs allow you to view the performance charts for one or more selected NFS file systems or buckets, respectively.

Displaying Protocol-Specific Metrics

By default, the performance charts display data for all supported protocols, as indicated by the **All** button in the upper-right corner of the **Array** tab. Click a protocol button to display only the performance information for that protocol.

The performance charts display data in a user-friendly view that is common for all protocols.

For HTTP and S3 protocol data, there are two display options:

- **General**: The Latency, IOPS, and Bandwidth graphs display operations that the array is doing in chunks. Default.
- **Protocol Specific**: The Latency, IOPS, and Bandwidth graphs display a protocol-specific number of user operations.

For example, for HTTP data, in the **General** view, operations such as listing a directory and deleting are both grouped in the Other line. In the **Protocol Specific** view, read and write operations on a directory are split out into their own line.

The file transfer can be performed in many chunks for one large file. The **General** view shows one operation per chunk transfer, while the **Protocol Specific** view shows only one operation per file. When transferring small files, the number of operations is the same in both the **General** and **Protocol Specific** views.

Displaying File System Performance

Note: Views of write operations may be reported as high because the system performs confirmation operations internaly within the metrics service.

Peformance metrics per file system are available from the **Analysis > Performance > File Systems** page.

Note: If Native SMB mode is enabled, the **All**, **NFS**, and **SMB** buttons are available. If not enabled, only NFS file systems are displayed.

The **Analysis > Performance > File Systems** page provides four panes: The **File Systems** pane provides a list of all file systems pre-sorted by most to least used. The **Latency**, **IOPs**, and **Bandwidth** panes provide performance charts that display the performance of file systems selected from the **File Systems** pane. If Native SMB mode is enabled, by default the **All** button is selected and metrics are displayed for all protocols; you can further filter by protocol by selecting **NFS** or **SMB**.

Figure 7.2 Displaying File System Performance (NFS)

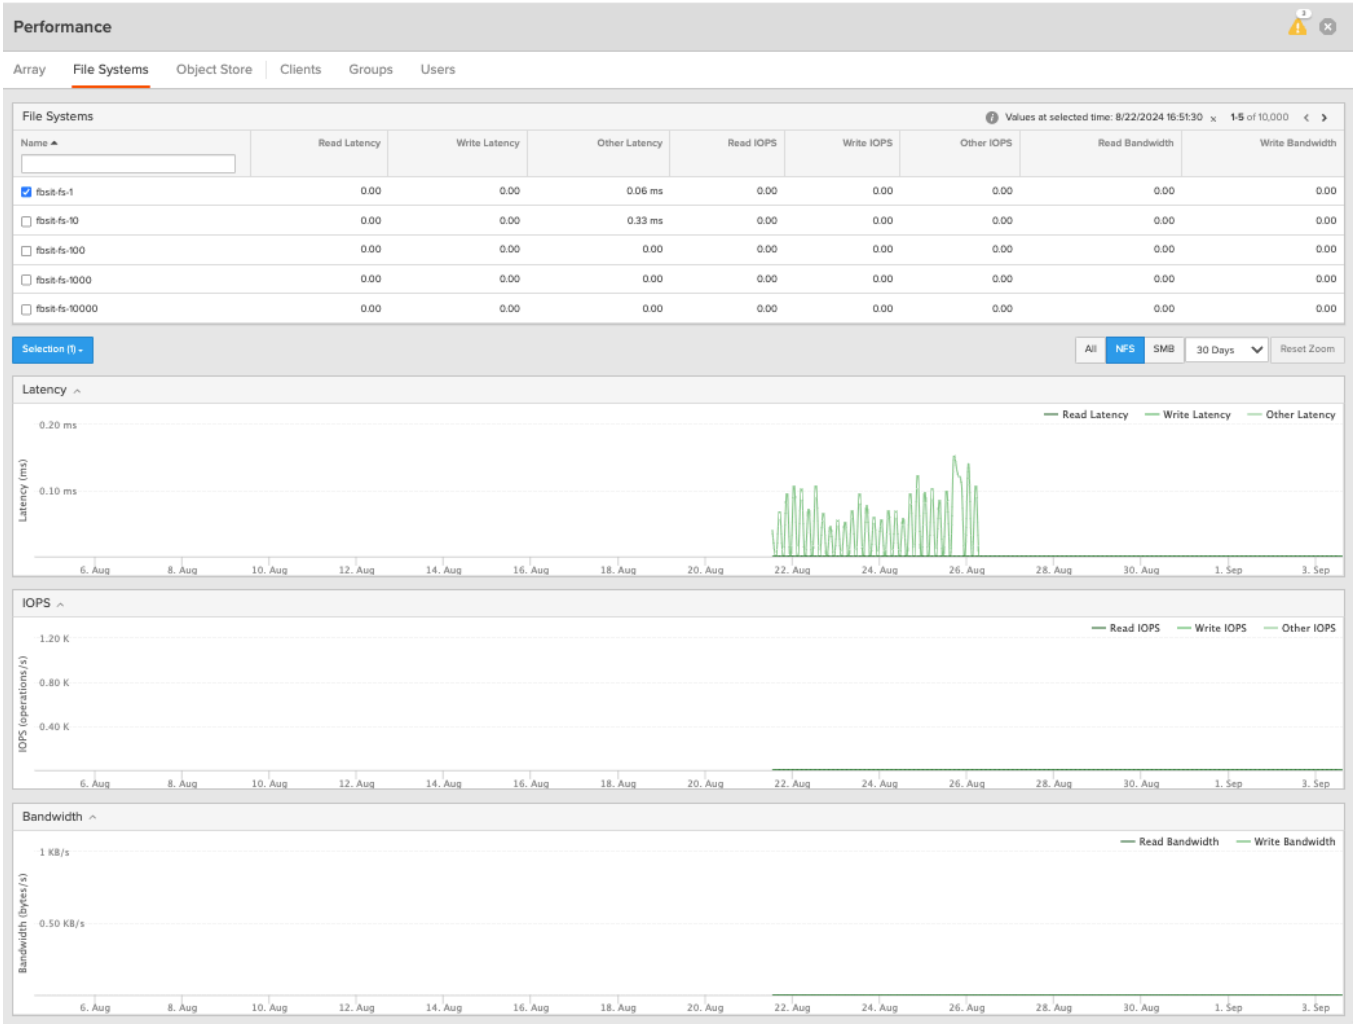

Identifying the Most Active File Systems

The **Performance** page allows you to view your FlashBlade's most active file systems. This can be useful in scenarios where you suspect some file systems may be impacting the workload on other file systems by allowing you to quickly view the IOPs, throughput, and latency of the most active file systems.

Note: The graphs on the **Analysis** page display rolling charts consisting of real-time performance metrics (summed) forfile systems, and buckets.

When workload performance issues are observed, the **Peformance** > **Analysis** page allows you to quickly assess if there are any spikes in the **Latency**, **IOPs**, or **Bandwidth** graphs and view the timestamp of the performance abnormality.

Once the timestamp is determined, the **Analysis > Performance > File Systems** page allows you to identify which file systems contributed to the performance issue and when the issue occurred. The **File Systems** pane provides a list of all file systems that are pre-sorted by most to least used. By default the performance values displayed in the table reflect the current time. The values automatically update and the each column can be sorted.

1. Select one or more file systems from the **File System** pane.
2. At the top of the performance graph panes, depending on the timestamp observed from the **Performance > Analysis** page, select the date range you wish to be displayed in the graphs.
3. In the performance graphs, click the spike in the graph to display the timestamp and performance metrics of the abnormality.

In the following example, file system `fbsit-fs-1` is determined to have experienced a 0.14ms latency spike not related to reads or writes.

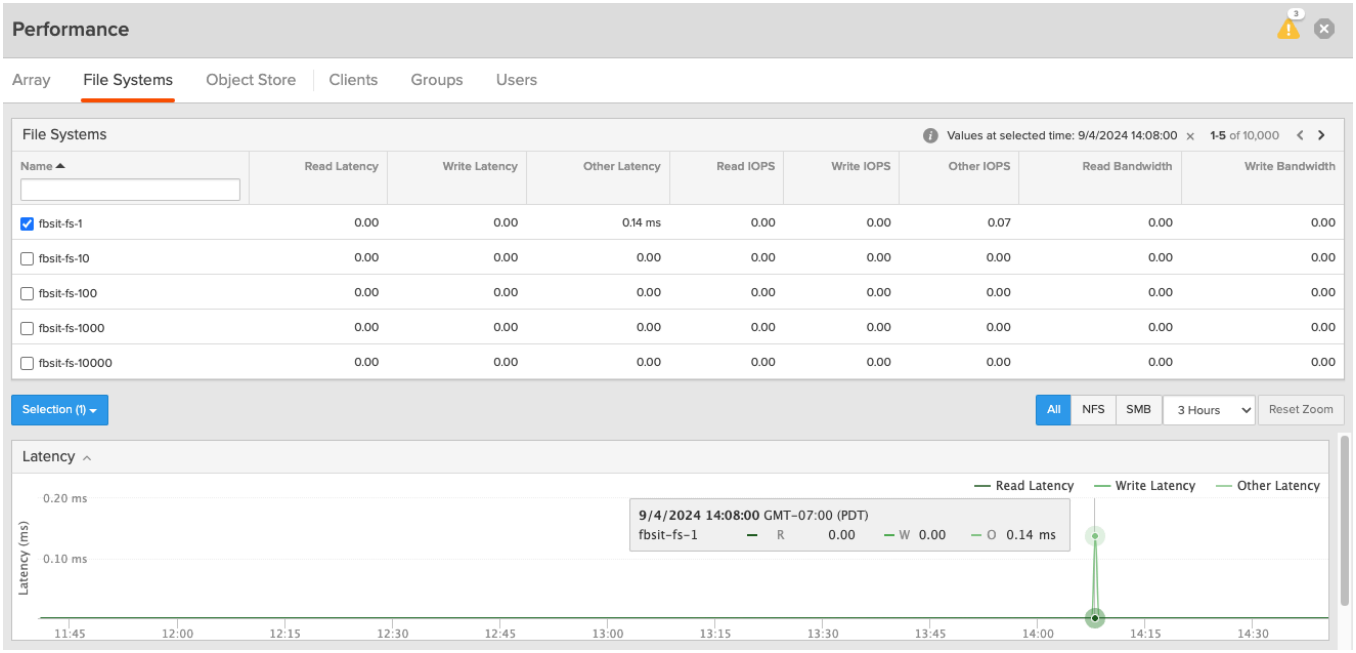

Note about the Performance Charts

The **Dashboard** and **Analysis** pages display the same latency, IOPS, and bandwidth performance charts, but the information is presented differently between the two pages.

In the **Dashboard** page:

- The performance charts are displayed for the array.

- The performance charts are updated once every second.
- The performance charts display up to 5 minutes of historical data.

In the **Analysis** page:

- The performance charts are provided for the array.
- At its shortest range (5m), the performance charts are updated once every second. As the range increases, the update frequency (and resolution) decreases.
- The performance charts display up to one year of historical data.
- The performance charts can be filtered to display metrics by protocol.

Displaying Client, Group, and User Performance Statistics

In addition to displaying array, file system, and object store performance charts, the **Performance** page also provides performance statistics for clients (NFS and S3), groups (NFS), and users (NFS).

Client, group, and user performance statistics are accessed by clicking the **Clients**, **Groups**, and **Users** tabs on the **Performance** page.

Note that client, group, and user performance statistics do not automatically refresh. To reload the performance data, click **Reload Cache**. Filters other than names are ignored when reloading the cache. When reloading the cache, the data is retrieved based on the specified sort order.

Viewing NFS Client Performance Statistics

To view statistics for NFS clients, click **Reload Cache** in **Performance by Clients** table header. The following information is displayed:

- **Name** - The NFS client name.
- **B/s (read)** - The average number of bytes per read operation.
- **B/s (write)** - The average number of bytes per write operation.
- **op/s (read)** - The average number of read operations processed per second.
- **op/s (write)** - The average number of write operations processed per second.
- **op/s (other)** - The average number of metadata operations processed per second.

- **us/op (read)** - The average time, measured in microseconds, to process a read request.
- **us/op (write)** - The average time, measured in microseconds, to process a wite request.
- **us/op (other)** - The average time, measured in microseconds, to process a metadata operation request.
- **B/op (read)** - The average bytes per read operation processed.
- **B/op (write)** - The average bytes per write operation processed.
- **B/op (all)** - The average byte size per read, write, and all operations processed.

As with other tables provided the FlashBlade GUI, you can search each column by entering a value in the field below each column header, as well as sort by clicking the column header.

The data can be downloaded in `.csv` format by clicking **Download CSV**.

Viewing S3 Client Performance Statistics

General or protocol specific S3 client performance statistics can be displayed.

To view general statistics for S3 clients:

1. Select **S3** and **General**.
2. Click **Reload Cache** in **Performance by Clients** table header.

The following information is displayed:

- **Name** - The S3 client name.
- **B/s (read)** - The average number of bytes per read operation.
- **B/s (write)** - The average number of bytes per write operation.
- **op/s (read)** - The average number of read operations processed per second.
- **op/s (write)** - The average number of write operations processed per second.
- **op/s (other)** - The average number of metadata operations processed per second.
- **us/op (read)** - The average time, measured in microseconds, to process a read request.
- **us/op (write)** - The average time, measured in microseconds, to process a write request.

- **us/op (other)** - The average time, measured in microseconds, to process a metadata operation request.
- **B/op (read)** - The average bytes per read operation processed.
- **B/op (write)** - The average bytes per write operation processed.
- **B/op (all)** - The average byte size per read, write, and all operations processed.

To view protocol specific S3 client performance:

1. Select **S3** and **Protocol Specific**.
2. Click **Reload Cache** in **Performance by Clients** table header.

The following information is displayed:

- **Name** - The S3 client name.
- **RB/s** - The average number of bytes per bucket read operation.
- **WB/s** - The average number of bytes per bucket write operation.
- **RO/s** - The average number of bytes per object read operation.
- **WO/s** - The average number of bytes per object write operation.
- **O/s** - The average number of bytes per non-data (metadata and delete) operations.
- **us/RB** - The average time to process a bucket read request.
- **us/WB** - The average time to process a bucket write request.
- **us/RO** - The average time to process an object read request.
- **us/WO** - The average time to process an object write request.
- **us/O** - The average time to process non-data (metadata and delete) operations.

As with other tables provided the FlashBlade GUI, you can search each column by entering a value in the field below each column header, as well as sort by clicking the column header.

The data can be downloaded in `.csv` format by clicking **Download CSV**.

Viewing Group and User Performance Statistics

To view statistics for groups and users:

1. Click **Select File System** above the **Performance by Groups (NFS)**, or **Performance by Users (NFS)** table header.
2. Select a file system, and click **Select**.
3. Click **Reload Cache**.

The following information is displayed:

- **Group ID** - For group statistics only. The group ID.
- **Group Name** - For group statistics only. The group name.
- **User Id** - For users statistics only. The user ID.
- **User Name** - For users statistics only. The user name.
- **B/s (read)** - The average number of bytes per read operation.
- **B/s (write)** - The average number of bytes per write operation.
- **op/s (read)** - The average number of read operations processed per second.
- **op/s (write)** - The average number of write operations processed per second.
- **op/s (other)** - The average number of metadata operations processed per second.
- **us/op (read)** - The average time, measured in microseconds, to process a read request.
- **us/op (write)** - The average time, measured in microseconds, to process a write request.
- **us/op (other)** - The average time, measured in microseconds, to process a metadata operation request.
- **B/op (read)** - The average bytes per read operation processed.
- **B/op (write)** - The average bytes per write operation processed.
- **B/op (all)** - The average byte size per read, write, and all operations processed.

As with other tables provided in the FlashBlade GUI, you can search each column by entering a value in the field below each column header, as well as sort by clicking the column header.

The data can be downloaded in `.csv` format by clicking **Download CSV**.

Capacity

By default, the **Capacity** graph displays information for all file systems on the array.

The **Capacity** graph displays the following array-wide capacity and consumption information:

- **Total Physical**: Total storage space used on the array.
- **Capacity**: Total storage capacity of the array.
- **Data Reduction**: Ratio of the size of the written data (i.e. Used) versus the amount of space the data occupies after data compression (i.e. Physical) and/or deduplication
- **Unique**: Total space that the written data occupies on the array's file systems. The value displayed does not include Destroyed space.
- **Destroyed**: The sum of destroyed unique space of all the destroyed file systems).
- **Snapshots**: Amount of space that the written data occupies on the array's snapshots after reduction via data compression.

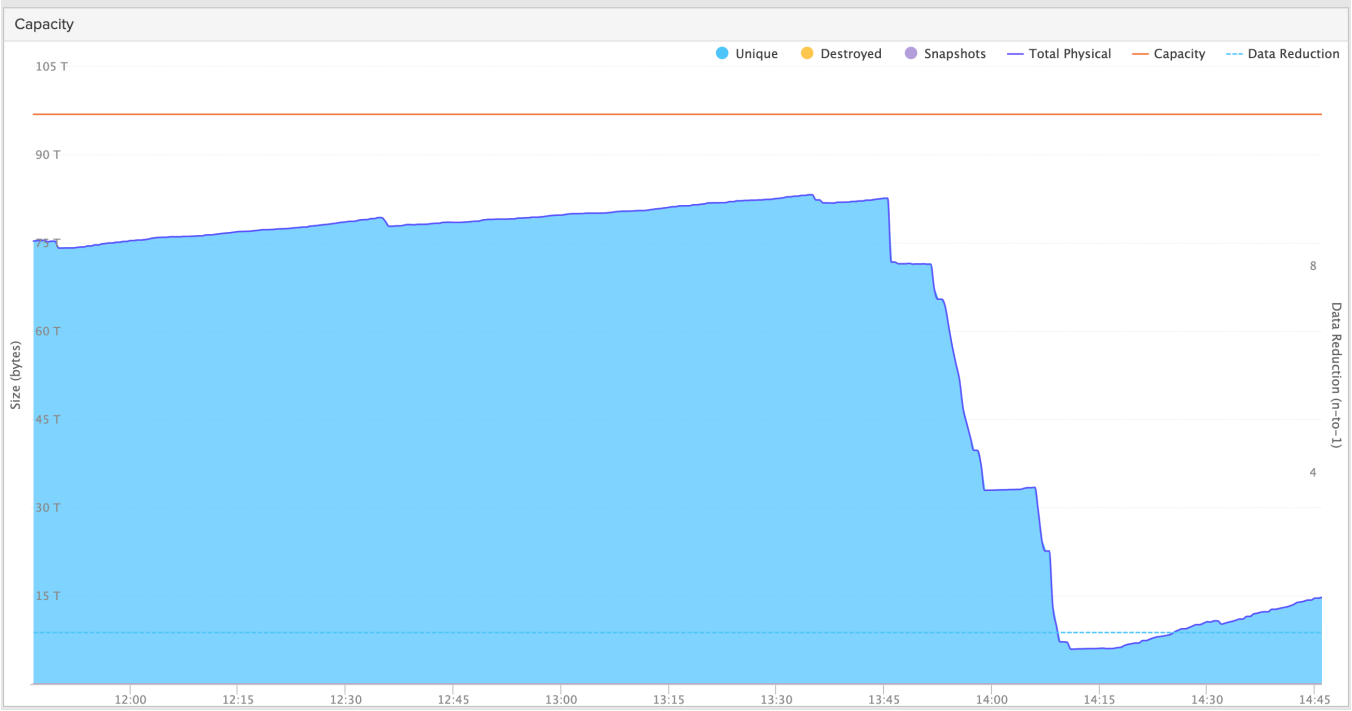

Hover over any part of a chart to display values for a specific point in time. The size values that appear in the point-in-time tooltips are rounded to two decimal places.

By default, the Capacity graph displays data for the past 3 hours, as indicated by the time range drop-down button. Click the drop-down button to view capacity data over a different time range.

Purity DeepReduce™

Purity DeepReduce™ is a data reduction feature running Global Nearline Similarity based algorithms to generate higher data reduction rates in the given dataset on FlashBlade//E systems. Built to handle today's surge of pre-reduced, immutable, and application-compressed datasets, this advanced similarity-based reduction engine helps unlock meaningful efficiency gains across challenging workloads. Engineered for reduction at scale, **DeepReduce**™ seamlessly augments as a background process and works in conjunction with FlashBlade//E's existing inline compression technologies to maximize effective capacity. With **significantly higher data reduction than conventional deduplication**, customers can now achieve greater efficiency, lower TCO, and more predictable storage economics at scale. Progress and effectiveness can be monitored, with metrics such as percentage of data blocks scanned by DeepReduce™ and achieved data reduction ratios reported.

Network

By default, the **Analysis > Network** page displays the **Array** tab. The **Array** tab displays historical array-level **Errors/s**, **Bandwidth**, and **Packets/s** charts.

Figure 7.3 Viewing the Network Page

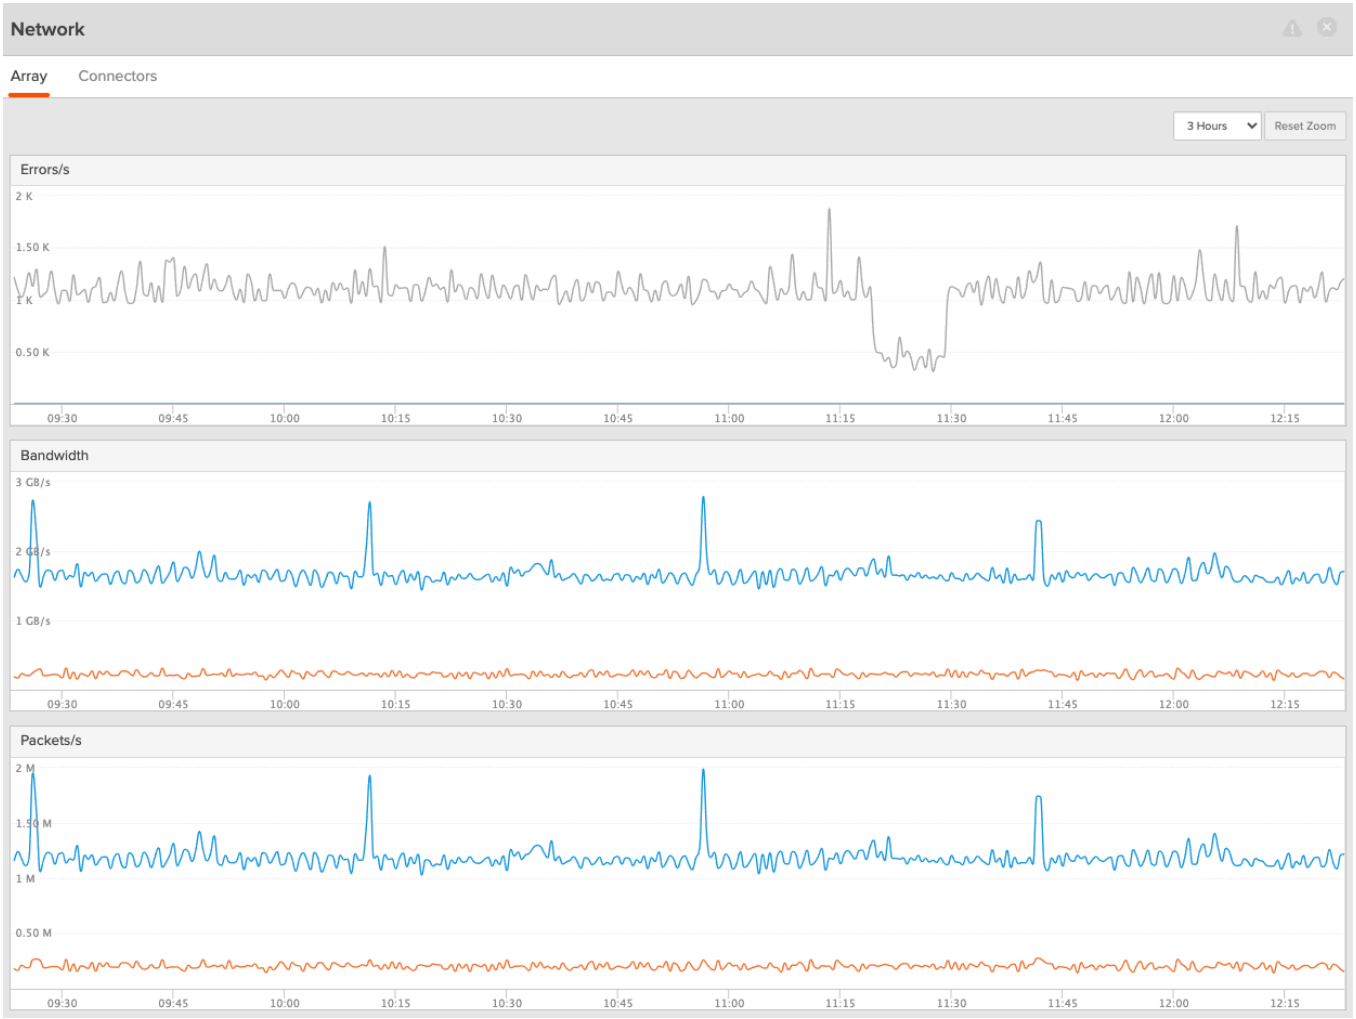

The **Errors/s** charts displays the following information:

- **CRC Error/ss**: Indicates the number of received packets with incorrect checksums. A cycle redundancy check (CRC) is an error-detecting code for data transmission.
- **Frame Errors/s**: Indicates the number of received packets per second with misaligned Ethernet frames.

- **Carrier Errors/s**: Indicates the number of transmitted packets per second with duplex mismatch or faulty hardware issues.
- **Dropped Errors/s**: Indicates the number of transmitted packets per second that were dropped.
- **Other Errors/s**: Indicates the number of packets per second with all other types of receive and transmit errors such as Jabber, InDiscards, InErrs, and MacTrasmitErrs.
- **Total Errors/s**: The error history over the selected range of time, annotated with the number of total errors per second for individual (or the sum of all) interfaces at a specific point in time.

Figure 7.4 Network Errors/s Chart

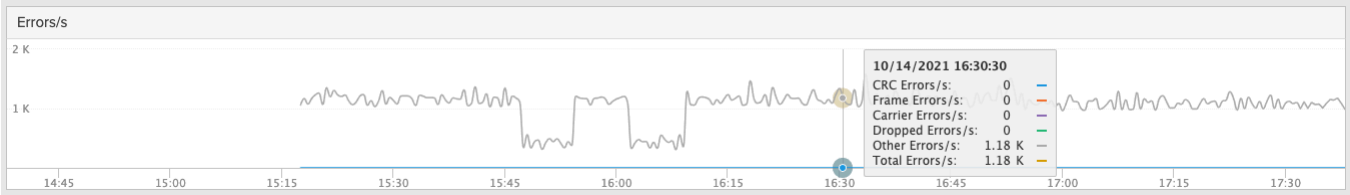

The **Bandwidth** chart displays the graphical representation of the bandwidth history over the selected range of time, annotated with the numbers of the transmitted bytes and received bytes per second for individual (or the sum of all) interfaces at a specific point in time.

Figure 7.5 Network Bandwidth Chart

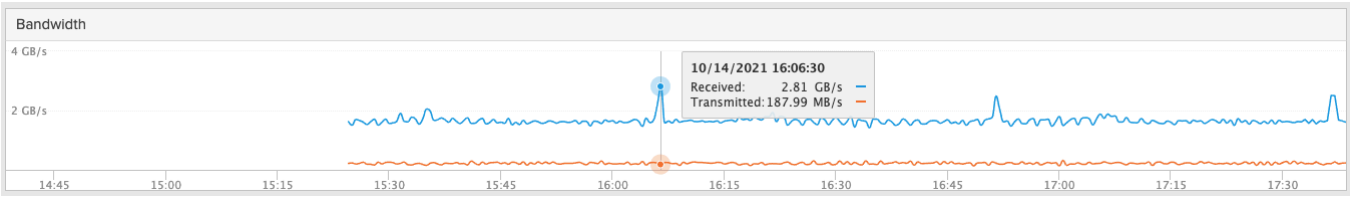

The **Packets/s** chart displays the graphical representation of the historical packet information over the selected range of time, annotated with the numbers of transmitted packets and received packets per second for individual (or the sum of all) interfaces at a specific point in time.

Figure 7.6 Network Packets/s Chart

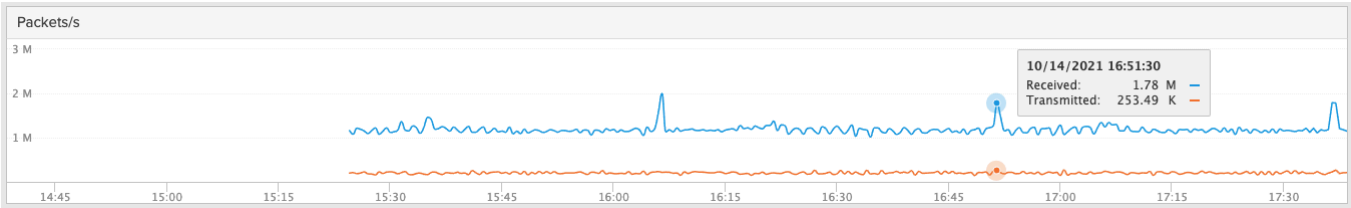

By default, the charts display data for the past 3 hours, as indicated by the time range drop-down button. Click the drop-down button to view data over a different time range.

Viewing Network Connector Statistics

The **Analysis > Network** page includes a **Connectors** tab which allows you to view performance statistics for the array's QSFP ports in tabular and graphical form. Unlike the historical statistics displayed in the **Array** tab, the **Connectors** tab displays instant values for each port.

Figure 7.7 Viewing Network Connector Statistics

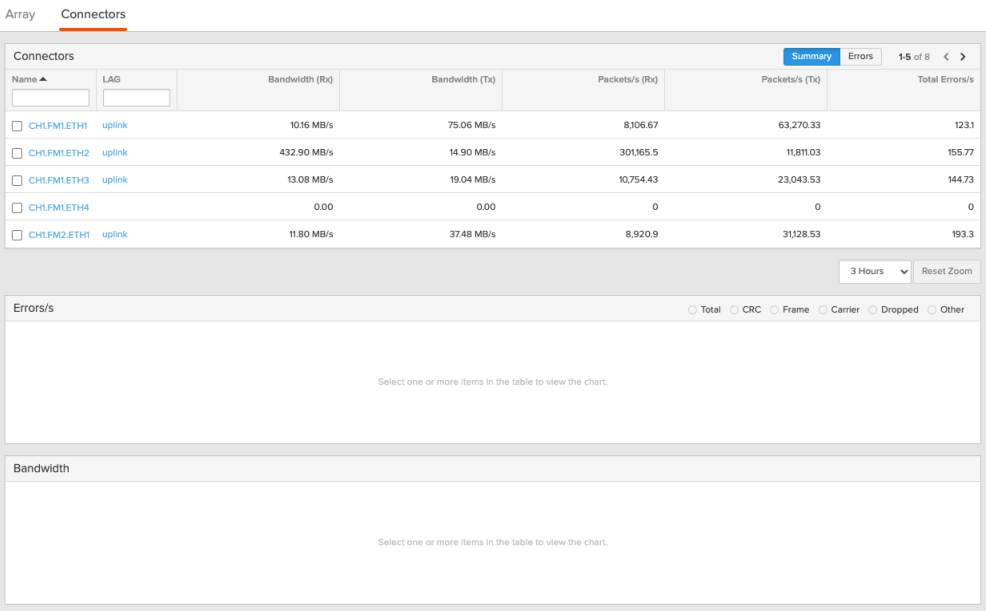

The **Connectors** panel provides a table that displays the following information:

- **Name**: The QSFP port name in the format `CHx.FMx.ETHx`, where `CHx` is the chassis number, `FMx` is the fabric module number, and `ETHx` is the ethernet port number.
- **LAG**: The subnet link aggregation group name.
- **Bandwidth (Rx)**: The number of received bytes per second.
- **Bandwidth (Tx)**: The number of transmitted bytes per second.
- **Packets/s (Rx)**: The number of received packets.
- **Packets/s (Tx)**: The number of transmitted packets.

- **Total Errors/s**: Displays the graphical representation of the error history over the selected range of time, annotated with the number of total errors per second for individual (or the sum of all) interfaces at a specific point in time.

You can select one or more connectors from the **Connectors** table to display a graphical representation of their error, bandwidth, and packet statistics in the **Errors/s**, **Bandwidth**, and **Packets/s** charts.

Figure 7.8 Viewing Network Connector Charts

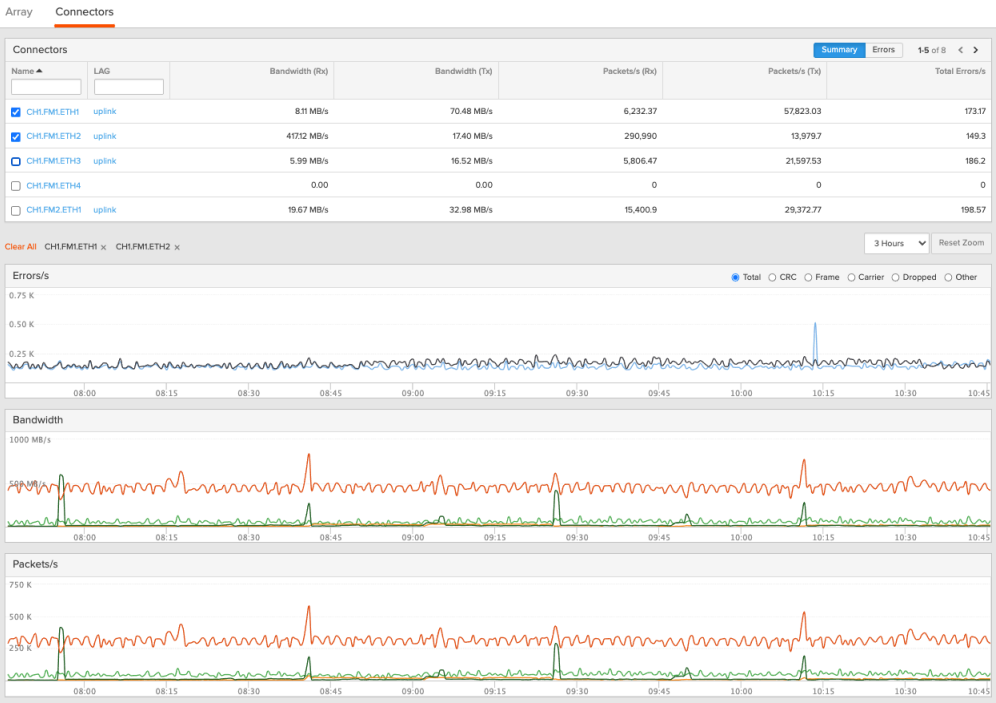

The **Errors/s** chart displays the graphical representation of the error history over the selected range of time, annotated with the number of total errors per second for each of the selected interfaces at a specific point in time. By default, the **Errors/s** chart displays **Total** errors. Buttons are provided at the top of the **Errors/s** chart allowing you to change the view of the chart to show:

- **Total**: The number of total errors per second for each selected interface.
- **CRC**: The number of received packets per second with incorrect checksums. A cyclic redundancy check (CRC) is an errordetecting code for data transmission.
- **Frame**: The number of received packets per second with misaligned Ethernet frames.
- **Carrier**: The number of transmitted packets per second with duplex mismatch or faulty hardware issues.

- **Dropped**: The number of transmitted packets per second that were dropped.
- **Other**: The number of packets per second with all other types of receive and transmit errors such as Jabber, InDiscards, InErrs, and MacTrasmitErrs.

The **Bandwidth** chart displays the graphical representation of the bandwidth history over the selected range of time, annotated with the numbers of the transmitted bytes and received bytes per second for each selected at a specific point in time.

The **Packets/s** chart displays the graphical representation of the historical packet information over the selected range of time, annotated with the numbers of transmitted packets and received packets per second for each selected interface at a specific point in time.

Like other FlashBlade charts, by default the charts display data for the past 3 hours, as indicated by the time range drop-down button. Click the drop-down button to view data over a different time range.

Displaying a Different Time Range

By default, the **Performance**, **Capacity**, and **Replication** charts display data for the past 3 hours, as indicated by the time range drop-down button. Click the drop-down button to view performance data over a different time range.

Figure 7.9 Selecting a Time Range

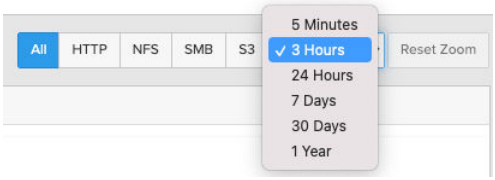

Chapter 8
The Purity//FB GUI Health Page

The **Health** page displays information to help gauge the health of the array.

Hardware

The **Health** > **Hardware** page contains information about the array hardware components.

The title bar of the hardware panel includes the array name, the raw capacity value, and parity information. The raw capacity value of the array is displayed in gibibyte (G), gigabyte (GB), tebibyte (TiB), terabyte (TB), pebibyte (P), and petabyte (PB) units. Note that usable capacity is reflected on the **Dashboard**. The parity value represents the percentage of data that is fully protected. The value will drop below 100% if the data is not fully protected, such as when a blade or drive is unhealthy leading to evacuation.

The hardware panel graphically displays the status of each hardware component, which is useful for diagnosing hardware-related problems.

Figure 8.1 Health Page

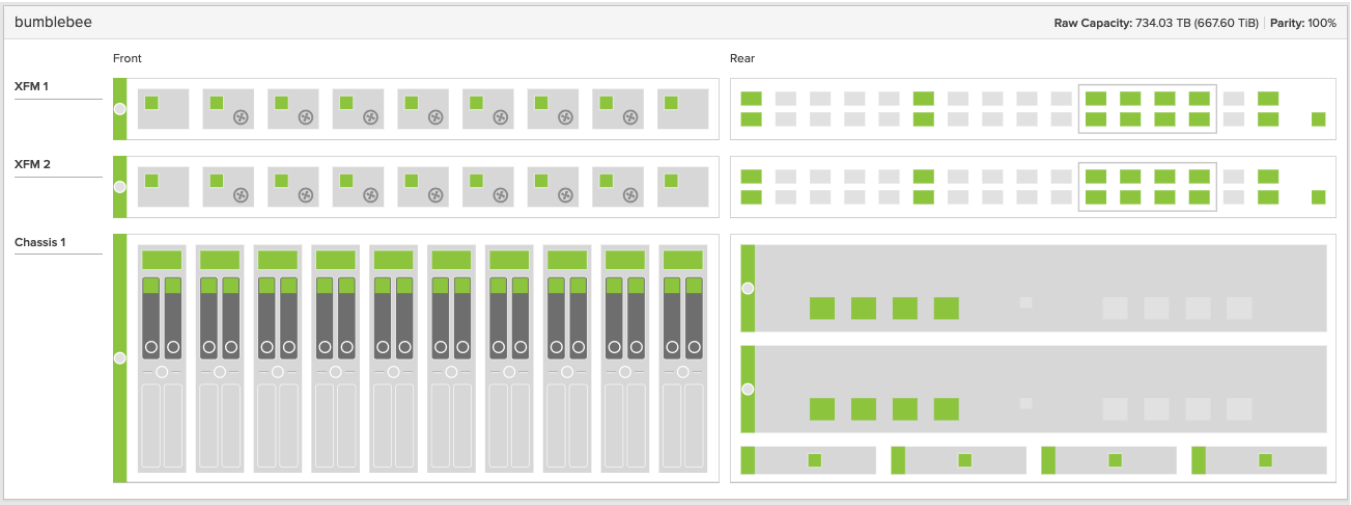

Note: The health status of the chassis reflects the overall physical health of its components. If a chassis component (blade, FM, PSU) is down, the chassis health will be displayed as unhealthy. A component's software health state is not reflected in the chassis health. If a component's software state is unhealthy, the chassis health will be displayed as healthy.

The view is a schematic representation of the array with colored indicators of each component's status:

- **Green**: Healthy and functioning properly.
- **Yellow**: Unhealthy, identifying, or unrecognized.
- **Red**: Critical or unknown.
- **Gray**: Disconnected, unused, or unclaimed.

Hover the mouse over a hardware component to display its status and details. For example, hovering over a power supply displays the power supply's name, its health status, slot number, and its serial and model number.

Below the graphic of the array is the **Data Nodes** panel. The raw capacity of the data nodes is displayed in gibibyte (G), gigabyte (GB), tebibyte (TiB), terabyte (TB), pebibyte (P), and petabyte (PB) units. The Data Nodes panel displays the health of the nodes with a colored bar:

- **Green**: Healthy and functioning properly.
- **Yellow**: Unhealthy, identifying, or unrecognized.
- **Red**: Critical or unknown.

Inside each bar is a number representing the number of nodes currently in that state. Clicking a health staus bar will open the the **Settings** > **Hardware** > **Data Nodes** page where you can then filter the Data Nodes panel to review the health status of a specific node or nodes.

Clicking the Edit button to the right of the raw capacity value opens the **Settings** > **Hardware** > **Data Nodes** page where details about each data node and data node group can be viewed.

Alerts

Purity//FB generates an alert when there is a change to the array or to one of the Purity//FB hardware or software components. The **Health** > **Alerts** page displays a list of alerts that have been generated on the array.

To conserve space, Purity//FB stores a reasonable number of alert records on the array. Older entries are deleted from the log as new entries are added. To access the complete list of messages, contact Pure Storage Support.

Purity//FB assigns a unique numeric ID to each alert as it is created. By default, alerts are sorted in chronological descending order by "Last Seen" date.

The flag icon represents an alert that has been flagged by Purity or the user. Purity automatically flags all alerts. An alert remains flagged until you have manually cleared the flag to indicate that the alert has been addressed. If there are further changes to the condition that caused the alert (for example, the health of a blade has changed), Purity will set the flag again.

The icons that appear along the left side of each alert in the list output represent the alert severity level:

- **Blue (INFO)** icons represent informational messages generated due to a change in state. INFO messages can be used for reporting and analysis purposes. No action is required.
- **Yellow (WARNING)** icons represent important messages warning of an impending error if action is not taken.
- **Red (CRITICAL)** icons represent urgent messages that require immediate attention.

Click any of the column headings to change the sort order.

Each alert in the list output includes the following information:

- **ID**: Unique number assigned by the array to the alert. ID numbers are assigned to alerts in chronological ascending order.
- **Severity**: Alert severity, categorized as `critical`, `warning`, or `info`.

  Critical alerts are typically triggered by service interruptions, major performance issues, or risk of data loss, and require immediate attention. For example, Purity//FB triggers a critical alert if a blade has been removed from the chassis.

  Warning alerts are of low to medium severity and require attention, though not as urgently as critical alerts. For example, Purity//FB triggers a warning alert if it detects an unhealthy blade that is still processing data.

  Informational (Info) alerts inform users of a general behavior change and require no action. For example, Purity//FB triggers an informational alert if the NFS service is unhealthy.
- **Code**: Pure Storage alert code number that identifies the type of alert event.
- **State**: Current state of the alert. Possible states include: `open`, `closed`, `closing`, and `waiting to downgrade`.

  An alert goes from `open` state to `close` state when the issue is completely resolved, such as when a blade has been removed from the chassis. If the blade is then reinserted and pulled again, a new alert will be generated.

  Certain alerts are generated due to intermittent issues. Examples include temperature sensors reporting values outside of normal operating range and space usage exceeding certain

capacity thresholds. These types of alerts can change from `open` state to `closing` or `waiting to downgrade` states.

And alert changes from `open` state to `closing` state when it is no longer an issue. After the alert remains in `closing` state for 24 hours without change, it changes into `closed` state.

An alert changes from `open` state to `waiting to downgrade` state when the issue becomes less severe. After the alert remains in `waiting to downgrade` state for 24 hours without change, the alert severity is downgraded. For example, when array capacity reaches 100%, a critical alert is issued with an `open` state. When array capacity drops below 100%, the alert changes to `waiting to downgrade` state. After the array has remained at less than 100%, but more than 90% capacity, for 24 hours, the alert severity is downgraded to `warning`.

- **Created**: Date and time the alert was first generated and initial alert email notifications were sent to alert watchers.
- **Last Seen**: Most recent date and time Purity//FB saw the issue that generated the alert.
- **Duration**: The time interval between the creation and closing of an alert. If not closed, the current time.
- **Summary**: Alert summary.

Click anywhere in an alert row to display the alert details. Some alert detail descriptions include a **more info** link providing additional information about the alert through the Pure Storage Knowledge Base.

The identify light for blades and fabric modules can be illuminated from the GUI by clicking that individual component from the image of the array from the **Health** > **Hardware** page. The identify light in the image will turn blue and the physical identify light will illuminate to assist with locating the physical component in the data center. Click the component again to turn off the indicator light.

Flagging Alert Messages

To flag alert messages:

1. Select **Health** > **Alerts**.
2. Locate the alert you want to flag.
3. Click the flag icon. When the icon becomes blue, the alert is flagged.

Unflagging Alert Messages

To unflag alert messages:

1. Select **Health** > **Alerts**.

2 Locate the alert you want to unflag.

3 Click the flag icon. When the icon becomes gray, the alert is unflagged.

Alerts can also be unflagged and dismissed from the **Recent Alerts** panel from the **Dashboard** page by clicking the X to the right of the alert.

Chapter 9
The Purity//FB GUI Settings Page

The **Settings** page is used to display and configure system, network, and security settings of an array.

Figure 9.1 Purity//FB Settings Page

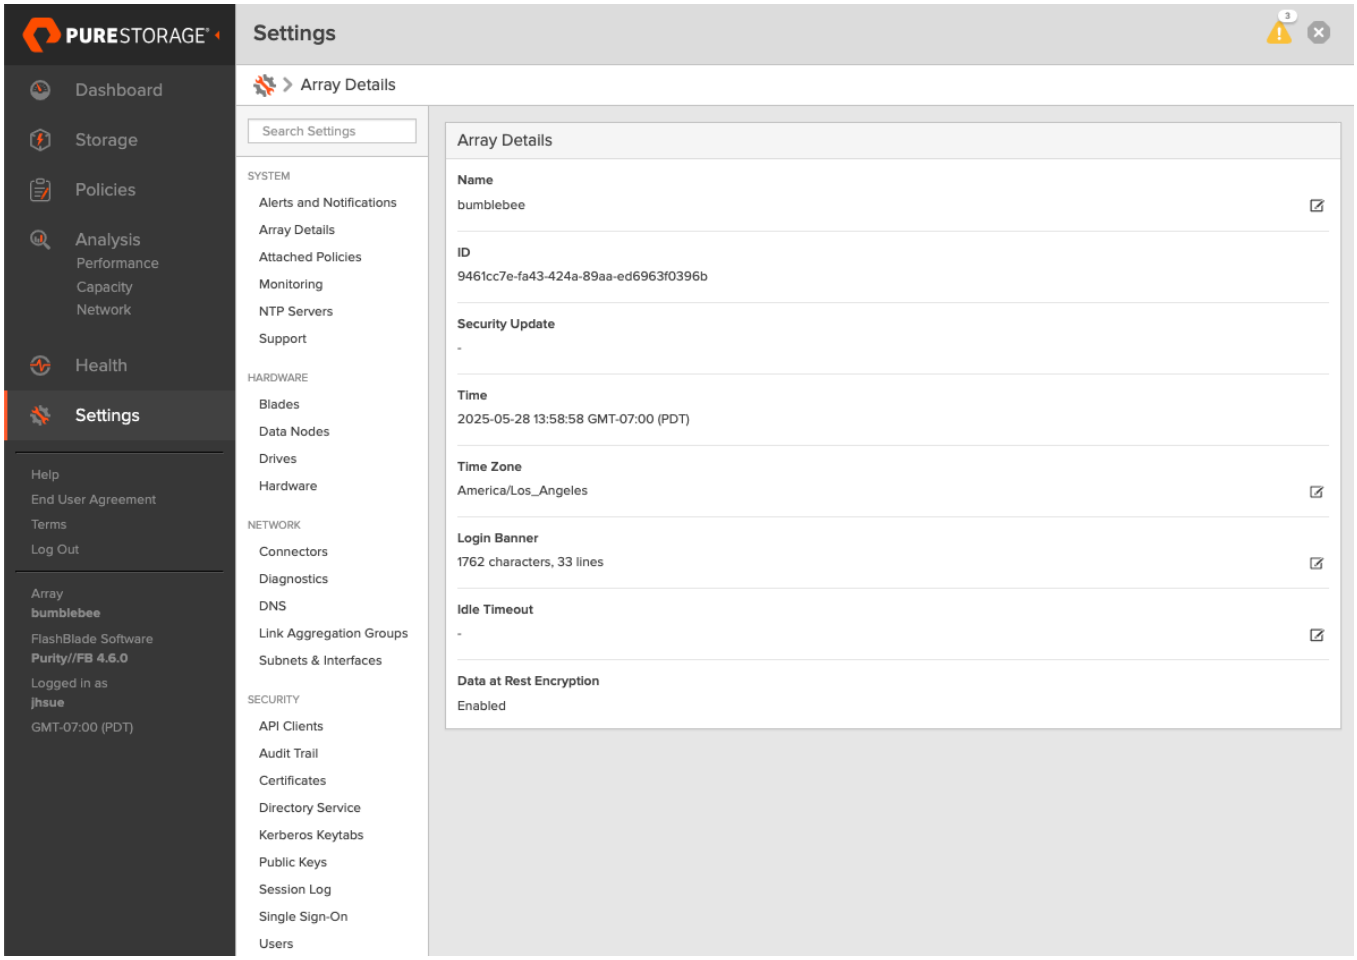

By default, the **Settings** page displays the **Array Details** panel (see *Array Details*). This panel displays basic information about your FlashBlade array.

Figure 9.2 **Array Details Panel**

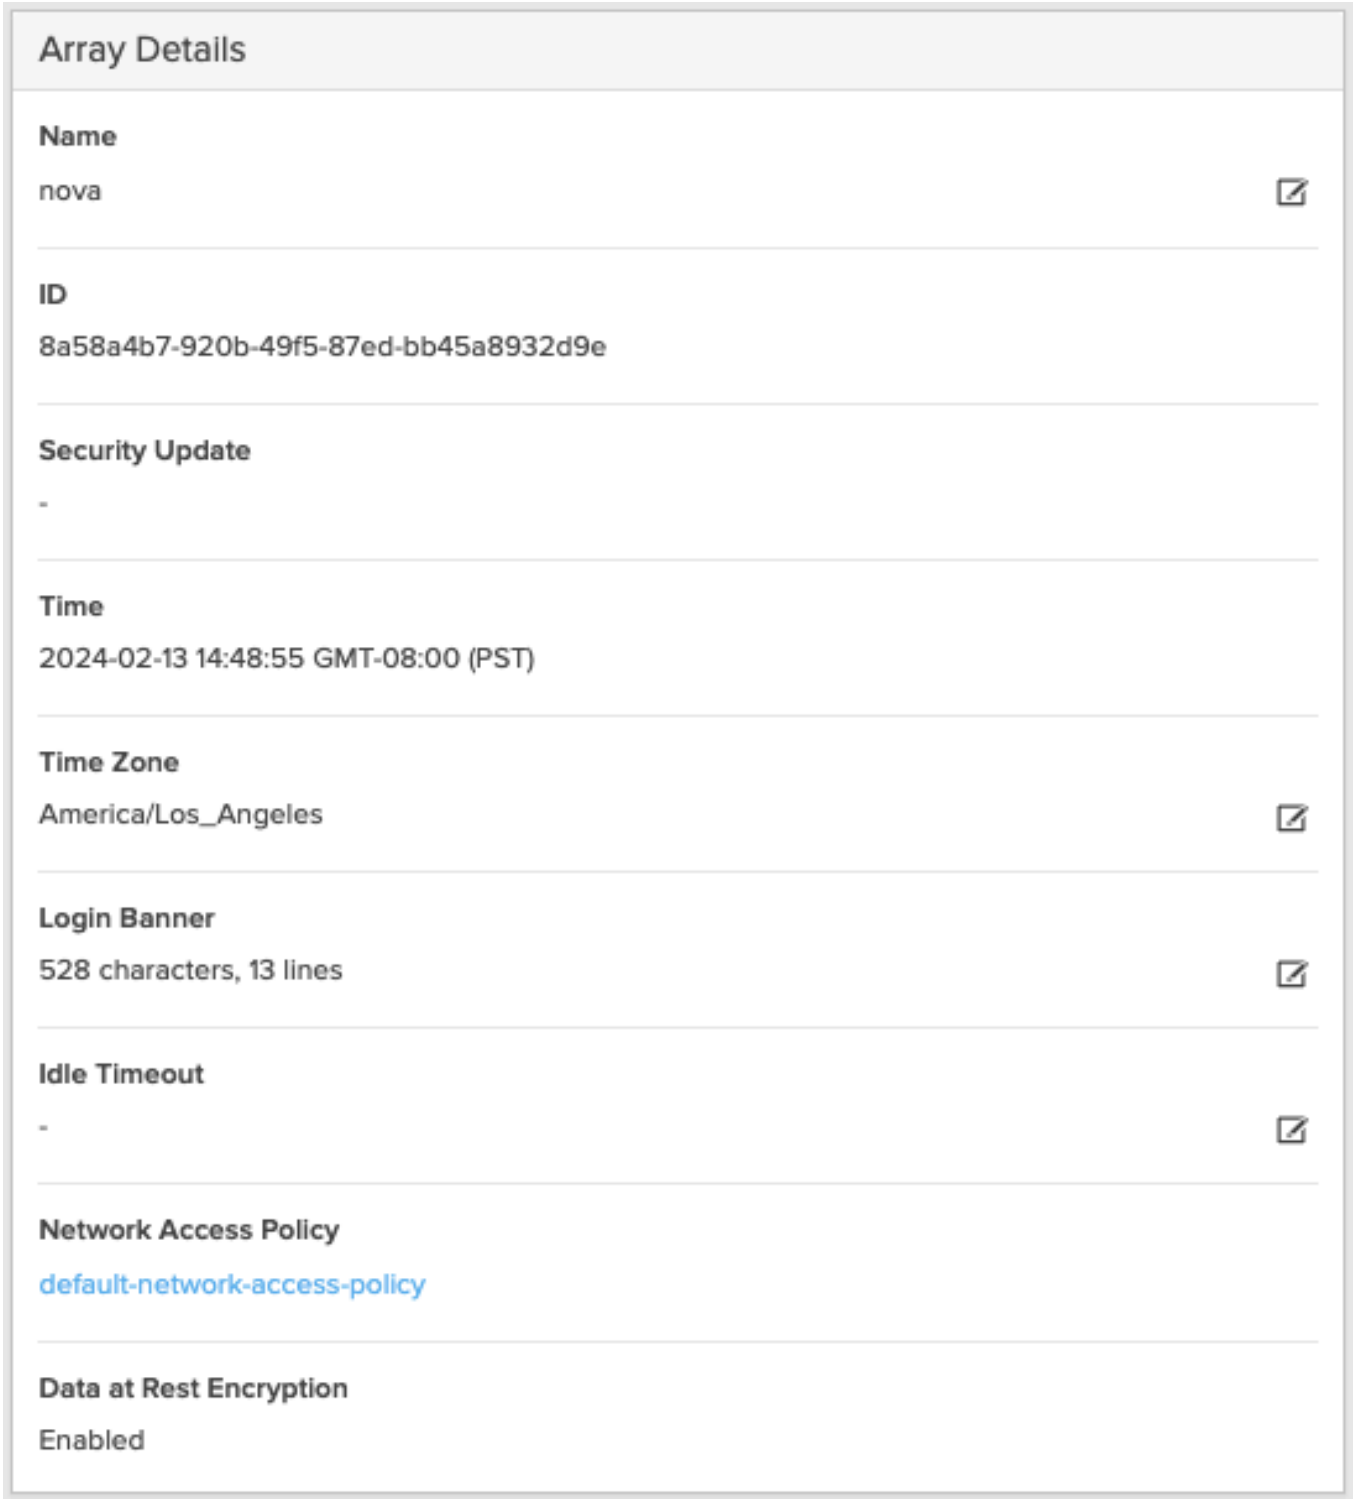

Side Navigation Bar

The **Settings** page provides a navigation sidebar, which allows you to view and manage various array settings.

The navigation side bar is divided into four categories: **System**, **Hardware**, **Network**, and **Security** settings.

The navigation sidebar also provides a search field, allowing you to search for specific system, network, and security items. For example, if you enter "CA Certificate", the navigation sidebar filters the list to just two items: **Certificates** and **Directory Service**.

Working with System Settings

The following **System** settings are available:

- **Alerts and Notifications**
  - **Alert Watchers** - Displays and manages recipients of alert notifications.
  - **Alert Routing** - Displays and allows configuration of how alerts and logs are managed.
  - **File System User/Group Quota Notifications** - Displays and manages recipients of quota notifications.
- **Array Details** - Displays and manages basic information about your FlashBlade array.
- **Attached Policies** - Displays array-level policies.
- **Monitoring**
  - **SNMP** - Displays and manages Simple Network Management Protocol (SNMP) manager information.
  - **Syslog Server Connections** - Displays and manages remote syslog servers.
- **NTP Servers** - Displays and manages the Network Time Protocol (NTP) servers that are currently being used by the array to maintain reference time.

- **Support** - Displays and manages the features used to communicate with Pure Storage Support, as well as view and update the FlashBlade's signed verification key.

Alert Watchers

Purity//FB generates an alert whenever the health of a component degrades or a capacity threshold is reached. Alerts can also be sent as email notifications to designated alert watchers.

The **Alert Watchers** panel displays the email addresses of designated alert watchers, the severity level of alert notifications that alert watchers will receive, and the status of each watcher. The sending account name for Purity//FB alert email notifications is the array name at the configured sender domain.

Once added, an alert watcher will start receiving alert email notifications.

By default, alert watchers receive all alert email notifications regardless of alert severity. You can optionally configure alert watchers to receive only critical and warning-level alert notifications, or only critical-level alert notifications.

To access the **Alert Watchers** panel, from the **Settings** page navigation sidebar click **Alerts and Notifications** under **System**.

Alert watchers can be in enabled (**True**) or disabled (**False**) state. Alert watchers that are enabled will receive alert email notifications. New alert watchers are enabled by default. Disabled alert watchers do not receive email notifications. Disabling an alert watcher does not delete the recipient's email address, it only stops the watcher from receiving email notifications.

Deleting an alert watcher completely removes the watcher from the list. Once an email address has been deleted, the corresponding alert watcher will no longer receive alert notifications.

Adding an Alert Watcher

Add an alert watcher to receive email alert notifications.

To add an alert watcher, perform the following:

1. From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.
2. In the **Alert Watchers** panel, click the Add (+) button. The **Add Alert Watchers** pop-up window appears.
3. Enter the email address of the alert watcher in the **Email** field.
4. (Optional) Select a severity level (**Info+**, **Warning+**, or **Critical**) from the **Notification Severity** menu.

- **Info+**: If selected, the alert watcher will receive all alert notifications for all severity levels.
- **Warning+**: If selected, the alert watcher will receive any alert with a severity of Warning or more severe. Info-level alert notifications will not be received.
- **Critical**: If selected, the alert watcher will only receive the most severe Critical-level alert notifications.

5 (Optional) Click **Test** to send a test email to the new alert watcher.

6 Click **Add**.

Enabling and Disabling an Alert Watcher

In order for an alert watcher to receive email alert notifications, the alert watcher must be enabled. By default, alert watchers are enabled at creation.

1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.

2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher for whom you wish to enable or disable and select **Edit**. The **Edit Alert Watchers** pop-up window appears.

3 Click the **Enabled** toggle switch to set the alert watcher to enabled (blue) or disabled (gray).

4 Click **Save**.

Setting the Severity Level for Alert Watcher Notifications

By default, an alert watcher receives all alert email notifications, regardless of severity. You can configure the severity level of alert that your alert watchers will receive.

To set the severity level for alert watcher notifications, perform the following:

1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.

2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher for whom you wish to set a notification severity level and select **Edit**. The **Edit Alert Watchers** pop-up window appears.

3 Select a severity level (**Info+**, **Warning+**, or **Critical**) from the **Notification Severity** menu.

- **Info+**: If selected, the alert watcher will receive all alert notifications for all severity levels.
- **Warning+**: If selected, the alert watcher will receive any alert with a severity of Warning or more severe. Info-level alert notifications will not be received.
- **Critical**: If selected, the alert watcher will only receive the most severe Critical-level alert notifications.

4 Click **Save**.

Sending a Test Email

You can test an array 's ability to send email notifications to alert watchers, designated or not. Two types of test messages can be issued:

- Before creating an alert watcher, send a test message to the alert watcher's email address to ensure alert notifications can reach the intended destination.
- At any time, send a test message to all of the enabled alert watchers to ensure alert notifications reach their intended destinations.

To send an test email, perform the following:

1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.

2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher for whom you wish to send a test email and select **Send Test Email**.

If the email is not valid, an error message appears.

Deleting an Alert Watcher

Once an email address has been deleted, the corresponding alert watcher will no longer receive alert notifications.

1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.

2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher you wish to delete and select **Delete**. The **Delete Alert Watcher** pop-up window appears.

3 Click **Delete**.

Alert Routing

The **Alert Routing** panel displays the ways in which alerts and logs are managed.

To access the **Alert Routing** panel, from the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.

The sender domain determines how logs are parsed and treated by Pure Storage Support and Escalations. The domain name is also used in the "from" address of outgoing alert email notifications. The domain name must be set to your company's domain name. For example, `mydomain.com`.

The relay host represents the hostname or IP address of the email relay server currently being used as a forwarding point for alert email notifications generated by the array. If a relay host is not configured, Purity//FB sends all alert email notifications directly to the recipient addresses rather than route them via the relay (mail forwarding) server.

FlashBlade currently supports encryption for alert email messages with StartTLS. Note that the external SMTP relay server must support TLS 1.2 or above for encrypted communication. Beginning with Purity//FB 4.5.0, encryption is Opportunistic TLS and enforced when StartTLS is configured for the relay host.

Configuring Alert Routing

Note: If a relay host is not configured, Purity//FB sends all alert email notifications directly to the recipient addresses rather than route them via the relay (mail forwarding) server.

To configure the sender domain and relay host for alert routing, perform the following.

1. From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.
2. Click the edit icon in the **Alert Routing** panel. The **Edit Alert Routing** pop-up window appears.
3. Enter the sender domain name. The domain name must be set to your company's domain name. For example, `mydomain.com`.
4. (Optional) Enter the host name or IP address of the email relay server that is to be used as the forwarding point for alert email notifications generated by the array. If specifying an IP address, enter the IPv4 or IPv6 address.

   For IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from `0` to `255` representing a group of 8 bits. If a port number is also specified, append it to the end of the address in the form `ddd.ddd.ddd.ddd:PORT`, where `PORT` represents the port number.

   For IPv6, specify the IP address in the form `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`) . If a port number is also specified, enclose the entire address in square brackets (`[]`) and append the port number to the end of the address. For example, `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]:PORT`, where `PORT` represents the port number.
5. Select the encryption mode, **None** or **starttls**. If **None** is selected, encryption is Opportunistic TLS. If **starttls** is selected, encryption is enforced. Note that the external SMTP relay server must support TLS 1.2 or above.

6 Click **Save**.

Array Details

The **Array Details** panel is displayed by default when navigating to the Purity//FB GUI's **Settings** page. It can also be accessed from the **Settings** page navigation sidebar by clicking **Array Details** under **System**

Figure 9.3 Viewing the Array Details Panel

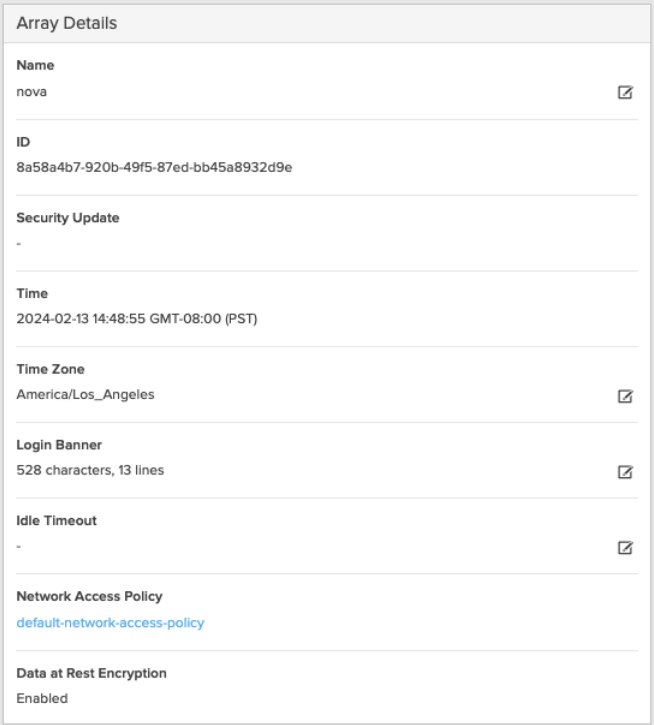

This panel displays basic information about your FlashBlade array, including:

- **Name** - The array name set during FlashBlade installation.
- **ID** - The array's unique ID.
- **Security Update** - The version of the last installed Purity//FB security update (if it still applies).
- **Time** - The array time is based on the time zone of the array, which is set during FlashBlade installation.
- **Time Zone** - The array time zone. The array time zone is not related to the user's local time zone that appears in the bottom-left corner of the Navigation panel. The user's local time zone is determined by the browser's local time.

Note that Purity//FB CLI dates and times are displayed in the time zone of the array, while Purity//FB GUI dates and times are displayed in the local user's time zone.

- **Login Banner** - The number of characters and lines currently used for the login banner.
- **Idle Timeout** - The idle timeout period currently set on the FlashBlade.
- **Network Access Policy** - The current network access policy governing access to the FlashBlade interfaces. Clicking the policy opens that policy's detail page where its network access rules can be viewed and edited. See *Network Access Policies* for additional information.
- **Data at Rest Encryption** - Indicates whether data at rest encryption is enabled or disabled. If enabled, inactive data is encrypted preventing unauthorized access.

Items displayed with an edit icon are configurable.

Renaming the Array

1. Navigate to the **Array Details** panel on the **Settings** page.
2. Click the edit icon next to the current array name. The **Rename Array** pop-up window appears.
3. Type the new array name.
4. Click **Rename**.

Array Time

The array time is based on the time zone of the array, which is set during FlashBlade installation.

The array time zone is not related to the user's local time zone that appears in the bottom-left corner of the Navigation panel. The user's local time zone is determined by the browser's local time.

Purity//FB CLI dates and times are displayed in the time zone of the array, while Purity//FB GUI dates and times are displayed in the local user's time zone.

Changing the Time Zone

To change the time zone of the array:

1. Navigate to the **Array Details** panel on the **Settings** page.
2. Click the edit icon next to **Time Zone**. The **Edit Array Time Zone** pop-up window appears.

3 Select the new time zone.

4 Click **Save**.

Creating and Editing a Login Banner

To help identify the FlashBlade system being logged into, you can set a login banner. If a login banner already exists, the **Array Details** panel displays the the number of characters and lines currently used for the login banner.

To create a login banner, perform the following:

1 Navigate to the **Array Details** panel on the **Settings** page.

2 Click the edit icon next to **Login Banner**. The **Edit Login Banner** pop-up window appears.

3 Enter the login banner text. The limit is 8,000 characters.

4 Click **Save**.

Setting the Idle Timeout Period

To automatically log out the user after a period of inactivity, you can set the array's idle timeout period.

Set the idle timeout period in the **Array Details** panel. If a timeout period has already been set, the timeout value is displayed next to **Idle Timeout**.

To set the array idle timeout period, perform the following:

1 Navigate to the **Array Details** panel on the **Settings** page.

2 Click the edit icon next to **Idle Timeout**. The **Edit Idle Timeout** pop-up window appears.

3 Enter the timeout value. The timeout value can be set from five minutes to three hours and must be specified in seconds, minutes, or hours. For example, 300s, 15m, 2h, etc.

4 Click **Save**.

Attached Policies

The **Attached Policies** panel displays policies that are attached at the array level. Each policy displayed in the panel can be clicked to display the policy's details.

The **Attached Policies** panel also allows you to set the array's default TLS policy (see *Setting the Array's Default TLS Policy*).

The panel also allows for the addition and removal of other array policies.

SNMP

The Simple Network Management Protocol (SNMP) is used by SNMP agents and SNMP managers to send and retrieve information. FlashBlade supports SNMP versions v2c and v3.

FlashBlade's built-in SNMP agent has local knowledge of the array. The agent collects and organizes this array information and translates it via SNMP to or from the SNMP managers. The agent cannot be deleted. The managers are defined by creating SNMP manager objects on the array. The managers communicate with the agent via the standard TCP port 161, and they receive notifications on port 162.

The SNMP agent has two functions: responding to GET-type SNMP requests and transmitting alert messages.

The agent responds to GET-type SNMP requests made by the SNMP managers, and returns values for an information block, such as purePerformance, or individual variables within the block, depending on the type of request issued. The following variables are supported:

Table 9.1 System Variables

Variable	OID	Description
pureSystem	iso.3.6.1.4.1.40482.1	System information
pureArrayName	iso.3.6.1.4.1.40482.1.1.0	Array name
pureArrayId	iso.3.6.1.4.1.40482.1.2.0	Array ID
pureArrayVersion	iso.3.6.1.4.1.40482.1.3.0	Array Purity//FB software version
pureArrayType	iso.3.6.1.4.1.40482.1.4.0	Array platform model version
pureArrayCurrentTime	iso.3.6.1.4.1.40482.1.5.0	Array time based on local time zone

Table 9.2 Array Performance Variables

Variable	OID	Description
purePerformance	iso.3.6.1.4.1.40482.4	Array performance information

Table 9.2 Array Performance Variables (continued)

Variable	OID	Description
pureArrayReadBandwidth	iso.3.6.1.4.1.40482.4.1.0	Array read bandwidth (bytes/s)
pureArrayWriteBandwidth	iso.3.6.1.4.1.40482.4.2.0	Array write bandwidth (bytes/s)
pureArrayReadIOPS	iso.3.6.1.4.1.40482.4.3.0	Array read IOPS (op/s)
pureArrayWriteIOPS	iso.3.6.1.4.1.40482.4.4.0	Array write IOPS (op/s)
pureArrayReadLatency	iso.3.6.1.4.1.40482.4.5.0	Array read latency (us/op)
pureArrayWriteLatency	iso.3.6.1.4.1.40482.4.6.0	Array write latency (us/op)
pureArrayAverageOperationSize	iso.3.6.1.4.1.40482.4.7.0	Array average operation size (bytes)
pureArrayAverageReadSize	iso.3.6.1.4.1.40482.4.8.0	Array average read size per read operation (bytes)
pureArrayAverageWriteSize	iso.3.6.1.4.1.40482.4.9.0	Array average write size per write operation (bytes)
pureArrayOtherIOPS	iso.3.6.1.4.1.40482.4.10.0	Array other operations process per second (op/s)
pureArrayOtherLatency	iso.3.6.1.4.1.40482.4.11.0	Array other latency (us/op)

Table 9.3 Hardware Variables

Variable	OID	Description
pureHardware	iso.3.6.1.4.1.40482.5	Hardware information
pureHardwareTable	iso.3.6.1.4.1.40482.5.1	Hardware table index (name field)
pureHardwareName	iso.3.6.1.4.1.40482.5.1.1.1	Hardware component name

Table 9.3 Hardware Variables (continued)

Variable	OID	Description
pureHardwareStatus	iso.3.6.1.4.1.40482.5.1.1.2	Hardware component status
pureHardwareSlot	iso.3.6.1.4.1.40482.5.1.1.3	Hardware component location
pureHardwareModel	iso.3.6.1.4.1.40482.5.1.1.4	Hardware component model
pureHardwareSerialNumber	iso.3.6.1.4.1.40482.5.1.1.5	Hardware component serial number
pureHardwareType	iso.3.6.1.4.1.40482.5.1.1.6	Hardware component type

Table 9.4 Network DNS Variables

Variable	OID	Description
pureNetwork	iso.3.6.1.4.1.40482.6	Network information
pureDNSTable	iso.3.6.1.4.1.40482.6.1	DNS table index (name field)
pureDNSName	iso.3.6.1.4.1.40482.6.1.1.1	DNS name
pureDNSDomain	iso.3.6.1.4.1.40482.6.1.1.2	DNS domain
pureDNSServers	iso.3.6.1.4.1.40482.6.1.1.3	DNS name servers

Table 9.5 LAG Variables

Variable	OID	Description
pureLAGTable	iso.3.6.1.4.1.40482.6.2	Link aggregation group table index (name field)
pureLAGName	iso.3.6.1.4.1.40482.6.2.1.1	Link aggregation group name
pureLAGStatus	iso.3.6.1.4.1.40482.6.2.1.2	Link aggregation group status

Table 9.5 LAG Variables (continued)

Variable	OID	Description
pureLAGPorts	iso.3.6.1.4.1.40482.6.2.1.3	Link aggregation group port
pureLAGPortSpeed	iso.3.6.1.4.1.40482.6.2.1.4	Link aggregation group port speed
pureLAGSpeed	iso.3.6.1.4.1.40482.6.2.1.5	Link aggregation group link speed

Table 9.6 Port LAG Variables

Variable	OID	Description
purePortToLAGTable	iso.3.6.1.4.1.40482.6.3	Port link aggregation group table index (name field)
purePortName	iso.3.6.1.4.1.40482.6.3.1.1	Port name
purePortLAGName	iso.3.6.1.4.1.40482.6.3.1.2	Port link aggregation group name

Table 9.7 Subnet Variables

Variable	OID	Description
pureSubnetTable	iso.3.6.1.4.1.40482.6.4	Subnet table index (name field)
pureSubnetName	iso.3.6.1.4.1.40482.6.4.1.1	Subnet name
pureSubnetEnabled	iso.3.6.1.4.1.40482.6.4.1.2	Subnet status
pureSubnetLAG	iso.3.6.1.4.1.40482.6.4.1.3	Subnet link aggregation group
pureSubnetServices	iso.3.6.1.4.1.40482.6.4.1.4	Service on subnet
pureSubnetInterfaces	iso.3.6.1.4.1.40482.6.4.1.5	Subnet interfaces

Table 9.8 Network Interface Variables

Variable	OID	Description
pureNetworkInterface	iso.3.6.1.4.1.40482.6.5	Network inferface information
pureNetworkInterfaceName	iso.3.6.1.4.1.40482.6.5.1.1	Network interface name
pureNetworkInterfaceEnabled	iso.3.6.1.4.1.40482.6.5.1.2	Network interface status
pureNetworkInterfaceSubnetName	iso.3.6.1.4.1.40482.6.5.1.3	Network interface subnet name
pureNetworkInterfaceServices	iso.3.6.1.4.1.40482.6.5.1.4	Network interface services

Table 9.9 Hardware Connector Variables

Variable		Description
pureHardwareConnectorTable	iso.3.6.1.4.1.40482.6.6	Hardware connector table index (name field)
pureConnectorName	iso.3.6.1.4.1.40482.6.6.1.1	Hardware connector name
pureConnectorType	iso.3.6.1.4.1.40482.6.6.1.2	Hardware connector type
pureConnectorLaneSpeed	iso.3.6.1.4.1.40482.6.6.1.3	Hardware connector lane speed
pureConnectorTransceiverType	iso.3.6.1.4.1.40482.6.6.1.4	Hardware connector transceiver type

Table 9.10 Network Performance Variables

Variable	OID	Description
pureNetworkPerformance	iso.3.6.1.4.1.40482.7	Network performance information
pureCRCErrors	iso.3.6.1.4.1.40482.7.1.0	CRC errors
pureFrameErrors	iso.3.6.1.4.1.40482.7.2.0	Frame errors
pureCarrierErrors	iso.3.6.1.4.1.40482.7.3.0	Carrier errors

Table 9.10 Network Performance Variables (continued)

Variable	OID	Description
pureDroppedErrors	iso.3.6.1.4.1.40482.7.4.0	Dropped errors
pureOtherErrors	iso.3.6.1.4.1.40482.7.5.0	Other errors
pureTotalErrors	iso.3.6.1.4.1.40482.7.6.0	Total number errors
pureReceivedBytes	iso.3.6.1.4.1.40482.7.7.0	Received bytes per second
pureTransmittedBytes	iso.3.6.1.4.1.40482.7.8.0	Transmitted bytes per second
pureReceivedPackets	iso.3.6.1.4.1.40482.7.9.0	Received packets per second
pureTransmittedPackets	iso.3.6.1.4.1.40482.7.10.0	Transmitted packets per second

The FlashBlade Management Information Base (MIB) describes these variables and can be downloaded from the array to your local machine using the GUI or displayed using the CLI or REST API.

The PURESTORAGE-MIB defines a number of tables. The Object Identifier (OID) structure of a table is demonstrated in the following table on a Name value in the pureHardwareTable for a component called CH1:

Path to the pureHardwareTable	.1.3.6.1.4.1.40482.5.**1**
Path to entry within the table:	.1.3.6.1.4.1.40482.5.1.**1**
Path to the pureHardwareName column for the entry	.1.3.6.1.4.1.40482.5.1.1.**1**
Path to the pureHardwareName value for CH1 component specified by an ASCII index	.1.3.6.1.4.1.40482.5.1.1.1.**67.72.49**
Path to the pureHardwareName value for CH1 component specified by a string index	.1.3.6.1.4.1.40482.5.1.1.1.**\'CH1\'**

Uploading the MIB on the manager side will translate numeric OIDs to variable names when using GET-type SNMP requests. A query to retrieve the pureHardwareName value for CH1 component could

then use the path, `pureHardwareName.\'CH1\'`, as opposed to the aforementioned numbered OID path, `.1.3.6.1.4.1.40482.5.1.1.1.\'CH1\'`.

Having the MIB uploaded on the manager side also enables the **`snmptable`** command that formats tables output for better readability.

Note: The **`snmptable`** command only works if bulk requests are disabled using the `-CB` flag.

SNMP managers are added to the array through the creation of SNMP manager objects. When creating an SNMP manager object, enter the Host, which represents the DNS hostname or IP address of the computer that hosts the SNMP manager. Also specify the SNMP version from the Version drop-down list. Valid versions are v2c and v3.

The SNMP agent generates and transmits messages to the SNMP manager as traps or inform requests (informs), depending on the notification type that is configured on the manager. An SNMP trap is an unacknowledged SNMP message, meaning the SNMP manager does not acknowledge receipt of the message. An SNMP inform is an acknowledged trap.

If the SNMP manager notification type is set to `trap`, the agent sends the SNMP message (trap) without expecting a response. If the SNMP manager is set to `inform`, the agent sends the SNMP message (inform) and waits for a reply from the manager confirming message retrieval. If the agent does not receive a response within a certain timeframe, it will retry until the inform has passed through successfully. If the notification type is not set, the manager defaults to `trap`.

SNMPv2 uses a type of password called a community string to authenticate the messages that are passed between the agent and manager. The community string is sent in clear text, which is considered an unsecured form of communication. SNMPv3 supports secure communication between the agent and manager through the use of authentication and privacy encryption methods. As such, SNMPv2c and SNMPv3 have different security attributes.

To configure the SNMPv2c agent and managers, set the **Community** field to the community string under which the agent is to communicate with the managers. The agent and manager must belong to the same community; otherwise, the agent will not accept requests from the manager. When setting the community, Purity prompts twice for the community string. To remove the agent or manager from the community, leave the field blank.

To configure the SNMPv3 agent and managers, in the **User** field, specify the user ID that Purity uses to communicate with the SNMP manager. Also set the authentication and privacy encryption security levels for the agent and managers. SNMPv3 supports the following security levels:

- **noAuthNoPriv**: Authentication and privacy encryption is not set. Similar to SNMPv2c, communication between the SNMP agent and managers is not authenticated and not encrypted. noAuthNoPriv security requires no configuration.

- **authNoPriv**: Authentication is set, but privacy encryption is not set. Communication between the SNMP agent and managers is authenticated but not encrypted. Password authentication is based on MD5 or SHA hash authentication.

  To configure authNoPriv security, in the **Auth Protocol** field, set the authentication protocol to MD5 or SHA, and in the **Auth Passphrase** field, enter an authentication passphrase.

- **authPriv**. Communication between the SNMP agent and managers is authenticated and encrypted. Password authentication is based on MD5 or SHA hash authentication. Traffic between the FlashBlade and SNMP manager is encrypted using encryption protocol AES or DES.

  To configure authPriv security, in the **Auth Protocol** field, set the authentication protocol to MD5 or SHA, and in the **Auth Passphrase** field, enter an authentication passphrase. Also, in the **Privacy Protocol** field, set the privacy protocol to AES5 or DES, and in the **Privacy Passphrase** field, enter a privacy passphrase.

  Note that privacy cannot be configured without authentication

Once an SNMP manager object is created on the array, the FlashBlade immediately starts transmitting SNMP messages and alerts to the manager.

The **SNMP** panel displays configured SNMP managers. The **SNMP** panel also allows you to configure and manage SNMP managers and agents. From the panel you can also download the SNMP MIB.

To access the **SNMP** panel, from the **Settings** page navigation sidebar, click **Monitoring** under **System**.

Figure 9.4 Viewing the SNMP Panel

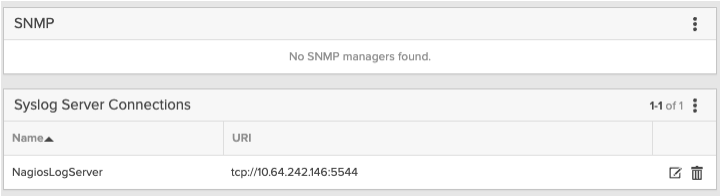

The **SNMP** panel displays a list of configured SNMP managers with the following information:

- **Name**: The name of the SNMP manager.

- **Host**: The host address configured with the SNMP manager.

- **Version**: The SNMP version.

Downloading the Management Information Base (MIB) File

To download the MIB file, perform the following:

1. From the **Settings** page navigation sidebar, click **Monitoring** under **System**.
2. From the header of the **SNMP** panel, click **More Options** and select **Download MIB**.

   The `PURESTORAGE-MIB.txt` immediately downloads to the default download location on your local machine.

Creating an SNMP Manager Object

SNMP managers are added to the array through the creation of SNMP manager objects.

To add an SNMP manager object perform the following:

1. From the **Settings** page navigation sidebar, click **Monitoring** under **System**.
2. From the header of the **SNMP** panel, click **More Options** and select **Add SNMP Manager**. The **Add SNMP Manager** pop-up window appears.
3. Complete the following fields:

   - **Name**: The SNMP manager's name.
   - **Host**: DNS hostname or IP address of a computer that hosts an SNMP manager to which Purity is to send messages when it generates alerts. If specifying an IP address, enter the IPv4 or IPv6 address.

     For IPv4, specify the IP address in the form ddd.ddd.ddd.ddd, where ddd is a number ranging from 0 to 255 representing a group of 8 bits. If a port number is also specified, append it to the end of the address in the format ddd.ddd.ddd.ddd:PORT, where PORT represents the port number.

     For IPv6, specify the IP address in the form xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx, where xxxx is a hexadecimal number representing a group of 16 bits. Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (::). If a port number is also specified, enclose the entire address in square brackets ([]) and append the port number to the end of the address. For example, [xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]:PORT, where PORT represents the port number.
   - **SNMP Version**: Version of the SNMP protocol to be used by Purity in communications with the specified managers. Valid values are v2c and v3 (default).
   - **Community**: SNMPv2c only. SNMP manager community ID under which Purity is to communicate with the specified managers.

- **User**: SNMPv3 only. User ID recognized by the specified SNMP managers that Purity is to use in communications with them.
- **SNMP Notification**: Notification type that determines whether the recipient (remote host) acknowledges receipt of SNMP messages. Valid options are **trap** and **inform**. An SNMP trap is an unacknowledged (asynchronous) SNMP message, meaning the recipient does not acknowledge receipt of the message. An SNMP inform request is an acknowledged trap. If not specified, the notification type defaults to trap.
- **Auth Protocol**: SNMPv3 only. Hash algorithm used to validate the authentication passphrase. Valid values are **MD5**, **SHA**, or **None**.
- **Auth Passphrase**: SNMPv3 only. Passphrase used by Purity to authenticate the array with the specified managers. Required if the **Auth Protocol** option is set to **MD5** or **SHA**.
- **Privacy Protocol**: SNMPv3 only. Encryption protocol for SNMP messages. Valid values are **AES**, **DES**, or **None**.
- **Privacy Passphrase**: SNMPv3 only if privacy protocol is set to **AES** or **DES**. Passphrase used to encrypt SNMP messages. The passphrase must be between 8 and 63 non-spaced ASCII characters.

4 Click **Save**.

Editing an SNMP Manager Object

To edit an SNMP manager object, perform the following:

1 Select **Settings** > **System**.

2 In **SNMP** panel, locate the SNMP manager you wish to edit and click **More Options** and select **Edit**. The **Edit SNMP Manager** pop-up window appears.

3 Modify the following fields:

- **Host**: DNS hostname or IP address of a computer that hosts an SNMP manager to which Purity is to send messages when it generates alerts. If specifying an IP address, enter the IPv4 or IPv6 address.

  For IPv4, specify the IP address in the form ddd.ddd.ddd.ddd, where ddd is a number ranging from 0 to 255 representing a group of 8 bits. If a port number is also specified, append it to the end of the address in the format ddd.ddd.ddd.ddd:PORT, where PORT represents the port number.

  For IPv6, specify the IP address in the form xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx, where xxxx is a hexadecimal number representing a group of 16 bits. Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (::). If a port number is also specified, enclose the entire address in square brackets ([]) and append the port number to the end of the address. For example, [xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]:PORT, where PORT represents the port number.

- **SNMP Version**: Version of the SNMP protocol to be used by Purity in communications with the specified managers. Valid values are v2c and v3 (default).
- **Community**: SNMPv2c only. SNMP manager community ID under which Purity is to communicate with the specified managers.
- **User**: SNMPv3 only. User ID recognized by the specified SNMP managers that Purity is to use in communications with them.
- **SNMP Notification**: Notification type that determines whether the recipient (remote host) acknowledges receipt of SNMP messages. Valid options are **trap** and **inform**. An SNMP trap is an unacknowledged (asynchronous) SNMP message, meaning the recipient does not acknowledge receipt of the message. An SNMP inform request is an acknowledged trap. If not specified, the notification type defaults to trap.
- **Auth Protocol**: SNMPv3 only. Hash algorithm used to validate the authentication passphrase. Valid values are **MD5**, **SHA**, or **None**.
- **Auth Passphrase**: SNMPv3 only. Passphrase used by Purity to authenticate the array with the specified managers. Required if the **Auth Protocol** option is set to **MD5** or **SHA**.
- **Privacy Protocol**: SNMPv3 only. Encryption protocol for SNMP messages. Valid values are **AES**, **DES**, or **None**.
- **Privacy Passphrase**: SNMPv3 only if privacy protocol is set to **AES** or **DES**. Passphrase used to encrypt SNMP messages. The passphrase must be between 8 and 63 non-spaced ASCII characters.

4 Click **Save**.

Renaming an SNMP Manager Object

To rename an SNMP manager object, perform the following:

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 In **SNMP** panel, locate the SNMP manager you wish to rename and click **More Options** and select **Rename**. The **Rename SNMP Manager** pop-up window appears.

3 Enter a new name in the **Name** field.

4 Click **Rename**.

Deleting an SNMP Manager Object

To delete an SNMP manager object, perform the following:

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 In **SNMP** panel, locate the SNMP manager you wish to delete and click **More Options** and select **Delete**. The **Delete SNMP Manager** pop-up window appears.

3 Click **Delete**.

Sending a Test Trap

To send a test trap to an SNMP manager, perform the following:

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 In **SNMP** panel, locate the SNMP manager to which you wish to send a test trap and click **More Options** and select **Send Test Trap**. The **Send Test Trap** pop-up window appears notifying you that a test trap has been sent to the SNMP manager.

3 Click **Close**.

Editing an SNMP Agent

To edit an SNMP agent, perform the following:

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 From the header of the **SNMP** panel, click **More Options** and select **Edit SNMP Agent**. The **Edit SNMP Agent** pop-up window appears.

3 Modify the following fields:

- **SNMP Version**: Version of the SNMP protocol to be used by Purity in communications with the specified managers. Valid values are v2c (default) and v3.
- **Community**: SNMPv2c only. SNMP manager community ID under which Purity is to communicate with the specified managers.
- **User**: User ID recognized by the specified SNMP managers that Purity is to use in communications with them.
- **Auth Protocol**: SNMPv3 only. Hash algorithm used to validate the authentication passphrase. Valid values are **MD5**, **SHA**, or **None**.
- **Auth Passphrase**: SNMPv3 only. Passphrase used by Purity to authenticate the array with the specified managers. Required if the **Auth Protocol** option is set to **MD5** or **SHA**.
- **Privacy Protocol**: SNMPv3 only. Encryption protocol for SNMP messages. Valid values are **AES**, **DES**, or **None**.
- **Privacy Passphrase**: SNMPv3 only if privacy protocol is set to **AES** or **DES**. Passphrase used to encrypt SNMP messages. The passphrase must be between 8 and 63 non-spaced ASCII characters.

4 Click **Save**.

Syslog Server

The Syslog Server feature enables you to forward syslog messages to remote syslog servers.

Up to three remote syslog servers can be connected. All connected syslog servers can be viewed from the **Syslog Server Connections** panel.

To access the **Syslog Server Connections** panel, from the **Settings** page navigation sidebar, click **Monitoring** under **System**.

Figure 9.5 Viewing the Syslog Server Connections Panel

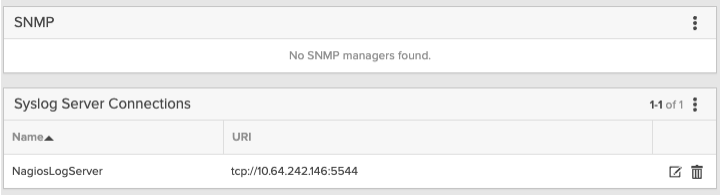

The Purity//FB syslog logging facility generates messages of major events within the FlashBlade and forwards the messages to remote servers. Purity//FB generates syslog messages for three types of events:

- Alerts (`purity.alert`)
- Audit (`purity.audit`)
- Tests (`purity.test`)

Alerts

Purity//FB generates alerts when there is a change to the array or to one of the Purity//FB hardware or software components. There are three alert severity levels:

- **INFO**: Informational messages that are generated due to a change in state. INFO messages can be used for reporting and analysis purposes. No action is required.
- **WARNING**: Important messages that warn of an impending error if action is not taken.
- **CRITICAL**: Urgent messages that require immediate attention.

Syslog alerts are broken down into the following format:

```
<Event Timestamp> <FM Host Name> purity.alert: <Alert Severity> <Alert State>
<Alert Details>
```

In the following example, Purity//FB generated a Warning alert due to an unhealthy blade:

```
May 15 13:38:08 irp125a-v08g08-sim-ch1-fm1 purity.alert: WARNING Open - irp125a-
v08g08-sim: Blade CH1.FB1 is unhealthy [1000] #012A blade is unhealthy. No
data has been lost, but performance may be impacted and additional wear may
be occurring. There will be daily notifications for 7 days.#012#012Severity:
Warning#012State: Open#012First Seen UTC Time: 2020 May 15 13:29:34 UTC#012First
Seen Array Time: 2020 May 15 06:29:34 PDT#012Array Name: irp125a-v08g08-
sim#012Array ID: 00000000-0000-4000-8000-000000000000#012Purity//FB Version:
3.1.99#012#012Suggested Action: This alert should automatically generate a
support case. Please review your support case status by logging into
Pure1 at https://pure1.purestorage.com/. #012Knowledge Base Article: https://
support.purestorage.com/?cid=Alert_1000#012
```

Alerts are also sent via the phone home facility to the Pure Storage Support team. If configured, alerts can also be sent to designated email recipients and SNMP trap managers. Alerts can also be viewed from the GUI **Health > Alerts** page and from the **`purealert list`** CLI command.

Audit

An audit trail represents a chronological history of the GUI, CLI, or REST API operations that a user has performed to modify the configuration of the array. Each message within an audit trail includes the name of the Purity//FB user who performed the operation and the Purity//FB operation that was performed.

Syslog audit trail messages are broken down into the following format:

```
<Event Timestamp> <FM Host Name> purity.audit: <GUID> <ID> <Timestamp> <User>
<Location> <User Interface> <Command Details>
```

In the following example, **pureuser** performed the **`purelog`** CLI command:

```
May 15 13:43:24 irp125a-v08g08-sim-ch1-fm1 purity.audit: - Adding
audit entry. GUID=b1936eff-c4da-4582-89ad-9f5ab2d57777, ID=2, Time=2020-05-15
13:43:24.584225, User=pureuser, Location=10.255.8.1, UI=CLI, Command=purelog,
Subcommand=delete, Arguments=test_syslog_alerts_remote
```

The audit trail can also be viewed from the GUI **Settings > User** page and from the the **`pureaudit list`** CLI command.

Tests

Test messages are generated by users to verify that the array can send messages to remote syslog servers. The message does not indicate whether or not the test message successfully reached the recipients; you must manually check the remote syslog server to confirm that the message arrived at the destination.

Syslog test messages are broken down into the following format:

```
<Event Timestamp> <FM Host Name> purity.test: <Test Message Details>
```

In the following example, a test was perfomed to determine if the array could send messages to a remote syslog server.

```
May 15 13:39:03 irp552-c01-ch1-fm2 purity.test: - This is a test message
generated by Pure Storage FlashBlade. UTC Time: 2020 May 15 13:39:03 from FM
Name: fm2
```

Creating a Syslog Server Connection

1. From the **Settings** page navigation sidebar, click **Monitoring** under **System**.
2. From the header of the **Syslog Server Connections** panel click **More Options** and select **Create Connection**. The **Create Syslog Server Connection** pop-up window appears.
3. In the **Name** field, enter the name for the remote syslog server connection.
4. In the **URI** field, enter the URI of the remote syslog server.

   Specify the URI in the format **PROTOCOL://HOSTNAME:PORT**, where

   - **PROTOCOL** is TCP, TLS, or UDP.

     Note that when establishing TLS communication, the syslog server must be configured with either a CA certificate or CA certificate group (for multiple remote syslog servers). See *Selecting CA Certificates for Syslog Server Connections*. Additionally, your firewall must be set to allow in-bound and out-bound traffic.

   - **HOSTNAME** is the syslog server's hostname or IP address. If specifying an IP address, for IPv4, specify the IP address in the form ddd.ddd.ddd.ddd, where ddd is a number ranging from 0 to 255 representing a group of 8 bits.

     For IPv6, specify the IP address in the form [xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx], where xxxx is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets ([]). Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (::).

- **PORT** is the port at which the server is listening. Append the port number after the end of the entire address. This is optional if the standard port number 514 for UDP or TCP, or the standard port number 6514 for TLS, is used.

5 Select the service, **data-audit** or **management**.

6 (Optional) Click the arrow in the **Source** field and select the network interface for the syslog server. By default the management network vip is used for syslog server traffic. If configuring for data-audit service, it is recommended to select a data network interface which can be used for routing audit traffic to reduce network resource contention with the management network interface.

7 Click **Create**.

The **Syslog Server Connections** panel is updated with the new syslog server

When establishing TLS communication, each remote syslog servers' certificate must have been signed by the configured CA certificate, or by one of the CA certificates in the configured CA certificate group. See *Selecting CA Certificates for Syslog Server Connections*. Additionally, your firewall must be set to allow in-bound and out-bound traffic.

Selecting CA Certificates for Syslog Server Connections

For TLS connections, you must select a CA certificate for the remote syslog server.

If you are setting up TLS connections for mutliple remote syslog servers, select a CA certificate group.

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 From the header of the **Syslog Server Connections** panel click **More Options** and select **Select Certificate**. The **Select CA Certificate** pop-up window appears.

3 Select either a CA certificate or CA certificate group.

As a best practice, it is recommended that you create a separate CA certificate group for syslog certificates.

4 Click **Save**.

Testing Syslog Server Connections

To test connected remote syslog servers, perform the following:

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 From the header of the **Syslog Server Connections** panel click **More Options** and select **Test**. The **Test Syslog Server Connection** pop-up window appears displaying the results of the syslog server connection test.

A test message is sent to the remote syslog server. You must manually check the remote syslog server to confirm that the message at the destination.

If the test message is not received, possible issues can include:

- URI typos
- Issues with the remote syslog server's configuration
- Issues with firewall settings
- Issues IPv6 connectivity (if the FlashBlade is configured with IPv6 support)
- Invalid certificate configurations

3 Click **Close** after you have reviewed the test results.

Editing a Syslog Server Connection

To edit a connected remote syslog server, perform the following:

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 In the **Syslog Server Connections** panel click the **Edit** button to the right of the remote syslog server you wish to edit. The **Edit Syslog Server Connection** pop-up window appears.

3 In the **URI** field, enter the new URI of the remote syslog server.

Specify the URI in the format **PROTOCOL://HOSTNAME:PORT**, where

- **PROTOCOL** is TCP, TLS, or UDP.

  Note that when establishing TLS communication, the syslog server must be configured with either a CA certificate or CA certificate group (for multiple remote syslog servers). See *Selecting CA Certificates for Syslog Server Connections*. Additionally, your firewall must be set to allow in-bound and out-bound traffic.

- **HOSTNAME** is the syslog server's hostname or IP address. If specifying an IP address, for IPv4, specify the IP address in the form ddd.ddd.ddd.ddd, where ddd is a number ranging from 0 to 255 representing a group of 8 bits.

  For IPv6, specify the IP address in the form [xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx], where xxxx is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets ([]). Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (::).

- **PORT** is the port at which the server is listening. Append the port number after the end of the entire address. This is optional if the standard port number 514 for UDP or TCP, or the standard port number 6514 for TLS, is used.

4 Select the service, **data-audit** or **management**.

5 (Optional) Click the arrow in the **Source** field and select the network interface for the syslog server. By default the management network vip is used for syslog server traffic. If configuring for data-audit service, it is recommended to select a data network interface which can be used for routing audit traffic to reduce network resource contention with the management network interface.

6 Click **Save**.

The **Syslog Server Connections** panel is updated with your changes.

Deleting a Syslog Server Connection

To delete a remote syslog server connection, perform the following:

1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.

2 In the **Syslog Server Connections** panel click the **Delete** button to the right of the remote syslog server whose connection you wish to delete. The **Delete Syslog Server Connection** pop-up window appears.

3 Click **Delete**.

The **Syslog Server Connections** panel is updated with your changes.

NTP Servers

The **NTP Servers** panel displays the hostnames or IP addresses of the Network Time Protocol (NTP) servers that are currently being used by the array to maintain reference time. The installation technician sets the proper time zone for an array when it is installed. During operation, arrays maintain time synchronization by interacting with the NTP server.

To access the **NTP Servers** panel, from the **Settings** page navigation sidebar, click **NTP Servers** under **System**.

Figure 9.6 Viewing the NTP Servers Panel

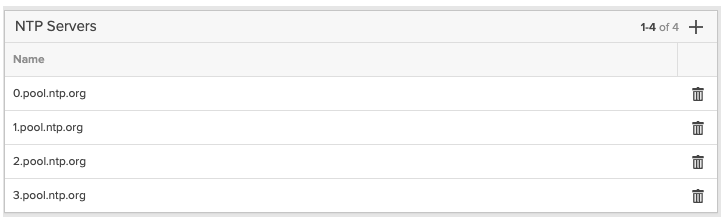

Configuring the NTP Server

The array maintains time synchronization by interacting with the NTP server.

1. From the **Settings** page navigation sidebar, click **NTP Servers** under **System**.
2. In the **NTP Servers** panel:
   - To add an NTP server, in the header of the **NTP Servers** panel, click the **Add** (+) button. The **Add NTP Server** pop-up window appears. Type the hostname or IP address of the NTP server, and then Enter multiple servers as comma-separated values.

     If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from `0` to `255` representing a group of 8 bits. For IPv6, specify the IP address in the form `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where `xxxx` is a hexadecimal number representing a group of 16 bits. When specifying an IPv6 address, consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`) .
   - To remove an NTP server, click the **Delete** button in the same row as the NTP server you wish to remove.

Support

The **Support** panel displays and manages the features used to communicate with Pure Storage Support, the **Challenge Response Verification Key** panel, which allows you to view and update the FlashBlade's signed verification key, and the Maintenance Windows panel, which allows you to schedule a maintenance window for your array.

To access the panels, from the **Settings** page navigation sidebar, click **Support** under **System**.

Figure 9.7 Viewing the Support Panels

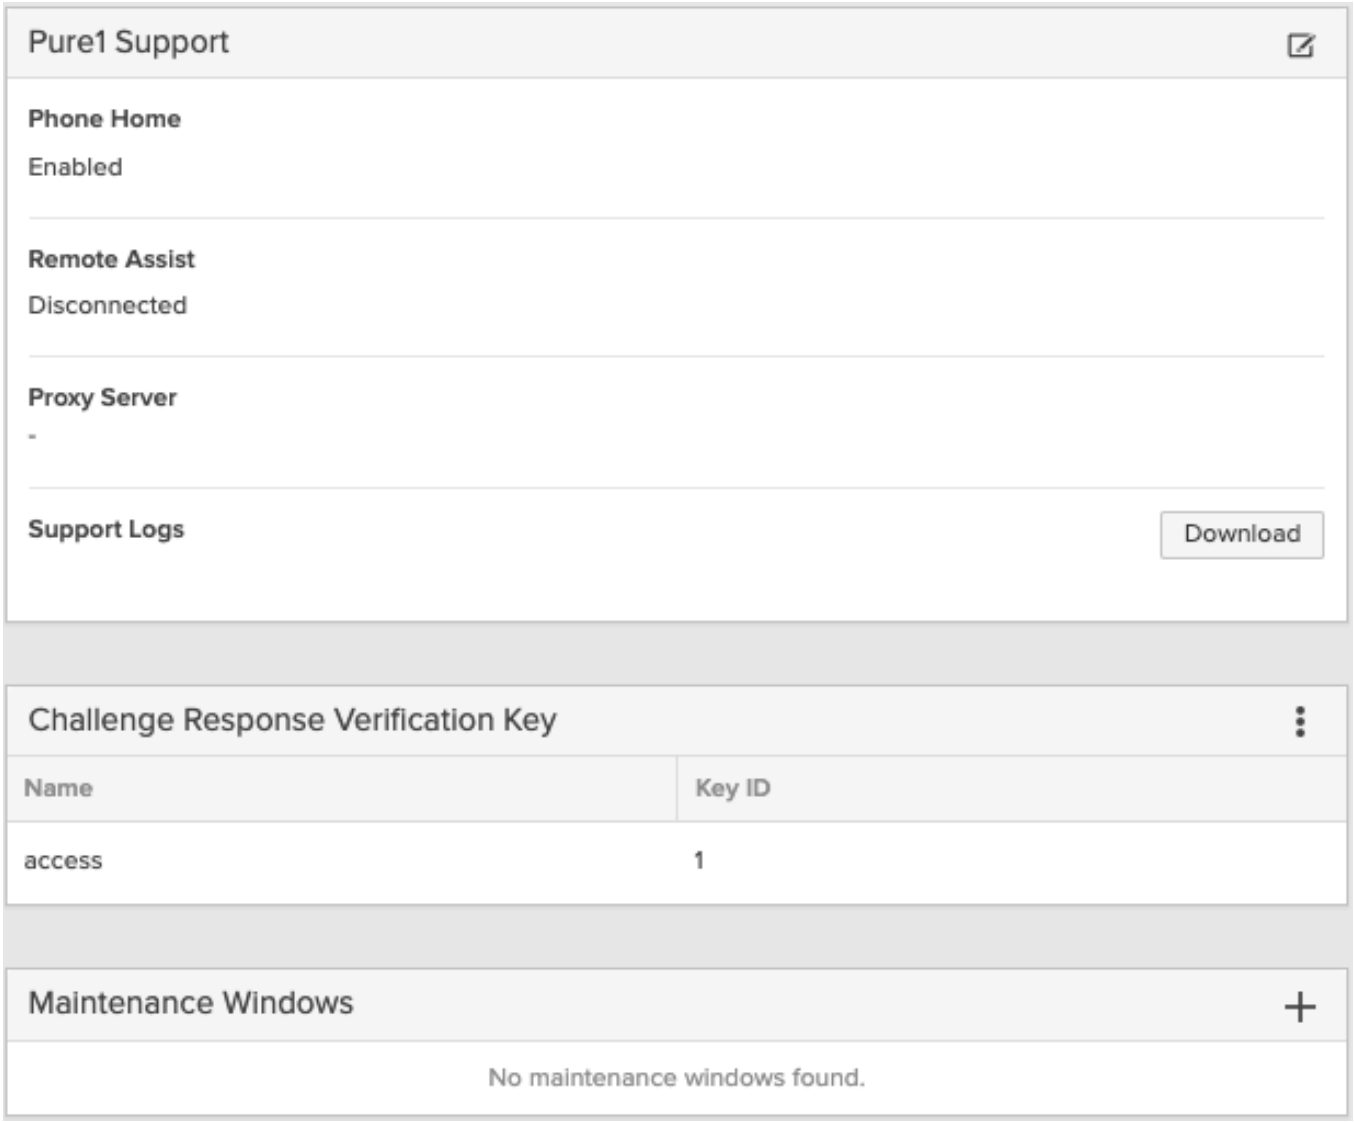

Note that if Remote Assist (RA) is enabled and a custom duration has been set for the RA session, the session expiration date and time will be displayed.

Phone Home

The phone home facility allows the array to transmit log and diagnostic information to Pure Storage Support via a secure network connection.

The phone home facility can be enabled and disabled at any time. The current phone home status appears just below the **Phone Home** label.

Enabling and Disabling Phone Home

Enabling phone home functionality allows the array to transmit log and diagnostic information to Pure Storage Support. It is highly recommended that you enable phone home.

To enable (or disable) phone home, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.
3. Select the **Phone Home Enabled** checkbox to enable.
4. Click **Save**.

Remote Assist

In some cases, the most efficient way for Pure Storage Support to service a FlashBlade array or diagnose problems is through direct access to the array. A remote assistance (RA) session grants Pure Storage Support direct and secure access to the array through a reverse tunnel which you, the administrator, open. This is a two-way communication.

Opening an RA session gives Pure Storage Support the ability to log into the array, effectively establishing an administrative session. Once the RA session is successfully established, the array returns connection details, including the date and time when the session was opened, the date and time when the session expires, and the proxy status (true, if configured).

After the Pure Storage Support team has performed all of the necessary diagnostic or maintenance functions, close the RA session to terminate the connection.

RA sessions can be opened and closed at any time. The current status of the remote assistance session appears just below the Remote Assist label.

An open RA session automatically terminates (closes) after two days have lapsed.

Opening and Closing a Remote Assistance (RA) Session

A remote assistance (RA) session grants Pure Storage Support direct and secure access to the array through a reverse tunnel which you, the administrator, open. This is a two-way communication. By default an RA session lasts 24 hours.

To open or close an RA session, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.
3. Select the **Remote Assist Connected** checkbox to enable an RA session. Opening an RA session gives Pure Storage Support direct and secure access to the array.
4. Click **Save**.

Configuring a Custom Duration for a Remote Assist Session

A remote assistance (RA) session grants Pure Storage Support direct and secure access to the array through a reverse tunnel which you, the administrator, open. This is a two-way communication. By default, an RA session lasts 24 hours. You can optionally set a custom duration for the RA session. The maximum duration is 14 days.

The option to configure a custom duration is only available if RA session is enabled. Additionally, if an RA session is currently running, you cannot set a custom RA duration (the option to set a custom duration will not be visible).

To set a custom duration for an RA session, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.
3. In the **Duration** field, enter the desired RA session duration.
4. Click **Save**.

Proxy Server

The proxy hostname, if set, represents the server to be used as the HTTP or HTTPS proxy. The format for the proxy host name is `http(s)://hostname:port`, where `hostname` is the name of the proxy host, and `port` is the TCP/IP port number used by the proxy host.

Configuring the Proxy Host Name

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.

3 In the **Proxy Server** field type the proxy host name.

The format for the host name is `http(s)://hostname:port`, where `hostname` is the name of the proxy host, and `port` is the TCP/IP port number used by the proxy host.

4 Click **Save**.

Deleting the Proxy Host Name

1 From the **Settings** page navigation sidebar, click **Support** under **System**.

2 In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.

3 In the **Proxy Server** field, delete the proxy host name.

4 Click **Save**.

Support Logs

Purity//FB continuously logs a variety of array activity, including performance metrics, hardware and software operations, and administrative actions. The logs contain the activity of both fabric modules and all blades. The time stamp of each log is based on the array time.

If Phone Home is enabled, the logs are periodically transmitted to Pure Storage. The logs are also saved to the array with up to 30 days' worth of activity available for manual download.

When downloading support logs, specify an event start and end date and time. The array generates a file containing a chronological list of array activity that occurred leading up to, during, and a little after the specified event time. Log files are password encrypted and can only be opened by Pure Storage Support.

Downloading Support Logs

Purity//FB continuously logs a variety of array activity. These logs contain the activity of both fabric modules and all blades. These logs are used by Pure Storage Support when assessing the array's performance and for diagnosing and troubleshooting any events or issues.

To download logs, perform the following:

1 From the **Settings** page navigation sidebar, click **Support** under **System**.

2 In the **Support Logs** section of the **Pure1 Support** panel, click **Download**. The **Download Support Logs** pop-up window appears.

3 Specify the event date and time, representing the approximate array time the activity of interest occurred, and the duration of time for which you wish log information to be collected (1 to 6 hours). Note that the time is based on your local browser's time zone.

4 Optionally select the component(s) for which you wish to prepare support logs.

Figure 9.8 Downloading Support Logs for Single-Chassis Systems

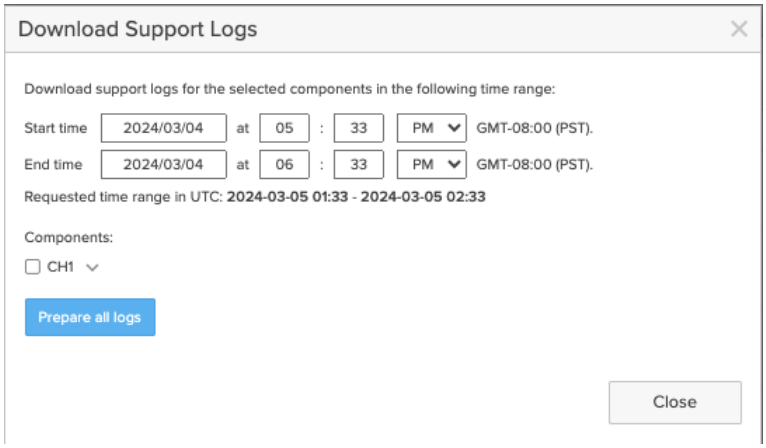

Figure 9.9 Downloading Support Logs for Multi-Chassis Systems

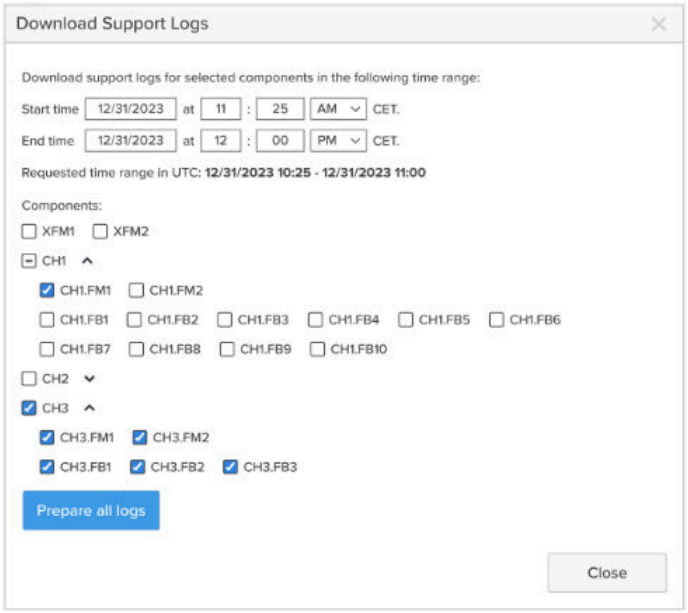

5 If you did not select a component, click **Prepare all logs**. If you selected one or more components, click **Prepare selected logs**.

6 Click **Download** to save the password encrypted file to your administrative workstation. The files can only be opened by Pure Storage Support.

Viewing the Verification Key

To view the FlashBlade's current key used for challenge response verification, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the header of the **Challenge Response Verification Key** panel, click the **More Options** button and then select **View Public Key**. The **Challenge Response Public Key** pop-up window appears displaying the current signed public key (verification key).

Updating the Verification Key

To update the FlashBlade's key used for challenge response verification, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the header of the **Challenge Response Verification Key** panel, click the **More Options** button and then select **Update Signed Verification Key**. The **Update Signed Verfication Key** pop-up window appears.
3. Paste the new signed public key (verification key) in the pop-up window and click **Save**.

Maintenance Windows

The **Maintenance Windows** panel displays the start and end times of any in-progress maintenance window. The **Name** column displays the maintenance type; currently `array` maintenance is supported. The panel also provides the options to modify the duration of an in-progress maintenance window and to delete an in-progress maintenance window.

Figure 9.10 Viewing the Maintenance Windows Panel

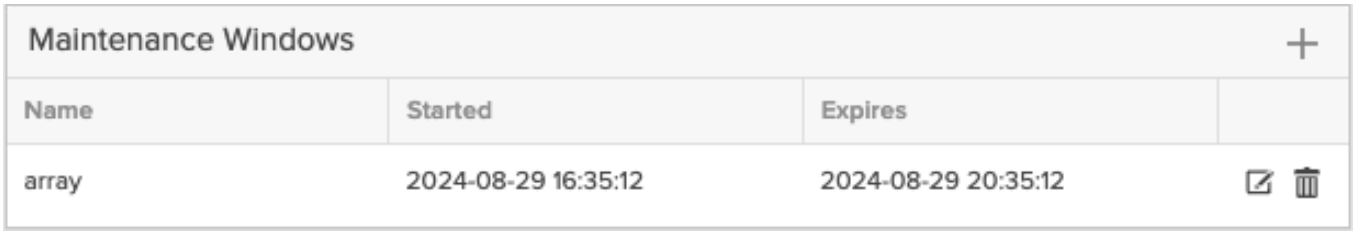

Array users with the roles of `array_admin` and `ops_admin` can create and delete maintenance windows for the array. Array users with the roles of `storage_admin` and `ops_admin` can view/list maintenance windows for the array.

Scheduling a Maintenance Window

A maintenance window is initiated upon completing and submitting the scheduling request. Maintenance windows cannot be scheduled for a future time.

To schedule and initiate a maintenance window for your array, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the header of the **Maintenance Windows** panel click the add button (+). The **Schedule Maintenance Window** pop-up window appears.
3. In the **Duration** field, enter the desired maintenance window, from 1 to 48 hours. By default, the maintenance window is 4 hours.
4. Click **Schedule**.

   The **Maintenance Windows** panel is updated with this window's start and end times, and the maintenance window begins.

Modifying a Maintenance Window

An in-progress maintenance window can be modified at any time before it expires. When modifying an in-progress maintenance window, a new maintenance window duration is applied from the maintenance window's current time, effectively updating the maintenance window. This is described in the following example:

- A 5 hour maintenance window was is at 13:45:45. Its expiration time is 18:45:45.
- At 14:45:45 the maintenance window is modified with a new duration of 1 hour.
- The new expiration time of the maintenance window is 15:45:45.

To modify an in-progress maintenance window, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the **Maintenance Windows** panel click the **Edit** button in the same row of the listed maintenance window. The **Update Maintenance Window** pop-up window appears.
3. In the **Duration** field, enter the desired duration that you wish to update the original maintenance window, from 1 to 48 hours. By default, the maintenance window is 4 hours.
4. Click **Update**.

   The in-progress maintenance window's duration is updated and the **Maintenance Windows** panel is updated to reflect the time the maintenance window was modified and the new expiration time.

Deleting a Maintenance Window

An in-progress maintenance window can be stopped and deleted at any time before it expires.

To stop and delete an in-progress maintenance window, perform the following:

1. From the **Settings** page navigation sidebar, click **Support** under **System**.
2. In the **Maintenance Windows** panel click the **Delete** button in the same row of the listed maintenance window. The **Delete Maintenance Window** pop-up window appears.
3. Click **Delete**.

   The in-progress maintenance window is stopped and deleted. The **Maintenance Windows** panel is cleared.

Working with Hardware Settings

The following **Hardware** settings are available:

- **Blades** - Displays information about all blades in the array.
- **Data Nodes** - Displays and manages data nodes and data node groups in the FlashBlade//EXA array.
- **Drives** - Displays information about all blades in the array.
- **Hardware** - Displays information about all hardware components in the array.

Viewing Blade Information

The **Blades** panel displays information about each blade used in the array, including status and capacity.

To access the **Blades** panel, from the **Settings** page navigation sidebar click **Blades** under **Hardware**.

Viewing and Managing Data Nodes

Information about the data nodes and data node groups used it the FlashBlade//EXA array are displayed in the **Data Node Groups** and **Data Nodes** panels, respectively.

To access the **Data Node Groups** and **Data Nodes** panel, from the **Settings** page navigation sidebar click **Data Nodes** under **Hardware**.

⚠ **Important!** By using your own third-party servers for hosting services, you acknowledge and agree that it is solely your responsibility to ensure the physical and network security of your network infrastructure. This includes, but is not limited to, securing access to your physical hardware, implementing appropriate firewalls, encryption, and other security measures to protect against unauthorized access, data breaches, and cyber threats. Pure Storage recommends you to always follow best practices for both physical and cyber security in order to maintain the integrity of your data and systems. Pure Storage is not liable for any damages, losses, or unauthorized access to your servers, data, or networks that may occur due to improper system maintenance or insufficient security measures.

Data Node Groups

A data node is a high availability server that is used as the FlashBlade//EXA system's data storage. Customers can use their own servers as data nodes but High Availability (HA) will not be supported. One or more data nodes can be added into a node group or resiliency group, and data nodes can be a member of one or more node groups.

A node group or resiliency group is a virtual bundling of one or more data nodes. Groups help organize files, improves availability, and manages compute resources. One or more data nodes can be combined into a group and one data node can belong to one or more groups.

Data node groups are collections of data nodes used to manage file data. The **Data Node Groups** panel displays all created data node groups in the FlashBlade//EXA system.

The **Data Node Groups** panel displays each data node group name and ID. Clicking a data node group name opens the details page for that group. The details page provides **Member Data Nodes** and **Used By** panels, which display the data nodes belonging to the group as well as a list of file systems currently using the group.

The **Data Node Groups** panel also provides functionality to create, rename, and delete data node groups.

Resiliency Groups

A resiliency group is a virtual grouping of data nodes designed to enhance fault tolerance and data protection, otherwise known as high availability. The main purpose of a resiliency group is to ensure that the system can tolerate certain types of hardware failures—such as the loss of one or more nodes—without data loss or service interruption.

Data nodes provided by Pure are in HA mode and will be added to resiliency groups in pairs. A resiliency group consists of two data nodes so for every two Pure data nodes added, a resiliency group will be automatically created. The resiliency group name and nodes in that group can be displayed by using the `purenode resiliency-group list` command.

Data Nodes

Data nodes are servers used to manage file data. The **Data Nodes** panel displays all servers used to manage file data in the FlashBlade//EXA system.

The **Data Nodes** panel displays each data node's name, management address, data addresses, serial number, capacity, and health status.

The **Data Nodes** panel also provides functionality to add, bulk add, edit, rename, and delete data nodes.

Creating a Data Node Group

To create a data node group, perform the following:

1. From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.
2. From the header of the **Data Node Groups** panel click **More Options** and select **Create**. The **Create Data Node Group** pop-up window appears.
3. In the **Name** field, enter the name for data node group.
4. Click **Create**.

Renaming a Data Node Group

To rename a data node group, perform the following:

1. From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.
2. From the header of the **Data Node Groups** panel, locate the data node group you wish to rename and click **More Options** > **Rename** at the end of same row. The **Rename Data Node Group** pop-up window appears.
3. In the **New Name** field, enter the name for data node group.
4. Click **Rename**.

Deleting a Data Node Group

To delete a data node group, perform the following:

1. From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.

2 From the header of the **Data Node Groups** panel, locate the data node group you wish to delete and click **More Options** > **Delete** at the end of same row. The **Delete Data Node Group** pop-up window appears.

3 Click **Delete**.

Adding a Data Node

To add a data node, perform the following:

⚠ **Important!** High Availability (HA) only supports Viking data nodes provided by Pure Storage.

1 From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.

2 From the header of the **Data Nodes** panel, click **More Options** and select **Add**. The **Add Data Node** pop-up window appears.

3 In the **Name** field, enter the name for data node.

4 In the **Management Address** field, enter the management IP address in IPv4 format.

5 In the **Serial Number** field, enter the serial number of the data node.

6 Check the **Add To Groups** option to enable High Availability (HA) on the data node and add to the Resiliency Group.

7 Click **Add**.

Adding Multiple Data Nodes

To add multiple data nodes, perform the following:

1 From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.

2 From the header of the **Data Nodes** panel, click **More Options** and select **Bulk Add**. The **Add Data Nodes** pop-up window appears.

3 In the **JSON** field, paste a valid JSON array containing one or more nodes in the following format:

```
[
  {
    "name": "abc",
    "management_address": "1.2.3.4",
    "serial_number": "xxx-yyy-zzz"
  },
  ...
]
```

4 Click **Add**.

Editing a Data Node

To edit a data node, perform the following:

1. From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.
2. In the **Data Nodes** panel, locate the data node you wish to edit and click **More Options** > **Edit** at the end of the same row. The **Edit Data Node** pop-up window appears.
3. Edit the information in the **Management Address** and **Serial Number** fields as needed.
4. Click **Save**.

Renaming a Data Node

To rename a data node, perform the following:

1. From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.
2. In the **Data Nodes** panel, locate the data node you wish to edit and click **More Options** > **Rename** at the end of the same row. The **Rename Data Node** pop-up window appears.
3. In the **New Name** field, enter the data node's new name.
4. Click **Rename**.

Deleting a Data Node

To delete a data node, perform the following:

1. From the **Settings** page navigation sidebar, click **Data Nodes** under **Hardware**.
2. In the **Data Nodes** panel, locate the data node you wish to edit and click **More Options** > **Delete** at the end of the same row. The **Delete Data Node** pop-up window appears.
3. Click **Delete**.

Viewing Drive Information

The **Drives** panel displays information about each blade used in the array, including status and capacity.

To access the **Drives** panel, from the **Settings** page navigation sidebar click **Drives** under **Hardware**.

Viewing Array Hardware Information

The **Hardware** panel displays information about all hardware components used in the array. By default, general information about each component is displayed, which includes information such as the hardware type, status, and location. Clicking **Specifications** at the top of the **Hardware** panel displays each component's model number, part number, and serial number.

To access the **Hardware** panel, from the **Settings** page navigation sidebar click **Hardware** under **Hardware**.

Working with Network Settings

The following **Network** settings are available:

- **Connectors** - Displays and manages array network interfaces.
- **Diagnostics** - Allows you to conduct network pings and traceroutes.
- **DNS** - Displays and manages DNS settings.
- **Object Store Virtual Hosts** - Displays and manages virtual hosts for object stores.
- **Link Aggregation Groups** - Displays and manages subnet link aggregation groups (LAGs).
- **Subnets & Interfaces** - Displays and manages all subnets and interfaces used for file exports.

Connectors

The **Connectors** panel displays and allows you to edit the array's connection interfaces.

To access the **Connectors** panel, from the **Settings** page navigation sidebar, click **Connectors** under **Network**. The **Connectors** panel displays the following:

- **Name** - The connector name in the format `CHx.FMx.ETHx`, where `CHx` is the chassis number, `FMx` is the fabric module number, and `ETHx` is the ethernet port number.

  Note that the Out-of-Band management ports for FlashBlade are as follows (these ports are not present in in-band clusters):

  - Single-chassis first-generation FlashBlade: CH1.FM1.ETH5, CH1.FM2.ETH5

  - Single-chassis FlashBlade//S: CH1.FM1.ETH9, CH1.FM2.ETH9
  - Multi-chassis FlashBlade: XFM1.ETH33, XFM2.ETH33

- **Connector Type** - The connector type. Possible values include **QSFP**, **RJ-45**, and **SFP**.
- **Port Count** - The number of ports for the connector.
- **Port Speed** - The speed (in Gb/s) at which the component is operating.
- **Transceiver Type** - The transceiver type.
- **External** - Whether the connector is external (**True**) or not (**False**).

Editing a Connector

The number of ports and lane speed for the array's connectors can be edited as they are changed.

To edit a connector, perform the following:

1. From the **Settings** page navigation sidebar, click **Connectors** under **Network**.
2. Click the edit button at the end of the same row of the connector that you wish to edit. The **Edit Connector** pop-up window appears.
3. Select a new port count and/or enter the new lane speed. The port count can be **1** or **4**. The lane speed can be specified as xM or xG. The minimum lane speed that can be specified is 0.
4. Click **Save**.

Diagnostics

Click **Diagnostics** under **Network** to access the **Diagnostics** page. The **Diagnostics** page provides two panels: **Network Diagnostics** and **Connection Statistics**.

Use the **Network Diagnostics** panel to perform network pings and traceroutes to determine the reachability of components and troubleshoot network issues.

Use the **Connection Statistics** panel to monitor TCP connections.

Monitoring TCP

TCP (Transmission Control Protocol) monitoring provides information on the state of the connection. The state of the local and remote ports and IP addresses are given so any issues can be diagnosed and debugged.

The impact of TCP states on network performance depends on factors like network congestion, packet loss, and the efficiency of handling retransmissions and acknowledgments.

The states of TCP connections can significantly impact network performance in various ways. For example, during the SYN-SENT and SYN-RECEIVED states, the initial handshake process may introduce delays, especially if there is network congestion or packet loss, resulting in retransmissions and increased latency. This can degrade performance by slowing down the establishment of connections.

In the ESTABLISHED state, TCP connections can handle data transmission efficiently. However, if the network is congested and experiences frequent packet loss, the retransmissions can affect overall throughput and latency.

In the FIN-WAIT-1 and FIN-WAIT-2 states, as connections are being terminated, delays can occur if the termination requests or acknowledgments are lost or delayed. Time spent in the TIME-WAIT state ensures proper closure and prevents issues with stray packets from old connections, but it can tie up resources and potentially impact performance if many connections are in this state simultaneously.

General network performance issues, such as high latency and low throughput, can be exacerbated by TCP states like CLOSE-WAIT and LAST-ACK, where connection termination is delayed and can affect the speed at which new connections can be established.

Use the **Connection Statistics** panel to debug connection issues. The following fields can be optionally populated to TCP connections of interest. If not populated (default) all TCP connections for all configured vips will be displayed.

- **Local Address**: The array configured data IP address. The IP address can be IPv4 or IPv6 in classful IP or CIDR format.
- **Local Port**: The array data service port number, such as for NFS or replication. Port ranges are 0-65535.
- **Remote Address**: The remote (client) IP address. The IP address can be IPv4 or IPv6 in classful IP or CIDR format.
- **Remote Port**: The remote (client) TCP port number. Port ranges are 0-65535.

- **State**: The TCP connection state. The TCP connection state can be one of the following:
  - CLOSE-WAIT: The remote side has closed the connection, and the local end needs to close it.
  - CLOSED: The connection is closed and no further communication is possible.
  - ESTABLISHED: The connection is established and communication is open.
  - ≠ESTABLISHED: All states, except ESTABLISHED, are selected.
  - FIN-WAIT-1: The socket is closing, and it has sent a connection termination request (FIN).
  - FIN-WAIT-2: The socket is waiting for a connection termination request from the remote side.
  - LAST-ACK: The socket is waiting for a final acknowledgment from the remote side.
  - LISTEN: The socket is listening for incoming connections.
  - SYN-RECEIVED: A connection request (SYN) received, and a response (SYN-ACK) is sent back.
  - SYN-SENT: A connection request (SYN) is sent and waiting for a matching connection request.
  - TIME-WAIT: The socket is waiting to ensure the remote TCP received the acknowledgment of its connection termination request.

Once populated, click **Update** to view the connection status.

Performing Pings and Traceroutes

Network pings and traceroutes are peformed in the **Network Diagnostics** panel.

1. Enter the name of the blade or controller you wish to ping or run a traceroute in the **Components** field.
2. Enter the IP address or fully qualified domain name (FQDN) of the ping or traceroute target in the **Target** field.
3. (Optional) Enter the subnet, interface name, or IP address used to run the test in the **Source** field.
4. For pings, enter the number of times you wish to ping the target in the **Count** field and then click **Ping**.
5. (Optional) For traceroutes, select the **Discover MTU** checkbox and then click **Traceroute**.

The window below the ping and trace parameters displays the results of each ping and traceroute.

Figure 9.11 Viewing Results

```
CH1.FB1:
PING 123.123.123 (123.123.0.123) 56(84) bytes of data.

--- 123.123.123 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1028ms


CH1.FM1:
PING 123.123.123 (123.123.0.123) 56(84) bytes of data.

--- 123.123.123 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1001ms
```

DNS Configurations

The **DNS Configurations** panel displays and manages the DNS attributes for an array's administrative network. The array has one default DNS configuration for management services and the option to add additional configurations for multi-tenancy.

The **DNS Configurations** panel can be accessed from the **Settings** page navigation sidebar, click **DNS** under **Network**.

Figure 9.12 Viewing the DNS Configuration Panel

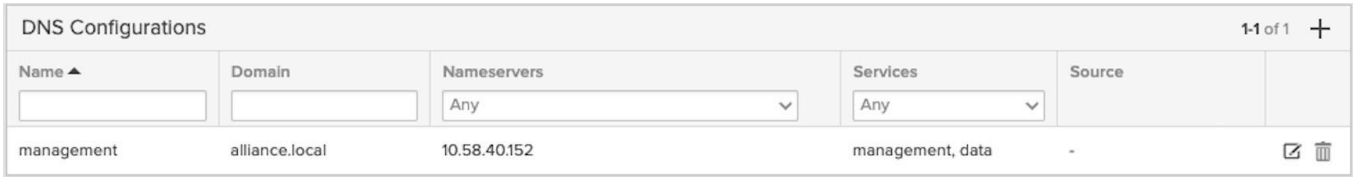

- **Name** - The administrative name for this DNS configuration.
- **Domain** - The DNS domain represents the domain suffix to be appended by the array when performing DNS lookups. For example, `mydomain.com`.
- **Nameservers** - The IP address(es) of the servers to be used when resolved names for the specified domain. Up to three can be entered and they will be queried in the order entered.
- **Services** - The type of service the DNS will handle. The services can be Data and/or Management.
- **Source** - The network interface name for outbound traffic for the DNS server(s) using that particular DNS entry. The Management service is the default DNS configuration and the Source is not selectable.

Creating a DNS Configuration

To create a new DNS, perform the following:

1. From the **Settings** page navigation sidebar, click **DNS** under **Network**.
2. In the **DNS Configuration** panel, click the **Add** interface button (+). The **Create DNS Configuration** pop-up window appears.
3. In the **Name** field, enter the DNS name.
4. In the **Domain** field, enter the domain name.
5. In the **Nameserver 1** field, enter the name server IP address.
6. (Optional) In the **Nameserver 2** field, enter the server IP address.
7. (Optional) In the **Nameserver 3** field, enter the server IP address.
8. In the **Selected Source** list, select the data service network interface.
9. Click **Create**.

Editing DNS Configurations

To edit DNS configurations, perform the following:

1. From the **Settings** page navigation sidebar, click **DNS** under **Network**.
2. Click the edit button in the header of the **DNS Settings** panel. The **Edit DNS Configuration** pop-up window appears.
3. (Optional) In the **Name** field, type a new name for the DNS server.
4. (Optional) In the **Domain** field, type the domain suffix to be appended by the array when doing DNS lookups. For example, `mydomain.com`.
5. (Optional) In the **Nameserver** fields, specify up to three DNS server IP addresses for Purity//FB to use to resolve hostnames to IP addresses. Enter one IP address in each Nameserver field. Purity//FB queries the DNS servers in the order that the IP addresses are listed.
6. (Optional) In the **Selected Source** field, select the network interface name for outbound traffic for the server. Note that the source cannot be changed for the default DNS configuration.
7. Click **Save**.

Link Aggregation Groups

A link aggregation group (LAG) is a collection of Ethernet ports on the array. This concept is also known as a port channel, or a bond group. Multiple LAGs can be used to connect to separate physical networks. By default, the array is preconfigured with all ports in a single LAG named **uplink**, which cannot be deleted. Ports can be removed from the uplink LAG and added to another LAG.

The **Link Aggregation Groups** panel displays and allows management of LAGs on the array.

To access the **Link Aggregation Groups** panel, from the **Settings** page navigation sidebar, click **Link Aggregation Groups** under **Network**.

Figure 9.13 Viewing the Link Aggregation Groups Panel

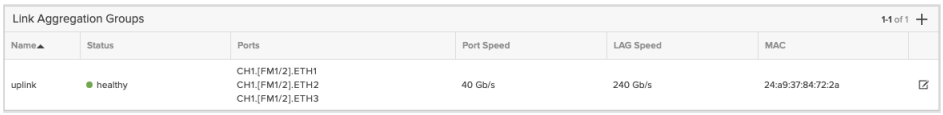

The **Link Aggregation Groups** panel displays the following information:

- **Name**: The assigned name of the LAG.
- **Status**: Condition of the LAG. If the status is **healthy**, then all ports in the LAG are functioning as expected.
- **Ports**: Port names associated with a LAG.
- **Port Speed**: Maximum speed of each port.
- **LAG Speed**: Combined speed of all ports in the LAG.
- **MAC**: Unique media access control (MAC) address assigned to the network interface. This field cannot be modified.

Creating a Link Aggregation Group

When creating a link aggregation group (LAG):

- All ports must be of the same speed.
- Ports must be created in pairs, mirroring FM1 and FM2. For example, ETH1.1 from FM1 and ETH1.1 from FM2 are a pair of mirrored ports from each FM.

- Ports can only be assigned to one LAG. Ports must be removed from one LAG before being added to another LAG.

To create a LAG, perform the following:

1. From the **Settings** page navigation sidebar, click **Link Aggregation Groups** under **Network**.
2. Click the Add (+) button in the header of the **Link Aggregation Groups** panel. The **Create Link Aggregation Group** pop-up window appears.
3. Select the ports to be assigned to the LAG.
4. Click **Create**.

   Once created, the new LAG will appear in the list of available LAGs when creating and editing subnets.

Adding and Removing Ports to and from a Link Aggregation Group

When adding ports to a link aggregation group (LAG):

- All ports must be of the same speed.
- Ports must be created in pairs, mirroring FM1 and FM2. For example, ETH1.1 from FM1 and ETH1.1 from FM2 are a pair of mirrored ports from each FM.
- Ports can only be assigned to one LAG. Ports must be removed from one LAG before being added to another LAG.

When removing ports from a LAG, ports must be removed in pairs, mirroring FM1 and FM2. For example, ETH1.1 from FM1 and ETH1.1 from FM2 are a pair of mirrored ports from each FM.

To add or remove ports, perform the following:

1. From the **Settings** page navigation sidebar, click **Link Aggregation Groups** under **Network**.
2. In the **Link Aggregation Groups** panel, click the **Edit** icon in the same row of the LAG to which you wish to add or remove ports. The **Edit Link Aggregation Group** pop-up window appears.
3. Select the ports you wish to add or remove.
4. Click **Save**.

Deleting a Link Aggregation Group

Deleting a link aggregation group (LAG) deletes the entire LAG configuration and unbinds the ports.

Note: The **uplink** LAG cannot be deleted

To delete a LAG, perform the following:

1. From the **Settings** page navigation sidebar, click **Link Aggregation Groups** under **Network**.
2. In the **Link Aggregation Groups** panel, click the **Delete** icon in the same row of the LAG you wish to delete. The **Delete Link Aggregation Group** pop-up window appears.
3. Click **Delete**.

Subnets and Interfaces

In the FlashBlade array, data virtual IP addresses (data VIPs) are organized into subnetworks (subnets). Exporting a file system requires a data vip which, in turn, must be attached to a subnet. Data vips that share the same network prefix and gateway are grouped into the same subnet.

Note: IPv6 is not supported with this release for multi-tenancy file access and file systems with NFS and SMB access enabled need to have access to UID/GID sources via LDAP or AD.

Note: Egress traffic for DNS/AD/LDAP/Kerberos/SMB-Audit-Syslog-Server goes over the management network with an exception where the external server IP and FlashBlade datavip is in the same subnet.

The **Subnets** panel displays the subnets on the array. Attached to each subnet are the data virtual IP addresses (data VIPs) that are used to export file systems, which are displayed in the **Interfaces** panel.

To access the **Subnets** and **Interfaces** panels, from the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.

Figure 9.14 Viewing the Subnets & Interfaces Panel

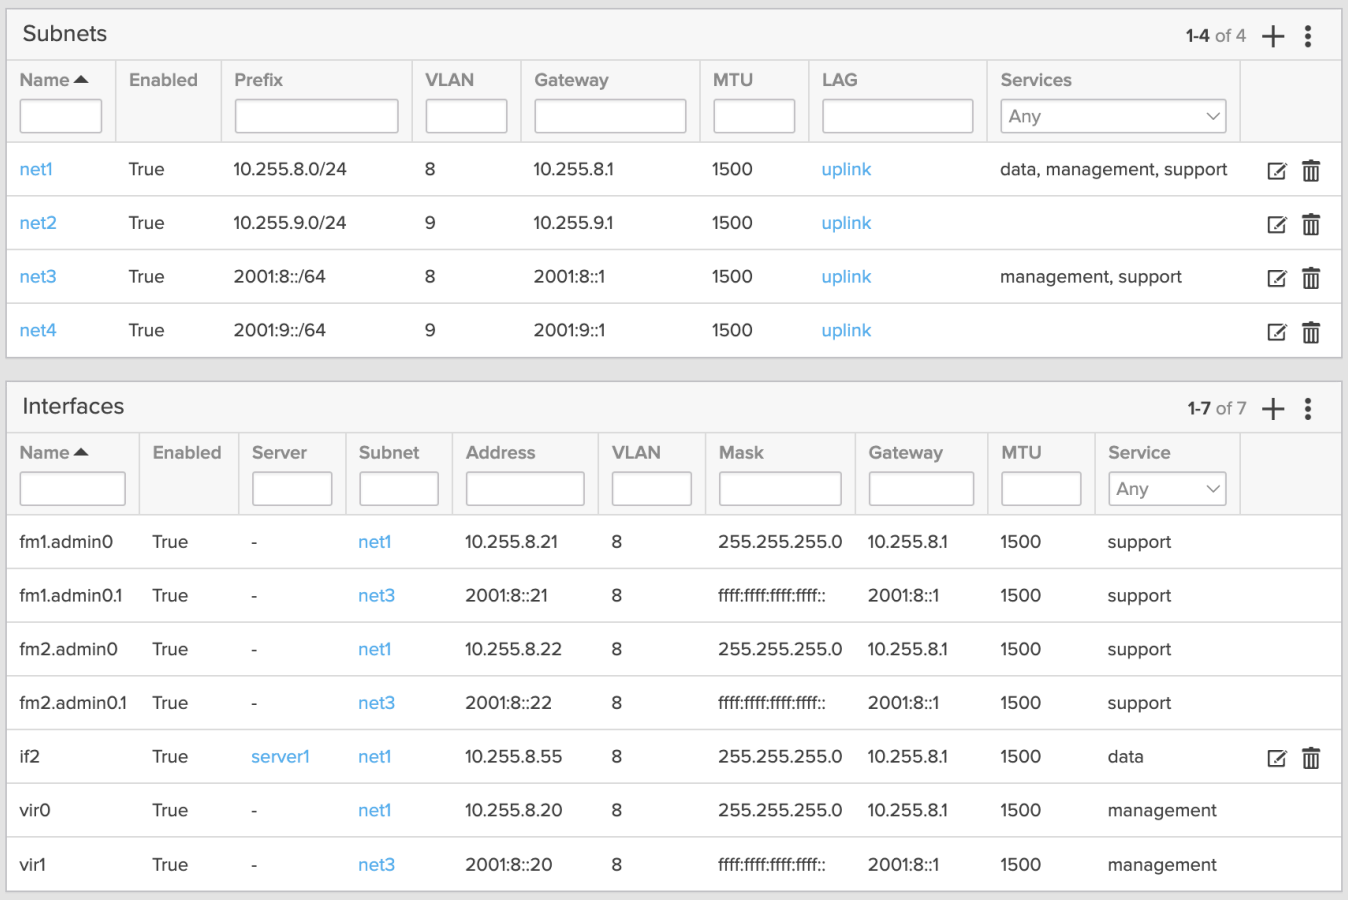

Subnets 1-4 of 4

Name	Enabled	Prefix	VLAN	Gateway	MTU	LAG	Services
net1	True	10.255.8.0/24	8	10.255.8.1	1500	uplink	data, management, support
net2	True	10.255.9.0/24	9	10.255.9.1	1500	uplink	
net3	True	2001:8::/64	8	2001:8::1	1500	uplink	management, support
net4	True	2001:9::/64	9	2001:9::1	1500	uplink	

Interfaces 1-7 of 7

Name	Enabled	Server	Subnet	Address	VLAN	Mask	Gateway	MTU	Service
fm1.admin0	True	-	net1	10.255.8.21	8	255.255.255.0	10.255.8.1	1500	support
fm1.admin0.1	True	-	net3	2001:8::21	8	ffff:ffff:ffff:ffff::	2001:8::1	1500	support
fm2.admin0	True	-	net1	10.255.8.22	8	255.255.255.0	10.255.8.1	1500	support
fm2.admin0.1	True	-	net3	2001:8::22	8	ffff:ffff:ffff:ffff::	2001:8::1	1500	support
if2	True	server1	net1	10.255.8.55	8	255.255.255.0	10.255.8.1	1500	data
vir0	True	-	net1	10.255.8.20	8	255.255.255.0	10.255.8.1	1500	management
vir1	True	-	net3	2001:8::20	8	ffff:ffff:ffff:ffff::	2001:8::1	1500	management

The **Subnets** panel includes the following information:

- **Name** - Subnet name.
- **Enabled** - Subnet status. When a subnet is enabled (**True**), the data VIPs within the subnet that are enabled are automatically available and ready to connect. Newly created subnets are automatically enabled.
- **Prefix** - The network prefix and prefix length.
- **VLAN** - VLAN ID number.
- **Gateway** - IP address of the gateway through which the data vip communicates with devices outisde the given subnet.
- **MTU** - Maximum transmission unit (MTU) for the data vips, in bytes.
- **LAG** - Link aggregation group (LAG) for the data vips.

The **Interfaces** panel includes the following information:

- **Name** - The data VIP name.
- **Enabled** - The interface status. If the subnet to which the interface is attached is enabled, the interface status will display as enabled (**True**).
- **Server** - The server which the data VIP is attached.
- **Subnet** - The subnet to which the data VIP is attached.
- **Addresses** - IP address associated with the data VIP.
- **VLAN** - VLAN ID number inherited from the subnet.
- **Mask** - The subnet mask.
- **Gateway** - IP address of the gateway through which the specified interface is to communicate with devices outside the subnet.
- **MTU** - Maximum transmission unit (MTU) for the data VIPs, in bytes.

Create a new subnet if the desired one does not appear in the **Subnets** panel.

Creating a subnet requires a prefix and VLAN interface. Specify the IP address of the network prefix and prefix length in the form `ddd.ddd.ddd.ddd/dd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx` for IPv6. Specify the VLAN ID to which the subnet is configured. Valid VLAN ID numbers are between 1 and 4094.

Optionally specify the gateway address in the form `ddd.ddd.ddd.ddd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx` for IPv6.

Optionally specify the maximum transmission unit (MTU), in bytes, of the data VIP. If the MTU is not specified during subnet creation, the value defaults to `1500`.

When a data VIP is attached to a subnet, it inherits the gateway, MTU and VLAN ID of the subnet. Likewise, the subnet inherits the `data` service type from the data VIP.

Delete a subnet if it is no longer needed. Note that exporting a file system and file system replication require at least one valid data VIP, which must be attached to a subnet. Therefore, there must always be at least one subnet (for service type `data` or `replication`) on the array.

Replication requires creating a replication VIP. Replication VIPs can be in the same or different subnets as data VIPs, and hence using different physical ports if required. The FlashBlade MTU setting applies

to a subnet (VLAN). The MTU configured for replication VLAN must be the minimum MTU along the link between the source array and the target array.

Creating a Subnet

Data VIPs that share the same network prefix and gateway are grouped into the same subnet. Exporting a file system requires a data VIP, which in turn, must be attached to a subnet. Create a subnet if none of the existing ones meet your requirements.

Replication VIPs can be in the same or different subnets as data VIPs, and hence using different physical ports if required. The FlashBlade MTU setting applies to a subnet (VLAN). The MTU configured for the replication VLAN must be the minimum MTU along the link between the source array and the target array. If you notice the arrays are connected, but replication shows little throughput, ensure that the MTU configured for the replication VLAN is the correct size.

It is recommended to create separate VLANs for data and replication so that data and replication can use different MTU settings.

1. From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
2. Click the **Add** (+) button in the header of the **Subnets** panel. The **Create Subnet** pop-up window appears.
3. In the **Name** field, type the name of the subnet.
4. In the **Prefix** field, type the IP address of the network prefix and prefix length in the form `ddd.ddd.ddd.ddd/dd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx` for IPv6.
5. In the **VLAN** field, specify the VLAN ID to which the subnet is configured. Valid VLAN ID numbers are between 1 and 4094.
6. In the **Gateway** field, type the IP address of the gateway through which the data vip communicates with devices outside this network in the form `ddd.ddd.ddd.ddd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx` for IPv6
7. In the **MTU** field, specify the maximum transmission unit (MTU) of the data VIP. If the MTU is not specified during subnet creation, the value defaults to `1500`.
8. Select a Link Aggregation Group from the **LAG** list.
9. Click **Create**.

Editing a Subnet

To edit a subnet, perform the following:

1. From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
2. In the **Subnets** panel, locate the subnet you wish to edit and click the **Edit** icon to the right of the subnet name. The **Edit Subnet** pop-up window appears.
3. In the **VLAN** field, specify the VLAN ID to which the subnet is configured. Valid VLAN ID numbers are between 1 and 4094.
4. In the **Gateway** field, type the IP address of the gateway through which the data VIP communicates with devices outside this network in the form `ddd.ddd.ddd.ddd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx` for IPv6
5. In the **MTU** field, specify the maximum transmission unit (MTU) of the data vip. If the MTU is not specified during subnet creation, the value defaults to `1500`.
6. Select a Link Aggregation Group from the **LAG** list.
7. Click **Save**.

Deleting a Subnet

A subnet can only be deleted if it has no interfaces.

To delete a subnet, perform the following:

1. From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
2. You can delete a single or multiple subnets.

   To delete a single subnet:

   a. In the **Subnets** panel, locate the subnet you wish to delete and click the **Delete** icon to the right of the subnet name. The **Delete Subnet** pop-up window appears.

   b. Click **Delete**.

   To delete mutliple subnets:

   a. Click the **More Options** button in the **Subnets** header and select **Delete**. The **Delete Subnets** pop-up window appears.

   b. Select one or more subnets from the **Existing Subnets** list.

   c. Click **Delete**.

Creating a Network Interface

Once you have created a subnet you can create and add network interfaces.

Note: IPv6 is not supported with this release for multi-tenancy file access and file systems with NFS and SMB access enabled need to have access to UID/GID sources via LDAP or AD.

Note: Egress traffic for DNS/AD/LDAP/Kerberos/SMB-Audit-Syslog-Server goes over the management network with an exception where the external server IP and FlashBlade datavip is in the same subnet.

To create a network interface, perform the following:

1. From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
2. In the **Interfaces** panel, click the Add button (+). The **Create Network Interface** pop-up window appears.
3. In the **Name** field, enter the interface name.
4. Select a subnet from the **Subnet** list.
5. In the address field, enter the network address.
6. From the **Services** list, select if the interface will be used for **data**, **replication**, or **sts**.
7. (Optional) From the **Select Server** list, select the desired server to be used for multi-tenancy.
8. Click **Create**.

The **Interfaces** list is updated with the new interface.

Editing a Network Interface

To edit a network interface, perform the following:

1. From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
2. In the **Interfaces** panel, locate the nettwork interface you wish to edit and click the **Edit** icon in the same row of the interface. The **Edit Network Interface** pop-up window appears.
3. If you wish to change the subnet, select a subnet from the **Subnet** list.
4. In the address field, enter the network address.
5. From the **Services** list, select if the interface will be used for **data**, **replication**, or **sts**.
6. Click **Save**.

Deleting a Network Interface

Note: Built-in interfaces cannot be deleted. Additionally, custom interfaces in use cannot be deleted. For example, an interface with replication service cannot be deleted if there are remote FB array/S3 target connections configured.

To delete a network interface, perform the following:

1. From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
2. You can delete a single or multiple network interfaces.

   To delete a single network interface:

   a. In the **Interfaces** panel, locate the network interface you wish to delete and click the **Delete** icon to the right of the interface name. The **Delete Network Interface** pop-up window appears.

   b. Click **Delete**.

   To delete multiple network interfaces:

   a. Click the **More Options** button in the **Interfaces** header and select **Delete**. The **Delete Network Interfaces** pop-up window appears.

   b. Select one or more interraces from the **Existing Network Interfaces** list.

   c. Click **Delete**.

Working with Security Settings

The following **Security** settings are available:

- **API Clients** - Displays and manages configured API clients.
- **Audit Trail** - Displays a chronological history of modification operations (for example, commands that create, update, and delete) that were run on the array.
- **Certificates**
  - **Array Certificates** - Displays, imports, and exports array certificates.
  - **External Certificates** - Displays and manages CA certificates.
  - **Certificate Groups** - Displays and manages certificate groups.

- **Directory Service** - Displays and manages the integration of FlashBlades with Active Directory and directory services. Also displays the FlashBlade's SMB authentication mode.
- **Kerberos Keytabs** - Displays and manages Kerberos keytab files.
- **Rapid Data Locking** - Displays and manages Rapid Data Locking (RDL) and KMIP servers.
- **Session Log** - Displays the history of each access and access attempts to the FlashBlade and duration of each FlashBlade session.
- **Single Sign-On** - Displays and manages single sign-on (SSO) configuration.
- **Users** - Displays and manages local and remote FlashBlade array users.

API Clients

The use of API clients is an alternate method of token authorization that is available on FlashBlade.

The **API Clients** panel displays and provides management of API clients used for authorization with the FlashBlade.

To access the **API Clients** panel, from the **Settings** page navigation sidebar, click **API Clients** under **Security**.

Figure 9.15 Viewing the API Clients Panel

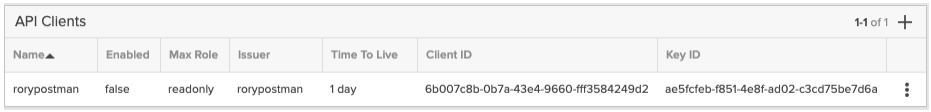

The following information is displayed:

- **Name**: The API client name.
- **Enabled**: Whether the API client is enabled (**true**) or disabled (**false**) to exchange its ID tokens for access tokens.
- **Max Role**: The role associated with the API client.
- **Issuer**: The name of the API client's issuer.
- **Client ID**: The UUID that uniquely identifies the API client.
- **Key ID**: The UUID that uniquely identifies the issuer's public key.

Creating an API Client

Note: A public key is required when creating an API client. The public key should be PEM formatted (Base64 encoded), include the " -----BEGIN PUBLIC KEY----- " and " -----END PUBLIC KEY----- " lines..

To create an API client, perform the following:

1. From the **Settings** page navigation sidebar, click **API Clients** under **Security**.
2. Click the Add (+) button in the header of the **API Clients** panel. The **Create API Client** pop-up window appears.
3. Enter the API client's name.
4. The **Issuer** defaults to the name entered for the API client.
5. Select the role (permissions) of the API client from the **Max Role** list.
6. In the **Time To Live** field, enter the duration that the API token will be valid. The value must be entered in the format N[s|m|h|d]. For example **45s**, **50m**, **6h**, **3d**, etc. By default, the value is one day.
7. Enter the public key into the **Public Key** field.
8. Click **Create**.

   Once the API client is created, it is displayed in the **API Clients** panel as disabled. You must manually enable it.

After creating and enabling an API client, you can use the Pure Unified 2.x SDK to make REST API calls to the FlashBlade. To create a session, you will need to specify the API client name, issuer, client ID, and key ID as shown in the GUI, along with the corresponding private key to the 2.x REST SDK.

Note: The FlashBlade 1.x REST SDK only supports API tokens. API clients for authentication are not supported. See the *FlashBlade Management REST API guide* for additional information.

Enabling an API Client

To enable an API client, perform the following:

1. From the **Settings** page navigation sidebar, click **API Clients** under **Security**.
2. In the **API Clients** panel, locate the API client you wish to enable. Click the **More Options** button in the same row and select **Enable API Client**. The **Enable API Client** pop-up window appears.
3. Click **Enable**.

Viewing an API Client's Public Key

To view an API client's public key perform the following:

1. From the **Settings** page navigation sidebar, click **API Clients** under **Security**.
2. In the **API Clients** panel, locate the API client whose public key you wish to view. Click the **More Options** button in the same row and select **View Public Key**. The **View Public Key** pop-up window appears displaying the key.
3. Click **Close** when finished.

Deleting an API Client

To delete an API client, perform the following:

1. From the **Settings** page navigation sidebar, click **API Clients** under **Security**.
2. In the **API Clients** panel, locate the API client you wish to delete. Click the **More Options** button in the same row and select **Delete API Client**. The **Delete API Client** pop-up window appears.
3. Click **Delete**.

Viewing the Audit Trail

The **Audit Trail** panel provides a chronological history of modification operations (for example, commands that create, update, and delete) that were run on the array.

Figure 9.16 Viewing the Audit Trail Panel

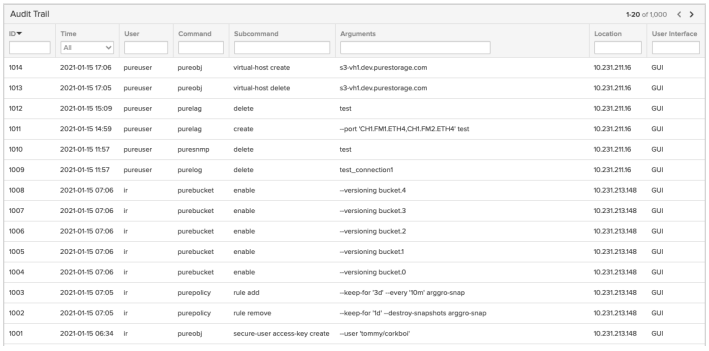

To view the history of modification operations issued on the array, from the **Settings** page navigation sidebar, click **Audit Trail** under **Security**.

The **Audit Trail** panel provides the following information:

- **ID**: The chronological number the command was used.

- **Time**: The date and time the command was run.
- **User**: The user who ran the command.
- **Command**: The name of the command.
- **Subcommand**: The subcommand used with the command.
- **Arguments**: Any arguments used with the command and their associated user-supplied input.
- **Location**: The IP address where the command was issued.
- **User Interface**: The interface from which the operation was run. For example, CLI, GUI, REST.

Array Certificates

Array certificates are used by clients to validate the array. Array certificates require private keys.

⚠ **Important!** The global certificate is the default certificate after setup of the system for both management GUI/REST and data (ObjectStore/S3) access. It is also the default certificate used by the Default TLS Policy. It can be updated but it cannot be removed.

Purity creates a self-signed certificate and private key when you start the system for the first time. The SSL Certificate panel allows you to view the certificate and import or export certificates and private keys. Intermediate certificates can also be imported to the system.

Imported SSL certificates must be PEM formatted (Base64 encoded), include the "`-----BEGIN CERTIFICATE-----`" and "`-----END CERTIFICATE-----`" lines.

The **Array Certificates** panel displays imported array and self-signed certificates on the FlashBlade.

To access the **Array Certificates** panel, from the **Settings** page navigation sidebar, click **Certificates** under **Security**.

Figure 9.17 Viewing the Array Certificates Panel

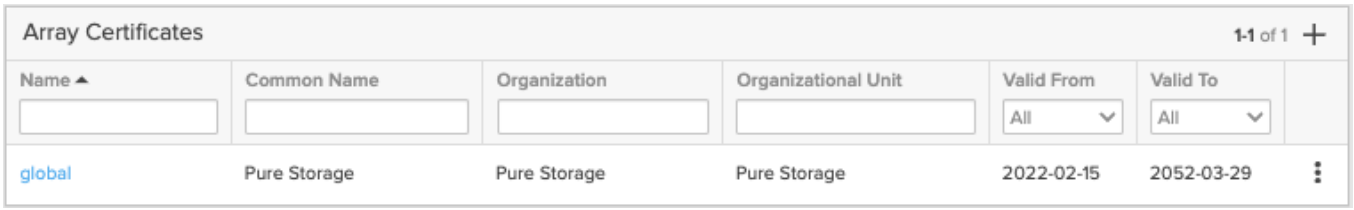

The following information is displayed:

- **Name**: The name of the array certificate. Clicking the array certificate name opens the **Certificate Detail**s panel.
- **Common Name**: The name of the website or entity using the certificate.
- **Organization**: The name of the company or organization.
- **Organizational Unit**: The specific department within the organization.
- **Valid From**: The date from which the certificate is valid.
- **Valid To**: The date the certificate expires.

Array certificates used for signing and decryption of the SAML2 SSO service provider can also be imported from the **Array Certificates** panel.

FlashBlade allows you to create and update self-signed certificate objects, as well as create certificate signing requests based on existing certificate objects.

Self-signed certificates are digital certificates that are signed using their own private key and are a quick and cost-effective option for internal testing and development environments where external trust is not a requirement.

Certificate signing requests (CSRs) are formal requests to a Certificate Authority (CA) to issue a digital certificate containing information about the requesting entity and public key.

Importing Array Certificates

To import array certificates, perform the following:

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
2. Click the Add + button in the **Array Certificates** panel header. The **Import Array Certificate** pop-up window appears.
3. Select the scope type.
4. In the **Name** field, enter the certificate name.
5. Complete or modify the following fields:
   - **Certificate** - Click **Choose File** and select the signed certificate. Verify the certificate is PEM formatted (Base64 encoded), include the " `-----BEGIN CERTIFICATE-----` " and " `-----END CERTIFICATE-----` " lines.
   - **Private Key** - Click **Choose File** and select the private key.

- **Intermediate Certificate** - (Optional) Click **Choose File** and select the intermediate certificate.
- **Key Passphrase** - (Optional) If the private key is encrypted with a passphrase, enter the passphrase.

6 Click **Import**.

Viewing Array Certificate Details

To view an array certificate's details, perform the following:

1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.

2 Click the certificate whose information you wish to view in the **Array Certificates** panel. The **Certificate Details** and **Used By** panels appear.

Exporting Array Certificates

To export array certificates, perform the following:

1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.

2 Locate the certificate you wish to export in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Export Certificate**. The **Export Certificate** pop-up window appears.

3 Click **Download**.

Updating Array Certificates with Importing Options

To update an array certificate with importing operations, perform the following:

1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.

2 Locate the certificate you wish to update in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Update Certificate with Importing Operations**. The **Update Certificate with Importing Operations** pop-up window appears.

3 Complete or modify the following fields:

- **Certificate** - Click **Choose File** and select the signed certificate. Verify the certificate is PEM formatted (Base64 encoded), include the " `-----BEGIN CERTIFICATE-----` " and " `-----END CERTIFICATE-----` " lines.
- **Private Key** - Click **Choose File** and select the private key.
- **Intermediate Certificate** - (Optional) Click **Choose File** and select the intermediate certificate.

- **Key Passphrase** - (Optional) If the private key is encrypted with a passphrase, enter the passphrase.

4 Click **Update**.

Deleting Array Certificates

To delete a non-global array certificate, perform the following:

⚠ **Important!** The global certificate is the default certificate after setup of the system for both management GUI/REST and data (ObjectStore/S3) access. It is also the default certificate used by the Default TLS Policy. It can be updated but it cannot be removed.

1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.

2 Locate the certificate you wish to delete in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Delete**. The **Delete Certificate** pop-up window appears.

3 Click **Delete**.

External Certificates

The **External Certificates** panel displays and allows you to import external (CA) certificates for the FlashBlade.

CA certificates are used to provide security authentication between two identities, allowing the FlashBlade to communicate securely with servers and clients.

To access the **External Certificates** panel, from the **Settings** page navigation sidebar, click **Certificates** under **Security**.

Figure 9.18 Viewing the External Certificates Panel

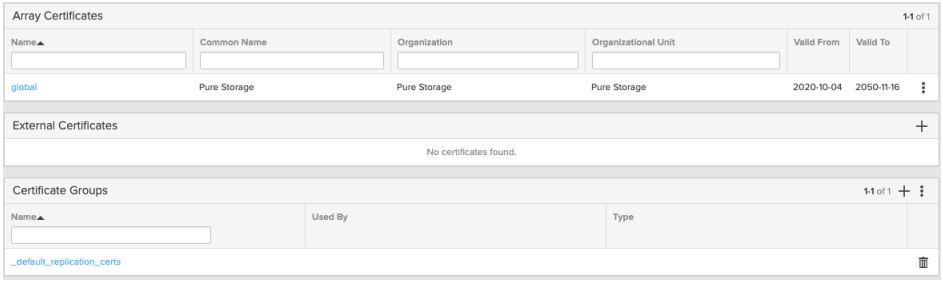

Importing an External Certificate

External (CA) certificates provide the security authentication between two identities. This allows the FlashBlade array to communicate securely with servers and clients.

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
2. Click the Add (+) button in the header of the **External Certificates** panel. The **Import CA Certificate** pop-up window appears.
3. In the **Name** field, type the name you wish to assign to the certificate.
4. In the **Type** field, **external** is displayed indicating the certificate used when the FlashBlade array acts as a client to validate a server.
5. In the **CA Certificate** field, click **Choose File** and select the certificate file.
6. Click **Import**.

Viewing Certificate Details

To view a certificate's details, perform the following:

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
2. In the **External Certificates** panel click the certificate whose details you wish to view.

   The **Certificate Details** panel appears displaying the certificate's information.

Exporting a Certificate

To export a certificate, perform the following:

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
2. In the **External Certificates** panel locate the certificate you wish to export. Click the **More Options** button in the same row and select **Export Certificate**. The **Export Certificate** pop-up window appears.
3. Click **Download**.

Deleting a Certificate

To delete a certificate, perform the following:

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.

2 In the **External Certificates** panel locate the certificate you wish to delete. Click the **More Options** button in the same row and select **Export Certificate**. The **Delete Certificate** pop-up window appears.

3 Click **Delete**.

Certificate Groups

LDAP over TLS (Transport Layer Security) can be optionally configured to encrypt communication between the FlashBlade array, acting as an LDAP client, and LDAP servers. In order to establish LDAP over TLS, the LDAP server and FlashBlade must have mutually supported versions of TLS. FlashBlade supports TLS versions 1.0, 1.1, 1.2, and 1.3. The LDAP server must support at least one of these versions. The highest matching TLS version on the LDAP server and FlashBlade is used.

⚠ **Important!** The LDAP over TLS configuration does not support NFS with NIS (Network Information Service).

LDAP over TLS requires that each LDAP server's CA certificate be imported to the FlashBlade where they can be used individually by the FlashBlade or grouped into CA certificate groups to verify the LDAP server's identity and establish communication over TLS. CA certificate groups are used to manage multiple trusted certificates issued to different servers as a directory of certificates. A CA certificate group can contain one or more CA certificates. A CA certificate can be used in multiple CA certificate groups. CA certificates can be added, removed, or deleted without disruption to the services using them as long as the certificates for any configured servers remain in the group. While a CA certificate group is in use, it cannot be deleted, nor can you remove or delete all CA certificates in that group.

⚠ **Important!** Adding multiple CA certificates in a certificate group which secure the same hostname can produce undefined behavior—there is no guarantee which certificate will be used.

The **Certificate Groups** panel displays all configured certificate groups for the FlashBlade.

To access the **Certificate Groups** panel, from the **Settings** page navigation sidebar, click **Certificates** under **Security**.

Figure 9.19 Viewing the Certificate Groups Panel

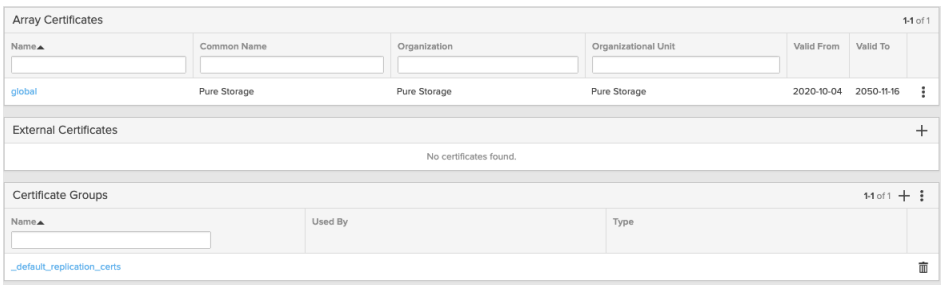

Creating a Certificate Group

A group can be created to house a collection of CA certificates. After you create a group, certificates can be added to that group.

To create a certificate group, perform the following:

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
2. Click the Add (+) button in the header of the **Certificate Groups** panel. The **Create Certificate Group** pop-up window appears.
3. In the **Name** field, type the name you want to assign to the certificate group.
4. Click **Create**.

Adding a Certificate to a Group

⚠ **Important!** Adding multiple CA certificates in a certificate group which secure the same hostname can produce undefined behavior—there is no guarantee which certificate will be used.

Once a certificate group is created, you may add certificates to that group.

To add a CA certificate to a certificate group, perform the following:

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
2. In the **Certificate Groups** panel, click the certificate group to which you wish to add certificates. The selected group's certificate details page opens.
3. In the **Certificates** panel, click the Add (+) button. The **Add Certificates** pop-up window appears.
4. Select the certificate(s) you wish to add to the group from the **Available Certificates** column.
5. Click **Add**.

Removing a Certificate from a Group

Certificates can be removed from a certificate group if the certificate has expired or is no longer used.

To remove a CA certificate from a certificate group, perform the following:

1. From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
2. In the **Certificate Groups** panel, click the certificate group from which you wish to remove certificates. The selected group's certificate details page opens.

3 In the **Certificates** panel, click the **More Options** button and select **Remove**. The **Remove Certificates** pop-up window appears.

4 Select the certificates you wish to remove from the group from the **Available Certificates** column.

5 Click **Remove**.

Deleting a Certificate Group

Note: While a certificate group is in use, it cannot be deleted, nor can you remove or delete all CA certificates in that group.

To delete a certificate group, perform the following:

1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.

2 In the **Certificate Groups** panel, locate the certificate group you wish to delete and click the **Delete** button in the same row. The **Delete Certificate Group** pop-up window appears.

3 Click **Delete**.

Directory Service

The **Directory Service** panel manages the integration of FlashBlades with an existing directory service.

To access the **Directory Service** panel, from the **Settings** page navigation sidebar, click **Directory Service** under **Security**.

Figure 9.20 Viewing the Directory Service Panel

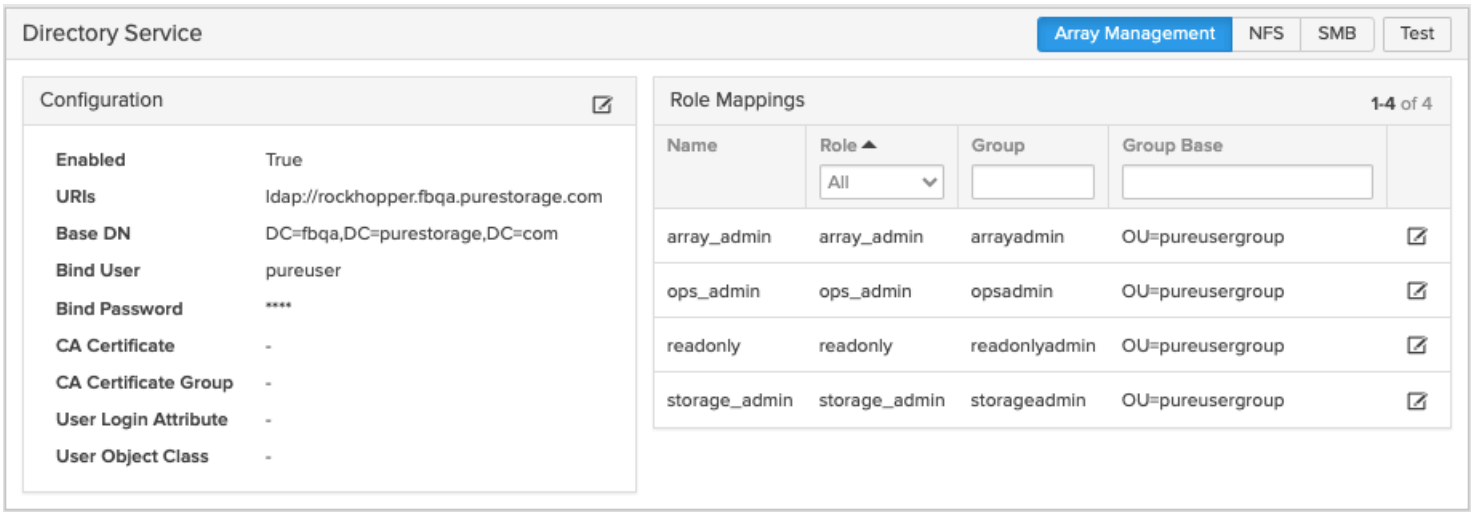

When the **Directory Service** panel's user roles and directory service protocols are configured and enabled, the FlashBlade leverages a directory service to perform user account and permission level searches.

The FlashBlade supports a single local user, named `pureuser`, with array-wide (array_admin) permissions. Users can be added to the array through Lightweight Directory Access Protocol (LDAP) by integrating the array with a directory service such as Active Directory or OpenLDAP and assigning roles for one or more groups. Users can sign into the array only if they are in a group that has been assigned a role.

To integrate the FlashBlade array with a directory service:

1. Configure the FlashBlade directory service.
2. Test the directory service configuration.
3. Enable the directory service.

When configuring the directory service for array management or NFS protocol integration with LDAP, specify the URIs, base DN, and bind username and password of the directory service. If selecting the NIS directory service for NFS, clear any existing LDAP configurations before specifying the NIS servers and NIS domain.

⚠ **Important!** Anonymous binding is supported for NFS and Management. If configuring anonymous binding for NFS or Management, neither the bind user nor the bind password should be specified.

A "bind" account must be configured to allow the array to read the directory. The bind account should be tied to a service account (rather than an individual), and the account password should never expire or be required to change periodically. The bind account must have read privileges for users and groups and the privileges should be configured at the Organizational Unit.

For Active Directory, users with disabled accounts are not able to access the file systems.

After the directory service has been configured, test the configuration to verify that the URIs can be resolved and the FlashBlade array can bind and query the tree using the bind user credentials.

Enable the directory service after it has passed the directory service test. The directory service can be disabled at any time. Disabling the directory service stops any users in the directory from accessing the file system. In the case of management, disabling directory service prevents users from logging into the array.

For instructions on configuring each directory service, see:

- *Configuring the Array Management Directory Service*

- *Configuring the NFS Directory Service with LDAP*
- *Configuring the NFS Directory Service with NIS*

Kerberos Keytabs

Purity//FB supports Kerberos authentication with directory services and Active Directory (AD) accounts.

Kerberos authentication is automatically enabled when configuring FlashBlade with an Active Directory account (see *Active Directory Overview*).

Kerberos network authentication protocols (krb5, krb5i, and krb5p) for NFSv3,NFSv4.1 and SMB clients, in addition to auth_sys, are supported. With Kerberos, NFSv3, NFSv4.1, and SMB clients have to be authenticated by a Key Distribution Center (KDC) to get a session ticket. They have to present this session ticket to the FlashBlade array in order to get access to the file system.

Keytab files can be generated for a FlashBlade array by a KDC, and can then be imported and managed by the FlashBlade. Keytab files contain list of Kerberos principal names and encryption keys to automatically authenticate clients without requiring human interaction or access to password stored in a plain-text file. Keytab file must be manually uploaded with NFS or SMB directory services. With AD accounts, keytab files are automatically generated.

The **Kerberos Keytabs** panel displays and allows management of Kerberos keytabs on the FlashBlade for use with directory services and AD accounts.

To access the **Kerberos Keytabs** panel, from the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.

Figure 9.21 Viewing the Kerberos Keytabs Panel

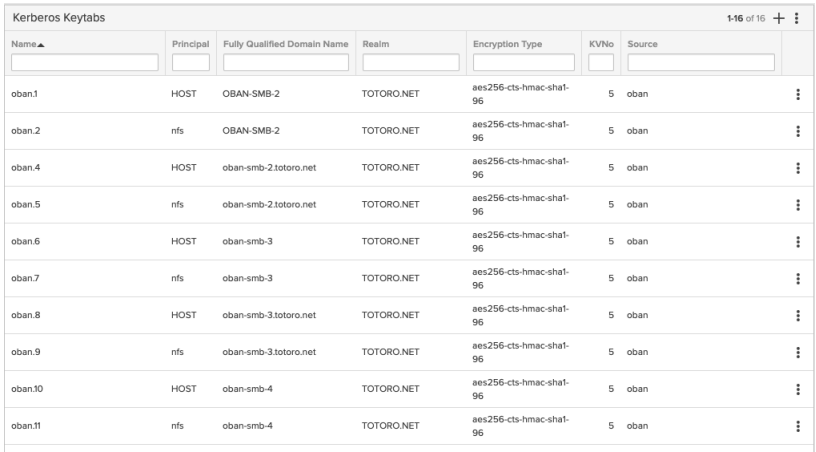

Keytabs are displayed from newest to oldest. The following information is displayed:

- **Name**: The name of the keytab.
- **Principal**: The service principal name.
- **Fully Qualified Domain Name**: The fully qualified domain name that the client will mount the file system over for NFS or SMB. This can be any of the FQDNs that have been registered for data VIPs on the FlashBlade.
- **Realm**: The Kerberos realm that issued the keytab.
- **Encryption Type**: The suite of encryption ciphers used for kerberized communication with clients whose Kerberos tickets are issued against this keytab.
- **KVNo**: The version number of the key used to issue the keytab.
- **Source**: The keytab source. If generated by an AD account, the AD account would be displayed.

Importing a Keytab File

To import a keytab file, perform the following:

1. From the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.
2. Click the **More Options** button in the header of the **Kerberos Keytabs** panel and select **Import**. The **Import Keytab File** pop-up window appears.
3. Enter a prefix in the **Name Prefix** field. Specifying a file entry prefix is required because a single keytab file can contain multiple keytab entries in multiple slots.
4. Click **Choose File** in the **Keytab File** field. Navigate to, and select the keytab file you wish to import.
5. Click **Import**.

Exporting a Keytab File

To export a keytab file that is currently present on your FlashBlade, perform the following:

1. From the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.
2. To export one or more keytab files:
   a. Click the **More Options** button in the header of the **Kerberos Keytabs** panel and select **Export**. The **Export Keytabs** pop-up window appears.

b Select one or more keytab files from the **Existing Keytabs** column on the left. Each selected keytab file appears under the **Selected Keytabs** column on the right.

c Click **Export**.

3 To export a single keytab:

a From the **Kerberos Keytabs** panel, click the **More Options** button in the same row as the keytab file you wish to export. and select **Export**. The **Export Keytab File** pop-up window appears.

b Click to export in **Binary** or **MIME** format.

Deleting a Keytab File

Note: Deleting keytab files can be a disruptive action if there were clients with active sessions that relied on the deleted keytab.

To delete a keytab file, perform the following:

1 From the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.

2 To delete one or more keytab files:

a Click the **More Options** button in the header of the **Kerberos Keytabs** panel and select **Delete**. The **Delete Keytabs** pop-up window appears.

b Select one or more keytab files from the **Existing Keytabs** column on the left. Each selected keytab file appears under the **Selected Keytabs** column on the right.

c Click **Delete**.

3 To delete a single keytab:

a From the **Kerberos Keytabs** panel, click the **More Options** button in the same row as the keytab file you wish to delete and select **Delete**. The **Delete Keytab** pop-up window appears.

b Click **Delete**

Session Log

The **Session Log** screen provides a layer of security to monitor and record who has accessed or attempted to access the FlashBlade and the length of their session. To access the **Session Log** screen, from the **Settings** page navigation sidebar, click **Session Log** under **Security**.

Figure 9.22 Viewing the Session Log

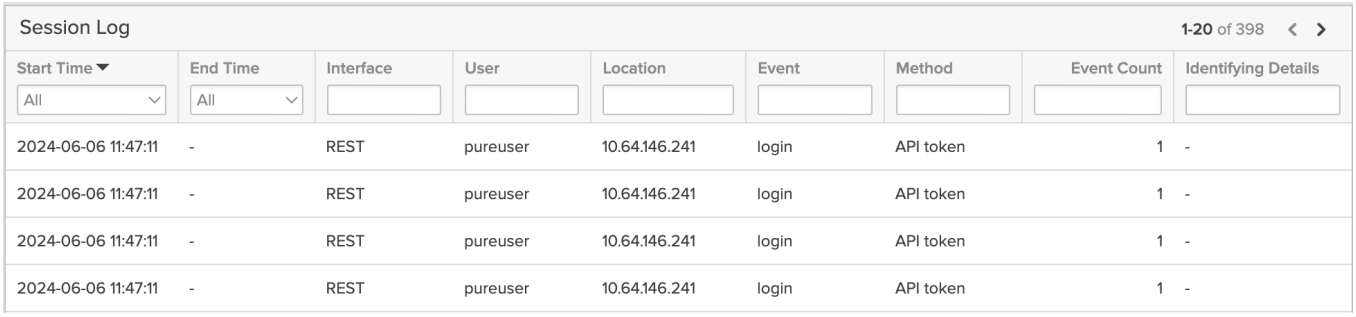

The **Session Log** displays the following inforrmation:

- **Start Time** - The timestamp of when the user logged into the array. Note that at most, 1,000 of the most recent login records are kept. If the number of records exceeds the memory threshold, the oldest start time(s) are removed and a dash (-) is displayed instead.
- **End Time** - The timestamp of when the user logged out of the array. If the user has not yet logged out, a dash (-) is displayed. Note that the log out time is only displayed for the FlashBlade CLI and REST API interfaces.
- **User** - Users include **ir**, **pureuser**, **pureeng**, **UNKNOWN**, and directory service users or bad user names.
  - Users **ir** and **pureeng** are used for support operations.
  - For security purposes, to avoid logging passwords entered in the wrong GUI box, **UNKNOWN** is used for failed logins via the GUI. Failed REST logins are not associated with users and are also displayed as **UNKNOWN**.
  - **UNKNOWN** is used when GUI and REST access has been denied by a network policy rule.
  - **UNKNOWN** is used when CLI access has been denied by a network policy rule as a result of remote SSH fully disabled for a client source. **ir** is used if only superuser CLI access is denied (governed by the network policy rule, local-network-superuser-password-access).
  - Successful logins will always show the user regardless of what interface was used to login.
- **Interface** - The type of user interface the user logged in from. Interface types are **CLI**, **GUI**, **REST**, and **CONSOLE**.
- **Location** - The location of the user, indicated as an IP address.
- **Event** - The action of the the user's session. Types of events include:

  - **login** - Displayed when the user has logged in and has not yet logged out.
  - **logout** - Displayed when the user has logged out.
  - **user session** - Displayed when the user has logged out. Note that due to memory limitations, when a **Start Time** event has been removed from memory, a **user session** record will change to **logout**.
  - **request without session** - Displayed when a REST request fails.
  - **failed authentication** - Displayed when attempts to access FlashBlade fail.
- **Method** - The access authentication used to log in. Types of authentication are **API token**, **certificate**, **password**, **one-time password**, and **public key**.
  - **one-time password** is only applicable for the user **pureeng**.
  - For network policy rule denial events, the method can be null, but will be **password** for superuser password access.
- **Event Count** - The number of times an action was attempted during a 15 minute time period. Note that the counter increments for unsuccessful events such as **failed authentication** or **request without session** only. The event counter will refresh every 15 minutes.
- **Identifying Details** - The detailed information of the user's SSH certificate and logins with an SSH public key.

Single Sign-On

FlashBlade supports Single Sign-On (SSO) authentication on the management GUI and AWS Security Token Service (STS) on the object store.

FlashBlade supports both the OpenID Connect (OIDC) and SAML (SAML2) authentication protocols to allow identity providers (IdP) and service providers (SPs) (i.e. the FlashBlade) to implement user authentication and access control.

- SAML2 SSO can be configured for both management and object services.
- OIDC SSO can be configured for object services.

SSO Authentication for Object Service

FlashBlade supports OpenID Connect (OIDC) and SAML (SAML2) authentication protocols for SSO to the object store.

OIDC and SAML are protocols used for identity federation in object services allowing the creation of temporary credentials for federated identities, which enables secure access to object storage. This is achieved through the use of roles and trust policies, which define the permissions and trusted entities that can assume these roles.

Purity//FB allows the configuration of OIDC to work with identity providers (IdPs) such as Okta, Google, X/Twitter, etc. The process involves setting up an OIDC identity provider configuration on the FlashBlade array, which includes providing the IdP URL and optionally uploading necessary certificates if the IdP URL points to a server running with certificates not signed by a well known certificate authority.. Once configured, roles can be created and associated with trust policies that specify the IdP users that are allowed to assume these roles and access the object storage.

Configuring SAML2 SSO for FlashBlade allows federated identities to assume roles and securely access object storage through the use of roles and trust policies, which define the permissions and trusted entities that can assume these roles. Configuration involves setting up the identity provider (IdP) on FlashBlade with the necessary certificates and URLs.

FlashBlade supports up to five SAML and five OIDC configurations for object service.

SSO Authentication for Management Service

⚠ **Important!** Do not enable Multi-factor Authentication with SAML2 SSO for management service before testing the configuration. If the configuration is enabled with errors, the user will be locked out. Go to Testing the SAML2 SSO Configuration section to perform the test.

Purity//FB supports single sign-on (SSO) integration with the Microsoft® Active Directory Federation Services (AD FS), Okta, Azure Active Directory (Azure AD), and Duo Security via the SAML2 protocol. When SAML2 SSO for management service is configured and enabled, all user logins to the Purity//FB GUI are redirected to the IdP login page for single sign-on.

For example, the AD FS identity provider (IdP) can optionally be configured to enable multi-factor authentication (MFA) and supports X.509 certificates, Microsoft Azure™ authentication, and other authentication methods, in addition to password authentication. This release is verified with X.509 certificate authentication as the multi-factor authentication method.

SSO Limitations

The following considerations apply to this release:

- Only one SAML2 SSO configuration for management service can be created at a time on an array.
- Up to five SAML and five OIDC configurations for object service can be created at a time on an array.
- A single SSO configuration for both object and management services is not allowed.
- SAML2 SSO for management service authentication applies only to GUI logins. SSH logins continue to use their existing password authentication or other authentication mechanism.

Viewing Identity Provider Configurations

ODIC and SAML2 SSO identity provider configurations are displayed and managed in the **ODIC Identity Providers** and **SAML Identity Providers** panels in the **Single Sign-On** screen. To access the **Single Sign-On** screen, from the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.

Figure 9.23 Viewing the Single Sign-On Screen

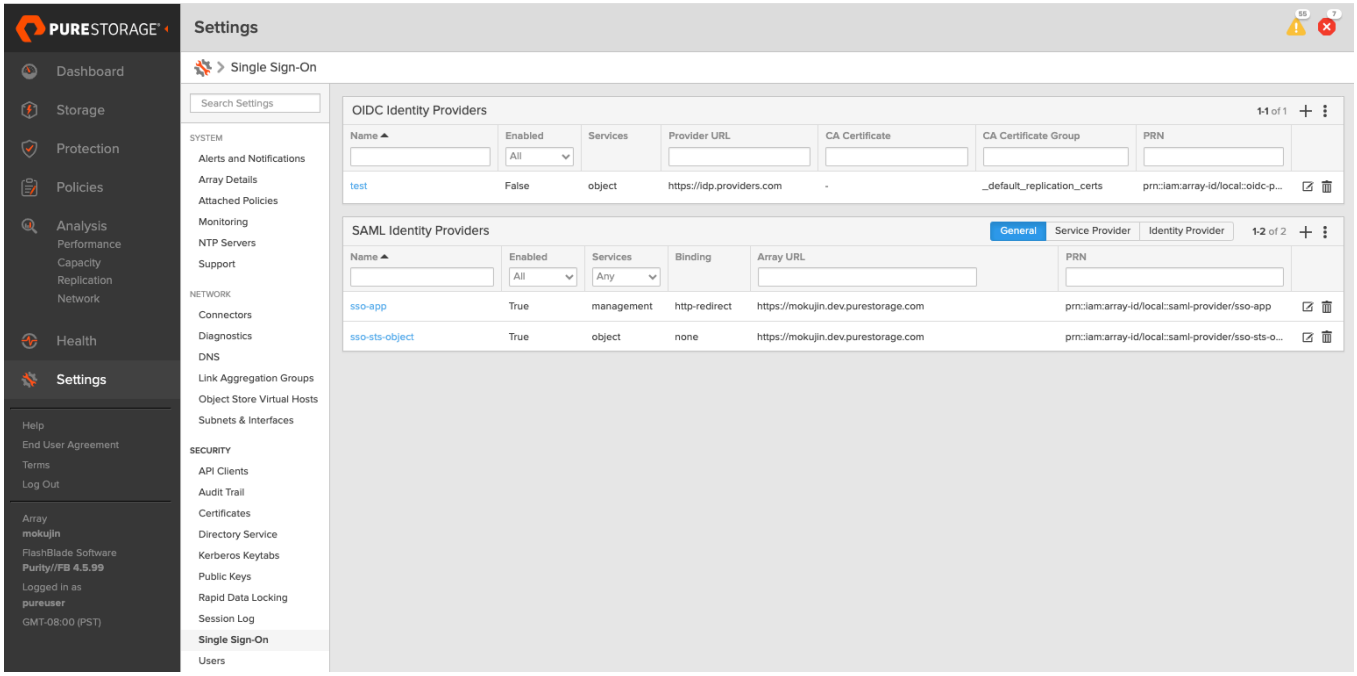

Two panels, **OIDC Identity Providers** and **SAML Identity Providers**, provide information about OIDC and SAML identity provider configurations. General (default), Service Provider and Identity Provider-specific information can be displayed in the **SAML Identity Providers** panel.

Clicking the name of an OIDC or SAML identity provider configuration opens a details screen for that configuration.

The **OIDC Identity Provider** details screen provides a summary of the OIDC identity provider configuration information.

Figure 9.24 Viewing OIDC Identity Provider Configuration Details

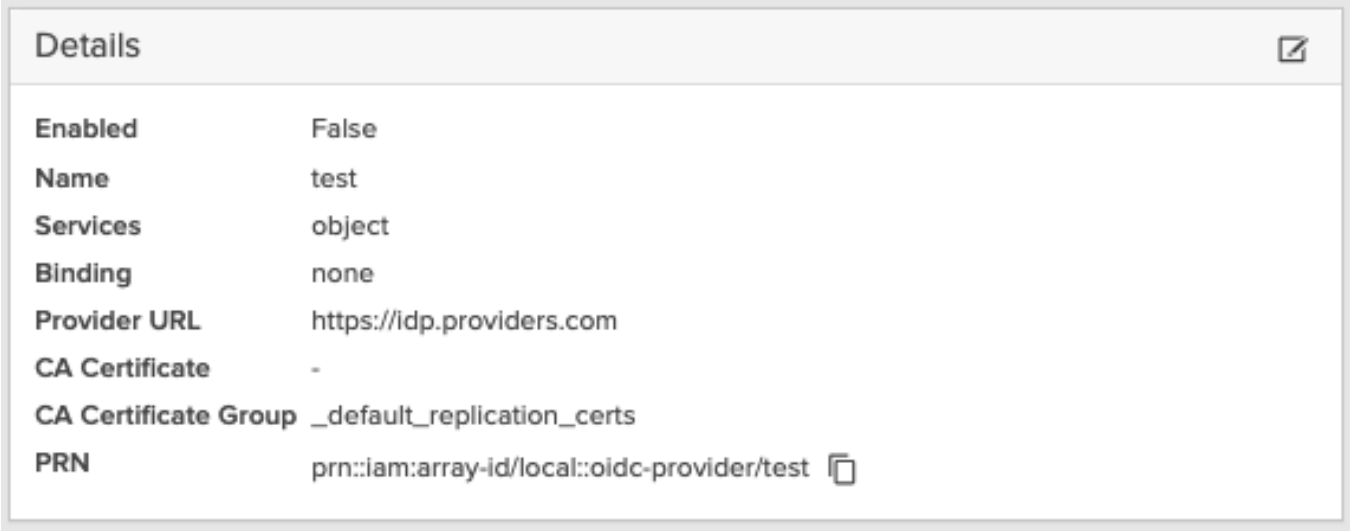

The **SAML Identity Provider** details screen provides two panels, **Configuration** and **Role Mappings**.

Figure 9.25 Viewing SAML Identity Provider Configuration Details

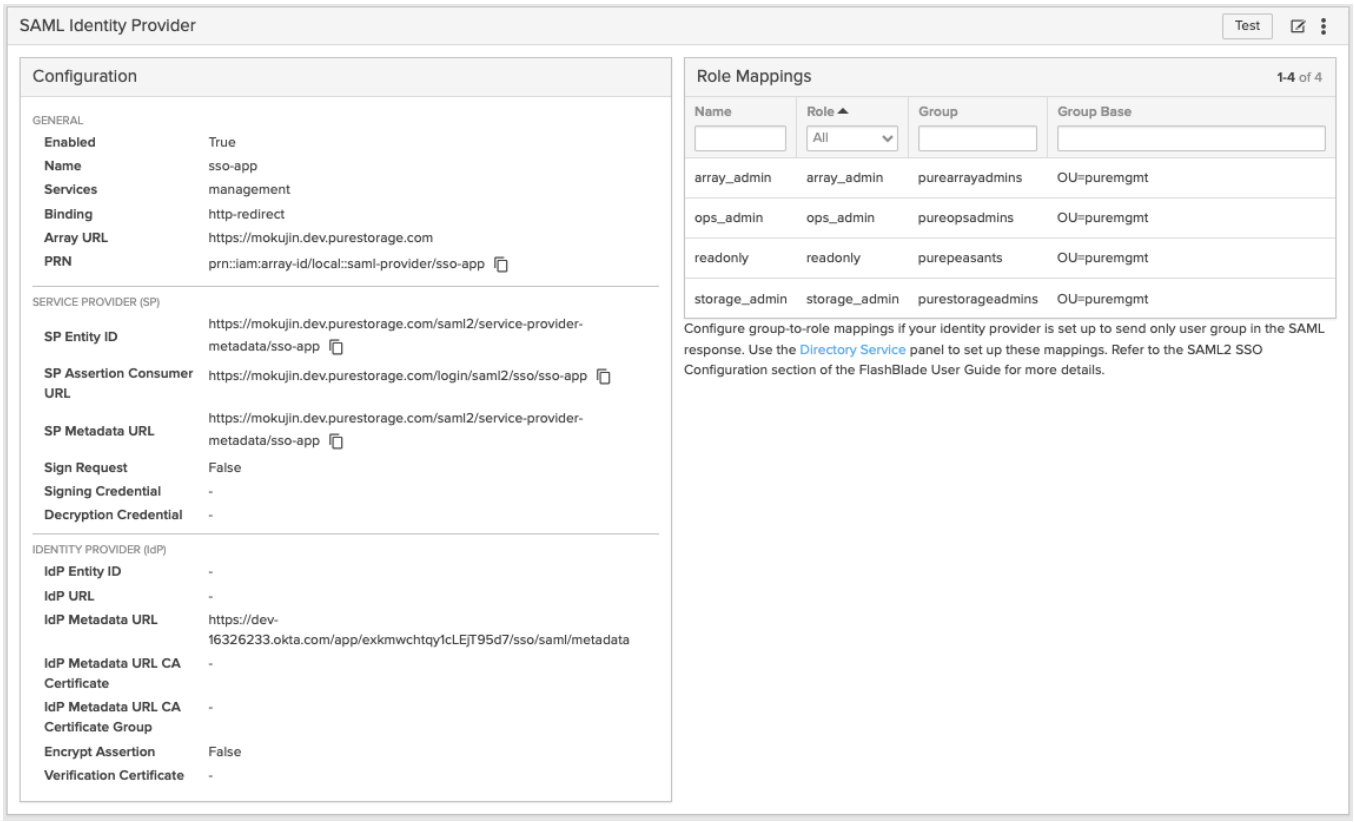

The **Configuration** panel displays the following:

- **General**
  - **Enabled** : If SSO is enabled (**True**) or not (**False**).
  - **Name**: The SSO configuration name.
  - **Services**: The service type, either **management** or **object**.
  - **Binding**: The binding method. For management service, the binding method is **http-direct**. For object service, the binding method is **none**.
  - **Array URL**: The FlashBlade array's URL.
- **Service Provider (SP)**
  - **SP Entity ID**: The URL of the SP (not used for object service).
  - **SP Assertion Consumer URL**: The URL where the identity provider will send its SAML response after authenticating a user (not used for object service).

  - **SP Metadata URL**: The URL to the SP metadata file (not used for object service).
  - **Sign Request**: If the sign request is enabled (**True**) or not (**False**). When enabled, the matching signing credential must be confifured at the identity provider (not used for object service).
  - **Signing Credential**: The credential used by the service provider to sign SAML requests (not used for object service).
  - **Decryption Credential**: The credential used by the service provider to decrypt encrypted SAML assertions from the identify provider.
- **Identity Provider (IdP)**
  - **IdP Entity ID**: The globally unique name for the IdP.
  - **IdP URL**: The URL of the IdP.
  - **IdP Metadata URL**: The URL to the IdP metadata file.
  - **IdP Metadata URL CA Certificate**: The URL to the CA certificate used to validate the authenticity of the metadata file.
  - **IdP Metadata URL CA Certificate Group**: The URL to the group of CA certificates used to validate the authenticity of the metadata file.
  - **Encrypt Assertion**: If the encryption assertion is enabled (**True**) or not (**False**). When enabled, certificates matching the decryption credential must be configured at the identity provider.
  - **Verification Certificate**: The certificate used by the SP to verify the signed SAML response from the IdP.

The **Role Mappings** panel displays current role mapping as configured from the **Directory Service** panel for management service SSO. Note that the **Role Mappings** panel is read only. Any changes to role mapping should be performed from the **Directory Service** panel.

- **Name**: The directory group name.
- **Role**: The role assigned to the group. Values can be array_admin, storage_admin, ops_admin, or read-only (see *Roles* for additional information).
- **Group**: The common name of the directory group.
- **Group Base**: The configured group in the directory tree comprised of the OU or CN combined with the base DN attribute.

Note: FlashBlade uses IdP cookies for SSO, and clearing browser cookies/redoing the IdP login may be necessary when IdP authorization settings change.

Configuring SSO Authentication for Object Service

OIDC and SAML2 identity Providers can be configured on the FlashBlade for SSO for object service.

Configuration for Object Service Prerequisites

Before attempting OIDC or SAML SSO configuration for object service, you must request a metadata document from the identity provider (IdP) that contains the following information.

For OIDC configuration, you must have the following:

- CA and CA Group Certificates (not needed if the IdP server certificates are signed by a well known CA).
- The IdP URL, which specifies the URL of the identity provider.

For SAML configuration, you must have the following:

- The identity provider (IdP) entity ID, which specifies the globally unique name for the identity provider.
- The IdP URL, which specifies the URL of the identity provider.
- The IdP verification certificate.

OR

- The IdP metadata URL
- (Optional) CA certificate or CA Certificate Group.

The IdP metadata file contains this information.

Note: The IdP metadata URL is sufficient without the IdP verification certificate.

Configuration for Object Service Overview

The following list summarizes the steps to configure and enable ODIC or SAML2 SSO for object service authentication on a FlashBlade. Configurations are required in both the service provider side (Purity//FB) and on the identity provider side, in order to complete the SSO configuration.

1. Create and OIDC or SAML2 identity provider on the FlashBlade.
2. Create a role and attach permissions to the role using access policies on the FlashBlade.
3. Configure the trust policy for the role to authenticate the IdP.

Configuring an OIDC Identity Provider for Object Service on the FlashBlade

To configure an OIDC identity provider for object service on the FlashBlade array, perform the following:

1. From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
2. Click the Add (+) button in the header of the **OIDC Identity Providers** panel. The **Create ODIC Identity Provider** pop-up window appears.
3. In the **Name** field, enter the OIDC identity provider name.
4. Click **Enabled** to enable the configuration after creation.
5. The **Service** and **Binding** fields are automatically populated with **object** and **none**, respectively, and cannot be changed.
6. In the **Provider URL** field, enter the identity provider URL.
7. (Optional) In the **CA Certificate** field, select the CA certificate that will be used to validate the identity provider's authenticity.
8. (Optional) In the **CA Certificate Group** field, select the group of CA certificates that will be used to validate the identity provider's authenticity.
9. Click **Create**.

Configuring a SAML2 Identity Provider for Object Service on the FlashBlade

To configure SAML2 identity provider for object service on the FlashBlade array, perform the following:

1. From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
2. Click the Add (+) button in the header of the **SAML Identity Providers** panel. The **Create SAML Identity Provider** pop-up window appears.
3. In the **Name** field, enter the SAML identity provider name.
4. Click **Enabled** to enable the configuration after creation.

5 From the **Service** list, select **object**. The **Binding** field is automatically populated with **none**.

6 In the **Array URL** field, enter the FlashBlade array's URL.

The **SP Entity ID**, **SP Assertion Consumer URL**, and **SP Metadata URL** fields are populated based on the information entered in the **Name** and **Array URL** fields.

7 (Optional) In the **Decryption Credential** field, select the credential from the drop-down list. Note that only array-type certificates are displayed.

Note: Enter either the IdP Metadata URL by itself, or enter the IdP Metadata URL, the optional IdP Metadata CA certificate, and the optional IdP Metadata CA certificate group.

8 In the **IdP Entity ID** field, enter the globally unique name for the IdP. Required if the IdP metadata URL is not provided.

9 In the **IdP URL** field, enter the IdP's URL. Required if the IdP metadata URL is not provided.

10 In the **IdP Metadata URL** field, enter the URL to the IdP metadata file. .

11 (Optional) In the **IdP Metadata CA Certificate** field, select the CA certificate to validate the authenticity of the IdP metadata file.

12 (Optional) In the **IdP Metadata CA Certificate Group** field, select the group of CA of certificates to validate the authenticity of the IdP metadata file.

13 (Optional) Enable **Encryption Assertion**.

14 In the **IdP Verification Certificate** field, select the certificate used by the SP to verify the signed SAML response from the IdP from the drop-down list. This is required if the IdP metadata URL is not provided.

15 Click **Create**.

Configuring an IdP for SSO using OIDC or SAML

Configuration of an IdP of SSO differs depending on the IdP. The following are general overviews of what is required for configuring an IdP for SSO using OIDC and SAML.

For OIDC, determine the Identity Provider's public URL from documentation or IdP administrative console. When appended with **/.well-known/openid-configuration**, you should see the OpenID Connect Discovery specification-compliant metadata.

For SAML:

- You must access your IdP's administrative console and navigate to the SAML configuration section to create a new SAML application within your IdP.

- Input details such as the SP's entity ID, ACS URL.
- Define which user attributes from your IdP should be sent to the SP.
- Metadata files must be downloaded and then uploaded on the SP for integration. Alternatively, the metadata URL can be passed to the SP.

Creating a Role

Once an OIDC or SAML2 identity provider has been created, you must create a role with access policies. A maximum of 10,000 roles are allowed per FlashBlade.

When creating a role, you can define the allowed session duration and add one or more access policies, which will grant the S3 permissions in the account.

To create a role, perform the following:

1. From the **Storage** > **Object Store** page, click the account name in the **Accounts** panel under which you wish to create a role. The details page for the account appears.
2. Click the Add (+) button in the header of the **Roles** panel. The **Step 1: Create Role** pop-up window appears.
3. In the **Role Name** field, enter a name for the role.
4. In the **Maximum Session Duration** field, enter the maximum session duration from 1 hour to 48 hours. If no value is entered, the default is 1 hour.
5. Click **Create**. The **Step 2: Add Access Policies to role <name>** appears.
6. From the **Available Policies** panel, select the access policies you wish to assign to the role. At least one policy must be added to grant the role permissions in the account.

   If a specific access policy does not exist in the list of available policies, click **Create a new Access Policy** (see *Creating Object Store Access Policies* for instructions).
7. Click **Add**.

Creating Trust Policy Rules

After creating a role, create a trust policy rule.

To create a trust policy rule, perform the following:

1. From the **Storage** > **Object Store** page, click the account name in the **Accounts** panel under which you created a role. The details page for the account appears.
2. Click the created role from the **Roles** panel. The details page for the role appears.
3. Click the Add (+) button in the header of the **Trust Policy Rules** panel. The **Create Trust Policy Rule** pop-up window appears.
4. In the **Rule Name** field, enter a name for the rule.
5. Select whether the action (**Effect**) of the rule will be to **Allow** or **Deny**. Note that explicit Deny rules take precedence over Allow rules.
6. Click the Add (+) button in the **Principals** field. The **Add Principals** pop-up window appears.
   a. Select the principals you wish to add from the **Available Principals** panel.
   b. Click **Add**.
7. Click the Add (+) button in the **Actions** field. The **Add Actions** pop-up window appears.
   a. Select the actions you wish to add from the **Available Actions** panel.
   b. Click **Add**.
8. (Optional) If you wish to add a condition to the rule, click **Add condition**. The condition key must have a format prefix:value. Supported key prefixes are: aws, jwt, saml, sts.
   a. In the Key field, enter the key prefix.
   b. Select **All** or **Any**, the operator, and select **If Exists** under the **Operator** field.
   c. Click the Add (+) button under the **Values** field. Enter a comma separated value string and then click **OK**.
9. Click **Create**.

Uploading a Trust Policy in IAM Format

To upload a trust policy in IAM format, perform the following:

1. From the **Storage** > **Object Store** page, click the account name in the **Accounts** panel with the role that you wish to upload a trust policy in IAM format. The details page for the account appears.
2. Click the role from the **Roles** panel. The details page for the role appears.
3. Click **More Options** > **Upload Trust Policy in IAM format** in the header of the **Trust Policy Rules** panel. The **Upload Trust Policy** pop-up window appears.

4 In the **Trust Policy** field, copy and paste the trust policy in IAM format. You can optionally load the policy from a file.

5 If you did not manually past the trust policy in IAM format in the previous field, you can load the policy from a file. In the **Load from file** field, click **Choose File** and select the file containing the trust policy in IAM format to upload.

6 Click **Upload**.

Configuring SSO Authentication with SAML2 SSO for Management Service

⚠ **Important!** Do not enable SSO Authentication with SAML2 SSO for management service before testing the configuration. If the configuration is enabled with errors, the user will be locked out. Go to *Testing the SAML2 SSO Configuration* section to perform the test.

SAML2 identity Providers can be configured on the FlashBlade for SSO for management service.

SAML2 SSO Configuration for Management Service Prerequisites

Before attempting SAML SSO configuration for management service, note the following prerequisites:

- The SAML2 SSO configuration for management service steps require either administrator access to the identity provider's management tool, or, for example, coordination with an AD FS administrator to create a relying party trust for the array and to provide the following information:
  - The identity provider (IdP) entity ID, which specifies the globally unique name for the identity provider.
  - The IdP URL, which specifies the URL of the identity provider.
  - The IdP verification certificate.

  The IdP metadata file contains this information.

  Note: The IdP metadata URL is sufficient without the IdP verification certificate.
- The directory service used with the array must be the same directory service instance as used by the identity provider in the relying party trust configuration for this array. Purity//FB does not support multiple or federated directory services.
- It is common for identity providers to use TLS 1.2 and above for securing communications and enabling services such as SP monitoring in AD FS.

⚠ **Important!** It is highly recommended that before configuring SSO, create a strong password for the **pureuser** and other array administrator accounts, and to save those passwords according to your organization's security policies.

Additionally, note that configurations are required on both the service provider side (Purity//FB) and on the identity provider (IdP) side to complete the SSO configuration.

SAML2 SSO Configuration for Management Service Overview

The following list summarizes the steps to configure and enable SAML2 SSO for management service authentication on a FlashBlade. Configurations are required in both the service provider side (Purity//FB) and on the identity provider side, in order to complete the SSO configuration.

1. Configure SAML2 SSO on the FlashBlade.
   a. Obtain IdP information from AD FS, Okta, or an administrator.
   b. Configure the service provider (SP) with the array and IdP information.
2. In the identity provider, set up SSO using the SP information from the FlashBlade.
3. On the FlashBlade, enable SAML2 SSO Authentication.

For information about Directory Services and Directory Services configuration instructions, see *Directory Services Overview* and its related topics.

Use the AD FS Management Tool on AD FS to configure the AD FS IdP for SSO and, optionally, multi-factor authentication (MFA) with the FlashBlade.

Configuration Notes

- The verification certificate (the AD FS primary token-signing certificate) must be an X.509 certificate in PEM format.
- The array name portion of the URL in the browser used to configure the SP must be consistent with the URL entered into the **Array URL** field in the SAML2 SSO panel. Whether the URLs are based on a FQDN (such as pure01.- mycompany.com) or on a hostname (such as pure01), the browser URL and the **Array URL** field configured in the SAML2 SSO panel must be consistent in the way the array name is specified. When there is a mismatch, the browser used to configure SAML SSO will not be able to complete the redirect and authentication process.

Group to Role Mapping

Group to Role Mapping on Purity//FB

Group to role mapping on Purity//FB is only certified with the AD FS server. To configure the group to role mapping, see *Adding Role Mapping* .

To configure the identity provider server to send the user group in the SAML response, refer to your identity provider's documentation.

Group to Role Mapping on IdPs

Note: Directory service configuration is not necessary for role mapping on IdPs.

Group to role mapping on IdPs is certified with AD FS and Okta. Each IdP has unique configuration steps to map the user group(s) to the Purity//FB roles (array_ admin, storage_admin, ops_admin, and readonly). The specified role is then sent in the SAML response with attribute name purity_roles. Contact Pure Storage Technical Services to create an SSO integration application on the IdP side and to configure group to role mapping for the specified IdP.

Configuring a SAML2 Identity Provider for Management Service on the FlashBlade Array

Note: It is not recommended to enter the service provider (SP) signing and decryption credentials or enable the identity provider (IdP) request or assertion until after the initial configuration passes the end-to-end test.

To configure SAML2 identity provider for management service on the FlashBlade array, perform the following:

1. From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
2. Click the Add (+) button in the header of the **SAML Identity Providers** panel. The **Create SAML Identity Provider** pop-up window appears.
3. In the **Name** field, enter the SAML identity provider name.
4. From the **Service** list, select **management**. The **Binding** field is automatically populated with **http-direct**.
5. In the **Array URL** field, enter the FlashBlade array's URL. The URL must use HTTPS.

The **SP Entity ID**, **SP Assertion Consumer URL**, and **SP Metadata URL** fields are populated based on the information entered in the **Name** and **Array URL** fields. The information in these fields is required to create a relaying party for the FlashBlade array in the AD FS identity provider.

6 (Optional) Enable **Sign Request**.

7 (Optional) In the **Signing Credential** field, select the credential from the drop-down list. Note that only array-type certificates are displayed.

8 (Optional) In the **Decryption Credential** field, select the credential from the drop-down list. Note that only array-type certificates are displayed.

Note: Enter either the IdP Metadata URL by itself, or enter the Idp Entity ID and IdP URL fields and the IdP Verification Certificate. If both the IdP Metadata URL and any of the Idp Entity ID, IdP URL, or IdP Verification Certificate are provided, the metadata file is read only for those options that are not provided.

9 In the **IdP Entity ID** field, enter the globally unique name for the IdP. For AD FS, this name is found under the AD FS Management Tool under `AD FS/Service/Federated Service Properties`.

10 In the **IdP URL** field, enter the IdP's URL. This value is found under the AD FS Management Tool under `AD FS/Service/Endpoints`.

11 In the **IdP Metadata URL** field, enter the URL to the IdP metadata file. This value is found under the AD FS Management Tool under `AD FS/Service/Endpoints`.

12 (Optional) Enable **Encryption Assertion**.

13 In the **IdP Verification Certificate** field, select the certificate used by the SP to verify the signed SAML response from the IdP from the drop-down list. For example, this value is found under the AD FS Management Tool under `AD FS/Service/Certificates/Token-signing`.

14 Click **Create**.

Configuring the IdP

Note: FlashBlade uses IdP cookies for SSO, and clearing browser cookies/redoing the IdP login may be necessary when IdP authorization settings change.

Once SAML2 SSO configuration is complete on the FlashBlade, you must register Purity//FB as a relying party trust on your IdP. Administrator-level access is required to access your IdP's management tool.

Configuration steps will be different depending on the IdP. As an example, the following steps describe how to register Purity//FB as a relying party trust on AD FS.

To register Purity//FB as a relying party trust on AD FS, perform the following:

1 On the machine running AD FS, click **Control Panel** > **System and Security** > **Administrative Tools** and select **AD FS Management.**

2 In the left panel, select **Relying Party Trusts**. In the right panel, select **Add Relying Part Trust**. The **Add Relying Part Trust Wizard** appears.

Note: The Access Control Policy **Permit everyone** corresponds to password authentication. Multi-factor authentication can be configured later if required. Consider waiting until after the SSO configuration passes with password authentication before enabling multi-factor authentication.

Table 9.12 Sample Configuration in the Add Relying Party Trust Wizard

Wizard Page	Action
Welcome	Ensure **Claims aware** is selected.
Select Data Source	Select **Enter data about the relying party manually**.
Specify Display Name	Enter a display name for the new relying party (for convenience, this name could match the SP configuration display name, as set in the FlashBlade).
Configure Certificate	No action.
Configure URL	Select **Enable support for the SAML 2.0 Web SSO protocol.** In the **Relying party SAML 2.0 SSO service URL** field, enter the Assertion Consumer URL from the Purity SAML2 SSO pane.
Configure Identifiers	In the **Relying party trust identifier** field, enter the SP entity ID from the Purity SAML2 SSO pane.
Choose Access Control Policy	Select **Permit everyone**.
Ready to Add Trust	No action.
Finish	Select **Configure claims insurance policy for this application.**

3 Once the new relying party appears in the Relying Party Trusts table, select the relying party and in the right panel select **Edit Claims Insurance Policy**. The **Edit Claims Insurance Policy** page opens.

4 Click **Add Rule** at the bottom right of the page. These rules specify the content the IdP returns in an assertion to the SP. Two rules are required: one for a unique identity for the user, the second for group information. The **Add Transform Claim Rule Wizard** opens.

5 In the **Claim rule template** field on the **Select Rule Template** page, select S**end LDAP Attributes as Claims** and then click **Next**.

6 Click **Choose Rule Type** and follow the instructions as described in the following table:

Table 9.13 Actions in the Add Transform Claim Rule Wizard for Users

Field	Action
Claim rule name	Enter a descriptive name for the rule, such as **map name id**.
Attribute store	Select **Active Directory**.
Mapping table LDAP Attribute column	Select **SAM-Account-Name**.
Outgoing Claim Type column	Select **Name ID** and then click **Finish**.

The new rule for name mapping appears in the **Edit Claims Insurance Policy** page **Insurance transform Rules** table.

7 Click Add Rule for the group information rule. The Add Transform Claim Rule Wizard appears.

8 In the **Select Rule Template** page, select **Send LDAP Attributes as Claims** from the **Claim rule template** field.

Table 9.14 Actions in the Add Transform Claim Rule Wizard for Groups

Field	Action
Claim rule name	Enter a descriptive name for the rule, such as **pass group info**.
Attribute store	Select **Active Directory**.
Mapping table LDAP Attribute column	Select **Is-Member-Of-DL**.
Outgoing Claim Type column	Select **Group** and then click **Finish**.

The new rule for group mapping appears in the **Edit Claims Insurance Policy** page **Insurance Transform Rules** table.

Enabling Multi-factor Authentication

⚠ **Important!** Whenever you modify the SAML2 SSO for management service configuration, it is recommended to re-test to verify the configuration information before enabling.

Note: The types of multi-factor authentication (MFA) available depend on the IdP. The AD FS IdP supports certificate authentication, authentication with Microsoft® Azure™ software, and others.

You can optionally enable MFA as described in the following:

1 In the AD FS **Management tool Relying Party Trust** page, select the relying party you created. In the right panel, select **Edit Access Control Policy**.

2 On the **Choose an access control policy** page, select **Permit everyone and require MFA** and then click **Next**.

3 The next steps, such as configuring a certificate, depends on the type of MFA selected.

Testing the SAML2 SSO for Management Service Configuration

⚠ **Important!** Whenever you modify the SAML2 SSO for management service configuration, you must re-test to verify the configuration information before enabling.

Note: Refer to the *SAML2 SSO Login Error Messages* table to assist with troubleshooting end-to-end test error messages. The table also includes actionable tasks to correct the error messages.

Once you have your SAML2 SSO for management service configuration set up, perform the basic test and end-to-end (E2E) test to verify all the data is valid. This test can only be administered by the array_admin role. If SAML2 SSO for management service is not configured then the test option will not be enabled. Use the basic test and E2E test to validate the configuration to be applied before enabling SSO. If you encounter issues, use the test result page with the SAML response payload from the IdP to correct any configuration errors.

To test the SAML2 SSO for management service configuration, perform the following:

1 From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.

2 Click the name of the SAML identity provider configured for SSO for management service you wish to test from the **SAML Identity Providers** panel. The details page for the configuration appears.

3 Click **Test** in the header of the **SAML Identity Provider** panel. The **Test SAML2 SSO Configuration** pop-up appears.

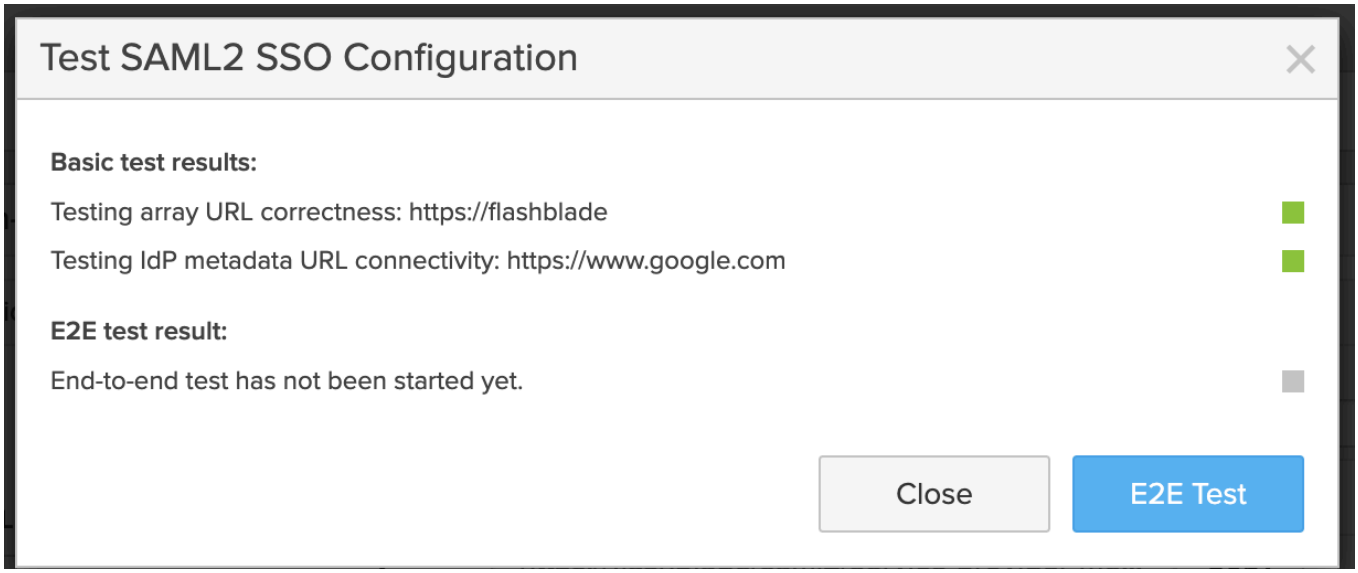

The IdP metadata URL in the example is only a demonstration that a non-viable site may pass the basic test because it only checks for connectivity and that the site is trusted and supports downloading of the page's content. Use the end-to-end test results to validate the IdP metadata URL.

If the basic test fails, review the error failure messages, close the window and edit your configuration then perform another test.

If you selected to enable your SAML2 SSO for management configuration and click **Save**,without testing, a **Test Reminder** pop-up will appear. It is not recommended to click **Proceed to Save** without testing. You should always test new configurations and modified configurations.

4 Click **E2E Test** when your basic test results have passed.

5 The ADFS (Active Directory Federation Services) window will appear. Enter your credentials and log in to ADFS and the E2E test will be initiated, as indicated by the **Test SAML2 SSO Configuration** pop-up with the **Check E2E Test Result** button.

6 Click the **Check E2E Test Result** button to review the test results when pop-up indicates the test is successful.

The E2E test results may not be ready immediately. If you click the test results before the E2E test completes, you will be notified that the test has not completed.

The SAML2 SSO for management configuration may be incorrect if the test does not complete or is stopped before completion. If the test takes more time than expected, you can close the test and restart it by reopening the **Test SAML SSO Configuration** dialogue box and clicking the **E2E Test** button.

7 Review the E2E test results. If your configuration failed, review the error messages and use the *SAML2 SSO Login Error Messages* table to troubleshoot your configuration.

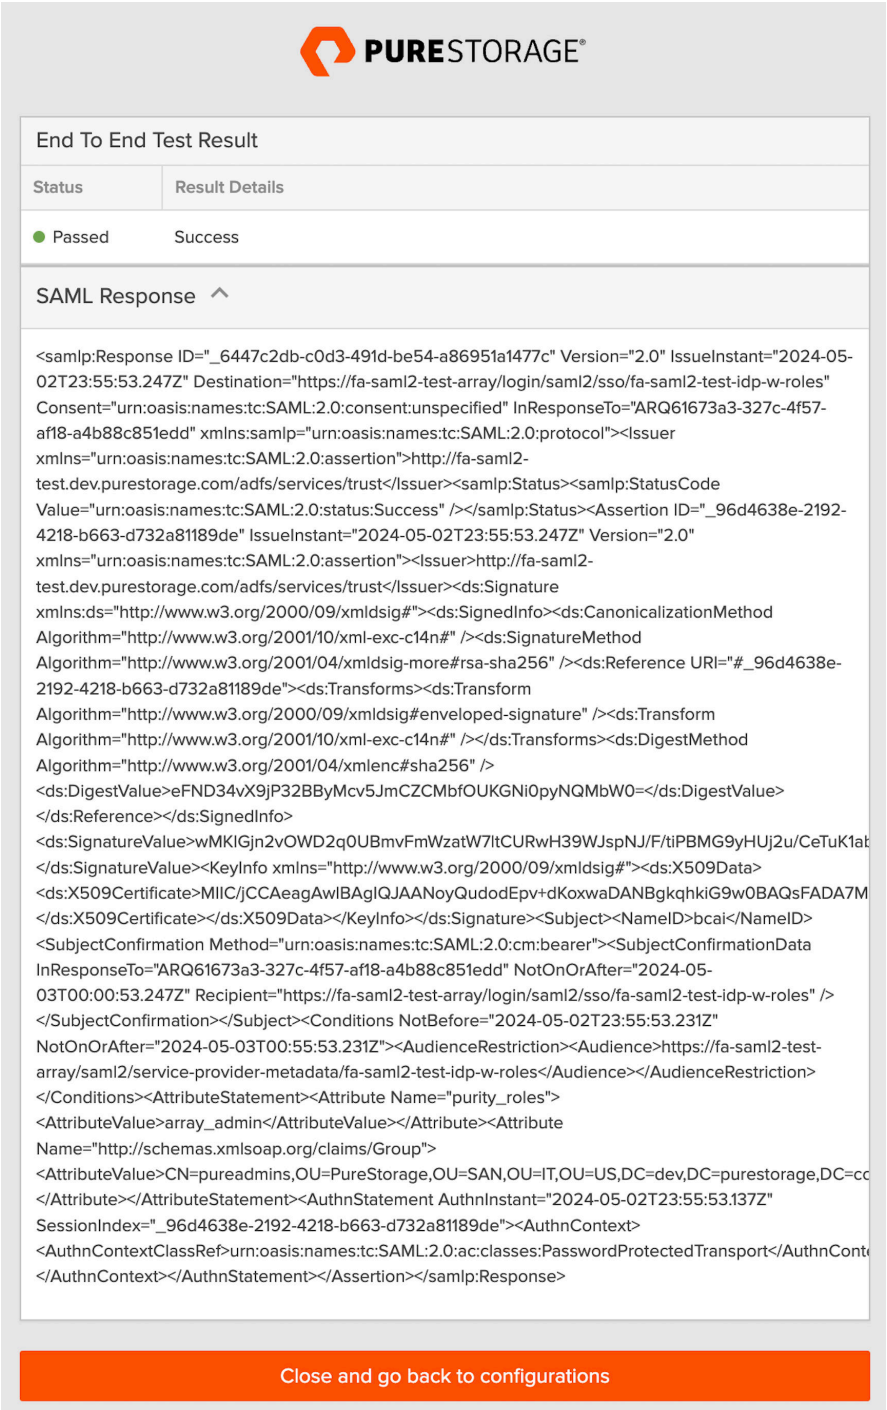

8 Click the **Close and go back to configurations** button and when ready, enable SAML2 SSO for management service.

Enabling SAML2 SSO for Management Service Authentication

⚠ **Important!** Whenever you modify the SAML2 SSO for management service configuration, it is recommended to re-test to verify the configuration information before enabling.

After the SAML2 SSO for management service configuration is enabled, SAML users are referred to the identity provider for authentication for all new login attempts to the Purity//FB management GUI. Existing user sessions are not affected.

To enable SAML2 SSO for management service authentication, perform the following:

1. (Optional) Log into Purity//FB in a second browser window if login access is interrupted.
2. From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
3. Click the Edit button in the same row as the SAML identity provider for management service in the **SAML Identity Providers** panel you wish to enable. The **Edit SAML Identity Provider** pop-up window appears.
4. Click **Enabled**.
5. Click **Save**.

   The **SAML Identity Providers** panel displays **True** under the **Enabled** column for the SAML identity provider for managmenet service.

   Once enabled, the Purity//FB login screen no longer prompts for a password. Upon a user's first log in, and also after an SSO session expires, the AD FS log in page will open and prompt for the user's AD FS credentials.

Other SAML2 SSO for Management Service Configuration Steps

Enabling IdP Sign Request and Encryption Assertion

⚠ **Important!** Once the sign request and encryption assertion are enabled, the IdP also needs to have matching certificates imported, if the SP metadata url is not monitored by IdP.

If not enabled during initial SAML2 SSO configuration on the FlashBlade, you can enable IdP Sign Request and Encryption Assertion.

1. From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
2. Click the **More Options** button in the header of the **Single Sign-On** panel and select **Edit**. The **Edit SAMLs SSO Configuration** pop-up window appears.

3. Enable **IdP Sign Request** and select the signing credential in the **SP Signing Credential** field.
4. Enable **IdP Encryption Assertion** and select the decryption credential in the **SP Decryption Credential** field.

SSO Session Timeout

When SAML2 SSO is enabled, the Purity//FB management GUI session timeouts are based on the idle timeout configured locally on the FlashBlade. After the GUI session has logged out (due to idle timeout), users can attempt to login to the GUI via SSO. When the IdP session is active, users are prompted to re-authenticate. For example, by default an AD FS SSO session times out after eight hours.

Refer to the IdP documentation for information on customizing the time out setting.

TLS 1.2 or 1.3 Support

The FlashBlade requires TLS 1.2 or 1.3 with strong authentication. The AD FS instance must also support TLS 1.2 or 1.3 with strong authentication in order to test the relying party's federation metadata URL in an AD FS instance. If TLS 1.2 or 1.3 with strong authentication is not configured, an error is seen when clicking the Test URL button when configuring monitoring in the Replying Party Configuration GUI in the AD FS instance.

Refer to Microsoft documentation for instructions on how to enable TLS 1.2/1.3 and strong authentication in AD FS.

SAML2 SSO for Management Service Runtime Notes

- If a user is deactivated in AD FS but is currently logged in, the current login session is not affected. The user is denied access at the next login attempt.
- In case the SAML2 SSO service is temporarily unavailable, users with array-admin role permissions can access the array through the Local Access link on the login page. This link provides emergency administrator access when GUI logins are unavailable.
- SSO session timeouts depend on the IdP. For example, by default AD FS times out after eight hours and Okta times out after 2 hours. A different time out length can be configured in the IdP. For additional information, refer to *SSO Session Timeout* in *Other SAML2 SSO for Management Service Configuration Steps*.

- When a user logs in through SAML2 SSO, the browser URL field reports its URL based on what is configured in the **Array URL** field, regardless of whether the user entered the array hostname or FQDN when directing the browser to the array.

Understanding SAML2 SSO Login Error Messages

The following table describes error messages that can be returned if login errors occur.

Note: We recommended checking the raw SAML response from the IdP to further understand the error details for errors showing up during the "verification of IdP response" stage. The SAML response XML text is available from the E2E test results page.

Table 9.15 SAML SSO Login Error Messages

Error Message	Cause	Stage	Actions
SSO is not available.	SSO is not enabled or failed to load to the GUI server. The error is displayed when a reserved path is accessed before authentication.	SSO setup.	Enable SSO in the **Settings** page.
SAML login is not supported for user xyz.	The remote user is trying to login through SAML with same name as a reserved internal user name space, such as ir, root.	Verification of IdP response.	The user is not allowed for SAML.
No user role or group information provided.	No role nor groups information are included in the identity provider's assertion response.	Verification of IdP response.	Check the group information from the identity provider, also use the management DS test function to ensure the group has a mapped role in the FlashBlade.
No valid user role information found.	The role information is provided, but is not a valid role defined in the FlashBlade.	Verification of IdP response.	Check the "purity_roles" definition from the identity provider side.
SSO login fails to get enabled because the identity provider's metadata is not accessible.	The IdP's metadata URL is provided, but is not accessible.	SSO setup.	Check and update the metadata URL of identity provider on the FlashBlade.
Necessary service provider(SP) cookie is not included in the request. Please ensure that cookies are enabled and verify that the accessed URL corresponds to the expected SP's consumer URL.	The SP site cookie is not sent back along with the IdP response, so SP cannot verify which request the response is associated with.	Verification of IdP response.	Check the cookie settings on the browser, also make sure the host name in the login page matches the array URL setting for SSO.
Incomplete SAML response: No subject found in the assertion. Please check the claim setups from the identity provider.	The configuration of the replying party on IdP does not add data such as user name and group/ role information to the assertion.	Verification of IdP response.	Ensure the claim setting on the IdP side has added the user name inside the assertion claim.

Table 9.15 SAML SSO Login Error Messages (continued)

Error Message	Cause	Stage	Actions
Failed to decrypt the data received. Check the encryption key in the SAML2 configuration.	The encryption key may not match between SP and IdP.	Verification of IdP response.	Check the encryption credential's expiration and ensure they match with the IdP record.
Assertion signature verification failed. Check the verification certificate in the SAML2 configuration.	The signature verification failed, possibly due to an incorrect verification cert for the SSO configuration.	Verification of IdP response.	Check the verification credential's expiration and ensure they match with the IdP record.
Failed to verify the assertion. Check the verification certificate in the SAML2 configuration.	The general error when assertion is invalid.	Verification of IdP response.	Check the verification certificate.
Failed to authenticate via SSO	Generic error when no explicit error is found.	Verification of IdP response.	n/a
Issuer from the SAML response doesn't match the record. Please check the entity ID details of the SSO config.	Occurs when the issuer of the IdP response does not match the record from the SP side.	Verification of IdP response.	Check and update the SP configuration.
Invalid status [urn:oasis:names:tc:SAML:2.0:status:Responder] for SAML response	Generic error from the IdP. Indicates something is wrong with the IdP configuration.	Verification of IdP response.	Check if the certificate for verifying SAML requests is properly configured when the FlashBlade includes signatures in the SAML request. Check if the rules in the response claim is properly configured.

Users

The **Users** panel manages local and remote FlashBlade array users.

Note: Remote users are listed in the **Users** panel only after API tokens have been explicitly configured for them.

To access the **Users** panel, from the **Settings** page navigation sidebar, click **Users** under **Security**.

Figure 9.26 Users - Users Panel

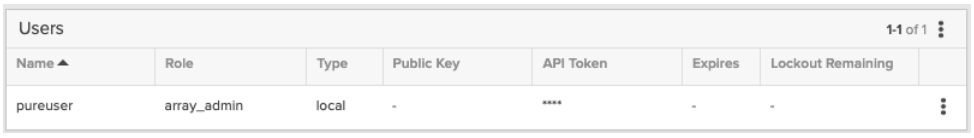

The panel displays the following:

- **Name**: The user name.
- **Role**: The role assigned to the user. Values can be array_admin, storage_admin, ops_admin, or read-only.
- **Type**: Displays `local` for users local to the array (pureuser), or `remote` for LDAP users.
- **Public Key**: The user's public key, if any, displayed as ****.
- **API Token**: Active API token, if any, displayed as ****.
- **Expires**: Expiration time, if any, of the API token.

In addition to displaying user attribute details, you can also create and manage local users, user API tokens, and user array authentication in the **Users** panel.

Local Users

Users can be added by configuring local users directly on the array. Alternatively, users can be added via Lightweight Directory Access Protocol (LDAP) by integrating the array with a directory service such as Active Directory or OpenLDAP (refer to the Directory Services Overview for additional information about integrating the array with a directory service).

On the array (locally), users can only be created by array administrators. Local users must have unique names and cannot be the same names as LDAP users. If an LDAP user appears with the same name as a local user, the local users will have priority. Up to 100 local users can be created.

Role-based access control (RBAC) restricts system access and capabilities to each user based on their assigned role in the array. All users are assigned a role on the array, as described in the following table.

Role	Description
readonly	Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array.
storage_admin	Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
array_admin	Array Admin users can perform all FlashBlade operations.

Role	Description
ops_admin	Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.

For local users, the role is set during user creation. For LDAP users, the role is set by configuring groups in the directory that correspond to the FlashBlade user roles. Each local user account on the array is password protected. The password is assigned during user creation and can be modified by array administrators. All local users can manage their own passwords, but only array administrators can manage the passwords of other users. Changing a local user's password requires knowledge of the current password. If the password of a local user is unknown, delete the account and recreate it with the desired password. Note that deleting a local user's account means deleting any public key associated with the user. If the password of the pureuser account is unknown, contact Pure Storage Technical Services to reset the account to the default pureuser password. Passwords of LDAP users are managed in the directory service.

Public keys and API tokens that have been created for users appear masked in the **Users** panel.

API Tokens

The REST API requires permission to run commands. To access the server to run REST API commands, an authentication API token must be generated for the user. Each API token is unique and has an optional expiration setting. When a token expires, it can no longer be used and a new token must be generated to continue running REST API commands. It is a best practice to set an expiration on a token for the duration of time it will be used.

Array Authentication

Note: Setting other users' public keys is available to administrators and users with admin permissions only. The public key should be in the OpenSSH format, which begins with `ssh-rsa`. A maximum of 100 users may have a public key set.

Administrators and users with administrator privileges can set passwords for array users. Additionally, public key authentication can be configured allowing users to SSH to the array.

Creating a Local User

To create a local user, perform the following:

1. From the **Settings** page navigation sidebar, click **Users** under **Security**.
2. Click the **More Options** button in the header of the **Users** panel and select **Create User**. The **Create Local User** pop-up window appears.
3. In the **User** field, enter the user name.
4. In the **Role** field, select the user's role.
   - **Read-Only**: Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array.
   - **Ops Admin**: Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.
   - **Storage Admin**: Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
   - **Array Admin**: Array Admin users can perform all FlashBlade operations.
5. In the **Password** field, enter the user's password.
6. In the **Confirm Password** field, re-enter the password.
7. Click **Create**.

Editing a Local User

To edit a local user, perform the following:

1. From the **Settings** page navigation sidebar, click **Users** under **Security**.
2. From the **Users** panel, click the **More Options** button in the the same row as the user you wish to edit and select **Edit User**. The **Edit User** pop-up window appears.
3. (Optional) In the **Role** field, select the user's role.
   - **Read-Only**: Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array.

- **Ops Admin**: Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.
- **Storage Admin**: Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
- **Array Admin**: Array Admin users can perform all FlashBlade operations.

4 (Optional) In the **Current Password** field, enter the user's password.

5 (Optional) In the **New Password** field, enter the user's new password.

6 (Optional) In the **Confirm Password** field, reenter the new password.

7 Click **Save**.

Deleting a Local User

Local users can be deleted at any time. Once deleted, the user will no longer be able to access the FlashBlade array. Deleted users cannot be restored.

To delete a local user, perform the following:

1 From the **Settings** page navigation sidebar, click **Users** under **Security**.

2 From the **Users** panel, click the **More Options** button in the the same row as the user you wish to delete and select **Delete User**. The **Delete Local User** pop-up window appears.

3 Click **Delete**.

Creating an API Token

API tokens allow users to run REST API commands. To access the server to run REST API commands, you must generate an authentication API token for the user. As a best practice set an expiration on a token for the duration of time it will be used.

To create an API token for a user, perform the following:

1 From the **Settings** page navigation sidebar, click **Users** under **Security**.

2 Click the **More Options** button in the header of the **Users** panel and select **Create API Token**. The **Create API Token** pop-up window appears.

3 Enter the user for whom the token will be created.

4 Enter the token's expiration duration in the format N[m/h/d/w]. For example, **15m**, **1h**, **2d**, **3w**. If not set, the API token will not expire.

5 Click **Create**.

The API Token pop-up window appears displaying the token and its creation and expiration dates.

Updating a User's Public Key

A public key can be added or updated for user authentication in order to allow users to SSH to the FlashBlade.

Note: Setting other users' public keys is available to administrators and users with admin permissions only. The public key should be in the OpenSSH format, which begins with `ssh-rsa`. A maximum of 100 users may have a public key set.

To add or update a user's public key, perform the following:

1 From the **Settings** page navigation sidebar, click **Users** under **Security**.

2 Click the **More Options** button in the header of the **API Clients** panel and select **Update Public Key**. The **Update Public Key** pop-up window appears.

3 Enter the user for whom you are adding a public key.

4 Enter the public key into the **Public Key** field.

5 Click **Update**.

Displaying an API Token

If you previously created an API token you can later display and copy the token for use with the REST API.

To display a user's API token for a user, perform the following:

1 From the **Settings** page navigation sidebar, click **Users** under **Security**.

2 In the **Users** panel, locate the user whose API token you wish to display. Click the **More Options** button in the same row, and select **Show API Token**. The **API Token** pop-up window appears.

3 Click **Copy** to copy the token if needed.

4 Click **Close**.

Recreating an API Token

API tokens can be recreated. Note that recreating an API token invalidates the old token and the recreated token must be authenticated.

To recreate an API token, perform the following:

1. From the **Settings** page navigation sidebar, click **Users** under **Security**.
2. In the **Users** panel, locate the user whose API token you wish to recreate. Click the **More Options** button in the same row, and select **Recreate API Token**. The **Recreate API Token** pop-up window appears.
3. Click **Recreate**.

Removing an API Token

To remove an API token from a user, perform the following:

1. From the **Settings** page navigation sidebar, click **Users** under **Security**.
2. In the **Users** panel, locate the user whose API token you wish to remove. Click the **More Options** button in the same row, and select **Remove API Token**. The **Remove API Token** pop-up window appears.
3. Click **Remove**.

Chapter 10
Choosing an Active Directory Account or Directory Services Configuration

FlashBlade supports both Active Directory (AD) account and directory services configuration. The following describes key differences between the two configuration types, as well as usage scenarios for each. Choose the configuration that provides the functionality you need and that best fits your environment.

⚠ **Important!** FlashBlade systems have to be joined to Active Directory for Native SMB. If NTLMv2 is required then you must enable arcfour-hmac encryption type apart from default AES256-sha1 during the join process or it can be enabled by editing the Active Directory account if the FlashBlade has already joined to the Active Directory.

Note: For additional information about Active Directory and directory services, see *Active Directory Overview* and *Directory Services Overview*.

Key Differences between Active Directory and Directory Services Configuration

Active Directory Accounts	Directory Services
Recommended for Native SMB. The computer account name or service principal name (SPN) should resolve to the data VIP. **Note:** A computer account should be associated with and resolved by DNS to a specific data VIP on the FlashBlade. **Note:** Refer to *Service Principal Names* for details on configuring AD accounts and DNS name registration. **Note:** Pure Technical Services is required to enable native SMB.	
Only AD can be used as the directory services for SMB and NFS.	AD can be used as a directory service for SMB, and OpenLDAP or OUD or NIS can be used for NFS services

Active Directory Accounts	Directory Services
AD account creation configures Kerberos authentication with AD and creates the keytabs automatically. Service principal names (SPNs) and encryption types for Kerberos are managed on the FlashBlade.	Keytab files must be uploaded separately and Kerberos configurations are managed off the FlashBlade in the AD domain.
AD account creation can be used as a single workflow to configure AD Kerberos authentication for NFS and SMB.	NFS and SMB directory services is authenticated separately. You must upload keytab files separately to use Kerberos authentication for NFS.
FlashBlade communicates with the AD domain using the created AD account and authenticates via Kerberos. No machine account passwords are stored locally on the FlashBlade for security.	FlashBlade communicates with the AD/LDAP domain as the configured bind user and authenticates via simple authentication unless authentication is configured for NFS services. The specified bind password is stored locally on FlashBlade. Note that anonymous binding can be used if supported by your server.
FlashBlade uses only LDAPS for secure communication with the AD domain.	Plaintext LDAP can be configured or LDAPS can be used.

Configuration Choice Examples Based on Environment

Scenario	Configuration	Additional Information
Domain environment exclusively supports unsecure communication	Directory Services	• AD account only supports secure communication • Use directory service with plaintext LDAP • Upload keytabs if your NFS clients need Kerberos
You want to use Native SMB with AD Kerberos	AD Account	• AD account supports secure communications. • You can also use NFS with AD Kerberos
You do not need to use Kerberos	Directory Services	AD account creation automatically sets up Kerberos
You want to use two different AD domains for NFS and SMB	Directory Services and AD Account	At least one AD domain must be configured through directory service configuration because FlashBlade can join only one AD domain for any of the services

Scenario	Configuration	Additional Information
The realm that you use for Kerberos is disjointed from your AD domain	Directory Services	AD account configuration is not supported for environments using a disjointed Kerberos realm. You must upload keytab files manually on the FlashBlade for NFS Kerberos authentication.
You want to configure role-based access control (RBAC) for LDAP users	Directory Services	Configure the array management directory service and then configure user roles on the FlashBlade

Configuration Choices Based on Security Considerations

In addition to environmental requirements, there may be security considerations with your environment that require the configuration of directory services versus AD account as described below:

Scenario	Configuration
• Your corporate security policies mandate the Kerberos SPN and encryption types are managed on the domain side and not allow storage device to control it. • You configure service accounts that are used as the bind user/bind password on your storage devices. Such accounts would be read-only. • You configure your domains to allow anonymous binding for read-only operations from storage devices.	Configure directory services and upload keytab files if Kerberos is needed for NFS.
• You are not allowed to provide modify permisssions to your storage device at any time. • Your corporate security policies mandate that the Kerberos SPNs and encryption types are managed exclusively on the domain side and not allow storage devices to control it. • You may set up a computer account in your AD domain in advance to which FlashBlade can connect, allowing an AD account to be configured.	• Configure a computer account to allow minimal permissions to the FlashBlade upon joining the domain. • Configure AD account, joining the FlashBlade to your domain by connecting it to the pre-created computer account.

Scenario	Configuration
• FlashBlade is able to communicate only with a read-only domain controller in the domain, with change-capable domain controllers being unreachable. • You may have disjoint data and management networks.	Configure directory services and upload keytabs If the change-capable domain controllers are on your data network, but you sync a read-only controller on the management network for FlashBlade.

Combining Active Directory Account Configuration with NFS Directory Service Configuration

An Active Directory (AD) account can be configured in combination with the NFS directory service for the same supported data services.

When both are configured, the configured AD account performs lookup and mapping. If no results are returned or the lookup fails because the AD domain is unreachable, the configured NFS directory service then attempts the lookup.

Configuring both AD and NFS directory service can provide the following advantages:

- If you previously configured the NFS directory service and wish to migrate to an AD account, you can do so non-disruptively. By leaving the NFS directory service configured, you can ensure there is no outage while creating and testing your AD account.
- You can non-disruptively migrate your servers for NFS over time with FlashBlade switching between them as records are added to one and removed from another.
- Failback if the domain's Kerberos servers go down if simple bind was used and directory service was configured with a read-only bind credential.
- Failback if DNS goes down is provided if the directory service configuration and AD account point to the same domain.

Note: Combining AD accounts with the SMB directory service is not supported. If combined, the SMB directory service configuration will be ignored.

Active Directory Overview

An Active Directory is a hierarchical directory of large groups of computers, objects, and users within a network.

Configuring an Active Directory (AD) account to join the FlashBlade to an AD domain automatically enables support for Kerberos authentication for NFS and SMB services on FlashBlade.

FlashBlade allows the creation of one AD computer account, or joining of an existing AD computer account.

When an AD account is created on the FlashBlade, Kerberos keytabs are automatically generated for each encryption type and service principal name (SPN) that is set on the account, allowing clients to authenticate using Kerberos the same way that they would if keytab files were uploaded manually on the FlashBlade. This computer account is also used by FlashBlade to perform Lightweight Directory Access Protocol (LDAP) queries to the AD domain, such as user/group membership, netgroup information, name/ID, mappings for quotas, etc, which are authenticated by the computer account using Kerberos.

An AD account can support NFS, SMB, or both depending on the SPN. If the account has any SPNs beginning with `nfs/`, then it will be used to support Kerberos authentication for NFS services. If the account has any service principle names beginning with `HOST/` or `cifs/` (both case insensitive), then it will be used to support Kerberos authentication for SMB service.

When creating an AD account, SPNs or fully qualified domain names (FQDNs) can be specified for the account. If you specify FQDNs (or do not specify either FQDNs or SPNs), the account will be assigned SPNs automatically for all services supported by FlashBlade. If native SMB is enabled at the time of account creation, the account will support both SMB and NFS. If native SMB is not enabled, only NFS is supported. By default, AES256 encryption is used, though AES128 and RC4-HMAC are also supported as other choices.

Active Directory Networking Requirements

Network Ports

FlashBlade uses a variety of different protocols that require a variety of ports. Ports must be reachable on the Active Directory (AD) domain controller from the FlashBlade management network.

The following ports must not be blocked:

- Port 636 - default port used for secure communication (LDAPS) between the FlashBlade and AD domain for core AD functionality. Note that all LDAP communication is over TCP.
- Port 88 - default port used for Kerberos authentication within the domain allowing FlashBlade to perform "Kerberized" binds with the AD domain.

Network Speed

Slow environments can affect Kerberos performance. Kerberos KDCs, which handle key generation and authentication, will reject a request if two identical requests are received in a short time frame. This can occur if your network is slow or unstable, or with large MTU sizes.

Time Sync

FlashBlade and AD domain controllers (DCs) must be synchronized with Network Time Protocol (NTP) servers to maintain the network reference time. The time on the FlashBlade must be the same as the AD DCs. It is recommended that FlashBlade and AD DCs use the same NTP servers. If the time on the FlashBlade differs from the time in the AD domain, the FlashBlade's client Kerberos authentication requests with the domain might be rejected.

DNS Name Resolution

Domain Name System (DNS) is necessary for discovering domain controllers, KDCs, and LDAP servers within an AD domain. It is also necessary for Kerberized communications between the FlashBlade and the AD domain.

FlashBlade checks for the following DNS service entries:

- Service record: `_ldap._tcp.dc._msdcs.<DomainName>` - used to discover domain controllers against which LDAP queries can be made.
- Service record: `_kerberos._tcp.dc._msdcs.<DomainName>` - used to discover domain controllers that can issue Kerberos client tickets.

Resolving a hostname (or FQDN) to an IP address requires that DNS records be configured. You must configure reverse DNS if resolving an IP address to a hostname (or FQDN). A reverse DNS zone must be created to map IP addresses within the normal DNS zone in a range, and then PTR records can be automatically created when normal *hostname-to-IP* records are created. This process varies depending on the software used for your DNS nameserver.

Though Kerberos is designed to authenticate with FQDNs, some Kerberos clients support using IP addresses if reverse DNS is configured for the IP address. When reverse DNS is configured, and IP addresses are used, such a Kerberos client resolves the IP address to the proper FQDN and then obtains a client ticket for that FQDN. If you wish to specify IP addresses for the directory server and/or Kerberos server preferences on the FlashBlade, reverse DNS must be configured for all such IP addresses.

If the FlashBlade AD account was configured using only FQDNs and not short host names, i.e. `array1.mydomain.com` and `array1`, then clients attempting a Kerberos mount will not be able to obtain a ticket from the domain of the short host name. This will result in AD being unable to map the client ticket request against the FlashBlade account. Note that many DNS servers do not allow short host names to be configured. When using such a DNS server, a client-side configuration where the DNS search domain is configured to use short host names is required. How you configure the DNS search domain is dependent on the client type, I.e. Windows versus Linux; contact your network administrator for assistance.

Computer Account Naming and Credentials

By default, your Active Directory (AD) account configuration name will become the name of the computer account that is created by FlashBlade. The computer account name will be used to generate the fully qualified domain names (FQDNs), which can then be used for mapping the SMB shares in UNC pathnames. This is performed in a way that conforms with AD standards.

For example, the default FQDN for any `computer-name` in `example.domain.com` will be:

- Short host name: `computer-name`
- FQDN: `computer-name.example.domain.com`

The following restrictions apply for the computer name:

- A separate common name vs sAMAccountName is not supported.
- If you configure `computer-name` for the computer name attribute, FlashBlade sets the common name to `computer-name` and the sAMAccountName to `computer-name$`.
- The computer name cannot exceed 15 characters.
- The computer name must be a valid DNS name if no alternative FQDNs or service principal names (SPNs) are specified.

- The computer name cannot be changed after creation.
- The computer account name should resolve to a data VIP.

When creating an account in an AD domain, you must provide user credentials for FlashBlade to use to create the account. These credentials are not stored. Credentials can take any of the following formats:

- Short username style (sAMAccountName): `username`
- Kerberos style: `username@domain.com`
- Windows logon style: `domain.com\username`
- LDAP distinguished name style: `CN=UserName,CN=...,DC=domain,DC=com`

Minimum Permissions for Joining an Active Directory Domain

When FlashBlade joins an Active Directory (AD) domain the credentials used are not stored. The minimum set of user permissions needed in order to be used by FlashBlade for joining an AD domain depends on the mechanism used for joining.

The following are the minimum set of permissions needed if FlashBlade creates its own AD account:

- Create computer objects within the OU being joined.
- (Optional) Read and write computer account permissions allowing the created computer account to delete itself and update its own SPNs.

The following are the minimum set of permissions needed if FlashBlade connects to an existing AD account that was pre-created for FlashBlade:

- Basic READ permissions used to validate the attributes of the account:
  - Read servicePrincipalNames
  - Read msDS-SupportedEncryptionTypes
  - Read userAccountControl
- Reset password. This allows FlashBlade to reset the password on the computer account when generating Kerberos keys for the account.

Joining an Organizational Unit

When creating an Active Directory (AD) account, you must specify the organization unit (Join OU) within the AD domain where the computer account will reside. The location of the computer account can influence applied domain policies and permissions needed by credentials during creation.

By default, the computer accounts will use the CN=Computers OU. When specifying the Join OU, the OU should be specified as a relative distinguished name (RDN) in the LDAP style, for example: `OU=location,OU=storage-machines,CN=ny-datacenter`.

Note: The Join OU cannot be changed at the same time the encryption types, SPNs, or FQDNs are changed. This is due to permission changes that can occur when changing the Join OU.

Joining with an Existing Computer Account

In some instances, it may be preferable for FlashBlade to join using an existing Active Directory (AD) computer account in an AD domain (see *Choosing an Active Directory Account or Directory Services Configuration* for scenarios). This can include when environments have strong security restrictions or when different admin teams manage DNS naming for the FlashBlade as well as client encryption.

In order to use a pre-created computer account to join the FlashBlade to the domain, the following attributes must be pre-configured in AD:

- msDS-SupportedEncryptionTypes - defines what encryption types can use Kerberos authentication with FlashBlade.
- servicePrincipalNames - defines what protocols clients can use to mount and what FQDNs the protocols can mount over. Service components must be valid for the given FlashBlade; nfs/ for non-native SMB systems and nfs/, cifs/, or HOST/ for native SMB systems).
- userAccountControl - dictates controls on an account, including its password restrictions, enablement status, etc.

For the attributes above:

- The user whose credentials are used to perform the action of joining the domain must have READ access for these attributes on the pre-created computer account.
- You may wish to configure the msDS-SupportedEncryptionTypes and servicePrincipleNames attributes before having the FlashBlade join as the AD account configuration will pull the values of these attributes from the computer account.

- The userAccountControl attribute must include the WORKSTATION_TRUST_ACCOUNT and DONT_EXPIRE_PASSWD flags.

When joining an existing AD account, you need only specify the AD domain, credentials, and your computer account name. You can optionally specify directory and Kerberos servers. The Join OU, encryption type, and SPNs/FQDNs are "inherited" from the existing AD computer account. FlashBlade searches and validates the computer account and then resets the account's password/Kerberos key in order to generate Kerberos keys for FlashBlade use.

For security purposes, when joining an existing AD account, multiple FlashBlade arrays or other appliances cannot share the same computer account.

Service Principal Names

When creating an Active Directory (AD) account using the FlashBlade GUI or CLI, you can specify what service principal names (SPNs) will be used by clients to mount FlashBlade file systems. Up to 32 SPNs can be specified per computer account. Only a single computer name per AD account can be configured.

An SPN is a unique identifier of a service instance and is used by Kerberos authentication to associate a service instance with a service account.

If an SPN is not specified during the creation of an AD account on FlahBlade, the computer name will be used as the SPN name. If an AD account is added with the computer name without specifying an SPN then you will have to manually register the data VIP against the computer account name in the DNS server.

SPNs follow two formats:

- Short name - `<SERVICE>/<NAME>`
- Long name - `<SERVICE>/<NAME>.<DOMAIN_NAME>`

`<SERVICE>` can have the following values:

- `nfs`: For NFSv4.1 Kerberized access. NFS must be specified in lowercase, i.e. `nfs`.
- `cifs`: Available if native SMB is enabled.
- `HOST`: Available if native SMB is enabled. This is the default SMB service with AD.

<NAME> should be the DNS name associated with the data VIP on the FlashBlade. When accessing the file system, this name should be used instead of the IP address for Kerberized authentication.

If SPNs were not specified during AD account creation, they can be added to the account at a later time by editing the AD account on the FlashBlade.

If you configured multiple data VIPs for different services, you can register them as extra SPNs with the <NAME> registered on a DNS server corresponding to the FlashBlade data VIP. For example, if you configured data VIP 10.1.1.2 with the DNS name `data-vip2` for NFS and data VIP 10.2.2.3 with the DNS name `data-vip3` for NFS and SMB, you would specify the following SPNs:

- `nfs/data-vip2`
- `nfs/data-vip2.mydomain.com`
- `nfs/data-vip3`
- `nfs/data-vip3.mydomain.com`
- `HOST/data-vip3`
- `HOST/data-vip3.mydomain.com`

Directory and Kerberos Servers

Up to five directory servers and/or five Kerberos servers can be specified for Active Directory (AD) account configuration.

During configuration, FlashBlade automatically discovers the five fastest/closest directory and Kerberos servers in the DNS. You may wish to manually configure the servers if:

- FlashBlade needs to connect to specific servers in the domain with very specific policies. For example, you may have policies to derive RFS2307 UID/GID mappings on only a subset of domain controllers; your FlashBlade configured for both NFS and SMB would need to be configured accordingly in order to communicate with the server.
- all updates to the domain will happen in a specific change-capable server and you would like the FlashBlade to pick up the updates without waiting for them to propogate across the domain.

When configuring servers, the following formats are accepted:

- DNS names

- IP addresses for Kerberos servers.
- IP addresses for directory servers if reverse DNS is configured for those IP addresses.

User and Group Name Mapping for Active Directory and LDAP and NIS

SMB/NFS cross-protocol access is supported for user and group name mapping between AD (Active Directory) users and LDAP (Lightweight Directory Access Protocol) and between AD and NIS (Network Information Service) users on the FlashBlade array. Name mapping is enabled by default for matching user and group names. Active Directory and OpenLDAP, and Active Directory and NIS as Directory Services are supported.

FlashBlade maps identities stored in separate directory service servers (Active Directory for SMB identity, and LDAP-capable server for NFS identity).

- The SMB user/group identity is expressed as a SID (long string such as S-1-5-21-35434234-242322-542423-1245).
- The NFS identity is expressed as a UID/GID (32 bit number such as 1000).

The SID to UID/GID mapping (and reverse) is performed by matching usernames used by either an AD and LDAP server, or AD and NIS server.

A query is issued against AD querying the SID belonging to the username. Another query is issued against the LDAP server or NIS server querying the UID/GID belonging to the same username. The SID and UID/GID returned by the queries are logically paired by Purity//FB as identifying the same user. This identity is used to evaluate permissions set by either protocol, and is expressed as SIDs for SMB permissions and UID/GID for NFS permissions.

Creating an Active Directory Account

Note: IPv6 is not supported with this release for multi-tenancy file access and file systems with NFS and SMB access enabled need to have access to UID/GID sources via LDAP or AD.

To create an Active Directory account, perform the following:

1. From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.
2. Click the Add (+) button. The **Create Active Directory Account** pop-up window appears.

3 Select **Create a new account**.

4 Enter the name of the new account in the **Name** field.

5 In the **Server** list, select the server.

6 In the **Computer Name** field, enter the name for the AD computer account that will be created in the AD domain. If a computer name is not entered, the computer name will default to the value entered in the **Name** field. The computer account name should resolve to the data VIP.

7 In the **Domain** field, enter the AD domain to which this computer account will be created.

8 In the **User Name** field, enter the name of the user who is capable of creating the computer account within the domain.

9 Enter the user's password in the **Password** field.

If you wish to enter additional configuration information, click **Additional Settings**.

10 In the **Join OU** filed, enter the organization unit (location) where the computer account will be created. For example, **OU=Arrays**, **OU=Storage**, etc. If an OU is not entered, the value defaults to **CN=Computers**.

11 Select the **Encryption Type**. The default encryption type is **aes256-sha1**. If NTLMv2 is required, **arcfour-hmac** encryption must be enabled.

12 Click the Add (+) button to the right of the **Service Principal Names** field and enter a comma-separated list of service principal names (SPNs) or fully qualified domain names (FQDNs) for registering services with the AD domain. If nothing is entered, the value defaults to **Computer Name Domain** as an FQDN.

13 Click the Add (+) button to the right of the **Kerberos Servers** field and enter a comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.

14 Click the Add (+) button to the right of the **Directory Servers** field and enter a comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.

15 Click **Create**.

Instructions and troubleshooting on accessing Native SMB file systems using NTLMv2 are available in *this Knowledge Base article*.

Creating an Active Directory Account using the CLI

Active Directory (AD) accounts can also be created using the **`puread account create`** CLI command. For full details about the **`puread`** command, see the *FlashBlade CLI Reference*.

To create an active directory account using the CLI, perform the following:

1 Run the **`puread account create account --domain DOMAIN ACCOUNT`** command, where `ACCOUNT` is the name of the new AD account. Note that the `--domain` argument, which specifies the AD domain to which the computer account will be created, is required. You can also specify the following:

- `--computer-name`: (Optional) The name for the AD computer account that will be created in the AD domain. If a computer name is not entered, the computer name will default to the name used for the `ACCOUNT-NAME`. The computer account name should resolve to the data VIP.
- `--join-ou`: (Optional) The organization unit (location) where the computer account will be created. For example, `OU=Arrays`, `OU=Storage, etc`. If an OU is not entered, the value defaults to `CN=Computers`.
- `--encryption-type`: (Optional) A comma-separated list of encryption types that will be supported for use by clients for Kerberos authentication. Defaults to `AES256-cts-hmac-sha1-96`.
- `--fqdn`: (Optional) A comma-separated list of fully qualified domain names (FQDNs) used for the creation of keys for Kerberos authentication. If FQDNs are specified, all supported service principals will be registered for them. If neither FQDNs nor SPNs are specified, all supported service princples with the `COMPUTER_NAME.DOMAIN` as the FQDN are registered.
- `--spn`: (Optional) A comma-separated list of service principal names (SPNs) or fully qualified domain names (FQDNs) for registering services with the AD domain. Up to 32 SPNs can be specified. If nothing is entered, all supported SPNs with the `COMPUTER_NAME.DOMAIN` as the FQDN are registered.
- `--kerberos-server`: (Optional) A comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- `--directory-server`: (Optional) A comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- `--global-catalog-server`: (Optional) A comma-separated list of directory servers that will be used for lookups related to user authorization. Servers can be listed as IP addresses or DNS names with the optional @domain suffix. If the suffix is ommited, the joined domain is assumed. All servers speciifed must be registered to the domain in the array's configured DNS and will be communicated with over the LDAPS protocol.

For example:

```
puread account create mynewadaccnt --domain mycorp.com --computer-name
flashblade01
```

2 When prompted enter your user name and password.

```
puread account create mynewadaccnt --domain mycorp.com --computer-name
flashblade01
Enter the user name used to join:
Enter password for user:
```

In this example:

- the created account is `mynewadaccnt`
- the account is created in the domain `mycorp.com`
- the account is created in the organization unit (OU) `CN=Computers` because no Join OU value was specified
- the AD computer name that is created in the AD domain is `flashblade01`
- the encryption type defaults to `AES-256-SHA1` because no encyrption type was specified
- all supported service principles with `COMPUTER_NAME.DOMAIN` as the FQDN are registered because no FQDNs or SPNs were specified
- Kerberos and directory servers are resolved for the domain in DNS

Joining Active Directory with an Existing Computer Account

When joining AD with an existing computer account, you need only specify the AD domain, credentials, and your computer account name. The Join OU, encryption type, and SPNs/FQDNs are "inherited" from the existing AD account.

To create a computer account to join an existing AD account, perform the following:

1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.

2 Click the Add (+) button. The **Create Active Directory Account** pop-up window appears.

3 Select **Connect to existing account**.

4 Enter the name of the new account in the **Name** field.

5 In the **Computer Name** field, enter the name for the AD computer account that will be created in the AD domain. If a computer name is not entered, the computer name will default to the value entered in the **Name** field.

6 In the **Domain** field, enter the AD domain to which this computer account will be created.

7 In the **User Name** field, enter the name of the user who is capable of creating the computer account within the domain.

8 Enter the user's password in the **Password** field.

If you wish to enter additional configuration information, click **Additional Settings**.

9 In the **Kerberos Servers** field, enter a comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.

10 In the **Directory Servers** field, enter a comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.

11 Click **Create**.

Joining Active Directory with an Computer Existing Account using the CLI

The CLI can also be used to join an existing Active Directory (AD) account.

Use the **`puread account create --join-existing-account`** command to create a computer account to join an existing AD account. For full details about the **`puread`** command, see the *FlashBlade CLI Reference*.

To create a computer account to join an existing AD account using the CLI, perform the following:

1 Run the **`puread account create account --domain --join-existing-account ACCOUNT-NAME`** command, where `ACCOUNT-NAME` is the name of the existing AD account. Note that the `--domain` argument, which specifies the AD domain to which the computer account will be created, is required. You can also specify the following arguments:

- `--computer-name`: (Optional) The name for the AD computer account that will be created in the AD domain. If a computer name is not entered, the computer name will default to the name used for the `ACCOUNT-NAME`.
- `--kerberos-server`: (Optional) A comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- `--directory-server`: (Optional) A comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.

For example:

```
puread account create mynewaccnt --computer-name newcomputer --domain
mycorp.com --join-existing-account
```

2 When prompted enter your user name and password.

```
puread account create mynewaccnt --computer-name newcomputer --domain
mycorp.com --join-existing-account
Enter the user name used to join:
Enter password for user:
```

In this example:

- the computer account `newcomputer` is created in the existing AD account `mynewaccnt`.
- the account is in the domain `mycorp.com`

Testing an Active Directory Configuration

An Active Directory (AD) configuration can fail if:

- there is a change to your DNS configuration (or nameserver failure) resulting in FlashBlade's auto-resolution of LDAP and Kerberos servers to fail.
- the server that FlashBlade is communicating with for user and group lookups is misconfigured or goes down.
- the AD account is moved, modified, reset, or deleted by an administrator.

If any of these have occurred, test your AD account to ensure that your configuration is valid.

To test an AD configuration, perform the following:

1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.

2 Click the AD account. The **Details** panel appears.

3 Click the **More Options** button and select **Test**. The **Test Active Directory Configuration** pop-up window appears displaying test results.

Figure 10.1 Viewing Active Directory Configuration Test Results

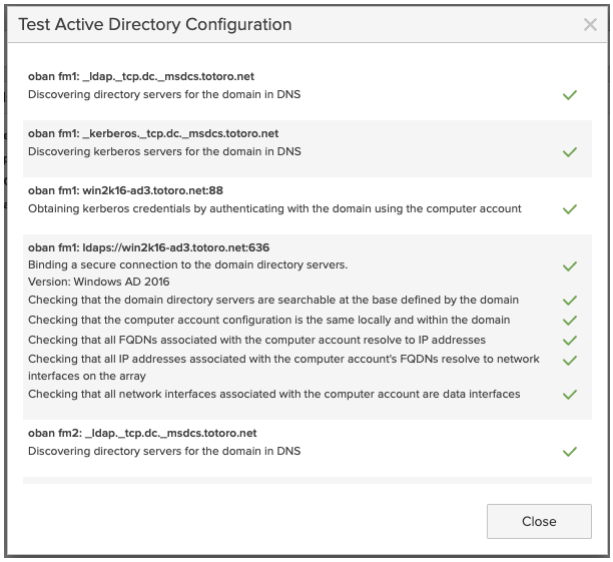

If there are items that did not pass, you may need to edit your AD account configuration.

4 Click **Close**.

Testing an Active Directory Configuration using the CLI

You can also use the CLI to test your AD configuration to ensure that your configuration is valid.

Run the **`puread account test`** command to test an AD using the CLI. For full details about the **`puread`** command, see the *FlashBlade CLI Reference*.

For example:

```
puread account test
Name  Component Name  Component Address  Destination
Test                                     Success  Details
oban  fm1             10.64.242.197      _ldap._tcp.dc._msdcs.totoro.net
dns-directory-server-discovery  True     -
                                         _kerberos._tcp.dc._msdcs.totoro.net
dns-kerberos-server-discovery   True     -
                                         win2k16-ad3.totoro.net:88
kerberos-initialization         True     -
                                         ldaps://win2k16-ad3.totoro.net:636
directory-server-binding        True     Version:...

base-dn-searching               True     -

account-consistency-check       True     -

network-resolution-check        True     -

network-interface-check         True     -

network-service-check           True     -
      fm2             10.64.242.198      _ldap._tcp.dc._msdcs.totoro.net
dns-directory-server-discovery  True     -
                                         _kerberos._tcp.dc._msdcs.totoro.net
dns-kerberos-server-discovery   True     -
                                         win2k16-ad3.totoro.net:88
kerberos-initialization         True     -
                                         ldaps://win2k16-ad3.totoro.net:636
directory-server-binding        True     Version:...

base-dn-searching               True     -

account-consistency-check       True     -

network-resolution-check        True     -

network-interface-check         True     -

network-service-check           True     -
```

Verify that all tests are successfully passed (`True`) in the `Success` column.

Additional options are availble to format the output of the command. Refer to **`puread`** in the *FlashBlade CLI Reference* for details.

Editing an Active Directory Account

To edit an Active Directory (AD) account, perform the following:

1. From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.

2 Click the AD account. The **Details** panel appears.

3 Click the **More Options** button and select **Edit Account**. The **Edit Active Directory Account** pop-up window appears.

4 In the **Join OU** field, enter the organization unit (location) where the computer account will be created. For example, **OU=Arrays**, **OU=Storage**, etc. If an OU is not entered, the value defaults to **CN=Computers**.

5 Select the **Encryption Type**. The default encryption type is **aes256-sha1**. If NTLMv2 is required, **arcfour-hmac** encryption must be enabled.

6 Click the Add (+) button to the right of the **Service Principal Names** field and enter a comma-separated list of service principal names (SPNs) for registering services with the AD domain. If you wish to remove previously entered SPNs, click the Remove (**x**) button next to the SPN you wish to remove.

7 Click the Add (+) button to the right of the **Kerberos Servers** field and enter a comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.

8 Click the Add (+) button to the right of the **Directory Servers** field and enter a comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.

9 Click **Save**.

Instructions and troubleshooting on accessing Native SMB file systems using NTLMv2 are available in *this Knowledge Base article*.

Editing an Active Directory Account using the CLI

The **`puread account setattr`** CLI command can also be used to edit the configuration of active directory accounts. For full details about the **`puread`** command, see the *FlashBlade CLI Reference*.

When running the command, the `ACCOUNT-NAME` must be specified. Similar to creating an AD account, the following can be set/changed:

- `--join-ou`: (Optional) The organization unit (location) where the computer account will be created. For example, `OU=Arrays`, `OU=Storage, etc`. If an OU is not entered, the value defaults to `CN=Computers`.

- `--encryption-type`: (Optional) A comma-separated list of encryption types that will be supported for use by clients for Kerberos authentication. Defaults to `AES256-sha1`.

- `--fqdn`: (Optional) A comma-separated list of fully qualified domain names (FQDNs) used for the creation of keys for Kerberos authentication. If FQDNs are specified, all supported service principals will be registered for them. If neither FQDNs nor SPNs are specified, all supported service principles with the `COMPUTER_NAME.DOMAIN` as the FQDN are registered.
- `--spn`: (Optional) A comma-separated list of service principal names (SPNs) or fully qualified domain names (FQDNs) for registering services with the AD domain. Up to 32 SPNs can be specified. If nothing is entered, all supported SPNs with the `COMPUTER_NAME.DOMAIN` as the FQDN are registered.
- `--kerberos-server`: (Optional) A comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- `--directory-server`: (Optional) A comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- `--global-catalog-server`: (Optional) A comma-separated list of directory servers that will be used for lookups related to user authorization. Servers can be listed as IP addresses or DNS names with the optional @domain suffix. If the suffix is ommited, the joined domain is assumed. All servers speciifed must be registered to the domain in the array's configured DNS and will be communicated with over the LDAPS protocol.

For example, an account previously created with the default encryption type AES256-SHA1 is updated to also use RC4-HMAC:

```
puread account setattr mynewacct --encryption-type AES256-SHA1,RC4-HMAC
```

Rotating Keytabs

FlashBlade provides the ability to rotate keytabs. Rotating keytabs creates new keytabs for your Active Directory (AD) account.

You may need to rotate your AD acocunt's keytabs if:

- your keytabs have been compromised.
- there is a security policy requiring your security implementation to change after a specific duration.

When keytabs are rotated, new keytabs are created and added to your AD account. You can specify a new prefix for the new keytabs. Your existing keytabs will still exist and can be deleted as needed to ensure that users are not locked out.

To rotate keytabs for your Active Directory (AD) account, perform the following:

1. From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.
2. Click the AD account. The **Details** panel appears (the **Kerberos Keytabs** panel appears below the **Details** panel).
3. Click the **More Options** button in the header of the **Details** panel and select **Rotate Keytabs**. The **Rotate Keytabs** pop-up window appears.
4. (Optional) In the **Keytab Name Prefix**, enter a new prefix to be used for the generated keytabs.
5. Click **Rotate**.

The new keytabs appear at the top of the **Kerberos Keytabs** panel.

Rotating Keytabs using the CLI

Keytabs can be rotated for Active Directory (AD) accounts using the CLI.

Use the **`purekeytab create --ad`** CLI command to rotate keytabs. For full details about the **`purekeytab`** command, see the *FlashBlade CLI Reference*

When rotating keytabs, new keytabs are created and added to your AD configuration, as identified with the `--ad` argument. You can optionally specify the `--prefix` argument. This allows you to specify a prefix used for naming the new keytab entries.

For example:

```
purekeytab create --ad myadaccnt --prefix myaccnt
```

In this example, `--ad myadaccnt` specifies the new keytabs are created for the `myadaccnt` AD configuration and `--prefix myaccnt` specifies that all new keytabs begin with the prefix `myaccnt`.

Your existing keytabs will still exist and can be deleted as needed to ensure that users are not locked out.

Importing a Keytab File

To import a keytab file, perform the following:

1. From the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.
2. Click the **More Options** button in the header of the **Kerberos Keytabs** panel and select **Import**. The **Import Keytab File** pop-up window appears.
3. Enter a prefix in the **Name Prefix** field. Specifying a file entry prefix is required because a single keytab file can contain multiple keytab entries in multiple slots.
4. Click **Choose File** in the **Keytab File** field. Navigate to, and select the keytab file you wish to import.
5. Click **Import**.

Importing a Keytab File using the CLI

Use the **`purekeytab create --keytab-file`** CLI command to import keytabs. For full details about the **`purekeytab`** command, see the *FlashBlade CLI Reference*

Specifying a file entry prefix with the `--prefix` argument is required because a single keytab file can contain multiple keytab entries in multiple slots. When prompted, enter the keytab file you wish to import. The keytab file must be in Base64 MIME-encoded format.

For example:

```
purekeytab create --prefix new --keytab-file
Please enter a keytab file in Base64 MIME-encoded format followed by ^D:
```

In this example, the entered keytab beginning with the prefix `new` is imported.

Exporting a Keytab File

To export a keytab file that is currently present on your FlashBlade, perform the following:

1. From the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.
2. To export one or more keytab files:
   a. Click the **More Options** button in the header of the **Kerberos Keytabs** panel and select **Export**. The **Export Keytabs** pop-up window appears.
   b. Select one or more keytab files from the **Existing Keytabs** column on the left. Each selected keytab file appears under the **Selected Keytabs** column on the right.

c Click **Export**.

3 To export a single keytab:

a From the **Kerberos Keytabs** panel, click the **More Options** button in the same row as the keytab file you wish to export. and select **Export**. The **Export Keytab File** pop-up window appears.

b Click to export in **Binary** or **MIME** format.

Exporting a Keytab File using the CLI

You cannot directly export a keytab file present on the FlashBlade using the CLI. However, the CLI does allow you to view the BASE64 MIME-encoded format of keytabs currently in use with the AD account that you can copy and use elsewhere.

Use the **`purekeytab list --keytab-file`** CLI command to view the BASE64 MIME-encoded format of a keytab file. For full details about the **`purekeytab`** command, see the *FlashBlade CLI Reference*

For example:

```
purekeytab list --keytab-file
Keytab File
BQIAAABNAAIAClRPVE9STy5ORVQABEhPU1QACk9CQU4tU01CLTIAAAABX5KrmgUAEgAgmJn3q0GV
ZuieCjlU8MmZWoLyO16EuQaWLB+GQNWC85MAAABMAAIAClRPVE9STy5ORVQAA25mcwAKT0JBTi1T
TUItMgAAAAFfkquaBQASACCYmferQZVm6J4KOVTwyZlagvI7XoS5BpYsH4ZA1YLzkwAAAFgAAgAK
VE9UT1JPLk5FVAAESE9TVAAVb2Jhbi1zbWItMi50b3Rvcm8ubmV0AAAAAV+Sq5oFABIAIJiZ96tB
lWbongo5VPDJmVqC8jtehLkGliwfhkDVgvOTAAAAVwACAApUT1RPUk8uTkVUAANuZnMAFW9iYW4t
c21iLTIudG90b3JvLm5ldAAAAAFfkquaBQASACCYmferQZVm6J4KOVTwyZlagvI7XoS5BpYsH4ZA
1YLzkwAAAE0AAgAKVE9UT1JPLk5FVAAESE9TVAAKb2Jhbi1zbWItMwAAAAFfkquaBQASACCYmfer
QZVm6J4KOVTwyZlagvI7XoS5BpYsH4ZA1YLzkwAAAEwAAgAKVE9UT1JPLk5FVAADbmZzAApvYmFu
LXNtYi0zAAAAAV+Sq5oFABIAIJiZ96tBlWbongo5VPDJmVqC8jtehLkGliwfhkDVgvOTAAAAWAAC
AApUT1RPUk8uTkVUAARIT1NUABVvYmFuLXNtYi0zLnRvdG9yby5uZXQAAAABX5KrmgUAEgAgmJn3
...
```

Copy and save the output for use elsewhere.

Deleting a Keytab File

Note: Deleting keytab files can be a disruptive action if there were clients with active sessions that relied on the deleted keytab.

To delete a keytab file, perform the following:

1 From the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.

2 To delete one or more keytab files:

a Click the **More Options** button in the header of the **Kerberos Keytabs** panel and select **Delete**. The **Delete Keytabs** pop-up window appears.

b Select one or more keytab files from the **Existing Keytabs** column on the left. Each selected keytab file appears under the **Selected Keytabs** column on the right.

c Click **Delete**.

3 To delete a single keytab:

a From the **Kerberos Keytabs** panel, click the **More Options** button in the same row as the keytab file you wish to delete and select **Delete**. The **Delete Keytab** pop-up window appears.

b Click **Delete**

Deleting Keytab Files using the CLI

Use the **`purekeytab delete KEYTAB`** CLI command to delete one or more keytabs from the FlashBlade. For full details about the **`purekeytab`** command, see the *FlashBlade CLI Reference*

`KEYTAB` is the name of the keytab file you wish to delete. Multiple keytab files can be specified in a comma-separated list.

For example:

```
purekeytab delete keytab1,keytab2,keytab3
```

In this example, keytab files `keytab1`, `keytab2`, and `keytab3` are deleted from the AD account.

Deleting an Active Directory Account

Deleting an Active Directory (AD) account deletes the computer account associated with the FlashBlade and all Kerberos keytabs associated with the account.

To delete an Active Directory (AD) account, perform the following:

1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.

2 Click the **Delete** icon in the row of the account. The **Delete Active Directory Account** pop-up window appears.

3 (Optional) If you wish to retain the computer account associated with the FlashBlade in your AD domain, enable (blue) **Local Only**. Note that if enabled, you will need to manually delete Kerberos keytabs.

4 Click **Delete**.

Deleting an Active Directory Account using the CLI

Deleting an Active Directory (AD) account can also be performed using the CLI.

Use the **`puread account delete`** command to delete an AD account using the CLI. For full details about the **`puread`** command, see the *FlashBlade CLI Reference*.

To delete an AD account using the CLI, issue the **`puread account delete ACCOUNT-NAME`** command, where `ACCOUNT-NAME` is the AD account name you wish to delete. For example:

```
puread account delete mynewacct
```

You can optionally specify the `--local-only` argument, which will only delete the AD account on the local array. The computer account will not be deleted in the AD domain.

Directory Services Overview

A directory service allows the configuration of a server that will be used to perform some form of lookup. For example, a management directory server is used to look up users and validate their passwords, as well as to search for groups to which the users belong.

By configuring a directory service, FlashBlade is able to look up relevant information to authenticate users along different data and management paths.

A directory service is defined by the following configuration:

- A list of one or more URIs that specify the directory servers that will be used for lookups by the directory service. These may be specified in the form of IP addresses.
- A Base Distinguished Name (Base DN) as the point from which a server will begin its searches.
- A bind user and bind password (user and password) of a single service account that is used by FlashBlade to search the directory server.

Generally, FlashBlade communicates with directory servers over the Lightweight Directory Access (LDAP) protocol.

FlashBlade has three directory services:

- Array Management (see *Array Management Directory Service* for more details)
- NFS (see *NFS Directory Service* for more details)
- SMB (see *SMB Directory Service* for more details)

If you want to enable Kerberos authentication for FlashBlade's clients when configuring NFS or SMB directory services, you must upload keytab files with keytabs for the appropriate encryption types and service principal names that clients will use. Doing so grants the same client functionality as joining an AD domain.

NFS and SMB Nested User Groups

Note: Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

A nested group is a group which is a member of another group. When using groups to manage permissions, you can create nested groups to allow inheritance of membership permissions from one group to its sub-groups. Similarly, users can belong to multiple group hierarchies and inherit those groups' membership permissions.

Purity//FB supports NFS and native SMB users belonging to a nested hierarchy of group membership (NFS and SMB nested groups). This allows access to users who are not the owners of particular files, but who instead belong to a group that owns the files. A user can belong to a group directly or indirectly (when groups are nested, i.e. when a group belongs to another group). Without NFS or SMB nested groups, Pure's NFS server allows only access to users that are direct members of the group that owns the file. With NFS and SMB nested groups, a user's membership in nested groups, as well as in direct groups, is taken into account for access checks. NFS and SMB interoperability with nested groups is supported; FlashBlade looks up both nested SIDs and GIDs.

The following limits apply to nested groups:

- Maximum number of groups, including direct and indirect groups, per user = 10,000
- Maximum number of nested groups to which a user can be a direct member = 500
- Maximum number of nested groups to which a group can be a direct member = 500
- Maximum number of nested levels = 5

Pure's implementation of NFS nested groups supports the following LDAP-based directory servers:

- OpenLDAP
- Oracle Unified Directory (OUD)
- Active Directory (AD)

In order to use NFS and SMB nested groups with directory services, nesting must be configured according to the RFC2307bis specification on your LDAP server.

Access from multiple domains with two-way trusted domains are supported within the same AD forest. Access across different AD forests is not supported.

Array Management Directory Service

The array management directory service configuration is used to allow users within an LDAP server to log in to a Flashblade (via the CLI, GUI, or REST) and manage it.

In order for a user to log in to the FlashBlade, they must be within a group that is configured to have a role on the array. If a group does not have a role, then its members cannot log in.

When configuring the array management directory service, a bind user and bind password can be optionally specified. If not specified, FlashBlade attempts to bind anonymously (without authenticating) with the configured LDAP servers. Note that if the configured LDAP servers do not grant read permissions on the needed pathways to unauthenticated machines, the anonymous bind will fail.

Roles

Directory services for array management can configure role-based access control (RBAC) for LDAP users by configuring groups in the directory that corresponds to the following permission groups (roles) on the array.

Role	Description
readonly	Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array
storage_admin	Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
array_admin	Array Admin users can perform all FlashBlade operations.

Role	Description
ops_admin	Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.

When an array management user connects to the FlashBlade, the array confirms the user's identity from the directory service. The response from the directory service includes the user's group, which Purity//FB maps to a role on the array, granting access accordingly.

If users are assigned to more than one role, an alert is issued. For security purposes, a user belonging to more than one FlashBlade role will inherit the permissions of the most restrictive role. For example, a user who has been assigned both the Storage Admin and Read Only roles will inherit the privileges of the more restrictive Read Only role. Note that the user will continue to have privileges from the original role until the directory service role cache is cleared (cache is automatically cleared every 8 hours). To make new permissions take immediate effect, run the **`pureadmin refresh`** command.

One exception to this rule is users assigned the Ops Admin role—users assigned the Ops Admin role and any other role are denied the ability to log into the FlashBlade. If there is an error with the directory service configuration on the FlashBlade array, contact a user with Array Admin privileges to resolve the issue. If the LDAP server is configured incorrectly, contact your LDAP server administrator. The alert issued in the case where a user is denied access due to being assigned the Ops Admin and another role, provides additional information to help resolve the issue.

As a best practice, after changing a user's role or management directory service configuration settings, use the **`pureadmin refresh`** command to refresh the user's role or clear directory service role cache to ensure all user privileges are up to date.

For directory service-enabled accounts, user passwords to the FlashBlade are managed through the directory service while API tokens are configured through Purity//FB. The password can be a maximum of 100 characters in length and include any character that is entered from a US keyboard.

In the FlashBlade GUI, buttons and menu items for unauthorized actions are disabled per the user's configured role.

Multiple Directory Service Group Mapping to Role

For an LDAP user to log into the array, the user must belong to a configured group in the LDAP directory, and that group must be mapped to a Purity//FB RBAC role (array_admin, ops_admin, storage_admin, or

readonly) on the array. Once an LDAP group is mapped to a role, all members of the group are assigned that role.

Multiple groups can be mapped to a single role, thereby allowing users from different groups to share the same role and log into the array. This can be useful for large enterprises with complex directory service configurations with many different groups spread across their organization. This is feature is enabled by default and allows array_admin users to configure and manage multiple groups to a single role.

- Each Purity//FB role can be mapped to up to 100 LDAP groups.
- Of those groups, different group bases are allowed.
- The permission level of an individual user is cached locally to prevent frequently binding and querying the directory. The cache entries expire after a time limit at which point the directory is queried again and the cache entry is updated. Cache entries are also automatically updated when a new login that requires communication with the configured directory server(s) occurs.

Configuring the Array Management Directory Service

Note: Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The array management directory service enables FlashBlade management access for LDAP users.

To configure the array management directory service using the GUI, perform the following:

1. Verify that the directory server is accessible through a management interface.
2. Verify that the DNS settings can be used to resolve the directory server domain and server.
3. From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
4. In the **Directory Service** panel, select **Array Management**.
5. Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.
6. In the **URIs** field, type the comma-separated list of up to 30 URIs of the directory servers.

   Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory

service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

If specifying a domain name, it should be resolvable by the configured DNS servers.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from `0` to `255` representing a group of 8 bits.

For IPv6, specify the IP address in the form `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`) . Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

Optionally specify a port. Append the port number after the end of the entire address. Default ports are `389` for LDAP, and `636` for LDAPS. Non-standard ports can be specified in the URI if they are in use.

7 In the **Base DN** field, type the base distinguished name (DN) of the directory service.
The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.

⚠ **Important!** Anonymous binding is supported for the array management directory service. If configuring anonymous binding, neither the bind user nor the bind password should be specified.

8 In the **Bind User** field, type the username used to bind to and query the directory. For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`. For Active Directory servers, you may instead choose to enter the username - often referred to as the sAMAccountName or user login name - of the account that is used to perform directory lookups. The username cannot contain the characters `" [] : ; | = + * ? < > / \`, and cannot exceed 104 characters in length.

9 In the **Bind Password** field, type the password for the bind user account.

10 (Optional) Select a CA certificate or certificate group to verify the authenticity of configured LDAP servers to establish a TLS (Transport Layer Security) communication for directory services.

11 The **User Login Attribute** defaults to the sAMAccountName for Active Directory, or UID for other server types.

12 The **User Object Class** defaults to User for AD, posixAccount or shadowAccount for OpenLDAP depending on the server's schema, posixAccount for posix-compliant servers, or person for all other server types.

13 Click **Test** to verify the configuration information. The **Test <Directory Service> Configuration** pop-up window appears, displaying the output of the test. During the directory service test, Purity//FB tests the directory service configuration to verify that the URI can be resolved and that the directory service can successfully bind and query the tree using the bind user credentials.

   If the test fails, update the configuration information and retest.

   Alternatively, you may save the configuration and enable the directory service at a later time.

14 If the test passes, select the **Enabled** checkbox to enable the directory service.

15 Click **Save**.

Adding Role Mapping

Note: Role mapping can only be performed by array_admin users.

Once the array management directory service is configured, you can add (or map) a group or groups to the FlashBlade roles used with array management directory service. See *Roles* for details about FlashBlade roles.

To map a group to a role, perform the following:

1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.

2 In the **Role Mappings** panel, click the **Add** (+) button. The **Add Role Mapping** pop-up window appears.

3 Enter a name. The name can be the same as the role.

4 Select a role from the **Role** list.

5 In the **Group** field, enter the common name (CN) of the directory group. The name should be the Common Name of the group without the "CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, "pureusers,OU=PureStorage", where pureusers is the common name of the directory service group.

6 In the **Group Base** field, enter the path to the configured groups in the directory tree. The Group Base consists of one or more path components (OUs and CNs) that, when combined with the base DN attribute and the configured group CNs, complete the full distinguished name of each

group. The group base should specify “OU=” for each OU and “CN=” for each CN. Multiple path components should be separated by commas. The order of the path components should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: “OU=PureGroups,OU=SANManagers”.

7 Click **Add**.

If you wish to map multiple groups to the single role, repeat this procedure for each group you wish to map.

Editing Role Mapping

Note: Role mapping editing can only be performed by array_admin users.

The group and group base of existing role mappings can be edited.

To edit an existing role mapping, perform the following:

1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.

2 In the **Role Mappings** panel, click the **Edit** button at the end of the same row of the role mapping you wish to edit. The **Edit Role Mapping** pop-up window appears.

3 In the **Group** field, enter the common name (CN) of the directory group. The name should be the Common Name of the group without the "CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, "pureusers,OU=PureStorage", where pureusers is the common name of the directory service group.

4 In the **Group Base** field, enter the path to the configured groups in the directory tree. The Group Base consists of one or more path components (OUs and CNs) that, when combined with the base DN attribute and the configured group CNs, complete the full distinguished name of each group. The group base should specify “OU=” for each OU and “CN=” for each CN. Multiple path components should be separated by commas. The order of the path components should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: “OU=PureGroups,OU=SANManagers”.

5 Click **Save**.

Deleting Role Mapping

Note: Role mapping deletion can only be performed by array_admin users.

To delete an existing role mapping, perform the following:

1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.

2 In the **Role Mappings** panel, click the **Delete Role Mapping** button at the end of the same row of the role mapping you wish to delete. The **Delete Role Mapping** pop-up window appears.

3 Click **Delete**.

Configuring the Array Management Directory Service using the CLI

Note: Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The array management directory service can also be configured using the **`pureds setattr management`** command. For full details about the **`pureds`** command, see the *FlashBlade CLI Reference*.

To configure the array management directory service using the CLI, perform the following:

1 Issue the **`pureds setattr management`** command with the following arguments:

- `--uri`: Up to 30 URIs can be specified in a comma-separated list.

  Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

  We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

  If specifying a domain name, it should be resolvable by the configured DNS servers.

  If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from `0` to `255` representing a group of 8 bits.

  For IPv6, specify the IP address in the form `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`) . Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

  If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

  Optionally specify a port. Append the port number after the end of the entire address. Default ports are `389` for LDAP, and `636` for LDAPS. Non-standard ports can be specified in the URI if they are in use.

- `--base-dn`: The base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.

  ⚠ **Important!** Anonymous binding is supported for the array management directory service. If configuring anonymous binding, neither the bind user nor the bind password should be specified.

- `--bind-user`: The username used to bind to and query the directory. For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`. For Active Directory servers, you may instead choose to enter the username - often referred to as the sAMAccountName or user login name - of the account that is used to perform directory lookups. The username cannot contain the characters `" [] : ; | = + * ? < > / \`, and cannot exceed 104 characters in length.

- `--bind-password`: The password for the bind user account.

- `--ca-certificate` or `--ca-certificate-group`: (Optional) A CA certificate or certificate group to verify the authenticity of configured LDAP servers to establish a TLS (Transport Layer Security) communication for directory services.

- `--user-login-attribute`: Defaults to the sAMAccountName for Active Directory, or UID for other server types.

- `--user-object-class`: Defaults to User for AD, posixAccount or shadowAccount for OpenLDAP depending on the server's schema, posixAccount for posix-compliant servers, or person for all other server types.

For example:

```
pureds setattr --uri ldaps://[2000:0db1:54a1::ae25:7b1f:0505:1112] --base dn
DC=mycompany,DC=com management
```

In this example:

- the URI is set to the IPv6 address `2000:0db1:54a1::ae25:7b1f:0505:1112`
- the scheme is set to `ldaps://` to enable SSL
- the base DN is set to be the correct domain components (DC)
- anonymous binding is used

2 Test the configuration using the **`pureds test management`** command. If the test fails, update the configuration information.

3 If the test passes, enable the directory service using the **`pureds enable management`** command.

4 Configure the role(s) you wish to assign the service by issuing the **`pureds role setattr --group --group-base ROLE`** command, where:

- `--group`: The common name (CN) of the directory group. The name should be the Common Name of the group without the "CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, "pureusers,OU=PureStorage", where pureusers is the common name of the directory service group.
- `--group-base`: The organizational unit (OU) to the configured groups in the directory tree. The Group Base consists of OUs that, when combined with the base DN attribute and the configured group CNs, complete the full distinguished name of each group. The group base should specify "OU=" for each OU and multiple OUs should be separated by commas. The order of OUs should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: "OU=PureGroups,OU=SANManagers"
- `ROLE`: The role assigned to the service. See *Roles* for details about FlashBlade roles.

For example:

```
pureds role setattr --group admins --group-base OU=Marketing storage_admin
```

In this example the role `storage_admin` is set as belonging the group `admins` in the group base `Marketing`.

NFS Directory Service

The NFS directory service is used for resolving addresses for netgroups, looking up the extended group membership of NFS data path clients, and looking up name/ID mappings.

Active Directory, OpenLDAP, and Oracle Unified Directory servers (that are RFC-2307 or RFC-2307bis compliant) are supported.

When configuring NFS directory service, you can optionally specify a bind user and bind password. If not specified, FlashBlade attempts to bind anonymously (without authenticating) with the configured LDAP servers. Note that if the configured LDAP servers do not grant read permissions on the needed pathways to unauthenticated machines, the anonymous bind will fail.

If a user belongs to more than 16 groups then you must configure and enable the NFS directory service in order to obtain proper user permissions. If NFS directory service is enabled, FlashBlade attempts to look up user IDs and determine the groups to which they belong.

NIS Domains and NIS Servers

NIS Domains

NIS domains can be used with quotas and normal group lookup.

NIS servers can be configured instead of LDAP servers. When NIS servers are configured, FlashBlade communicates with them over the NIS protocol in order to convert between user IDs and user names, as well as to convert between group IDs and group names.

NIS Servers

NIS servers can also be used with quotas and normal group lookup.

NIS servers can be configured instead of, or in addition to, LDAP servers. The servers can be specified as either host names or IP addresses.

NIS servers are required in order to configure the NFS directory service to use NIS.

Configuring the NFS Directory Service with LDAP

Note: Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The NFS directory service enables file access for users who are defined in a directory server.

Configuring the NFS directory service is a one-time process and is only required if you are integrating the FlashBlade array with a directory service for protocol support.

To export a file system over NFS where a user might belong to more than 16 groups, the FlashBlade array must be integrated with a directory service.

To configure the NFS directory service with LDAP for protocol support, perform the following:

1. Verify that the directory server is accessible through a management interface.
2. Verify that the DNS settings can be used to resolve the directory server domain and server.
3. From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
4. In the **Directory Service** panel, select **NFS**.

5 Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.

6 Under the **Service Name** field, select **LDAP**.

7 In the **URIs** field, type the comma-separated list of up to 30 URIs of the directory servers.

Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

If specifying a domain name, it should be resolvable by the configured DNS servers.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from `0` to `255` representing a group of 8 bits.

For IPv6, specify the IP address in the form `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`) . Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

Optionally specify a port. Append the port number after the end of the entire address. Default ports are `389` for LDAP, and `636` for LDAPS. Non-standard ports can be specified in the URI if they are in use.

8 In the **Base DN** field, type the base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.

⚠ **Important!** Anonymous binding is supported for the NFS directory service. If configuring anonymous binding, neither the bind user nor the bind password should be specified.

9 In the **Bind User** field, type the username used to bind to and query the directory. For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`. For Active Directory servers, you may instead choose to enter the username - often referred to as the sAMAccountName or user

login name - of the account that is used to perform directory lookups. The username cannot contain the characters `" [] : ; | = + * ? < > / \`, and cannot exceed 104 characters in length.

10 In the **Bind Password** field, type the password for the bind user account.

11 (Optional) Select a CA certificate or certificate group to verify the authenticity of configured LDAP servers to establish a TLS (Transport Layer Security) communication for directory services.

12 Click **Test** to verify the configuration information. The **Test <Directory Service> Configuration** pop-up window appears, displaying the output of the test. During the directory service test, Purity//FB tests the directory service configuration to verify that the URI can be resolved and that the directory service can successfully bind and query the tree using the bind user credentials.

If the test fails, update the configuration information and retest.

Alternatively, you may save the configuration and enable the directory service at a later time.

13 If the test passes, set the **Enabled** toggle to enable (blue) to enable the directory service.

14 Click **Save**.

Configuring the NFS Directory Service with LDAP using the CLI

Note: Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The NFS directory service with LDAP can also be configured using the **`pureds setattr nfs`** command. For full details about the **`pureds`** command, see the *FlashBlade CLI Reference*.

To configure the NFS directory service using the CLI, perform the following:

1 Run the **`puredns list`** command to verify that DNS settings can be used to resolve the directroy server domain and server.

2 Issue the **`pureds setattr nfs`** command with the following arguments:

- `--uri`: Up to 30 URIs can be specified in a comma-separated list.

  Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

  We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

If specifying a domain name, it should be resolvable by the configured DNS servers.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from `0` to `255` representing a group of 8 bits.

For IPv6, specify the IP address in the form `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`) . Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

Optionally specify a port. Append the port number after the end of the entire address. Default ports are `389` for LDAP, and `636` for LDAPS. Non-standard ports can be specified in the URI if they are in use.

- `--base-dn`: The base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.

  ⚠ **Important!** Anonymous binding is supported. If configuring anonymous binding, neither the bind user nor the bind password should be specified.

- `--bind-user`: The username used to bind to and query the directory. For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`. For Active Directory servers, you may instead choose to enter the username - often referred to as the sAMAccountName or user login name - of the account that is used to perform directory lookups. The username cannot contain the characters `" [] : ; | = + * ? < > / \`, and cannot exceed 104 characters in length.

- `--bind-password`: The password for the bind user account.

- `--ca-certificate` or `--ca-certificate-group`: (Optional) A CA certificate or certificate group to verify the authenticity of configured LDAP servers to establish a TLS (Transport Layer Security) communication for directory services.

- `--user-login-attribute`: Defaults to the sAMAccountName for Active Directory, or UID for other server types.

- `--user-object-class`: Defaults to User for AD, posixAccount or shadowAccount for OpenLDAP depending on the server's schema, posixAccount for posix-compliant servers, or person for all other server types.

For example:

```
pureds setattr --uri ldap://ds.mycompany.com --base dn DC=mycompany,DC=com nfs
```

In this example:

- the URI is set to `ds.company.com`, where `ds` is the host name in the domain `mycompany.com`
- the scheme is set to unencrypted `ldap://`
- the base DN is set to be the correct domain components (DC)
- anonymous binding is used

3 Test the configuration using the **`pureds test nfs`** command. If the test fails, update the configuration information.

4 If the test passes, enable the directory service using the **`pureds enable nfs`** command.

Removing an NFS Directory Service with LDAP Configuration

If you previously configured the NFS directory service with LDAP and wish to swtich to NIS, you must first remove the LDAP configuration.

To remove an LDAP configuration, perform the following:

1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.

2 In the **Directory Service** panel, select **NFS**.

3 Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.

4 Click the **Enabled** toggle (grey) to disable NFS directory service.

5 Clear all fields.

6 Click **Save**.

The LDAP configuration is now removed. If you wish to configure NFS directory service with NIS, see *Configuring the NFS Directory Service with NIS*.

Removing an NFS Directory Service with LDAP Configuration using the CLI

The NFS directory service with LDAP configuration can alternately be removed using the CLI.

Remove the NFS with LDAP is configuration using the **`pureds setattr nfs`** command. For full details about the **`pureds`** command, see the *FlashBlade CLI Reference*.

To remove the NFS with LDAP configuration using the CLI, perform the following:

1. Issue the **`pureds disable nfs`** command to disable NFS.
2. Issue the **`pureds setattr --uri "" --base-dn "" --bind-user "" --bind-user "" --bind-password "" nfs`** command. Entering `""` sets the values to null. When prompted for passwords, press ENTER to input empty values.
3. Issue the **`pureds list nfs`** command and verify that the LDAP configuration is cleared.

The LDAP configuration is now removed. If you wish to configure NFS directory service with NIS, see *Configuring the NFS DIrectory Service with NIS using the CLI*.

Configuring the NFS Directory Service with NIS

The NFS directory service enables file access for users who are defined in a directory server.

Configuring the NFS directory service is a one-time process and is only required if you are integrating the FlashBlade array with a directory service for protocol support.

To export a file system over NFS where a user might belong to more than 16 groups, the FlashBlade array must be integrated with a directory service. For NFS with NIS, clear any existing LDAP configurations before specifying the NIS servers and NIS domain.

To configure the NFS directory service with NIS for protocol support, perform the following:

1. Verify that the directory server is accessible through a management interface.
2. Verify that the DNS settings can be used to resolve the directory server domain and server.
3. From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
4. In the **Directory Service** panel, select **NFS**.
5. Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.
6. Under the **Service Name** field, select **NIS**.
7. In the **NIS Servers** field, type the comma-separated list of server host names or IP addresses. For example, `myserver.example.com,188.121.4.56`.
8. In the **NIS Domain** field, type the domain name. For example, `yp-example-domain`.
9. Click **Test** to verify the configuration information. The **Test <Directory Service> Configuration** pop-up window appears, displaying the output of the test. During the directory service test, Purity//FB

tests the directory service configuration to verify that the URI can be resolved and that the directory service can successfully bind and query the tree using the bind user credentials.

If the test fails, update the configuration information and retest.

Alternatively, you may save the configuration and enable the directory service at a later time.

10 If the test passes, set the **Enabled** toggle to enable (blue) to enable the directory service.

11 Click **Save**.

Configuring the NFS DIrectory Service with NIS using the CLI

NFS directory service with NIS can alternately be configured using the CLI.

The NFS directory service with NIS is configured using the **`pureds setattr nfs`** command. For full details about the **`pureds`** command, see the *FlashBlade CLI Reference*.

To configure the NFS directory service using the CLI, perform the following:

1 Issue the **`pureds setattr nfs`** command with the following arguments:

- `--nis-servers`: A comma-separated list of server host names or IP addresses. For example, `myserver.example.com,188.121.4.56`.
- `--nis-domain`: The NIS domain name. For example, `yp-example-domain`.

For example:

```
pureds setattr --nis-servers my-server.com,188.121.25.4 --nis-domain my-nis-domain nfs
```

In this example:

- the NIS domain is set to `my-nis-domain`
- the NIS servers are set to `my-server.com` and `188.121.25.4`

2 Test the configuration using the **`pureds test nfs`** command. If the test fails, update the configuration information.

3 If the test passes, enable the directory service using the **`pureds enable nfs`** command.

Removing an NFS Directory Service with NIS Configuration

If you previously configured the NFS directory service with NIS and wish to swtich to LDAP, you must first remove the NIS configuration.

To remove a NIS configuration, perform the following:

1. From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
2. In the **Directory Service** panel, select **NFS**.
3. Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.
4. Click the **Enabled** toggle (grey) to disable NFS directory service.
5. Clear the **NIS Servers** and **NIS Domain** fields.
6. Click **Save**.

The NIS configuration is now removed. If you wish to configure NFS directory service with LDAP, see *Configuring the NFS Directory Service with LDAP*.

Removing an NFS Directory Service with NIS Configuration using the CLI

The NFS directory service with NIS configuration can alternately be removed using the CLI.

Remove the NFS with NIS is configuration using the **`pureds setattr nfs`** command. For full details about the **`pureds`** command, see the *FlashBlade CLI Reference*.

To remove the NFS with NIS configuration using the CLI, perform the following:

1. Issue the **`pureds disable nfs`** command to disable NFS.
2. Issue the **`pureds setattr nfs --nis-server "" --nis-domain ""`** command. Entering `""` sets the values to null.
3. Issue the **`pureds list --nfs-specific nfs`** command and verify that the previously configured NIS server and domain are cleared.

The NIS configuration is now removed. If you wish to configure NFS directory service with LDAP, see *Configuring the NFS Directory Service with LDAP using the CLI*.

SMB Directory Service

The SMB directory service is compatible with only Active Directory (AD) servers.

When joining a domain, a machine account is created with the AD domain. FlashBlade then issues a keytab to that machine account, which is used for future communication with that domain. The default Organizational Unit (OU) where the machine account is created within the AD domain is "Computers".

Because SMB authentication is inherently a Kerberos authentication setup, Kerberos ports must be open for the SMB directory service configuration to function properly.

The following issues can result in loss of SMB functionality and file access:

- If domain controllers delete or otherwise invalidate FlashBlade's keytab or machine account, SMB will stop functioning.
- SMB clients do not have IDs; they have names. FlashBlade resolves the client's name to an ID using the array's SMB mode and the configured SMB directory service. If the SMB directory service is unhealthy or disabled, SMB clients will lose all file access because the name/ID cannot be resolved.

Pure Storage, Inc.
X: @purestorage

2555 Augustine Drive, 1st Floor
Santa Clara, CA 95054
800-379-7873

Sales: sales@purestorage.com
Support: support@purestorage.com
Media: pr@purestorage.com
General: info@purestorage.com