

FlashBlade User Guide

Version 4.3.6

Copyright Statement

© 2024 Pure Storage® (“Pure”), Portworx® and its associated trademarks can be found [here](#) and its virtual patent marking program can be found [here](#). Third party names may be trademarks of their respective owners.

The Pure Storage products and programs described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc. 2555 Augustine Dr, Santa Clara, CA 95054 <http://www.purestorage.com>

Direct comments to DocumentFeedback@purestorage.com.

Table of Contents

- Chapter 1 About this Guide..... 11**
 - Organization of this Guide 11
 - A Note on Format and Content..... 12
 - Documentation and Support..... 12
- Chapter 2 Purity//FB Overview..... 13**
 - Purity//FB Conventions..... 13
 - Object Names.....13
 - Array and File System Sizes..... 13
 - Time Zones..... 14
 - IP Addresses..... 14
 - SMB Port Assignments.....15
 - Management Network Ports.....15
 - Firewall Ports..... 15
 - Data Network Ports.....16
 - Pure File SafeMode.....16
 - Pure Object SafeMode.....17
 - About NFSv4.1 Support.....17
 - Top Level Mounting.....18
 - Troubleshooting..... 19
- Chapter 3 Purity//FB GUI Overview..... 20**
 - Purity//FB GUI Navigation.....21
 - Purity//FB GUI Login.....24
 - Logging in to the Purity//FB GUI.....25
 - Viewing and Accepting the End User Agreement.....26
- Chapter 4 The Purity//FB GUI Dashboard Page..... 27**
 - Capacity..... 28
 - Performance.....28
 - Note about the Performance Charts..... 30
 - Recent Alerts.....31
 - Replication Bandwidth.....31
 - Hardware Health.....32
- Chapter 5 The Purity//FB GUI Storage Page.....33**

Array.....	33
Connecting Arrays.....	35
Editing a Connected Array.....	37
Disconnecting a Connected Array.....	38
Connecting an S3 Target.....	38
Updating an S3 Target.....	39
Disconnecting an S3 Target.....	39
Bandwidth Throttling.....	39
File Systems.....	41
About Creating and Exporting File Systems.....	45
File System Size.....	47
File System Usage Limits.....	47
User and Group Usage Limits.....	48
File System Snapshots.....	49
ACLs.....	49
Multi-Domain Trust.....	52
NFS and File Locking.....	54
Fast Remove.....	54
File Replication Failover, Reprotect, and Failback.....	55
Creating and Exporting a File System.....	56
Creating a File System Snapshot.....	65
Enabling and Disabling Fast Remove.....	72
Promoting and Demoting File Systems.....	72
Enabling and Disabling File System Writes.....	73
Viewing File System Performance.....	73
Destroying a File System.....	75
Recovering a File System.....	76
Eradicating a File System.....	77
Performing File Replication Failover and Reprotect.....	77
Performing File Replication Failback.....	79
Object Store.....	81
Object Store Usage Limits.....	85
Object Store Account Information.....	86
Creating an Account.....	87
Editing an Account.....	88
Deleting an Account.....	89
Users.....	89

Buckets.....	96
Chapter 6 The Purity//FB GUI Protection Page.....	117
Snapshots.....	117
Destroying File System Snapshots.....	122
Restoring a File System to a Snapshot	123
Removing Snapshot Policies from Snapshots.....	125
Recovering File System Snapshots.....	126
Eradicating File System Snapshots.....	127
File Replication.....	128
Creating a File System Replica Link.....	129
Adding a Policy to a File Replica Link.....	130
Removing a Policy from a File Replica Link.....	130
Creating Fan Out File System Replica Links.....	131
Deleting a File System Replica Link.....	131
Object Replication.....	132
Object Replication Write Conflict Continuation.....	135
Creating a Bucket Replica Link.....	136
Editing a Bucket Replica Link.....	136
Pausing a Bucket Replica Link.....	137
Resuming a Bucket Replica Link.....	137
Deleting a Bucket Replica Link.....	138
Creating Remote Credentials.....	138
Editing Remote Credentials.....	139
Deleting Remote Credentials.....	139
Chapter 7 Replication and Recovery.....	140
File System Replication Overview.....	140
Setting Up File System Replication Using the GUI.....	144
Creating Replication Network Interfaces.....	145
Creating a Connection Key.....	145
Connecting the Source and Target Arrays.....	145
Creating a Policy.....	147
Creating a File System Replica Link.....	147
Setting Up File System Replication Using the CLI.....	149
Creating Replication Network Interfaces.....	149
Creating a Connection Key.....	149
Connecting the Source and Target Arrays.....	149

Creating a Policy.....	151
Creating a File System Replica Link.....	151
Performing File System Replication Failover Using the GUI.....	152
Performing File System Replication Failback Using the GUI.....	154
Performing File System Replication Failover Using the CLI.....	156
Performing File System Replication Failback Using the CLI.....	158
Performing File System Fan Out Replication Failover Using the GUI.....	159
Performing File System Fan Out Replication Failback Using the GUI.....	161
Performing File System Fan Out Replication Failover Using the CLI.....	163
Performing File System Fan Out Replication Failback Using the CLI.....	165
Object Replication Overview.....	166
Setting Up Array to Array Object Replication Using the GUI.....	169
Creating Replication Network Interfaces.....	169
Creating a Connection Key.....	170
Connecting the Source and Target Arrays.....	170
Creating Remote Credentials.....	171
Creating an Bucket Replica Link.....	172
Enabling Bucket Versioning.....	172
Setting Up Array to Array Object Replication Using the CLI.....	173
Creating Replication Network Interfaces.....	173
Creating a Connection Key.....	173
Connecting the Source and Target Arrays.....	174
Creating Remote Credentials.....	175
Creating a Bucket Replica Link.....	175
Enabling Bucket Versioning.....	176
Setting Up FlashBlade to Amazon S3 Object Replication Using the GUI.....	177
Creating a Replication Network Interface.....	177
Connecting the FlashBlade Array to the S3 Target.....	177
Creating Remote Credentials.....	178
Creating an Bucket Replica Link.....	178
Enabling Bucket Versioning.....	179
Setting Up FlashBlade to Amazon S3 Object Replication Using the CLI.....	179
Creating a Replication Network Interface.....	179
Connecting the FlashBlade Array to the S3 Target.....	180
Creating Remote Credentials.....	180
Creating a Bucket Replica Link.....	180
Enabling Bucket Versioning.....	181

Chapter 8 The Purity//FB GUI Policies Page.....	182
NFS Export Policies.....	183
NFS Export Rules.....	186
Creating NFS Export Policies and Rules.....	192
Attaching File Systems to an NFS Export Policy.....	193
Moving File Systems to Another Policy.....	193
Removing File Systems from an NFS Export Policy.....	194
Adding Additional NFS Export Rules to NFS Export Policies.....	195
Editing NFS Export Rule Clients.....	196
Editing NFS Export Rules.....	196
Reordering NFS Export Rules.....	197
Removing NFS Export Rules.....	197
Renaming NFS Export Policies.....	198
Enabling and Disabling NFS Export Policies.....	198
Deleting NFS Export Policies.....	199
Object Store Access Policies.....	200
About Object Store Access Policies.....	200
Rules: Actions, Resources, and Conditions.....	203
Viewing Object Store Policy Details.....	208
Creating Object Store Access Policies.....	209
Adding Object Store Access Policy Rules.....	210
Editing Object Store Access Policy Rules.....	212
Copying Object Store Access Policy Rules to Other Policies.....	213
Removing Object Store Access Policy Rules.....	214
Adding Users to an Object Store Access Policy.....	214
Removing Users from an Object Store Access Policy.....	215
Deleting Object Store Policies.....	215
SMB Client and Share Policies.....	216
SMB Client Policies.....	217
SMB Share Policies.....	226
File System Snapshot Policies.....	235
Viewing Snapshot Policy Details.....	236
Creating Snapshot Policies.....	236
Adding Snapshot Policy Rules.....	237
Adding Multiple Snapshot Policy Rules.....	238
Editing Snapshot Policy Rules.....	239
Removing Snapshot Policy Rules.....	239

Adding File Systems to a Snapshot Policy.....	240
Removing File Systems from a Snapshot Policy.....	240
Adding Replica Links to a Snapshot Policy.....	241
Removing Replica Links from a Snapshot Policy.....	241
Enabling and Disabling Snapshot Policies.....	242
Removing a Snapshot from a Policy.....	242
Removing a Policy from a Snapshot.....	243
Deleting Snapshot Policies.....	243
Chapter 9 The Purity//FB GUI Analysis Page.....	245
Performance.....	245
The Performance Chart.....	245
Displaying Protocol-Specific Metrics.....	247
Note about the Performance Charts.....	249
Displaying NFS Client, User, and Group Performance Statistics.....	249
Capacity.....	251
Replication.....	253
Network.....	254
Viewing Network Connector Statistics.....	257
Displaying a Different Time Range	259
Chapter 10 The Purity//FB GUI Health Page.....	260
Hardware.....	260
FlashBlade//S System Resiliency.....	262
Alerts.....	262
Flagging Alert Messages.....	264
Unflagging Alert Messages.....	264
Chapter 11 The Purity//FB GUI Settings Page.....	266
Side Navigation Bar.....	267
Working with System Settings.....	268
Alert Watchers.....	269
Alert Routing.....	272
File System User/Group Quota Notifications.....	273
Array Details.....	274
SNMP.....	277
Syslog Server.....	289
NTP Servers.....	294
Support.....	295

Working with Network Settings.....	300
Connectors.....	300
Diagnostics.....	302
DNS Settings.....	303
Link Aggregation Groups.....	304
Object Store Virtual Hosts.....	306
Subnets and Interfaces.....	311
Working with Security Settings.....	316
API Clients.....	317
Viewing the Audit Trail.....	319
Array Certificates.....	320
External Certificates.....	322
Certificate Groups.....	324
Viewing the SMB Mode.....	326
Active Directory Accounts.....	327
Directory Service.....	329
Kerberos Keytabs.....	331
Rapid Data Locking.....	334
Session Log.....	342
Users.....	343
Chapter 12 Choosing an Active Directory Account or Directory Services Configuration..	348
Active Directory Overview.....	352
Active Directory Networking Requirements.....	352
Computer Account Naming and Credentials.....	354
Minimum Permissions for Joining an Active Directory Domain.....	355
Joining an Organizational Unit.....	356
Joining with an Existing Computer Account.....	356
Service Principal Names.....	357
Directory and Kerberos Servers.....	358
User and Group Name Mapping for Active Directory and LDAP and NIS.....	359
Creating an Active Directory Account.....	359
Joining Active Directory with an Existing Computer Account.....	362
Testing an Active Directory Configuration.....	364
Editing an Active Directory Account.....	366
Rotating Keytabs.....	368
Importing a Keytab File.....	369
Exporting a Keytab File.....	370

Deleting a Keytab File.....	371
Deleting an Active Directory Account.....	372
Directory Services Overview.....	373
NFS and SMB Nested User Groups.....	374
Array Management Directory Service.....	375
Configuring the Array Management Directory Service.....	376
NFS Directory Service.....	381
Configuring the NFS Directory Service with LDAP.....	382
Configuring the NFS Directory Service with NIS.....	386
SMB Directory Service.....	389
Configuring the SMB Directory Service.....	390

Chapter 1

About this Guide

The Pure Storage FlashBlade® User Guide provides information about the first-generation FlashBlade, FlashBlade//E, and FlashBlade//S, their features, and how to use them. This guide is written for array administrators who view and manage the Pure Storage first-generation FlashBlade, FlashBlade//E, and FlashBlade//S storage systems. First-generation FlashBlade, FlashBlade//E, and FlashBlade//S 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.

The FlashBlade//S consists of an //S200 and //S500 model. The //S200 is a system built for ultimate efficiency and the //S500 is for extreme performance. As the customer requirements and workloads evolve overtime, they have the ability to non-disruptively upgrade from an //S200 to an //S500. Contact Pure Technical Services to inquire about upgrading.

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.

[Replication and Recovery](#) - Provides step-by-step instructions on how to set up file and object replication, as well as how to perform failover and failback, using the Purity//FB GUI and CLI.

[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 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 (Gi), gigabyte (GB), tibibyte (Ti), terabyte (TB), pebibyte (Pi), 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.

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 Technical Services.

Pure Object SafeMode

 **Note:** If enabling Pure Object SafeMode and Object Replication, arrange for Object SafeMode to be enabled on all arrays involved in replication before enabling replication. Contact Pure Technical Services for assistance.

 **Important!** When Pure Object SafeMode is enabled, the NTP server(s) cannot be modified.

Pure Object SafeMode is a data protection mechanism for S3 that is enabled by Pure Technical Services. Object SafeMode is designed to protect the system from accidental changes that can result in data loss. With Object SafeMode:

- Pure Technical services can configure an object modification protection period preventing objects from being deleted or overwritten if their age is less than the configured retention period.
- Buckets can be destroyed with assistance from Pure Technical Services. If a destroyed bucket is not recovered within the modification protection period, then that bucket will be eradicated.
- Destroyed buckets can only be immediately eradicated by Pure Technical Services.
- Buckets cannot be deleted through the S3 API.
- Objects that have expired per configured object lifecycle rules are still protected within the configured object modification protection period.

Object SafeMode can be enabled with object replication on both the source and target. Note that Object SafeMode can be enabled when all existing buckets are in the versioning state NONE and there is no in- or outbound replica link.

For more information, contact Pure Technical Services.

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. An **Object Store** tab is provided to quickly access object store accounts, buckets, and users.

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. Object store access policies and snapshot policies can be viewed and managed from the **Object Store Access** and **Snapshot** tabs.

Analysis

Displays historical array information, including space and I/O performance metrics, from various viewpoints. Latency, IOPs, and bandwidth data for file systems and the object store, as well as clients and NFS users and groups, can be viewed from the **File Systems**, **Object Store**, **Clients**, **Groups**, and **Users** tab. Additionally, you can view charts for capacity and file and object replication.

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. Additionally, 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.

File Systems										
General Space 1-10 of 65 < > + ⋮										
Name▲	Source Location	Source Name	Size	Used	Hard Limit	Created	Protocols	Replica State	Writable	
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

The image shows the Purity//FB GUI login screen. At the top left is the Pure Storage logo. Below it are two input fields: 'Username' and 'Password'. Below these fields is a large orange 'Log In' button. At the bottom of the login area is a small blue link labeled 'Terms'.

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.

Note: The **Hardware Panel** in the image below depicts a single-chassis FlashBlade//S array. The graphic displayed in this panel differs for single and multi-chassis first-generation FlashBlade arrays, single and multi-chassis FlashBlade//E arrays, and single and multi-chassis FlashBlade//S arrays.

Figure 4.1 Dashboard



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


- **Capacity**
- **Performance Charts**
- **Recent Alerts**
- **Replication Bandwidth** (only for replication setups)

- **Hardware Health**

Capacity

The **Capacity** panel displays array size and storage consumption details. All capacity values are rounded to two decimal places.

Figure 4.2 Dashboard - Capacity Panel



The capacity wheel displays the percentage of array space occupied by file system 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.
- **Empty:** Unused space.
- **Used:** Amount of space that the written data occupies on the array after reduction via data compression.
- **Capacity:** Total usable 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).

Performance

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

Figure 4.3 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.

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 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**.

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


In the **Analysis** page:

- 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 1 year of historical data.
- The performance charts can be filtered to display metrics by protocol.

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 4.4 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.

Replication Bandwidth

The **Replication Bandwidth** panel provides a graph that displays the total replication bandwidth. Note that this panel appears only if replication has been set up.

Figure 4.5 Dashboard - Replication Bandwidth Panel




Clicking the graph opens the **Analysis > Replication** page where finer detail about bandwidth usage can be viewed (see [Replication](#)).

Hardware Health

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

Figure 4.6 Hardware Health Panel - FlashBlade//E



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 file systems and object store accounts, which contain users and buckets.

Accessing a file system or bucket 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, a **FlashBlade Array Connections** panel, and an **S3 Target Connections** 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 System Snapshots** count opens the **Protection > Snapshots** tab in the GUI.

Note that the number of objects displayed under **Object Store Objects** is approximate. The units used for the approximate object count are **K** (thousands), **M** (millions), **B** (billions), and **T** (trillions).

Array attributes such as array time and the array name are configured through the **Settings > System** page.

FlashBlade Array Connections

The **FlashBlade Array Connections** panel displays all arrays connected to the current FlashBlade array for replication.

Figure 5.3 FlashBlade Array Connections Panel


A connection must be established between two arrays in order for data replication to occur. FlashBlade supports replication, which allows data to be replicated from one array to another. When two arrays are connected for asynchronous replication, the array where data is being transferred from is called the local (source) array, and the array where data is being transferred to is called the remote (target) array. FlashBlade supports up to three array connections; however, only one connection between two arrays can be used for file replication. All three connections can be used for object replication.

Details about the array connection are provided including the current connection status, Purity software version, management and replication IP addresses, encryption status, applied CA certificate group (if any), and if bandwidth throttling is enabled. If you click **Paths**, the connected array and all connection paths (source and destination IP address) are displayed.

If the array connection is not healthy, the **Paths** view can be useful in determining which network paths are having problems.

S3 Target Connections

The **S3 Target Connections** panel displays all S3 targets connected to the current FlashBlade array for object replication.

Figure 5.4 S3 Target Connections Panel

S3 Target Connections ⓘ					+	
Name	Status	Address	CA Certificate Group			
aws-shuy1	connected	s3.amazonaws.com	_default_replication_certs	🔍	✕	
aws-shuy1-2	connected	aws.amazon.com	_default_replication_certs	🔍	✕	
aws-shuy1-ca	connected	s3-us-west-1.amazonaws.com	_default_replication_certs	🔍	✕	
aws-shuy1-w2	connected	s3-us-west-2.amazonaws.com	_default_replication_certs	🔍	✕	

The names of connected S3 targets, connection status, connection address, and applied CA certificate groups are displayed.

In addition to displaying details about connected S3 targets, this panel also allows you to rename existing S3 targets and connect new S3 targets.

Connecting Arrays

- Replication network interfaces on the source and target arrays.
- A network path from the source array to the target array's management vip.

 **Note:** Although multiple array connections are supported, file replication is only allowed between two arrays.

 **Note:** When connecting arrays, it is recommended to create the array connection from the array running the most current Purity version to the array running an older Purity version.

Connecting two arrays is performed in the **FlashBlade Array Connections** panel of the **Array** page.

 **Note:** Data can be optionally encrypted. However, file replication is not supported over an encrypted connection.

 **Important!** When connecting the source array to the target array, a secure connection is required even if encryption is not enabled. Certificates with small key sizes, weak signature algorithms, or that only support weak TLS algorithms are not supported.

 **Important!** It is highly recommended that you add the `pure:policy/full-access` access policy to the user (whose credentials are used for replication) on the target cluster, otherwise replication may fail due to denied access. See [Adding Access Policies](#) for instructions on how to add access policies to users.

To connect two arrays, perform the following:

- 1 From the **Storage > Array** page, click the add (+) button in the **FlashBlade Array Connections** panel. The **Connect FlashBlade Array** pop-up window appears.
- 2 Enter the target array's hostname or management address (unless using NAT) in the **Management Address** field. The address can be located from the **Subnets** table by navigating to the **Settings > Network** page.
- 3 Enter a connection key for the target array in the **Connection Key** field.
 - a On the target array, create the connection key (see [Creating a Connection Key](#)).
 - b Copy the new connection key.
 - c On the source array, paste the connection key in the **Connection Key** field in the **Connect FlashBlade Array** pop-up window.
- 4 The source and target arrays each need a replication network interface. If not already created, create the replication network interfaces on the source and target arrays (see [Creating a Network Interface](#)). The replication address is automatically populated on both the source and target cluster unless manually specified. If manually specified, then the target cluster will not have a replication address populated and will need to manually entered. It is recommended to allow the arrays to automatically discover the replication address, unless using NAT. Enter the target array's replication network address in the **Replication Address** field.
- 5 (Optional) Enable encryption by setting the **Encrypted** toggle to on. If set, the `_default_replication_certs` CA certificate group is applied. Before toggling, make sure to import the CA certificate to the source array if the target array's array-certificate is signed by a CA and add it to the certificate group. Note that encryption is set on the source array only. If enabled, the encryption setting displays as enabled on the target array.

Warning: Use of CA certificates to establish trust and secure communication for object replication is the recommended configuration. If the target array's array-certificate is self-signed, directly importing the target array certificate to the source and adding it to the certificate group will not provide as strict validation around aspects of the peer, such as certificate expiry.

- 6 Click **Connect**.
- 7 If replication address was specified on the source cluster while creating the array connection in step 4, on the target array, enter the source array's replication address.
 - a On the target array, in the **FlashBlade Array Connections** panel on the **Storage > Array** page, click the edit button at the end of the row displaying the connected source array. The **Edit Connected Array** pop-up window appears.
 - b Enter the replication address for the source array in the **Replication Address** field and click **Save**.

Creating a Connection Key

In order to connect two arrays for replication, a connection key must be created on the target array to be used on the source array during array connection.

Connection keys are created from the **FlashBlade Array Connections** panel of the **Array** page.

To create a connection key, perform the following:

- 1 On the target array, navigate to the **Storage > Array** page.
- 2 In the **FlashBlade Array Connections** panel, click **More Options > Create Connection Key**. The **Connection Key** pop-up window appears displaying the new connection key.

Editing a Connected Array

A connected array must be edited if its management or replication address information has changed, or if you wish to enable or disable encryption (for object replication).

Editing a connected array is performed in the **FlashBlade Array Connections** panel of the **Array** page.

To edit a connected array, perform the following:

- 1 On the source array, navigate to the **Storage > Array** page.
- 2 In the **FlashBlade Array Connections** panel, click the **More Options** button located at the end of the row of the array you wish to edit and select **Edit Connected Array**. The **Edit Connected Array** pop-up window appears displaying the new connection key.

- 3 Enter a new management and/or replication address as needed.
- 4 Enable (blue) or disable (gray) encryption.
- 5 Click **Save**.

Disconnecting a Connected Array

Disconnecting an array is performed in the **FlashBlade Array Connections** panel of the **Array** page.

An array cannot be disconnected while in use by file system replica links. You must delete all file system replica links to the remote array. Once file system replic links are deleted, the array can then be disconnected. Note that replica link propagation to the remote array is asynchronous and can be delayed.

To disconnect a connected target array, perform the following:

- 1 On the source array, navigate to the **Storage > Array** page.
- 2 In the **FlashBlade Array Connections** panel, click the **More Options** button located at the end of the row of the array you wish to disconnect and select **Disconnect**. The **Disconnect FlashBlade Array** pop-up window appears.
- 3 Click **Disconnect**.

Connecting an S3 Target

A FlashBlade array can be connected to an S3 target for object replication.

Connecting a FlashBlade to an S3 target is performed in the **S3 Target Connections** panel of the **Array** page.

To connect an S3 target, perform the following:

- 1 From the **Storage > Array** page, click the add (+) button in the **S3 Target Connections** panel. The **Connect S3 Target** pop-up window appears.
- 2 Enter the name of the S3 target and its address.
- 3 Click **Connect**.

Updating an S3 Target

The names and addresses S3 targets can be updated in the **S3 Target Connections** panel of the **Array** page.

To update an S3 target, perform the following:

- 1 On the source array, navigate to the **Storage > Array** page.
- 2 In the **S3 Target Connections** panel, click the **Edit** button located at the end of the row of the S3 target you wish to update. The **Edit S3 Target** pop-up window appears.
- 3 Enter a new name and/or new address.
- 4 Click **Save**.

Disconnecting an S3 Target

Disconnecting an S3 target from a FlashBlade array is performed in the **S3 Target Connections** panel of the **Array** page.

To disconnect an S3 target, perform the following:

- 1 On the source array, navigate to the **Storage > Array** page.
- 2 In the **S3 Target Connections** panel, click the Disconnect (x) button located at the end of the row of the S3 target you wish to disconnect. The **Disconnect S3 Target** pop-up window appears.
- 3 Click **Disconnect**.

Bandwidth Throttling

Bandwidth throttling for file replication manages data transfer capacity of the FlashBlade. By limiting the bandwidth for replication operations during peak times it will reduce the competition for bandwidth with priority operations, thus allowing priority operations to complete sooner.

Throttling can be set with a default limit and a narrower window limit. The default limit can control the file replication throttling 24 hours a day and the window limit can control a smaller range of time under a different throttle limit. For both the default and window throttle limit, the value entered is the maximum bandwidth threshold for outbound traffic in bytes. Once exceeded, bandwidth throttling occurs. Note that the bandwidth used by replication will be throttled such that the bandwidth achieved may be lower than or equal to the set threshold.

The window throttle limit is set similarly to the default throttle limit and can be a smaller or larger data transfer size than the default throttle limit. Once the window throttle limit is set, you can set a window time range which will supersede the default throttle limit for the specified time range. The window time range is a daily time range and this “window” is only open once per day. For example, you can set the window throttle limit to 10G and the window time range from 9pm-5am, and everyday between 9pm and 5am, file replication can transfer data at a rate of up to 10 GB/s, provided the network bandwidth is available. Replication between 5am to 9pm, outside of the 9pm to 5am range, will follow the default throttle limit.

Default and/or window bandwidth throttling status (**True** if enabled or **False** if disabled) is displayed in the **Throttled** column in the **FlashBlade Array Connections** panel. Hovering the mouse pointer over the column displays the configured default and window bandwidth throttling values.

Figure 5.5 Viewing Bandwidth Throttling Status

Name	Status	Version	Management Address	Replication Address	Encrypted	CA Certificate Group	Throttled
Oban	connected	3.3.0	-	10.64.128.79	False	-	True

Setting the Default Bandwidth Throttling Limit

The default bandwidth throttling limit controls the array’s overall replication throttling 24 hours a day

To set a default replication bandwidth throttling limit for the array, perform the following:

- 1 On the source array, navigate to the **Storage > Array** page.
- 2 In the **FlashBlade Array Connections** panel, click the **More Options > Edit Bandwidth Throttling** button located at the end of the row of the array to which you wish to apply a default bandwidth throttling limit. The **Edit Bandwidth Throttling** pop-up window appears.
- 3 Enable (blue) **Default Throttle**.
- 4 Enter a value in the **Limit** field for the maximum bandwidth threshold in bytes. Bandwidth will be throttled up to the entered value. The value must be specified in format n[M|G], where M is MB/s and G is GB/s. For example 5M, 10G, etc. The minimum limit that can be entered is 5 MB/s. The maximum limit is 28 GB/s. To set a limit greater than 28 GB/s, contact Pure Technical Services.
- 5 Click **Save**.

Setting the Window Bandwidth Throttling Limit

The window bandwidth throttling limit controls the array’s replication throttling for a smaller range of time within the array’s default bandwidth throttling limit.

To set a replication window bandwidth throttling limit for the array, perform the following:

- 1 On the source array, navigate to the **Storage > Array** page.
- 2 In the **FlashBlade Array Connections** panel, click the **More Options > Edit Bandwidth Throttling** button located at the end of the row of the array to which you wish to apply a window bandwidth throttling limit. The **Edit Bandwidth Throttling** pop-up window appears.
- 3 Enable (blue) **Window Throttle**.
- 4 Select the time range for the bandwidth throttling window.
- 5 Enter a value in the **Limit** field for the maximum bandwidth threshold in bytes. Bandwidth will be throttled up to the entered value. The value must be specified in format n[MIG], where M is MB/s and G is GB/s. For example 5M, 10G, etc. The minimum limit that can be entered is 5 MB/s. The maximum limit is 28 GB/s. To set a limit greater than 28 GB/s, contact Pure Technical Services.
- 6 Click **Save**.

File Systems

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

Figure 5.6 File Systems Tab

Storage									
Array File Systems Object Store									
File Systems Size: 1.03 P Virtual: 196.13 T Data Reduction: 1.3 to 1 Unique: 149.76 T Destroyed: 0.00 Snapshots: 49.98 T Total: 199.75 T									
File Systems									
Name	Source Location	Source Name	Size	Virtual	Hard Limit	Created	Protocols	Promotion Status	Writable
fbst-fs-1	-	-	150 G	2.34 G	False	2023-08-17 05:53:01	HTTP, NFSv3, NFSv4.1, SMB	promoted	True
fbst-fs-10	-	-	150 G	2.34 G	False	2023-08-17 05:53:15	HTTP, NFSv3, NFSv4.1, SMB	promoted	True
fbst-fs-100	-	-	150 G	150.70 G	True	2023-01-11 12:28:16	HTTP, NFSv3, NFSv4.1, SMB	promoted	True

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.

You can alternately 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:** Ratio of the size of the written data versus the amount of space the data occupies after data compression.
- **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 5.7 File System Details Page



Array **File Systems** Object Store

File Systems > EDA

Size 100.00 T Virtual 0.00 Data Reduction - Unique 0.00 Snapshots 0.00 Total 0.00

Details

Provisioned Size 100 T Hard Limit False Promotion Status promoted
 Default Group Quota - Protocols SMB Writable True
 Default User Quota - Source

Attached Policies 1-2 of 2

Name Type
 15minutes snapshot
 test_jh nfs-export

User Quotas +
 No quotas found.

Group Quotas +
 No quotas found.

File System Snapshots General Transfer 1-10 of 121 < > +

Name	Source Location	Source Name	Policy	Created
EDA.2021_10_14_15_34	Oban	EDA	15minutes	2021-10-14 15:34:50
EDA.2021_08_30_16_14	Oban	EDA	-	2021-08-30 16:14:20
EDA.2021_08_30_15_14	Oban	EDA	-	2021-08-30 15:14:20
EDA.2021_08_30_14_14	Oban	EDA	-	2021-08-30 14:14:20

The file systems detail page provides the following panels:

- **Details:** Provides an overview of the file system's configured parameters, including: the file system's provisioned size, default group and quota sizes, if hard limits is enabled (True or False), export

protocols, if ACLs safeguarding is enabled (True or False), the applied access control style, the file system's promotion status, if the file system is writable, and its source (for replication setups).

- **Attached Policies:** Lists all snapshot and NFS export policies applied to the file system and provides management of these policies.
- **User Quotas:** Lists all configured user quotas and provides management of user quotas.
- **Group Quotas:** Lists all group quotas and provides management of group quotas.
- **File System Snapshots:** Lists all file system snapshots and provides management of snapshots.
- **Destroyed Snapshots:** Lists all destroyed snapshots and allows you to recover or eradicate destroyed snapshots.

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:** rovisioned 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.

About Creating and Exporting File Systems

Creating and exporting a file system involves creating the file system and enabling protocol support for the file system. FlashBlade file 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. 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 8191P. 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 Snapshots

A file system snapshot is an immutable, point-in-time image of a file system. A snapshot can be created for any file system at any time. Snapshots are named after the file system it is copying with an assigned suffix to differentiate it from the source file system. A suffix is automatically assigned to a snapshot if one is not given at the time of creation.

Each file system has special directories named `.snapshot` that by default are enabled at file system creation to provide access to stored snapshots at the root and sub-directory level of that file system. The `.snapshot` directory provides a virtual repository for snapshots of a specified file system. By enabling a file system's snapshot directory, the `.snapshot` directory will be visible in the root directory and accessible at the sub-directory level of that file system. The `.snapshot` directories allow access to their parent-level directories. All snapshots for a file system will be accessible in the `.snapshot` directory at the root level. The `.snapshot` directories at the sub-directory level allow access to snapshots taken after those sub-directories were created. For example, all snapshots are accessible through the root `/ .snapshots` directory because it was created before any snapshots were taken (during the creation of the file system). At the `/lolcats/ .snapshot` sub-directory, snapshots taken before `/lolcats/` was created are not accessible.

The `.snapshot` directory can be optionally disabled. Once disabled, historical snapshots cannot be viewed and will not be accessible from the mount points until the directory feature is re-enabled. Disabling the directory will not destroy any snapshots, they must be manually destroyed.

The `.snapshot` directories cannot be renamed, deleted, or moved into other directories. The directory name is reserved and you cannot create a regular directory with the `.snapshot` name, regardless of the enabled/disabled status.

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.

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:** Multi-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.

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.

File Replication Failover, Reprotect, and Failback

In file system replication setups you can stop writes to the local file system and redirect writes to the target file system (failover). Failover can be used if there is an issue with the local file system. During failover, replication stops until you reprotect.

Failover requires promotion of the target file system and redirection of all clients to the target array to return to operation.

Note that when demoting the local file system, any writes performed since the last replication will be discarded. If a local snapshot was taken since last replication, demoting the file system requires that snapshot be destroyed and eradicated. You must remove any data since the last snapshot of the file system was taken. The file system must either have no snapshots, or the most recent snapshot must

be the last replicated snapshot. Additionally, it is suggested that any policies attached to the local file system and replica link be removed in order to expedite failover.

Once the local file system comes back up, re-enable the scheduled replication policy that was disabled before demoting the file system. Replication resumes as scheduled between the promoted remote file system and the demoted local file system. Data is written to the promoted remote file system, which is then replicated to the demoted local file system.

For fan out replication setups, once scheduled replication policies for the local file system are disabled across all the clusters in the fan out file system configuration, replica links related to that file system on the target clusters must be deleted. Once the file system on the new source cluster is promoted, replica links must be created between the promoted file system and the other clusters in your fan out file system replication configuration, and the original source file system must be demoted. Note that any policies that were disabled on the former source file system cannot be re-enabled on the promoted file system. You must create new policies on the promoted file system.

To resume the original replica link direction (failback), manually stop writes on the remote file system and send the remaining data to the local file system. Once data is sent to the local file system, promote the local file system and demote the remote file system.

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. A file system can support multiple protocols.

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.

 **Important!** When enabling the SMB protocol in Windows environments you must modify the default and ownership permissions from your SMB client. Refer to the Knowledge Base article, [Initial SMB Share Permissions on FlashBlade](#), for instructions.

For the SMB protocol in Native mode, 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.

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 **Provisioned Size** field, specify the provisioned size allocated to the file system. The size is a quota of space that helps gauge the fullness of the file system. If left blank, the provisioned size will default to an unlimited size. To enable a hard limit for the file system to prevent exceeding the file system's provisioned space, set the **Hard Limit** toggle button to enable (blue). For more information about the provisioned size, refer to [File System Size](#).

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

 **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.

- 4 (Optional for NFS users and groups) In the **Default User Quota** field, specify the user quota size.
- 5 (Optional for NFS users and groups) In the **Default Group Quota** field, specify the group quota size.
- 6 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](#).
- 7 In the **Special Directories** section, to enable the ability to access stored snapshots in a snapshot directory, select **Snapshot**.
- 8 Under **Protocols**, perform the following protocol-specific configurations:

a **NFS**. Define the NFS export rules so the file system is accessible to the appropriate clients.

- 1 Click the **NFSv3** or **NFSv4.1** checkbox to enable.
- 2 Select **Use Policy** if you have already created an NFS export policy with export rules that you wish to apply to the file system, or **Use Rules** to configure NFS export rules.

If you selected **Use Policy**, select the NFS export rule from the **Export Policy** list.

 **Note:** You can order the rules in an NFS export rule string as you wish, but the order is not honored and will always be interpreted with the following prioritization: IP address > Subnet > Netgroup > Wildcard.

If you selected **Use Rules**, in the **Export Rules** field 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 not supported for NFSv3.
- b SMB.** Click the **SMB** checkbox to enable the SMB protocol adapter. The **SMB native** mode is displayed under the **SMB** checkbox.
- 1** Select a client policy from the **Client Policy** list.
 - 2** Select a share policy from the **Share Policy** list.
 - 3** Select the **Continuous Availability** checkbox to ensure the file system remains available to Windows clients during disruptions that can occur from events such as blade failures, network disruptions, and upgrades.
- c HTTP.** Click the **HTTP** checkbox to enable the HTTP protocol for HTTP and HTTPS access to a file system.
- d Multi-protocol.** Select if you wish to use multiple protocols. Set the access control style to govern the interaction between multiple protocols.

 **Note:** With Purity//FB 3.1.1 and later, by default, all new file systems created are created with the **Shared** access control style with ACL safeguarding enabled.

Available access control styles are:

- **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**.

(Optional) If **Shared** or **NFS** are selected, you can select the **Safeguard ACLs** checkbox to enable ACLs safeguarding.

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 the following:

- If the access control style was set to **SMB**, ACL safeguarding is automatically enabled. Disabling ACL safeguarding is not allowed.
- ACL safeguarding is not available for **Independent** or **Mode Bits** access control styles.

9 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 If SMB is enabled and the **Continuous Availability** status of the checkbox is changed, the SMB workloads and all SMB sessions on this file system will be disrupted. This can impact multiple file systems if the affected SMB sessions are shared across them.
- 4 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 Click **NFS** in the **Protocols** section.
- 3 If you wish to change the export rules, perform the following:

- a Select **Use Rules**.
- b In the **Export Rules** field, configure the NFS export rules to define the access rights and privileges a client has to the file system.

The format is `host (options)`.

For IPv4 addresses, `host` represents `ddd.ddd.ddd.ddd` for IP address, `ddd.ddd.ddd.ddd/dd` for netmask, or `*(options)` for all clients.

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 not supported for NFSv3.

- 4 If you wish to change the NFS export policy, perform the following:
 - a Select **Use Policy**.
 - b Select the policy you wish to apply to the file system from the **Export Policy** list.
- 5 Click **Save**.

Changing the Provisioned Size of a File System

To change the provisioned size of a file system:

⚠ Important! If you are reducing a file system's provisioned size, but had previously enabled the hard limit for the file system and have already exceeded the new size, a warning will appear indicating subsequent data writes will fail. You must click **Update** to proceed with the operation.

- 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 In the **Provisioned Size** field, set the size to any value between 0 and 8191P. If the field is left blank or set to 0, the provisioned size will default to an unlimited size and a hard limit cannot be set.
- 3 Click **Save**.

Modifying the Default User and Group Quotas

⚠ 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.

- 1 From the **Storage > File Systems** page, click the **More Options** button in the row of the file system you want to edit. and click **Edit** The **Edit File System** pop-up window appears.
- 2 In the **Default User Quota** field, enter a new user quota size.
- 3 In the **Default Group Quota** field, enter a new group quota size.
- 4 Click **Save**.

Creating an Explicit User Quota

⚠ Important! If a default user quota was created during file system creation, the explicit user quota will take priority over the default quota and will be used for that file system.

⚠ 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.

- 1 From the **Storage > File Systems** page, click the file system for which you wish to set an explicit user quota. The file system's details page appears.
- 2 In the **User Quotas** panel, click the add button (+). The **Add User Quota** pop-up window appears.
- 3 Enter the user ID or user name, and quota size.
- 4 Click **Add**.

Modifying an Explicit User Quota

 **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.

- 1 From the **Storage > File Systems** page, click the file system whose explicit user quota(s) you wish to modify. The file system's details page appears.
- 2 From the **User Quotas** panel, click the **Edit** button on the same row as the explicit user quota you wish to modify. The **Edit User Quota** pop-up window appears.

 **Note:** By default, the user quotas 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 modify multiple explicit user quotas, click **More Options > Edit Quotas**. The **Edit User Quotas** pop-up window appears. From the **Existing Quotas** list, select the user quotas you wish to modify. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all user quotas by clicking the checkbox next to the search box.

All selected user quotas appear in the **Selected Quotas** list.

- 3 Enter a new quota size.
- 4 Click **Save**.

Deleting an Explicit User Quota

 **Important!** If a default user quota was created during file system creation, the default user quota will take effect once the explicit user quota is deleted.

- 1 From the **Storage > File Systems** page, click the file system whose explicit user quota(s) you wish to delete. The file system's details page appears.
- 2 From the **User Quotas** panel, click the trash can icon on the same row as the explicit user quota you wish to delete. The **Delete User Quota** pop-up window appears.

 **Note:** By default, the user quotas 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 delete multiple explicit user quotas, click **More Options > Delete Quotas**. The **Delete User Quotas** pop-up window appears. From the **Existing Quotas** list, select the user quotas 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 user quotas by clicking the checkbox next to the search box. All selected user quotas appear in the **Selected Quotas** list.

- 3 Click **Delete**.

Creating an Explicit Group Quota

 **Important!** If a default group quota was created during file system creation, the explicit group quota will take priority over the default quota and will be used for that file system.

 **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.

- 1 From the **Storage > File Systems** page, click the file system for which you wish to set an explicit group quota. The file system's details page appears.
- 2 In the **Group Quotas** panel, click the add button (+). The **Add Group Quota** pop-up window appears.
- 3 Enter the group ID or group name, and quota size.
- 4 Click **Add**.

Modifying an Explicit Group Quota

 **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.

- 1 From the **Storage > File Systems** page, click the file system whose explicit group quota(s) you wish to modify. The file system's details page appears.
- 2 From the **Group Quotas** panel, click the **Edit** button on the same row as the explicit group quota you wish to modify. The **Edit Group Quota** pop-up window appears.

 **Note:** By default, the group quotas 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 modify multiple explicit group quotas, click **More Options > Edit Quotas**. The **Edit Group Quotas** pop-up window appears. From the **Existing Quotas** list, select the group quotas you wish to modify. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all group quotas by clicking the checkbox next to the search box. All selected group quotas appear in the **Selected Quotas** list.

- 3 Enter a new quota size.
- 4 Click **Save**.

Deleting an Explicit Group Quota

 **Important!** If a default group quota was created during file system creation, the default group quota will take effect once the explicit group quota is deleted.

- 1 From the **Storage > File Systems** page, click the file system whose explicit group quota(s) you wish to delete. The file system's details page appears.
- 2 From the **Group Quotas** panel, click the trash can icon on the same row as the explicit group quota you wish to delete. The **Delete Group Quota** pop-up window appears.

 **Note:** By default, the group quotas 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 delete multiple explicit group quotas, click **More Options > Delete Quotas**. The **Delete Group Quotas** pop-up window appears. From the **Existing Quotas** list, select the group quotas 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 group quotas by clicking the checkbox next to the search box. All selected group quotas appear in the **Selected Quotas** list.

- 3 Click **Delete**.

Enabling and Disabling the Hard Limit of a File System

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

- 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** dialog box appears.
- 2 Set the **Hard Limit** toggle button to enable (blue) or disable (gray) the feature.

Creating a File System Snapshot

Snapshots are immutable, point-in-time images of the contents of one or more file systems. Snapshot tasks include creating and deleting file system snapshots.

To create a snapshot of a file system:

- 1 Select **Storage > File Systems**.
- 2 Click the file system. The snapshot list appears for that file system.
- 3 Click add (+) on the right. The **Create Snapshot** pop-up window appears.
- 4 Enter a suffix name for the snapshot. A suffix will be automatically assigned if this field is left empty.
- 5 If you have file replication set up, select the target(s) to which you wish to replicate the snapshot.
- 6 Click **Create**.

Enabling and Disabling the Snapshot Directory

A snapshot directory is a special directory that can be enabled or disabled to store snapshots. By default, snapshot directories are enabled at file system creation.

- 1 To enable or disable the snapshot directory:
 - 1 Select **Storage > File Systems**.
 - 2 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.
 - 3 Click the **Snapshot** toggle button to switch between enabled (blue) and disabled (gray) status.
 - 4 Click **Save**.
 - 5 If you are enabling the snapshot directory, in the root directory of the NFS mount, confirm that a pseudo-directory named `.snapshot` has been created.

The `.snapshot` directory name is reserved. You cannot create a directory with the `.snapshot` name, regardless of the enabled/disabled status. The directory cannot be renamed, deleted, or moved into other directories. It can only be removed by disabling the directory.

Removing a Policy from a Snapshot

Policies can be removed from snapshots at any time.

 **Note:** If a policy is removed from a snapshot, that snapshot is no longer managed by that policy and must be destroyed manually.

You can remove a policy from a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To remove a policy from a snapshot, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots.
- 2 If you are on the **Protection** page, select **Protection > Snapshots**.
- 3 From the table in the **File System Snapshots** panel, locate the file system snapshot from which you wish to remove a policy.
- 4 Click **More Options > Remove Policy** at the end of that snapshot's row . The **Remove Policy from Snapshot** pop-up window appears.
- 5 If there are multiple policies attached to the file system snapshot, click **More Options > Remove Policies** at the end of that snapshot's row. The **Remove Policies from Snapshot** pop-up window appears.
- 6 From the **Policies** panel, select one or more snapshot policies from the list. The selected policy (or policies) appears in the **Selected Policies** panel.
- 7 Click **Remove**.

Destroying File System Snapshots

You can destroy a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To destroy one or more file system snapshots, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 From the table in the **File System Snapshots** panel, locate the file system snapshot you either wish to destroy, or is the point from which earlier snapshots are to be destroyed.
- 3 If you wish to destroy the snapshot, perform the following:

- a Click the **More Options** button located at the end of that snapshot's row and select **Destroy**. The **Destroy Snapshot** pop-up window appears.

- b Click **Destroy**

4 If you wish to destroy multiple snapshots, perform the following:

 **Note:** By default, the snapshots 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 **More Options > Destroy** in the **File System Snapshots** panel title bar. The **Destroy File System Snapshots** pop-up window appears.
- a From the **File System Snapshots** list, select the snapshots you wish to destroy. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- b Click **Destroy**.

5 If you wish to use the snapshot as the point from which earlier snapshots are destroyed, perform the following:

- a Click the **More Options** button located at the end of that snapshot's row and select **Destroy older**. The **Destroy Snapshots older than...** pop-up window appears.

 **Note:** By default, the buckets 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 Select the snapshots you wish to destroy from the **Older File System Snapshots** panel. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- c Click **Destroy**.

The destroyed file system snapshot appears under the **Destroyed Snapshots** panel on the **Snapshots** tab. Destroyed snapshots can be recovered within a 24-hour period. Once the 24-hour period has passed, the snapshot is permanently destroyed (eradicated).

Restoring a File System to a Snapshot

Purity//FB allows you to restore a file system to a given point in time from a snapshot (via file system rollback). Restoring a file system to a snapshot overwrites the contents of that file system with data from the snapshot.

File system rollback is available for snapshots created on a FlashBlade array running Purity//FB 3.0.0 or later.

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

 **CAUTION:** When restoring to a snapshot, the file system's configured quota, as well as any configured user/group quotas, do not change to reflect the limit when the snapshot was taken. This can cause the live space usage to increase, resulting in subsequent write failures due to exceeding the quota. You may need to re-adjust your quota settings accordingly.

 **Note:** It is highly recommended that there be no I/O on the file system when performing a restore.

You can restore a file system from a snapshot from either the **Storage** or **Protection** page.

 **Note:** If there are snapshots that are newer than the snapshot you wish to use as the restore point, the newer snapshots must be first destroyed and eradicated before you can proceed with the file system restoration. See [Destroying and Eradicating Newer Snapshots](#).

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To restore a file system from a snapshot, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 From the table in the **File System Snapshots** panel, locate the file system snapshot to be used to restore the file system.
- 3 Click the **More Options** button located at the end of row of the snapshot to be used for restore and select **Restore**. The **Restore File System from Snapshot** pop-up window appears.
- 4 Click **Restore**.

Destroying and Eradicating Newer Snapshots

 **Note:** Destroyed snapshots can be recovered within a 24-hour period. Once the 24-hour period has passed, the snapshot is permanently destroyed (eradicated).

When restoring a file system from a snapshot, if there are snapshots that are newer than the snapshot you wish to use as the restore point, the newer snapshots must be first destroyed and eradicated before you can proceed with the file system restoration.

You can destroy and eradicate snapshots from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 From the table in the **File System Snapshots** panel, locate the file system snapshot to be used to restore the file system.
- 3 Click the **More Options** button located at the end of that snapshot's row and select **Destroy newer**. The **Destroy Snapshots newer than...** pop-up window appears.
- 4 In the **Newer File System Snapshots** panel, select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- 5 Click **Destroy**. The snapshot(s) appear in the **Destroyed Snapshots** panel (directly below the **File System Snapshots** panel).
- 6 Click **More Options > Eradicate** in of the **Destroyed Snapshots** panel. The **Eradicate File System Snapshots** pop-up window appears.

 **Note:** By default, the snapshots 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.

- 7 From the **Destroyed File System Snapshots** list, select the snapshots you wish to eradicate. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the checkbox next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- 8 Click **Eradicate**.

Recovering File System Snapshots

When a file system snapshot is destroyed, you have 24 hours to recover the data. This is known as the eradication pending state. When the eradication pending state expires, the snapshot will be eradicated and you will no longer be able to retrieve the snapshot data.

You can recover a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To recover a previously destroyed file system snapshot, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 If you wish to recover a single snapshot:
 - a From the **Destroyed Snapshots** panel, locate the file system snapshot you wish to recover.
 - b Click **More Options > Recover** at the end of the row of the file system snapshot you wish to recover.
 - c Click **Recover**.
- 3 If you wish to recover multiple file system snapshots:

 **Note:** By default, the snapshots 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 **More Options > Recover** in the **Destroyed Snapshots** panel title bar. The **Recover File System Snapshots** pop-up window appears.
- b From the **Destroyed File System Snapshots** list, select the snapshots you wish to recover. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- c Click **Recover**.

Eradicating File System Snapshots

During the eradication pending period, you can manually eradicate the destroyed file system to reclaim physical storage space occupied by the destroyed snapshot. Once reclamation starts, the destroyed snapshots can no longer be recovered.

You can eradicate a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

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

To eradicate a destroyed snapshot:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's destroyed snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 If you wish to eradicate a single snapshot:
 - a In the **Destroyed Snapshots** panel under **File System Snapshots**, click **More Options > Recover** at the end of the row of the snapshot you want to eradicate. The **Eradicate Snapshot** pop-up window appears.
 - b Click **Eradicate**. The array will begin reclaiming the eradicated space of the snapshot.
- 3 If you wish to eradicate multiple snapshots:

 **Note:** By default, the snapshots 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 **More Options > Eradicate** in of the **Destroyed Snapshots** panel. The **Eradicate File System Snapshots** pop-up window appears.
- b From the **Destroyed File System Snapshots** list, select the snapshots you wish to eradicate. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the checkbox next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- c Click **Eradicate**. The array will begin reclaiming the eradicated space of the snapshots.

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

Promoting and demoting file systems is required when performing file system failover and failback. In a file replication setup, data is written to the promoted local file system and then that data is replicated to the demoted remote file system on a remote array. During failover, the local file system is demoted and the remote file system is promoted. During failback, the original promoted/demoted filesystem designations are restored.

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

For complete failover and failback instructions, see [Performing File Replication Failover and Reprotect](#) and [Performing File Replication Failback](#).

To promote and demote file systems, perform the following:

- 1 From 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 statistics, perform 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 performance 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 NFS Client, User, and Group 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 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.

Performing File Replication Failover and Reprotect

To initiate failover in a file replication setup:

- The target file system must first be promoted.
- All clients must then be directed to the target array.
- To reprotect, the local file system is then demoted. Note that when a file system is demoted, any writes performed since the last replication will be discarded.

To initiate failover, perform the following:

- 1 On the target array, from the **Storage > File Systems** page, locate the remote file system that you wish to promote from the **File Systems** panel.
- 2 Click **More Options > Promote** at the end of the row in which the remote file system appears.
- 3 Click **Promote** when prompted for confirmation.
- 4 Manually redirect clients to the target array.

To reprotect, you must remove snapshots since the last replicated snapshot. This can be difficult if the policy engine is creating new snapshots. You can stop the policy engine from creating new snapshots by removing the policies from the file system and replica link.

- 5 On the local array, remove policies attached to the to-be-demoted local file system and replica link.

To remove the policy from the local file system:

- a From the **Storage > File Systems** page, locate and click local file system that you wish to demote from the **File Systems** panel. The file system's details page appears.
- b In the **Policies** panel, click the remove button (X) in the same row as the policy you wish to remove.
- c Click **Remove** when prompted for confirmation.

To remove the policy from the replica link:

- a From the **Protection > File Replica Links** page, click **More Options > Remove Policies** at the end of the row of the replica link you wish to edit. The **Remove Policies** pop-up window appears.
- b Select the policy or policies you wish to remove from the replica link from the **Available Policies** list.
- c Click **Remove**.

- 6 Navigate back to the **Storage > File Systems** page and locate the local file system that you wish to demote from the **File Systems** panel.
- 7 Click **More Options > Demote** at the end of row in which the local file system appears.
- 8 Click **Demote** when prompted for confirmation.

Writes to the local file system are stopped and are redirected to the promoted remote file system.

- 9 If you removed policies in step 5, add the removed policy or policies back to the file system and replica link.

To add the policy to the file system:

- a From the **Storage > File Systems** page, locate and click the file system from the **File Systems** panel. The file system's details page appears.

- b** In the **Policies** panel, click the add button (+). The **Add Policies** pop-up window appears.
- c** Select the policy or policies you wish to add to the file system from the **Available Policies** list.
- d** Click **Add**.

To add the policy to the replica link:

- a** From the **Protection > File Replica Links** page, click **More Options > Add Policies** at the end of the row of the replica link you wish to edit. The **Add Policies** pop-up window appears.
- b** Select the policy or policies you wish to add back to the replica link from the **Available Policies** list.
- c** Click **Add**.

Once the local file system comes back up, replication resumes as scheduled between the newly promoted target file system and demoted local file system because the policy was re-added to the replica link. Because clients have been redirected to the target, data is written to the promoted remote file system, which is then replicated to the demoted local file system.

Performing File Replication Failback

 **Note:** This procedure continues from the previous failover procedure. During the failover procedure, the remote file system was promoted and the local file system was demoted.

As noted in the failover procedure, if the replication schedule policy was re-added to the replica link in the replication setup, once the local file system comes back up, replication resumes as scheduled between the promoted target file system and demoted local file system. Performing failback restores the original local/target relationship between arrays and file systems in a replication setup. During failback, once the remote file system is demoted and the local file system is promoted and clients redirected to the local file system, the original replication setup resumes; data is written to the local file system and then replicated to the remote file system.

To initiate failback in a file replication setup:

- Stop writes on the remote file system and send the remaining data to the local file system.
- Promote the local file system.
- Add back any removed policy to the local file system.
- Demote the remote file system.

To initiate failback, perform the following:

- 1 On the target array, from the **Storage > File Systems** page, locate the remote file system that is to be demoted from the **File Systems** panel.
- 2 Click **More Options > Disable Writes** at the end of row in which the remote file system appears.
- 3 Click **Disable** when prompted for confirmation.
- 4 Ensure all writes are sent to the source file system by creating a snapshot on the target and send it to the source.
- 5 On the local array, from the **Storage > File Systems** page, locate the local file system that you wish to promote from the **File Systems** panel.
- 6 Click **More Options > Promote** at the end of row in which the local file system appears.
- 7 Manually redirect clients back to the to the local array.
- 8 Remove policies attached to the to-be-demoted remote file system and replica link.

To remove the policy from the file system:

- a From the **Storage > File Systems** page, locate and click local file system that you wish to demote from the **File Systems** panel. The file system's details page appears.
- b In the **Policies** panel, click the remove button (**X**) in the same row as the policy you wish to remove.
- c Click **Remove** when prompted for confirmation.

To remove the policy from the replica link:

- a From the **Protection > File Replica Links** page, click **More Options > Remove Policies** at the end of the row of the replica link you wish to edit. The **Remove Policies** pop-up window appears.
 - b Select the policy or policies you wish to remove from the replica link from the **Available Policies** list.
 - c Click **Remove**.
- 9 On the target array, from the **Storage > File Systems** page, locate the remote file system that you wish to demote from the **File Systems** panel.
 - 10 Click **More Options > Demote** at the end of row in which the remote file system appears.
 - 11 Click **Demote** when prompted for confirmation.

Writes to the remote file system are stopped and are redirected to the promoted local file system.

- 12 If you removed policies in step 8, add the removed policy or policies back to the file system and replica link.

To add the policy to the file system:

- a** From the **Storage > File Systems** page, locate and click the file system from the **File Systems** panel. The file system's details page appears.
- b** In the **Policies** panel, click the add button (+). The **Add Policies** pop-up window appears.
- c** Select the policy or policies you wish to add to the file system from the **Available Policies** list.
- d** Click **Add**.

To add the policy to the replica link:

- a** From the **Protection > File Replica Links** page, click **More Options > Add Policies** at the end of the row of the replica link you wish to edit. The **Add Policies** pop-up window appears.
- b** Select the policy or policies you wish to add back to the replica link from the **Available Policies** list.
- c** Click **Add**.

The original replication setup has been restored.

Object Store

The **Object Store** tab is used to manage the array's object store accounts, users, and buckets.

Figure 5.8 Object Store Tab

Storage

Array
File Systems
Object Store

Object Store

Quota Limit
Virtual
Data Reduction
Unique
Destroyed
Total
Object Count
User Count

-
137.59 T
1.0 to 1
136.04 T
0.00
136.04 T
~ 2.41 B
1,016

Accounts

General
Space
Public Access
1-10 of 1,014

Name	Quota Limit	Available	Virtual	Hard Limit	Created	
boto-vms	-	-	8.80 T	False	2022-06-09 04:33:44	✎ ✕
boto3	-	-	0.00	False	2022-11-30 05:46:08	✎ ✕
btest	-	-	512.00 B	False	2022-07-18 06:23:53	✎ ✕
cflo	-	-	475.46 G	False	2022-08-08 00:10:08	✎ ✕
cflo2000-1002	-	-	0.00	False	2022-08-11 00:23:26	✎ ✕
cflo2000-1004	-	-	0.00	False	2022-08-11 00:23:52	✎ ✕
cflo2000-1006	-	-	0.00	False	2022-08-11 00:24:18	✎ ✕
cflo2000-1007	-	-	0.00	False	2022-08-11 00:24:31	✎ ✕
cflo2000-1008	-	-	0.00	False	2022-08-11 00:24:44	✎ ✕
cflo2000-1009	-	-	0.00	False	2022-08-11 00:24:57	✎ ✕

Access Keys

1-10 of 1,015

Access Key ID	Enabled	User	Created	
PSFBSAZRMAGHIIHAJDIBGODAJOAEPHEPELOJJBCHGJ	True	boto-vms/boto-vm	2022-06-09 04:34:00	✎

At the top of the **Object Store** tab, the following general information is displayed:

- **Quota Limit:** The user-specified size limit.
- **Virtual:** The amount of logically written data, excluding destroyed bucket virtual space.
- **Data Reduction:** Ratio of mapped sectors within an account versus the amount of physical space the data occupies after data compression for all objects.
- **Unique:** The unique physical space occupied by customer data, excluding the destroyed unique space from destroyed buckets.
- **Destroyed:** The sum of destroyed unique space of all destroyed buckets.
- **Total:** The provisioned size.
- **Object Count:** The approximate number of objects in the bucket. The units used for the approximate object count are **K** (thousands), **M** (millions), **B** (billions), and **T** (trillions). Note that there can be a temporary inconsistency in the number of objects shown in the GUI versus the client if an HA event (for example an NDU upgrade, software restart, failover event, etc.) occurred within a 12-hour period.

Objects deleted when bucket versioning is enabled will still be included in the object count total if the objects still have one or more previous versions.

- **User Count:** The total number of accounts.

The **Object Store** tab is divided into the following panels:

- **Accounts:** Displays all object store accounts.
- **Access Keys:** Displays all user access keys.
- **Buckets:** Displays all buckets.

Accounts Panel

The **Accounts** panel can be toggled between displaying **General** information (default), **Space** information, and **Public Access** information.

- **Name:** The account name.
- **Quota Limit:** The user-specified size limit for the account in bytes.
- **Available:** The available virtual space at the account level when an account-level quota is configured. If no hard limit is applied, null is displayed.
- **% Available:** Displayed only with **Space** information. The percentage of available space.
- **Virtual:** The total amount of pre-compressed data written to the account.
- **Destroyed:** Displayed only with **Space** information. The sum of destroyed virtual space of all the destroyed buckets of the account.
- **Hard Limit:** Displayed only with **General** information. Whether the quota limit is used as a hard limit quota (**True**) or not (**False**).
- **Created:** Displayed only with **General** information. The account creation date.
- **Data Reduction:** Displayed only with **Space** information. The ratio of mapped sectors within an account versus the amount of physical space the data occupies after data compression.
- **Total:** Displayed only with **Space** information. The total physical space, including physical space occupied by destroyed buckets.
- **Object Count:** Displayed only with **Space** information. The total approximate number of objects.

- **Block Public Access:** Displayed with only **Public Access** information. Whether public access is blocked (**True**) or not (**False**).
- **Block New Public Policies:** Displayed with only **Public Access** information. Whether creating new public access policies are blocked (**True**) or not (**False**).

You can search the accounts list by entering search criteria in the search box under the column names.

Access Keys Panel

The **Access Keys** panel displays the following information about user access keys:

- **Access Key ID:** The access key ID.
- **Enabled:** Whether the access key is enabled (**True**) or disabled (**False**).
- **User:** The user name in the format <account name>/<user name>.
- **Created:** The access key creation time.

You can search the access keys list by entering search criteria in the search box under the column names

Buckets Panel

The **Buckets** panel displays the following information about buckets:

- **Name:** The bucket name.
- **Account:** The account to which the bucket belongs.
- **Quota Limit:** The user-specified size limit for the bucket in bytes.
- **Available:** The following can be displayed:
 - If there is no account-level quota, or the account-level quota does not have a hard limit:
 - If there is a bucket-level quota with a hard limit, the available bucket space is displayed.
 - If there is no bucket-level quota, or the bucket-level quota does not have a hard limit, null is displayed.
 - If there is an account-level quota with a hard limit:

- If there is no bucket-level quota, or the bucket-level quota does not have a hard limit, the available writeable space in the account is displayed.
- If there is a bucket-level quota with a hard limit, the available space of either the account or bucket, whichever is lower, is displayed.
- **Virtual:** The total amount of pre-compressed data written to the bucket.
- **Hard Limit:** Whether the quota limit is used as the bucket's hard limit quota (**True**) or not (**False**).
- **Created:** The bucket creation time.
- **Versioning:** Whether bucket versioning is enabled or disabled.
- **Bucket Type:** The bucket type, **classic** or **multi-site-writable**. Refer to [Types of Buckets](#) for additional information about classic and multi-site writable buckets.

By default the **Buckets** panel is set to display the **General** view, which provides the information described above. You can also select **Space**, **Eradication Config**, **Object Lock**, and **Public Access** views for specific information about each configuration. See [Buckets](#) for additional information.

You can search the buckets list by entering search criteria in the search box under the column names.

Object Store Usage Limits

Usage limits, or quotas, can be set to limit the provisioned object store space consumption at the account and bucket levels. *Soft* and *hard* limits can be set.

A soft limit is the usage threshold at which an alert is triggered. When a soft limit usage value is set, the following alerts are raised:

- above 80% usage and below 90% usage - informational alert
- 90% usage and above - WARNING alert

A hard limit also triggers alerts, but will also halt writes at 100% usage:

- above 80% usage and below 90% usage - informational alert
- above 90% usage and below 100% usage - WARNING alert
- 100% usage and above - CRITICAL alert is triggered and all writes are halted until usage decreases below the hard limit size.

During account creation, you can set the account-level soft or hard limit and set a default soft or hard limit to be applied to any bucket subsequently created under that account. When creating a bucket, this default limit value can be overridden by specifying a new lower or higher limit for the bucket. Further, the limit type can also be overridden during bucket creation. For example, if a default soft limit of 15G for buckets was set during account creation, during bucket creation, you can specify a bucket hard limit of 10G.

When a hard limit is set, a limit check is performed for each write operation to the account or bucket. Prior to the write operation, the system checks that the current account or bucket usage is below its set hard limit. If the write operation is greater than 32M, the system performs a second check at the end of the write operation to ensure the account or bucket usage is still below its set hard limit. Once verified, the system marks the operation is complete. If the system detects that the account or bucket's hard limit has been exceeded, an error is returned.

Hard limits that are set at the account level supersede hard limits set for any buckets in that account—if all buckets under an account (including destroyed buckets) have collectively exceeded the account's hard limit, writes to the buckets, regardless of their hard limit, will halt.

Object Store Account Information

Per object store account information is available by clicking an account name from the **Object Store** tab.

Figure 5.9 Viewing Object Store Account Information

Array
File Systems
Object Store

Object Store
boto-vms

Quota Limit
Available
Virtual
Data Reduction
Total
Object Count
User Count

-
-
9.00 T
1.1 to 1
8.33 T
~1.25 B
1

Access Keys
1-1 of 1

Access Key ID
Enabled
User
Created

PSFBSAZRMAGHIHAJDIBGODAJOAFEPHEPELOJJBCHGJ
True
boto-vms/boto-vm
2022-06-09 04:34:00

Buckets
General
Space
Eradication Config
Object Lock
Public Access
1-3 of 3

Name
Quota Limit
Available
Virtual
Hard Limit
Created
Versioning
Bucket Type

boto-vm1
-
-
59.60 G
False
2022-07-21 09:00:48
none
multi-site-writable

boto-vm2
-
-
6.70 T
False
2022-07-18 03:40:13
none
multi-site-writable

boto-vm3
-
-
2.24 T
False
2022-07-18 03:40:20
none
multi-site-writable

Destroyed (0)

Details

Hard Limit
False

Bucket Default Quota Limit
-

Bucket Default Hard Limit
False

Block Public Access
True

Block New Public Policies
True

The account information is divided into three panels:

- **Access Keys** - displays the access key(s) configured with the account.
- **Buckets** - displays all buckets created with the account.
- **Details** - displays a summary of the account's configuration.

Creating an Account

When creating an account, you can optionally set the usage limit (quota) for the account and default usage limit for bucket's created in the account.

Note that if the hard limit option is selected for an account, writes will stop once the usage limit is reached. Further, if all buckets under the account have collectively exceeded the account's hard limit, writes to the buckets, regardless of their hard limit, will halt.

To create a new account:

- 1 In the **Accounts** panel of the **Storage > Object Store** page, click the add (+) button. The **Create Account** pop-up window appears.
- 2 Enter the new account's name in the **Name** field.
- 3 (Optional) Enter the usage limit, in bytes, of the **Quota Limit** for the account. Select **Hard Limit** if you wish writes to stop once the entered usage limit is reached (summed from all buckets in the account). If not selected, an alert is raised when usage reaches 80%, 90%, and 100%.
- 4 (Optional) Enter the usage limit, in bytes, of the **Bucket Default Quota Limit** for buckets created in this account. The bucket default quota limit applies to newly created buckets if a specific usage limit is not specified during the bucket's creation. Select **Hard Limit** if you wish writes to stop once the entered usage limit is reached. If not selected, an alert is raised when usage reaches 80%, 90%, and 100%.
- 5 Click **Create**.

By default, newly created accounts are not publicly accessible. If you wish to make the account publicly accessible, edit the account (see [Editing an Account](#)) and edit the buckets (see [Creating a Bucket](#) and [Editing a Bucket](#)).

Editing an Account

When editing an account, you can optionally set the usage limit (quota) for the account and default usage limit for bucket's created in the account. Note that if the hard limit option is selected, writes will stop once the usage limit is reached. Further, if all buckets under the account have collectively exceeded the account's hard limit, writes to the buckets, regardless of their hard limit, will halt. You can also set the public accessibility for buckets in the account.

To edit an account:

- 1 In the **Accounts** panel of the **Storage > Object Store** click the Edit button in the same row as the account you wish to edit. The **Edit Account** pop-up window appears.
- 2 (Optional) Enter the usage limit, in bytes, of the **Quota Limit** for the account. Select **Hard Limit** if you wish writes to stop once the entered usage limit is reached (summed from all buckets in the account). If not selected, an alert is raised when usage reaches 80%, 90%, and 100%.
- 3 (Optional) Enter the usage limit, in bytes, of the **Bucket Default Quota Limit** for buckets created in this account. The bucket default quota limit applies to newly created buckets if a specific usage limit is not specified during the bucket's creation. Select **Hard Limit** if you wish writes to stop once the entered usage limit is reached. If not selected, an alert is raised when usage reaches 80%, 90%, and 100%.
- 4 (Optional) If you wish to restrict access to buckets in this account to only authenticated users within the account, select **Block Public Access**.

- 5 (Optional) If you wish to block the addition of policies that grant public access to buckets in this account, select **Block New Public Policies**.
- 6 Click **Save**.

Deleting an Account

To delete an account:

- 1 In the **Accounts** section of the **Storage > Object Store** page, click the **Delete** button on the same row as the account you wish to delete. The **Delete Account** pop-up window appears.
- 2 Click the **Delete** button to confirm the account deletion.

Users

To view information about an existing user, from the **Storage > Object Store** page, click the account to which the user you wish to view belongs. The **Object Store > Accounts** information page appears which provides a **Users** panel, **Access Keys** panel, and **Buckets** panel.

The **Users** panel lists all account users. The following information is displayed about each account user:

- **Name:** The user name in the format <account name>/<user name>.
- **Access Key Age:** The age of the oldest access key for this user.
- **# Policies:** The number of access policies assigned to the user.
- **Created:** The user creation time.

The **Access Keys** panel lists all access keys for the account users:

- **Access Key ID:** The access key ID.
- **Enabled:** Whether the access key is enabled (True) or disabled (False).
- **User:** The user name in the format <account name>/<user name>.
- **Created:** The access key creation time.

The **Buckets** panel lists all buckets associated with the selected account and displays the following information:

- **Name:** The bucket name.

- **Created:** The bucket creation time.
- **Versioning:** Whether bucket versioning is enabled or disabled.

Access Keys

An access key allows the user to administer the object store. Users can have a maximum of two sets of keys at any single time. A set of keys consists of an access key ID and Secret Access Key.

Access keys can be created or imported for users. This can occur at the time of user creation, or any time after a user has been created.

When creating an access key you must copy or download your key information before closing the pop-up. The secret key can only be accessed during key creation confirmation and cannot be viewed or accessed after the confirmation pop-up is closed.

Credentials generated on one FlashBlade can be imported on another FlashBlade for a single account, with data on both arrays being accessible using the same credentials. This can be useful in FlashBlade replication setups, or in environments where multiple FlashBlades are used in multi-site deployments of Splunk SmartStore.

 **Note:** FlashBlade does not validate if the access key being imported from another FlashBlade was a previously deleted access key during or after import.

When importing access keys, the following should be noted:

- Only new access keys created after Purity//FB 3.0.0 can be imported on another FlashBlade.
- The access key you are importing cannot already exist on the FlashBlade to which it is being imported.
- If there are two access keys present for a user, you cannot import a third access key.
- Credentials are not synchronized or validated against the FlashBlade from which they were imported; deleting or disabling a set of credentials on the source cluster will not delete or disable the credentials on any other clusters where they were imported.

As a best practice, Pure Storage recommends that customers assign a unique access key for each individual user and each application, and avoid sharing an access key across two FlashBlade arrays except when specifically required by your application configuration (for example, Splunk SmartStore configurations on FlashBlade that use a second FlashBlade for disaster recovery).

Creating an Access Key

Creating an access key also creates the user's associated secret key.

Be sure to save off or download the access key and secret key information when you create a key.

To create an access key:

- 1 From the **Storage > Object Store** page, click the account to which you recently added a new user.
- 2 From the **Users** panel, click the user name for whom you wish to create a new access key. The user details page appears.
- 3 In the **Access Keys** panel, click the **More Options** button and select **Create**. The **Access Key** pop-up window appears displaying the new access key ID and secret access key.
- 4 Make note of, or copy, the access key and secret key information displayed in the confirmation pop-up. Optionally use the **CSV** or **JSON** button to download the key information.

 **Important!** You must copy or download your key information before closing the pop-up. The secret key can only be accessed during key creation confirmation and cannot be viewed or accessed after the confirmation pop-up is closed.

- 5 Use the access key and secret key in your `.s3cfg` or `.s3curl` files, for example.

Importing an Access Key

To import an access key:

- 1 From the **Storage > Object Store** page of the FlashBlade to which you are importing the key, click the account containing the user to which you wish to import the access key.
- 2 From the **Users** panel, click the user name for whom you wish to import a new access key. The user details page appears.
- 3 In the **Access Keys** panel, click the **More Options** button and select **Import**. The **Import Access Key** pop-up window appears.
- 4 Enter the access key ID and secret access key.
- 5 Click **Import**.

Deleting an Access Key

Deleting an access key also deletes the user's associated secret key. Once an access key has been deleted, it cannot be recovered.

To delete an access key:

- 1 From the **Storage > Object Store** page, click the account containing the user who's access key you wish to delete.
- 2 From the **Access Keys** panel, click the **More Options** button on the same row as the user for which you wish to delete the access key and select **Delete**.
- 3 When prompted to confirm the access key deletion, click the **Delete** button.

Disabling an Access Key

Access keys are enabled when they are created. To disable an access key:

- 1 From the **Storage > Object Store** page, click the account containing the user who's access key you wish to disable.
- 2 From the **Access Keys** panel, click the **More Options** button on the same row as the user for which you wish to disable the access key and select **Disable**.
- 3 When prompted to confirm disabling the access key, click the **Disable** button.

Enabling an Access Key

To enable a previously disabled an access key:

- 1 From the **Storage > Object Store** page, click the account containing the user who's access key you wish to enable.
- 2 From the **Access Keys** panel, click the **More Options** button on the same row as the user for which you wish to enable the access key and select **Enable**.
- 3 When prompted to confirm enabling the access key, click the **Enable** button.

Access Policies

FlashBlade Object Store allows you to manage what actions can be carried out by what users when accessing objects and buckets via the S3 API.

By default, new users created in Purity//FB 3.2.0 have no permissions assigned to them. You can assign a user permissions by choosing from a set of predefined, Pure-managed simple user policies, or *access policies*. Adding an access policy to a user in Object Store grants that user permissions to make specific types of S3 API calls on any object or bucket within the Object Store account where the user was created.

FlashBlade's access policies are pre-defined, static, and can only be added to or removed from users. Policies can be added during user creation, or any time thereafter. Similarly, policies can be removed from a user at any time.

 **Note:** The administrative users cannot create or delete the static policies.

The following table provides a list of the available access policies:

Table 5.2 Access Policies for Object Store

Access Policy	Policy Definition	Grant*
Storage Admin Roles		
pure:policy/bucket-list	Permissions to list all bucket names for the account.	s3:ListAllMyBuckets
pure:policy/bucket-info	Permissions to read bucket configurations.	s3:GetBucketVersioning s3:GetLifecycleConfiguration s3:GetBucketAcl s3:GetBucketLocation
pure:policy/bucket-configure	Permissions to configure buckets such as versioning and lifecycle rules on buckets.	s3:PutBucketVersioning s3:PutLifecycleConfiguration
pure:policy/bucket-create	Permissions to create buckets.	s3:CreateBucket
pure:policy/bucket-delete	Permissions to delete empty buckets.	s3>DeleteBucket
Application Roles		
pure:policy/object-list	Permissions to list objects and versions and verify their sizes, entity tags (ETags), and timestamps (but not to read their data or custom metadata).	s3:ListBucket s3:ListBucketVersions s3:ListBucketMultipartUploads

Table 5.2 Access Policies for Object Store (continued)

Access Policy	Policy Definition	Grant*
pure:policy/object-read	Permissions to read object and version data and custom metadata.	s3:GetObject s3:GetObjectVersion s3:GetObjectAcl
pure:policy/object-write	Permissions to create and overwrite objects when versioning is not enabled and to create new versions when versioning is enabled.	s3:PutObject s3:AbortMultipartUpload s3:ListMultipartUploadsParts
pure:policy/object-delete	Permission to eradicate objects when versioning is not enabled on the bucket and to create object delete markers when versioning is enabled.	s3:DeleteObject
pure:policy/version-delete	Permissions to eradicate objects and versions.	s3:DeleteObjectVersion
Special Roles		
pure:policy/full-access	Full Access is a special role. Permissions for all current and future S3 APIs and any non-standard FlashBlade S3 API extensions. (Triggers a warning in the GUI if selected)	s3:*
pure:policy/safemode-configure**	Permissions to configure safemode.	s3:ExtendSafemodeRetentionPeriod

*Subject to change. AWS occasionally updates the permissions defined in an AWS managed policy. When AWS does this, the update affects all principal entities (users, groups, and roles) to which the policy is attached. This does not affect security of FlashBlade access policies.

**If Object Safemode is disabled and this policy is not in use, this policy will not be visible.

Adding Access Policies

Access policies can be added to a user at any time. See [Access Policies](#) for a list of policies and their definitions.

To add access policies to a user:

- 1 From the **Storage > Object Store** page, click the account with the user for which you wish to add policies.
- 2 From the **Users** panel, click the user name for whom you wish to add policies. The user details page appears.
- 3 In the **Attached Access Policies** panel, click the **More Options** button and select **Add**. The **Add Access Policies** pop-up window appears displaying all available policies that can be added.
- 4 Select the policies you wish to add to the user from the **Available Policies** column.
- 5 Click **Add**.

The **Access Policies** panel is updated with the added policies.

Removing Access Policies

Access policies can be removed from a user at any time.

To remove access policies from a user:

- 1 From the **Storage > Object Store** page, click the account with the user for which you wish to remove policies.
- 2 From the **Users** panel, click the user name for whom you wish to remove policies. The user details page appears.
- 3 In the **Attached Access Policies** panel, click the **More Options** button and select **Remove**. The **Remove Access Policies** pop-up window appears displaying all existing policies that can be removed from the user.
- 4 Select the policies you wish to remove to the user from the **Existing Policies** column.
- 5 Click **Remove**.

The removed policies no longer appear in the **Access Policies** panel.

Creating a User

To create a new user:

- 1 From the **Storage > Object Store** page, click the account to which you wish to add a new user.
- 2 In the **Users** panel, click the add (+) button. The **Step 1: Create User** pop-up window appears.
- 3 Enter the new user's name in the **User Name** field.
- 4 Click **Create**. The **Step 2: Add Policies to user...** pop-up window appears.
- 5 (Optional) Select the simple user policy or policies you wish to assign the new user from the **Available Policies** column (see [Access Policies](#) for a list of policies and their definitions). You can skip this step and add policies at a later date. Note that the user will not have any access until policies have been added.
- 6 Click **Add**. The **Step 3: Add Access Key to user...** pop-up window appears.
- 7 (Optional) Select to **Create a new key** or **Import an existing key**. You can skip this step and an access key can be created for the user at a later date.
 - a If you select **Create a new key**, click **Create**. The **Access Key** pop-up window appears allowing you to copy the **Access Key ID** and **Secret Access Key** or download them in JSON or CSV format.
 - b If you select **Import an existing key**, enter the **Access Key ID** and **Secret Access Key**, and click **Import**.

 **Important!** You must copy or download your key information before closing the pop-up. The secret key can only be accessed during key creation confirmation and cannot be viewed or accessed after the confirmation pop-up is closed.

The access key is automatically enabled.

Deleting a User

Deleting a user also deletes all access keys associated with that user. To delete a user:


- 1 From the **Storage > Object Store** page, click the account from which you wish to delete a user.
- 2 From the **Users** panel, click the **More Options > Delete** button on the same row as the user you wish to delete.
- 3 When prompted to confirm the deletion, click **Delete**.

Buckets

Buckets and their contents can be created and managed with the FlashBlade GUI, CLI, and management REST API. Buckets and their contents can also be managed with the object store REST API, with open source tools such as s3curl and s3cmd, and with Java, Python, or other languages.

To view buckets per account, from the **Storage > Object Store** page, click an account name. The buckets for that account are displayed in the **Buckets** panel of the account's details page.

Figure 5.10 Viewing Buckets



Name	Quota Limit	Available	Virtual	Hard Limit	Created	Versioning	Bucket Type
boto-vm1	-	-	59.60 G	False	2022-07-21 09:00:48	none	multi-site-writable
boto-vm2	-	-	6.69 T	False	2022-07-18 03:40:13	none	multi-site-writable
boto-vm3	-	-	2.05 T	False	2022-07-18 03:40:20	none	multi-site-writable

The **Buckets** panel displays the following information about buckets:

- **Name:** The bucket name.
- **Quota Limit:** The user-specified size limit for virtual data on the bucket in bytes.
- **Available:** The following can be displayed:
 - If there is no account-level quota, or the account-level quota does not have a hard limit:
 - If there is a bucket-level quota with a hard limit, the available bucket space is displayed.
 - If there is no bucket-level quota, or the bucket-level quota does not have a hard limit, null is displayed.
 - If there is an account-level quota with a hard limit:
 - If there is no bucket-level quota, or the bucket-level quota does not have or hard limit, the available writeable space in the account is displayed.
 - If there is a bucket-level quota with a hard limit, the available space of either the account or bucket, whichever is lower, is displayed.
- **Virtual:** The total amount of pre-compressed data written to the bucket.
- **Hard Limit:** Whether the quota limit is used as the bucket's hard limit quota (**True**) or not (**False**).

- **Created:** The bucket creation time.
- **Versioning:** Whether bucket versioning is enabled or disabled.
- **Bucket Type:** The bucket type, **classic** or **multi-site-writable**.

By default the **Buckets** panel is set to display the **General** view, which provides the information described above. You can also select **Space**, **Eradication Config**, **Object Lock**, and **Public Access** views for specific information about each parameter.

The **Space** view includes the following:

- **% Available:** The percentage of available space.
- **Data Reduction:** The ratio of mapped sectors within an account versus the amount of physical space the data occupies after data compression.
- **Total:** The physical space occupied by the bucket.
- **Object Count:** The total approximate number of objects.

The **Eradication Config** view includes the following:

- **Retention Lock:** If Retention Lock is enabled (**ratcheted**) or disabled (**unlocked**).
- **Manual Eradication:** If manual eradication is **enabled** or **disabled**.
- **Eradication Delay:** The time delay from when a bucket is destroyed to when it is eradicated.

The **Object Lock** view includes the following:

- **Retention Lock:** If Retention Lock is enabled (**ratcheted**) or disabled (**unlocked**).
- **Object Lock:** If Object Lock is **enabled** or **disabled**.
- **Freeze Locked Objects:** If Freeze Locked Objects is enabled (**True**) or disabled (**False**). If enabled, locked objects cannot be overwritten.
- **Default Retention:** The retention period used for an Object Lock.
- **Default Retention Mode:** The Object Lock retention mode of **governance**, or **compliance**.

The **Public Access** view includes the following:

- **Block Public Access:** Whether public access is blocked (**True**) or not (**False**).

- **Block New Public Policies:** Whether new public policies are blocked (**True**) or not (**False**).
- **Public Status:** The status of bucket, which can be:
 - **public:** The bucket can be accessed without restrictions.
 - **bucket-and-object-not-public:** The bucket and its objects are not accessible.
 - **only-authorized-users-of-this-account:** The bucket and its objects are accessible to only specific users and roles based on the user access policies of the account.

In each view, you can search the buckets list by entering search criteria in the search box under the column names.

Refer to [Types of Buckets](#) for additional information about classic and multi-site writable buckets.

Refer to [Object Store Usage Limits](#) and [Bucket Limits](#) for additional information about Object Store account and bucket limits (quotas).

Refer to [Per-Bucket SafeMode](#) for additional information about Object Lock and Retention Lock.

Refer to [Public Buckets](#) for additional information about buckets.

To view details about a bucket, click its name from the **Buckets** panel.

Figure 5.11 Viewing Bucket Details

Array
File Systems
Object Store

Object Store > testaccount > testbucket2

Quota Limit - Available - Virtual 0.00 Data Reduction - Total 0.00 Object Count ~ 0

Bucket Replica Links

General Backlog +

No bucket replica links found.

Lifecycle Rules

+

No lifecycle rules found.

Details

Bucket Type

multi-site-writable

Hard Limit

False

Retention Lock

unlocked

Versioning

none

Object Lock

Default Retention

-

Default Retention Mode

-

Object Lock

disabled

Freeze Locked Objects

False

Eradication

Manual Eradication

enabled

Eradication Delay

1d

Public Access

Block Public Access

False

Block New Public Policies

False

Public Status

public

Bucket Policies

1-1 of 1

Name	Rules count	Enabled	Type	
testbucket2/acce...	1	True	bucket-access	

The bucket details page includes **Bucket Replica Links**, **Lifecycle Rules**, **Details**, and **Bucket Policies** panels.

Note: The information displayed in the **Bucket Replica Links** panel on this page is identical to the information displayed in the **Protection > Bucket Replica Links** panel.

For each replica link, the source bucket name, replication direction, remote bucket name, remote credentials, replication status, lags, and recovery point are displayed. See [Object Replication](#) for information.

For each lifecycle rule, the rule ID, rule status, object prefix filter, and number of days to keep non-current objects are displayed. See [Object Lifecycle Configuration Rules](#) for information.

The **Bucket Policies** panel displays the bucket access policies and CORS policies for public buckets.

The **Details** panel displays a summary of bucket's configuration.

Types of Buckets

Two types of buckets can be created: classic and multi-site writable.

With classic buckets, new objects are added to the top of the version stack. When classic buckets are connected by bi-directional replication, if the application writes different data to the same Object Key on multiple FlashBlades, there is a chance that the order of versions is not maintained across both systems.

When applications write to the same Object Key on multiple FlashBlades, a multi-site writeable bucket which is setup on each FlashBlade and connected via bi-directional replication, helps resolve write conflicts and manages a consistency across systems, so both FlashBlades are viewing the same version order of the object.

 **Note:** For replication setups, buckets must be the same type.

Bucket Limits

Hard and soft limits (quotas) can be set for buckets to limit space consumption.

A soft limit is the usage threshold at which an alert is triggered. When a soft limit usage value is set, the following alerts are raised:

- above 80% usage and below 90% usage - informational alert
- 90% usage and above - `WARNING` alert

A hard limit also triggers alerts, but will also halt writes at 100% usage:

- above 80% usage and below 90% usage - informational alert
- above 90% usage and below 100% usage - `WARNING` alert
- 100% usage and above - `CRITICAL` alert triggered and all writes are halted until usage decreases below the hard limit size.

When the system detects that the bucket's provisioned size has been exceeded, all future writes are halted until usage decreases below the provisioned size. This occurs within 120 seconds 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 20M write is in-process for a bucket whose provisioned size is 30M and current usage is 12M. When the system halts the write, the bucket'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.

 **Note:** A hard limit cannot be enabled or disabled if the bucket's provisioned size is unlimited.

During account creation, you can specify a default soft or hard limit for buckets that are created under that account. This default limit type and limit value can be overridden during bucket creation. However, hard limits that are set at the account level supersede hard limits set for any buckets in that account—if the account's hard limit is reached, writes to all buckets in the account are stopped regardless if the hard limits of the buckets in the account have been reached.

Bucket Versioning

A bucket's versioning status is displayed in the **Buckets** panel after clicking an account name from the **Storage > Object Store** page. Versioning status can also be displayed in the CLI using the `purebucket list` command.

Bucket versioning allows multiple variants of an object in the same bucket. Buckets can be in three different versioning states: **None**, **Enabled**, and **Suspended**. All buckets are initially created in a non-versioned (**None**) state. Once you enable versioning on a bucket, every object in the bucket automatically retains the entire history on every update, and that bucket can never return to the non-versioned state. However, you can suspend versioning on that bucket.

In a non-version bucket, each object name has only one version. This is always the current version. Deleting an existing object name permanently destroys the object.

In a versioning-enabled bucket, an object can have multiple versions for the same object name. Each object name has one current version and either zero or multiple non-current versions.

The versioning-suspended state can prevent the number of versions from growing.

Lifecycle rules can also prevent the number of versions from growing. Configuring a lifecycle rule, either for all objects or for objects matching a key name prefix, will cause noncurrent object versions older than the specified number of days to be permanently deleted from the bucket.

Per-Bucket SafeMode

Per-bucket SafeMode provides enhanced security of objects in the form of *Object Lock* and *Retention Lock*. Object Lock provides the ability to lock objects and prevents locked objects from being overwritten or deleted. Retention Lock allows you to lock and limit the ability to modify Object Lock settings.

Object Lock and Retention Lock can be enabled and configured when editing a bucket.

 **Note:** If both Object Lock and Retention Lock are enabled, non-empty buckets cannot be destroyed.

Object Lock

 **Note:** Replication to an Object Lock-enabled bucket is supported. However, replication *from* an Object Lock-enabled bucket is not supported.

When using Object Lock you can configure the default retention period and retention mode for a bucket's objects. Object Lock can only be enabled on an empty bucket. Once enabled, Object Lock cannot be disabled.

Object Lock's default retention period is set in days. Objects are retained for the number of days specified and cannot be eradicated or destroyed during this period.

Object Lock provides the following retention mode options:

- **governance:** Objects are protected from deletion. Only users with special permissions can alter the retention settings or delete the object. Otherwise, users cannot overwrite or delete an object, or alter the object's lock settings.
- **compliance:** Objects cannot be overwritten or deleted by any user. This mode ensures that an object cannot be overwritten or deleted for the duration of the retention period.

Additionally, Object Lock provides the Freeze Locked Objects option. When this option is enabled, additional security is added that prevents locked objects from being overwritten. In order to use the Freeze Locked Objects option, bucket versioning must not be enabled. If enabled, Pure Technical Services is required to disable this option.

Retention Lock

When enabled, Retention Lock limits the ability to modify Object Lock settings and certain functions. A Retention Lock can only be unlocked by Pure Technical Services.

When Retention Lock is enabled on a bucket, permission to modify settings is ratcheted (locked). When ratcheted, the ability to manually eradicate the bucket is disabled. Additionally, the following Object Lock settings are affected when Retention Lock is ratcheted:

- The Object Lock retention mode can be set to **governance** or **compliance**, and can be changed from **governance** to **compliance** or **governance** to **none**. Note that once set to **compliance**, the mode can no longer be changed.
- If the Object Lock retention mode is set to **compliance**, the retention period can only be increased and the mode cannot be changed.

Public Buckets

Access to content stored in an S3 bucket requires client authentication with the credentials and permissions of S3 users. FlashBlade supports public buckets, which are standard S3 buckets with the additional property of being public, allowing all clients to access the content in the public bucket.

 **Note:** When setting up a public bucket, both the bucket and the account that owns the bucket must be configured to allow public access.

When creating a bucket, you can optionally set the bucket to be publicly available by selecting the **Make Bucket Public** option. The account that owns the bucket can be edited to allow public access (if not already configured to do so) by disabling the **Block Public Access** and **Block New Public Policies** options. For instructions on creating buckets and editing accounts, refer to:

- [Creating a Bucket](#)
- [Editing an Account](#)

You can also edit a pre-existing bucket to be publicly available by disabling its **Block Public Access** and **Block New Public Policies** options (see [Editing a Bucket](#) for instructions). As noted, the owner account must also be configured to allow public access.


Additional access parameters are set using bucket access policies (see [Bucket Access Policies](#)).

Bucket Access Policies

A bucket access policy contains the access rules for a public bucket. When a public bucket is created, a default bucket access policy is also created and attached.

Currently all bucket access policies allow the s3:GetObject action to be performed on all objects in the bucket. Additional actions as well as flexibility to Principals may be provided in a future release. Bucket access policies are displayed in the **Bucket Policies** panel from a bucket's details page.

Figure 5.12 Viewing Bucket Access Policies



Bucket Policies				1-1 of 1
Name	Rules count	Enabled	Type	
testbucket/access...	1	True	bucket-access	

Clicking the name of the bucket access policy opens the **Bucket Access Policy Rules** panel, which displays the policy's rule details.

Figure 5.13 Viewing the Bucket Access Policy Rules Panel

Bucket Access Policy Rules					1-1 of 1 + ⋮
Name	Effect	Actions	Resources	Principals	
default	allow	s3:GetObject ⓘ	testbucket2/*	All: True	🗑️

Cross-Origin Resource Sharing

Cross-Origin Resource Sharing (CORS) protocol is also supported with buckets. CORS is an HTTP header-based security mechanism that defines how client web applications in one domain interact with resources in a different domain, preventing the loading of content from third-party sites.

Allowed HTTP methods, headers, and origins are defined in a CORS policy. Currently all headers, methods, and origins are applied with the creation of a CORS policy. Additional control may be provided in a future release.

CORS policies are displayed and created in the **Bucket Policies** panel from a bucket's details page.

Figure 5.14 Viewing CORS Policies

Bucket Policies					1-2 of 2 ⋮
Name	Rules count	Enabled	Type		
testbucket2/access-policy	1	True	bucket-access		🗑️
testbucket2/cors-policy	1	True	cross-origin-resource-sharing		🗑️

Clicking the name of the CORS policy opens the **Cross Origin Resource Sharing Policy Rules** panel, which displays the CORS policy's rule details.

Figure 5.15 Viewing the Cross Origin Resource Sharing Policy Panel

Cross Origin Resource Sharing Policy Rules					1-1 of 1 + ⋮
Name	Allowed headers	Allowed methods	Allowed origins		
testcors	*	GET, PUT, HEAD, POST, DELETE	*		🗑️

Creating a Bucket

 **Note:** For replication setups, the source and target buckets connected with a replica link must be the same type.

To create a bucket:

- 1 From the **Storage > Object Store** page, click the account to which you wish to add a bucket.
- 2 From the **Buckets** panel, click the add (+) button. The **Create Bucket** pop-up window appears.
- 3 Enter a name for the bucket.
- 4 From the **Bucket Type** list, select bucket type for created bucket.
- 5 (Optional) In the **Quota Limit** field, enter the bucket's usage limit, in bytes. Select **Hard Limit** if you wish writes to stop once the entered usage limit is reached. If not selected, an alert is raised when usage reaches 80%, 90%, and 100%. Note that hard limits that are set at the account level supersede hard limits set for any buckets in that account—if all buckets under an account have collectively exceeded the account's hard limit, writes to the buckets, regardless of their hard limit, will halt.
- 6 (Optional) To make the bucket publicly available, select **Make Bucket Public**. Note that the bucket's parent account settings must allow public access in order to set public access for the bucket.
- 7 Click **Create**.

When a public bucket is created, a default bucket access policy is also created and attached. The default bucket access policy allows unauthenticated public access to objects in the bucket.

Editing a Bucket

To edit a bucket:

- 1 From the **Storage > Object Store** page, click the account in which the bucket you wish to edit resides.
- 2 From the **Buckets** panel, click **More Options > Edit**. The **Edit Bucket** pop-up window appears.
- 3 Enter a name for the bucket.
- 4 (Optional) In the **Quota Limit** field, enter the bucket's usage limit, in bytes. Select **Hard Limit** if you wish writes to stop once the entered usage limit is reached. If not selected, an alert is raised when usage reaches 80%, 90%, and 100%. Note that hard limits that are set at the account level supersede hard limits set for any buckets in that account—if all buckets under an account have collectively exceeded the account's hard limit, writes to the buckets, regardless of their hard limit, will halt.
- 5 Select if you wish to enable **Versioning**.
- 6 To edit the bucket's retention security settings, click the arrow next to **Retention Settings** and optionally set each of the following:
 - a Select the **Object Lock** checkbox to enable Object Lock. Note that once enabled, Object Lock cannot be disabled.

- b** If **Object Lock** is enabled, select the **Freeze Locked Objects** checkbox to prevent locked objects from being overwritten. Note that once enabled, only Pure Technical Services can disable Freeze Locked Objects.
 - c** If **Object Lock** is enabled, enter the number of days objects are retained in the **Default Retention** field.
 - d** Enable the **Default Retention Mode** by selecting either **governance** or **compliance**.
 - e** Enable **Retention Lock** by selected **ratcheted (locked)**. Note that once enabled, only Pure Technical Services can disable Retention Lock.
- 7** To edit the bucket's public accessibility settings, click the arrow next to **Public Access Settings** and optionally set each of the following:
 - a** Select or deselect **Block Public Access** to restrict or allow access to the bucket to only authenticated users within the account or all users.
 - b** Select or deselect **Block New Public Policies** to block or allow the addition of policies that grant public access to the bucket.
- 8** Click **Save**.

Enabling Bucket Versioning

To enable bucket versioning, perform the following:

- 1** From the **Storage > Object Store** page, click the account to which the bucket you wish to version belongs.
- 2** From the **Buckets** panel, click **More Options > Enable versioning** on the same row as the bucket you wish to version. The **Enable Versioning** pop-up window appears.
- 3** Click **Enable**.

Suspending Bucket Versioning

To suspend bucket versioning, perform the following:

- 1** From the **Storage > Object Store** page, click the account to which the bucket whose versioning you wish to suspend belongs.
- 2** From the **Buckets** panel, click **More Options > Suspend versioning** on the same row as the bucket whose versioning you wish to suspend.. The **Suspend Versioning** pop-up window appears.
- 3** Click **Suspend**.

Object Lifecycle Configuration Rules

When bucket versioning is enabled, multiple versions of an object in the same bucket are created that retain the object's entire history on every update. In order to help manage space usage, you can create object lifecycle configuration rules to automatically delete old, unneeded versions of the object. Without an object lifecycle rule, an object's versions are retained indefinitely until manually deleted.

 **Note:** If bucket versioning was previously enabled, but currently suspended, the versions created while versioning was enabled are still subject to any lifecycle rules created after versioning was suspended.

When an object lifecycle rule is created, that rule manages the amount of time that the current and non-current version of the object is retained. For example, you can specify that current versions be retained for 2 days. Once the 2 day retention period is reached, the versions are deleted from the system within 24 hours. Similarly, you can also specify that a non-current version be retained for 10 days. When the 10 day retention period is reached, that version will be deleted from the system within 24 hours.

When a lifecycle policy is applied, then for all objects in the bucket or, if specified, for those objects matching the key prefix, all versions that have been noncurrent for more than the specified retention period are subject to immediate deletion.

Up to 1,000 object lifecycle rules can be created per bucket.

Object lifecycle rules can be viewed, created, edited, enabled, disabled, and deleted using the Purity//FB GUI, CLI, and Management REST API, as well as the Object Store REST API.

Object lifecycle rules are listed in the **Lifecycle Rules** list on a bucket's detail screen. To navigate to a bucket's detail screen, click **Storage > Object Store > Bucket > bucket name**.

Figure 5.16 Viewing the Lifecycle Rules List

Lifecycle Rules							
Rule ID▲	Enabled	Prefix	Keep Current Version For	Keep Current Version Until	Keep Previous Version For	Abort Multipart Uploads After	Cleanup Expired Object Delete Marker
1	True	abc	5 days	-	2 days	-	False
2	True	abc	-	-	3 days	-	False

The **Lifecycle Rules** list displays the following information:

- **Rule ID** - The object lifecycle rule ID. This ID is either created by the system or specified by the user. If system-assigned, the ID is in the format `fbRuleIdX`, where X increases by one with each rule. The rule ID must be unique for the given bucket.

- **Enabled** - Indicates if the object lifecycle rule is enabled (**True**) or disabled (**False**).
- **Prefix** - An optional field that causes the rule to only be applied to objects beginning with this prefix.

 **Note:** Only one of **Keep Current Version For** and **Keep Current Version Until** value can be configured per rule. Entering a value in one causes the value for the other to be cleared.


- **Keep Current Version For** - The duration that the current object version is retained expressed in a whole number day, for example: 10 days, 4 weeks, etc.
- **Keep Current Version Until** - The date that the current object version is retained expressed in the format YYYY-MM-DD, for example: 2021-05-31.
- **Keep Previous Version For** - The duration that the non-current object version is retained expressed in a whole number day, for example: 10 days, 4 weeks, etc.
- **Abort Multipart Uploads After** - The duration that incomplete multi-part uploads are aborted expressed in a whole number day, for example: 10 days, 4 weeks, etc.
- **Cleanup Expired Objects Delete Marker** - - Displays whether object versions are deleted (**True**) or not (**False**) after expiration.

Creating an Object Lifecycle Rule

 **Note:** The object lifecycle rule is automatically enabled upon creation.

To create an object lifecycle rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account that contains the bucket to which you wish to create a lifecycle configuration.
- 2 From the **Buckets** panel, select the bucket for which you wish to create an object lifecycle rule.
- 3 Click the **Add (+)** button in the heading of the **Lifecycle Rules** list. The **Create Lifecycle Rule** pop-up window appears.



4 (Optional) Enter a unique ID for the rule in the **Rule ID** field. The ID cannot exceed 255 characters.

5 (Optional) Enter an object prefix in the **Prefix** field. The prefix cannot exceed 1,024 characters.

An object prefix filter objects within the bucket to which the rule is applied. For example, if `logs` is entered, the created rule will apply to only objects with the prefix `logs`, such as `logs/20201020` or `logstoday`, but not `/logs`.

 **Note:** Only one of **Keep Current Version For** and **Keep Current Version Until** value can be configured per rule. Entering a value in one causes the value for the other to be cleared.

6 In the **Keep Current Version For** field, enter the retention period for the current object version. The value must be entered in the format `N[m/h/d/w]`, where `N` is an integer followed by `m` (minutes), `h` (hours), `d` (days), or `w` (weeks). The minimum valid retention period is one day. Note that the value entered must be in a multiple of 1 day increments. For example: `24h`, `1440m`, etc.

7 In the **Keep Current Version Until** field, enter the current object version's retention end date. The value must be entered in the format `YYYY-MM-DD`. For example `2021-05-31`.

8 In the **Keep Previous Version For** field, enter the retention period for the non-current object version. The value must be entered in the format `N[m/h/d/w]`, where `N` is an integer followed by `m` (minutes), `h` (hours), `d` (days), or `w` (weeks). The minimum valid retention period is one day. Note that the value entered must be in a multiple of 1 day increments. For example: `24h`, `1440m`, etc.

For example, entering `15d` results in the retention of the non-current object version for 15 days; when the 15 day retention period is reached, that version will be deleted from the system within 24 hours.

9 In the **Abort Incomplete Multipart Upload After** field, enter the duration after which incomplete multi-part uploads are aborted. The value must be entered in the format `N[m/h/d/w]`, where `N` is an integer followed by `m` (minutes), `h` (hours), `d` (days), or `w` (weeks). The minimum valid period is one day. Note that the value entered must be in a multiple of 1 day increments. For example: `24h`, `1440m`, etc.

- 10 Review the summary and click **Create**.

Editing an Object Lifecycle Rule

To edit an object lifecycle rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account that contains the bucket to which you wish to edit a lifecycle configuration.
- 2 From the **Buckets** panel, select the bucket with the object lifecycle rule you wish to edit.
- 3 From the **Lifecycle Rules** list, click the **More Options** button in the row of the object lifecycle rule you wish to edit and select **Edit**. The **Edit Lifecycle Rule** pop-up window appears.
- 4 (Optional) Enter or change the object prefix in the **Prefix** field. The prefix cannot exceed 1,024 characters.

An object prefix filter objects within the bucket to which the rule is applied. For example, if `/tmp/` is entered, the created rule will apply to only objects with the prefix `/tmp/`.

- 5 In the **Keep Current Version For** field, enter the retention period for the current object version. The value must be entered in the format `N[m/h/d/w]`, where `N` is an integer followed by `m` (minutes), `h` (hours), `d` (days), or `w` (weeks). The minimum valid retention period is one day. Note that the value entered must be in a multiple of 1 day increments. For example: 24h, 1440m, etc.
- 6 In the **Keep Current Version Until** field, enter the current object version's retention end date. The value must be entered in the format `YYYY-MM-DD`. For example 2021-05-31.
- 7 (Optional) In the **Keep Previous Version For** field, enter the period for the non-current object version will be retained. The value must be entered in the format `N[m/h/d/w]`, where `N` is an integer followed by `m` (minutes), `h` (hours), `d` (days), or `w` (weeks). Note that the value entered must be in a multiple of 1 day increments. For example: 24h, 1440m, etc.

For example, entering 15d results in the retention of the non-current object version for 15 days; when the 15 day retention period is reached, that version will be deleted from the system within 24 hours.

- 8 In the **Abort Incomplete Multipart Upload After** field, enter the duration after which incomplete multi-part uploads are aborted. The value must be entered in the format `N[m/h/d/w]`, where `N` is an integer followed by `m` (minutes), `h` (hours), `d` (days), or `w` (weeks). The minimum valid period is one day. Note that the value entered must be in a multiple of 1 day increments. For example: 24h, 1440m, etc.
- 9 Review the summary and click **Save**.

The object lifecycle rule is updated with your changes.

Enabling and Disabling Object Lifecycle Rules

Object lifecycle rules are enabled by default when created. You can disable or enable lifecycle rules at any time.

To enable or disable an object lifecycle rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account that contains the bucket to which you wish to edit a lifecycle configuration.
- 2 From the **Buckets** panel, select the bucket with the object lifecycle rule you wish to edit.
- 3 From the **Lifecycle Rules** list, click the **More Options** button in the row of the object lifecycle rule you wish to enable or disable.
 - a If you selected **Enable**, the **Enable Lifecycle Rule** pop-up window appears. Click **Enable**.
 - b If you selected **Disable**, the **Disable Lifecycle Rule** pop-up window appears. Click **Disable**.

Deleting an Object Lifecycle Rule

To delete an object lifecycle rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account to which the bucket whose object lifecycle rule you wish to delete belongs.
- 2 From the **Buckets** panel, select the bucket with the object lifecycle rule you wish to delete.
- 3 From the **Lifecycle Rules** list, click the **More Options** button in the row of the object lifecycle rule you wish to delete and select **Delete**. The **Delete Lifecycle Rule** pop-up window appears.
- 4 When prompted for confirmation, click **Delete**.

The object lifecycle rule is deleted from the bucket.

Creating a Bucket Access Policy

Currently you can create access policy rules that allow the s3:GetObject action to be performed on all objects in the bucket. Additional actions may be provided in a future release.

To create a bucket access policy for a bucket, perform the following:

- 1 From the **Storage > Object Store** page, click the account in which the public bucket resides.
- 2 From the **Buckets** panel, click the public bucket to which you wish to add a bucket access policy. The bucket's detail page appears.

- 3 In the header of the **Bucket Policies** panel, click **More Options > Create Bucket Access Policy**. The **Create Bucket Access Policy with Rule for <bucket name>** pop-up window appears.
- 4 Enter a name for the policy.
- 5 Click **Create**.

The **Bucket Policies** panel is updated with the new policy.

Deleting a Bucket Access Policy

Bucket access policies can be deleted from buckets. Note that if a bucket access policy is deleted from a bucket, that bucket will no longer be publicly available.

To delete a bucket access policy, perform the following:

- 1 From the **Storage > Object Store** page, click the account in which the bucket with the bucket access policy resides.
- 2 From the **Buckets** panel, click the public bucket from which you wish to delete a bucket access policy. The bucket's detail page appears.
- 3 In the **Bucket Policies** panel, click the **Delete** button on the same row as the bucket access policy you wish to delete.

The **Bucket Policies** panel is updated with the new policy.

Deleting Bucket Access Policy Rules

The rules in a bucket access policy can be deleted. Note that if the rules are deleted from a bucket access policy, that bucket may no longer be publicly available.

To delete the rules from a bucket access policy, perform the following:

- 1 From the **Storage > Object Store** page, click the account in which bucket with the bucket access policy resides.
- 2 From the **Buckets** panel, click the public bucket from which you wish to remove the rules from a bucket access policy. The bucket's detail page appears.
- 3 Click the name of the bucket access policy. The **Bucket Access Policy Rules** panel appears.
- 4 Click the **Delete** button on the same row as the bucket access policy rules you wish to delete.
- 5 Click **Remove** when prompted.

Creating a CORS Policy

Cross-Origin Resource Sharing (CORS) policies can be created for buckets.

Currently all headers, methods, and origins are allowed with the creation of a CORS policy. Additional actions may be provided in a future release.

To create a CORS policy for a bucket, perform the following:

- 1 From the **Storage > Object Store** page, click the account in which the bucket resides.
- 2 From the **Buckets** panel, click the bucket to which you wish to add a CORS policy. The bucket's detail page appears.
- 3 In the header of the **Bucket Policies** panel, click **More Options > Create Cors Policy**. The **Create CORS Policy with Rule for <bucket name>** pop-up window appears.
- 4 Enter a name for the policy.
- 5 Click **Create**.

The **Bucket Policies** panel is updated with the new policy.

Deleting a CORS Policy

To delete a CORS policy, perform the following:

- 1 From the **Storage > Object Store** page, click the account in which bucket with the CORS policy resides.
- 2 From the **Buckets** panel, click the bucket from which you wish to delete a CORS policy. The bucket's detail page appears.
- 3 In the **Bucket Policies** panel, click the **Delete** button on the same row as the CORS policy you wish to delete.

The **Bucket Policies** panel is updated with the new policy.

Deleting CORS Policy Rules

To delete the rules from a CORS policy, perform the following:

- 1 From the **Storage > Object Store** page, click the account in which bucket with the bucket access policy resides.
- 2 From the **Buckets** panel, click the bucket from which you wish to remove the rules from a CORS policy. The bucket's detail page appears.

- 3 Click the name of the CORS policy. The **Cross Origin Resource Sharing Policy Rules** panel appears.
- 4 Click the **Delete** button on the same row as the CORS policy rules you wish to delete.
- 5 Click **Remove** when prompted.

Destroying a Bucket

When a bucket is destroyed, it enters a 24-hour eradication pending period. During this time, the bucket and its contents can be recovered. Note that also during this time, data in the bucket is inaccessible. When the 24-hour eradication pending period has lapsed, Purity//FB starts reclaiming storage space. During this time, the bucket can no longer be recovered.

To destroy a bucket:

- 1 From the **Storage > Object Store** page, click the account from which you wish to destroy a bucket.
- 2 From the **Buckets** panel, click the trash can icon on the same row as the bucket you wish to destroy. The **Destroy Bucket** pop-up window appears.

 **Note:** By default, the buckets 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 destroy multiple buckets, click **More Options > Destroy**. The **Destroy Buckets** pop-up window appears. From the **Existing Buckets** list, select the buckets 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 buckets by clicking the checkbox next to the search box. All selected file systems appear in the **Selected Buckets** list.

- 3 Click **Destroy**.

Recovering a Bucket

Destroyed buckets have 24 hours to be recovered. Once 24 hours has expired, the system begins eradicating the bucket, at which point the bucket is no longer recoverable.

To recover a destroyed bucket:

- 1 From the **Storage > Object Store** page, click the account from which you wish to recover a destroyed bucket.

- 2 From the **Destroyed Buckets** panel, click the clock icon on the same row as the bucket you wish to recover. The **Recover Bucket** pop-up window appears.

 **Note:** By default, the buckets 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 buckets, click **More Options > Recover**. The **Recover Buckets** pop-up window appears. From the **Destroyed Buckets** list, select the buckets 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 buckets by clicking the checkbox next to the search box. All selected file systems appear in the **Selected Buckets** list.

- 3 Click **Recover**.

Eradicating a Bucket

During the eradication pending period, you can manually eradicate the destroyed buckets to reclaim physical storage space. Once reclamation starts, the destroyed buckets can no longer be recovered.

To eradicate a destroyed bucket:

- 1 From the **Storage > Object Store** page, click the account from which you wish to eradicate a destroyed bucket.
- 2 From the **Destroyed Buckets** panel, click the clock icon on the same row as the bucket you wish to eradicate. The **Eradicate Bucket** pop-up window appears.

 **Note:** By default, the buckets 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 buckets, click **More Options > Eradicate**. The **Eradicate Buckets** pop-up window appears. From the **Destroyed Buckets** list, select the buckets 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 buckets by clicking the checkbox next to the search box. All selected file systems appear in the **Selected Buckets** list.

- 3 Click **Eradicate**.

Chapter 6

The Purity//FB GUI Protection Page

The **Protection** page displays and manages the array's file system snapshots and replica links.

The **Protection** page provides the following four tabs:

- **Snapshots:** Displays and allows you to manage snapshots from all file systems.
- **File Replication:** Displays and allows you to manage all file replica links used to schedule file system snapshot replication.
- **Object Replication:** Displays and allows you to manage all object replica links used to schedule object data replication.

Figure 6.1 Purity//FB Protection Page

Name	Source Location	Source Name	Policies	Created
fbsit-fs-5940.2023_11_04_22_59	durin	fbsit-fs-5940	fbsit-fanout-extended-5	2023-11-04 22:59:40
fbsit-fs-5941.2023_11_04_22_59	durin	fbsit-fs-5941	fbsit-fanout-extended-1	2023-11-04 22:59:40
fbsit-fs-5907.2023_11_04_22_46	durin	fbsit-fs-5907	fbsit-fanout-extended-2	2023-11-04 22:46:32
fbsit-fs-5909.2023_11_04_22_46	durin	fbsit-fs-5909	fbsit-fanout-extended-4	2023-11-04 22:46:32
fbsit-fs-5912.2023_11_04_22_46	durin	fbsit-fs-5912	fbsit-fanout-extended-2	2023-11-04 22:46:32
fbsit-fs-5548.2023_11_04_22_36	durin	fbsit-fs-5548	fbsit-fanout-extended-3	2023-11-04 22:36:31
fbsit-fs-5542.2023_11_04_22_30	durin	fbsit-fs-5542	fbsit-fanout-extended-2	2023-11-04 22:30:31
fbsit-fs-5515.2023_11_04_21_52	durin	fbsit-fs-5515	fbsit-fanout-extended-5	2023-11-04 21:52:29
fbsit-fs-5516.2023_11_04_21_52	durin	fbsit-fs-5516	fbsit-fanout-extended-1	2023-11-04 21:52:29
fbsit-fs-5517.2023_11_04_21_52	durin	fbsit-fs-5517	fbsit-fanout-extended-2	2023-11-04 21:52:29

Snapshots

A file system snapshot is a point-in-time copy of a file system. File system snapshots can be created and managed manually or by using snapshot policies. A file system snapshot policy consists of one or more rules that govern when snapshots are created or replicated, and how long they are retained. When a

policy is attached to a file system, it will create local file system snapshots. When one or more policies is attached to a file system replica link, it will replicate file system snapshots to a connected array.

For maximum flexibility, file system replication links can have different policies attached for the replication of specific file systems. In instances where multiple policies attached to a replication link are triggered at the same time, a single *shared file system snapshot* is created for the link.

For example, a replica link for file system `fs1` from `array1` to `array2` has the following policies attached:

- `Policy1` - This policy's rule creates a snapshot every 2 hours and 4 hours
- `Policy2` - This policy's rule creates a snapshot every 4 hours

A snapshot is created after 2 hours for policy `Policy1`. Because its next snapshot is scheduled to be created 2 hours later and `Policy2` is set to create a snapshot in 4 hour intervals, at the next snapshot creation interval for `Policy1` and the snapshot creation interval for `Policy2`, a single shared snapshot is created. After 2 hours, another snapshot will be created per `Policy1`. At the next 2 hour interval, another shared snapshot will be created.

In fan-out replication setups, file system snapshots and common file system snapshots can either be manually created, or created and propagated to targets via file system replica links where there is at least one snapshot policy attached. A *common* snapshot is a file system snapshot that is created on all eligible fan-out file replication target arrays. Because a common snapshot is an identical snapshot that is created on all target arrays, a common snapshot provides optimal disaster recovery and migration through snapshot policy synchronization, which reduces the time needed for the resumption of replication from the new source to targets.

When the rules for snapshot policies are created, you must specify two values (see [Creating Snapshot Policies](#)):

- **Create or replicate 1 snapshot every** - the frequency at which snapshots are created or replicated.
- **And keep for** - the retention period for the snapshot.
- You can optionally specify an **At** value for the time of day for snapshot creation.

In fan-out replication setups, a common snapshot is created and synchronized based on the longest **every** value (the frequency at which snapshots are created or replicated) among the shortest rules.

Example 1

A replica link (`link1`) for file system `fs1` from `array1` to `array2` has a policy (`Policy1`) with the following rules attached:

- Rule 1 - This rule creates a snapshot every 15 minutes with a retention period of 30 minutes.
- Rule 2 - This rule creates a snapshot every 60 minutes with a retention period of 15 minutes.

A replica link (`link2`) for file system `fs1` from `array1` to `array3` has a policy (`Policy2`) with the following rules attached:

- Rule 1 - This rule creates a snapshot every 5 minutes with a retention period of 30 minutes.
- Rule 2 - This rule creates a snapshot every 30 minutes with a retention period of 15 minutes.

For this example:

- The shortest rules of `Policy1` and `Policy2` create snapshots every 15 minutes and every 5 minutes, respectively.
- Between these two policies, `Policy1` has the longest **every** value at 15 minutes.
- A common snapshot of file system `fs1` is created on, and synchronized across, each target every 15 minutes using Rule 1 from `Policy1`.
- Because the replica link `link2` from `array1` to `array3` does not detect the policies and rules of the replica link `link1` from `array1` to `array2`, the retention of snapshots created from the policies on `link2` will follow its own rule, in this case 30 minutes.

Example 2

A replica link (`link1`) for file system `fs1` from `array1` to `array2` has the following policies and rules attached :

- `Policy1` - This policy has a single rule that creates a snapshot every 15 minutes with a retention period of 30 minutes.
- `Policy2` - This policy has a single rule that creates a snapshot every 60 minutes with a retention period of 20 minutes.

A replica link (`link2`) for file system `fs1` from `array1` to `array3` has the following policies and rules attached:

- `Policy3` - This policy has a single rule that creates a snapshot every 5 minutes with a retention period of 30 minutes.
- `Policy4` - This policy has a single rule that creates a snapshot every 30 minutes with a retention period of 30 minutes.

For this example:

- The system checks for the shortest rule for each policy. Because there is only one rule per policy, by default, each rule is the shortest rule for its respective policy. The system then determines the longest **every** value.
- Between these policies, `Policy2` has the longest **every** value at 60 minutes.
- A common snapshot of file system `fs1` is created on, and synchronized across, each target every 60 minutes using `Policy2`.
- Because the replica link `link2` does not detect the policies and rules of the replica link `link1`, the retention of snapshots created from the policies on `link2` will follow its own rule, in this case 30 minutes.

 **Note:** The retention of snapshots on the target is based on the target's local state. There is no guarantee that a common snapshot will be treated with the longest retention time. To ensure the protection of common snapshots, the shortest retention time must be set up with sufficient overlap within the policy rules.

Viewing Snapshots

By default, the **Protection** page displays the **Snapshots** tab. File system snapshots are created from the **Storage > File Systems** page (see [Creating a File System Snapshot](#)).

The top of the **Snapshots** tab contains the **File System Snapshots** panel, which displays general information about file system snapshots in a table with the following columns:

- **Name:** The name of the file system snapshot.
- **Source Location:** The name of the array from where the source file system was replicated.
- **Source Name:** The name of the source file system.

- **Policy** or **Policies**: The snapshot policy or policies attached to the source file system. For file systems with multiple attached policies, click the down arrow next to the policy to expand and view the list of all attached policies.
- **Created**: The creation date and time of the snapshot.

Each column can be sorted by clicking the column title.

To view information about the replication transfer status of the file system snapshots, click **Transfer** from the title bar of the **File System Snapshots** panel. Note that replication transfer status information is shown only if replication has been configured.

Once selected, the table displayed in the **File System Snapshots** panel changes to display the following columns:

- **Name**: The name of the file system snapshot.
- **Direction**: The direction of replication, either inbound (←) or outbound (→).
- **Remote**: The name of the connected array.
- **Remote Snapshot**: The name of the remote (target) file system snapshot.
- **Started**: The date and time replication started.
- **Completed**: The date and time replication completed.
- **Progress**: The percentage of replication completed.
- **Data Transferred**: The amount of data replicated.
- **Status**: The current replication status.

As with the **General** table, each column in the **Transfer** table can be sorted by clicking the column title.

The bottom of the **Snapshots** tab contains the **Destroyed Snapshots** panel, which displays all destroyed file system snapshots. Destroyed snapshots can be recovered within a 24-hour period. Once the 24-hour period has passed, the snapshot is permanently destroyed (eradicated). The **Time Remaining** column displays the time remaining for the snapshot to be restored.

Note that with scheduled file system snapshot replication, if there is a delay with a snapshot completing the transfer of one snapshot, the next snapshot (per the policy schedule) is not created until the previous snapshot transfer is completed. For example, if snapshots are scheduled every 30 minutes and

a snapshot completes transferring in 31 minutes, the next snapshot is created once that snapshot is completely transferred. The creation times as shown in the **File System Snapshots** panel will reflect this.

Destroying File System Snapshots

You can destroy a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To destroy one or more file system snapshots, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 From the table in the **File System Snapshots** panel, locate the file system snapshot you either wish to destroy, or is the point from which earlier snapshots are to be destroyed.
- 3 If you wish to destroy the snapshot, perform the following:
 - a Click the **More Options** button located at the end of that snapshot's row and select **Destroy**. The **Destroy Snapshot** pop-up window appears.
 - b Click **Destroy**
- 4 If you wish to destroy multiple snapshots, perform the following:

 **Note:** By default, the snapshots 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 **More Options > Destroy** in the **File System Snapshots** panel title bar. The **Destroy File System Snapshots** pop-up window appears.
 - a From the **File System Snapshots** list, select the snapshots you wish to destroy. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
 - b Click **Destroy**.
- 5 If you wish to use the snapshot as the point from which earlier snapshots are destroyed, perform the following:

- a Click the **More Options** button located at the end of that snapshot's row and select **Destroy older**. The **Destroy Snapshots older than...** pop-up window appears.

 **Note:** By default, the buckets 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 Select the snapshots you wish to destroy from the **Older File System Snapshots** panel. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- c Click **Destroy**.

The destroyed file system snapshot appears under the **Destroyed Snapshots** panel on the **Snapshots** tab. Destroyed snapshots can be recovered within a 24-hour period. Once the 24-hour period has passed, the snapshot is permanently destroyed (eradicated).

Restoring a File System to a Snapshot

Purity//FB allows you to restore a file system to a given point in time from a snapshot (via file system rollback). Restoring a file system to a snapshot overwrites the contents of that file system with data from the snapshot.

File system rollback is available for snapshots created on a FlashBlade array running Purity//FB 3.0.0 or later.

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

 **CAUTION:** When restoring to a snapshot, the file system's configured quota, as well as any configured user/group quotas, do not change to reflect the limit when the snapshot was taken. This can cause the live space usage to increase, resulting in subsequent write failures due to exceeding the quota. You may need to re-adjust your quota settings accordingly.

 **Note:** It is highly recommended that there be no I/O on the file system when performing a restore.

You can restore a file system from a snapshot from either the **Storage** or **Protection** page.

 **Note:** If there are snapshots that are newer than the snapshot you wish to use as the restore point, the newer snapshots must be first destroyed and eradicated before you can proceed with the file system restoration. See [Destroying and Eradicating Newer Snapshots](#).

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To restore a file system from a snapshot, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 From the table in the **File System Snapshots** panel, locate the file system snapshot to be used to restore the file system.
- 3 Click the **More Options** button located at the end of row of the snapshot to be used for restore and select **Restore**. The **Restore File System from Snapshot** pop-up window appears.
- 4 Click **Restore**.

Destroying and Eradicating Newer Snapshots

 **Note:** Destroyed snapshots can be recovered within a 24-hour period. Once the 24-hour period has passed, the snapshot is permanently destroyed (eradicating).

When restoring a file system from a snapshot, if there are snapshots that are newer than the snapshot you wish to use as the restore point, the newer snapshots must be first destroyed and eradicated before you can proceed with the file system restoration.

You can destroy and eradicate snapshots from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 From the table in the **File System Snapshots** panel, locate the file system snapshot to be used to restore the file system.
- 3 Click the **More Options** button located at the end of that snapshot's row and select **Destroy newer**. The **Destroy Snapshots newer than...** pop-up window appears.
- 4 In the **Newer File System Snapshots** panel, select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- 5 Click **Destroy**. The snapshot(s) appear in the **Destroyed Snapshots** panel (directly below the **File System Snapshots** panel).

- 6 Click **More Options > Eradicate** in of the **Destroyed Snapshots** panel. The **Eradicate File System Snapshots** pop-up window appears.

 **Note:** By default, the snapshots 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.

- 7 From the **Destroyed File System Snapshots** list, select the snapshots you wish to eradicate. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the checkbox next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- 8 Click **Eradicate**.

Removing Snapshot Policies from Snapshots

 **Note:** Once all snapshot policies are removed from a file system snapshot, the snapshot will no longer be managed by the policies and will have to be manually destroyed.

To remove a snapshot policy from a file system snapshot, perform the following:

- 1 To remove a single snapshot policy from a file system snapshot:
 - a From the **Protection > Snapshots** tab, locate the file system snapshot from which you wish to remove the snapshot policy.
 - b Click **More Options > Remove Policy** at the end of that snapshot's row. The **Remove Policy from Snapshot** pop-up window appears.
 - c Click **Remove**.
- 2 To remove one or more snapshot policies from a file system snapshot:
 - a From the **Protection > Snapshots** tab, locate the file system snapshot from which you wish to remove one or more snapshot policies.
 - b Click **More Options > Remove Policies** at the end of that snapshot's row. The **Remove Policies from Snapshot** pop-up window appears.
 - c From the **Policies** panel, select one or more snapshot policies from the list. The selected policy (or policies) appears in the **Selected Policies** panel.
 - d Click **Remove**.

Recovering File System Snapshots

When a file system snapshot is destroyed, you have 24 hours to recover the data. This is known as the eradication pending state. When the eradication pending state expires, the snapshot will be eradicated and you will no longer be able to retrieve the snapshot data.

You can recover a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To recover a previously destroyed file system snapshot, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 If you wish to recover a single snapshot:
 - a From the **Destroyed Snapshots** panel, locate the file system snapshot you wish to recover.
 - b Click **More Options > Recover** at the end of the row of the file system snapshot you wish to recover.
 - c Click **Recover**.
- 3 If you wish to recover multiple file system snapshots:

 **Note:** By default, the snapshots 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 **More Options > Recover** in the **Destroyed Snapshots** panel title bar. The **Recover File System Snapshots** pop-up window appears.
- b From the **Destroyed File System Snapshots** list, select the snapshots you wish to recover. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the check box next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- c Click **Recover**.

Eradicating File System Snapshots

During the eradication pending period, you can manually eradicate the destroyed file system to reclaim physical storage space occupied by the destroyed snapshot. Once reclamation starts, the destroyed snapshots can no longer be recovered.

You can eradicate a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

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

To eradicate a destroyed snapshot:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's destroyed snapshots. If you are on the **Protection** page, select **Protection > Snapshots**.
- 2 If you wish to eradicate a single snapshot:
 - a In the **Destroyed Snapshots** panel under **File System Snapshots**, click **More Options > Recover** at the end of the row of the snapshot you want to eradicate. The **Eradicate Snapshot** pop-up window appears.
 - b Click **Eradicate**. The array will begin reclaiming the eradicated space of the snapshot.
- 3 If you wish to eradicate multiple snapshots:


 **Note:** By default, the snapshots 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 **More Options > Eradicate** in of the **Destroyed Snapshots** panel. The **Eradicate File System Snapshots** pop-up window appears.
- b From the **Destroyed File System Snapshots** list, select the snapshots you wish to eradicate. You can optionally enter a partial name in the search box to filter the snapshot list. You can also optionally select all snapshots by clicking the checkbox next to the search box. All selected snapshots appear in the **Selected File System Snapshots** list.
- c Click **Eradicate**. The array will begin reclaiming the eradicated space of the snapshots.

File Replication

File replica links are required for replicating file system snapshots between two or more connected FlashBlades. The **File Replication** tab is used to view and manage file replica links.

Figure 6.2 File Replication



Local File System	Direction	Remote	Remote File System	Policy	Status	Recovery Point	Lag
acl-b-1	←	moonraker	acl-b-1	moonraker:daily	idle	2022-01-03 11:57:20	1 h 18 m
rep-sf-1-tgt	←	moonraker	rep-sf-1	moonraker:daily	unhealthy	2021-02-26 12:28:16	311 d 00 h 47 m
smb-share	→	moonraker	smb-share	daily	idle	2022-01-03 12:05:20	1 h 10 m

The **File Replication** tab displays a list of all configured file replica links:

- **Local File System:** The name of the local (source) file system.
- **Direction:** The direction of replication, either inbound (←) or outbound (→).
- **Remote:** The name of the connected array.
- **Remote File System:** The name of the remote (target) file system.
- **Policy:** The name of the applied replication policy.
- **Status:** The current replication status.
 - **idle:** No errors, no current replication.
 - **replicating:** No errors, currently attempting replication.
 - **unhealthy:** No connectivity.
- **Recovery Point:** The creation time of the last snapshot that was completely replicated.
- **Lag:** The amount of time, if any, the replication target is behind the source, and is the time difference between the current time and the recovery point.

Clicking the file replica link on the local file system name directs you to the **Storage > File Systems** page, which provides detailed information about the file system to which the replica link is attached.

Creating a File System Replica Link

 **Note:** When creating a file system replica link, you have the option to connect a FlashBlade array if you have not already done so. In order to connect an array you must have previously created a replication network interface and a connection key. See [Creating a Network Interface](#) and [Creating a Connection Key](#) for instructions.

File system replica links are used to configure file system replication between two connected FlashBlade arrays.

Fan out file system replication allows file system replication to more than one target. In order to replicate to multiple targets, file system replica links to each target must be created. The following are supported:

- A maximum of 1,000 file systems can have replica links.
- A file system can have a maximum of 10,000 outgoing replica links and a maximum of 1,000 incoming replica links.

To create a file system replica link, perform the following:

- 1 From the **Protection > File Replication** page, click the add (+) button in the heading of the **File System Replica Links** list. The **Step 1: Create File System Replica Link** pop-up window appears.
- 2 In the **Step 1: Create File System Replica Link** pop-up window, perform the following:
 - a From the **Local File System** list, select the local (source) file system to be replicated.
 - b From the **Remote Array** list, select the remote (target) array.
 - c (Optional) Enter a name for the remote file system to which data from the local file system will be replicated in the **Remote File System** field.
 - d (Optional) Click **Connect FlashBlade Array** if you did not previously connect an array. See [Connecting Arrays](#) for instructions.
 - e Click **Create**. The **Step 2: Create Fanout Replica Links** pop-up window appears
- 3 In the **Step 2: Create Fanout Replica Links** pop-up window, perform the following:
 - a If you are not setting up fan out file replication, click **Skip**. The **Add Snapshot Policies** pop-up window appears (see Step 4).

- b Select the connected FlashBlade arrays from the **Available FlashBlade Arrays** column to which you wish to create replication links. Selected arrays appear in the **Selected FlashBlade Arrays** column.
 - c If there is an existing demoted remote file system with the same name as the source file system on the remote array, selecting the **Relink existing file systems** option will reconnect and promote the file system.
 - d Click **Create Links**. The **Add Snapshot Policies** pop-up window appears.
- 4 In the **Add Snapshot Policies** pop-up window, perform the following:
 - a Select one or more snapshot policy from the **Available Policies** column. Selected policies appear in the **Selected Policies** column. If a policy does not exist, click **Create a New Snapshot Policy**. Follow the policy creation instructions as described in [Creating Snapshot Policies](#). If you do not add a policy, the replica link can still be created. However, replication will not occur until a policy is added to the link.
 - b Click **Add**.

Adding a Policy to a File Replica Link

A snapshot policy can be added to any outbound (➔) replica link.

To add a snapshot policy to a replica link, perform the following:

- 1 From the **Protection > File Replication** page, locate the file replica link to which you wish to add a policy and click the **More Options** button located at the end of that replica link's row.
- 2 Select **Add Policies**. The **Add Policies** pop-up window appears.
- 3 Select the policy or policies you wish to add to the replica link from the **Available Policies** panel.
- 4 Click **Add**.

Removing a Policy from a File Replica Link

A snapshot policy can be removed from any outbound (➔) replica link.

To remove a snapshot policy from a replica link, perform the following:

- 1 From the **Protection > File Replication** page, locate the file replica link from which you wish to remove a policy and click the **More Options** button located at the end of that replica link's row.
- 2 Select **Remove Policies**. The **Remove Policies** pop-up window appears.
- 3 Select the policy or policies you wish to remove from the replica link from the **Available Policies** panel.

- 4 Click **Remove**.

Creating Fan Out File System Replica Links

Fan out file system replication allows file system replication to more than one target. In order to replicate to multiple targets, file system replica links to each target must be created.

To create fan out file system replica links, perform the following:

- 1 From the **Protection > File Replication** page, locate the local file system that you wish to replicate to multiple targets from the **File System Replica Links** panel, and click **More Options > Create Fanout Replica Links**. The **Create Fanout Replica Links** pop-up window appears.
- 2 Select the connected FlashBlade arrays from the **Available FlashBlade Arrays** column to which you wish to create replication links. Selected arrays appear in the **Selected FlashBlade Arrays** column.
- 3 If there is an existing demoted remote file system on the remote array, selecting the **Relink existing file systems** option will reconnect and promote the file system.
- 4 Click **Create Links**.

Deleting a File System Replica Link

File system replica links are used to configure file system replication between two connected FlashBlade arrays. When a file system replica link is deleted, replication of the file system to the target FlashBlade stops and any in-progress transfers are cancelled.

An array cannot be disconnected while in use by file system replica links. You must delete all file system replica links to the remote array. Once file system replica links are deleted, the array can then be disconnected. Note that replica link propagation to the remote array is asynchronous and can be delayed.

- 1 To delete a file system replica link:
 - a From the **Protection > File Replication** tab, locate the file system replica link you wish to delete from the **File System Replica Links** panel.
 - b From the same row as the file system replica link to be deleted, click **More Options > Delete**. The **Delete File Replica Link** dialog box appears.
 - c Click **Delete**.
- 2 To delete multiple file system replica links:
 - a From the **Protection > File Replication** tab, click **More Options > Delete** from the **File System Replica Links** panel. The **Delete File Replica Links** pop-up window appears.

- b** From the **Available Links** panel, select the file system replica links you wish to delete. Each selected file system replica link appears in the **Selected Links** panel.
- c** Click **Delete**.

Object Replication


Object replica links and remote credentials are required for replicating object data between FlashBlade arrays and S3 targets.

 **Note:** Objects created after the creation of the object replica link are guaranteed to be replicated from the source bucket to the target bucket. Objects that were created shortly before the creation of the object replica link may also be replicated.

Objects can be replicated from a bucket on a FlashBlade (source) to a remote target bucket residing on another FlashBlade or an S3 target using object replica links. A replica link is the relationship between a source bucket and target bucket, and describes the properties of that connection. During replication, all newly written objects in the source bucket are replicated to the target bucket.

In object replication, up to two replica links can be created from the source bucket to different target buckets. If replicating to multiple targets, you can set up links in *fan out* or *cascading* configurations. In a fan out configuration, a single source bucket has replica links to different target buckets. For example, A→B and A→C. In a cascading configuration, the source bucket has a replica link to a target bucket, which in turn has a replica link to another target bucket. For example, A→B→C. Note that in this example, only B requires the cascading option to be enabled.

The **Object Replication** tab is used to view and manage object replica links and remote credentials.

Figure 6.3 Object Replication


Local Bucket	Direction	Remote Bucket	Remote Credentials	Cascading	Status	Recovery Point	Lag
ko0	→	ko0	robin/robin	False	● paused	2021-12-07 20:53:29	26 d 16 h 24 m
ko10	→	ko10	robin/robin	False	● paused	2021-11-19 01:18:43	45 d 11 h 59 m
ko11	→	ko11	robin/robin	False	● paused	2021-11-30 05:36:04	34 d 07 h 41 m
ko111	→	ko111	robin/robin	False	● paused	2021-12-06 19:46:15	27 d 17 h 31 m
ko112	→	ko112	robin/robin	False	● paused	2021-12-06 17:08:46	27 d 20 h 09 m
ko113	→	ko113	robin/robin	False	● paused	2021-12-07 18:20:43	26 d 18 h 57 m
ko114	→	ko114	robin/robin	False	● paused	2021-12-07 18:20:37	26 d 18 h 57 m

The **Object Replication** tab contains two panels, **Bucket Replica Links** and **Remote Credentials**.

Clicking **General** (default view) in the **Bucket Replica Links** panel displays a list of all configured object replica links:

- **Local Bucket:** The name of the local (source) bucket.
- **Direction:** The direction of replication, outbound (→).
- **Remote Bucket:** The name of the of the remote (target) bucket.
- **Remote Credentials:** The name of the remote credentials is the combined remote array name or S3 target with the remote credential name. For example `array2/cred1`.
- **Cascading:** Whether the cascading option was selected (**True**) or not (**False**) during link creation.
- **Status:** The current replication status.
 - **replicating:** No errors, currently attempting replication.
 - **paused:** Data transfer is currently halted due to administrator action.
 - **unhealthy:** Replication encountered errors. For example, connectivity issues, denial of access, target bucket versioning not enabled, etc.
- **Recovery Point:** A point in time where all object creations, updates, and deletions that occurred after the creation of the replica link and before this point in time are guaranteed to have been replicated.

Object creations, updates, and deletions that occurred before the creation of a replica link are not replicated. Those that occurred after the recovery point may or may not have been replicated.

- **Lag:** The amount of time, if any, the replication target is behind the source, and is the time difference between the current time and the recovery point.

Clicking **Backlog** in the **Object Replica Links** panel displays information about all object data pending replication. The information provided in the panel is also summarized for quick reference directly below the **Object Replication** tab header.

Figure 6.4 Object ReplicationHeader




Figure 6.5 Object Replica Links Backlog View

Bucket Replica Links							General	Backlog	1-10 of 872	<	>	+
Local Bucket ▲	Direction	Remote Bucket	Lag	Backlog Size ?	PUT Ops	DELETE Ops	Other Ops					
ko0	→	ko0	26 d 16 h 26 m	0.00	~ 0	~ 0	~ 0					
ko10	→	ko10	45 d 12 h 01 m	85.22 G	~ 19.86 M	~ 0	~ 0					
ko11	→	ko11	34 d 07 h 43 m	111.04 G	~ 499.83 K	~ 0	~ 0					
ko111	→	ko111	27 d 17 h 33 m	8.14 M	~ 711.20 K	~ 0	~ 0					
ko112	→	ko112	27 d 20 h 11 m	801.72 G	~ 820.96 K	~ 0	~ 0					
ko113	→	ko113	26 d 18 h 59 m	0.00	~ 0	~ 0	~ 0					
ko114	→	ko114	26 d 18 h 59 m	0.00	~ 0	~ 0	~ 0					
ko115	→	ko115	46 d 23 h 35 m	40.66 G	~ 2.55 M	~ 0	~ 0					
ko116	→	ko116	45 d 09 h 42 m	87.25 G	~ 14.00 M	~ 0	~ 0					
ko117	→	ko117	46 d 02 h 19 m	22.88 T	~ 23.99 M	~ 0	~ 0					

The following information is displayed:

- **Local Bucket:** The name of the local (source) bucket.
- **Direction:** The direction of replication, outbound (→).
- **Remote Bucket:** The name of the remote (target) bucket.

- **Lag:** The amount of time, if any, the replication target is behind the source, and is the time difference between the current time and the recovery point.
- **Backlog Size:** The total size of object data in bytes to be replicated. Note that this does not contain custom metadata. The size may be larger than the value displayed depending on custom metadata usage.
- **PUT Ops:** The total number of objects to be replicated.
- **DELETE Ops:** The total number of delete operations to be replicated.
- **Other Ops:** The total number of other operations to be replicated.

The **Remote Credentials** panel displays:

- **Name:** The name of the array and user is the combined remote array name or S3 target with the remote credential name. For example `array2/cred1`.
- **Access Key ID:** The access key ID created for the specific user.
- **Secret Access Key:** The secret key is encrypted and only revealed when first created.

Object Replication Write Conflict Continuation

In SafeMode object replication, replicating an object from the source bucket to the target bucket which contains an object with the same key, but different bodies will result in a conflict and render the replica link unhealthy, halting the replication process.

This feature resolves the conflict by automatically selecting the older object of them as the winner of the conflict. In the meantime, storing both objects in a special bucket, which is called replication conflict virtual bucket. You may manually re-select the desired object version in conflicts whenever you find the result of the automatic resolution is not satisfactory (see the *FlashBlade Object Store S3 REST API* reference).

After automatic resolution, the older one of the incoming object and the existing object will be exposed in the target bucket and both the object versions of the incoming object and the existing object can be found in the virtual bucket of the target bucket.

An alert message will appear on the source and target system to notify you when a conflict occurs and you should review the results of automatic resolution and make any necessary corrections.

 **Note:** Alerts may not appear on both systems due to HA events.

Creating a Bucket Replica Link

 **Note:** Objects created after the creation of the object replica link will be replicated from the source bucket to the target bucket. Any objects that were created prior to the creation of the bucket replica link are not replicated.

 **Note:** When creating an bucket replica link, you have the option to connect a FlashBlade array or S3 target if you have not already done so. In order to connect an array or S3 target, you must have previously created a replication network interface and a connection key (for array connection). See [Creating a Network Interface](#) and [Creating a Connection Key](#) for instructions.

Bucket replica links allow object data to be replicated between FlashBlade arrays and S3 targets.

To create an bucket replica link, perform the following:

- 1 From the **Protection > Object Replication** page, click the **Add (+)** button in the heading of the **Bucket Replica Links** list. The **Create Bucket Replica Link** pop-up window appears.
- 2 From the **Local Bucket Name** list, select the local (source) bucket from which data will be replicated.
- 3 From the **Remote Array or S3 Target** list, select the remote (target) array or S3 target.
- 4 (Optional) Click **Connect FlashBlade Array** or **Connect S3 Target** if you did not previously connect an array or S3 target. See [Connecting Arrays](#) and [Connecting an S3 Target](#) for instructions.
- 5 From the **Remote Bucket** field, enter the name of the remote (target) bucket to which data will be replicated. Versioning must be enabled on the target bucket for replication to proceed.
- 6 From the **Remote Credential** list, select the proper remote credential to allow replication.
- 7 Select **Cascading** if the replica link will be used to connect to multiple targets in a cascade configuration. Note that once set, this setting cannot be changed after the replica link is created.
- 8 Click **Create**.

Editing a Bucket Replica Link

To edit (modify) a bucket replica link, perform the following:

- 1 From the **Protection > Object Replication** page, locate the bucket replica link you wish to modify and click the **More Options** button located at the end of that replica link's row.
- 2 Click **Edit**. The **Edit Bucket Replica Link** pop-up window appears.
- 3 From the **Remote Credential** list, select the proper remote credential to allow replication.

Pausing a Bucket Replica Link

Replication between a local bucket and target bucket can be paused at any time by pausing the replica link.

- 1 To pause a single bucket replica link, perform the following:
 - a From the **Protection > Object Replication** page, locate the bucket replica link you wish to pause and click the **More Options** button located at the end of that replica link's row.
 - b Click **Pause**. The **Pause Bucket Replica Link** pop-up window appears.
 - c Click **Pause** to confirm that you wish to pause the replica link.
- 2 To pause multiple bucket replica links, perform the following:
 - a From the **Protection > Object Replication** page, click the **More Options** button in the header of the **Bucket Replica Links** list and select **Pause**. The **Pause Bucket Replica Links** pop-up window appears.
 - b From the **Active Replica Links** column, select the bucket replica links that you wish to pause. Each selected replica link appears in the **Selected Replica Links** column.
 - c Click **Pause**.

Resuming a Bucket Replica Link

Paused replication between a local bucket and target bucket can be resumed at any time by resuming the replica link.

- 1 To resume a paused bucket replica link, perform the following:
 - a From the **Protection > Object Replication** page, locate the paused bucket replica link you wish to resume and click the **More Options** button located at the end of that replica link's row.
 - b Click **Resume**. The **Resume Bucket Replica Link** pop-up window appears.
 - c Click **Resume** to confirm that you wish to resume the replica link.
- 2 To resume multiple paused bucket replica links, perform the following:
 - a From the **Protection > Object Replication** page, click the **More Options** button in the header of the **Bucket Replica Links** list and select **Resume**. The **Resume Bucket Replica Links** pop-up window appears.
 - b From the **Paused Replica Links** column, select the bucket replica links that you wish to resume. Each selected replica link appears in the **Selected Replica Links** column.
 - c Click **Resume**.

Deleting a Bucket Replica Link

When a bucket replica link is deleted, replication between the local bucket and target bucket is suspended.

- 1 To delete a single bucket replica link, perform the following:
 - a From the **Protection > Object Replication** page, locate the bucket replica link you wish to delete in the **Bucket Replica Links** list and click the **More Options** button located at the end of that replica link's row.
 - b Click **Delete**. The **Delete Bucket Replica Link** pop-up window appears.
 - c Click **Delete** to confirm that you wish to delete the replica link.
- 2 To delete multiple bucket replica links, perform the following:
 - a From the **Protection > Object Replication** page, click the **More Options** button in the header of the **Bucket Replica Links** list and select **Delete**. The **Delete Bucket Replica Links** pop-up window appears.
 - b From the **Replica Links** column, select the bucket replica links that you wish to delete. Each selected replica link appears in the **Selected Replica Links** column.
 - c Click **Delete**.

Creating Remote Credentials

In addition to a replica link, a remote credential is required for replication between a local and target bucket. In order to create a remote credential, the following are required:

- The source and target array, or S3 target, must be connected.
- A previously created user access key and secret access key (see [Creating a User](#)).

To create remote credentials for a replica link, perform the following:

- 1 From the **Protection > Object Replication** page, click the **Add (+)** button in the heading of the **Remote Credentials** list. The **Create Remote Credentials** pop-up window appears.
- 2 From the **Remote Array or Target** list, select the connected remote (target) array or S3 target.
- 3 In the **Name** field, enter the credential name.
- 4 Enter the user access key from the target array or S3 target in the **Access Key ID** field.
- 5 Enter the user secret access key from the target array or S3 target in the **Secret Access Key** field.

- 6 Click **Create**.

Editing Remote Credentials

The access key ID and secret access key for remote credentials for object replication can be edited. In order to edit a remote credential, you must have created a new user access key and secret access key (see [Creating a User](#)).

To edit remote credentials, perform the following:

- 1 From the **Protection > Object Replication** page, click the edit button in the same row of the remote credential you wish to edit in the **Remote Credentials** list. The **Edit Remote Credentials** pop-up window appears.
- 2 Enter the user access key from the target array or S3 target in the **Access Key ID** field.
- 3 Enter the user secret access key from the target array or S3 target in the **Secret Access Key** field.
- 4 Click **Save**.

Deleting Remote Credentials

Deleting remote credentials in an object replication setup suspends replication between the array and its target.

- 1 To delete a single remote credentials, perform the following:
 - a From the **Protection > Object Replication** page, click the delete button in the same row of the remote credential you wish to delete in the **Remote Credentials** list.
 - b Click **Delete**.
- 2 To delete multiple remote credentials, perform the following:
 - a From the **Protection > Object Replication** page, click the **More Options** button in the **Remote Credentials** list header and select **Delete**. The **Delete Remote Credentials** pop-up window appears.
 - b From the **Remote Credentials** column, select the remote credentials that you wish to delete. Each selected remote credential appears in the **Selected Remote Credentials** column.
 - c Click **Delete**.

Chapter 7

Replication and Recovery

FlashBlade supports both file system replication and object replication.

File systems can be replicated from one FlashBlade array to another FlashBlade array. File system replication requires two connected arrays and a file system replica link between the source file system and a target file system.

Objects can be replicated from a bucket on a FlashBlade (source) to a remote target bucket residing on another FlashBlade or an S3 target using object replica links. An object replica link is the relationship between a source bucket and target bucket, and describes the properties of that connection.

File System Replication Overview

File System Replication

File systems can be replicated from one FlashBlade array to another FlashBlade array. File system replication requires two connected arrays and a file system replica link between the source file system and a target file system. A file system replica link is the connection between a local file system and target file system.

Multiple arrays can be connected. However, file system replication is allowed between only one array pair. Multiple file system replica links on a single array are allowed as long as they are using different file systems.

File System Replication Fan Out Configuration

In a file system replication fan out configuration, a single source file system has file system replica links to different target file systems. For example, A→B and A→C.

During fan out configuration set up:

- Connect the source FlashBlade array to the first remote FlashBlade.
- Connect the source FlashBlade array to the second remote FlashBlade.

- Create file replica links on the source FlashBlade array to enable file replication from the source FlashBlade array to the two remote targets.

File System Replication Prerequisites

- Replication traffic requires a replication vip using port 8117; before arrays can replicate, they must first be connected over the management vip/port 443.
- Ports 8117 and 80 are used for non-encrypted file replication.

File System Recovery

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

In file system replication setups you can stop writes to the local file system and redirect writes to the target file system (failover). Failover can be used if there is an issue with the local file system. During failover, replication stops until you reprotect.

Failover requires promotion of the target file system and redirection of all clients to the target array to return to operation. To reprotect, demote the (former) source file system. Once the file system is demoted, replication policies will start replicating in the new direction.

Note that when demoting the local file system, any writes performed since the last replication will be discarded. If a local snapshot was taken since last replication, demoting the file system requires that snapshot be destroyed and eradicated. You must remove any data since the last snapshot of the file system was taken. The file system must either have no snapshots, or the most recent snapshot must be the last replicated snapshot. Additionally, it is suggested that any policies attached to the local file system and replica link be removed in order to expedite failover.

Once the local file system comes back up, re-enable the scheduled replication policy that was disabled before demoting the file system. Replication resumes as scheduled between the promoted remote file system and the demoted local file system. Data is written to the promoted remote file system, which is then replicated to the demoted local file system.

For fan out replication setups, once scheduled replication policies for the local file system are disabled across all the clusters in the fan out file system configuration, replica links related to that file system on the target clusters must be deleted. Once the file system on the new source cluster is promoted, replica links must be created between the promoted file system and the other clusters in your fan out file system replication configuration, and the original source file system must be demoted. Note that any policies

that were disabled on the former source file system cannot be re-enabled on the promoted file system. You must create new policies on the promoted file system.

To resume the original replica link direction (failback), manually stop writes on the remote file system and send the remaining data to the local file system. Once data is sent to the local file system, promote the local file system and demote the remote file system.

Getting Started

File system replication setup and recovery can be performed using either the Purity//FB GUI or CLI.

The following sections provide an overview of the tasks required for file replication setup and for performing recovery. Complete instructions using the GUI and CLI are provided for each in subsequent topics.

File System Replication Setup

To set up file system replication, the following tasks must be completed:

- 1 Create replication network interfaces on the source and target arrays.
- 2 Create a connection key on the target array.
- 3 Connect the source and target arrays.
- 4 Create a policy to govern snapshot creation and retention.
- 5 Create a file system replica link on the local array.

Complete instructions for completing each of these tasks are provided in:

- [*Setting Up File System Replication Using the GUI*](#)
- [*Setting Up File System Replication Using the CLI*](#)

File System Replication Failover

To perform file system replication failover, the following tasks must be completed:

- 1 Promote the remote file system.
- 2 Manually direct clients to the remote array.

- 3 Remove any policies from the local file system.
- 4 Demote the local file system.

Complete instructions for completing each of these are provided in:

- [*Performing File System Replication Failover Using the GUI*](#)
- [*Performing File System Replication Failover Using the CLI*](#)

File System Replication Failback

To perform file system replication failback, the following tasks must be completed:

- 1 Disable writes on the remote file system.
- 2 Promote the local file system.
- 3 Re-add any policies to the local file system.
- 4 Manually redirect clients back to the local array.
- 5 Demote the target file system.

Complete instructions for completing each of these tasks are provided in:

- [*Performing File System Replication Failback Using the GUI*](#)
- [*Performing File System Replication Failback Using the CLI*](#)

File System Fan Out Replication Failover

To perform file system replication failover for fan out configurations, the following tasks must be completed:

- 1 Disable policies related to the source file system on all array (of the policy owner).
- 2 Delete file system replica links to the source file system on the target arrays.
- 3 Promote the remote file system.
- 4 Create file system replica links between the promoted file system and other arrays.
- 5 Demote the original source file system.

- 6 Re-create any replication policies on the promoted file system.

Complete instructions for completing each of these are provided in:

- [*Performing File System Fan Out Replication Failover Using the GUI*](#)
- [*Performing File System Fan Out Replication Failover Using the CLI*](#)

File System Fan Out Replication Failback

To perform file system replication failback for fan out configurations, the following tasks must be completed:

- 1 Disable writes on the remote file system.
- 2 Promote the local file system.
- 3 Delete file system replica links between the remote file system and other remote file systems.
- 4 Re-create file system replica links between the local file systems and the other remote file systems.
- 5 Re-add any policies to the local file system.
- 6 Manually redirect clients back to the local array.
- 7 Demote the target file system.

Complete instructions for completing each of these are provided in:

- [*Performing File System Fan Out Replication Failback Using the GUI*](#)
- [*Performing File System Fan Out Replication Failback Using the CLI*](#)

Setting Up File System Replication Using the GUI

The following instructions describe how to set up file system replication from one FlashBlade array to another FlashBlade array using the Purity//FB GUI. These instructions should be followed in the order that they are presented.

Creating Replication Network Interfaces

Perform the following on both your source and target FlashBlade arrays:

- 1 From the **Settings** sidebar, click **Subnets** under **Network**. The **Subnets** panel appears.
- 2 In the **Subnets** panel, locate the subnet to which you wish to add a network interface and click the **Add interface (+)** button. The **Create Network Interface** pop-up window appears.
- 3 In the **Name** field, enter the interface name.
- 4 In the address field, enter the network address.
- 5 From the **Services** list, select `replication`.
- 6 The subnet defaults to the subnet in which you are creating the interface. Click **Create**.

Creating a Connection Key

 **Important!** The connection key is active for two hours. If the source and target arrays have not been connected within that two hour period, a new connection key must be created.

Once you have created replication network interfaces, you must create a connection key on the target array.

- 1 On the target array, select **Storage > Array** page.
- 2 In the **FlashBlade Array Connections** panel, click **More Options > Create Connection Key**. The **Connection Key** pop-up window appears displaying the new connection key.
- 3 Copy the new connection key.

Connecting the Source and Target Arrays

 **Note:** Multiple file system replica links on a single array are allowed as long as they are using different file systems. Multiple file system replica links are not allowed from a single file system or when using multiple array connections.

 **Note:** When connecting arrays, it is recommended to create the array connection from the array running the most current Purity version to the array running an older Purity version.

With the creation of replication network interfaces on the source and target arrays, and the creation of a connection key on the target array, you can now connect the source and target arrays.

 **Note:** Encryption is only valid for object replication.

 **Important!** When connecting the source array to the target array, a secure connection is required even if encryption is not enabled. Certificates with small key sizes, weak signature algorithms, or that only support weak TLS algorithms are not supported.

 **Important!** The management IP address required during configuration cannot be a loopback IP address.

 **Important!** It is highly recommended that you add the `pure:policy/full-access` access policy to the user (whose credentials are used for replication) on the target cluster, otherwise replication may fail due to denied access. See [Adding Access Policies](#) for instructions on how to add access policies to users.

- 1 On the source array, select **Storage > Array** page, click the add (+) button in the **FlashBlade Array Connections** panel. The **Connect FlashBlade Array** pop-up window appears.
- 2 Enter the target array's hostname or management address in the **Management Address** field. The address can be located from the **Subnets** table by navigating to the **Settings > Network** page.
- 3 Paste the copied connection key for the target array in the **Connection Key** field.
- 4 The replication address is automatically populated on both the source and target cluster unless manually specified. If manually specified, then the target cluster will not have a replication address populated and will need to be manually entered. It is recommended to allow the arrays to automatically discover the replication address, unless using NAT. If using NAT, enter the target array's replication network address in the **Replication Address** field.
- 5 (Optional) Enable encryption by setting the **Encrypted** toggle to on. If set, the `_default_replication_certs` CA certificate group is applied. Before toggling, make sure to import the CA certificate to the source array if the target array's array-certificate is signed by a CA and add it to the certificate group. Note that encryption is set on the source array only. If enabled, the encryption setting displays as enabled on the target array.

 **Warning:** Use of CA certificates to establish trust and secure communication for object replication is the recommended configuration. If the target array's array-certificate is self-signed, directly importing the target array certificate to the source and adding it to the certificate group will not provide as strict validation around aspects of the peer, such as certificate expiry.

- 6 Click **Connect**.
- 7 If replication address was specified on the source cluster while creating the array connection in step 4, on the target array, enter the source array's replication address.
 - a On the target array, in the **FlashBlade Array Connections** panel on the **Storage > Array** page, click the edit button at the end of the row displaying the connected source array. The **Edit Connected Array** pop-up window appears.

- b** Enter the replication address for the source array in the **Replication Address** field and click **Save**.

Creating a Policy

Policies set the rules that govern the frequency of replication and the duration that snapshots are retained.

- 1** From the source array, select **Protection > Policies**, and click the add (+) button in the heading of the **Snapshot Policies** list. The **Step 1: Create Policy** pop-up window appears.
- 2** Enter a policy name.
- 3** Click **Create**. The **Step 2: Add Rule to Policy <name>** pop-up window appears.
- 4** In the **Create or replicate 1 snapshot every** field, enter the frequency at which snapshots are created or replicated. The value must be entered in the format n{mlhldlw}. For example, 15m, 3h, 2d, 1w, etc.
- 5** In the **At** field, enter the time of day the snapshot is created. The value must be entered in the format n[am|pm], where n is a value of 1-12. If entered without am or pm, n must be a value of 0-23. This field is only editable if the value specified in the **Create or replicate 1 snapshot every** field is in days. For example, 24h, 48h, 72h, 1d, 2d, 3d, 1w, etc.
- 6** In the **And keep for** field, enter the retention period for the snapshot. The value must be entered in the format n{mlhldlw}. For example, 30m, 1h, 2d, 1w, etc. The retention period cannot be less than the snapshot interval.
- 7** Click **Add**.

Creating a File System Replica Link

 **Note:** When creating a file system replica link, you have the option to connect a FlashBlade array if you have not already done so. In order to connect an array you must have previously created a replication network interface and a connection key. See [Creating a Network Interface](#) and [Creating a Connection Key](#) for instructions.

File system replica links are used to configure file system replication between two connected FlashBlade arrays.

Fan out file system replication allows file system replication to more than one target. In order to replicate to multiple targets, file system replica links to each target must be created. The following are supported:

- A maximum of 1,000 file systems can have replica links.

- A file system can have a maximum of 10,000 outgoing replica links and a maximum of 1,000 incoming replica links.

To create a file system replica link, perform the following:

- 1 From the **Protection > File Replication** page, click the add (+) button in the heading of the **File System Replica Links** list. The **Step 1: Create File System Replica Link** pop-up window appears.
- 2 In the **Step 1: Create File System Replica Link** pop-up window, perform the following:
 - a From the **Local File System** list, select the local (source) file system to be replicated.
 - b From the **Remote Array** list, select the remote (target) array.
 - c (Optional) Enter a name for the remote file system to which data from the local file system will be replicated in the **Remote File System** field.
 - d (Optional) Click **Connect FlashBlade Array** if you did not previously connect an array. See [Connecting Arrays](#) for instructions.
 - e Click **Create**. The **Step 2: Create Fanout Replica Links** pop-up window appears
- 3 In the **Step 2: Create Fanout Replica Links** pop-up window, perform the following:
 - a If you are not setting up fan out file replication, click **Skip**. The **Add Snapshot Policies** pop-up window appears (see Step 4).
 - b Select the connected FlashBlade arrays from the **Available FlashBlade Arrays** column to which you wish to create replication links. Selected arrays appear in the **Selected FlashBlade Arrays** column.
 - c If there is an existing demoted remote file system with the same name as the source file system on the remote array, selecting the **Relink existing file systems** option will reconnect and promote the file system.
 - d Click **Create Links**. The **Add Snapshot Policies** pop-up window appears.
- 4 In the **Add Snapshot Policies** pop-up window, perform the following:
 - a Select one or more snapshot policy from the **Available Policies** column. Selected policies appear in the **Selected Policies** column. If a policy does not exist, click **Create a New Snapshot Policy**. Follow the policy creation instructions as described in [Creating Snapshot Policies](#). If you do not add a policy, the replica link can still be created. However, replication will not occur until a policy is added to the link.
 - b Click **Add**.

Setting Up File System Replication Using the CLI

The following instructions describe how to set up file system replication from one FlashBlade array to another FlashBlade array using the Purity//FB CLI. These instructions should be followed in the order that they are presented.

Creating Replication Network Interfaces

On both the source and target arrays, issue the

purenetwork vip create --address ADDRESS --servicelist replication NAME

command, where ADDRESS is the IP address associated with the replication vip and NAME is the data vip.

```
# purenetwork vip create --address 123.123.123.123 --servicelist replication  
datavip1
```

Creating a Connection Key

 **Important!** The connection key is active for two hours. If the source and target arrays have not been connected within that two hour period, a new connection key must be created.

Once you have created replication network interfaces on the source and target arrays, you must create a connection key on the target array.

- 1 Issue the **purearray create --connection-key** command.

```
# purearray create --connection-key  
Connection Key          Created          Expires  
T-71a63c01-80aa-4173-842c-8ec9a1285d5b 2020-01-09 14:49:08 PST 2020-01-09  
16:49:08 PST
```

- 2 Copy the connection key.

Connecting the Source and Target Arrays

 **Note:** Multiple file system replica links on a single array are allowed as long as they are using different file systems. Multiple filesystem replica links are not allowed from a single file system or when using multiple array connections.

 **Note:** When connecting arrays, it is recommended to create the array connection from the array running the most current Purity version to the array running an older Purity version.

With the creation of replication network interfaces on the source and target arrays, and the creation of a connection key on the target array, you can now connect the source and target arrays.

 **Important!** It is highly recommended that you add the `pure:policy/full-access` access policy to the user (whose credentials are used for replication) on the target cluster, otherwise replication may fail due to denied access. See the `pureobj` command in the *FlashBlade CLI Reference* for information on how to add access policies to users.

 **Important!** When connecting the source array to the target array, a secure connection is required even if encryption is not enabled. Certificates with small key sizes, weak signature algorithms, or that only support weak TLS algorithms are not supported.

 **Important!** The management IP address required during configuration cannot be a loopback IP address.

- 1 On the target array, issue the **purenetwork list** command and make note of the array's management IP address.

```
# purenetwork list
Name      Enabled Subnet Address  ... Services
vir1      True   admin  10.11.123.12 ... management
```

- 2 On the source array, issue the **purearray connect --management-address MANAGEMENT_ADDRESS** command, where `MANAGEMENT_ADDRESS` is the management IP address of the target array.

- If the optional `--replication-address` argument is used, the specified replication address is set on the local array connection. The remote array connection will have no replication address set and must be manually entered.
- If not used, the replication address available on the target array is used, and the replication address of the local array is automatically populated on the target array connection. This is recommended, unless using NAT. If using NAT, specify the target array's replication address using the `--replication address` argument.

```
# purearray connect --management-address 10.11.123.12
Enter the connection key of the target array:
```

- 3 When prompted, enter (paste) the target array's connection key.
- 4 If using NAT, or if the replication address was specified on the source cluster while creating the array connection, on the target array issue the

purearray setattr --connect --replication address REPLICATION_ADDRESS
command, where **REPLICATION_ADDRESS** is the replication address of the source array.

- 5 The connection between two arrays in an object replication set up can be encrypted. To enable encryption between two arrays, on the source array run the **purearray enable --connect --encrypted NAME** command, where **NAME** is the name of the target array.

Creating a Policy

Policies set the rules that govern the frequency of replication and the duration that snapshots are retained.

- 1 From source array, issue the **purepolicy create POLICY** command, where **POLICY** is the name of the snapshot policy.

```
# purepolicy create snap_poll
```

- 2 Add rules to the new policy by issuing the **purepolicy rule add** command. You must specify when the snapshots are created, the interval at which they are created, and the duration they are retained using the **--at**, **--every**, and **--keep-for** arguments. In the following example, snapshots are created at 12 AM each day and retained for one day for the policy **snap_poll**.

```
# purepolicy rule add --at 12am --every 1d --keep-for 1d snap_poll
```

Creating a File System Replica Link

Once the source and target arrays are connected, you must create a file system replica link on the source array. The file system replica link contains the details and specifies the rules of the replication relationship between the local and remote file systems.

Fan out file replication allows file system replication to more than one target. In order to replicate to multiple targets, file replica links to each target must be created.

From the source array, issue the **purefs replica-link create LOCAL_FILE-SYSTEM** command, where **LOCAL_FILE-SYSTEM** is the name of the source file system. Additionally, specify the remote array as well as the remote file system to which the source will be replicating, using the **--remote** and **--remote-fs** arguments.

```
# purefs replica-link create --remote Array2 --remote-fs fs1-R --policy pol_poll  
fs1
```

In this example, the replica link is created between the file system `fs1` on the local array and the remote file system `fs1-R` (also created with this command) on the remote array `Array2`. Replication occurs per the schedule defined in the policy `pol_pol1`. Note that if the file system `fs1-R` already existed on the remote array, this command would fail.

Performing File System Replication Failover Using the GUI

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

To initiate failover in a file system replication setup:

- The target file system must first be promoted.
- All clients must then be directed to the target array.
- The local file system is then demoted. Note that when a file system is demoted, any writes performed since the last replication will be discarded.

To initiate failover, perform the following:

- 1 On the target array, from the **Storage > File Systems** page, locate the remote file system that you wish to promote from the **File Systems** panel.
- 2 Click **More Options > Promote** at the end of the row in which the remote file system appears.
- 3 Click **Promote** when prompted for confirmation.
- 4 Manually redirect clients to the target array.

To reprotect, you must remove snapshots since the last replicated snapshot. This can be difficult if the policy engine is creating new snapshots. You can stop the policy engine from creating new snapshots by removing the policies from the file system and replica link.

- 5 On the local array, remove policies attached to the to-be-demoted local file system and replica link.

To remove the policy from the local file system:

- a From the **Storage > File Systems** page, locate and click local file system that you wish to demote from the **File Systems** panel. The file system's details page appears.

- b** In the **Policies** panel, click the remove button (**X**) in the same row as the policy you wish to remove.
- c** Click **Remove** when prompted for confirmation.

To remove the policy from the replica link:

- a** From the **Protection > File System Replica Links** page, click **More Options > Remove Policies** at the end of the row of the replica link you wish to edit. The **Remove Policies** pop-up window appears.
 - b** Select the policy or policies you wish to remove from the replica link from the **Available Policies** list.
 - c** Click **Remove**.
- 6** Navigate back to the **Storage > File Systems** page and locate the local file system that you wish to demote from the **File Systems** panel.
- 7** Click **More Options > Demote** at the end of row in which the local file system appears.
- 8** Click **Demote** when prompted for confirmation.

Writes to the local file system are stopped and are redirected to the promoted remote file system.

- 9** If you removed policies in step 5, add the removed policy or policies back to the file system and replica link.

To add the policy to the file system:

- a** From the **Storage > File Systems** page, locate and click the file system from the **File Systems** panel. The file system's details page appears.
- b** In the **Policies** panel, click the add button (+). The **Add Policies** pop-up window appears.
- c** Select the policy or policies you wish to add to the file system from the **Available Policies** list.
- d** Click **Add**.

To add the policy to the replica link:

- a** From the **Protection > File System Replica Links** page, click **More Options > Add Policies** at the end of the row of the replica link you wish to edit. The **Add Policies** pop-up window appears.
- b** Select the policy or policies you wish to add back to the replica link from the **Available Policies** list.
- c** Click **Add**.

Once the local file system comes back up, replication resumes as scheduled between the newly promoted target file system and demoted local file system because the policy was re-added to the

replica link. Because clients have been redirected to the target, data is written to the promoted remote file system, which is then replicated to the demoted local file system.

Performing File System Replication Failback Using the GUI

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

 **Note:** In this procedure, the remote file system was promoted and the local file system was demoted during the failover procedure.

As noted in the failover procedure, if the replication schedule policy was re-added to the replica link in the replication setup, once the local file system comes back up, replication resumes as scheduled between the promoted target file system and demoted local file system. Performing failback restores the original local/target relationship between arrays and file systems in a replication setup. During failback, once the remote file system is demoted and the local file system is promoted and clients redirected to the local file system, the original replication setup resumes; data is written to the local file system and then replicated to the remote file system.

To initiate failback in a file replication setup:

- Stop writes on the promoted remote file system.
- Optionally manually create a snapshot of the file system and send it to the demoted local file system to ensure that data written since the last snapshot resides on the to-be-promoted file system.
- Promote the previously demoted local file system.
- Add back any removed policy to the local file system and replica link.
- Demote the remote file system.

To initiate failback, perform the following:

- 1 On the target array, from the **Storage > File Systems** page, locate the remote file system that is to be demoted from the **File Systems** panel.
- 2 Click **More Options > Disable Writes** at the end of row in which the remote file system appears.

- 3 Click **Disable** when prompted for confirmation.
- 4 Ensure all writes are sent to the source file system by creating a snapshot on the target and send it to the source.
- 5 On the local array, from the **Storage > File Systems** page, locate the local file system that you wish to promote from the **File Systems** panel.
- 6 Click **More Options > Promote** at the end of row in which the local file system appears.
- 7 Manually redirect clients back to the to the local array.
- 8 Remove policies attached to the to-be-demoted remote file system and replica link.

To remove the policy from the file system:

- a From the **Storage > File Systems** page, locate and click local file system that you wish to demote from the **File Systems** panel. The file system's details page appears.
- b In the **Policies** panel, click the remove button (X) in the same row as the policy you wish to remove.
- c Click **Remove** when prompted for confirmation.

To remove the policy from the replica link:

- a From the **Protection > File System Replica Links** page, click **More Options > Remove Policies** at the end of the row of the replica link you wish to edit. The **Remove Policies** pop-up window appears.
 - b Select the policy or policies you wish to remove from the replica link from the **Available Policies** list.
 - c Click **Remove**.
- 9 On the target array, from the **Storage > File Systems** page, locate the remote file system that you wish to demote from the **File Systems** panel.
 - 10 Click **More Options > Demote** at the end of row in which the remote file system appears.
 - 11 Click **Demote** when prompted for confirmation.

Writes to the remote file system are stopped and are redirected to the promoted local file system.

- 12 If you removed policies in step 8, add the removed policy or policies back to the file system and replica link.

To add the policy to the file system:

- a From the **Storage > File Systems** page, locate and click the file system from the **File Systems** panel. The file system's details page appears.

- b** In the **Policies** panel, click the add button (+). The **Add Policies** pop-up window appears.
- c** Select the policy or policies you wish to add to the file system from the **Available Policies** list.
- d** Click **Add**.

To add the policy to the replica link:

- a** From the **Protection > File System Replica Links** page, click **More Options > Add Policies** at the end of the row of the replica link you wish to edit. The **Add Policies** pop-up window appears.
- b** Select the policy or policies you wish to add back to the replica link from the **Available Policies** list.
- c** Click **Add**.

The original replication setup has been restored.

Performing File System Replication Failover Using the CLI

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

In a replication setup, failover requires promoting the file system on target (remote) array and then demoting the file system on the local (source) array.

Note that when demoting the local file system, any writes performed since the last replication will be discarded. If a local snapshot was taken since last replication, demoting the file system requires that snapshot be destroyed and eradicated. You must use the `--discard-non-snapshotted-data` option to remove any data since the last snapshot of the file system was taken. The file system must either have no snapshots, or the most recent snapshot must be a replication snapshot.

In the procedure below:

- The local array is `array1`
- The local file system is `fs1`
- The target array is `array2`
- The target file system is `fs1-R`

- The policy attached to the local file system `fs1` is `fs_snap_poll`.
 - The policy attached to the replica link is `rl_poll`.
- 1 On the target array, `array2`, promote the target file system. In the following example, the target file system `fs1-R` is promoted:

```
# purefs promote fs1-R
```

- 2 Manually redirect clients to the the target array, `array2`.
- 3 On the local array, `array1`, remove the policy `fs_snap_poll` that attached to the local file system `fs1`

```
# purepolicy snapshot remove fs_snap_poll
```

- 4 Remove the policy from the replica link using the **`purepolicy snapshot remove --remote`** command, where `--remote` specifies the remote (target) array. In the following example the policy `rl_poll` is removed from the replica link:

```
# purepolicy snapshot remove --remote array2 rl_poll
```

- 5 On the local array, `array1`, demote the local file system :

```
# purefs demote --discard-non-snapshotted-data fs1
```

- 6 Add the policy back to the replica link using the **`purepolicy snapshot add --remote`** command, where `--remote` specifies the demoted array In the following example, the policy `rl_poll` is added back to the replica link:

```
# purepolicy snapshot add --remote array1 rl_poll
```

- 7 On the local array, verify that replication is disabled and that the target file system is promoted:

```
# purefs replica-link list
Name  Direction  Remote  Remote File System  Policy      Status      Recov...
fs1   <--        array2  fs1-R               rl_poll     unhealthy   2019-...

purefs replica-link list --status-details
Name  Direction  Remote  Remote File System  Policy      Status      Details
fs1   <--        array2  fs1-R               rl_poll     unhealthy   The
target file system is promoted.
```

Performing File System Replication Failback Using the CLI

⚠ CAUTION: If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

As noted in the failover procedure, if the replication schedule policy was re-added to the replica link in the replication setup, once the local file system comes back up, replication resumes as scheduled between the promoted target file system and demoted local file system. Performing failback restores the original local/target relationship between arrays and file systems in a replication setup. During failback, once the remote file system is demoted and the local file system is promoted and clients redirected to the local file system, the original replication setup resumes; data is written to the local file system and then replicated to the remote file system.

In the procedure below:

- The local array is `array2`
 - The local file system is `fs1-R`
 - The target array is `array1`
 - The target file system is `fs1`
 - The policy re-attached to the local file system `fs1` is `fs_snap_poll`.
 - The policy attached to the replica link is `rl_poll`.
- 1 On the local array, `array2`, stop all writes on the local file system `fs1-R` and then transfer all remaining data to the target file system `fs1`:

```
# purefs disable --writable fs1-R
# purefs snap --send --target array1 fs1
```

- 2 Create a snapshot of the file system and send it to the target file system `fs1` to ensure that data since the last snapshot resides on the to-be-promoted file system (`fs1`). Skip this step if you do not need data written since the last snapshot.
- 3 On the target array, `array1`, promote the target file system `fs1`:

```
# purefs promote fs1
```

- 4 Add back any previously removed file system policy. In the following example, the policy `fs_snap_poll` is added back to file system `fs1`.

```
# purepolicy snapshot add fs_snap_poll
```

- 5 Remove the policy from the replica link using the **`purepolicy snapshot remove --remote`** command, where `--remote` specifies the remote (target) array. In the following example the policy `rl_poll` is removed from the replica link:

```
# purepolicy snapshot remove --remote array1 rl_poll
```

- 6 Manually direct all clients back to the promoted file system `fs1`.

- 7 On the local array, `array2`, demote the local file system `fs1-R`:

```
# purefs demote --discard-non-snapshotted-data fs1-R
```

- 8 Add the policy back to the replica link using the **`purepolicy snapshot add --remote`** command, where `--remote` specifies the demoted array. In the following example, the policy `rl_poll` is added back to the replica link:

```
# purepolicy snapshot add --remote array2 rl_poll
```

- 9 On the (now) local array `array1`, verify that the replica link has returned to its original direction (from file system `fs1` to the target file system `fs1-R`):

```
# purefs replica-link list
Name Direction Remote Remote File System Policy Status Lag
Recovery P...
fs1 --> array2 fs1-R rl_poll healthy 10m
2019-11-11 12...
```

Performing File System Fan Out Replication Failover Using the GUI

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

There are two types of failover: planned and unplanned. A planned failover can occur in situations such as planned maintenance. It is predetermined which remote file system will be promoted while the local file system is taken offline. An unplanned failover, as the name suggests, occurs without preparation.

such as loss of power to a data center. In such events, when the local file system goes offline, an analysis is required to determine the remote file system most suitable to be promoted. The following procedure describes a planned failover.

To initiate failover in a file system fan out replication setup, perform the following:

- 1** Disable the policies related to the local file system on each array.
 - a** From the **Policies > Snapshot** tab, click the policy you wish to disable. The policy's detail screen appears.
 - b** In the **Details** panel, click **Disable**. The **Disable Policy** pop-up window appears.
 - c** Click **Disable**.
- 2** Delete the file system replica links to the local file system on each target array.
 - a** From the **Protection > File Replication** tab, locate the file system replica link you wish to delete from the **File System Replica Links** panel.
 - b** From the same row as the file system replica link to be deleted, click **More Options > Delete**. The **Delete File Replica Link** dialog box appears.
 - c** Click **Delete**.
- 3** Promote the file system on the remote target array.
 - a** On the target array, from the **Storage > File Systems** page, locate the remote file system that you wish to promote from the **File Systems** panel.
 - b** Click **More Options > Promote** at the end of the row in which the remote file system appears.
 - c** Click **Promote** when prompted for confirmation.
- 4** Create file system replica links between the promoted file system and the other arrays.
 - a** From the **Protection > File Replication** page, click the add (+) button in the heading of the **File System Replica Links** list. The **Step 1: Create File System Replica Link** pop-up window appears.
 - b** From the **Local File System** list, select the promoted file system.
 - c** From the **Remote Array** list, select the remote (target) array.
 - d** (Optional) Enter a name for the remote file system to which data from the promoted file system will be replicated in the **Remote File System** field.
 - e** Click **Create**. The **Step 2: Create Fanout Replica Links** pop-up window appears.
 - f** Select the connected FlashBlade arrays from the **Available FlashBlade Arrays** column to which you wish to create replication links. Selected arrays appear in the **Selected FlashBlade Arrays** column.

- g If there is an existing demoted remote file system with the same name as the source file system on the remote array, selecting the **Relink existing file systems** option will reconnect and promote the file system.
 - h Click **Create Links**. The **Add Snapshot Policies** pop-up window appears.
 - i Select one or more snapshot policy from the **Available Policies** column. Selected policies appear in the **Selected Policies** column. If a policy does not exist, click **Create a New Snapshot Policy**. Follow the policy creation instructions as described in [Creating Snapshot Policies](#). If you do not add a policy, the replica link can still be created. However, replication will not occur until a policy is added to the link.
 - j Click **Add**.
- 5 On the original local array, demote the local file system.
- a Navigate to the **Storage > File Systems** page and locate the local file system that you wish to demote from the **File Systems** panel.
 - b Click **More Options > Demote** at the end of row in which the local file system appears.
 - c Click **Demote** when prompted for confirmation. Note that when a file system is demoted, any writes performed since the last replication will be discarded.

Performing File System Fan Out Replication Failback Using the GUI

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

 **Note:** In this procedure, the remote file system was promoted and the local file system was demoted during the failover procedure.

Performing failback restores the original local/target relationship between arrays and file systems in a replication setup. During failback, once the remote file system is demoted and the local file system is promoted and clients redirected to the local file system, the original replication setup resumes; data is written to the local file system and then replicated to the remote file system.

In the following procedure, *remote file system* refers to the file system that was promoted during failover, and *local file system* refers to the file system that was demoted during failover.

- 1 Disable writes on the remote file system.

- a On the target array, from the **Storage > File Systems** page, locate the remote file system that was previously promoted during failover from the **File Systems** panel.
 - b Click **More Options > Disable Writes** at the end of row in which the remote file system appears.
 - c Click **Disable**.
- 2 Promote the local file system that was demoted during failover.
 - a On the local array, from the **Storage > File Systems** page, locate the file system that you wish to promote from the **File Systems** panel.
 - b Click **More Options > Promote** at the end of the row in which the file system appears.
 - c Click **Promote** when prompted for confirmation.
- 3 Delete the file system replica links to the promoted remote file system on each target array.
 - a From the **Protection > File Replication** tab, locate the file system replica link you wish to delete from the **File System Replica Links** panel.
 - b From the same row as the file system replica link to be deleted, click **More Options > Delete**. The **Delete File Replica Link** dialog box appears.
 - c Click **Delete**.
- 4 Re-create file system replica links between the newly promoted local file system and the other remote file systems.
 - a From the **Protection > File Replication** page, click the add (+) button in the heading of the **File System Replica Links** list. The **Step 1: Create File System Replica Link** pop-up window appears.
 - b From the **Local File System** list, select the promoted file system.
 - c From the **Remote Array** list, select the remote (target) array.
 - d (Optional) Enter a name for the remote file system to which data from the promoted file system will be replicated in the **Remote File System** field.
 - e Click **Create**. The **Step 2: Create Fanout Replica Links** pop-up window appears
 - f Select the connected FlashBlade arrays from the **Available FlashBlade Arrays** column to which you wish to create replication links. Selected arrays appear in the **Selected FlashBlade Arrays** column.
 - g If there is an existing demoted remote file system with the same name as the source file system on the remote array, selecting the **Relink existing file systems** option will reconnect and promote the file system.
 - h Click **Create Links**. The **Add Snapshot Policies** pop-up window appears.

- i Select one or more snapshot policy from the **Available Policies** column. Selected policies appear in the **Selected Policies** column. If a policy does not exist, click **Create a New Snapshot Policy**. Follow the policy creation instructions as described in [Creating Snapshot Policies](#). If you do not add a policy, the replica link can still be created. However, replication will not occur until a policy is added to the link.
 - j Click **Add**.
- 5 Manually redirect clients back to the to the local array.
- 6 Demote the remote file system.
- a On the target array, from the **Storage > File Systems** page, locate the file system that you wish to demote from the **File Systems** panel.
 - b Click **More Options > Demote** at the end of row in which the remote file system appears.
 - c Click **Demote** when prompted for confirmation.

Performing File System Fan Out Replication Failover Using the CLI

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

There are two types of failover: planned and unplanned. A planned failover can occur in situations such as planned maintenance. It is predetermined which remote file system will be promoted while the local file system is taken offline. An unplanned failover, as the name suggests, occurs without preparation such as loss of power to a data center. In such events, when the local file system goes offline, an analysis is required to determine the remote file system most suitable to be promoted. The following procedure describes a planned failover.

Note that when demoting the local file system, any writes performed since the last replication will be discarded. If a local snapshot was taken since last replication, demoting the file system requires that snapshot be destroyed and eradicated. You must use the `--discard-non-snapshotted-data` option to remove any data since the last snapshot of the file system was taken. The file system must either have no snapshots, or the most recent snapshot must be a replication snapshot.

In the procedure below:

- The local array is `array1`

- The local file system is `fs1`
- Fan out file system replication is set up from the local array `array1` to target arrays `array2` and `array3`.
- The target array chosen for failover is `array2` with the remote file system `fs1-R`.
- The policy attached to the local file system `fs1` is `fs_snap_poll`.
- The policy attached to the replica link is `rl_poll`.

- 1 Disable policies related to the source file system on all arrays. On `array2` and `array3`, issue the `purepolicy snapshot disable` command specifying the policy name(s).

```
array2# purepolicy snapshot disable fs_snap_poll
array3# purepolicy snapshot disable fs_snap_poll
```

- 2 Delete file system replica links to the source file system on the target array. On `array3`, issue the `purefs replica-link delete` command specifying the name of the source file system, `fs1`.

```
array3# purefs replica-link delete fs1
```

- 3 On the target array, `array2`, promote the remote file system `fs1-R`.

```
array2# purefs promote fs1-R
```

- 4 Create a file system replica link between the promoted file system (`fs1-R`) and the other target array (`array3`) by issuing the `purefs replica-link create --remote <target_array> --remote-fs <file_system>` command. You must specify the source file system (`fs1-R`).

```
array3# purefs replica-link create --remote-fs fs1 --remote array3 fs1-R
```

- 5 On the local array, `array1`, demote the local file system:

```
array1# purefs demote --discard-non-snapshotted-data fs1
```

- 6 On the promoted file system, recreate the replication policy by issuing the `purepolicy snapshot create` command. You must specify the policy name.

```
array2# purepolicy snapshot create rl_poll
```

- 7 Add the policy back to the replica link using the `purepolicy add --fs <name> --remote` command, where `<name>` is the file system to which the policy is attached and `--remote` specifies

the demoted array. In the following example, the policy `rl_poll` is added back to the replica link and is attached to file system `fs1-R`:

```
array2# purepolicy add --fs fs1-R --remote array1 rl_poll
```

Performing File System Fan Out Replication Failback Using the CLI

 **CAUTION:** If Pure File SafeMode is enabled, contact Pure Technical Services to promote and demote file systems.

Performing failback restores the original local/target relationship between arrays and file systems in a replication setup. During failback, once the remote file system is demoted and the local file system is promoted and clients redirected to the local file system, the original replication setup resumes; data is written to the local file system and then replicated to the remote file system.

In the following procedure, *remote file system* refers to the file system that was promoted during failover, and *local file system* refers to the file system that was demoted during failover.

In the procedure below:

- The remote array is `array2`
 - The remote file system is `fs1-R`
 - The local array is `array1`
 - The local file system is `fs1`
 - The policy re-attached to the local file system `fs1` is `fs_snap_poll`.
 - The policy attached to the replica link is `rl_poll`.
- 1 On the remote array, `array2`, stop all writes on the remote file system `fs1-R` and then transfer all remaining data to the local file system `fs1`:

```
array2# purefs disable --writable fs1-R  
# purefs snap --send --target array1 fs1
```

- 2 Create a snapshot of the file system and send it to the local file system `fs1` to ensure that data since the last snapshot resides on the to-be-promoted file system (`fs1`). Skip this step if you do not need data written since the last snapshot.

- 3 On the local array, `array1`, promote the local file system `fs1`:

```
array1# purefs promote fs1
```

- 4 Delete file system replica links to the remote file system (`fs1-R`) on the target array. On array 3, issue the `purefs replica-link delete` command, specifying the name of the remote file system, `fs1-R`.

```
array3# purefs replica-link delete fs1-R
```

- 5 Create file system replica links between the promoted file system (`fs1`) and the other target arrays (`array 2` and `array3`) by issuing the `purefs replica-link create --remote-fs <file_system> --remote <target_array>` command. You must specify the source file system (`fs1`).

```
array1# purefs replica-link create --remote-fs fs-R --remote array2 fs1
# purefs replica-link create --remote-fs fs-R --remote array3 fs1
```

- 6 Add back any previously removed file system policy. In the following example, the policy `fs_snap_poll` is added back to file system `fs1`.

```
array1# purepolicy snapshot add fs_snap_poll
```

- 7 Manually direct all clients back to the promoted file system `fs1`.

- 8 On the remote array, `array2`, demote the remote file system `fs1-R`:

```
array2# purefs demote --discard-non-snapshotted-data fs1-R
```

Object Replication Overview

Object Replication

Objects can be replicated from a bucket on a FlashBlade (source) to a remote target bucket residing on another FlashBlade or an S3 target using object replica links. A replica link is the relationship between a source bucket and target bucket, and describes the properties of that connection. During replication, all objects in the source bucket are replicated to the target bucket.

In object replication, up to three arrays can be connected, and up to two replica links can be created from the source bucket to different target buckets (i.e. one source and two targets). If replicating to multiple targets, you can set up links in *fan out* or *cascading* configurations.

Object Replication Fan Out Configuration

In an object replication fan out configuration, a single source bucket has replica links to different target buckets. For example, A→B and A→C. A maximum of two outbound replica links are allowed from a single bucket.

During fan out configuration set up:

- Connect the source FlashBlade array to the first remote FlashBlade or S3 target.
- Connect the source FlashBlade array to the second remote FlashBlade or S3 target.
- Create remote credentials and object replica links on the source FlashBlade array to enable object replication from the source FlashBlade array to the two remote targets.

Object Replication Cascading Configuration

In an object replication cascading configuration, the source bucket has a replica link to a target bucket, which in turn has a replica link to another target bucket. For example, A→B→C.

During cascading configuration set up:

- Connect the source FlashBlade array to the first remote FlashBlade.
- Connect the first remote FlashBlade to the second remote FlashBlade or S3 target.
- Create remote credentials and an object replica link on the source FlashBlade array to the first remote target to enable object replication from the source FlashBlade array to the first remote target.
- Create remote credentials and an object replica link from the first remote target to the second remote target to enable object replication from the first remote target to the second remote target.

Note that for cascading object replication configurations, when creating object replica links between the two remote targets, you must specify the `cascading` option when creating the replica link. If you create the object replica link without specifying this option, objects replicated from another FlashBlade array will not be replicated to the target. Note that no alert is emitted in this scenario. If the object link was created incorrectly, contact Pure Technical Services to baseline your data prior to deleting and re-creating the object replica link.

Object Replication Prerequisites

- Replication traffic requires a replication vip using port 8117; before arrays can replicate, they must first be connected over the management vip/port 443.
- Ports 8117 and 443 are used for encrypted object replication.
- Versioning must be enabled on the target bucket for object replication to proceed, and a lifecycle rule may need to be configured to limit space used by noncurrent versions.

Getting Started

Object replication setup and recovery can be performed using either the Purity//FB GUI or CLI.

The following sections provide an overview of the tasks required for object replication setup. Complete instructions using the GUI and CLI are provided for each in subsequent topics.

Replication Setup

To set up object replication from a bucket on a source FlashBlade array to a remote target bucket on another FlashBlade array, the following tasks must be completed:

- 1 Create replication network interfaces on the source and target arrays.
- 2 Create a connection key on the target array.
- 3 Connect the source and target arrays.
- 4 Create remote credentials for use with the replica link.
- 5 Create a bucket replica link on the local array.
- 6 Enable versioning on the target bucket.
- 7 Optionally configure a lifecycle rule on the target bucket to control space usage of non-concurrent versions.

Complete instructions for completing each of these tasks are provided in:

- [*Setting Up Array to Array Object Replication Using the GUI*](#)
- [*Setting Up Array to Array Object Replication Using the CLI*](#)

To set up object replication from a bucket on a source FlashBlade array to a remote target bucket on an Amazon S3 target, the following tasks must be completed:

- 1 Connect the source FlashBlade array to the S3 target.
- 2 Create remote credentials for use with the replica link.
- 3 Create a bucket replica link on the local array.
- 4 Enable versioning using the AWS Management Console or the Amazon S3 API.
- 5 Optionally enable a lifecycle rule using the AWS Management Console or the Amazon S3 API.

Complete instructions for completing each of these tasks are provided in:

- [Setting Up FlashBlade to Amazon S3 Object Replication Using the GUI](#)
- [Setting Up FlashBlade to Amazon S3 Object Replication Using the CLI](#)

Setting Up Array to Array Object Replication Using the GUI

The following instructions describe how to set up object replication from one FlashBlade array to another FlashBlade array using the Purity//FB GUI. These instructions should be followed in the order that they are presented.

Creating Replication Network Interfaces

Perform the following on both your source and target FlashBlade arrays:

- 1 From the **Settings** sidebar, click **Subnets** under **Network**. The **Subnets** panel appears.
- 2 In the **Subnets** panel, locate the subnet to which you wish to add a network interface and click the **Add interface (+)** button. The **Create Network Interface** pop-up window appears.
- 3 In the **Name** field, enter the interface name.
- 4 In the address field, enter the network address.
- 5 From the **Services** list, select `replication`.
- 6 The subnet defaults to the subnet in which you are creating the interface. Click **Create**.

Creating a Connection Key

 **Important!** The connection key is active for two hours. If the source and target arrays have not been connected within that two hour period, a new connection key must be created.

Once you have created replication network interfaces, you must create a connection key on the target array.

- 1 On the target array, select **Storage > Array** page.
- 2 In the **FlashBlade Array Connections** panel, click **More Options > Create Connection Key**. The **Connection Key** pop-up window appears displaying the new connection key.
- 3 Copy the new connection key.

Connecting the Source and Target Arrays

 **Note:** Multiple file system replica links on a single array are allowed as long as they are using different file systems. Multiple file system replica links are not allowed from a single file system or when using multiple array connections.

 **Note:** When connecting arrays, it is recommended to create the array connection from the array running the most current Purity version to the array running an older Purity version.

With the creation of replication network interfaces on the source and target arrays, and the creation of a connection key on the target array, you can now connect the source and target arrays.

 **Note:** Encryption is only valid for object replication.

 **Important!** When connecting the source array to the target array, a secure connection is required even if encryption is not enabled. Certificates with small key sizes, weak signature algorithms, or that only support weak TLS algorithms are not supported.

 **Important!** The management IP address required during configuration cannot be a loopback IP address.

 **Important!** It is highly recommended that you add the `pure:policy/full-access` access policy to the user (whose credentials are used for replication) on the target cluster, otherwise replication may fail due to denied access. See [Adding Access Policies](#) for instructions on how to add access policies to users.

- 1 On the source array, select **Storage > Array** page, click the add (+) button in the **FlashBlade Array Connections** panel. The **Connect FlashBlade Array** pop-up window appears.
- 2 Enter the target array's hostname or management address in the **Management Address** field. The address can be located from the **Subnets** table by navigating to the **Settings > Network** page.
- 3 Paste the copied connection key for the target array in the **Connection Key** field.
- 4 The replication address is automatically populated on both the source and target cluster unless manually specified. If manually specified, then the target cluster will not have a replication address populated and will need to be manually entered. It is recommended to allow the arrays to automatically discover the replication address, unless using NAT. If using NAT, enter the target array's replication network address in the **Replication Address** field.
- 5 (Optional) Enable encryption by setting the **Encrypted** toggle to on. If set, the `_default_replication_certs` CA certificate group is applied. Before toggling, make sure to import the CA certificate to the source array if the target array's array-certificate is signed by a CA and add it to the certificate group. Note that encryption is set on the source array only. If enabled, the encryption setting displays as enabled on the target array.

! Warning: Use of CA certificates to establish trust and secure communication for object replication is the recommended configuration. If the target array's array-certificate is self-signed, directly importing the target array certificate to the source and adding it to the certificate group will not provide as strict validation around aspects of the peer, such as certificate expiry.
- 6 Click **Connect**.
- 7 If replication address was specified on the source cluster while creating the array connection in step 4, on the target array, enter the source array's replication address.
 - a On the target array, in the **FlashBlade Array Connections** panel on the **Storage > Array** page, click the edit button at the end of the row displaying the connected source array. The **Edit Connected Array** pop-up window appears.
 - b Enter the replication address for the source array in the **Replication Address** field and click **Save**.

Creating Remote Credentials

Object replica links, which allow object data to be replicated between arrays and S3 targets, require remote credentials. In order to create a remote credential, you must have previously created an access key ID and secret access key (see [Creating a User](#) for instructions).

- 1 From the **Protection > Object Replica Links** page, click the add (+) button in the heading of the **Remote Credentials for Object Replication** list. The **Create Remote Credentials** pop-up window appears.
- 2 From the **Remote Array or Target** list, select the connected remote (target) array.

- 3 In the **Name** field, enter the credential name.
- 4 Enter the user access key from the target array in the **Access Key ID** field.
- 5 Enter the user secret access key from the target array in the **Secret Access Key** field.
- 6 Click **Create**.

Creating an Bucket Replica Link

Once the source and target are connected, you must create a bucket replica link on the source array. The bucket replica link contains the details about the source array and bucket from which objects are replicated, and the target bucket to which objects are replicated.

- 1 From the **Protection > Object Replication** page, click the add (+) button in the heading of the **Bucket Replica Links** list. The **Create Bucket Replica Link** pop-up window appears.
- 2 From the **Local Bucket Name** list, select the local (source) bucket from which data will be replicated.
- 3 From the **Remote Array or Target** list, select the remote (target) array.
- 4 (Optional) Click **Connect FlashBlade Array** or **Connect S3 Target** if you did not previously connect an array (for array to array object replication) or S3 (for array to S3 object replication) target.
- 5 From the **Remote Bucket** field, enter the name of the remote (target) bucket to which data will be replicated.
- 6 From the **Remote Credential** list, select the remote credential that was previously created to allow replication.
- 7 Select **Cascading** if the replica link will be used to connect to multiple targets in a cascading configuration. Once the replica link has been created, this setting cannot be changed.
- 8 Click **Create**.

Enabling Bucket Versioning

Bucket versioning allows multiple variants of an object in the same bucket.

Object replication requires versioning to be enabled on the target bucket to prevent write inconsistencies. It is optional to enable versioning on the source bucket. On buckets with versioning enabled, you can use lifecycle rules to control the space used by non-current object versions.

In a versioning-enabled bucket, an object can have multiple versions for the same object name. Each object name has one current version and either zero or multiple non-current versions. Modifying an

existing object results in that instance of the object name becoming the most current version and the previous version becoming a non-current version.

- 1 From the **Storage > Object Store** page, locate the bucket for which you wish to enable versioning in the **Buckets** panel.
- 2 Click the more options button at the end of the row in which the bucket appears and select **Enable versioning**.
- 3 Click **Enable** when prompted for confirmation.

Setting Up Array to Array Object Replication Using the CLI

The following instructions describe how to set up object replication from one FlashBlade array to another FlashBlade array using the Purity//FB CLI. These instructions should be followed in the order that they are presented.

Creating Replication Network Interfaces

On both the source and target arrays, issue the **purenetwork vip create --address ADDRESS --servicelist replication NAME** command, where ADDRESS is the IP address associated with the replication vip and NAME is the data vip.

```
# purenetwork vip create --address 123.123.123.123 --servicelist replication  
datavip1
```

Creating a Connection Key

⚠ Important! The connection key is active for two hours. If the source and target arrays have not been connected within that two hour period, a new connection key must be created.

Once you have created replication network interfaces on the source and target arrays, you must create a connection key on the target array.

- 1 Issue the `purearray create --connection-key` command.

```
# purearray create --connection-key
Connection Key      Created              Expires
T-71a63c01-80aa-4173-842c-8ec9a1285d5b  2020-01-09 14:49:08 PST  2020-01-09
16:49:08 PST
```

- 2 Copy the connection key.

Connecting the Source and Target Arrays

Note: Multiple file system replica links on a single array are allowed as long as they are using different file systems. Multiple filesystem replica links are not allowed from a single file system or when using multiple array connections.

Note: When connecting arrays, it is recommended to create the array connection from the array running the most current Purity version to the array running an older Purity version.

With the creation of replication network interfaces on the source and target arrays, and the creation of a connection key on the target array, you can now connect the source and target arrays.

Important! It is highly recommended that you add the `pure:policy/full-access` access policy to the user (whose credentials are used for replication) on the target cluster, otherwise replication may fail due to denied access. See the `pureobj` command in the *FlashBlade CLI Reference* for information on how to add access policies to users.

Important! When connecting the source array to the target array, a secure connection is required even if encryption is not enabled. Certificates with small key sizes, weak signature algorithms, or that only support weak TLS algorithms are not supported.

Important! The management IP address required during configuration cannot be a loopback IP address.

- 1 On the target array, issue the `purenetwork list` command and make note of the array's management IP address.

```
# purenetwork list
Name      Enabled Subnet Address  ... Services
vir1      True    admin  10.11.123.12  ... management
```

- 2 On the source array, issue the `purearray connect --management-address MANAGEMENT_ADDRESS` command, where `MANAGEMENT_ADDRESS` is the management IP address of the target array.

- If the optional `--replication-address` argument is used, the specified replication address is set on the local array connection. The remote array connection will have no replication address set and must be manually entered.
- If not used, the replication address available on the target array is used, and the replication address of the local array is automatically populated on the target array connection. This is recommended, unless using NAT. If using NAT, specify the target array's replication address using the `--replication address` argument.

```
# purearray connect --management-address 10.11.123.12
Enter the connection key of the target array:
```

- 3 When prompted, enter (paste) the target array's connection key.
- 4 If using NAT, or if the replication address was specified on the source cluster while creating the array connection, on the target array issue the **`purearray setattr --connect --replication address REPLICATION_ADDRESS`** command, where `REPLICATION_ADDRESS` is the replication address of the source array.
- 5 The connection between two arrays in an object replication set up can be encrypted. To enable encryption between two arrays, on the source array run the **`purearray enable --connect --encrypted NAME`** command, where `NAME` is the name of the target array.

Creating Remote Credentials

Object replica links, which allow object data to be replicated between arrays, require remote credentials. In order to create a remote credential, you must have previously created an access key ID and secret access key using the **`pureobj user access-key create --user`** command.

To create remote credentials, issue the **`pureobj remote-credentials create`** command. You must specify the `--access-key-id` argument to set the access key ID and also specify the name of the remote credentials. The remote credentials name must be specified in the format of `<remote_array_name>/<remote_credential_name>`.

In the following example, the credentials `array1/credential1` are created.

```
# pureobj remote-credentials create --access-key-id ABCDE12345 array1/credential1
```

Creating a Bucket Replica Link

Once the source and target are connected, and you have created remote credentials, you must create a bucket replica link on the source array. The bucket replica link contains the details about the

source array and bucket from which objects are replicated, and the target bucket to which objects are replicated.

From the source array, issue the **purebucket replica-link create BUCKET** command, where **BUCKET** is the name of the local bucket from which object data is replicated. You must also specify the name of the remote bucket and the remote credentials using the **--remote-bucket** and **--remote-credentials** arguments.

```
# purebucket replica-link create --remote-bucket r_bucket1 --remote-credentials  
array1/credential1 bucket1
```

In this example, the replica link is created between the local bucket `bucket1` on the local array and the remote bucket `r_bucket1`.

If setting up a cascading configuration, specify **--cascading** when running the **purebucket replica-link create BUCKET** command. Once the replica link has been created, this setting cannot be changed.

Enabling Bucket Versioning

Bucket versioning allows multiple variants of an object in the same bucket.

Object replication requires versioning to be enabled on the target bucket. Versioning is not required on the source bucket. On buckets with versioning enabled, you can use lifecycle rules to control the space used by non-current object versions.

In a versioning-enabled bucket, an object can have multiple versions for the same object name. Each object name has one current version and either zero or multiple non-current versions. Modifying an existing object results in that instance of the object name becoming the most current version and the previous version becoming a non-current version.

Issue the **purebucket enable --versioning BUCKET** command, where `bucket` is the name of the remote bucket.

```
# purebucket enable --versioning r_bucket1
```

In this example, versioning is enabled on the remote bucket `r_bucket1`.

Setting Up FlashBlade to Amazon S3 Object Replication Using the GUI

The following instructions describe how to set up object replication from one FlashBlade array to an Amazon S3 target using the Purity//FB GUI. These instructions should be followed in the order that they are presented.

Creating a Replication Network Interface

Perform the following on the source FlashBlade array:

- 1 Select **Settings > Network**.
- 2 In the **Subnets** list, locate the subnet to which you wish to add a network interface and click the **Add interface (+)** button. The **Create Network Interface** pop-up window appears.
- 3 In the **Name** field, enter the interface name.
- 4 In the address field, enter the network address.
- 5 From the **Services** list, select `replication`.
- 6 The subnet defaults to the subnet in which you are creating the interface. Click **Create**.

Connecting the FlashBlade Array to the S3 Target

With the creation of replication network interfaces on the source array, you can now connect the source array and S3 target.

- 1 On the source array, select **Storage > Array** page, click the add (+) button in the **S3 Target Connections** panel. The **Connect S3 Target** pop-up window appears.
- 2 Enter the S3 target name in the **Name** field.
- 3 Enter the S3 target's address in the **Address** field. For example, `s3.amazonaws.com`.
- 4 Click **Connect**.

Creating Remote Credentials

Object replica links, which allow object data to be replicated between arrays and S3 targets, require remote credentials. In order to create a remote credential, you must have previously created an access key ID and secret access key (see [Creating a User](#) for instructions).

- 1 From the **Protection > Object Replica Links** page, click the add (+) button in the heading of the **Remote Credentials for Object Replication** list. The **Create Remote Credentials** pop-up window appears.
- 2 From the **Remote Array or Target** list, select the connected remote (target) array.
- 3 In the **Name** field, enter the credential name.
- 4 Enter the user access key from the target array in the **Access Key ID** field.
- 5 Enter the user secret access key from the target array in the **Secret Access Key** field.
- 6 Click **Create**.

Creating an Bucket Replica Link

Once the source and target are connected, you must create a bucket replica link on the source array. The bucket replica link contains the details about the source array and bucket from which objects are replicated, and the target bucket to which objects are replicated.

- 1 From the **Protection > Object Replication** page, click the add (+) button in the heading of the **Bucket Replica Links** list. The **Create Bucket Replica Link** pop-up window appears.
- 2 From the **Local Bucket Name** list, select the local (source) bucket from which data will be replicated.
- 3 From the **Remote Array or Target** list, select the remote (target) array.
- 4 (Optional) Click **Connect FlashBlade Array** or **Connect S3 Target** if you did not previously connect an array (for array to array object replication) or S3 (for array to S3 object replication) target.
- 5 From the **Remote Bucket** field, enter the name of the remote (target) bucket to which data will be replicated.
- 6 From the **Remote Credential** list, select the remote credential that was previously created to allow replication.
- 7 Select **Cascading** if the replica link will be used to connect to multiple targets in a cascading configuration. Once the replica link has been created, this setting cannot be changed.
- 8 Click **Create**.

Enabling Bucket Versioning

Bucket versioning allows multiple variants of an object in the same bucket.

Object replication requires versioning to be enabled on the target bucket. Versioning is not required on the source bucket. On buckets with versioning enabled, you can use lifecycle rules to control the space used by non-current object versions.

In a versioning-enabled bucket, an object can have multiple versions for the same object name. Each object name has one current version and either zero or multiple non-current versions. Modifying an existing object results in that instance of the object name becoming the most current version and the previous version becoming a non-current version.

Bucket versioning can be enabled on an Amazon S3 bucket via the AWS Management Console, the Amazon S3 API, or the AWS CLI.

Optionally configure a lifecycle rule on the target bucket to control space usage of non-concurrent versions.

Setting Up FlashBlade to Amazon S3 Object Replication Using the CLI

The following instructions describe how to set up object replication from one FlashBlade array to an Amazon S3 target using the Purity//FB CLI. These instructions should be followed in the order that they are presented.

Creating a Replication Network Interface

On the source array, issue the

purenetwork vip create --address ADDRESS --servicelist replication command, where ADDRESS is the IP address associated with the replication vip.

```
# purenetwork vip create --address 123.123.123.123 --servicelist replication
```

Connecting the FlashBlade Array to the S3 Target

On the FlashBlade array, issue the **puretarget s3 create TARGET** command, where **TARGET** is the name of the S3 target. You must use the **--address** argument to specify the address of the S3 target.

```
# puretarget s3 create --address s3.amazonaws.com s3r1
```

In this example, the S3 target name is `s3r1` and its address is `s3.amazonaws.com`.

Creating Remote Credentials

Object replica links, which allow object data to be replicated between arrays, require remote credentials. In order to create a remote credential, you must have previously created an access key ID and secret access key using the **pureobj user access-key create --user** command.

To create remote credentials, issue the **pureobj remote-credentials create** command. You must specify the **--access-key-id** argument to set the access key ID and also specify the name of the remote credentials. The remote credentials name must be specified in the format of `<remote_array_name>/<remote_credential_name>`.

In the following example, the credentials `array1/credential1` are created.

```
# pureobj remote-credentials create --access-key-id ABCDE12345 array1/credential1
```

Creating a Bucket Replica Link

Once the source and target are connected, and you have created remote credentials, you must create a bucket replica link on the source array. The bucket replica link contains the details about the source array and bucket from which objects are replicated, and the target bucket to which objects are replicated.

From the source array, issue the **purebucket replica-link create BUCKET** command, where **BUCKET** is the name of the local bucket from which object data is replicated. You must also specify the name of the remote bucket and the remote credentials using the **--remote-bucket** and **--remote-credentials** arguments.

```
# purebucket replica-link create --remote-bucket r_bucket1 --remote-credentials array1/credential1 bucket1
```

In this example, the replica link is created between the local bucket `bucket1` on the local array and the remote bucket `r_bucket1`.

If setting up a cascading configuration, specify `--cascading` when running the **`purebucket replica-link create BUCKET`** command. Once the replica link has been created, this setting cannot be changed.

Enabling Bucket Versioning

Bucket versioning allows multiple variants of an object in the same bucket.

Object replication requires versioning to be enabled on the target bucket. Versioning is not required on the source bucket. On buckets with versioning enabled, you can use lifecycle rules to control the space used by non-current object versions.

In a versioning-enabled bucket, an object can have multiple versions for the same object name. Each object name has one current version and either zero or multiple non-current versions. Modifying an existing object results in that instance of the object name becoming the most current version and the previous version becoming a non-current version.

Bucket versioning can be enabled on an Amazon S3 bucket via the AWS Management Console, the Amazon S3 API, or the AWS CLI.

Optionally configure a lifecycle rule on the target bucket to control space usage of non-concurrent versions.

Chapter 8

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 **Policies** page provides the following four tabs:

- **NFS** : Displays and allows you to manage NFS export policies and NFS export rules. See [NFS Export Policies](#) for additional information.
- **Object Store** : Displays and allows you to manage object store policies. See [Object Store Access Policies](#) for additional information.
- **SMB**: Displays and allows you to manage SMB share policies. See [SMB Client and Share Policies](#) for additional information.
- **Snapshot**: Displays and allows you to manage file system snapshot replication policies. See [File System Snapshot Policies](#) for additional information.

Figure 8.1 Purity//FB Policies Page

The screenshot shows the Purity//FB GUI Policies page. The sidebar on the left contains navigation links: Dashboard, Storage, Protection, Policies (selected), Analysis, Performance, Capacity, Replication, Network, Health, and Settings. Below these are links for Help, End User Agreement, Terms, Log Out, Array, Oban, FlashBlade Software, Purity//FB 4.1.0, and a status bar showing 'Logged in as pureuser' and 'GMT-08:00 (PST)'. The main content area is titled 'Policies' and has tabs for NFS, Object Store, SMB, and Snapshot. The 'NFS' tab is selected, showing a sub-tab '> NFS'. The main content area displays two tables: 'NFS Export Policies' and 'NFS Export Rules'. The 'NFS Export Policies' table has columns for Name, Type, and Enabled. The 'NFS Export Rules' table has columns for Policy, Client, Permission, Security, Access, Anon UID, Anon GID, Secure, fileid-32bit, atime, and Index.

Name	Type	Enabled
fs-scale-policy	nfs-export	True
Gollum-Clients-ESX	nfs-export	True
iozone-clients	nfs-export	True
JH_test12	nfs-export	False
krb5	nfs-export	True
krb5-krbi	nfs-export	True
krb5-krbi-krbp	nfs-export	True
noatime	nfs-export	True
oban-nfsv3-secure-clients	nfs-export	True
replica_oban_1	nfs-export	True

Policy	Client	Permission	Security	Access	Anon UID	Anon GID	Secure	fileid-32bit	atime	Index
Gollum-Clients-ESX	10.64.129.60	rw	sys	no-squash	-	-	False	False	False	1


NFS Export Policies

NFS export policies are made up of 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.

The **NFS** tab allows you to view, create, and manage NFS export policies and NFS export rules.

Figure 8.2 NFS Tab



NFS Export Policies

Name	Type	Enabled	
fs-scale-policy	nfs-export	True	
Gollum-Clients-ESX	nfs-export	True	
iozone-clients	nfs-export	True	
JH_test2	nfs-export	False	
krb5	nfs-export	True	
krb5-krbi	nfs-export	True	
krb5-krbi-krbp	nfs-export	True	
noatime	nfs-export	True	
oban-nfsv3-secure-clients	nfs-export	True	
replica_oban_1	nfs-export	True	

NFS Export Rules

Policy	Client	Permission	Security	Access	Anon UID	Anon GID	Secure	fileid-32bit	atime	Index
Gollum-Clients-ESX	10.64.129.60	rw	sys	no-squash	-	-	False	False	False	1

The **NFS** tab 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.
- **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 NFSv4). 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.

Figure 8.3 Viewing NFS Export Policy Details

Policies

2

16

NFS
Object Store
SMB
Snapshot

NFS > Gollum-Clients-ESX

NFS Export Rules

1-10 of 63 < > + :

Client	Permission	Security	Access	Anon UID	Anon GID	Secure	fileid-32bit	atime	Index	
10.64.129.60	rw	sys	no-squash	-	-	False	False	False	1	:
10.64.137.223	rw	sys	no-squash	-	-	False	False	False	2	:
10.64.139.141	rw	sys	no-squash	-	-	False	False	False	3	:
10.64.137.243	rw	sys	no-squash	-	-	False	False	False	4	:
10.64.140.20	rw	sys	no-squash	-	-	False	False	False	5	:
10.64.141.246	rw	sys	no-squash	-	-	False	False	False	6	:
10.64.135.229	rw	sys	no-squash	-	-	False	False	False	7	:
10.64.133.119	rw	sys	no-squash	-	-	False	False	False	8	:
10.64.139.119	rw	sys	no-squash	-	-	False	False	False	9	:
10.64.137.193	rw	sys	no-squash	-	-	False	False	False	10	:

Details

Name

Gollum-Clients-ESX

Type

nfs-export

Enabled

True

Attached File Systems

No file systems found.

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

The **Attached File Systems** panel displays all file systems governed by the policy. It also allows you to add and remove file systems 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 not supported for NFSv3.

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 8.1 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. Supported with NFSv4.1 only.
sec=krb5i	Enables Kerberos security authentication and integrity checking of NFS operations. Supported with NFSv4.1 only.
sec=krb5p	Enables Kerberos security authentication, integrity checking of NFS operations, and NFS traffic encryption. Supported with NFSv4.1 only.
fileid_32bit	Allows 32-bit inode support for clients.

Table 8.1 Export Rules (continued)

Export Rules	Rules Description
<code>nofileid_32bit</code>	Disables 32-bit inode support for clients.
<code>root_squash</code>	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 <code>anonuid</code> . All users with GID 0 will have their GID mapped to <code>anongid</code> .
<code>no_root_squash</code>	Allows root users and groups to access the file system with root privilege.
<code>all_squash</code>	Maps all UIDs (including root) to <code>anonuid</code> and all GIDs (including root) to <code>anongid</code> .
<code>no_all_squash</code>	Prevents the remapping of user and group IDs to <code>anonuid 65534</code> or <code>anongid 65534</code> . All users and groups will retain their IDs unless <code>root_squash</code> is also specified.
<code>anonuid</code>	Any user whose UID is affected by <code>root_squash</code> or <code>all_squash</code> will have their UID mapped to <code>anonuid</code> . The default <code>anonuid</code> is 65534.
<code>anongid</code>	Any user whose GID is affected by <code>root_squash</code> or <code>all_squash</code> will have their GID mapped to <code>anongid</code> . The default <code>anongid</code> is 65534.
<code>atime</code>	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.
<code>noatime</code>	Disables the update of inode access times after read operations.

Table 8.1 Export Rules (continued)

Export Rules	Rules Description
secure	Prevents NFS access to client connections coming from non-reserved ports. Applies to NFSv3, NFSv4.1, and auxiliary protocols MOUNT and NLN.
insecure	Allows NFS access to client connections coming from non-reserved ports. Applies to NFSv3, NFSv4.1, and auxiliary protocols MOUNT and NLN.

 **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 | Netgroup > *. 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 8.1](#).

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 8.1](#).

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 <code>ro</code> access and <code>root_squash</code> privileges to the file system.
* ()	All clients have default <code>ro</code> access and <code>root_squash</code> 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 <code>rw</code> access and <code>root_squash</code> (default) privileges.

Export Rules	Rules Description
2620::/16(no_root_squash)	Client 2620::/16 has <code>ro</code> (default) access and <code>no_root_squash</code> privileges.
*(rw,no_root_squash)	All clients have <code>rw</code> access and <code>no_root_squash</code> privileges. This is the default rule used when creating a file system.
*(rw,root_squash,fileid_32bit)	All clients have <code>rw</code> access and <code>root_squash</code> 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 <code>rw</code> 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** tab, 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 (Optional) Select to restrict client ports to secure ports only.
- 12 (Optional) Select to use 32-bit inodes instead of 64-bit inodes.
- 13 (Optional) Select to update the timestamp when accessed.
- 14 Click **Save**.

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

Attaching File Systems to an NFS Export Policy

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

To attach file systems to an NFS export policy, perform the following:

- 1 From the **Policies > NFS > NFS Export Policies** panel, click the NFS export policy to which you wish to attach file systems. The detail page for that policy appears.
- 2 Click **More Options > Add File Systems** in the heading of the **Attached File Systems** panel. The **Add File Systems** pop-up window appears.

 **Note:** By default, the file systems 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.

- 3 From the **Available File Systems** panel, select one or more file systems you wish to attach to the policy. If you wish to attach 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 Systems** panel.
- 4 Click **Add**.

You can alternatively attach NFS export policies to file systems from the **Storage > File Systems** page.

- 1 From the **Storage > File Systems** page, select the file system to which you wish to attach an NFS export policy. The details page for the that file system appears.
- 2 In the **Attached Policies** panel, click **More Options > Set NFS Export Policy**. The **Set NFS Export Policy** pop-up window appears.
- 3 From the **Selected Policy** list, select the NFS export policy that you wish to attach.
- 4 Click **Save**.

Moving File Systems to Another Policy

One or more file systems can be moved from one policy to another.

To move file systems from one policy to another, perform the following:

- 1 From the **Policies > NFS > NFS Export Policies** panel, click the NFS export policy from which you wish to move file systems. The detail page for that policy appears.
- 2 Click **More Options > Move File Systems** in the heading of the **Attached File Systems** panel. The **Move File Systems to Another Policy** pop-up window appears.

 **Note:** By default, the file systems 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.

- 3 From the **Current File Systems** panel, select one or more file systems you wish to move to a different policy. If you wish to move 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 Systems** panel.
- 4 Select the policy to which the file systems will be moved from the **Target Policy** list.
- 5 Click **Move**.

Removing File Systems from an NFS Export Policy

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

- 1 To remove a single file system from an NFS export policy:
 - a From the **Policies > NFS > NFS Export Policies** panel, click the NFS export policy from which you wish to remove a file system. The detail page for that policy appears.
 - b In the **Attached File Systems** panel, click the remove (X) button on the same row as the file system you wish to remove. The **Remove File System from Policy** pop-up window appears.
 - c Click **Remove**.
- 2 To remove multiple file systems from an NFS export policy:
 - a From the **Policies > NFS > NFS Export Policies** panel, click the NFS export policy from which you wish to remove a file system. The detail page for that policy appears.
 - b Click **More Options > Remove File Systems** in the heading of the **Attached File Systems** panel. The **Remove File Systems** pop-up window appears.

 **Note:** By default, the file systems 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 **Current File Systems** panel, select one or more file systems 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 Systems** panel.
- d Click **Remove**.

You can alternatively remove NFS export policies from file systems from the **Storage > File Systems** page.

- 1 From the **Storage > File Systems** page, select the file system from which you wish to remove an NFS export policy. The details page for the that file sytem appears.
- 2 To remove a single NFS export policy, click the remove (X) button in the same row as the policy you wish to remove.
- 3 To remove multiple NFS export policies:
 - a In the **Attached Policies** panel, click **More Options > Remove**. The **Remove Policies** pop-up window appears.
 - b Select the policies you wish to remove from the **Available Policies** panel. Each selected policy willl appear in the **Selected Policies** panel.
 - c Click **Remove**.

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 > 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 (Optional) Select to restrict client ports to secure ports only.
- 10 (Optional) Select to use 32-bit inodes instead of 64-bit inodes.
- 11 (Optional) Select to update the timestamp when accessed.

- 12 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.
- 13 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 > 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 > 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 (Optional) Select to restrict client ports to secure ports only.
- 10 (Optional) Select to use 32-bit inodes instead of 64-bit inodes.
- 11 (Optional) Select to bypass updating the timestamp when accessed.
- 12 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 > 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 > 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 > 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 > 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 > 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**.

Object Store Access Policies

Object store access policies are made up of Pure-managed policies and customer-managed policies, which are rules that govern object accessibility, creation, and modification.

Figure 8.4 Object Store Tab

Policies 3				
NFS	Object Store	SMB	Snapshot	
Object Store				
Object Store Access Policies 1:11 of 11				
Name ▲	Created	Updated	Description	
pure:policy/bucket-configure	2020-10-07 17:00:00	2020-10-07 17:00:00	Permissions to configure buckets -- for example, to configure versioning and lifecycle rules, which determine when and how objects are permanently deleted.	
pure:policy/bucket-create	2020-10-07 17:00:00	2020-10-07 17:00:00	Permissions to create buckets.	
pure:policy/bucket-delete	2020-10-07 17:00:00	2020-10-07 17:00:00	Permissions to delete buckets, if empty.	
pure:policy/bucket-info	2020-10-07 17:00:00	2020-10-07 17:00:00	Permissions to read bucket configuration, such as bucket versioning state and lifecycle configuration.	
pure:policy/bucket-list	2020-10-07 17:00:00	2020-10-07 17:00:00	Permission to list bucket names and bucket creation timestamps in the account.	
pure:policy/full-access	2020-10-07 17:00:00	2020-10-07 17:00:00	Permissions to use all APIs supported by FlashBlade Object Store, including all S3 APIs.	

The **Object Store** tab displays a list of all object store access policies:

- **Name:** The policy name.
- **Created:** The date the policy was created.
- **Updated:** The date the policy was last modified.
- **Description:** A description of the permissions of the policy.

For additional information about object store access policies and rules see [About Object Store Access Policies](#) and [Rules: Actions, Resources, and Conditions](#).

About Object Store Access Policies

FlashBlade Object Store allows you to manage what actions can be carried out by what users when accessing objects and buckets via the S3 API. Object store access policies consist of Pure-managed user policies and customer-managed policies.

You can assign a user permissions by choosing from a set of predefined, Pure-managed user policies, or *access policies*. Adding an access policy to a user in Object Store grants that user permissions to make specific types of S3 API calls on any object or bucket within the Object Store account where the user was created.

Pre-defined policies cannot be modified or deleted. They can only be added or removed from users.

The following table provides a list of the available Pure-managed policies:

Table 8.3 Pure-managed Policies

Access Policy	Policy Definition	Grant*
Storage Admin Roles		
pure:policy/bucket-list	Permissions to list all bucket names for the account.	s3:ListAllMyBuckets
pure:policy/bucket-info	Permissions to read bucket configurations.	s3:GetBucketVersioning s3:GetLifecycleConfiguration s3:GetBucketAcl s3:GetBucketLocation s3:GetObjectLockConfiguration
pure:policy/bucket-configure	Permissions to configure buckets such as versioning and lifecycle rules on buckets.	s3:PutBucketVersioning s3:PutLifecycleConfiguration s3:PutObjectLockConfiguration
pure:policy/bucket-create	Permissions to create buckets.	s3:CreateBucket
pure:policy/bucket-delete	Permissions to delete empty buckets.	s3>DeleteBucket
Application Roles		
pure:policy/object-list	Permissions to list objects and versions and verify their sizes, entity tags (ETags), and timestamps (but not to read their data or custom metadata).	s3:ListBucket s3:ListBucketVersions s3:ListBucketMultipartUploads

Table 8.3 Pure-managed Policies (continued)

Access Policy	Policy Definition	Grant*
pure:policy/object-lock	Permissions to configure Object Lock.	s3:GetObjectLockRetention s3:GetObjectLockLegalHold s3:PutObjectRetention s3:PutObjectLegalHold
pure:policy/object-lock-bypass-governance	Permissions to bypass Object Lock governance retention.	s3:BypassGovernanceRetention
pure:policy/object-read	Permissions to read object and version data and custom metadata.	s3:GetObject s3:GetObjectVersion s3:GetObjectAcl s3:GetObjectRetention s3:GetObjectLegalHold
pure:policy/object-write	Permissions to create and overwrite objects when versioning is not enabled and to create new versions when versioning is enabled.	s3:PutObject s3:AbortMultipartUpload s3:ListMultipartUploadsParts
pure:policy/object-delete	Permission to eradicate objects when versioning is not enabled on the bucket and to create object delete markers when versioning is enabled.	s3:DeleteObject
pure:policy/version-delete	Permissions to eradicate objects and versions.	s3:DeleteObjectVersion
Special Roles		

Table 8.3 Pure-managed Policies (continued)

Access Policy	Policy Definition	Grant*
pure:policy/full-access	Full Access is a special role. Permissions for all current and future S3 APIs and any non-standard FlashBlade S3 API extensions. (Triggers a warning in the GUI if selected)	s3:*
pure:policy/safemode-configure**	Permissions to configure safemode.	s3:ExtendSafemodeRetentionPeriod

*Subject to change. AWS occasionally updates the permissions defined in an AWS managed policy. When AWS does this, the update affects all principal entities (users, groups, and roles) to which the policy is attached. This does not affect the security of FlashBlade access policies.

**This policy will only be visible if Object Safemode is enabled or this policy is in use.

As noted in the table above, each policy allows specific S3 grants (actions). For more detailed information about the S3 grants (actions) listed above, refer to [Rules: Actions, Resources, and Conditions](#).

In addition to the twelve pre-configured Pure-managed user policies above, you can also create and modify your own policies to manage your objects and buckets by specifying specific actions, resources, and conditions. (see [Creating Object Store Access Policies](#)).

Rules: Actions, Resources, and Conditions

A user-created object store access policy can have S3 actions attached to the policy's rules. Additionally, resources and conditions can be specified in the rule to further narrow the focus of a policy's actions.

- An object store access policy is a collection of rules.
- Each rule can then match any given S3 request made by the user if the request is consistent with the rule's action, resources, and conditions (if relevant). Each rule is an independent permission grant on a specific combination of actions, resources, and optionally, conditions.
- If the request matches a rule, then the request is allowed.
- If a request does not match any rules, then the request is implicitly denied.

Actions

During rule configuration for a user-created object store access policy, one or more of the following S3 actions can be attached to the rule:

Table 8.4 S3 Actions

S3 Action	Permission
s3:*	Permission to use all S3 APIs supported by FlashBlade Object Store. (Triggers a warning in the GUI if selected)
s3:ExtendSafemodeRetentionPeriod	Permissions to configure safemode.
s3:ListAllMyBuckets	Permission to list all bucket names for the account.
s3:GetBucketVersioning	Permission to read bucket configurations.
s3:GetLifecycleConfiguration	Permission to read a bucket's lifecycle configuration.
s3:GetBucketAcl	Permission to read a bucket's access control list. Bucket ACLs on FlashBlade are always "FULL_CONTROL" since user policies control access.
s3:GetBucketLocation	Permission to read the location where a bucket is stored. Bucket locations on FlashBlade are always "on-prem".
s3:PutBucketVersioning	Permission to read a bucket's versioning state.
s3:PutLifecycleConfiguration	Permission to set a bucket's lifecycle configuration. Lifecycle configuration controls whether objects, old object versions, and delete markers in a bucket are permanently deleted after a period of time.
s3:CreateBucket	Permission to create buckets.
s3>DeleteBucket	Permission to delete empty buckets.

Table 8.4 S3 Actions (continued)

S3 Action	Permission
s3:ListBucket	Permission to list metadata about current object versions in a bucket or "folder" (object prefix), such as object "names" (keys), sizes, timestamps, hashes (entity tags), and prefixes, but not delete markers or old object versions.
s3:ListBucketVersions	Permission to list metadata about all object versions in a bucket or "folder" (object prefix), such as object "names" (keys), sizes, timestamps, hashes (entity tags), and prefixes, including delete markers and old object versions.
s3:ListBucketMultipartUploads	Permission to list uploaded parts of an in-progress multipart upload, for example, when resuming an upload.
s3:GetObject	Permission to read object data and metadata of current object versions and delete markers, but not old object versions.
s3:GetObjectVersion	Permission to read object data and metadata of all object versions, including delete markers and old object versions.
s3:GetObjectAcl	Permission to read an object's access control list (ACL). Object ACLs on FlashBlade are always "FULL_CONTROL" since user policies control access.
s3:PutObject	Permission to create new objects, permanently overwrite objects when versioning is not enabled, and create new object versions when versioning is enabled. Permanent deletion of old object versions may be controlled by the bucket's lifecycle configuration.
s3:AbortMultipartUpload	Permission to cancel an in-progress multipart upload.
s3:ListMultipartUploadParts	Permission to list uploaded parts of an in-progress multipart upload, for example, when resuming an upload.

Table 8.4 S3 Actions (continued)

S3 Action	Permission
s3:DeleteObject	Permission to permanently delete objects when versioning is not enabled, or to mark objects as deleted when versioning is enabled. Permanent deletion may also be controlled by the bucket's lifecycle configuration.
s3:DeleteObjectVersion	Permission to permanently delete objects, old object versions, and delete markers, whether or not versioning is enabled. Permanent deletion may also be controlled by the bucket's lifecycle configuration.
s3:GetObjectLockConfiguration	Permission to read the Object Lock configuration for a bucket. (Part of pure:policy/bucket-info)
s3:PutObjectLockConfiguration	Permission to set the Object Lock configuration for a bucket. The rule specified in the Object Lock configuration will be applied by default to every new object placed in the specified bucket. (Part of pure:policy/bucket-configure)
s3:GetObjectRetention	Permission to read an object's retention settings. (Part of pure:policy/object-read and pure:policy/object-lock)
s3:PutObjectRetention	Permission to create an Object Retention configuration on an object. (Part of pure:policy/object-lock)
s3:GetObjectLegalHold	Permission to read an object's legal hold settings. (Part of pure:policy/object-read and pure:policy/object-lock)
s3:PutObjectLegalHold	Permission to create a Legal Hold configuration on an object. (Part of pure:policy/object-lock)
s3:BypassGovernanceRetention	Permission to change an object's retention period, if the object is currently protected by Governance retention.(Part of pure:policy/object-lock-bypass-governance)

Resources

Resources are the buckets/objects to which a rule's actions are applied.

When configuring a rule, you must specify the bucket (and optionally the bucket's object) to which the rule will apply. The bucket and object components must be specified in the format `bucket_name/object_name`; for example, `bucket1/obj1`. More than one bucket/object can be specified and wildcards (`?` to match any single character, and `*` to match any number of characters) may be used.

For example:

- If a simple rule has been created specifying `bucket1` as the resource with the action `s3:CreateBucket`, then the rule would allow a bucket named `bucket1` to be created.
- A rule created with specifying `dev*` as the resource with the action `s3:BucketCreation` would allow the creation of any bucket whose names start with `dev`, for instance `devops`, `devA`, `devB`, etc.
- By further specifying an object component with the bucket resource, you can limit where actions occur. For example, a rule specifying the bucket/object resources `dev/devobj` with the actions `s3:CreateBucket` and `s3:PutObject` will allow the creation of a bucket named `dev` and allow the update of an object named `devobj` in the `dev` bucket.
 - A rule instead specifying the bucket/object resource `dev*` with the same actions will allow the creation of buckets whose names match `dev*` and allow the update of any objects in matching buckets.
 - A rule instead specifying the bucket/object resource `dev*/devobj` with the same actions will not allow any bucket creation, as no bucket names can match the resource provided. However, the rule will allow updating objects named `devobj` or which end in `/devobj` in buckets whose names match `dev*`.
 - A rule specifying both `dev*` and `dev*/devobj` as bucket/object resources, with the same actions, would behave exactly as the rule specifying just `dev*`, as the former grants a superset of the permissions which the latter grants.

Conditions

When configuring a rule, you also have the option to specify conditions that must be met in order for the actions in the rule to execute. You can optionally specify the following:

- **Source IP Addresses** - A list of IP addresses and subnets in standard CIDR format from which the rule will allow requests. This condition can be used with any action. IPv4 and IPv6 are allowed.
- **S3 Prefixes** - A list of object prefixes to restrict access and results based on the prefix of the relevant objects. This can include a trailing delimiter character and an optional wildcard. This condition can only be used with the `s3:ListBucket` and `s3:ListBucketVersions` actions.
- **S3 Delimiters** - A list of valid prefix delimiter characters to group result objects by the specified delimiter. Only top-level groupings will be returned. This condition can only be used with the `s3:ListBucket`, `s3:ListBucketMultipartUploads`, and `s3:ListBucketVersions` actions.

For example:


- If you create a rule specifying `dev` as the resource with the action `s3:CreateBucket` and a condition limiting bucket creation to the IP address `168.12.145.12`, then creation of a bucket named `dev` is allowed only from the IP address `168.12.145.12`.
- If you create a rule specifying `bucket1` as the resource with the action `s3:ListBucket` and specify the S3 prefix condition of `obj*`, requests will be allowed showing only information about `bucket1` objects matching `obj`.
- In a scenario with the following objects:
 - `eng1/obj1`
 - `eng1/obj2`
 - `eng2/obj1`
 - `eng2/obj2`

If you create a rule with an `s3:ListBucket` action and specify an S3 prefix condition of `eng`, then requests for object information are allowed for all four objects. However if `/` is also specified as an S3 delimiter condition, then requests for object information are limited to the two objects `eng1` and `eng2` only (the top-level groupings are returned).

Viewing Object Store Policy Details

To view details about each object store access policy in the array, from the **Policies > Object Store** tab, click a policy name to open a detail page about that policy.

Figure 8.5 Object Store Policy Details



Policies 3 ×

NFS **Object Store** SMB Snapshot

 > Object Store > pure:policy/bucket-configure

Rules 1-1 of 1

Name ▲	Effect	Actions	Resources	Conditions	
bucketconfigure1	allow	s3:PutBucketVersioning ⓘ s3:PutLifecycleConfiguration ⓘ	*	-	⋮

Details

Name pure:policy/bucket-configure
Type object-store-access
Enabled True
Created 2020-10-07 17:00:00
Updated 2020-10-07 17:00:00
ARN pure:policy/bucket-configure
Description Permissions to configure buckets -- for example, to configure versioning and lifecycle rules, which determine when and how objects are permanently deleted.

Users ⋮

No users found.

The policy's members, rules, and status are displayed in the **Users**, **Rules**, and **Details** panels, respectively.

Creating Object Store Access Policies

To create an object store access policy, perform the following:

- 1 From the **Policies > Object Store** tab, click the **Add (+)** button in the heading of the **Object Store Access Policies** list. The **Step 1: Create Object Store Access Policy** pop-up window appears.
- 2 From the **Account** list select the object store account to which the policy will be applied.
- 3 In the **Name** field, enter a name for the policy.
- 4 (Optional) In the **Description** field, enter a brief description of the the policy's permissions.
- 5 Click **Create**. The **Step 2: Add Rule to Policy <name>** pop-up window appears. You can click **Skip** to exit the pop-up window, or proceed and add rules to your policy.
- 6 In the **Name** field, enter a name for the policy rule.

- 7 In the **Effect** field, select to **Allow** or **Deny** S3 requests that match the actions and parameters set with this rule. Note that if a policy contains both an Allow rule and a Deny rule, the Deny rule will take precedent over the Allow rule (this is also known as an explicit Deny rule).
- 8 In the **Actions** field, click the **Add (+)** button. The **Add Actions** pop-up window appears.
- 9 From the **Available Actions** panel, select the actions that you wish the rule to allow. All selected actions appear in the **Selected Actions** panel.

If you wish to select all actions, select the checkbox directly beneath the **Available Actions** heading.
- 10 Click **Add**.
- 11 (Optional) Select the **Ignore Action Restriction Enforcement** checkbox if you wish the rule to allow combinations of actions, resources, and conditions which would be otherwise prevented.
- 12 In the **Resources** field, click the **Add (+)** button. Enter the object store resources (bucket component optionally followed by an object component) to which the selected action(s) will apply. The bucket and object components must be specified in the format `bucket_name/object_name`; for example, `bucket1/obj1`. For example, if you enter `bucket1` as the resource and selected the action of `s3:CreateBucket`, the rule will only allow the creation of a bucket named `bucket1`.
- 13 For additional optional rule parameters specifying the conditions in which the rule will be applied, click **Optional Conditions**.
- 14 In the **Source IPs** field, click the **Add (+)** button. Enter the IP addresses or subnets to which the rule will be applied. Multiple entries can be entered with a comma-separated list. Each entry must be in standard CIDR format.
- 15 In the **S3 Prefixes** field, click the **Add (+)** button. Enter the object prefixes, including a trailing delimiter character and optional wildcard, on which access and results will be restricted for the `s3:ListBucket` and `s3:ListBucketVersions` actions. Multiple entries can be entered with a comma-separated list.
- 16 In the **S3 Delimiters** field, click the **Add (+)** button. Enter the prefix delimiter characters (such as a trailing slash) to be used to group result objects for the `s3:ListBucket` and `s3:ListBucketVersions` actions. Multiple entries can be entered with a comma-separated list.
- 17 Click **Add**.

Adding Object Store Access Policy Rules

If you did not add rules to your object store access policy during its creation, you can add rules to the policy as follows:

- 1 From the **Policies > Object Store** tab, click the policy to which you wish to add a rule. The policy's detail page appears.

- 2 In the **Rules** panel, click the **Add (+)** button. The **Add Rule to Policy <name>** pop-up window appears.
- 3 In the **Name** field, enter a name for the policy rule.
- 4 In the **Effect** field, select to **Allow** or **Deny** S3 requests that match the actions and parameters set with this rule. Note that if a policy contains both an Allow rule and a Deny rule, the Deny rule will take precedent over the Allow rule (this is also known as an explicit Deny rule).
- 5 In the **Actions** field, click the **Add (+)** button. The **Add Actions** pop-up window appears.
- 6 From the **Available Actions** panel, select the actions that you wish the rule to allow. All selected actions appear in the **Selected Actions** panel.

If you wish to select all actions, select the checkbox directly beneath the **Available Actions** heading.
- 7 Click **Add**.
- 8 (Optional) Select the **Ignore Action Restriction Enforcement** checkbox if you wish the rule to allow combinations of actions, resources, and conditions which would be otherwise prevented.
- 9 In the **Resources** field, click the **Add (+)** button. Enter the object store resources (bucket component optionally followed by an object component) to which the selected action(s) will apply. The bucket and object components must be specified in the format `bucket_name/object_name`; for example, `bucket1/obj1`. For example, if you enter `bucket1` as the resource and selected the action of `s3:CreateBucket`, the rule will only allow the creation of a bucket named `bucket1`.
- 10 For additional optional rule parameters specifying the conditions in which the rule will be applied, click **Optional Conditions**.
- 11 In the **Source IPs** field, click the **Add (+)** button. Enter the IP addresses or subnets to which the rule will be applied. Multiple entries can be entered with a comma-separated list. Each entry must be in standard CIDR format.
- 12 In the **S3 Prefixes** field, click the **Add (+)** button. Enter the object prefixes, including a trailing delimiter character and optional wildcard, on which access and results will be restricted for the `s3:ListBucket` and `s3:ListBucketVersions` actions. Multiple entries can be entered with a comma-separated list.
- 13 In the **S3 Delimiters** field, click the **Add (+)** button. Enter the prefix delimiter characters (such as a trailing slash) to be used to group result objects for the `s3:ListBucket` and `s3:ListBucketVersions` actions. Multiple entries can be entered with a comma-separated list.
- 14 Click **Add**.

If you wish to add additional rules to the policy, in the **Rules** panel, click the **Add (+)** button and then repeat steps 3 through 14 above.

Editing Object Store Access Policy Rules

To edit an object store access policy rule, perform the following:

- 1 From the **Policies > Object Store** tab, click the policy whose rules you wish to edit. The policy's detail page appears.
- 2 In the **Rules** panel, click **More Options > Edit** on the row of the rule you wish to edit. The **Edit Rule** pop-up window appears.
- 3 In the **Effect** field, select to **Allow** or **Deny** S3 requests that match the actions and parameters set with this rule. Note that if a policy contains both an Allow rule and a Deny rule, the Deny rule will take precedent over the Allow rule (this is also known as an explicit Deny rule).
- 4 In the **Actions** field, click the **Add (+)** button. The **Add Actions** pop-up window appears.
- 5 From the **Available Actions** panel, select the actions that you wish the rule to allow. All selected actions appear in the **Selected Actions** panel.

If you wish to select all actions, select the checkbox directly beneath the **Available Actions** heading.

- 6 Click **Add**.
- 7 (Optional) Select the **Ignore Action Restriction Enforcement** checkbox if you wish the rule to allow combinations of actions, resources, and conditions which would be otherwise prevented.
- 8 In the **Resources** field, click the **Add (+)** button. Enter the object store resources (bucket component optionally followed by an object component) to which the selected action(s) will apply. The bucket and object components must be specified in the format `bucket_name/object_name`; for example, `bucket1/obj1`. For example, if you enter `bucket1` as the resource and selected the action of `s3:CreateBucket`, the rule will only allow the creation of a bucket named `bucket1`. To remove a resource, click the **X** next the resource you wish to remove.
- 9 In the **Source IPs** field, click the **Add (+)** button. Enter the IP addresses or subnets to which the rule will be applied. Multiple entries can be entered with a comma-separated list. Each entry must be in standard CIDR format. To remove a source IP address or subnet, click the **X** next the source IP address or subnet you wish to remove.
- 10 In the **S3 Prefixes** field, click the **Add (+)** button. Enter the object prefixes, including a trailing delimiter character and optional wildcard, on which access and results will be restricted for the `s3:ListBucket` and `s3:ListBucketVersions` actions. Multiple entries can be entered with a comma-separated list. To remove a prefix, click the **X** next the prefix you wish to remove.
- 11 In the **S3 Delimiters** field, click the **Add (+)** button. Enter the prefix delimiter characters (such as a trailing slash) to be used to group result objects for the `s3:ListBucket` and `s3:ListBucketVersions` actions. Multiple entries can be entered with a comma-separated list. To remove a delimiter, click the **X** next the delimiter you wish to remove.
- 12 Click **Save**.

Copying Object Store Access Policy Rules to Other Policies

Rules created for one object store access policy can be copied to other object store access policies. You can copy the rule as-is to be applied to the other policy, or modify the rule as needed for the other policy.

To copy a policy rule to another policy, perform the following:

- 1 From the **Policies > Object Store** tab, click the policy whose rule you wish to copy to another policy. The policy's detail page appears.
- 2 In the **Rules** panel, click **More Options > Copy** on the row of the rule you wish to edit. The **Copy Rule to Selected Policy** pop-up window appears.
- 3 Select the policy to which the rule will be copied from the **Policy** list.
- 4 In the **Effect** field, select to **Allow** or **Deny** S3 requests that match the actions and parameters set with this rule. Note that if a policy contains both an Allow rule and a Deny rule, the Deny rule will take precedent over the Allow rule (this is also known as an explicit Deny rule).
- 5 (Optional) In the **Name** field, enter a name for the copied rule.
- 6 (Optional) Add additional actions to the rule. In the **Actions** field, click the **Add (+)** button. The **Add Actions** pop-up window appears.
- 7 From the **Available Actions** panel, select the actions that you wish the rule to allow. All selected actions appear in the **Selected Actions** panel.

If you wish to select all actions, select the checkbox directly beneath the **Available Actions** heading.
- 8 Click **Add**.
- 9 (Optional) Select the **Ignore Action Restriction Enforcement** checkbox if you wish the rule to allow combinations of actions, resources, and conditions which would be otherwise prevented.
- 10 (Optional) Remove and/or add additional resources. In the **Resources** field, click the **Add (+)** button. Enter the object store resources (bucket component optionally followed by an object component) to which the selected action(s) will apply. For example, if you enter `bucket1` as the resource and selected the action of `s3:CreateBucket`, the rule will only allow the creation of a bucket named `bucket1`. To remove a resource, click the **X** next the resource you wish to remove.
- 11 (Optional) Remove and/or add additional source IPs. In the **Source IPs** field, click the **Add (+)** button. Enter the IP addresses or subnets to which the rule will be applied. Multiple entries can be entered with a comma-separated list. Each entry must be in standard CIDR format. To remove a source IP address or subnet, click the **X** next the source IP address or subnet you wish to remove.
- 12 (Optional) Remove and/or add additional S3 prefixes. In the **S3 Prefixes** field, click the **Add (+)** button. Enter the object prefixes, including a trailing delimiter character and optional wildcard, on which access and results will be restricted for the `s3:ListBucket` and `s3:ListBucketVersions`

actions. Multiple entries can be entered with a comma-separated list. To remove a prefix, click the **X** next the prefix you wish to remove.

- 13** (Optional) Remove and/or add additional S3 delimiters. In the **S3 Delimiters** field, click the **Add (+)** button. Enter the prefix delimiter characters (such as a trailing slash) to be used to group result objects for the `s3:ListBucket` and `s3:ListBucketVersions` actions. Multiple entries can be entered with a comma-separated list. To remove a delimiter, click the **X** next the delimiter you wish to remove.
- 14** Click **Save**.

Removing Object Store Access Policy Rules

- 1** To remove a single rule:
 - a** From the **Policies > Object Store** tab, click the policy from which you wish to remove a rule. The policy's detail page appears.
 - b** In the **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** In the **Rules** panel, click **More Options > Remove** in the the heading of the **Rules** panel. The **Remove Object Store Access 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.
 - b** 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.
 - c** Click **Remove**.

Adding Users to an Object Store Access Policy

To add a user to an object store access policy:

- 1** From the **Policies > Object Store** tab, click the policy from which you wish to add a user. The policy's detail page appears.
- 2** Click **More Options > Add Users** in the header of the **Users** panel. The **Add Users** pop-up window appears.

- 3 From the **Available Users** panel, select one or more users from the list. The selected user(s) appear in the **Selected Users** panel. To select all users, select the checkbox directly below the **Available Users** heading.
- 4 Click **Add**.

Removing Users from an Object Store Access Policy

To remove a user from an object store access policy:

- 1 From the **Policies > Object Store** tab, click the policy from which you wish to remove a user. The policy's detail page appears.
- 2 Click **More Options > Remove Users** in the header of the **Users** panel. The **Remove Users** pop-up window appears.
- 3 From the **Current Users** panel, select one or more users from the list. The selected user(s) appear in the **Selected Users** panel. To select all users, select the checkbox directly below the **Current Users** heading.
- 4 Click **Remove**.

Deleting Object Store Policies

 **Note:** Only object store policies that were created by you and not attached to any user may be deleted. Pre-configured object store policies (appended with `pure:policy`) cannot be deleted.

- 1 To delete a single object store policy:
 - a From the **Policies > Object Store** tab, locate the policy you wish to delete from the object store policy list.
 - b From the same row as the policy to be deleted, click the **Delete** button. Note that the button is greyed out and inaccessible for system policies. The **Delete Object Store Policy** dialog box appears.
 - c Click **Delete**.
- 2 To delete multiple object store policies:
 - a From the **Policies > Object Access Policies** tab, click **More Options > Delete**. 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 the policies you wish to delete. Only user-created policies are listed. Each selected policy appears in the **Selected Policies** panel.

For arrays with a large number of policies, you can search for a policy by entering a full or partial policy name in the search box to filter the policy list.

If you wish to delete all policies, from the **Available Policies** panel, select the checkbox above the list of policies to select all policies.


- c** Click **Delete**.

SMB Client and Share Policies

The **SMB** tab allows you to view, create, and manage SMB client and share policies.

SMB client policies manage the read and write privileges that an SMB client has to the file system. SMB share policies control SMB users and groups permissions to read, write, and delete files in the filesystem to which they are attached.

Figure 8.6 SMB Tab



The screenshot shows the 'Policies' section of the Pure Storage GUI. The 'SMB' tab is selected, and the 'SMB Client Policies' and 'SMB Share Policies' panels are visible. Each panel has a search box, a table of policies, and a '1 of 1' indicator.

Name	Type	Enabled
_smb_client_allow_everyone	smb-client	True

Name	Type	Enabled
_smb_share_allow_everyone	smb-share	True

The **SMB** tab displays the **SMB Client Policies** and **SMB Share Policies** panels.

The **SMB Client Policies** panel displays the following information:

- **Name:** The name of the SMB client policy.
- **Type:** The SMB client type, **smb-client**.

- **Enabled:** Whether the policy is enabled (**True**) or disabled (**False**).

The **SMB Share Policies** panel displays the following information:

- **Name:** The name of the SMB share policy.
- **Type:** The SMB policy type, **smb-share**.
- **Enabled:** Whether the policy is enabled (**True**) or disabled (**False**).

SMB Client Policies

SMB client policies allow you to restrict the read and write permission for a file system based on the SMB client's IP address. Permissions are configured in rules that are attached to the policies. SMB client policies work in conjunction with file ACLs to grant read and write permission.

An SMB client policy consists of two permission levels, which include **Read-Only** and **Read-Write**. These permissions are set when attaching rules to the policy.

Additionally, when creating a client policy rule, you must specify the client's IP address, subnets in CIDR notation, FQDNs, host names, or *. Wildcards are accepted in host names. Specifying star (*) indicates the SMB client is anonymous (everyone). Note that the following prioritization of rules is enforced: IP address > (host name | FQDN | wildcard | netmask) > star (*).

Purity//FB includes the built-in, immutable SMB client policy, **_smb_client_allow_everyone**. By default, this SMB client policy is attached to all file systems. This policy provides any SMB user read-write access to the file system.

SMB Encryption

SMB encryption is a feature of the SMB protocol and provides encryption of data transmitted between a client and server through the use of a negotiated encryption key. Encryption is done using Advanced Encryption Standard (AES) and Galois/Counter (GCM) or Counter with CBC-MAC (CCM) mode. FlashBlade supports all ciphers defined in the SMB specification.

FlashBlade supports SMB 3.1.1 and allows SMB encryption to be set globally or at the per-file share and client level.

To set global-level SMB encryption, contact Pure Technical Services.

With per-file share level SMB encryption, client policies that enforce encryption for specific clients are applied to specific file shares, allowing greater encryption precision. If encryption is required for a

specific file share/client, and the client does not use SMB 3.1.1, does not support SMB encryption, or cannot negotiate the encryption cipher, then all requests to that file share are declined. Setting SMB encryption for a file share/client is performed using the Purity//FB GUI, CLI, or REST API.


The following three strategies are provided for configuring per-file share level SMB encryption:

- **disabled** - Traffic is unencrypted. Communication targeting the file share, to which the SMB client policy is attached, cannot use any encryption.
- **optional** - Traffic is encrypted if the clients are configured to request encryption. Clients that are not configured to request encryption, or do not have the functionality, can still access the data, but the traffic will not be encrypted.
- **required** - Traffic is encrypted. Unencrypted traffic to the file share, to which the SMB client policy is attached, is rejected.

Viewing SMB Client Policy Details

From the **Policies > SMB** tab, clicking a policy in the **SMB Client Policies** panel will open the details page for SMB client rules associated with that SMB client policy.

Figure 8.7 Viewing SMB Client Policy Details



Policies

NFS Object Store **SMB** Snapshot

SMB Client > _smb_client_allow_everyone

Client	Permission	Encryption	Index
*	rw	optional	1

Details

Name: _smb_client_allow_everyone
Type: smb-client
Enabled: True

Attached File Systems

No file systems found.

Information about each of the policy's rules is displayed in the **SMB Client Rules** panel. The panel also allows you to add, remove, edit, move, and set encryption for the SMB client rules. The following columns are displayed:

- **Client** - The SMB client rule name.

- **Permission** - The SMB client rule's read and write permission, either **ro** (read-only) or **rw** (read-write).
- **Encryption** - The SMB client rule's configured encryption.
 - **disabled** - Traffic is unencrypted.
 - **optional** - Traffic is encrypted if the clients possess the necessary capability to do so. Clients lacking encryption capability can access data, but in unencrypted form.
 - **required** - Traffic is encrypted. Any unencrypted traffic is rejected.
- **Index** - The order in which each group's rules will be executed.

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

The **Attached File Systems** panel displays all file systems governed by the policy. It also allows you to add, move, and remove file systems from the policy.

Creating SMB Client Policies and Rules

To create an SMB client policy and its rule, perform the following:

- 1 From the **Policies > SMB** tab, click the **Add (+)** button in the heading of the **SMB Client Policies** panel. The **Step 1: Create SMB Client 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 SMB Client Rule** pop-up window appears. You can continue to configure the rule or click **Skip** to configure the rule at later time.
- 5 In the **Client** field, enter the SMB host name, FQDN, IP address/subnet, or star (*) to which the rule will apply. Wild cards are allowed for the host name and FQDNs. To input the subnet, use CIDR notation. Entering a star (*) indicates the SMB client is anonymous.
- 6 Set rule permission as **Read-Only** or **Read-Write**.
- 7 Set the rule encryption as follows.
 - **disabled** - Traffic is unencrypted.
 - **optional** - Traffic is encrypted if the clients possess the necessary capability to do so. Clients lacking encryption capability can access data, but in unencrypted form.

- **required** - Traffic is encrypted.

8 Click **Add**.

Once created, the policy will appear in the **SMB Client Policies** panel and rule will appear in the **SMB Client Rules** panel of the **SMB Client Policy** detail page.

Attaching File Systems to an SMB Client Policy

To attach file systems to an SMB client policy, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy to which you wish to attach file systems. The detail page for that policy appears.
- 2 Click **More Options > Add File Systems** in the heading of the **Attached File Systems** panel. The **Add File Systems** 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.

- 3 From the **Available File Systems** panel, select one or more file systems you wish to attach to the policy. If you wish to attach 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 Systems** panel.
- 4 Click **Add**.

You can alternatively attach SMB client policies to file systems from the **Storage > File Systems** page.

- 1 From the **Storage > File Systems** page, select the file system to which you wish to attach an SMB client policy. The details page for the that file sytem appears.
- 2 In the **Attached Policies** panel, click **More Options > Set SMB Client Policy**. The **Set SMB Client Policy** pop-up window appears.
- 3 From the **Selected Policy** list, select the SMB client policy that you wish to attach.
- 4 Click **Save**.

Moving File Systems to Another SMB Client Policy

One or more file systems can be moved from one SMB client policy to another.

To move file systems from one policy to another, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy from which you wish to move file systems. The detail page for that policy appears.
- 2 Click **More Options > Move File Systems** in the heading of the **Attached File Systems** panel. The **Move File Systems to Another Policy** 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.

- 3 From the **Current File Systems** panel, select one or more file systems you wish to move to a different policy. If you wish to move 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 Systems** panel.
- 4 Select the policy to which the file systems will be moved from the **Target Policy** list.
- 5 Click **Move**.

Removing File Systems from an SMB Client Policy

 **Note:** When removing a file system from an SMB client policy, some clients may lose access to the removed file system. If not client policy is attached to the file system, all clients are denied.

- 1 To remove a single file system from an SMB client policy:
 - a From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy from which you wish to remove a file system. The detail page for that policy appears.
 - b In the **Attached File Systems** panel, click the remove (X) button on the same row as the file system you wish to remove. The **Remove File System from Policy** pop-up window appears.
 - c Click **Remove**.
- 2 To remove multiple file systems from an SMB client policy:
 - a From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy from which you wish to remove a file system. The detail page for that policy appears.
 - b Click **More Options > Remove File Systems** in the heading of the **Attached File Systems** panel. The **Remove File Systems** 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.

- c From the **Current File Systems** panel, select one or more file systems 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 Systems** panel.
- d Click **Remove**.

You can alternatively remove SMB client policies from file systems from the **Storage > File Systems** page.

- 1 From the **Storage > File Systems** page, select the file system from which you wish to remove an SMB client policy. The details page for the that file sytem appears.
- 2 To remove a single SMB client policy, in the **Attached Policies** panel, click the remove (X) button in the same row as the policy you wish to remove.
- 3 To remove multiple SMB client policies:

 **Note:** By default, the SMB client policies 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 In the **Attached Policies** panel, click **More Options > Remove**. The **Remove Policies** pop-up window appears.
- b Select the policies you wish to remove from the **Available Policies** panel. Each selected policy willll appear in the **Selected Policies** panel.
- c Click **Remove**.

Adding Additional SMB Client Rules to SMB Client Policies

To add an SMB client rule to an existing SMB client policy, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client 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 **SMB Client Rules** panel. The **Add SMB Client Rule** pop-up window appears.
- 3 In the **Client** field, enter the SMB host name, FQDN, IP address/subnet, or star (*) to which the rule will apply. To input the subnet, use CIDR notation. Wildcards are allowed in host names and FQDNs. Entering star (*) indicates the SMB client is anonymous.
- 4 Set rule permission as **Read-Only** or **Read-Write**.
- 5 Click **Add**.

Setting SMB Client Rules Encryption

Encryption for each SMB client rule is set during its creation. When your encryption requirements change, rather than resetting the encryption for each SMB client rule one at time, you can set some or all rules at once.

To bulk-set the encryption for multiple rules, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy whose rules' encryption you wish to modify. The detail page for that policy appears.
- 2 In header of the **SMB Client Rules** panel, click **More Options > Encryption**. The **Set Encryption for SMB Client Rules** pop-up window appears.
- 3 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.
- 4 Select one of the following encryption type for the selected rules.
 - **disabled**- Traffic is unencrypted.
 - **optional** - Traffic is encrypted if the clients possess the necessary capability to do so. Clients lacking encryption capability can access data, but in unencrypted form.
 - **required** - Traffic is encrypted.
- 5 Click **Set Encryption**.

Editing SMB Client Rules

To edit SMB client rules, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy whose rules you wish to modify. The detail page for that policy appears.
- 2 In the **SMB Client Rules** panel, click **More Options > Edit** in the row of the rule you wish to modify. The **Edit SMB Client Rule** pop-up window appears.
- 3 In the **Client** field, enter the SMB host name, FQDN, IP address/subnet, or star (*) to which the rule will apply. To input the subnet, use CIDR notation. Wildcards are allowed in host names and FQDNs. Entering star (*) indicates the SMB client is anonymous.
- 4 Set rule permission as **Read-Only** or **Read-Write**.
- 5 Set the rule encryption as follows.
 - **disabled**- Traffic is unencrypted.

- **optional** - Traffic is encrypted if the clients possess the necessary capability to do so. Clients lacking encryption capability can access data, but in unencrypted form.
- **required** - Traffic is encrypted.

6 Click **Save**.

Reordering SMB Client Rules

The order that SMB client rules are applied is displayed in the **Index** column of the **SMB Client Rules** panel.

Note that the following prioritization of SMB client rules is enforced: IP address > (host name | FQDN | wildcard | netmask) > star (*).

To move an SMB client rule, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy whose rules you wish to reorder. The detail page for that policy appears.
- 2 In the **SMB Client Rules** panel, click **More Options > Move** at the end of the row in which the rule you wish to move appears. The **Move SMB Client Rule** pop-up window appears.
- 3 In the **Before** list, select the new order of the rule.
- 4 Click **Move**.

Once moved, the **SMB Client Rules** panel is updated reflecting the new rule order.

Removing SMB Client Rules

You can remove SMB client rules from SMB client policies.

 **Note:** If all SMB client rules are removed from an SMB client policy, the file systems to which the client policy is attached will deny all clients.

- 1 To remove a single SMB client rule:
 - a From the **Policies > SMB > SMB Client Policies** panel, click the policy from which you wish to remove a rule. The detail page for that policy appears.
 - b In the **SMB Share Rules** panel, click **More Options > Remove** on the row of the rule you wish to remove. The **Remove SMB Client Rule** pop-up window appears.
 - c Click **Remove**.
- 2 To remove multiple rules from an SMB client policy:

- a From the **Policies > SMB > SMB Client 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 **SMB Client Rules** panel. The **Remove SMB Client Rules** pop-up window appears.

 **Note:** By default, the rules 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.

- 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 SMB Client Policies

SMB client policies can be renamed at any time after their creation.

To rename an SMB client policy, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client 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 SMB Client 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 SMB Client Policies

Unless explicitly set to disabled, SMB client policies are enabled by default when they are created.

SMB client policies can be disabled or enabled at any time after their creation.

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

To enable or disable an SMB client policy and its rules, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client 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 SMB Client Policies

1 To delete a single SMB client policy:

- a** From the **Policies** > **SMB** > **SMB Client Policies** panel, click delete button at the end of the row in which the policy you wish to delete appears. The **Delete SMB Client Policy** pop-up window appears.
- b** Click **Delete**.

2 To delete multiple SMB client policies:

- a** From the **Policies** > **SMB** > **SMB Client Policies** panel, click **More Options** > **Delete** in the heading of the **SMB Client Policies** panel. The **Delete Policies** pop-up window appears.

 **Note:** By default, the policies 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.

- 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**.

SMB Share Policies

SMB share policies control SMB users and groups permissions to read, write, and delete files in the filesystem to which they are attached. Permissions are configured in rules that are attached to the policies. Only SMB users with full control permission are allowed to change a file's ACL. SMB share policies apply to users and groups who use SMB to access the file system; they do not apply to users and groups who use NFS to access the file system.

 **Note:** When an SMB share policy denies a user, the user will be denied. However, if the SMB share policy allows a user, access depends on the file ACL.

An SMB share policy consists of three permission levels, which include **Read**, **Change**, and **Full Control**. Only **Full Control** permissions allow users to modify file ACLs. These permissions are set when attaching rules to the policy. A file system may only have one policy of each type (SMB share policy and NFS export policy) attached. Backward compatibility is maintained so existing customers can still use their shares.

Purity//FB includes the built-in, immutable SMB share policy, **_smb_share_allow_everyone**. By default, this SMB share policy is attached to all file systems. This policy provides any SMB user full access to the file system. This policy can also be detached, thereby removing all access to the file system.

A principal name SMB user or group is called **everyone** and can be used to grant file system access to all SMB users and groups with an AD/LDAP configuration. Without creating an AD/LDAP configuration, you are not allowed to create a rule with any principal besides **everyone**.


SMB share policy rules can be created with the following permissions:

- **Read:** Permission to read files only. Set to **allow** or **deny**. If **Full Control** or **Change** are set to a specific value, **Read** must also be set to the same value. If **Full Control** or **Change** are unset, **Read** resolves to **allow** or **deny**.
- **Change:** Permission to read and write to files, add new files, and delete files. **Change** cannot set file ACLs. Set to **allow** or **deny**. If **Change** is set to a specific value, **Read** must also be set to the same value. If **Change** is unset, then **Full Control** must also be unset.
- **Full Control:** Permission to read and change files, and set file ACLs. Set to **allow** or **deny**, or optionally leave unset. If **Full Control** is set to a specific value, **Read** and **Change** must also be set to the same value.

Viewing SMB Share Policy Details

From the **Policies > SMB** tab, clicking a policy in the **SMB Share Policies** panel will open the details page for rules associated with that policy.

Figure 8.8 Viewing SMB Share Policy Details



The screenshot displays the 'Policies' section of the Pure Storage GUI, specifically the 'SMB' tab. The breadcrumb navigation shows 'SMB Share' > '_smb_share_allow_everyone'. The main content area is divided into three panels:

- SMB Share Rules:** A table showing the permissions for the policy. It has columns for Principal, Full Control, Change, Read, and Index. The first rule is for 'everyone' with 'allow' permissions for Full Control, Change, and Read, and an Index of 1.
- Details:** A panel showing the policy's name ('_smb_share_allow_everyone'), type ('smb-share'), and enabled status ('True').
- Attached File Systems:** A list of file systems governed by the policy. It shows 10 file systems (IDs: 0020b8ff79ee, 002382983b2b, 002a5e1ea8e4, 003ed4b4c7c0, 003fc4c98a38, 004bebe4b446, 004ccb20eeb2, 005c49f955b6, 005d373ac093, 0062e883740f) with a delete icon (X) next to each.

Information about each of the policy's rules is displayed in the **SMB Share Rules** panel. The panel displays the settings for the rule's **Full Control**, **Change**, and **Read** permissions. The **Index** column displays the order in which each group's rules will be executed.

The **Details** panel displays the policy name and if it is currently enabled.

The **Attached File Systems** panel displays all file systems governed by the policy. It also allows you to add and remove file systems from the policy.

Creating SMB Share Policies and Rules

To create an SMB share policy and its rules, perform the following:

- 1 From the **Policies > SMB** tab, click the **Add (+)** button in the heading of the **SMB Share Policies** panel. The **Step 1: Create SMB Share 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 SMB Share Rule** pop-up window appears. You can continue to configure the rule or click **Skip** to configure the rule at later time.
- 5 In the **Principal** field, enter user or group to which the rule will apply.
- 6 Set rule permissions.
 - **Read**: Permission to read files only. Set to **Allow** or **Deny**. If **Full Control** or **Change** are set to a specific value, **Read** must also be set to the same value.
 - **Change**: Permission to read and write to files, add new files, and delete files. Set to **Allow** or **Deny**. If **Change** is set to a specific value, **Read** must also be set to the same value. If **Change** is unset, then **Full Control** must also be unset.
 - **Full Control**: Permission to read and change files, and change permissions for files. Set to **Allow** or **Deny**, or optionally leave unset. If **Full Control** is set to a specific value, **Read** and **Change** must also be set to the same value.
- 7 Click **Add**.

Once created, the policy and rule will appear in the **SMB Share Policies** and **SMB Share Rules** panels.

Attaching File Systems to an SMB Share Policy

To attach file systems to an SMB share policy, perform the following:

- 1 From the **Policies > SMB > SMB Share Policies** panel, click the SMB share policy to which you wish to attach file systems. The detail page for that policy appears.
- 2 Click **More Options > Add File Systems** in the heading of the **Attached File Systems** panel. The **Add File Systems** 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.

- 3 From the **Available File Systems** panel, select one or more file systems you wish to attach to the policy. If you wish to attach 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 Systems** panel.
- 4 Click **Add**.

You can alternatively attach SMB share policies to file systems from the **Storage > File Systems** page.

- 1 From the **Storage > File Systems** page, select the file system to which you wish to attach an NFS export policy. The details page for the that file sytem appears.
- 2 In the **Attached Policies** panel, click **More Options > Set SMB Share Policy**. The **Set SMB Share Policy** pop-up window appears.
- 3 From the **Selected Policy** list, select the SMB share policy that you wish to attach.
- 4 Click **Save**.

Moving File Systems to Another SMB Share Policy

One or more file systems can be moved from one SMB share policy to another.

To move file systems from one policy to another, perform the following:

- 1 From the **Policies > SMB > SMB Share Policies** panel, click the SMB share policy from which you wish to move file systems. The detail page for that policy appears.
- 2 Click **More Options > Move File Systems** in the heading of the **Attached File Systems** panel. The **Move File Systems to Another Policy** 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.

- 3 From the **Current File Systems** panel, select one or more file systems you wish to move to a different policy. If you wish to move 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 Systems** panel.
- 4 Select the policy to which the file systems will be moved from the **Target Policy** list.
- 5 Click **Move**.

Removing File Systems from an SMB Share Policy

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

- 1 To remove a single file system from an SMB share policy:
 - a From the **Policies > SMB > SMB Share Policies** panel, click the SMB share policy from which you wish to remove a file system. The detail page for that policy appears.

- b In the **Attached File Systems** panel, click the remove (X) button on the same row as the file system you wish to remove. The **Remove File System from Policy** pop-up window appears.
- c Click **Remove**.

2 To remove multiple file systems from an SMB share policy:

- a From the **Policies > SMB > SMB Share Policies** panel, click the SMB share policy from which you wish to remove a file system. The detail page for that policy appears.
- b Click **More Options > Remove File Systems** in the heading of the **Attached File Systems** panel. The **Remove File Systems** 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.

- c From the **Current File Systems** panel, select one or more file systems 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 Systems** panel.
- d Click **Remove**.

You can alternatively remove SMB share policies from file systems from the **Storage > File Systems** page.

- 1 From the **Storage > File Systems** page, select the file system from which you wish to remove an SMB share policy. The details page for that file system appears.
- 2 To remove a single SMB share policy, in the **Attached Policies** panel, click the remove (X) button in the same row as the policy you wish to remove.
- 3 To remove multiple SMB share policies:
 - a In the **Attached Policies** panel, click **More Options > Remove**. The **Remove Policies** pop-up window appears.
 - b Select the policies you wish to remove from the **Available Policies** panel. Each selected policy will appear in the **Selected Policies** panel.
 - c Click **Remove**.

Adding Additional SMB Share Rules to SMB Share Policies

To add an SMB share rule to an existing SMB share policy, perform the following:

- 1 From the **Policies > SMB > SMB Share Policies** panel, click the SMB share 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 **SMB Share Rules** panel. The **Add SMB Share Rule** pop-up window appears.
- 3 In the **Principal** field, enter user or group to which the rule will apply.
- 4 Set rule permissions.
 - **Read:** Permission to read files only. Set to **Allow** or **Deny**. If **Full Control** or **Change** are set to a specific value, **Read** must also be set to the same value.
 - **Change:** Permission to read and write to files, add new files, and delete files. Set to **Allow** or **Deny**. If **Change** is set to a specific value, **Read** must also be set to the same value. If **Change** is unset, then **Full Control** must also be unset.
 - **Full Control:** Permission to read and change files, and change permissions for files. Set to **Allow** or **Deny**, or optionally leave unset. If **Full Control** is set to a specific value, **Read** and **Change** must also be set to the same value.
- 5 Click **Add**.

Editing SMB Share Rules

To edit SMB share rules, perform the following:

- 1 From the **Policies > SMB > SMB Share Policies** panel, click the SMB share policy whose rules you wish to modify. The detail page for that policy appears.
- 2 In the **SMB Share Rules** panel, click **More Options > Edit** in the row of the rule you wish to modify. The **Edit SMB Share Rule** pop-up window appears.
- 3 If changing the principal, enter the new user or group in the Principal field.
- 4 Set rule permissions.
 - **Read:** Permission to read files only. Set to **Allow** or **Deny**. If **Full Control** or **Change** are set to a specific value, **Read** must also be set to the same value.
 - **Change:** Permission to read and write to files, add new files, and delete files. Set to **Allow** or **Deny**. If **Change** is set to a specific value, **Read** must also be set to the same value. If **Change** is unset, then **Full Control** must also be unset.
 - **Full Control:** Permission to read and change files, and change permissions for files. Set to **Allow** or **Deny**, or optionally leave unset. If **Full Control** is set to a specific value, **Read** and **Change** must also be set to the same value.
- 5 Click **Save**.

Removing SMB Share Rules

- 1 To remove a single SMB share rule:
 - a From the **Policies > SMB > SMB Share Policies** panel, click the policy from which you wish to remove a rule. The detail page for that policy appears.
 - b In the **SMB Share 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 SMB share policy:
 - a From the **Policies > SMB > SMB Share 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 **SMB Share Rules** panel. The **Remove SMB Share Rules** pop-up window appears.

 **Note:** By default, the rules 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.
 - 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 SMB Share Policies

SMB share policies can be renamed at any time after their creation.

To rename an SMB share policy, perform the following:

- 1 From the **Policies > SMB > SMB Share Policies** panel, click the SMB share 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 SMB Share 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 SMB Share Policies

Unless explicitly set to disabled, SMB share policies are enabled by default when they are created.

SMB share policies can be disabled or enabled at any time after their creation.

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

To enable or disable an SMB share policy and its rules, perform the following:

- 1 From the **Policies > SMB > SMB Share Policies** panel, click the SMB share 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 SMB Share Policies

- 1 To delete a single SMB share policy:
 - a From the **Policies > SMB > SMB Share Policies** panel, click delete button at the end of the row in which the policy you wish to delete appears. The **Delete SMB Share Policy** pop-up window appears.
 - b Click **Delete**.
- 2 To delete multiple SMB share policies:
 - a From the **Policies > SMB > SMB Share Policies** panel, click **More Options > Delete** in the heading of the **SMB Share Policies** panel. The **Delete Policies** pop-up window appears.


 **Note:** By default, the policies 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.

- 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**.

File System Snapshot Policies

A file system snapshot is a point-in-time copy of a file system. File system snapshots can created and managed manually or by using file system snapshot policies. A file system snapshot policy consists of one or more rules that govern when snapshots are created or replicated, and how long they are retained. When a policy is attached to a file system, it will create local file system snapshots. When one or more policies is attached to a file system replica link, it will replicate file system snapshots to a connected array. The **Snapshot** tab is used to manage the array's snapshot policies.

Figure 8.9 Snapshot Tab



Name	Type	Enabled	Retention Lock	Rules	# Retained Snapshots / Member
15minutes	snapshot	True	unlocked	Create or replicate 1 snapshot every 15 minutes for 1 hour.	~ 4
every-4-hours	snapshot	True	unlocked		
Every3Hours	snapshot	True	unlocked	Create or replicate 1 snapshot every 3 hours for 1 day.	~ 8
fscale-4hrs	snapshot	True	unlocked	Create or replicate 1 snapshot every 4 hours for 1 day.	~ 6
serenity:1d	snapshot	True	unlocked	Create or replicate 1 snapshot every day for 1 week.	~ 7
serenity:snap_every_hour	snapshot	True	unlocked	Create or replicate 1 snapshot every hour for 1 day.	~ 24

The **Snapshot** tab displays a list of all snapshot policies and the following rule information for each snapshot policy:


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

- **Retention Lock:** Displays whether a policy and its rules can be modified (**unlocked**) or are read only (**locked**). By default, policies are **locked** if SafeMode is enabled. If SafeMode is enabled and you wish to change the retention lock configuration, SafeMode must be disabled by Pure Technical Services.
- **Rules:** The snapshot creation and retention details of the policy rules.
- **# Retained Snapshots/Member:** The total number of snapshots retained per member per policy. If multiple rules exist for a policy, a summary is shown of the number of retained snapshots per member for each rule.

Viewing Snapshot Policy Details

To view details about each snapshot policy in the array, from the **Policies > Snapshot** tab, click a policy name to open a detail page about that snapshot policy.

Figure 8.10 Snapshot Policy Details



Policies

NFS Object Store SMB **Snapshot**

Snapshot > 15minutes

Rules	+				
Action	Every	At	Time Zone	Keep For	# Retained Snapshots / Member
Create or replicate 1 snapshot	15 minutes	-	-	1 hour	4

Details

Type: snapshot

Enabled: True Disable

Retention Lock: unlocked

Members (File Systems)

File Systems File System Replica Links Snapshots

No file system members found.

The members, rules, and policy status are displayed in the **Members**, **Rules**, and **Details** panels, respectively.

Creating Snapshot Policies

To create a policy, perform the following:

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

- 3 By default, the policy is set to enabled. If you wish to disable the policy, click the **Enabled** checkbox to disabled (gray).
- 4 Click **Create**. The **Step 2: Add Rule to Policy <name>** pop-up window appears.
- 5 In the **Create or replicate 1 snapshot every** field, enter the frequency at which snapshots are created or replicated. The value must be entered in the format n{mlhldlw}. For example, 15m, 3h, 2d, 1w, etc. The minimum value is 5m.
- 6 In the **At** field, enter the time of day the snapshot is created. The value must be entered in the format n:[m][am|pm], where n is a value of 0-23. If entering a minute (m) value, m must be a two digit number between 00 and 59; for example 12:00am or 12:15pm. If entered without am or pm, n must be a value of 0-23. For example, 23, 5pm, 2:25am, 14:45, etc.

This field is only editable if the value specified in the **Create or replicate 1 snapshot every** field is a multiple of a day. For example, 24h, 48h, 72h, 1d, 2d, 3d, 1w, etc
- 7 In the **And keep for** field, enter the retention period for the snapshot. The value must be entered in the format n{mlhldlw}. For example, 30m, 1h, 2d, 1w, etc. The retention period cannot be less than the snapshot interval.
- 8 Click **Add**.

Adding Snapshot Policy Rules

After creating a snapshot policy, add snapshot creation and retention rules.

To add a policy rule, perform the following:

- 1 From the **Policies > Snapshot** tab, click the policy to which you wish to add a rule. The policy's detail page appears.
- 2 In the **Rules** panel, click the **Add (+)** button. The **Add Rule to Policy <name>** pop-up window appears.
- 3 In the **Create or replicate 1 snapshot every** field, enter the frequency at which snapshots are created or replicated. The value must be entered in the format n{mlhldlw}. For example, 15m, 3h, 2d, 1w, etc.
- 4 In the **At** field, enter the time of day the snapshot is created. The value must be entered in the format n:[m][am|pm], where n is a value of 0-23. If entering a minute (m) value, m must be a two digit number between 00 and 59; for example 12:00am or 12:15pm. If entered without am or pm, n must be a value of 0-23. For example, 23, 5pm, 2:25am, 14:45, etc. This field is only editable if the value specified in the **Create or replicate 1 snapshot every** field is a multiple of a day. For example, 24h, 48h, 72h, 1d, 2d, 3d, 1w, etc.

 **Note:** This field only appears if during the creation of the policy a value equaling a multiple of a day was specified in the **Create or replicate 1 snapshot every** field.

- 5 In the **And keep for** field, enter the retention period for the snapshot. The value must be entered in the format n{mlhldlw}. For example, 30m, 1h, 2d, 1w, etc. The retention period cannot be less than the snapshot interval.
- 6 Click **Add**.

Adding Multiple Snapshot Policy Rules

If you added a snapshot creation and retention rule when you created the policy, you can add up to four additional rules to the policy to save and retain snapshots from the original rule.

Consider the following example where a second rule is added to a policy with an existing rule:

- Rule 1 (existing) - One snapshot is created each hour and retained for a day.
- Rule 2 (new) - One snapshot per day will be retained for a week.

Rule 1 results in 24 snapshots being created in a 24-hour period. After 24 hours, the first snapshot that was created is destroyed and another snapshot is created; this cycle is repeated for each snapshot resulting in 24 snapshots each day. Per Rule 2, since one snapshot from Rule 1 is saved each day and retained for a week, six additional snapshots are retained for an additional six days. After one week, a total of 30 snapshots will be retained. This snapshot count is shown in the **#Retained Snapshots/Member** column of the **Protection > Policies** page.

 **Note:** If you are adding an additional retention rule to an existing rule, the values allowed for frequency of snapshot retention and the snapshot retention period are dependent on the snapshot creation rules that were initially created for the policy. For example, if a snapshot is created every 15 minutes and kept for one day, any additional retention frequency rule must be in multiples of 15 minutes and any additional retention periods must be greater than one day.

To add an additional rule, perform the following:

- 1 From the **Policies > Snapshot** tab, click the policy to which you wish to add a rule. The policy's detail page appears.
- 2 In the **Rules** panel, click the **Add (+)** button. The **Add Rule to Policy <name>** pop-up window appears.
- 3 In the **Create or replicate 1 snapshot every** field, enter the frequency at which snapshots are created or replicated. The value must be entered in the format n{mlhldlw}. For example, 15m, 3h, 2d, 1w, etc.
- 4 In the **And keep for** field, enter the retention period for the snapshot. The value must be entered in the format n{mlhldlw}. For example, 30m, 1h, 2d, 1w, etc. The retention period cannot be less than the snapshot interval.

- 5 Click **Add**.

Editing Snapshot Policy Rules

 **Note:** Existing snapshots managed by a modified policy will be destroyed if they do not comply with the policy's new snapshot schedule rules.

- 1 From the **Policies > Snapshot** tab, click the policy whose rules you wish to edit. The policy's detail page appears.
- 2 In the **Rules** panel, click the **Edit** button on the row of the rule you wish to edit. The **Edit Rule for Policy <name>** pop-up window appears.
- 3 In the **Create or replicate 1 snapshot every** field, enter the frequency at which snapshots are created. The value must be entered in the format n{mlhldlw}, for example 15m.
- 4 In the **At** field, enter the time of day the snapshot is created. The value must be entered in the format n:[m][am|pm], where n is a value of 0-23. If entering a minute (m) value, m must be a two digit number between 0 and 59; for example 12:00am or 12:15pm. If entered without am or pm, n must be a value of 0-23. For example, 23, 5pm, 2:25am, 14:45, etc.

This field is only editable if the value specified in the **Create or replicate 1 snapshot every** field is a multiple of days. For example, 24h, 48h, 72h, 1d, 2d, 3d, 1w, etc.

- 5 In the **And keep for** field, enter the retention period for the snapshot. The value must be entered in the format n{mlhldlw}, for example 30m.
- 6 Review the summary and consequence of modifying the rule.
- 7 Select **I understand these consequences of modifying this rule**.
- 8 Click **Save**.

Removing Snapshot Policy Rules

 **Note:** When a policy rule is removed, all snapshots managed by that rule are destroyed. Those destroyed snapshots will be eradicated after 24 hours.

- 1 From the **Policies > Snapshot** page, click the policy whose rules you wish to remove. The policy's detail page appears.
- 2 In the **Rules** panel, click the **Remove** button on the row of the rule you wish to remove. The **Remove Rule** pop-up window appears.
- 3 Review the summary of consequences for removing the rule. Click to confirm.
- 4 Click **Remove**.

Adding File Systems to a Snapshot Policy

- 1 From the **Policies > Snapshot** tab, click the policy to which you wish to add file systems. The policy's detail page appears.
- 2 In the **Members** panel, click **File Systems**.
- 3 Click **More Options > Add File Systems**. The **Add File Systems** pop-up window appears.
- 4 From the **Available File Systems** panel, select one or more file systems from the list. The selected file system(s) appears in the **Selected File Systems** panel.
- 5 Click **Add**.

Removing File Systems from a Snapshot Policy

 **Important!** After removing a file system from a policy, the policy will no longer create snapshots of that file system, but the policy will continue to destroy and eradicate the snapshots it created.

- 1 To remove a single file system:
 - a From the **Policies > Snapshot** tab, click the policy from which you wish to remove a file system. The policy's detail page appears.
 - b In the **Members** panel, click **File Systems**.
 - c Click the **Remove (X)** button in the row of the file system you wish to remove. The **Remove Member from Policy** dialog box appears.
 - d Click **Remove**.

- 2 To remove 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 From the **Policies > Snapshot** page, click the policy from which you wish to remove file systems. The policy's detail screen appears.
- b In the **Members** panel, click **More Options > Remove Members**. The **Remove Members** pop-up window appears.
- c From the **Current Members** panel, select one or more file systems from the list. The selected file system(s) appears in the **Selected File Systems** panel.
- d Click **Remove**.

Adding Replica Links to a Snapshot Policy

- 1 From the **Policies > Snapshot** tab, click the policy to which you wish to add file systems. The policy's detail page appears.
- 2 In the **Members** panel, click **File System Replica Links**.
- 3 Click **More Options > Add File System Replica Links**. The **Add File System Replica Links** pop-up window appears.
- 4 From the **Available Replica Links** panel, select one or more replica links from the list. The selected replica link(s) appears in the **Selected Replica Links** panel.
- 5 Click **Add**.

Removing Replica Links from a Snapshot Policy

- 1 To remove a single replica link from a policy:
 - a From the **Policies > Snapshot** tab, click the policy from which you wish to remove a replica link. The policy's detail page appears.
 - b In the **Members** panel, click **File System Replica Links**.
 - c Click the **Remove (X)** button in the row of the replica link you wish to remove. The **Remove File System Replica Link** dialog box appears.
 - d Click **Remove**.
- 2 To remove multiple replica links:
 - a From the **Policies > Snapshot** tab, click the policy from which you wish to remove replica links. The policy's detail screen appears.
 - b In the **Members** panel, click **File System Replica Links**.
 - c Click **More Options > Remove File System Replica Links**. The **Remove File System Replica Links** pop-up window appears.
 - d From the **Available Replica Links** panel, select one or more replica links from the list. The selected replica link(s) appears in the **Selected Replica Links** panel.
 - e Click **Remove**.

Enabling and Disabling Snapshot Policies

⚠ Important! After a policy is disabled, that policy will no longer create snapshots or eradicate expired snapshots.

- 1** From the **Policies > Snapshot** tab, click the policy you wish to enable or disable. The policy's detail screen appears.
- 2** In the **Details** panel, click **Enable** or **Disable**. The **Enable Policy** or **Disable Policy** pop-up window appears.
- 3** Click **Enable** or **Disable**.

Removing a Snapshot from a Policy

Snapshots can be removed from policies at any time. When a snapshot is removed from a policy, that snapshot is no longer managed by the policy.

- 1** To remove a single snapshot from a policy:
 - a** From the **Policies > Snapshot** tab, click the policy from which you wish to remove a snapshot. The policy's detail page appears.
 - b** In the **Members** panel, click **Snapshots**.
 - c** Click the **Remove (X)** button in the row of the snapshot you wish to remove. The **Remove Member from Policy** dialog box appears.
 - d** Click **Remove**.
- 2** To remove multiple snapshots from a policy:
 - a** From the **Policies > Snapshot** tab, click the policy from which you wish to remove snapshots. The policy's detail screen appears.
 - b** In the **Members** panel, click **Snapshots**.
 - c** Click **More Options > Remove Snapshots**. The **Remove Snapshots** pop-up window appears.
 - d** From the **Current Members** panel, select one or more snapshots from the list. The selected snapshot(s) appears in the **Selected Members** panel.
 - e** Click **Remove**.

Removing a Policy from a Snapshot

Policies can be removed from snapshots at any time.

 **Note:** If a policy is removed from a snapshot, that snapshot is no longer managed by that policy and must be destroyed manually.

You can remove a policy from a snapshot from either the **Storage** or **Protection** page.

The **Storage** page allows you to view snapshots per specific file system. The **Protection** page allows you to view all file system snapshots on the array.

To remove a policy from a snapshot, perform the following:

- 1 If you are on the **Storage** page, select **Storage > File Systems** and then click the file system from the table in the **File Systems** panel to access that file system's snapshots.
- 2 If you are on the **Protection** page, select **Protection > Snapshots**.
- 3 From the table in the **File System Snapshots** panel, locate the file system snapshot from which you wish to remove a policy.
- 4 Click **More Options > Remove Policy** at the end of that snapshot's row . The **Remove Policy from Snapshot** pop-up window appears.
- 5 If there are multiple policies attached to the file system snapshot, click **More Options > Remove Policies** at the end of that snapshot's row. The **Remove Policies from Snapshot** pop-up window appears.
- 6 From the **Policies** panel, select one or more snapshot policies from the list. The selected policy (or policies) appears in the **Selected Policies** panel.
- 7 Click **Remove**.

Deleting Snapshot Policies

- 1 To delete a single snapshot policy:
 - a From the **Policies > Snapshot** tab, locate the policy you wish to delete from the snapshot policy list.
 - b From the same row as the policy to be deleted, click the **Delete** button. The **Delete Policy** dialog box appears.
 - c Review the summary of consequences for removing the policy. Click to confirm.
 - d Click **Delete**.

2 To delete multiple snapshot policies:

- a** From the **Policies > Snapshot** tab, click **More Options > Delete**. The **Delete Policies** pop-up window appears.

 **Note:** By default, the policies 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.

- b** From the **Existing Policies** panel, select the policies you wish to delete. Each selected policy appears in the **Selected Policies** panel.

For arrays with a large number of policies, you can search for a policy by entering a full or partial policy name in the search box to filter the policy list.

If you wish to delete all policies, from the **Existing Policies** panel, select the checkbox above the list of policies to select all policies.

- c** Click **Delete**.


Chapter 9

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 displayed in the chart.

Figure 9.1 Chart Legend



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, its file systems, and buckets. By default, performance metrics are displayed 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.

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 and IOPS graphs display operations that the array is doing in chunks. Default.
- **Protocol Specific:** The Latency and IOPS 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

Performance 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 a two panes: The left pane provides a list of all file systems and the right pane displays **Latency**, **IOPs**, and **Bandwidth** performance charts. Select one or more file systems from the left pane to display their performance charts.

Figure 9.2 Displaying File System Performance (NFS)



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**. The following figure displays the selection of a single SMB file system and its **Latency** performance chart.

Figure 9.3 Displaying File System Performance (Native SMB Mode Enabled)



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 NFS Client, User, and Group Performance Statistics

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

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

To view statistics for clients, click **Reload Cache** in **Performance by Clients (NFS)** table header

To view statistics for groups and users:

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

Figure 9.4 Selecting a File System




Figure 9.5 Viewing NFS Groups Performance Statistics

Performance

Array File Systems Object Store Clients **Groups** Users

Performance by Groups (NFS) 1 of 1 Reload Cache Download CSV

Selected File System: exsdatastore Select

Showing cached response retrieved at 2021-06-01 15:26:02 sorted by B/s (read) and containing 1 rows (seconds ago).

Name	Group Id	Group Name	B/s (read)	B/s (write)	op/s (read)	op/s (write)	op/s (other)	us/op (read)	us/op (write)	us/op (other)	B/op (read)	B/op (write)	B/op (all)
exsdatastore/0	0	randomgroup	385.77 KB/s	100%	86.12	2.26	100%	2.510.94	626.02	-	4.586.77	4.586.77	
(total)	-	-	385.77 KB/s	-	86.12	2.26	-	2.510.94	626.02	-	4.586.77	4.586.77	

Figure 9.6 Viewing NFS Users Performance Statistics

Performance

Array File Systems Object Store Clients Groups **Users**

Performance by Users (NFS) 1 of 1 Reload Cache Download CSV

Selected File System: exsdatastore Select

Showing cached response retrieved at 2021-06-01 15:28:12 sorted by B/s (read) and containing 1 rows (seconds ago).

Name	User Id	User Name	B/s (read)	B/s (write)	op/s (read)	op/s (write)	op/s (other)	us/op (read)	us/op (write)	us/op (other)	B/op (read)	B/op (write)	B/op (all)
exsdatastore/0	0	administrator	589.02 KB/s	100%	109.7	6.79	100%	2.442.48	1.058.53	-	5.498.4	5.498.4	
(total)	-	-	589.02 KB/s	-	109.7	6.79	-	2.442.48	1.058.53	-	5.498.4	5.498.4	

The following information is displayed in the tables under each tab:

- **Name** - For client statistics only. The client name.
- **Group ID** - For groups statistics only. The group ID.
- **Group Name** - For groups 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 the FlashBlade GUI, you can search each column by entering a value in field below each column header, as well as sort by clicking the column header.

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

Note that NFS 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.


Capacity

By default, the **Capacity** graph displays information for all storage on the array. Click **File Systems** or **Object Store** to display information for only that storage type.

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

- **Empty Space:** Total storage space available on the array (Only shown if viewing the capacity graph for all storage).
- **Total Physical:** Total storage space used on the array, file systems, or buckets when viewing the capacity graph for all storage, file systems, or object store.

- **Capacity:** Total storage capacity of the array, file systems, or buckets when viewing the capacity graph for all storage, file systems, or object store.
- **Total 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). Data reduction is given for all storage on the array, all file systems on the array, and all buckets on 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) for all file systems on the array when viewing the File systems capacity graph, or for all buckets on the array when viewing the Object Store capacity graph.
- **Unique:** Total space that the written data occupies on the array's file systems or buckets (Only shown if viewing the File systems or Object Store capacity graph). The value displayed does not include Destroyed space.
- **Destroyed:** The sum of destroyed unique space of all the destroyed file systems or buckets of the account (Only shown if viewing the File systems or Object Store capacity graph).
- **Snapshots:** Amount of space that the written data occupies on the array's snapshots after reduction via data compression (Only shown if viewing the File systems capacity graph).



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.

Replication


The **Analysis > Replication** page provides replication **Bandwidth** and replication **Backlog** charts.

The replication **Bandwidth** chart displays historical replication bandwidth on the array. The replication **Bandwidth** chart displays the number of bytes of replication data transferred over the storage network per second between this array and its source arrays, target arrays, and external storage systems (such as S3 buckets), at certain points in time. Note that only connected arrays are displayed. If the replicaion link is removed, the historical replication performance data is also removed.

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.

The replication **Backlog** chart displays historical object replication backlog information. The replication **Backlog** chart includes two Y axes that display pending operations (left) and size in bytes (right). Like the replication **Bandwidth** chart, hovering over any part of the chart displays values for a specific point in time.

Figure 9.7 Bandwidth and Backlog Charts



The point-in-time pop-up in the replication **Bandwidth** chart displays the following metrics:

- **Transmitted:** The bandwidth of data transmitted to connected arrays and S3 targets.
- **Received:** The bandwidth of data received from connected arrays.
- **Total:** The total bandwidth of transmitted and received data.

The point-in-time pop-up in the replication **Backlog** chart displays the following metrics:

- **Size:** The total size of object data in bytes to be replicated. Note that this does not contain custom metadata. The size may be larger than the value displayed depending on custom metadata usage.
- **PUT Ops:** The total amount of object data to be replicated.
- **DELETE Ops:** The total number of delete operations to be replicated.
- **Other Ops:** The total number of other operations to be replicated.


Click **File Systems** or **Object Store** to display information for only that storage type.

By default, the **Bandwidth** and **Backlog** charts are displayed for the array. Click the **Connected Arrays** or **Connected S3 Targets** tabs to view array or target-specific bandwidth and backlog charts for one or more selected arrays or targets, respectively.

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 9.8 Viewing the Network Page




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

- **CRC Error/s:** 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 MacTransmitErrs.


- **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 9.9 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 9.10 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 9.11 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 9.12 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 9.13 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 9.14 Selecting a Time Range



Chapter 10

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.

 **Note:** For first-generation FlashBlade arrays: the array displayed will vary depending upon if your FlashBlade array consists of a single or a multi-chassis configuration. For multi-chassis configurations, in addition to front and rear views of the chassis, the display also includes the array's external fabric modules (XFM's).

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, 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. The graphics that are displayed reflect your FlashBlade type (first generation, FlashBlade//E, or FlashBlade//S) and configuration (single-chassis versus multi-chassis). The following figure shows the **Hardware** page displayed for a multi-chassis FlashBlade//S system.

Figure 10.1 Health - FlashBlade//S Hardware Page (Multi-Chassis)



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.

The first-generation FlashBlade graphics allow you to physically turn on/off identity LEDs on each blade and fabric module. Similarly, you can also turn on/off identity LEDs on each blade, DirectFlash Module, and fabric module from the FlashBlade//S graphics.

- A white circle indicates the LED is off. Click the white circle to turn on the LED.

- A blue circle indicates the LED is on. Click the blue circle to turn off the LED.

FlashBlade//S System Resiliency

A FlashBlade//S storage system provides a high-level of resiliency and cluster high-availability during unforeseen operational disruptions.

Both FlashBlade//S S200 and S500 models can tolerate the following failures:

- Any single blade (N+1 availability), plus any single DFM
- Any two DFMs (N+2 resiliency)
- Any two power supply units (PSU) (2+2 redundancy)


 **Note:** For FlashBlade//S to tolerate power failures, electrical power sources must be independent.

- A single Fabric I/O Module (FIOM)

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.

Figure 10.2 Health - Alerts Page



Flag	ID	Severity	Code	State	Created	Last Seen	Summary
	118	Warning	1113	Open	2019-10-23 09:36	2019-10-25 13:49	A file replication link requires attention
	129	Warning	1113	Open	2019-10-23 09:36	2019-10-25 13:49	A file replication link requires attention
	140	Warning	1113	Open	2019-10-23 09:36	2019-10-25 13:49	A file replication link requires attention
	144	Warning	1113	Open	2019-10-23 09:36	2019-10-25 13:49	A file replication link requires attention
	145	Warning	1113	Open	2019-10-23 09:36	2019-10-25 13:49	A file replication link requires attention
	155	Warning	1113	Open	2019-10-23 11:36	2019-10-25 13:49	A file replication link requires attention
	156	Warning	1113	Open	2019-10-23 11:40	2019-10-25 13:49	A file replication link requires attention

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.
- **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** 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 11

The Purity//FB GUI Settings Page

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

Figure 11.1 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 11.2 Array Details Panel

Array Details	
Name morpheus	
ID 10bbefef1-166f-4153-87e6-f894b54280ea	
Security Update -	
Time 2022-10-24 10:47:48 GMT-07:00 (PDT)	
Time Zone US/Pacific	
Login Banner -	
Idle Timeout 30 minutes	
Data at Rest Encryption Enabled	

Side Navigation Bar

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

Figure 11.3 Settings Page Navigation Sidebar



The navigation side bar is divided into three categories: **System**, **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.
- **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.

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 In the **Sender Domain** section of the **Alert Routing** panel, click the edit icon. The **Edit Alert Routing** pop-up window appears.
- 3 Type the sender domain name. The domain name must be set to your company's domain name. For example, `mydomain.com`.

- 4 (Optional) Type 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 Click **Save**.

File System User/Group Quota Notifications

The **File System User/Group Quota Notifications** panel displays whether email notifications regarding quota usage are enabled as well as contact information for notification recipients, if any.

 **Note:** End user quota notification is supported for LDAP only.

Direct email notifications can optionally be enabled to send emails to users and groups when they approach or exceed their file system quota. An email with a warning message is sent when a quota reaches 80% capacity and again at 90% capacity. When a quota is at 100% capacity, a critical email is sent. Emails are sent every 24-hours until the quota falls below the 80% capacity threshold.

A direct email notification can optionally be configured to include the administrator's contact information which the user and group can use to inquire about their quota and manageability options, such as requesting an increase of the file system quota. The FlashBlade administrator can include a free-form string. The contact information can include a name, phone number, or email. For example **John Doe, Storage Admin Team. jdoe@example.com (555-867-5309).**

To access the **File System User/Group Quota Notifications** panel, from the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.

Enabling File System User/Group Quota Notifications

If you wish for users and/or groups to receive notification when they approach or exceed their file system quota, configure (enable) quota notifications.

To enable quota notifications, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.
- 2 In the **File System User/Group Quota Notifications** panel, click the edit icon. The **Edit Quota Notifications** pop-up window appears.
- 3 Set the **Direct Email Notifications Enabled** toggle button to enable (blue) quota notifications.
- 4 In the **Contact Information** field, optionally enter the contact information of FlashBlade administrator that users and groups should reach out to if they have questions about their quota. This may be an email, a phone number, a name, or some combination thereof. The contact field cannot exceed 256 characters in length.
- 5 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 11.4 Viewing the Array Details Panel

Array Details	
Name morpheus	☒
ID 10bbefe1-166f-4153-87e6-f894b54280ea	
Security Update -	
Time 2022-10-24 10:47:48 GMT-07:00 (PDT)	
Time Zone US/Pacific	☒
Login Banner -	☒
Idle Timeout 30 minutes	☒
Data at Rest Encryption Enabled	

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.

- **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**.

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 11.1 System Variables

Variable	OID	Description
pureSystem	iso.3.6.1.4.1.40482.1	System information
pureArrayName	iso.3.6.1.4.1.40482.11.0	Array name
pureArrayId	iso.3.6.1.4.1.40482.12.0	Array ID
pureArrayVersion	iso.3.6.1.4.1.40482.13.0	Array Purity//FB software version
pureArrayType	iso.3.6.1.4.1.40482.14.0	Array platform model version
pureArrayCurrentTime	iso.3.6.1.4.1.40482.15.0	Array time based on local time zone

Table 11.2 Array Performance Variables

Variable	OID	Description
purePerformance	iso.3.6.1.4.1.40482.4	Array performance information
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)

Table 11.2 Array Performance Variables (continued)

Variable	OID	Description
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 11.3 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
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 11.4 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)

Table 11.4 Network DNS Variables (continued)

Variable	OID	Description
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 11.5 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
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 11.6 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 11.7 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 11.8 Network Interface Variables

Variable	OID	Description
pureNetworkInterface	iso.3.6.1.4.1.40482.6.5	Network interface 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 11.9 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

Table 11.9 Hardware Connector Variables (continued)

Variable		Description
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 11.10 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
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, `.1.3.6.1.4.1.40482.5.1.1.1.1. \'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 11.5 Viewing the **SNMP** Panel

SNMP	
No SNMP managers found.	
Syslog Server Connections	
1-1 of 1	
Name ▲	URI
NagiosLogServer	tcp://10.64.242.146:5544

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 11.6 Viewing the Syslog Server Connections Panel

SNMP		⋮
No SNMP managers found.		
Syslog Server Connections		1 of 1 ⋮
Name▲	URI	
NagiosLogServer	tcp://10.64.242.146:5544	✎ 🗑

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 performed 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 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 multiple 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 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 11.7 Viewing the NTP Servers Panel

NTP Servers 1-4 of 4 +	
Name	
0.pool.ntp.org	
1.pool.ntp.org	
2.pool.ntp.org	
3.pool.ntp.org	

Configuring the NTP Server

The array maintains time synchronization by interacting with the NTP server.

- From the **Settings** page navigation sidebar, click **NTP Servers** under **System**.
- 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, and the **Challenge Response Verification Key** panel, which allows you to view and update the FlashBlade's signed verification key

To access the **Pure 1 Support** and **Challenge Response Verification Keys** panel, from the **Settings** page navigation sidebar, click **Support** under **System**.

Figure 11.8 Viewing the Pure1 Support and Challenge Response Verification Key Panels <<update screenshot>>



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 Click the **Phone Home** toggle button to switch between enabled (blue) and disabled (gray) status.

- 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.

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 Click the **Remote Assist** toggle button to connect (blue) and disconnect (gray) an RA session. Opening an RA session gives Pure Storage Support direct and secure access to the array.
- 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 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** dialog box 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 11.9 Downloading Support Logs for Single-Chassis Systems

Download Support Logs

Download support logs for the selected components starting on 12/19/2021 at 10 AM PST.

Include 1 hour of logs after start time.

Components: ☐ XFM ☐ CH1

Prepare all logs

Close

Figure 11.10 Downloading Support Logs for Multi-Chassis Systems

Download Support Logs

Download support logs for the selected components starting on 04/05/2020 at 10 AM PDT.

Include 1 hour of logs after start time.

Components: ☐ XFM ☐ CH1 ☒ CH2 ☒ CH3 ☒ CH4 ☐ CH5 ☐ CH6 ☐ CH7 ☐ CH8 ☐ CH9 ☐ CH10

Prepare selected logs

Close

- 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 Verification Key** pop-up window appears.
- 3 Paste the new signed public key (verification key) in the pop-up window and click **Save**.

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**.

Figure 11.11 Viewing the Connectors Panel

Connectors 1-10 of 10					
Name▲	Connector Type	Port Count	Lane Speed	Transceiver Type	
CH1.FM1.ETH1	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM1.ETH2	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM1.ETH3	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM1.ETH4	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM1.ETH5	RJ-45	1	1 Gb/s	AUTO	☒
CH1.FM2.ETH1	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM2.ETH2	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM2.ETH3	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM2.ETH4	QSFP	1	10 Gb/s	AUTO	☒
CH1.FM2.ETH5	RJ-45	1	1 Gb/s	AUTO	☒

- **Name** - The connector name in the format CH_x.FM_x.ETH_x, where CH_x is the chassis number, FM_x is the fabric module number, and ETH_x is the ethernet port number.
- **Connector Type** - The connector type. Possible values include **QSFP**, **RJ-45**, and **SFP**.
- **Port Count** - The number of ports for the connector.
- **Lane Speed** - The speed (in Gb/s) at which the component is operating.
- **Transceiver Type** - The transceiver type.

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

Use the **Diagnostics** panel to perform network pings and traceroutes to determine the reachability of components and troubleshoot network issues.

To access the **Diagnostics** panel, from the **Settings** page navigation sidebar, click **Diagnostics** under **Network**.

Figure 11.12 Viewing the Diagnostics Panel



Network Diagnostics

Ping Traceroute

Components: CH1.FB1,CH1.FM1

Target: IP address or fully qualified domain name

Source: Subnet or interface name or IP address (optional)

Count: 1 (Ping)

Discover MTU: ☐ (Trace)

Performing Pings and Traceroutes

- 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 11.13 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 Settings

The **DNS Settings** panel displays and manages the DNS attributes for an array's administrative network.

To access the **DNS Settings** panel, from the **Settings** page navigation sidebar, click **DNS** under **Network**.

Figure 11.14 Viewing the DNS Settings Panel

DNS Settings	
Domain	totoro.net
DNS Server(s)	10.64.24.91

- **Domain** - The DNS domain represents the domain suffix to be appended by the array when performing DNS lookups. For example, `mydomain.com`.
- **DNS Server(s)** - The DNS servers manages the DNS domains that are configured for the array. Each DNS domain can include up to three static DNS server IP addresses. Purity//FB queries the DNS servers in the order in which the IP addresses are listed in this option.

Editing DNS Settings

To edit DNS settings, 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 Settings** pop-up window appears.
- 3 (Optional) In the **Domain** field, type the domain suffix to be appended by the array when doing DNS lookups. For example, `mydomain.com`.

- 4 (Optional) In the **DNS #** 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 DNS # field. Purity//FB queries the DNS servers in the order that the IP addresses are listed.
- 5 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 11.15 Viewing the Link Aggregation Groups Panel

Link Aggregation Groups						1 of 1 +
Name	Status	Ports	Port Speed	LAG Speed	MAC	
uplink	healthy	CH1[FM1/2].ETH1 CH1[FM1/2].ETH2 CH1[FM1/2].ETH3	40 Gb/s	240 Gb/s	24:a9:37:84:72:2a	

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**.

Object Store Virtual Hosts

The **Object Store Virtual Hosts** panel displays and allows management of virtual host-style addressing for object buckets.

To access the **Object Store Virtual Hosts** panel, from the **Settings** page navigation sidebar, click **Object Store Virtual Hosts** under **Network**.

Figure 11.16 Viewing the Object Store Virtual Hosts Panel

Object Store Virtual Hosts		1 of 1
Name		
s3.amazonaws.com		

Note the following requirements and restrictions for virtual hosted-style access:

- FlashBlade supports a virtual hosted-style by using a proxy and the fixed host name `s3.amazonaws.com`. This cannot be deleted.
- In order to use virtual host-style addressing, your DNS must be configured so that requests to your host name are routed to the data vip of the FlashBlade.
- Up to 10 host names can be configured, including the default `s3.amazonaws.com`.
- The virtual host cannot be an IP address.
- If your client setup uses host names, set the host names as the virtual hosts.
- If your client setup uses FQDNs, set the FQDNs as the virtual hosts.

- If you wish to use host names and FQDNs interchangeably, set both the host names and FQDNs as the virtual hosts.
- Host names can be numeric, however this is not recommended. If using numeric hosts, ensure that `https://bucket_name.domain/` does not appear like an IP address as this could result in routing to the wrong destination.
- Supersets and subsets of existing host names with the same root are not allowed. For example, if `buckethost.example.com` is registered, then `prefix.buckethost.example.com` or `example.com` are not allowed.
- Supported characters are ASCII letters a through z (case-insensitive), 0-9, and the hyphen (-) character.
- The host name character limit is 255. It is recommended to limit the host name length to 191 or fewer characters. Since bucket names are limited to 63 characters, host name lengths of 191 characters guarantees that any bucket name plus a dot (".") will be less than the 255 character limit.

Virtual Hosted-style and Path-style S3 Requests

Applications that support S3 must support at least one of S3's two request types: path-style or virtual hosted-style. All versions of FlashBlade support path-style requests and, beginning with Purity//FB 3.2.0, FlashBlade also supports virtual hosted-style requests.

To enable applications to use virtual hosted-style S3 requests, an extra configuration step is typically required in both the FlashBlade and DNS, as compared to applications that use path-style S3 requests. Virtual hosted-style addressing allows applications to request objects using a bucket's virtual host name for example `mybucket.b1.s3.amazonaws.com`. Because bucket names may contain dots (".") and because FlashBlade does not know what parent domain name or DNS suffix has been set up in your on-prem DNS for S3 buckets, in order for FlashBlade to correctly map a virtual hosted-style request from your application to the correct bucket (for example, a request for `mybucket.b1.yourcompany.com` could map to bucket `mybucket` or `mybucket.b1`), the bucket virtual host name suffix used by an application and configured in your DNS must also be configured on FlashBlade settings under **Virtual Hosts**. In addition, you may need to configure DNS so that any bucket's virtual hostname and the parent domain of the buckets' hostnames resolve to the IP address of your FlashBlade's data interface

Applications, tools, and developer libraries sometimes allow customers and developers to configure whether the application should use virtual hosted-style requests or path-style requests. Instead of configuring a virtual host parent domain in FlashBlade and DNS, you may prefer to configure your

application to use path-style addressing, which enables you to use the same hostname or IP address to access any bucket on FlashBlade

Virtual hosted-style access uses the format

`https://somebucket.flashblade-data.somecompany.com/someobject` in S3 requests to read or write an object named “someobject” in bucket “somebucket” on FlashBlade. Applications that access objects in `somebucket` or access the bucket configuration of `somebucket` in this way expect your DNS server to resolve `somebucket.flashblade-data.somecompany.com` to the IP address of your FlashBlade’s data VIP and (if using HTTPS) expect your FlashBlade’s certificate to include `somebucket.flashblade-data.somecompany.com` in a Subject Alternative Name (SAN) field. Additionally, applications that list the names of all buckets in an S3 account on FlashBlade require your DNS server to resolve `flashblade-data.somecompany.com` (without the bucket name) to your FlashBlade’s data VIP and your FlashBlade’s certificate to include `flashblade-data.somecompany.com` in another SAN field.

Enabling FlashBlade to support applications that use virtual hosted-style S3 access requires configuring the host name or domain name for FlashBlade’s data VIP that applications will use when making S3 requests. This can be configured in the **Object Store Virtual Hosts** panel. In the example above, the parent domain name for buckets on FlashBlade is `flashblade-data.somecompany.com`.

Path-style access uses the format,

`https://flashblade-data.somecompany.com/somebucket/someobject` in S3 requests. Applications that access objects or bucket configuration in this way expect your DNS server to resolve `flashblade-data.somecompany.com` to the IP address of your FlashBlade’s data VIP and expect your FlashBlade’s certificate to include `flashblade-data.somecompany.com` in a SAN field.

Supporting a variety of applications that use both types of S3 access can be done by creating a wildcard entry in DNS for `*.flashblade-data.somecompany.com` and another entry in DNS for `flashblade-data.somecompany.com` that both resolve to a data VIP on your FlashBlade, and by including both of these as SAN fields in your FlashBlade certificate. Note that bucket names with dots (“.”) in their name may not match wildcard certificates and may not work with DNS wildcard entries, so using HTTPS to access buckets with dots in their name using virtual hosted-style requests requires a separate DNS entry and separate FlashBlade certificate SAN entry for each bucket name.

Use of path-style access is required when no hostname is configured in DNS to resolve to FlashBlade’s data VIP and its IP address must be specified, e.g.

`https://IP_OF_FLASHBLADE_DATA_VIP/somebucket/someobject`. In order to use HTTPS with S3 path-style requests, the certificate configured on FlashBlade must specify `flashblade-data.somecompany.com` in the certificate’s SAN field. To handle situations where an

application does not support path-style addressing or does not support specifying an IP address for an S3 endpoint, and where a hostname for FlashBlade's data VIP cannot be configured in DNS, refer to this [KB article](#).

FlashBlade includes `s3.amazonaws.com` as a default entry on the **Object Store Virtual Hosts** panel. If you have an application that only supports Amazon S3 and does not support configuring a non-AWS hostname or IP address for an S3 server, then if it makes virtual-hosted style S3 requests to AWS region us-east-1 using `s3.amazonaws.com` as the parent domain name, and if the application uses HTTP rather than HTTPS, then configuring the local `/etc/hosts` file on your application's host to resolve `s3.amazonaws.com` to a data VIP on your FlashBlade will enable your application to work with FlashBlade Object Store.

 **Note:** S3 use of path-style-access will be deprecated in the future by Amazon S3. If you use the same applications on both FlashBlade Object Store and Amazon S3, you may wish to switch from path-style access to virtual hosted-style access.

Creating a Virtual Host

Before creating virtual host-style addressing for object buckets, note the following requirements and restrictions:

- In order to use virtual host-style addressing, your DNS must be configured so that requests to your host name are routed to the data vip of the FlashBlade.
- Up to 10 host names can be configured, including the default `s3.amazonaws.com`.
- The virtual host cannot be an IP address.
- If your client setup uses host names, set the host names as the virtual hosts.
- If your client setup uses FQDNs, set the FQDNs as the virtual hosts.
- If you wish to use host names and FQDNs interchangeably, set both the host names and FQDNs as the virtual hosts.
- Host names can be numeric, however this is not recommended. If using numeric hosts, ensure that `https://bucket_name.domain/` does not appear like an IP address as this could result in routing to the wrong destination.

- Supersets and subsets of existing host names with the same root are not allowed. For example, if `buckethost.example.com` is registered, then `prefix.buckethost.example.com` or `example.com` are not allowed.
- Supported characters are ASCII letters a through z (case-insensitive), 0-9, and the hyphen (-) character.
- The host name character limit is 255. It is recommended to limit the host name length to 191 or fewer characters. Since bucket names are limited to 63 characters, host name lengths of 191 characters guarantees that any bucket name plus a dot (".") will be less than the 255 character limit.

Your DNS must be configured so that requests to your virtual host are routed to the data vip of the FlashBlade.

For example, if `example.com` is set as the virtual host, then requests to `example.com` and `$(bucket).example.com` must be routed to the data vip. To achieve this, you can configure both `example.com` and `*example.com` in the DNS.

Note that your DNS settings are not visible to the FlashBlade. You can use **dig** (Linux) or **nslookup** (Windows) on the host where S3 clients run to ensure the DNS setting is correct and that sent requests are resolved to the data vip.

To create a virtual host, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Object Store Virtual Hosts** under **Network**.
- 2 Click the Add (+) button in the header of the **Object Virtual Hosts** panel. The **Create Virtual Hosts** pop-up window appears.
- 3 Enter the new host name or domain.
- 4 Click **Create**.

Deleting a Virtual Host

 **Important!** When deleting a virtual host, subsequent requests made to the virtual host will be treated as path style requests. Applications that rely on the deleted virtual host may lose access or may access unintended resources.

To delete a virtual host, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Object Store Virtual Hosts** under **Network**.
- 2 To delete a single virtual host:

- a In the **Object Virtual Hosts** panel, locate the virtual host you wish to delete and click the **Delete** icon in the same row. The **Delete Virtual Host** pop-up window appears.
 - b Click **Delete**.
- 3 To delete multiple virtual hosts:
- a Click the **More Options** button in the header of the **Object Virtual Hosts** panel and select **Delete**. The **Delete Virtual Hosts** pop-up window appears.
 - b From the **Existing Hosts** column, select the virtual hosts you wish to delete. The selected hosts will appear in the **Selected Hosts** column.
 - c 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.

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 11.17 Viewing the Subnets & Interfaces Panel

Subnets								
1-2 of 2 + ⋮								
Name▲	Enabled	Prefix	VLAN	Gateway	MTU	LAG	Services	
data	True	10.64.6.0/23	1006	10.64.6.1	1500	uplink	data	🔗 🗑
net1	True	10.64.240.0/20	0	10.64.240.1	1500		management, support	🔗 🗑

Interfaces								
1-7 of 7 + ⋮								
Name▲	Enabled	Subnet	Address	VLAN	Mask	Gateway	MTU	Service
datavip1	True	data	10.64.6.236	1006	255.255.254.0	10.64.6.1	1500	data
datavip2	True	data	10.64.6.237	1006	255.255.254.0	10.64.6.1	1500	data
datavip3	True	data	10.64.6.238	1006	255.255.254.0	10.64.6.1	1500	data
datavip4	True	data	10.64.6.239	1006	255.255.254.0	10.64.6.1	1500	data
fm1.admin0	True	net1	10.64.242.69	0	255.255.240.0	10.64.240.1	1500	support
fm2.admin0	True	net1	10.64.242.70	0	255.255.240.0	10.64.240.1	1500	support
vir0	True	net1	10.64.242.68	0	255.255.240.0	10.64.240.1	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 outside the given subnet.
- **MTU** - Maximum transmission unit (MTU) for the data vips, in bytes.
- **LAG** - Link aggregation group (LAG) for the data vips.
- **Services** - Service types for which the data VIP is configured. The supported service for data VIPs is **data** and the supported service for replication vips is **replication**.

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**).
- **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.
- **Service** - Service types for which the data VIP is configured. The supported service for data vips is **data** and the supported service for replication VIPs is **replication**.

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 multiple 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.

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** interface 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** or **replication**.

7 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 network 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** or **replication**.
- 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 interfaces 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.
- **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 11.18 Viewing the API Clients Panel

API Clients							1 of 1 +
Name▲	Enabled	Max Role	Issuer	Time To Live	Client ID	Key ID	
rorypostman	false	readonly	rorypostman	1 day	6b007c8b-0b7a-43e4-9660-ff3584249d2	ae5fcfeb-f851-4e8f-ad02-c3cd75be7d6a	⋮

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|ml|hd]. 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 11.19 Viewing the Audit Trail Panel

Audit Trail							120 of 1,000	<	>
ID	Time	User	Command	Subcommand	Arguments	Location	User Interface		
1014	2021-01-15 17:06	pureuser	pureobj	virtual-host create	s3-virt.dev.purestorage.com	10.231.213.16	GUI		
1013	2021-01-15 17:05	pureuser	pureobj	virtual-host delete	s3-virt.dev.purestorage.com	10.231.213.16	GUI		
1012	2021-01-15 16:09	pureuser	purelog	delete	test	10.231.213.16	GUI		
1011	2021-01-15 14:59	pureuser	purelog	create	--port 'CH5FM2ETH4,CH5FM2ETH4' test	10.231.213.16	GUI		
1010	2021-01-15 11:57	pureuser	pureosnp	delete	test	10.231.213.16	GUI		
1009	2021-01-15 11:57	pureuser	purelog	delete	test_connectant	10.231.213.16	GUI		
1008	2021-01-15 07:06	ir	purebucket	enable	--versioning bucket.4	10.231.213.148	GUI		
1007	2021-01-15 07:06	ir	purebucket	enable	--versioning bucket.3	10.231.213.148	GUI		
1006	2021-01-15 07:06	ir	purebucket	enable	--versioning bucket.2	10.231.213.148	GUI		
1005	2021-01-15 07:06	ir	purebucket	enable	--versioning bucket.1	10.231.213.148	GUI		
1004	2021-01-15 07:06	ir	purebucket	enable	--versioning bucket.0	10.231.213.148	GUI		
1003	2021-01-15 07:05	ir	purepolicy	rule add	--keep-for '3d' --every '10m' arggro-snap	10.231.213.148	GUI		
1002	2021-01-15 07:05	ir	purepolicy	rule remove	--keep-for '1d' --destroy-snapshots arggro-snap	10.231.213.148	GUI		
1001	2021-01-15 06:34	ir	pureobj	secure-user access-key create	--user 'tommy@cariboo'	10.231.213.148	GUI		

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.

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 certificates on the FlashBlade.

To access the **Array Certificates** panel, from the **Settings** page navigation sidebar, click **Certificates** under **Security**.

Figure 11.20 Viewing the Array Certificates Panel

Array Certificates						1 of 1
Name▲	Common Name	Organization	Organizational Unit	Valid From	Valid To	
global	Pure Storage	Pure Storage	Pure Storage	2020-10-04	2050-11-16	⋮
External Certificates						+
No certificates found.						
Certificate Groups						1 of 1 + ⋮
Name▲	Used By	Type				
_default_replication_certs						⋮

The following information is displayed:

- **Name:** The name of the array certificate. Clicking the array certificate name opens the **Certificate Details** panel.
- **Common Name:**
- **Organization:** IP
- **Organizational Unit:**
- **Valid From:** The date from which the certificate is valid.
- **Valid To:** The date the certificate expires.

Importing Array Certificates

To import array certificates, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Click the **More Options** button in the **Array Certificates** panel. The **Import Array Certificate** 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 Phrase** - (Optional) If the private key is encrypted with a passphrase, enter the passphrase.

4 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** panel appears.

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**.


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 11.21 Viewing the External Certificates Panel



Name	Common Name	Organization	Organizational Unit	Valid From	Valid To
global	Pure Storage	Pure Storage	Pure Storage	2020-10-04	2050-11-16

External Certificates

No certificates found.

Name	Used By	Type
_default_replication_certs		

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 11.22 Viewing the Certificate Groups Panel

Array Certificates						1 of 1
Name▲	Common Name	Organization	Organizational Unit	Valid From	Valid To	
global	Pure Storage	Pure Storage	Pure Storage	2020-10-04	2050-11-16	⋮
External Certificates						+
No certificates found.						
Certificate Groups						1 of 1 + ⋮
Name▲	Used By	Type				
_default_replication_certs						⋮

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**.


Viewing the SMB Mode

The **SMB Mode** panel displays the global SMB mode setting on the FlashBlade for Active Directory (AD) and directory services.

FlashBlade supports SMB Native mode.

To access the **SMB Mode** panel, from the **Settings** page navigation sidebar, click **Directory Service** under **Security**.

Figure 11.23 Viewing the SMB Mode Panel (SMB Native Mode)



- The **Mode** field displays the current global SMB mode setting.
- The **Description** field provides your options for configuring an AD account or directory service. With SMB native mode, you are advised to create an AD account. You can alternatively configure SMB directory service.


Active Directory Accounts

FlashBlade allows the management of a single Active Directory (AD) account, which is tied to a single AD domain. The **Active Directory Accounts** panel displays and manages AD accounts for the FlashBlade.

To access the **Active Directory Accounts** panel, from the **Settings** page navigation sidebar, click **Directory Service** under **Security**.

If the global SMB mode is set to SMB native, the **Active Directory** panel directs you to join an AD domain using Kerberos authentication for NFS and SMB as shown in the following figure.

Figure 11.24 Viewing the Active Directory Accounts Panel (SMB Native Mode)



The following information is displayed:

- **Name:** The name of the AD computer account. Clicking the name directs you to the **Details** panel about the account (see [Viewing Active Directory Account Details](#)).
- **Computer Name:** The name of the computer account in the AD domain that represents the FlashBlade.
- **Join OU:** The organizational unit within the AD domain in which the AD computer account was created.

- **Domain:** The AD domain in which the AD computer account for FlashBlade was created.

For an overview about AD, see [Active Directory Overview](#).


For instructions on creating an AD account, see [Creating an Active Directory Account](#).

Viewing Active Directory Account Details

The Active Directory (AD) Accounts **Details** panel provides a summary of the configuration details about the AD account.

To access the **Details** panel, from the **Active Directory Accounts** panel, click the AD account name.

Figure 11.25 Viewing the Active Directory Accounts Details Panel



Details			
Name	oban	Encryption Types	aes256-cts-hmac-sha1-96
Computer Name	oban-smb-1	Service Principal Names	HOST/oban-smb-1
Join OU	CN=Computers		HOST/oban-smb-1.totoro.net
Domain	totoro.net		nfs/oban-smb-1
			nfs/oban-smb-1.totoro.net
			HOST/oban-smb-2
			HOST/oban-smb-2.totoro.net
			nfs/oban-smb-2
			nfs/oban-smb-2.totoro.net
			HOST/oban-smb-3
			HOST/oban-smb-3.totoro.net
			nfs/oban-smb-3
			nfs/oban-smb-3.totoro.net
			HOST/oban-smb-4
			HOST/oban-smb-4.totoro.net
			nfs/oban-smb-4
			nfs/oban-smb-4.totoro.net
			HOST/oban-smb-5
			HOST/oban-smb-5.totoro.net
			nfs/oban-smb-5
			nfs/oban-smb-5.totoro.net

Managing AD accounts is described in detail under [Active Directory Overview](#).

Viewing Active Directory Account Keytabs

The Active Directory (AD) Accounts **Kerberos Keytabs** panel displays all Kerberos keytabs in use with the AD account.

To access the **Kerberos Keytabs** panel, from the **Active Directory Accounts** panel, click the AD account name.

Figure 11.26 Viewing Active Directory Account Kerberos Keytabs

Kerberos Keytabs

1-16 of 16

Name	Principal	Fully Qualified Domain Name	Realm	Encryption Type	KVNo	
oban.1	HOST	OBAN-SMB-2	TOTORO.NET	aes256-cts-hmac-sha1-96	5	
oban.2	nfs	OBAN-SMB-2	TOTORO.NET	aes256-cts-hmac-sha1-96	5	
oban.4	HOST	oban-smb-2.totoro.net	TOTORO.NET	aes256-cts-hmac-sha1-96	5	
oban.5	nfs	oban-smb-3.totoro.net	TOTORO.NET	aes256-cts-hmac-sha1-96	5	

By default, keytabs are displayed newest to oldest. The following information is displayed:

- **Name:** The name of the keytab.
- **Principal:** The Service Principal Name (SPN).
- **Fully Qualified Domain Name:** The fully qualified domain name that the client will mount the file system over for NFS. 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.


Managing Kerberos Keytabs for AD is described in detail under [Active Directory Overview](#).

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 11.27 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.

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)

For an overview about directory services, see [Directory Services Overview](#).

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, SMB, or NFS protocol integration with LDAP, specify the URIs (optional for SMB), base DN, and bind username and password of the directory service. Note that for file sharing over SMB, the base DN of the directory service is used to represent the domain that is joined, and the URIs can be configured for SMB to control what servers are communicated with when joining the domain. 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](#)
- [Configuring the SMB Directory Service](#)

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 NFSv4.1 and SMB clients, in addition to auth_sys, are supported. With Kerberos, 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 11.28 Viewing the Kerberos Keytabs Panel

Kerberos Keytabs							1-16 of 16 +
Name▲	Principal	Fully Qualified Domain Name	Realm	Encryption Type	KVNo	Source	
oban.1	HOST	OBAN-SMB-2	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.2	nfs	OBAN-SMB-2	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.4	HOST	oban-smb-2.totoro.net	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.5	nfs	oban-smb-2.totoro.net	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.6	HOST	oban-smb-3	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.7	nfs	oban-smb-3	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.8	HOST	oban-smb-3.totoro.net	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.9	nfs	oban-smb-3.totoro.net	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.10	HOST	oban-smb-4	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮
oban.11	nfs	oban-smb-4	TOTORO.NET	aes256-cts-hmac-sha1-96	5	oban	⋮

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**

Rapid Data Locking

 **CAUTION:** Enabling RDL is an irreversible operation. Once enabled, RDL cannot be disabled without erasing all data on the FlashBlade.

 **Note:** Rapid Data Locking (RDL) can only be enabled if there are no file systems or buckets created on the FlashBlade.

FlashBlade Rapid Data Locking (RDL) allows you to quickly lock your data remotely if local physical access to your FlashBlade is lost, thereby enhancing the data security of your array.

RDL works in conjunction with Enterprise Key Management (EKM) servers as an off-array option to manage access to the FlashBlade encryption key using Key Management Interoperability Protocol (KMIP) for communication. This server implementation, with regard to FlashBlade, is referred to as a *KMIP server*.

When RDL is enabled, data from the FlashBlade can only be read if the KMIP server registered with the FlashBlade provides the correct secret key during boot. If, for example, physical access to your FlashBlade has been compromised, you can revoke the FlashBlade's access on the KMIP server—this denies access to the decryption information on the KMIP server, thereby making it impossible to read the data on the FlashBlade. Once the threat has been addressed, re-enable the KMIP connection and then power cycle the FlashBlade to resume normal operation. Note that for security, rotation of external keys on the KMIP server should be regularly scheduled (see [Rotating Keys](#) for procedure).

 **Important!** Do not rotate external keys if you are performing an EKM (Enterprise Key Management) HA upgrade.

Communication between the FlashBlade and KMIP server requires that either CA certificates or CA certificate groups for the KMIP server be imported to the FlashBlade (see [External Certificates](#) and [Certificate Groups](#) and their related topics for information about how to import external certificates and groups). Likewise, you must also import FlashBlade's array certificate to the KMIP server (see [Array Certificates](#) and its related topics for information about how to export array certificates). These certificates must exist on the FlashBlade and KMIP server(s) in order for secure communication. This

certificate "exchange" must occur prior to configuring KMIP servers on the FlashBlade. If communication to the KMIP server is faulty, an alert will be issued.

 **Note:** The KMIP server used to enable RDL cannot be changed. For example, if you configure two KMIP servers, `kmip1` and `kmip2`, and enable RDL using server `kmip1`, RDL cannot be changed to use `kmip2`.

Viewing and Managing RDL

Rapid Data Locking (RDL) is managed from the **Rapid Data Locking** screen in the FlashBlade GUI. To access the **Rapid Data Locking** screen, from the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.

Figure 11.29 Viewing the Rapid Data Locking Screen

Rapid Data Locking			
Enabled	False	Server	-

KMIP Servers			
1 of 1 +			
Name	URI	CA Certificate	CA Certificate Group
kmip-c14-1	kmip-c14-1	kmip-c14-1	

The **Rapid Data Locking** screen provides two panels.

- The **Rapid Data Locking** panel displays the current RDL status. If disabled, the panel allows you to enable RDL. If enabled, the panel allows you to test RDL on the FlashBlade, as well as rotate KMIP server keys. The **Rapid Data Locking** panel also displays the following information:
 - **Enabled:** If RDL is currently enabled (**True**) or disabled (**False**).
 - **Server:** The configured KMIP server in use.
- The **KMIP Servers** panel allows you to configure and manage KMIP servers. It also displays the following for all configured KMIP servers:
 - **Name:** The KMIP server name.
 - **URI:** The KMIP server's URI.
 - **CA Certificate:** The CA certificate used to validate the KMIP server.
 - **CA Certificate Group:** The CA certificate group used to validate the KMIP server.

RDL and KMIP servers can also be viewed and managed using the FlashBlade CLI.

To view configured KMIP servers, issue the **purekmip list** command.

```
$ purekmip list
Name      URI      CA Certificate  CA Certificate Group
kmip-cl4-1 kmip-cl4-1 kmip-cl4-1-cert -
```

KMIP servers are configured and managed using the **purekmip** command with the **create**, **delete**, **setattr**, and **test** sub-commands. Their usage is described in subsequent topics. For complete details about the **purekmip** command, refer to the *FlashBlade CLI Reference*.

To view the status of RDL, issue the **puredatalock list** command. Note that if enabled (True), the associated KMIP server is also displayed.

```
$ puredatalock list
Name  Enabled
-     False
```

RDL is managed using the **puredatalock** command and **enable**, **rotate**, and **test** sub-commands. Their usage is described in subsequent topics. For complete details about the **puredatalock** command, refer to the *FlashBlade CLI Reference*.

Creating a KMIP Server Configuration

To create a KMIP server configuration, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.
- 2 Click the **Add (+)** button in the header of the **KMIP Servers** panel. The **Create KMIP Server Configuration** pop-up window appears.
- 3 Enter the KMIP server's name in the **KMIP Server Name** field.
- 4 Click the **Add (+)** button to the right of the **URIs** field. Enter one or more URIs for the KMIP server and click **OK**. Note that multiple URIs are recommended for high availability.
- 5 Perform one of the following: In the **CA Certificate** field select the CA certificate to validate the KMIP server, or in the **CA Certificate Group** field, select the CA certificate group to validate the KMIP server.
- 6 Click **Create**.

The **KMIP Servers** panel is updated with the newly created KMIP server configuration.

Creating a KMIP Server Configuration using the CLI

You can alternately create a KMIP server configuration using the FlashBlade CLI.

To create a KMIP server configuration, issue the **purekmip create** CLI command.

In addition to the KMIP server name, you must specify the following:

- **--uri** - One or more (in a comma separated list) server URIs. Note that multiple URIs are recommended for high availability.
- **--ca-certificate** OR **--ca-certificate-group** - The name of server's CA certificate, or name of the CA certificate group for all external certificates.

In the following example, the configuration for KMIP server `kmip1` is created:

```
$ purekmip create kmip1 --uri kmip1.company.com:5696,kmip2.company.com:5696 --ca-certificate kmip1-cert
```

For complete details about the **purekmip** command, refer to the *FlashBlade CLI Reference*.

Testing a KMIP Server Configuration

Once you have created a KMIP server configuration, test it to verify secure connectivity with the FlashBlade.

To test a KMIP server's connectivity, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.
- 2 In the **KMIP Servers** panel, click the **More Options** button in the same row of the server you wish to test and select **Test**. The **Test KMIP Server Configuration** pop-up window appears displaying the success and/or failure of the configuration test.

If any portion of the the test fails, edit your KMIP server configuration.

- 3 Click **Close**.

Testing a KMIP Server Configuration using the CLI

You can alternately test a KMIP server configuration using the FlashBlade CLI.

To test a KMIP server configuration, issue the **purekmip test** CLI command.

You must specify the KMIP server name.

In the following example, the configuration for KMIP server `kmip1` is tested:

```
$ purekmip test kmip1
```

If any portion of the test fails, edit your KMIP server configuration using the **`purekmip setattr`** CLI command.

For complete details about the **`purekmip`** command, refer to the *FlashBlade CLI Reference*.

Editing a KMIP Server Configuration

To edit a KMIP server configuration, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.
- 2 In the **KMIP Servers** panel, click the **More Options** button in the same row of the server you wish to edit and select **Edit**. The **Edit KMIP Server Configuration** pop-up window appears.
- 3 Click the **Add (+)** button to the right of the **URIs** field. Enter one or more URI for the KMIP server and click **OK**.
- 4 Perform one of the following: In the **CA Certificate** field select the CA certificate to validate the KMIP server, or in the **CA Certificate Group** field, select the CA certificate group to validate the KMIP server.
- 5 Click **Save**.

The edited KMIP server is updated with your changes in the **KMIP Servers** panel.

Editing a KMIP Server Configuration using the CLI

You can alternately edit a KMIP server configuration using the FlashBlade CLI.

To edit a KMIP server configuration, issue the **`purekmip setattr`** CLI command.

Editing a KMIP server configuration is similar to creating a KMIP server configuration. You must specify the KMIP server name. You can optionally specify the following:

- **`--uri`** - One or more (in a comma separated list) server URIs. Note that multiple URIs are recommended for high availability.
- **`--ca-certificate`** OR **`--ca-certificate-group`** - The name of server's CA certificate, or the name of the CA certificate group for all external certificates.

In the following example, the current configuration for the KMIP server is edited to add the additional server URI `kmip3.company.com:5696`:

```
$ purekmip create kmip1 --uri  
kmip1.company.com:5696,kmip2.company.com:5696,kmip3.company.com:5696
```

For complete details about the **purekmip** command, refer to the *FlashBlade CLI Reference*.

Deleting a KMIP Server Configuration

 **Note:** The KMIP server used to enable Rapid Data Locking (RDL) cannot be deleted.

To delete a KMIP server configuration, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.
- 2 In the **KMIP Servers** panel, click the **More Options** button in the same row of the server you wish to delete and select **Delete**. The **Delete KMIP Server Configuration** pop-up window appears.
- 3 Click **Delete**.

Deleting a KMIP Server Configuration using the CLI

You can alternately delete a KMIP server configuration using the FlashBlade CLI.

To delete a KMIP server configuration, issue the **purekmip delete** CLI command.

You must specify the KMIP server name.

In the following example, KMIP server `kmip1` is deleted:

```
$ purekmip delete kmip1
```

For complete details about the **purekmip** command, refer to the *FlashBlade CLI Reference*.

Enabling RDL

 **Note:** In order to enable Rapid Data Locking (RDL), the KMIP server cluster must be available.

 **CAUTION:** Enabling RDL is an irreversible operation. Once enabled, RDL cannot be disabled without erasing all data on the FlashBlade.

To enable RDL, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.
- 2 Click the **More Options** button in the header of the **Rapid Data Locking** panel and select **Enable**. The **Enable RDL** pop-up window appears.
- 3 Select a KMIP server.
- 4 Select the checkbox verifying that you understand the consequences of enabling RDL.
- 5 Click **Enable**.
- 6 Click **Create**.

The **Rapid Data Locking** panel is updated displaying the KMIP server and that RDL is enabled (**True**).

Enabling RDL using the CLI

 **Note:** In order to enable Rapid Data Locking (RDL), the KMIP server cluster must be available.

 **CAUTION:** Enabling RDL is an irreversible operation. Once enabled, RDL cannot be disabled without erasing all data on the FlashBlade.

You can alternately enable RDL using the FlashBlade CLI.

To enable RDL, issue the **puredatalock enable** CLI command.

You must specify the name of the KMIP server using the `--kmip-server` argument.

In the following example, RDL is enabled between the FlashBlade and the KMIP server `kmip1`:

```
$ puredatalock enable --kmip-server kmip1
```

For complete details about the **puredatalock** command, refer to the *FlashBlade CLI Reference*.

Testing the RDL Configuration

Once you have enabled Rapid Data Locking (RDL), you can test connectivity and authentication with the KMIP server. The test verifies if there is a secure server connection and if the server contains the expected authentication keys.

To test your RDL configuration, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.

- 2 Click the **More Options** button in the header of the **Rapid Data Locking** panel and select **Test**. The **Test RDL Configuration** pop-up window appears displaying the success and/or failure of the configuration test.
- 3 Click **Close**.

Testing the RDL Configuration using the CLI

You can alternately test FlashBlade's connectivity and authentication with the KMIP server using the FlashBlade CLI.

To test your RDL configuration, issue the **puredatalock test** CLI command.

Test failure is likely the result of an issue with your KMIP server configuration. Edit your KMIP server configuration using the **purekmip setattr** command to resolve the issue and then re-run the **puredatalock test** command.

For complete details about the **puredatalock** command, refer to the *FlashBlade CLI Reference*.

Rotating Keys

 **Important!** Do not rotate external keys if you are performing an EKM (Enterprise Key Management) HA upgrade.

You can change the secret key used with your KMIP server by rotating your RDL configuration's keys. Note that for security, it is recommended that rotation of external keys on the KMIP server should be regularly scheduled.

In order to rotate keys, RDL must be enabled and the KMIP server cluster must be available. Interruption during key rotation will result in key rotation to fail.

To rotate keys for your RDL configuration, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Rapid Data Locking** under **Security**.
- 2 Click the **More Options** button in the header of the **Rapid Data Locking** panel and select **Rotate keys**. The **Rotate Keys** pop-up window appears.
- 3 Click **Rotate**.

Rotating Keys using the CLI

You alternately can change the secret key used with your KMIP server by rotating your RDL configuration's keys using the FlashBlade CLI.

To rotate your RDL configuration's secret key, issue the **`puredatalock rotate`** CLI command.

For complete details about the **`puredatalock`** command, refer to the *FlashBlade CLI Reference*.

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 11.30 Viewing the Session Log

Session Log							
1-20 of 80 < >							
Start Time	End Time	Interface	User	Location	Event	Method	Event Count
All	All						
-	2021-09-29 15:25:28	REST	UNKNOWN	10.64.241.122	logout		1
-	2021-09-29 15:25:29	REST	UNKNOWN	10.64.241.122	logout		1
-	2021-09-29 15:25:30	REST	UNKNOWN	10.64.241.122	logout		1
-	2021-09-29 15:40:48	REST	UNKNOWN	10.64.241.122	logout		1
-	2021-09-29 15:40:48	REST	UNKNOWN	10.64.241.122	logout		1
-	2021-09-29 15:40:49	REST	UNKNOWN	10.64.241.122	logout		1
2021-10-04 13:27:14	-	GUI	pureuser	10.231.216.25	login	password	1
2021-10-04 13:15:00	2021-10-04 13:30:00	GUI	UNKNOWN	10.231.216.25	failed authentication	password	1
2021-10-04 12:15:00	2021-10-04 12:30:00	CLI	lr	172.16.231.35	failed authentication	password	3
2021-10-04 10:21:53	-	GUI	lr	10.231.211.93	login	password	1
2021-10-04 10:15:00	2021-10-04 10:30:00	GUI	UNKNOWN	10.231.211.93	failed authentication	password	3
2021-10-01 11:19:33	-	GUI	lr	10.231.211.93	login	password	1
2021-10-01 11:14:03	2021-10-01 12:28:32	CLI	lr	10.231.211.93	user session	password	1
2021-09-30 15:20:18	-	GUI	lr	10.231.210.116	login	password	1

The **Session Log** displays the following information:

- **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**. 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**, and **REST**.
- **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**, **password**, **one-time password**, and **publickey**. Note that **one-time password** is only applicable for the user **pureeng**.
- **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.

Users

The **Users** panel manages local and remote FlashBlade array users.

To access the **Users** panel, from the **Settings** page navigation sidebar, click **Users** under **Security**.

The panel displays the following:

- **Name:** The user name.

- **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 user API tokens and user array authentication 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 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 **API Clients** panel and select **Create API Token**. The **Create API Token** pop-up window appears.
- 3 Enter 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**.

Editing Users

When editing a user, you can set and change the user's password and manage the user's public key configuration.

 **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 edit a user, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 In the **Users** panel, locate the user you wish edit. Click the **More Options** button in the same row, and select **Edit User**. The **Edit User** pop-up window appears.

- 3 (Optional) Enter and confirm a new password.
- 4 If a user's public key is set up you can optionally **Retain**, **Overwrite**, or **Remove** the public key. If you select **Overwrite**, enter the new public key in the displayed field.
- 5 Click **Save**.

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 12

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 permissions 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 `servicePrincipalNames` 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_PASSWORD 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 FlashBlade, 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 propagate 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

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 **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.
- 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 **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**.
- 10 Select the **Encryption Type**. The default encryption type is **aes256-sha1**. If NTLMv2 is required, **arcfour-hmac** encryption must be enabled.
- 11 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.
- 12 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 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 omitted, the joined domain is assumed. All servers specified 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 mynewadacct --domain mycorp.com --computer-name  
flashblade01
```

2 When prompted enter your user name and password.

```
puread account create mynewadacct --domain mycorp.com --computer-name  
flashblade01  
Enter the user name used to join:  
Enter password for user:
```

In this example:

- the created account is `mynewadacct`

- 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 encryption 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 mynewacct --computer-name newcomputer --domain  
mycorp.com --join-existing-account
```

- 2 When prompted enter your user name and password.

```
puread account create mynewacct --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 `mynewacct`.
- 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 12.1 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 available 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 **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 omitted, the joined domain is assumed. All servers specified 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 account'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 myadacct --prefix myacct
```

In this example, **--ad myadacct** specifies the new keytabs are created for the **myadacct** AD configuration and **--prefix myacct** specifies that all new keytabs begin with the prefix **myacct**.

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
ZuieCj1U8MmZWOLyO16EuQaWLB+GQNWC85MAAABMAAIAClRPVE9STy5ORVQAA25mcwAKT0JBti1T
TUItMgAAAAFfkquaBQASACCYmferQZVm6J4KOVtwyZlagvI7XoS5BpYsH4ZA1YLzkwAAAFgAAgAK
VE9UT1JPLk5FVAAESE9TVAABv2Jhbi1zbWItMi50b3RvcM8ubmV0AAAAAV+Sq5oFABIAIJIz96tB
lWbongo5VPDjMvQc8jteHLkGliwfhkDVgvOTAAAVwACAAPUT1RPuk8uTkVUAANuZnMAFW9iYW4t
c2liLTiudG90b3JvLm5ldAAAAFfkquaBQASACCYmferQZVm6J4KOVtwyZlagvI7XoS5BpYsH4ZA
1YLzkwAAAE0AAgAKVE9UT1JPLk5FVAAESE9TVAABk2Jhbi1zbWItMwAAAAFfkquaBQASACCYmfer
QZVm6J4KOVtwyZlagvI7XoS5BpYsH4ZA1YLzkwAAAEwAAgAKVE9UT1JPLk5FVAADbmZzAApYmFu
LXNtYi0zAAAAAV+Sq5oFABIAIJIz96tBlWbongo5VPDjMvQc8jteHLkGliwfhkDVgvOTAAAWAAC
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

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.
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.

Configuring the Array Management Directory Service

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 **URLs** field, type the comma-separated list of up to 30 URLs 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, set the **Enabled** toggle to enable (blue) to enable the directory service.
- 15 Click **Save**.
- 16 Click the **Edit** icon next to **Roles** to configure the role(s) you wish to assign the service (see [Roles](#) for details about FlashBlade roles). The **Edit Directory Service Roles** pop-up window appears. Enter the following for each role you wish to assign:
 - a 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.
 - b 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".
 - c Click **Save**.

Configuring the Array Management Directory Service using the CLI

The array management directory service can also be configured using the **puredds setattr management** command. For full details about the **puredds** 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

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

The NFS directory service with LDAP can also be configured using the **puredns setattr nfs** command. For full details about the **puredns** 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 directory server domain and server.
- 2 Issue the **puredns 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 switch 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 **puredds setattr nfs** command. For full details about the **puredds** command, see the *FlashBlade CLI Reference*.

To remove the NFS with LDAP configuration using the CLI, perform the following:

- 1 Issue the **puredds disable nfs** command to disable NFS.
- 2 Issue the **puredds 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 **puredds 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 switch 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 **purefs setattr nfs** command. For full details about the **purefs** command, see the *FlashBlade CLI Reference*.

To remove the NFS with NIS configuration using the CLI, perform the following:

- 1 Issue the **purefs disable nfs** command to disable NFS.
- 2 Issue the **purefs setattr nfs --nis-server "" --nis-domain ""** command. Entering "" sets the values to null.
- 3 Issue the **purefs 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.

For file sharing over SMB, the base DN of the directory service is used to represent the joined domain and the URIs can be configured for SMB to control what servers are communicated with when joining the domain.

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.

Configuring the SMB Directory Service

Configuring the SMB 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 SMB, the FlashBlade array must be integrated with an Active Directory service. The SMB account must have read privileges for users and groups, including write permissions to add computer objects in Active Directory.

To configure the SMB directory service for protocol support, perform the following:

- 1 Verify that the Active 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 Verify that the FlashBlade array can access the directory service through a management interface. Interfaces are managed through the **Subnets** panel.
- 4 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
- 5 In the **Directory Service** panel, select **SMB**.
- 6 Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.
- 7 Under the **Service Name** field, select **LDAP**.
- 8 In the **URIs** field, type the comma-separated list of up to 30 URIs of the directory servers.

For file sharing over SMB, this field is optional. If not specified, the base DN of the SMB directory service is used to try to discover domain controllers within an Active Directory domain. If URIs are configured for SMB, they will be treated as the preferences for the servers with which FlashBlade will communicate when joining the domain.

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.

- 9 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`.
- 10 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.
- 11 In the **Bind Password** field, type the password for the bind user account.
- 12 In the **Join OU (SMB)** field, enter the relative DN of the organizational unit (OU) within your domain where the system machine account should be created when joining the domain for SMB. For example, `OU=Arrays,OU=Storage,OU=ServiceMachines`.
- 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, set the **Enabled** toggle to enable (blue) to enable the directory service.
- 15 Click **Save**.

Configuring the SMB Directory Service using the CLI

The SMB directory service can also be configured using the CLI.

The SMB directory service is configured using the **puredds setattr smb** command. For full details about the **puredds** command, see the *FlashBlade CLI Reference*.

To configure the SMB 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 directory server domain and server. Refer to the *FlashBlade CLI Reference* for complete details about the **puredns** command.
- 2 Run the **purenetwork list** command to verify that the FlashBlade can access the directory service through a management interface. Refer to the *FlashBlade CLI Reference* for complete details about the **purenetwork** command.
- 3 Issue the **puredds setattr smb** command with the following arguments:

- **--uri:** (Optional) Up to 30 URIs can be specified in a comma-separated list.

For file sharing over SMB, this field is optional. If not specified, the base DN of the SMB directory service is used to try to discover domain controllers within an Active Directory domain. If URIs are configured for SMB, they will be treated as the preferences for the servers with which FlashBlade will communicate when joining the domain.

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`.
- `--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.
- `--join-ou`: The relative DN of the organizational unit (OU) within your domain where the system machine account should be created when joining the domain for SMB. For example, `OU=Arrays,OU=Storage,OU=ServiceMachines`.

For example:

```
pureds setattr --uri ldaps://ad1.mycompany.com,ldaps://ad2.mycompany.com --  
base-dn DC=mycompany,DC=com --bind-user CN=John,OU=Users,DC=mycompany,DC=com  
smb
```

In this example:

- the URI is set to `ad1.mycompany.com` and `ad2.mycompany.com`
 - the scheme is set to `ldaps://` to enable SSL
 - the base DN is set to be the correct domain components (DC)
 - the bind username is set as the username used to bind to and query the directory
- 4 Test the configuration using the **pureds test smb** command. If the test fails, update the configuration information.
 - 5 If the test passes, enable the directory service using the **pureds enable smb** command.



Pure Storage, Inc.

Twitter: @purestorage

2555 Augustine Drive, 1st Floor

Santa Clara, CA 95054 USA

T: 800-379-7873

Sales: sales@purestorage.com

Support: support@purestorage.com

Media: pr@purestorage.com

General: info@purestorage.com