

FlashBlade Administration Guide

Version 4.6.7

Copyright Statement

© 2026 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.....	16
Organization of this Guide	16
A Note on Format and Content.....	17
Documentation and Support.....	17
Chapter 2 Purity//FB Overview.....	18
Purity//FB Conventions.....	18
Object Names.....	18
Array and File System Sizes.....	18
Time Zones.....	19
IP Addresses.....	19
FlashBlade Zero Move Tiering.....	20
SMB Port Assignments.....	21
Management Network Ports.....	21
Firewall Ports.....	21
Data Network Ports.....	22
About NFSv4.1 Support.....	22
Top Level Mounting.....	23
Troubleshooting.....	23
Chapter 3 Purity//FB GUI Overview.....	25
Purity//FB GUI Navigation.....	26
Purity//FB GUI Login.....	29
Logging in to the Purity//FB GUI.....	30
Viewing and Accepting the End User Agreement.....	30
Chapter 4 The Purity//FB GUI Dashboard Page.....	31
Capacity.....	32
Performance.....	34
Note about the Performance Charts.....	36
Recent Alerts.....	36
Replication Bandwidth.....	37
Hardware Health.....	37
Chapter 5 The Purity//FB GUI Storage Page.....	40
Array.....	40

Connecting Arrays.....	42
Editing a Connected Array.....	45
Disconnecting a Connected Array.....	45
Connecting an S3 Target.....	46
Updating an S3 Target.....	46
Disconnecting an S3 Target.....	46
Bandwidth Throttling.....	47
Realms.....	48
Creating a Realm.....	50
Destroying a Realm.....	50
Renaming a Realm.....	51
Create a Management Access Policy for Realms.....	51
Creating a Realm User.....	51
Sharing a Subnet with a Realm.....	52
Creating a Network Interface for a Realm.....	52
Creating a DNS Configuration for a Realm.....	53
Creating a Server for a Realm.....	53
Configuring the Network Interface for a Realm.....	54
Creating a File System in a Realm.....	54
Creating a File System Export for a Realm.....	54
Servers.....	55
Creating a Server.....	57
Exporting File Systems.....	58
Editing a File System Export.....	59
Deleting a File System Export.....	59
File Systems.....	60
Pure File SafeMode.....	65
WORM File Systems.....	65
About Creating and Exporting File Systems.....	67
File System Size.....	69
File System Usage Limits.....	70
User and Group Usage Limits.....	71
File System Group Ownership.....	71
File System Snapshots.....	72
ACLs.....	73
Group Policy Objects.....	75
Multi-Domain Trust.....	76

NFS and File Locking.....	78
Fast Remove.....	79
File Replication Failover, Reprotect, and Failback.....	81
Creating and Exporting a File System.....	82
Setting a File System's Storage Class.....	91
Setting a File System's Storage Class Tiering Policy.....	92
Changing a File System's Storage Class Tiering Policy.....	92
Removing a Storage Class Tiering Policy from a File System.....	93
Renaming a File System.....	93
Creating a File System Snapshot.....	93
Enabling and Disabling Fast Remove.....	100
Promoting and Demoting File Systems.....	101
Enabling and Disabling File System Writes.....	102
Applying a Legal Hold to a File System.....	102
Releasing a Legal Hold from a File System.....	103
Viewing File System Performance.....	103
Destroying a File System.....	105
Recovering a File System.....	106
Eradicating a File System.....	107
Performing File Replication Failover and Reprotect.....	107
Performing File Replication Failback.....	108
Open Files and Sessions Management with the Microsoft Management Console.....	110
Object Store.....	112
Object Store Usage Limits.....	116
Object Store Account Information.....	117
Creating an Account.....	118
Editing an Account.....	118
Deleting an Account.....	119
Creating an Account for a Realm.....	119
Creating an Account Export.....	120
Creating a Role.....	120
Editing a Role.....	121
Deleting a Role.....	121
Creating Trust Policy Rules.....	121
Editing a Trust Policy Rule.....	122
Removing a Trust Policy Rule.....	123
Uploading a Trust Policy in IAM Format.....	123

Downloading a Trust Policy in IAM Format.....	124
Users.....	124
Buckets.....	131
Chapter 6 The Purity//FB GUI Protection Page.....	157
Snapshots.....	157
Destroying File System Snapshots.....	162
Restoring a File System to a Snapshot	163
Removing Snapshot Policies from Snapshots.....	165
Recovering File System Snapshots.....	166
Eradicating File System Snapshots.....	167
File Replication.....	168
Creating a File System Replica Link.....	169
Adding a Policy to a File Replica Link.....	171
Removing a Policy from a File Replica Link.....	171
Creating Fan Out File System Replica Links.....	171
Forcing Synchronization for Partial File System Replica Links.....	172
Deleting a File System Replica Link.....	172
Deleting Partial File System Replica Links.....	173
Object Replication.....	173
Object Replication Write Conflict Continuation.....	176
Creating a Bucket Replica Link.....	176
Editing a Bucket Replica Link.....	177
Pausing a Bucket Replica Link.....	178
Resuming a Bucket Replica Link.....	178
Deleting a Bucket Replica Link.....	179
Creating Remote Credentials.....	179
Editing Remote Credentials.....	180
Deleting Remote Credentials.....	180
Legal Holds.....	181
Creating a Legal Hold.....	183
Editing a Legal Hold.....	183
Applying a Legal Hold.....	183
Releasing a Legal Hold.....	184
Deleting a Legal Hold.....	185
Chapter 7 Replication and Recovery.....	186
File System Replication Overview.....	186

Setting Up File System Replication Using the GUI.....	191
Creating Replication Network Interfaces.....	191
Creating a Connection Key.....	192
Connecting the Source and Target Arrays.....	192
Creating a Policy.....	194
Creating a File System Replica Link.....	194
Setting Up File System Replication Using the CLI.....	196
Creating Replication Network Interfaces.....	196
Creating a Connection Key.....	197
Connecting the Source and Target Arrays.....	197
Creating a Policy.....	198
Creating a File System Replica Link.....	199
Performing File System Replication Failover Using the GUI.....	200
Performing File System Replication Failback Using the GUI.....	202
Performing File System Replication Failover Using the CLI.....	204
Performing File System Replication Failback Using the CLI.....	206
Performing File System Fan Out Replication Failover Using the GUI.....	208
Performing File System Fan Out Replication Failback Using the GUI.....	210
Performing File System Fan Out Replication Failover Using the CLI.....	212
Performing File System Fan Out Replication Failback Using the CLI.....	214
Object Replication Overview.....	216
Setting Up Array to Array Object Replication Using the GUI.....	219
Creating Replication Network Interfaces.....	219
Creating a Connection Key.....	219
Connecting the Source and Target Arrays.....	220
Creating Remote Credentials.....	221
Creating a Bucket Replica Link.....	222
Enabling Bucket Versioning.....	222
Setting Up Array to Array Object Replication Using the CLI.....	223
Creating Replication Network Interfaces.....	223
Creating a Connection Key.....	223
Connecting the Source and Target Arrays.....	224
Creating Remote Credentials.....	225
Creating a Bucket Replica Link.....	225
Enabling Bucket Versioning.....	226
Setting Up FlashBlade to Amazon S3 Object Replication Using the GUI.....	227
Creating a Replication Network Interface.....	227

Connecting the FlashBlade Array to the S3 Target.....	227
Creating Remote Credentials.....	228
Creating a Bucket Replica Link.....	228
Enabling Bucket Versioning.....	229
Setting Up FlashBlade to Amazon S3 Object Replication Using the CLI.....	229
Creating a Replication Network Interface.....	229
Connecting the FlashBlade Array to the S3 Target.....	230
Creating Remote Credentials.....	230
Creating a Bucket Replica Link.....	230
Enabling Bucket Versioning.....	231
Chapter 8 The Purity//FB GUI Policies Page.....	232
Data Eviction Policies.....	233
Creating a Data Eviction Policy.....	234
Editing a Data Eviction Policy.....	234
Renaming a Data Eviction Policy.....	234
Enabling and Disabling Data Eviction Policies.....	235
Deleting Data Eviction Policies.....	235
Attaching a Data Eviction Policy to a File System from the Data Eviction Policy Details Page.....	236
Removing a Data Eviction Policy from a File System from the Data Eviction Policy Details Page.....	236
Attaching a Data Eviction Policy to a File System from the File System Details Page.....	236
Removing a Data Eviction Policy from a File System from the File System Details Page....	237
File System Audit.....	237
Creating a File System Audit Policy.....	240
Editing a File System Audit Policy.....	240
Attaching a File System Audit Policy to a File System.....	241
Removing a File System Audit Policy from a File System.....	241
Enabling and Disabling a File System Audit Policy.....	242
Deleting a File System Audit Policy.....	242
Deleting a Syslog Log Target.....	243
NFS Export Policies.....	243
NFS Export Rules.....	245
Creating NFS Export Policies and Rules.....	251
Creating File System Exports.....	252
Editing File System Exports.....	252
Deleting File System Exports.....	253

Adding Additional NFS Export Rules to NFS Export Policies.....	254
Editing NFS Export Rule Clients.....	254
Editing NFS Export Rules.....	255
Reordering NFS Export Rules.....	256
Removing NFS Export Rules.....	256
Renaming NFS Export Policies.....	257
Enabling and Disabling NFS Export Policies.....	257
Deleting NFS Export Policies.....	258
Object Store Access Policies.....	258
About Object Store Access Policies.....	259
Rules: Actions, Resources, and Conditions.....	261
Viewing Object Store Policy Details.....	267
Creating Object Store Access Policies.....	267
Adding Object Store Access Policy Rules.....	269
Editing Object Store Access Policy Rules.....	270
Copying Object Store Access Policy Rules to Other Policies.....	271
Removing Object Store Access Policy Rules.....	272
Adding Users to an Object Store Access Policy.....	273
Removing Users from an Object Store Access Policy.....	273
Deleting Object Store Policies.....	273
Object Store Audit.....	274
Creating an Object Store Log Target.....	279
Editing an Object Store Log Target.....	279
Creating an Object Store Audit Policy.....	279
Editing an Object Store Audit Policy.....	280
Adding Buckets to an Object Store Audit Policy.....	280
Removing Buckets from an Object Store Audit Policy.....	281
Enabling and Disabling an Object Store Audit Policy.....	281
Renaming an Object Store Audit Policy.....	281
Deleting an Object Store Audit Policy.....	282
Deleting an Object Store Log Target.....	282
File System Quality of Service Policies	282
Creating a File System QoS Policy.....	283
Editing a File System QoS Policy.....	284
Enabling a File System QoS Policy.....	284
Disabling a File System QoS Policy.....	284
Renaming a File System QoS Policy.....	284

Attaching a QoS Policy to File Systems from the File Systems Details Page.....	285
Attaching a QoS Policy to File Systems from the QoS Policy Details Page.....	285
Removing a Single QoS Policy from a File System from the File Systems Details Page....	285
Removing QoS Policies from File Systems from the QoS Policy Details Page.....	286
Deleting a File System QoS Policy.....	286
Object Store Quality of Service Policies.....	287
Creating an Object Store QoS Policy.....	287
Editing an Object Store QoS Policy.....	288
Enabling an Object Store QoS Policy.....	288
Disabling an Object Store QoS Policy.....	288
Renaming an Object Store QoS Policy.....	289
Attaching an Object Store QoS Policy to a Bucket from the Bucket Details Page.....	289
Attaching an Object Store QoS Policy to Buckets from the QoS Policy Details Page.....	289
Removing a Single Object Store QoS Policy from a Bucket from the Bucket Details Page	290
Removing Buckets from an Object Store QoS Policy from the QoS Policy Details Page...	290
Deleting an Object Store QoS Policy.....	290
Adding Members to a QoS Policy.....	291
Removing Members from a QoS Policy.....	291
S3 Export Policies.....	292
Creating an S3 Export Policy.....	292
Adding a Rule to an S3 Export Policy.....	293
Renaming an S3 Export Policy.....	293
Disabling an S3 Export Policy.....	293
Deleting an S3 Export Policy.....	294
SMB Client and Share Policies.....	294
SMB Client Policies.....	295
SMB Share Policies.....	304
File System Snapshot Policies.....	312
Viewing Snapshot Policy Details.....	313
Creating Snapshot Policies.....	313
Adding Snapshot Policy Rules.....	313
Adding Multiple Snapshot Policy Rules.....	314
Editing Snapshot Policy Rules.....	315
Removing Snapshot Policy Rules.....	316
Adding File Systems to a Snapshot Policy.....	316
Removing File Systems from a Snapshot Policy.....	316
Adding Replica Links to a Snapshot Policy.....	317

Removing Replica Links from a Snapshot Policy.....	317
Enabling and Disabling Snapshot Policies.....	318
Removing a Snapshot from a Policy.....	318
Removing a Policy from a Snapshot.....	319
Deleting Snapshot Policies.....	320
Storage Class Tiering Policies.....	321
Creating a Storage Class Tiering Policy.....	322
Editing a Storage Class Tiering Policy.....	322
Enabling and Disabling Storage Class Tiering Policies.....	323
Deleting a Storage Class Tiering Policy.....	323
Attaching File Systems to Storage Class Tiering Policies.....	323
Removing File Systems from Storage Class Tiering Policies.....	324
File System WORM Policies.....	325
Viewing File System WORM Policy Details.....	325
Creating File System WORM Policies.....	326
Editing File System WORM Policies.....	327
Deleting File System WORM Policies.....	328
Management Access Policies.....	328
Creating a Management Access Policy.....	329
Deleting a Management Access Policy.....	329
Enabling and Disabling a Management Access Policy.....	330
Renaming a Management Access Policy.....	330
Management Authentication Policies.....	330
Creating a Management Authentication Policy.....	332
Adding Members to a Management Authentication Policy.....	332
Removing Members from a Management Authentication Policy.....	333
Editing a Management Authentication Policy.....	333
Renaming a Management Authentication Policy.....	334
Enabling and Disabling Management Authentication Policies.....	334
Deleting a Management Authentication Policy.....	335
Password Policies.....	335
Editing Password Policies.....	336
Disabling Password Policies.....	338
Enabling Password Policies.....	338
SSH Certificate Authority Policies.....	338
Adding a Public Key.....	339
Creating an SSH Certificate Authority Policy.....	339

Adding Users to an SSH Certificate Authority Policy.....	340
Adding an Array to an SSH Certificate Authority Policy.....	340
Network Access Policies.....	340
Renaming a Network Access Policy.....	343
Adding a Network Access Rule to a Network Access Policy.....	343
Editing a Network Access Rule.....	344
Reordering Network Access Rules.....	344
Removing a Network Access Rule from a Network Access Policy	345
TLS Policies.....	346
Creating a TLS Policy.....	348
Editing a TLS Policy.....	349
Deleting a TLS Policy.....	349
Viewing a TLS Policy's Effective Values.....	350
Setting the Array's Default TLS Policy.....	350
Attaching TLS Policies to Network Interfaces.....	351
Removing TLS Policies from Network Interfaces.....	351
Telemetry Policies.....	351
Updating an Attached Metrics Target Policy.....	353
Editing a Telemetry Metrics Policy.....	353
Creating a Telemetry Target.....	353
Renaming a Telemetry Target.....	354
Disabling a Telemetry Target.....	354
Editing a Telemetry Target.....	355
Deleting a Telemetry Target.....	355
Chapter 9 The Purity//FB GUI Analysis Page.....	356
Performance.....	356
The Performance Chart.....	357
Displaying Protocol-Specific Metrics.....	358
Note about the Performance Charts.....	361
Displaying Client, Group, and User Performance Statistics.....	362
Capacity.....	366
Replication.....	368
Network.....	369
Viewing Network Connector Statistics.....	372
Displaying a Different Time Range	374
Chapter 10 The Purity//FB GUI Health Page.....	375

Hardware.....	375
FlashBlade//S and //E System Resiliency.....	377
Alerts.....	378
Flagging Alert Messages.....	380
Unflagging Alert Messages.....	380
Chapter 11 The Purity//FB GUI Settings Page.....	381
Side Navigation Bar.....	383
Working with System Settings.....	383
Alert Watchers.....	384
Alert Routing.....	386
File System User/Group Quota Notifications.....	388
Array Details.....	389
Attached Policies.....	391
SNMP.....	392
Syslog Server.....	404
NTP Servers.....	409
Support.....	410
Working with Hardware Settings.....	418
Viewing Blade Information.....	418
Viewing Drive Information.....	418
Viewing Array Hardware Information.....	419
Working with Network Settings.....	419
Connectors.....	419
Diagnostics.....	420
DNS Configurations.....	423
Link Aggregation Groups.....	424
Object Store Virtual Hosts.....	427
Subnets and Interfaces.....	432
Working with Security Settings.....	438
API Clients.....	439
Viewing the Audit Trail.....	441
Array Certificates.....	442
External Certificates.....	447
Certificate Groups.....	449
Viewing the SMB Mode.....	451
Active Directory Accounts.....	452
Directory Service.....	454

Kerberos Keytabs.....	455
Rapid Data Locking.....	458
Session Log.....	466
Single Sign-On.....	468
Configuring SSO Authentication for Object Service.....	473
Configuring SSO Authentication with SAML2 SSO for Management Service.....	478
Users.....	491
Chapter 12 The Purity//FB Fusion.....	498
Fusion Management.....	498
Requirements and Restrictions.....	499
Secure Communication in Arrays.....	499
Fleet Keys.....	499
Fleet Creation and Membership.....	500
Generating a Fleet Key.....	500
Joining an Existing Fleet.....	501
Removing a Remote Array from the Fleet.....	501
Search for Member Arrays.....	501
Summary of Steps to Create a Fleet with Multiple Array Members.....	502
Presets and Workloads.....	502
Creating a Preset.....	502
Configuring a Workload.....	503
Portable Presets.....	503
Remote Resource Management.....	505
Fleet Data Protection.....	505
Searching Across a Fleet.....	506
Viewing Audit Entries for a Remote Array.....	506
Chapter 13 Choosing an Active Directory Account or Directory Services Configuration..	507
Active Directory Overview.....	511
Active Directory Networking Requirements.....	511
Computer Account Naming and Credentials.....	513
Minimum Permissions for Joining an Active Directory Domain.....	514
Joining an Organizational Unit.....	515
Joining with an Existing Computer Account.....	515
Service Principal Names.....	516
Directory and Kerberos Servers.....	517
User and Group Name Mapping for Active Directory and LDAP and NIS.....	518

Creating an Active Directory Account.....	518
Joining Active Directory with an Existing Computer Account.....	521
Testing an Active Directory Configuration.....	523
Editing an Active Directory Account.....	525
Rotating Keytabs.....	527
Importing a Keytab File.....	529
Exporting a Keytab File.....	529
Deleting a Keytab File.....	530
Deleting an Active Directory Account.....	531
Directory Services Overview.....	532
NFS and SMB Nested User Groups.....	533
Array Management Directory Service.....	534
Configuring the Array Management Directory Service.....	536
NFS Directory Service.....	542
Configuring the NFS Directory Service with LDAP.....	543
Configuring the NFS Directory Service with NIS.....	548
SMB Directory Service.....	551

Chapter 1

About this Guide

The Pure Storage FlashBlade® Administration 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 //S100, //S200 and //S500 models. While //S100 provides a lower entry point to leverage all the advantages of scale-out flash storage in an efficient manner, 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.

Additionally, Pure Storage provides specific configurations of FlashBlade //S500 with expansion chassis, which leverage Zero Move Tiering allowing you to pin entire workloads to two storage classes: speed and archive. Contact Pure Technical Services for additional information about Zero Move Tiering configurations.

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 Storage Technical Services at support@purestorage.com.

Chapter 2

Purity//FB Overview

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

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

Purity//FB Conventions

Purity//FB follows certain conventions.

Object Names

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

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

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

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

Array and File System Sizes

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

Note that the raw capacity value displayed on the **Hardware > Health** page displays gibibyte (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
```

FlashBlade Zero Move Tiering

FlashBlade Zero Move Tiering introduces the notion of two storage classes: speed and archive.

FlashBlade Zero Move Tiering allows you to commit entire file system and bucket workloads to these two storage classes, enabling high performance for some data and archiving data for future use. These storage classes are represented in the GUI as S500X-S for speed (hot data) and S500X-A for archive (cold data).

FlashBlade Zero Move Tiering is available for hybrid system configurations consisting of a FlashBlade//S500 chassis with Expansion chassis in a single cluster and namespace (FB//S500+EX with Zero Move Tiering).

- The //S500 chassis contains FB-S500 blades with 37TB drives and is used for the speed storage class. It is designed to handle frequently accessed data requiring high-speed performance.
- The EX (Expansion) chassis contains FB-EX blades with 75TB drives for the archive storage class. It is designed for infrequently accessed data that can tolerate lower performance.

For systems using FlashBlade Zero Move Tiering, by default, file systems and buckets are set to the S500X-S (speed) storage class upon creation. File systems and buckets can be easily moved from one storage class to the other.

For file systems, FlashBlade allows for the creation of storage class tiering policies that can be applied to file systems. Storage class tiering policies allow you to specify the criteria for which data should be moved between hot and cold storage.

- For information about creating file systems, see [Creating and Exporting a File System](#).
- For information about viewing capacity usage for storage classes, see [Capacity](#).
- For information about setting a file system's storage class, see [Setting a File System's Storage Class](#).

For additional information about FlashBlade Zero Move Tiering configurations, contact Pure Technical Services.

SMB Port Assignments

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

Management Network Ports

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

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

Firewall Ports

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

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

Data Network Ports

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

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

About NFSv4.1 Support

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

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

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

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


- pNFS
- Delegation
- Referrals
- Shared reservation

⚠ Important! NFSv4.0 is not supported.

Top Level Mounting

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

Figure 2.1 Top Level Mounting



Troubleshooting

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

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

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

Error Messages

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

Service Temporarily Unavailable

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

Unexpected Error

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


Chapter 3

Purity//FB GUI Overview

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

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

Figure 3.1 Purity//FB GUI



The Purity//FB GUI contains the following pages:

Dashboard

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

Storage

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

Servers tab displays information for all servers on the array. An **Object Store** tab is provided to quickly access object store accounts, buckets, and users.

Policies

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

Analysis

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

Fleet

Displays member arrays that are connected to the FlashBlade. This allows you to manage multiple arrays through a single array GUI.

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										
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 an orange 'Log In' button. At the bottom of the login area is a 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

Capacity is displayed in the Dashboard in **Array Capacity** and **Capacity by Storage Class** panels.

 **Note:** The **Capacity Storage Class** panel is only visible for Zero Move Tiering systems.

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

Figure 4.2 Dashboard - Array Capacity Panel



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


The capacity data is broken down into the following components:

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

- **Capacity:** Total usable capacity of the array.

For FlashBlade Zero Move Tiering systems, the **Capacity by Storage Class** panel is divided into two halves displaying capacity wheels. One half displays the size and consumption of the speed storage class (**S500X-S**) and the other half displays the size and consumption of the archive storage class (**S500X-A**).

Figure 4.3 Dashboard - Capacity by Storage Class Panel



The speed and archive class capacity data is broken down into the following components:

- **Unique:** Unique physical space occupied by customer data, excluding snapshots and destroyed space.
- **Destroyed:** Unique physical space (excluding snapshots) occupied by destroyed data within the child containers. For buckets and file systems, the destroyed space will be 0 as they cannot have child containers. For the whole array, the space will be the sum of all destroyed buckets and file systems.
- **Snapshots:** Amount of space that snapshots occupy (other than unique) after data compression.
- **Shared:** Physical space occupied by deduplicated data, meaning that the space is shared with other filesystems/buckets and snapshots as a result of data deduplication.
- **Empty:** Unused space.
- **Data Reduction:** Ratio of the size of the written data (i.e. Used) versus the amount of space the data occupies after data compression (i.e. Physical) and/or deduplication.
- **Used:** Amount of space that the written data occupies in each storage class after reduction via data compression.
- **Capacity:** Total usable capacity of each storage class.


Purity DeepReduce™

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

Performance

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

Figure 4.4 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 4.5 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 4.6 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.

Depending on the configuration and type of your array, the **Hardware Health** panel displays either a single or multi-chassis first-generation FlashBlade, FlashBlade//S, FlashBlade Zero Move Tiering configurations, or FlashBlade//E graphic.

Figure 4.7 Hardware Health Panel - First-Generation FlashBlade Single Chassis




Figure 4.8 Hardware Health Panel - First-Generation FlashBlade Multi-Chassis




Figure 4.9 Hardware Health Panel - FlashBlade//S Single Chassis




Figure 4.10 Hardware Health Panel - FlashBlade//S Multi-Chassis




Figure 4.11 Hardware Health Panel - FlashBlade Zero Move Tiering Configuration





Figure 4.12 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 servers, 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 Systems** count opens the **Storage > File Systems** 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.

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 10 array connections; however, only one connection between two arrays can be used for file replication. All 10 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. Each array connection shows three paths. Two paths are for destination port 8117 and one path is for destination port 80 for disabled encryption or 443 for enabled encryption. For object replication, these three paths must all be in the healthy (connected) state.

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.

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.

The maximum number of target arrays is 10.

 **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. File replication over an encrypted connection is supported only if both the source and the remote arrays are running Purity//FB 4.5.2 or higher.

 **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.

 **Important!** If both the source and target arrays are running Purity//FB 4.4.2 or above, before enabling encryption, the CA certificate of the source array, or array certificate itself in case of self signed certificate, must be imported to the target array and added to the target array's `_default_replication_certs` CA certificate group. Likewise, any CA certificates

on the target array must be imported to the source array and added to the source array's `_default_replication_certs` CA certificate group.

In the case of mismatch where one side is running a lower software version, for example the source is running Purity//FB 4.4.2 and the target is running 4.4.0, the CA certificate on the target array must be imported to the source array and added to the source array's `_default_replication_certs` CA certificate group for encrypted object replication.

- 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. You must also import the source array's CA certificate to the target array and add it to its `_default_replication_certs` CA certificate group. Note that encryption is set on the source array only. If enabled, the encryption setting displays as enabled on the target array.

 **Note:** File replication over an encrypted connection is supported only if both the source and remote arrays are running Purity//FB 4.5.1 or higher.

 **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.

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.

 **Note:** File replication over an encrypted connection is supported only if both the source and remote arrays are running Purity//FB 4.5.2 or higher.

-
-
-
-
- 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 5.3 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

Default Limit: 250 M
Window Limit: -

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[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.
- 6 Click **Save**.

Realms

 **Important!** In a multi-tenancy environment, the FlashBlade system has an internal queuing system to manage all concurrent tenant requests, such as legal holds, file replication, and snapshots. Because tenants are isolated from the actions of other tenants, the realm administrator will not have visibility to the status of their requests in the queue or if snapshot limits have been reached, preventing the of taking new snapshots, which may lead to the impact of their data protection goals. Pure recommends that

you contact the FlashBlade administrator for additional information on the queuing status of extended pending requests and snapshot limits.

A realm is a unit of management and accounting for data, a feature that is used mostly by storage providers and others in multitenancy environments. Realms offer a way of creating storage areas that are fully separate and independent of each other, even on the same array.

Realms add the following support:

- **Delegated administration.** Administrators can be limited to accessing and modifying file systems within one or more realms.
- **Consumption controls.** Array administrators can set provisioning quota limits as well as IOPS and bandwidth QoS limits on each realm.
- **Capacity footprint metrics.** Footprint is the consumed capacity for data within the realm as if that were the only data on the array, taking into account any applicable data reduction.

Creating a Realm

When creating a realm, the default will be without default access list. This creates a realm that does not utilize any default access list that may have been configured. If no default access list is configured, commands will be executed without default access being applied, bypassing any access controls defined by that list. If an access list has been designated as the array-default and is not needed for the current command, the realm will not have default access listed in that access list.

Destroying a Realm

When a realm is no longer needed, it can be destroyed if the realm is empty. Ensure all pods, file systems, and network connections inside the realm have been either moved out of the realm or destroyed.

The destruction of realms and their contents is not immediate. Instead, the realm is placed in an eradication pending period. The realm can be manually recovered or permanently eradicated within the eradication pending period. After the eradication pending period has elapsed, the realm is completely eradicated and not recoverable.

The length of eradication pending period is set by the enabled delay value if the protection group is not protected by SafeMode, or by the disabled delay value if the protection group is protected by SafeMode.

Recovering a Realm

A destroyed realm can be recovered at any time during the eradication pending period. Once a realm has been recovered, its destroyed volumes and protection groups can also be recovered to bring the realm and its contents back to their previous state.

Eradicating a Realm

A destroyed realm can be permanently eradicated at any time during the eradication pending period (unless the SafeMode manual eradication prevention feature is enabled). Run the `purepod eradicate` command to eradicate a realm.

Eradication of a realm is possible only if its destroyed file systems are already eradicated. Therefore, before eradicating a realm, first eradicate all its destroyed files systems and objects.

Once a realm has been eradicated, it and its member file systems can no longer be recovered.

Creating a Realm

To create a realm, perform the following:

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

Destroying a Realm

To destroy a realm, perform the following:

 **Important!** Destroying a realm is not immediate and must meet certain conditions. Retention settings, lock policies, SafeMode, and links to file systems and buckets must be removed a realm can be destroyed.

- 1 From the **Storage** page navigation sidebar, click the **Realms** tab.
- 2 In the **Realms** panel, click the **More Options** button and select **Destroy**. The **Destroy Realms** pop-up window appears.
- 3 Select one or more realms from the **Available Realms** panel.
- 4 Verify the selected realms listed in the **Selected Realms** panel to be destroyed.

- 5 Click **Destroy**.

Renaming a Realm

To change the name of an existing realm, perform the following:

- 1 From the **Storage** page navigation sidebar, click the **Realms** tab.
- 2 In the **Realms** panel, click the realm that you would like to change the name.
- 3 Click the **More Options** button in the **Details** panel and select **Rename**. The **Rename Realm** pop-up window appears.
- 4 Enter a new name in the **New Name** field.
- 5 Click **Rename**.

Create a Management Access Policy for Realms

To create a management access policy for a realm, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Access**.
- 2 In the **Management Access Policies** panel, click the **Add** interface button (+). The **Create Management Access Policy** pop-up window appears.
- 3 In the **Name** field, enter a name for the policy.
- 4 Select the **Enable** checkbox to ensure that the policy is enabled once you are finished editing.
- 5 In the **Role** list, select the access role of the policy.
- 6 In the **Scope Type** list, select **realm**.
- 7 In the **Scope Name** field, enter the name of a realm.
- 8 In the **Aggregation Strategy** field, select **all-permissions** or **least-common-permission**.
- 9 Click **Create**.

Creating a Realm User

To create a realm user, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 In the **Users** panel, click the **More Options** button, and click **Create User**. The **Create Local User** pop-up appears.

- 3 In the **User** field, enter a name for the realm user.
- 4 In the **Access Policies** field, click the plus (+) icon for the **Add Management Access Policies** pop-up, then select a policy and click **Add**.
- 5 In the **Password** field, enter a password for the realm user.
- 6 In the **Confirm Password** field, re-enter the password for the realm user.
- 7 Click **Create**.

Sharing a Subnet with a Realm

To share a subnet with a realm, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 2 In the **Subnets** panel, click the Subnet you want to share with a realm. The selected Subnet details page appears.
- 3 In the **Accessible by** panel, click the **More Options** button and click **Add**. The **Add Realms** pop-up window appears.
- 4 In the **Existing Realms** panel, select the realm to be placed in the **Selected Realms** panel.
- 5 Click **Add**.

Creating a Network Interface for a Realm

 **Important!** This procedure is performed as the realm user.

 **Important!** Once a Server and DNS has been created for the realm, return to edit the Network Interface services type from **egress-only** to **data** and select the server.

To create a network interface for a realm, perform the following:

- 1 Log into the Purity GUI of your realm.
- 2 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 3 In the **Network Interface** panel, click the **Add** interface button (+). The **Create Network Interface** pop-up window appears.
- 4 In the **Name** field, enter the realm and network, using double colons (::) to separate the realm name and network. For example **myrealm::vip1**.
- 5 In the **Subnet** field, select the subnet name.

- 6 In the **Address** field, enter the IP address.
- 7 In the **Services** field, select **egress-only**.
- 8 Click **Create**.

Creating a DNS Configuration for a Realm

To create a DNS configuration for a realm, perform the following:

- 1 From the **Settings** page navigation sidebar, click **DNS** under **Network**.
- 2 In the **DNS Configuration** panel, click **Add** interface button (+). The **Create DNS Configuration** pop-up window appears.
- 3 In the **Name** field, enter the
- 4 In the **Scope Type** field, select **realm**.
- 5 In the **Scope Name** field, enter the DNS name.
- 6 In the **Domain** field, enter the domain of the DNS.
- 7 In the **Nameserver 1** field, enter the name server IP address (IPv4)
- 8 (Optional) In the **Nameserver 2** field, enter the server IP address.
- 9 (Optional) In the **Nameserver 3** field, enter the server IP address.
- 10 In the **Source** field, select the realm network interface.
- 11 Click **Create**.

Creating a Server for a Realm

To create a server for a realm, perform the following:

- 1 From the **Storage** page navigation sidebar, click the **Servers** tab.
- 2 In the **Servers** panel, click the **Add** interface button (+). The **Create Server** pop-up window appears.
- 3 In the **Scope Type** section, select the **realm**.
- 4 In the **Scope Name** field, enter the realm name.
- 5 In the **Name** field, enter the server name.
- 6 In the **Selected DNS** field, select the realm and DNS, using double colons (::) to separate the realm name and DNS. For example, **myrealm::dns1**.

- 7 Click **Create**.

Configuring the Network Interface for a Realm

To enable the realm for outbound traffic, the network interface's services configuration needs to be changed from **egress-only** to **data**. To edit the network interface configuration, perform the following:

- 1 Log into the Purity GUI of your realm.
- 2 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 3 In the **Interfaces** panel, find the interface you want to edit then click **Edit** icon. The **Edit Network Interface** pop-up window appears.
- 4 In the **Services** field, select **data**.
- 5 In the **Selected Server** field, select the realm and server name.
- 6 Click **Save**.

Creating a File System in a Realm

 **Important!** This procedure is performed as the realm user.

To create a file system in a realm, perform the following:

- 1 Log into your realm GUI.
- 2 From the **Storage > File Systems**, click the plus (+) icon on the File Systems panel. The **Create File System** pop-up window appears.
- 3 In the **Name** field, enter the realm and file system name, using double colons (::) to separate the realm name and file system name. For example, **myrealm::filesystem1**.
- 4 Click **Create**.

Creating a File System Export for a Realm

 **Important!** This procedure is performed as the realm user.

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

To create a file system export for a realm, perform the following:

- 1 Log into the realm GUI.

- 2 From the **Storage > File Systems** navigation page, in the **File System Exports** panel, click the **Add** interface button (+). The **Create File System Export** pop-up window appears.
- 3 In the **Server** field,
- 4 In the **File System** field,
- 5 In the **Export Name** field,
- 6 In the **Type** field, select **NFS** or **SMB**, which is the type of policy the file system is attached.
- 7 (When the NFS type is selected) In the **NFS Export Policy** list, select the NFS policy.
- 8 (When the SMB type is selected) In the **SMB Client Policy** list, select the policy.
- 9 (When the SMB type is selected) In the **SMB Share Policy** list, select the policy.
- 10 Click **Create**.

Servers


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

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

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

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

Figure 5.4 Servers Tab



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

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

Servers Panel

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

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


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

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

Server Details

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

Figure 5.5 Server Details Page



The screenshot displays the 'Storage' section of the Pure Storage GUI, specifically the 'Servers' tab. The server selected is '_array_server'. The page is divided into several panels:

- File Systems Exports:** A table listing exported file systems. It includes columns for Export Name, File System, Type, Policy, Share Policy, Enabled, and Status. Two entries are shown: 'fs1' for NFS and 'fs1' for SMB.
- Interfaces:** A section for configuring network interfaces. It includes fields for Name, Enabled, Subnet, Address, VLAN, Mask, Gateway, MTU, and Service. A message states 'No network interfaces found.'
- DNS Settings:** A panel for configuring DNS settings, including Name, Domain, Nameservers, Source, and Services.
- Directory Service:** A panel for configuring the Directory Service, including a 'Test All' button and an 'NFS Configuration' sub-panel.
- Active Directory Account:** A panel for configuring an Active Directory account, including fields for Name, Computer Name, Join OU, and Domain.

The server detail page provides the following panels:

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

Creating a Server

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

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

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

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

Exporting File Systems

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

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

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

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

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

- 10 Click **Create**.

Editing a File System Export

To edit a file system export, perform the following:

- 1 Select **Storage > Servers**.
- 2 Click the server used for export.
- 3 Click the **Edit** button in the row of the file system export you want to edit. The **Edit File System Export** pop-up window appears.
- 4 Enter a new export name.
- 5 Select a policy to attach.
- 6 Click **Save**.

Deleting a File System Export

To delete a file system export, perform the following:

- 1 Select **Storage > Servers**.
- 2 Click the server used for export.
- 3 To delete a single file system export:
 - a Click the **Delete** button in the row of the file system export you want to delete. The **Delete File System Export Configuration** pop-up window appears.
 - b Click **Delete**.
- 4 If you wish to delete multiple file system exports:
 - a Click the **More Options** button in the header of the **File Systems Exports** panel and select **Delete**. The **Delete File System Exports** pop-up window appears.
 - b From the **Exsting File System Exports** list, select the file systems you wish to delete. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all file systems by clicking the checkbox next to the search box. All selected file systems appear in the **Selected File System Exports** list.
 - c Click **Delete**.

File Systems

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

Figure 5.6 File Systems Tab

Storage 4									
Array Servers File Systems Object Store									
File Systems									
Size	Virtual	Data Reduction	Unique	Destroyed	Snapshots	Total			
-	0.00	1.0 to 1	0.00	0.00	0.00	0.00			

File Systems									
General Space Storage Class Storage Class Space 1-1 of 1 + ⋮									
Name ▲	Source Location	Source Name	Size	Virtual	Hard Limit	Created	Protocols	Promotion Status	Writable
fs1	-	-	-	0.00	False	2025-05-05 10:39:54	-	promoted	True
Destroyed (0) ▾									

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

File Systems Panel

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

- **Name:** File system name. Clicking a file system name displays any snapshots belonging to that file system, and also provides the option of creating snapshots.
- **Source Location:** The location where the file system was created.
- **Source Name:** If created on the target array, the name of the source file system.
- **Size:** Provisioned size of the file system. If not set, the provisioned size is unlimited.
- **Virtual:** Total amount of pre-compressed data written to the file system.
- **Hard Limit:** Whether the file system's provisioned size has a hard limit (true) or not (false).
- **Created:** The file system creation date.
- **Protocols:** One or more protocols configured for the file system. A dash (-) means the file system has no protocols assigned.

- **Promotion Status:** Whether the source or replicated file system is promoted or demoted.
- **Writable:** Whether the file system is writeable (true) or not (false). When the file system is changed from writable disabled to enabled, it is recommended to remount the file system to avoid stale client cache issues.

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

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

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

 **Note:** The **Storage Class** view is only visible for FlashBlade//S500 with Expansion chassis systems using Zero Move Tiering.

Click **Storage Class** to view file system storage class information. The columns in the **File Systems** panel display the following information:

- **Name:** File system name.
- **Storage Class:** The file system's storage class, either S500X-S (speed) or S500X-A (archive).
- **Status:** The transition status from one storage class (S500X-S/S500X-A) to the requested class (S500X-A/S500X-S). If a transition is ongoing, this field will display In-Progress or Queued. If no transition is active, it will show -.

- **Status Details:** Provides additional information about the transition status. If there is no ongoing transition, this field will display -.

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

- **Name:** File system name.
- **Storage Class:** The file system's storage class, either S500X-S (speed) or S500X-A (archive).
- **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 (i.e. Used) versus the amount of space the data occupies after data compression (i.e. Physical) and/or deduplication.
- **Unique:** The unique physical space occupied by customer data.
- **Snapshots:** Amount of space occupied by snapshots after reduction from data compression.
- **Total:** The total used space (physical space and snapshots).

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


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

In addition to displaying file systems, the **File Systems** panel allows you to create and maintain file systems, as well as assign a file system's storage class (for FlashBlade Zero Move Tiering systems only).

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



Details

Created: 2025-05-05 10:39:54 Quotas »

Group Ownership: creator Default Group Quota -

Hard Limit: False Default User Quota -

Promotion Status: promoted

Source: -

Writable: True

Storage Class: S500X-S

Protocols »

Enabled Protocols: -

Access Control Style: shared

Safeguard ACLs: True

Attached Policies

No policies found.

File System Exports 1-2 of 2

Export Name	Server	Type	Policy	Share Policy	Enabled	Status	
fs1	_array_server	NFS	-	-	False	Protocol Disabled	☒ ☐
fs1	_array_server	SMB	-	-	False	Protocol Disabled	☒ ☐

File System Snapshots General Transfer

No file system snapshots found.

Destroyed (0) ▾

The file systems detail page provides the following panels:

- **Details:** Provides an overview of the file system's configured parameters.
- **File System Exports:** Lists all file system exports.
- **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.
- **Legal Hold Entities:** Lists all legal holds applied to the file system.
- **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.

File Systems Exports Panel

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

- **Export Name:** Exported file system name.
- **Server:** The server used for multi-tenancy. The default server on the array is **_array_server**.
- **File System:** The file system name.
- **Type:** The directory service type, NFS or SMB.
- **Policy:** The NFS export policy. See [NFS Export Policies](#) for additional information.
- **Share Policy:** The SMB share policy. See [SMB Share Policies](#) for additional information.
- **Enabled:** Whether the export is enabled (**True**) or not (**False**).
- **Status:** The current export status.

File Locks Panel

Refer to [Viewing and Managing File Locks](#).

Destroyed File Systems Panel

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

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

- **Name:** File system name.
- **Size:** Provisioned size of the file system.

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

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

Pure File SafeMode

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

WORM File Systems

A Write Once Read Many (WORM) file system is a protected file system that contains WORM files. A WORM file system can be optionally created when creating a file system.

A WORM policy governs file system retention periods, retention lock settings, and retention mode. A WORM policy is attached to a WORM file system and then a file is committed to the WORM state in the WORM file system. Once a file is committed to WORM state, data can be read, but not changed or deleted. A WORM file system can be created after a WORM policy has been created (see [File System WORM Policies](#) for additional information).

Note the following restrictions:

- Fast remove directory cannot be enabled after a WORM file system is created.
- WORM file systems cannot be demoted.
- WORM file systems cannot be restored.
- WORM file systems with WORM Policies can be destroyed, however they will have no expiration.

- WORM file systems cannot be eradicated. Note that files with expired retention periods can be deleted or their WORM state can be extended, and snapshots can be deleted. Eradication of empty WORM file systems will be addressed in a future release.
- WORM file system replication is not supported with this release.
- When creating a WORM file system, a single WORM policy can be attached. Once attached, the policy cannot be detached.

WORM Files

A file can be protected from modification by committing it to the Write Once Read Many (WORM) state, thereby making it read-only for the duration of its set retention period. Once the retention period expires, the file's WORM status can be extended or the file can be deleted.

Until a file is committed, it is modifiable and does not have WORM protection. Clients can commit a file to WORM status via NFS or SMB. The committed file must be within a WORM file system.

Committing WORM Files

When committing a file, you can optionally set its access time (atime). The file's atime should fall within the range of the minimum and maximum retention periods plus the current server time (when the commit occurs), as defined by the WORM policy that is attached to the WORM file system. If the atime is less than the current server time, or no atime is specified explicitly by the client, the default retention period defined in the WORM policy is used to commit the file. If the commit falls outside of the minimum and maximum retention periods, the commit will fail.

Once you have verified the file resides within the WORM file system, perform the following on the client:

- 1 (Optional) Set the file's atime by running the `touch -a -t <retention_time> <file>` command.
- 2 Remove the file's write permissions by running the `chmod -w <file>` command.

Extending WORM Retention

When the WORM retention period is committed, the WORM retention period for the file cannot be deleted or shortened, but can only be extended. To extend WORM protection for a file, run the `touch -a -t <retention_time> <file>` command.

WORM File Interaction with Front End Operations

There are three states for files integrated with WORM protection:

- **UNCOMMITTED:** The file has not ever run into WORM-committed status.
- **IN_PROTECTION:** The file currently is in WORM-committed status.
- **EXPIRED:** The file WORM retention date is expired.

The following table summarizes front end operations with different WORM states:

State	File Append/Overwrite/ Truncate	File Rename	File Delete	Directory Rename
UNCOMMITTED	Allow	Allow	Allow	Disallow
IN_PROTECTION	Disallow	Disallow	Disallow	Disallow
EXPIRED	Disallow	Disallow	Allow	Disallow

About Creating and Exporting File Systems

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

Creating and exporting a file system involves creating the file system and enabling protocol support for the file system. 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 or files will be accessed via SMB and NFS at the same time. This assumes that you already have a directory service, such as OpenLDAP or Active Directory set up.

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

- 1 Configure the directory service for protocol support. The directory service is managed through the **Settings > Security > Directory Service** panel.
 - 2 Create the data vip. Data vips are managed through the **Settings > Network > Subnets** panel.
 - 3 Create the file system. Optionally set a provisioned size for the file system.
 - 4 Define the NFS export rules.
 - 5 Enable the NFS protocol.
- **SMB with Active Directory.** FlashBlade offers SMB Native authentication mode. It is recommended for customers requiring Kerberos authentication.

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

- 1 Prepare Active Directory for multi-protocol support.

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

- 2 Configure the directory service for protocol support.

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

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

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

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

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

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

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

File System Size

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

When the amount of data written to the file system reaches a certain percentage of its quota, alerts are generated notifying administrators of the threshold reached. For example, Purity//FB generates a **WARNING** alert when the amount of data written to the file system reaches 90% and then 100% of its provisioned size. This behavior differs slightly if a hard limit is set on the file system. When the file system usage reaches 90% of its provisioned size, the system generates a **WARNING** alert. When the file system

usage reaches 100% of its provisioned size, the system generates a **CRITICAL** alert. See **File System Usage Limits** for more information. To view alert details, select **Health > Alerts**.

The file system size can be blank or set to any value between 0 and 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 Group Ownership

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

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

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

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

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

Figure 5.8 Viewing a File System's Details Panel

Details		
Created	2024-01-11 13:05:59	Quotas »
Group Ownership	creator	Default Group Quota -
Hard Limit	False	Default User Quota -
Promotion Status	promoted	
Source	-	
Writable	True	
Protocols »		
Enabled Protocols	-	
Access Control Style	shared	
Safeguard ACLs	True	

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.

The following file system snapshot limits are supported:

- Maximum total number of snapshots: 250,000
- Maximum number of daily snapshots: 700,000

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 5.2 Access Control Style Transitions

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

ACL Safeguarding

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

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

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

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

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

Group Policy Objects

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

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

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

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

Multi-Domain Trust

Multi-Domain Overview

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

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

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

Multi-Domain Trust Setup

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

 **Note:** 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.

Viewing and Managing File Locks

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

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

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

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

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

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

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

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

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

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

Fast Remove

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

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

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

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

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

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

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

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

For example,

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

The `.fast-remove` directory cannot be renamed, deleted, or moved into other directories.

The `.fast-remove` directory name is a reserved name, so you cannot create a regular directory named `.fast-remove`, regardless of the fast remove enabled/disabled status.

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

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

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

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

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

File Replication Failover, Reprotect, and Failback

⚠ Important! For simplified management of file system snapshots and replication for data protection, Pure Storage recommends the use of a single policy with multiple rules instead of multiple policies with different rules for file system replication requirements.

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.

⚠ Warning: During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

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.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems
- Promotion or demotion of file systems is a disruptive operation that changes the writability of the file system, so all ongoing SMB sessions and open file handles will be destroyed. The client must recreate its SMB sessions and file handles.

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.

To ensure your data protection SLAs are met and consistent for failover and failback operations, Pure Storage recommends that you **retain the existing snapshot policies** associated with the file system under File Replication for the source and target array.

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.

 **Note:** For Zero Move Tiering systems, created file systems will be assigned the performance storage class (S500X-S). Once created you can change the storage class to archive (S500X-A).

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 **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 In the **Special Directories** section, to enable the ability to quickly remove large directories using the fast remove feature, select **Fast Remove**. For more information about the fast remove feature, refer to [Fast Remove](#).
- 5 In the **Special Directories** section, to enable the ability to access stored snapshots in a snapshot directory, select **Snapshot**.
- 6 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 **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 supported for NFSv3 and NFSv4.1.
- 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.

7 If you wish to set additional quota and group ownership parameters, click **Additional Settings**.

- (Optional for NFS users and groups) In the **Default User Quota** field, specify the user quota size.
- (Optional for NFS users and groups) In the **Default Group Quota** field, specify the group quota size.
- Under **Group Ownership**, select **Creator** (default) or **Parent Directory**.
 - **Creator:** The owning group of new files and directories is the primary group of the user who creates them.
 - **Parent Directory:** The owning group of new files and directories is the primary group of the parent directory. New files and directories inherit the parent directory's group ID.

8 (Optional) Select the **WORM** checkbox if you wish to create a WORM file system. From the **Policy** list, select the WORM policy you wish to attach to the file system.

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 supported for NFSv3 and NFSv4.1.

- 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.

Setting a File System's Storage Class

 **Note:** Changing a file system's storage class from speed to archive is immediate. Changing a file system's storage class from archive to speed takes a minimum of 15 minutes.

File systems in FlashBlade Zero Move Tiering systems can belong to one of two storage classes: speed (S500X-S) or archive (S500X-A). When file systems are created, they are set to the speed storage class by default. A file system's storage class can be changed from speed to archive, or archive to speed, at any time.

To set a file system's storage class, perform the following:

- 1 Navigate to the **Storage > File Systems** screen.
- 2 Click **Storage Class** in the **File Systems** panel.
- 3 To set the storage class for a single file system:
 - a Click the **More Options** button and select either **Set S500X-A Storage** or **Set S500X-S Storage** in the same row of the file system whose storage class you wish to change. If the file system's storage class is S500X-A (archive), only the **Set S500X-S Storage** option is visible, and vice-versa. The **Set Storage Class** pop-up window appears.
 - b Click **Set**.
- 4 If you wish to set the S500X-A (archive) storage class for multiple file systems:
 - a In the **File Systems** panel header, click the **More Options** button and select **Set S500X-A Storage**. The **Set Storage Class to S500X-A** pop-up window appears.
 - b From the **File Systems with S500X-S Storage** list, select the file systems you wish to change to the S500X-A (archive) storage class. 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.
 - c Click **Set**.
- 5 If you wish to set the S500X-S (speed) storage class for multiple file systems:

- a In the **File Systems** panel header, click the **More Options** button and select **Set S500X-S Storage**. The **Set Storage Class to S500X-S** pop-up window appears.
- b From the **File Systems with S500X-A Storage** list, select the file systems you wish to change to the S500X-S (speed) storage class. 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.
- c Click **Set**.

Setting a File System's Storage Class Tiering Policy

File systems in FlashBlade Zero Move Tiering systems can belong to one of two storage classes: speed (S500X-S) or archive (S500X-A). Storage class tiering policies can be created to automatically manage moving the file system's data from hot (speed) to cold (archive), and vice versa.

To set a file system's storage class tiering policy, perform the following:

- 1 Navigate to the **Storage > File Systems** screen.
- 2 Click **Storage Class** in the **File Systems** panel.
- 3 Locate the file system to which you wish to set a storage class tiering policy and click **More Options > Set Tiering Policy** in the same row. The **Set Tiering Policy** pop-up window appears.
- 4 Select the policy to apply to the file system from the **Selected Policy** list.
- 5 Click **Save**.

Changing a File System's Storage Class Tiering Policy

To change a file system's storage class tiering policy, perform the following:

- 1 Navigate to the **Storage > File Systems** screen.
- 2 Click **Storage Class** in the **File Systems** panel.
- 3 Locate the file system whose storage class tiering policy you wish to change and click **More Options > Set Tiering Policy** in the same row. The **Set Tiering Policy** pop-up window appears.
- 4 Select the policy to apply to the file system from the **Selected Policy** list.
- 5 Click **Save**.

Removing a Storage Class Tiering Policy from a File System

To remove a storage class tiering policy from a file system, perform the following:

- 1 Navigate to the **Storage > File Systems** screen.
- 2 Click **Storage Class** in the **File Systems** panel.
- 3 Locate the file system from which you wish to remove a storage class tiering policy and click **More Options > Remove Tiering Policy** in the same row. The **Remove Tiering Policy** pop-up window appears.
- 4 Click **Remove**.

Renaming a File System

To rename a file system:

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

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

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

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 Storage 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 Storage 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

Warning: During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

Warning: Any writes performed since the last snapshot will be destroyed and if the file system does not have any snapshots, all data on it will be permanently deleted.

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 file system designations are restored.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Manual refresh for GUI based SMB applications may be required to get the latest view of demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems
- Promotion or demotion of file systems is a disruptive operation that changes the writability of the file system, so all ongoing SMB sessions and open file handles will be destroyed. The client must recreate its SMB sessions and file handles.

CAUTION: If Pure File SafeMode is enabled, contact Pure Storage 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**.

Applying a Legal Hold to a File System

Once a legal hold has been created, it can be applied to a file or path in a file system (see [Legal Holds](#)).

Note the following:

- A single legal hold can be applied to no more than 100 different paths.
- Adding a legal hold path to a file system with fast remove enabled is not allowed.
- A legal hold cannot be applied on a .snapshot directory.

- Replication of file systems with legal hold applied is not allowed.
- The same legal hold cannot be applied twice on a given inode, however applying a different legal hold to the same inode is allowed.

To apply a legal hold, perform the following:

- 1 From the **Storage > File Systems** page, click the file system to which you wish to apply to a legal hold from the **File Systems** panel. The details page for that file system appears.
- 2 Click the Add (+) button from the header of the **Legal Hold Entities** panel. The **Apply Legal Hold** pop-up window appears.
- 3 Select a legal hold from the **Legal Hold** list.
- 4 In the **Path** field, enter the file or path to which the legal hold will apply.
- 5 Click **Apply**.

Releasing a Legal Hold from a File System

When a file system no longer requires a legal hold, the legal hold can be released. Once released, the file system can be modified or deleted.

Note that a legal hold cannot be released if there is an ongoing task that has the same legal hold on the path.

To release a legal hold, perform the following:

- 1 From the **Storage > File Systems** page, click the file system from which you wish to release from a legal hold from the **File Systems** panel. The details page for that file system appears.
- 2 In the **Legal Hold Entities** panel, click the Remove (X) button in the same row of the legal hold you wish to release. The **Release Legal Hold** pop-up window appears.
- 3 Click **Release**.

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 Client, Group, and User Performance Statistics](#) for details).

Destroying a File System

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

To destroy a file system, perform the following:

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



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

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

- a Click the **More Options** button in the header of the **File Systems** panel and select **Destroy**. The **Destroy File Systems** pop-up window appears.
 - b From the **Destroyable File Systems** list, select the file systems you wish to destroy. You can optionally enter a partial name in the search box to filter the list. You can also optionally select all file systems by clicking the checkbox next to the search box. All selected file systems appear in the **Selected File Systems** list. If protocols are enabled on a file system, that file system will not appear in the **Destroyable File Systems** list.
 - c Click **Destroy**.

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

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

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

Recovering a File System

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

To recover a destroyed file system:

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

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

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

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

Eradicating a File System

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

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

To eradicate a destroyed file system:

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

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

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

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

Performing File Replication Failover and Reprotect

 **Important!** For simplified management of file system snapshots and replication for data protection, Pure Storage recommends the use of a single policy with multiple rules instead of multiple policies with different rules for file system replication requirements.

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 ensure your data protection SLAs are met and consistent for failover and failback operations, Pure Storage recommends that you **retain the existing snapshot policies** associated with the file system under File Replication for the source and target array.

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

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

Open Files and Sessions Management with the Microsoft Management Console

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

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

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

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

- View open files.
- Close open files.

The following actions can be executed for sessions using MMC:

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

For more information about MMC, refer to [Microsoft MMC documentation](#).

Setting Up the Microsoft Management Console

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

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

- 1 Click **Start > Run**. The **Run** pop-up window appears.
- 2 Enter **mmc** and click **OK**. The Console appears.
- 3 Click **File > Add/Remove Snap-in**. The **Add or Remove Snap-ins** pop-up window appears.
- 4 Select **Shared Folders** from the **Available snap-ins** panel and then click **Add**. The **Shared Folders** pop-up window appears.
- 5 Select **Another computer** and enter the name or IP address of the AD account computer name of the FlashBlade.
- 6 Under **View**, select **All** and then click **Finish**. The **Shared Folders** pop-up window closes.
- 7 In the **Add or Remove Snap-ins** pop-up window, click **OK**.

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

Managing Open Files with the Microsoft Management Console

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

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

Perform the following to manage open files:

- 1 Open MMC on the Windows client.
- 2 Select **Open files** under **Shared Folders**.
- 3 To close one or more files, right-click the file (or selection of files) and click **Close File**. When prompted for confirmation, click **Yes**.

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

Managing Sessions using the Microsoft Management Console

Once the **Shared Folders** Microsoft Management Console (MMC) snap-in has been set up, the **Sessions** folder displays all active SMB sessions. Active SMB sessions can be closed at any time.

⚠ Warning: Closing a session will force-close all open files in that session, which can result in loss of data.

Perform the following to view and close SMB sessions:

- 1 Open MMC on the Windows client.
- 2 Select **Sessions** under **Shared Folders**. All active SMB sessions are displayed.
- 3 Right-click the SMB sessions you wish to close and select **Close Session**.

Object Store

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

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 and/or deduplication 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. In versioned buckets, this includes objects covered by deleted markers. 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.
- **Role Count:** The total number of roles.

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. Ratio of the size of the written data (i.e. Used) versus the amount of space the data occupies after data compression (i.e. Physical) and/or deduplication.
- **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. In versioned buckets, this includes objects covered by deleted markers.
- **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 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**. 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**, **Public Access**, and

Storage Class 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

PSFBSAZRMAGHIIHAJDIBGODAJOAPEPHEPELOJJBCHGJ
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 and edit the buckets.

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.

Creating an Account for a Realm

To create an account for a realm, perform the following:

- 1 Navigate to **Storage > Object Store**. In the **Accounts** panel, click Add (+). The **Create Account** dialogue window appears.
- 2 In the **Scope Type** field, select realm.
- 3 In the **Scope Name** field, select the realm from the drop-down list.
- 4 In the **Name** field, enter a name for the account.
- 5 (Optional) In the **Quota Limit** field, enter a quota limit size in Kilobytes, Megabyte, Gigabytes, or Terabytes. For example, **500M**, **50G**, **1T**. Check the **Hard Limit** box if you want your quota limit to stop at the specified limit.
- 6 (Optional) In the **Bucket Default Quota Limit**, enter a bucket quota limit size in Kilobytes, Megabyte, Gigabytes, or Terabytes. For example, **500M**, **50G**, **1T**. Check the **Hard Limit** box if you want your bucket limit to stop at the specified limit.

- 7 Click **Create**.

Creating an Account Export

To export an object account, perform the following:

- 1 Navigate to **Storage > Object Store**. In the **Accounts** panel, click an account to export.
- 2 In the **Account Export** panel, click Add (+). The **Create Account Export** dialogue window appears.
- 3 In the **Scope Type** field, select realm.
- 4 In the **Scope Name** field, select the realm from the drop-down list.
- 5 In the **Account** field, select the realm and account from the drop-down list.
- 6 In the **Server** field, select the server from the drop-down list or enter a known server name.
- 7 In the **Policy** field, select the policy you want applied to the account.
- 8 In the **Enabled** field, click to enable now or leave unchecked to enable at a later time.
- 9 Click **Create**.

Creating a Role

Once an OIDC or SAML2 identity provider has been created, you must create a role with access policies. A maximum of 10,000 roles are allowed per FlashBlade.

When creating a role, you can define the allowed session duration and add one or more access policies, which will grant the S3 permissions in the account.

To create a role, perform the following:

- 1 From the **Storage > Object Store** page, click the account name in the **Accounts** panel under which you wish to create a role. The details page for the account appears.
- 2 Click the Add (+) button in the header of the **Roles** panel. The **Step 1: Create Role** pop-up window appears.
- 3 In the **Role Name** field, enter a name for the role.
- 4 In the **Maximum Session Duration** field, enter the maximum session duration from 1 hour to 48 hours. If no value is entered, the default is 1 hour.
- 5 Click **Create**. The **Step 2: Add Access Policies to role <name>** appears.

- 6 From the **Available Policies** panel, select the access policies you wish to assign to the role. At least one policy must be added to grant the role permissions in the account.

If a specific access policy does not exist in the list of available policies, click **Create a new Access Policy** (see [Creating Object Store Access Policies](#) for instructions).

- 7 Click **Add**.

Editing a Role

To edit a role, perform the following:

- 1 From the **Storage > Object Store** page, click the Edit button in the same row as the account in the **Accounts** panel you wish to edit. The **Edit Role** pop-up window appears.
- 2 In the **Maximum Session Duration** field, enter the maximum session duration from 1 hour to 48 hours. If no value is entered, the default is 1 hour.
- 3 Click **Save**.

Deleting a Role

 **Note:** Any access policies must be detached from the role before deleting the role.

To edit a role, perform the following:

- 1 From the **Storage > Object Store** page, click the Delete button in the same row as the account in the **Accounts** panel you wish to delete. The **Delete Role** pop-up window appears.
- 2 Click **Delete**.

Creating Trust Policy Rules

After creating a role, create a trust policy rule.

To create a trust policy rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account name in the **Accounts** panel under which you created a role. The details page for the account appears.
- 2 Click the created role from the **Roles** panel. The details page for the role appears.
- 3 Click the Add (+) button in the header of the **Trust Policy Rules** panel. The **Create Trust Policy Rule** pop-up window appears.

- 4 In the **Rule Name** field, enter a name for the rule.
- 5 Select whether the action (**Effect**) of the rule will be to **Allow** or **Deny**. Note that explicit Deny rules take precedence over Allow rules.
- 6 Click the Add (+) button in the **Principals** field. The **Add Principals** pop-up window appears.
 - a Select the principals you wish to add from the **Available Principals** panel.
 - b Click **Add**.
- 7 Click the Add (+) button in the **Actions** field. The **Add Actions** pop-up window appears.
 - a Select the actions you wish to add from the **Available Actions** panel.
 - b Click **Add**.
- 8 (Optional) If you wish to add a condition to the rule, click **Add condition**. The condition key must have a format prefix:value. Supported key prefixes are: aws, jwt, saml, sts.
 - a In the Key field, enter the key prefix.
 - b Select **All** or **Any**, the operator, and select **If Exists** under the **Operator** field.
 - c Click the Add (+) button under the **Values** field. Enter a comma separated value string and then click **OK**.
- 9 Click **Create**.

Editing a Trust Policy Rule

To edit a trust policy rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account name in the **Accounts** panel with the role whose trust policy you wish to edit. The details page for the account appears.
- 2 Click the role from the **Roles** panel. The details page for the role appears.
- 3 In the same row as the rule you wish to edit, click **More Options > Edit**. The **Edit Trust Policy Rule** pop-up window appears.
- 4 Select whether the action (**Effect**) of the rule will be to **Allow** or **Deny**. Note that explicit Deny rules take precedence over Allow rules.
- 5 In the **Principals** field, remove any existing principals. Click the Add (+) button in the **Principals** field. The **Add Principals** pop-up window appears.
 - a Select the principals you wish to add from the **Available Principals** panel.
 - b Click **Add**.

- 6 In the **Actions** field, remove any existing actions. Click the Add (+) button in the **Actions** field. The **Add Actions** pop-up window appears.
 - a Select the actions you wish to add from the **Available Actions** panel.
 - b Click **Add**.
- 7 In the Conditions field, remove any existing conditions by clicking the Delete button. If you wish to add a condition to the rule, click **Add condition**. The condition key must have a format prefix:value. Supported key prefixes are: aws, jwt, saml, sts.
 - a In the Key field, enter the key prefix.
 - b Select **All** or **Any**, the operator, and select **If Exists** under the **Operator** field.
 - c Click the Add (+) button under the **Values** field. Enter a comma separated value string and then click **OK**.
- 8 Click **Save**.

Removing a Trust Policy Rule

To remove a trust policy rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account name in the **Accounts** panel with the role whose trust policy rule you wish to remove. The details page for the account appears.
- 2 Click the role from the **Roles** panel. The details page for the role appears.
- 3 In the same row as the rule you wish to edit, click **More Options > Remove**. The **Remove Trust Policy Rule** pop-up window appears.
- 4 Click **Remove**.

Uploading a Trust Policy in IAM Format

To upload a trust policy in IAM format, perform the following:

- 1 From the **Storage > Object Store** page, click the account name in the **Accounts** panel with the role that you wish to upload a trust policy in IAM format. The details page for the account appears.
- 2 Click the role from the **Roles** panel. The details page for the role appears.
- 3 Click **More Options > Upload Trust Policy in IAM format** in the header of the **Trust Policy Rules** panel. The **Upload Trust Policy** pop-up window appears.

- 4 In the **Trust Policy** field, copy and paste the trust policy in IAM format. You can optionally load the policy from a file.
- 5 If you did not manually past the trust policy in IAM format in the previous field, you can load the policy from a file. In the **Load from file** field, click **Choose File** and select the file containing the trust policy in IAM format to upload.
- 6 Click **Upload**.

Downloading a Trust Policy in IAM Format

To download a trust policy in IAM format, perform the following:

- 1 From the **Storage > Object Store** page, click the account name in the **Accounts** panel with the role that you wish to download a trust policy in IAM format. The details page for the account appears.
- 2 Click the role from the **Roles** panel. The details page for the role appears.
- 3 Click **More Options > Download Trust Policy in IAM format** in the header of the **Trust Policy Rules** panel.

The trust policy is downloaded to your local Downloads location.

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 5.3 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 5.3 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.

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

- **Name:** The bucket name.
- **Account:** The account to which the bucket belongs.
- **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**, **Public Access**, and **Storage Class** 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. In versioned buckets, this includes objects covered by deleted markers.

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.

The **Storage Class** view includes the following:

- **Storage Class:** The bucket's storage class (**S500X-A**) or **S500X-S**.
- **Status:** The current storage class transfer status of the bucket (**Queued** or **In-Progress**). If no transfer action has been taken on the bucket, - is displayed.
- **Status Details:** A detailed status of the storage class transfer. If no transfer action has been taken on the bucket, - is displayed.

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.

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

Destroyed Buckets

By default, this panel is minimized. To expand the panel and display its contents, click the expand button at the bottom next to the **Destroyed** count at the bottom of the **Buckets** panel.

The **Destroyed Buckets** panel displays the following information:


- **Name:** The bucket name.
- **Account:** The account to which the bucket belongs.
- **Quota Limit:** The user-specified size limit for virtual data on the bucket in bytes.
- **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**).
- **Time Remaining:** The time remaining until the bucket is eradicated, which is the difference between the computed eradication timestamp and the current time. If a specific time is not listed then one of the following values is displayed:

- **(calculating):** Indicates that the destroyed bucket's remaining time to eradication is being calculated. This value will be replaced with a specific time once system computations are completed.
- **Infinite (legal hold):** Indicates that the bucket cannot be automatically eradicated and its remaining time to eradication is infinite.

Bucket Details

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

Figure 5.10 Viewing Bucket Details



Array File Systems **Object Store**

Object Store > testaccount > testbucket2

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

Bucket Replica Links General Backlog +

No bucket replica links found.

Lifecycle Rules +

No lifecycle rules found.

Details ⋮

Bucket Type	multi-site-writable	Eradication »
Hard Limit	False	Manual Eradication enabled
Retention Lock	unlocked	Eradication Delay 1d
Versioning	none	
Object Lock »		Public Access »
Default Retention	-	Block Public Access False
Default Retention Mode	-	Block New Public Policies False
Object Lock	disabled	Public Status public
Freeze Locked Objects	False	

Bucket Policies 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.

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.

Bucket Security

FlashBlade provides enhanced security of objects in the form of *Pure Object SafeMode*, *Object Lock*, *Retention Lock*, and *Per-Bucket Safemode*.

- Pure Object SafeMode is a data protection mechanism for S3 that is enabled by Pure Technical Services to provide array-wide object protection against accidental changes that can result in data loss.
- 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.
- Per-Bucket Safemode is similar to Pure Object SafeMode, but provides more granular protection at the bucket level.

Object Lock, Retention Lock, and Per-Bucket SafeMode can be enabled and configured when editing a bucket.

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

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 Storage Technical Services.

Object Lock

 **Note:** Replication to and from an Object Lock-enabled bucket is supported. For replication from an Object Lock-enabled bucket, the target bucket must be running Purity//FB 4.5.2 or higher and have Object Lock enabled.

 **Important!** When creating a replica link from an Object Lock-enabled bucket to an AWS/third party S3 target, Object Lock metadata is not replicated. It is highly recommended that a default retention is configured on the bucket.

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 be disabled.

If bucket versioning and Object Lock are required, Freeze Locked Objects must be disabled. Contact Pure Technical Services for assistance in setting this configuration. Note that if this configuration is set, it cannot be changed.

For use cases requiring the retention of protected objects, such as compliance with security regulations SEC Rule 17a-4 and SEC Rule 18a-6, an eradication mode can be set on buckets that permanently disables the eradication of non-empty buckets. The following two eradication mode options are available:

- **permission-based:** This is a bucket's default eradication mode. When enabled, a bucket is automatically eradicated only after the array-wide eradication delay time has expired. Manual eradication is controlled by per-bucket SafeMode or Retention Lock. Empty buckets may be eradicated.
- **retention-based:** Manual and automatic eradication of the bucket is disabled, unless the bucket is empty. Note that once set to **retention-based**, the bucket cannot be set back to **permission-based**.

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 (see [Object Lock](#) for more information about Object Lock settings):

- 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.

Retention Lock also affects bucket eradication.

- If enabled, non empty buckets can be destroyed, but cannot be manually eradicated. Eradication occurs after the array-wide timeout period expires.
- If enabled, empty buckets can be destroyed and eradicated at any time.

Per-Bucket SafeMode

Similar to Pure Object SafeMode, Per-Bucket SafeMode is a data protection mechanism designed to protect the system from accidental changes that can result in data loss. Where Object SafeMode provides array-wide protection, Per-Bucket SafeMode is more granular allowing SafeMode protection to be configured at the bucket level. Per-Bucket SafeMode's increased granularity allows you to:

- set different retention durations for different buckets.
- use Per-Bucket SafeMode on a subset of buckets in combination with having unprotected buckets, versioning-enabled buckets, or object-lock-enabled buckets.

When considering enabling Per-Bucket SafeMode on a bucket, note the following rules:

- Per-Bucket SafeMode can only be enabled on buckets where versioning is not enabled.
- Per-Bucket SafeMode and Object Lock cannot both be enabled on the same bucket.
- Per-Bucket SafeMode can only be enabled if Object SafeMode (array wide SafeMode) is disabled.
- Replication between Per-Bucket SafeMode-enabled source and target buckets is allowed if their retention periods are the same.

Additionally, the ratcheted (locked) and unlocked state of Retention Lock affects the ability to change Per-Bucket SafeMode configuration by the FlashBlade admin as described in the following table:

Retention Lock State	Per-Bucket SafeMode Configuration	Retention Behavior	Bucket Manual Eradication Behavior	Retention Lock Behavior
Unlocked	Can be turned off	Can be increased or decreased	Allowed	Can be be changed to ratcheted
Racheted (locked)	Cannot be turned off	Can only be increased	Not allowed unless empty	Can be changed to unlocked only by Pure Technical Services

Per-Bucket SafeMode is enabled by Pure Technical Services. Contact Pure Technical Services for additional information.

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 5.11 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 5.12 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 5.13 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 5.14 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

 **Note:** Replication to and from an Object Lock-enabled bucket is supported. For replication from an Object Lock-enabled bucket, the target bucket must be running Purity//FB 4.5.2 or higher and have Object Lock enabled.

 **Important!** When creating a replica link from an Object Lock-enabled bucket to an AWS/third party S3 target, Object Lock metadata is not replicated. It is highly recommended that a default retention is configured on the 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 In the **Versioning** field, select if you wish versioning to be **enabled** or **suspended**.
- 6 To edit the bucket's retention security settings, 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.
 **Note:** The **safemode** option under **Default Retention Mode** is only available through activation by Pure Technical Services. Contact Pure Technical Services for additional information.
 - d Enable the **Default Retention Mode** by selecting **none**, **governance**, **compliance**, or **safemode**.
 **Note:** Eradication Mode and Retention Lock require Object Lock to be enabled.
 - e Set the **Eradication Mode** as either **permission-based** (default) or **retention-based**. If set to **permission-based**, the bucket is automatically eradicated after the array-wide eradication delay time expires. If set to **retention-based**, manual eradication is blocked. Note that if set to **retention-based**, the bucket's eradication mode cannot be set back to **permission-based**.
 - f 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 5.15 Viewing the Lifecycle Rules List

Lifecycle Rules 1-2 of 2 +								
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.

Setting a Bucket's Storage Class

 **Note:** Changing a bucket's storage class from speed to archive is immediate. Changing a bucket's storage class from archive to speed requires more time; for buckets larger than 900 GB a maximum 15 minute delay applies. For buckets smaller than 900 GB the penalty time scales linearly with the bucket size.

Buckets in FlashBlade Zero Move Tiering systems can belong to one of two storage classes: speed (S500X-S) or archive (S500X-A). When buckets are created, they are set to the speed storage class by default. A bucket's storage class can be changed from speed to archive, or archive to speed, at any time.

To set a bucket's storage class, perform the following:

- 1 Navigate to the **Storage > Object Store** screen.
- 2 Click **Storage Class** in the **Buckets** panel.
- 3 To set the storage class for a single bucket:
 - a Click the **More Options** button and select either **Set S500X-A Storage** or **Set S500X-S Storage** in the same row of the bucket whose storage class you wish to change. If the bucket's storage class is S500X-A (archive), only the **Set S500X-S Storage** option is visible, and vice-versa. The **Set Storage Class** pop-up window appears.
 - b Click **Set**.
- 4 If you wish to set the S500X-A (archive) storage class for multiple buckets:
 - a In the **Buckets** panel header, click the **More Options** button and select **Set S500X-A Storage**. The **Set Storage Class to S500X-A** pop-up window appears.
 - b From the **Buckets with S500X-S Storage** list, select the buckets you wish to change to the S500X-A (archive) storage class. 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 buckets appear in the **Selected Buckets** list.

c Click **Set**.

5 If you wish to set the S500X-S (speed) storage class for multiple buckets:

- a In the **Buckets** panel header, click the **More Options** button and select **Set S500X-S Storage**. The **Set Storage Class to S500X-S** pop-up window appears.
- b From the **Buckets with S500X-A Storage** list, select the buckets you wish to change to the S500X-S (speed) storage class. 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 buckets appear in the **Selected Buckets** list.

c Click **Set**.

Canceling an In-progress Storage Class Transition

In FlashBlade Zero Move Tiering systems, in-progress storage class transitions for buckets can be canceled if their status in the **Buckets** panel shows that they are **Queued**.

To cancel a bucket's in-progress storage class transition, perform the following:

- 1 Navigate to the **Storage > Object Store** screen.
- 2 Click **Storage Class** in the **Buckets** panel.
- 3 Locate the bucket whose storage class transition status shows **Queued**.
- 4 Click the **More Options** button and select either **Set S500X-A Storage** or **Set S500X-S Storage** in that bucket's row. The **Set Storage Class** pop-up window appears.
- 5 Confirm that you wish to cancel the transition by selecting **Cancel In Progress Storage Class Transition**.
- 6 Click **Set**.

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.
- **Legal Holds:** Displays and allows you to manage all legal hold instances.

Figure 6.1 Purity//FB Protection Page



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.

File System Snapshot Limits

The following file system snapshot limits are supported:

- Maximum total number of snapshots: 250,000
- Maximum number of daily snapshots: 700,000

Viewing File System 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 Storage 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**.

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 Storage 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 system replica links are required for replicating file system snapshots between two or more connected FlashBlades. There are two types of file snapshot replication, traditional and rapid file replication.

 **Note:** Rapid file replication is not suitable for data protection use cases because it only replicates file data on demand, even though file system metadata replicates periodically. For file system data protection with replication use cases, avoid using partial replication. For business continuity and disaster recovery requirements, all the file data must be replicated to the target site and is protected from data loss with file system replication.

Traditional file system replication involves the replication of entire snapshots from the source to a target. Rapid file replication is a variation of file system replication where, instead of the replication of a full snapshot, file snapshot metadata and empty files are replicated to the target. File data is pulled from the source during file read attempts by the user, allowing for more efficient use of network bandwidth and storage space as only the necessary data is transferred on-demand.

The **File Replication** tab is used to view and manage file replica links, and 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.
- **Link Type:** The type of file replica link, either **full-replica** or **partial-replica**.

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* in the *FlashBlade Administration Guide* for instructions.

File system replica links are used to configure file system replication between two connected FlashBlade arrays. Two types of file system replica links can be created: full-replica links and partial-replica links. A full-replica link replicates an entire snapshot from the source to target. A partial-replica link is used for rapid file replication. Instead of the replication of an entire snapshot, metadata and empty files are replicated from the source to target, allowing for more efficient use of network bandwidth and storage space. A maximum of 3,500 file system replica links per array connection are supported.

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 10,000 file systems can have replica links.
- A file system can have a maximum of 10,000 outgoing replica links and a maximum of 10,000 incoming replica links.
- A maximum of 2,000 fan-out/fan-in replication links.
- A maximum of 1,000 source file systems with partial-replica links.
- A maximum of 1,000 target partial-replica links.

 **Warning:** During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Manual refresh for GUI based SMB applications may be required to get the latest view of demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems

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 **Partial Replica** if you wish to create a partial file system replica link for use with rapid file replication. If selected, metadata will be periodically replicated and file data will be replicated on demand.
 - e (Optional) Click **Connect FlashBlade Array** if you did not previously connect an array. See *Connecting Arrays* in the *FlashBlade Administration Guide* for instructions.
 - f 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**.

Forcing Synchronization for Partial File System Replica Links

Forced synchronization of partial file system replica links transfers metadata immediately if there is no on-going transfer on the link.

To force the synchronization of a partial file system replica link, perform the following:

- 1 From the **Protection > File Replication** page, locate the partial file system replica link for which you wish to force synchronization in the **File System Replica Links** list.
- 2 In the same row as the partial file system replica link, click **More Options > Replicate Now**.

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**.

Deleting Partial File System Replica Links

Partial file system replica links can only be deleted by destroying and eradicating the replica file system on the target array.

See [Destroying a File System](#) and [Eradicating a File System](#).

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.

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 6.2 Object ReplicationHeader




Figure 6.3 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.

 **Note:** Replication to and from an Object Lock-enabled bucket is supported. For replication from an Object Lock-enabled bucket, the target bucket must be running Purity//FB 4.5.2 or higher and have Object Lock enabled.

 **Important!** When creating a replica link from an Object Lock-enabled bucket to an AWS/third party S3 target, Object Lock metadata is not replicated. It is highly recommended that a default retention is configured on the bucket.

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**.

Legal Holds

A legal hold (or litigation hold) is a process used to preserve relevant information when litigation is anticipated or pending. A legal hold prevents the intentional or unintentional destruction, alteration, or loss of data that may be needed for legal proceedings. Files or directories put into a legal hold cannot be deleted, destroyed, or altered for the duration of the legal hold. FlashBlade allows the creation of legal holds and the application of legal holds to files and directories.

Note the following legal hold limits:

- Up to 10,000 legal holds can be created system-wide.
- A single legal hold can be applied to no more than 100 different paths.
- A total of 1 million legal hold-held entities are allowed.
- Up to 1,000 legal hold entities in the apply/release state are allowed.

Note the following restrictions:

- A legal hold cannot be applied on a .snapshot directory.
- Restoring a file system with one or more legal hold paths is not allowed.
- Demoting a file system with one or more legal hold paths is not allowed.
- Enabling Fast Remove for a file system with one or more legal hold paths is not allowed.
- Adding a legal hold path to a file system with Fast Remove enabled is not allowed.
- Destroying a file system with one or more legal hold paths is not allowed.
- Replication of file systems with legal hold applied is not allowed.
- A legal hold cannot be deleted while it is applied to any path.
- The same legal hold cannot be applied twice on a given inode, however applying a different legal hold to the same inode is allowed.
- A legal hold cannot be released if there is an ongoing task that has the same legal hold on the path.

- A legal hold cannot be applied to a file's parent directory if a legal hold is already applied to that file.
- Files cannot be added or moved under directories with legal holds applied.
- Legal hold actions (create, update, delete, apply, release, etc.) can only be triggered by array admins.

 **Note:** Legal holds cannot be attached to parents or children of directories that already have legal holds attached. Removing an reapplying legal holds (moving a legal hold up and down a directory tree) results in a brief period of non-protection.

 **Note:** Files protected by legal hold can be locally overwritten on file systems mounted using NFS on MacOS. Locally overwritten files will not change any content persisted on the server side, and remounting will refresh the view. However, it is not recommended to mount using NFS on MacOS.

Viewing Legal Holds

The **Protection > Legal Holds** tab provides the **Legal Holds** panel, which is used to view, create, and manage legal holds. The **Legal Holds** panel displays the following:

- **Name:** The name of the legal hold.
- **Description:** A description of the legal hold (entered during legal hold creation).

Viewing Legal Hold Details

Clicking a legal hold from the Legal Holds panel opens the details page for that legal hold. A legal holds details page provides a **Details** panel and a **Legal Hold Entities** panel.

The **Details** panel provides summary information of the legal hold.

The **Legal Holds Entities** panel displays all file system entities to which the legal hold is applied.

- **File System:** The file system name.
- **Path:** The path to which the legal hold is applied.
- **Status:** The current legal hold status.
 - **applying:** The legal hold task has been submitted and is in process of being applied.
 - **applied:** The legal hold task has been completed and is applied.
 - **releasing:** The legal hold removal task has been submitted and is in process of releasing.

This panel also allows you to apply and manage legal holds.

Creating a Legal Hold

Note the following legal hold limits:

- Up to 10,000 legal holds can be created system-wide.
- A single legal hold can be applied to no more than 100 different paths.
- A total of 1 million legal hold-held entities are allowed.

To create a legal hold, perform the following:

- 1 From the **Protection > Legal Holds** page, click the **Add (+)** button in the heading of the **Legal Holds** panel. The **Create Legal Hold** pop-up window appears.
- 2 Enter a name for the legal hold. The name must be between 1 and 63 alphanumeric characters, including underscores (_) and dashes (-), and must begin and end with a letter or number. At least one letter must be included.
- 3 (Optional) Enter a description for the legal hold. The description can be up to 127 characters, including spaces, with no new lines.
- 4 Click **Create**.

Editing a Legal Hold

To edit a legal hold, perform the following:

- 1 From the **Protection > Legal Holds** page, locate the legal hold you wish to edit from the **Legal Holds** panel.
- 2 Click the Edit button in the same row as the legal hold you wish to edit. The **Edit Legal Hold** pop-up window appears.
- 3 (Optional) Enter a new description for the legal hold. The description can be up to 127 characters, including spaces, with no new lines.
- 4 Click **Save**.

Applying a Legal Hold

Once a legal hold has been created, it can be applied to a file or path in a file system.

Note the following:

- A single legal hold can be applied to no more than 100 different paths.
- Adding a legal hold path to a file system with fast remove enabled is not allowed.
- A legal hold cannot be applied on a .snapshot directory.
- Replication of file systems with legal hold applied is not allowed.
- The same legal hold cannot be applied twice on a given inode, however applying a different legal hold to the same inode is allowed.

To apply a legal hold, perform the following:

- 1 From the **Protection > Legal Holds** page, click the legal hold you wish to apply to a file system from the **Legal Holds** panel. The details page for that legal hold appears.
- 2 Click the Add (+) button from the header of the **Legal Hold Entities** panel. The **Apply Legal Hold** pop-up window appears.
- 3 Select a file system from the **File System** list.
- 4 In the **Path** field, enter the file or path to which the legal hold will apply.
- 5 Click **Apply**.

Releasing a Legal Hold

When a file system no longer requires a legal hold, the legal hold can be released. Once released, the file system can be modified or deleted.

Note that a legal hold cannot be released if there is an ongoing task that has the same legal hold on the path.

To release a legal hold, perform the following:

- 1 From the **Protection > Legal Holds** page, click the legal hold you wish to release from a file system from the **Legal Holds** panel. The details page for that legal hold appears.
- 2 In the **Legal Holds** panel, click the Remove (X) button in the same row of the file system from which you wish to release from the legal hold. The **Release Legal Hold** pop-up window appears.
- 3 Click **Release**.

Deleting a Legal Hold

When a legal hold is no longer needed, it can be deleted.

Note that a legal hold cannot be deleted while it is applied to any path, including the snapshots.

To delete a legal hold, perform the following:

- 1 From the **Protection > Legal Holds** page, locate the legal hold you wish to delete from the **Legal Holds** panel.
- 2 Click the Delete button in the same row as the legal hold you wish to delete. The **Delete Legal Hold** pop-up window appears.
- 3 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 (full or partial) 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 (a *full link*) between the source file system and a target file system. A file system replica link is the connection between a source 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.

Rapid File Replication

Rapid file replication is a variation of traditional file replication. Instead of the replication of a full snapshot, with rapid file replication metadata is periodically replicated from the source to the target file system and file data is replicated on demand, allowing for more efficient use of network bandwidth and storage space. Like traditional file replication, rapid file replication requires two connected arrays and a file system replica link (a *partial link*) between the source and target file systems.

- Attaching policies to partial file system links is not allowed.

- Partial file system replica links can only be deleted by destroying and eradicating the replica file system on the target array.
- File systems with a partial file system replica link cannot be promoted or restored.

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.

Encrypted File System Replication

Encrypted file system replication data traffic between local and target arrays is supported if both the source and target FlashBlade arrays are running Purity//FB 4.5.1 or higher. If the source and target are running incompatible versions of Purity//FB, existing file system replication links and traffic (unencrypted replication) will not be interrupted. However new file system replication links cannot be created (for encrypted replication) until both the source and target are running at least Purity//FB 4.5.1.

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 on the target array, replica file systems can be promoted, but cannot be demoted back. Contact Pure Storage Technical Services to 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.

! Warning: During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

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.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Manual refresh for GUI based SMB applications may be required to get the latest view of demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems
- Promotion or demotion of file systems is a disruptive operation that changes the writability of the file system, so all ongoing SMB sessions and open file handles will be destroyed. The client must recreate its SMB sessions and file handles.

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:** Each array connection has three paths. Two paths are for destination port 8117 and one path is for destination port 80 for disabled encryption or 443 for enabled encryption. For object replication, these three paths must all be in the the healthy (connected) state.

 **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* in the *FlashBlade Administration Guide* 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 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* in the *FlashBlade Administration Guide* for instructions.

File system replica links are used to configure file system replication between two connected FlashBlade arrays. Two types of file system replica links can be created: full-replica links and partial-replica links. A full-replica link replicates an entire snapshot from the source to target. A partial-replica link is used for rapid file replication. Instead of the replication of an entire snapshot, metadata and empty files are replicated from the source to target, allowing for more efficient use of network bandwidth and storage space. A maximum of 3,500 file system replica links per array connection are supported.

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 10,000 file systems can have replica links.
- A file system can have a maximum of 10,000 outgoing replica links and a maximum of 10,000 incoming replica links.
- A maximum of 2,000 fan-out/fan-in replication links.
- A maximum of 1,000 source file systems with partial-replica links.
- A maximum of 1,000 target partial-replica links.

Warning: During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Manual refresh for GUI based SMB applications may be required to get the latest view of demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems

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 **Partial Replica** if you wish to create a partial file system replica link for use with rapid file replication. If selected, metadata will be periodically replicated and file data will be replicated on demand.
 - e (Optional) Click **Connect FlashBlade Array** if you did not previously connect an array. See *Connecting Arrays* in the *FlashBlade Administration Guide* for instructions.

- # Setting Up File System Replication Using the CLI

Creating Replication Network Interfaces

```
# 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: Each array connection has three paths. Two paths are for destination port 8117 and one path is for destination port 80 for disabled encryption or 443 for enabled encryption. For object replication, these three paths must all be in the the healthy (connected) state.

📋 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_pol1
```

- 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_pol1**.

```
# purepolicy rule add --at 12am --every 1d --keep-for 1d snap_pol1
```

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_pol1  
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 Storage Technical Services to promote and demote file systems.

 **Warning:** During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

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.

 **Warning:** When a file system is demoted, any writes performed since the last replication will be discarded and if the file system does not have any snapshots, all data on it will be permanently deleted.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems
- Promotion or demotion of file systems is a disruptive operation that changes the writability of the file system, so all ongoing SMB sessions and open file handles will be destroyed. The client must recreate its SMB sessions and file handles.

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 Storage 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 Storage Technical Services to promote and demote file systems.

 **Warning:** During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

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.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems
- Promotion or demotion of file systems is a disruptive operation that changes the writability of the file system, so all ongoing SMB sessions and open file handles will be destroyed. The client must recreate its SMB sessions and file handles.

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_pol1`.
 - The policy attached to the replica link is `rl_pol1`.
- 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_pol1` that attached to the local file system `fs1`

```
# purepolicy snapshot remove fs_snap_pol1
```

- 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 Storage 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 Storage Technical Services to promote and demote file systems.

 **Warning:** During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

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 the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems

- Promotion or demotion of file systems is a disruptive operation that changes the writability of the file system, so all ongoing SMB sessions and open file handles will be destroyed. The client must recreate its SMB sessions and file handles.

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.

- # Performing File System Fan Out Replication Failback Using the GUI

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

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* in the *FlashBlade Administration Guide*. 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 Storage Technical Services to promote and demote file systems.

 **Warning:** During file system replication failover, the local file system is demoted and the target file system is promoted. Use of the HTTP protocol on a demoted file system will result in a read-only file system error and is not recommended.

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.

Note the following SMB behavior for demoted file systems:

- Continuous Availability (CA) is not supported on demoted file systems
- Manual refresh for GUI based SMB applications may be required to get the latest view of demoted file systems
- Leases and OpLocks will not be given for any files on demoted file systems
- Promotion or demotion of file systems is a disruptive operation that changes the writability of the file system, so all ongoing SMB sessions and open file handles will be destroyed. The client must recreate its SMB sessions and file handles.

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_pol1`.
 - The policy attached to the replica link is `rl_pol1`.
- 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_pol1
array3# purepolicy snapshot disable fs_snap_pol1
```

- 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 Storage 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, $A \rightarrow B$ and $A \rightarrow C$. 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, $A \rightarrow B \rightarrow C$.

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.

 **Note:** Each array connection has three paths. Two paths are for destination port 8117 and one path is for destination port 80 for disabled encryption or 443 for enabled encryption. For object replication, these three paths must all be in the the healthy (connected) state.

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:** Each array connection has three paths. Two paths are for destination port 8117 and one path is for destination port 80 for disabled encryption or 443 for enabled encryption. For object replication, these three paths must all be in the healthy (connected) state.

 **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* in the *FlashBlade Administration Guide* 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 a 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: Each array connection has three paths. Two paths are for destination port 8117 and one path is for destination port 80 for disabled encryption or 443 for enabled encryption. For object replication, these three paths must all be in the healthy (connected) state.

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

- **Data Eviction:** Displays and allows you to manage data eviction policies. See [Data Eviction Policies](#) for additional information.
- **File System Audit:** Displays and allows you to manage file system audit policies and syslog log targets. See [File System Audit](#) for additional information.
- **NFS Export:** Displays and allows you to manage NFS export policies and NFS export rules. See [NFS Export Policies](#) for additional information.
- **Object Store Access Policies:** Displays and allows you to manage object store access policies. See [Object Store Access Policies](#) for additional information.
- **Object Store Audit:** Displays and allows you to manage object store audit policies and audit policy log targets. See [Object Store Audit](#) for additional information.
- **QoS:** Displays and allows you to manage Quality of Service policies. See [File System Quality of Service Policies](#) and [Object Store Quality of Service Policies](#) for additional information.
- **S3 Export:** Displays and allows you to manage S3 export policies. See [S3 Export 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.
- **Storage Class Tiering:** Displays and allows you to manage file system storage class tiering policies. See [Storage Class Tiering Policies](#) for additional information.
- **WORM:** Displays and allows you to manage WORM file system policies. See [File System WORM Policies](#) for additional information.

- **Object Store Access:** Displays and allows you to manage object store policies. See [Object Store Access Policies](#) for additional information.
- **Management Access Policies:** Displays and allows you to manage the array's management access policies. See [Management Access Policies](#) for additional information.
- **Password:** Displays and allows you to manage the array's password policy. See [Password Policies](#) for additional information.
- **SSH Certificate Authority:** Displays and allows you to manage the SSH login rules for users, including local and remote users on the array. See [SSH Certificate Authority Policies](#) for additional information.
- **Network Access:** Displays and allows you to manage network access policies and network access rules. See [Network Access Policies](#) for additional information.
- **TLS:** Displays and allows you to manage TLS policies on the system. See [TLS Policies](#) for additional information.

Data Eviction Policies

Rapid file replication allows file data to be retrieved from a remote source array on demand when file read requests are received. Once retrieved, file data is retained until the data is deleted by the source or the entire rapid replica is deleted, which can introduce space management challenges. FlashBlade provides the ability to create space-based data eviction policies allowing data to be deleted (evicted) independently of file lifecycle once the user-configured space threshold is reached. Purity//FB periodically checks if eviction needs to be run for each file system with an attached data eviction policy. This can result in a delay (maximum 10 minute) until the data eviction takes effect.

Data eviction policies are viewed and managed from the **Policies > Storage > Data Eviction > Data Eviction Policies** panel. The **Data Eviction Policies** panel displays the following information:

- **Name:** The name of the data eviction policy.
- **Enabled:** Whether the policy is enabled (**True**) or disabled (**False**).
- **Keep Size:** The configured physical space threshold that when reached, data eviction occurs.

Clicking a data eviction policy opens that policy's details page, which provides **Details** and **Attached File Systems** panels.

- The **Details** panel provides summary information about the policy.
- The **Attached File Systems** panel allows you to attach and remove file systems to the policy.

Creating a Data Eviction Policy

To create a data eviction policy, perform the following:

- 1 From the **Policies > Data Eviction** page, click the **Create (+)** button in the heading of the **Data Eviction Policies** panel. The **Create Data Eviction Policy** pop-up window appears.
- 2 Select if the policy is for an array or realm (for a multi-tenancy environment).
- 3 Enter the policy name.
- 4 (Optional) By default, the policy will be enabled upon creation. To disable the policy, deselect the **Enabled** checkbox.
- 5 Enter the data threshold size in the **Keep Size** field. This is the physical space usage threshold at which data eviction occurs. The minimum size is 10 GB.
- 6 Click **Create** to complete adding the new audit policy.

Editing a Data Eviction Policy

To edit a data eviction policy, perform the following:

- 1 From the **Policies > Data Eviction** page, click **More Options > Edit** in the same row of the policy you wish to edit in the **Data Eviction Policies** panel. The **Edit Data Eviction Policy** pop-up window appears.
- 2 Select or deselect the **Enabled** checkbox to enable or disable the policy.
- 3 Enter the data threshold size in the **Keep Size** field. This is the physical space usage threshold at which data eviction occurs. The minimum size is 10 GB.
- 4 Click **Save** to complete adding the new audit policy.

Renaming a Data Eviction Policy

To rename a data eviction policy, perform the following:

- 1 From the **Policies > Data Eviction** page, click **More Options > Rename** in the same row of the policy you wish to rename in the **Data Eviction Policies** panel. The **Rename Data Eviction Policy** pop-up window appears.

- 2 Enter the policy's new name.
- 3 Click **Rename** to complete adding the new audit policy.

Enabling and Disabling Data Eviction Policies

Unless explicitly set to disabled, data eviction policies are enabled by default when they are created.

Data eviction policies can be disabled or enabled at any time after their creation.

 **Note:** When disabling a policy, the policy will stop evicting data from attached file systems. When the policy is re-enabled, data eviction will resume.

To enable or disable a data eviction policy, perform the following:

- 1 From the **Policies > Data Eviction > Data Eviction Policies** panel, locate the policy you wish to enable or disable.
- 2 To enable the policy:
 - a Click **More Options > Enable** at the end of the policy's row.
 - b Click **Enable** when prompted for confirmation.
- 3 To disable the policy:
 - a Click the **More Options > Disable** at the end of the policy's row.
 - b Click **Disable** when prompted for confirmation.

Deleting Data Eviction Policies

Before deleting a data eviction policy, ensure that no file systems are attached.

- 1 To delete a single policy:
 - a From the **Policies > Data Eviction > Data Eviction Policies** panel, click **More Options > Delete** at the end of the policy's row. The **Delete Data Eviction Policy** pop-up window appears.
 - b Click **Delete**.
- 2 To delete multiple policies:
 - a From the **Policies > Data Eviction > Data Eviction Policies** panel, click the **More Options > Delete** in the heading of the **Data Eviction Policies Policies** panel. The **Delete Data Eviction Policies** pop-up window appears.

- 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**.

Attaching a Data Eviction Policy to a File System from the Data Eviction Policy Details Page

To attach a data eviction policy to a file system, perform the following:

- 1 From the **Policies > Data Eviction > Data Eviction Policies** panel, click the policy you wish to attach to a file system. The policy's details page appears.
- 2 In the **Attached File System** panel, click **More Options > Add File Systems**. The **Add File Systems to Data Eviction Policy** pop-up window appears.
- 3 From the **Available File Systems** panel, select one or more file systems to add to the policy.
- 4 Click **Add**.

Removing a Data Eviction Policy from a File System from the Data Eviction Policy Details Page

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

- 1 From the **Policies > Data Eviction > Data Eviction Policies** panel, click the policy you wish to remove from a file system. The policy's details page appears.
- 2 In the **Attached File System** panel, click **More Options > Remove**. The **Remove File Systems from Data Eviction Policy** pop-up window appears.
- 3 From the **Available File Systems** panel, select one or more file systems to remove from the policy.
- 4 Click **Remove**.

Attaching a Data Eviction Policy to a File System from the File System Details Page

To attach a data eviction policy to a file system, perform the following:

- 1 Navigate to **Storage > File Systems**. In the **File Systems** panel, click the file system to which you wish to attach a data eviction policy. The details page for the file system appears.

- 2 In the **Attached Policies** panel, click **More Options > Set Data Eviction Policy**. The **Set Data Eviction Policy** pop-up window appears.
- 3 From the **Selected Policy** list, select the policy you wish to attach to the file system.
- 4 Click **Save**.

Removing a Data Eviction Policy from a File System from the File System Details Page

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

- 1 Navigate to **Storage > File Systems**. In the **File Systems** panel, click the file system from which you wish to remove a data eviction policy. The details page for the file system appears.
- 2 In the **Attached Policies** panel, click the **Remove** button in the same row as the policy you wish to remove from the file system. The **Remove Policy from File System** pop-up window appears.
- 3 Click **Remove**.

File System Audit

File system audit policies are comprised of log targets. A log target works as an endpoint where the audit logs for file system access operations will get stored. Attach an audit policy to any file system and audit logs for data access operations on that file system will be stored in the customer's designated log target.

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

 **Note:** Currently only one audit policy attachment is supported on a file system. To change the audit policy on a file system, you must first detach an audit policy and then attach the new audit policy. During this detach and attach policy operation, some audit events might get lost on the source file system. To avoid this, log targets on the attached audit policy can be edited to achieve the same requirement.

Users with GPO (Group Policy Object) privileges can set up the System Access Control Lists (SACLs) to specify the operations on the file system that need to be audited.

A group policy allows administrators to define security policies for users. A Group Policy Object (GPO) is a collection of group policy settings. When joining an Active Directory (AD) domain, FlashBlade applies two specific group policies as defined by the GPO within the AD domain: backup operator and

the privilege to set system access control lists (SACLs). GPOs, group policies and their privileges are configured on the client.

- **backup operator:** Allows select users full read access to file systems in order to create backup copies of them. This grants access to a file regardless of if the user has the appropriate Access Control Element (ACE) set in the Access Control List (ACL). Full read/write access for this policy requires the SeBackupPrivilege and SeRestorePrivilege privileges. Note that this provides identical permissions to the `fs.smb_backup_operators` tunable (tunable enabled by Pure Technical Services).
- **privilege to set up SACLs:** Allows select users to perform security filtering based on SACLs attached to file system folders and files that are served by FlashBlade. Access is granted with the SeSecurityPrivilege privilege.

 **Note:** You must configure the log targets before creating an audit policy with it.

The following are currently supported:

- Only syslog server is supported as a log target.
- Only one log target of any one type is supported within an audit policy.
- Only one policy for a single file system is supported.
- 100 audit policies are supported per FlashBlade cluster.
- 1,000 file systems can be audited per FlashBlade cluster.
- 3 syslog servers can be configured for auditing.

Audit logs contain common actions such as create, open, read, write, set_info, query_info, and close/delete. Below are the common key terms. Note that frequently occurring operations such as read and write are rate-limited to avoid overload of events. For a given file system that was open and closed, read and writes are reported once.

Table 8.1 Audit Log Key

Key	Detail
pf	Platform (FlashBlade)
v	Version

Table 8.1 Audit Log Key (continued)

Key	Detail
eid	unique/sequence number for each record
time	Time stamp
us_sid	User SID
user	Username
client_ip	Client IP
proto	Protocol (SMB)
inodee	Inode number
fs	File system name
path	Full path name
op_code	Protocol agnostic operation specific unique code
op_name	Protocol agnostic operation name
result	success/failure

From the **Policies > File System Audit** page, the **File System Audit Policies** and **Syslog Log Targets** panels allow you to view and manage file system audit polices and syslog log targets.

The **File System Audit Policies** panel lists the names of all created audit policies and if those policies are enabled or disabled. Clicking an audit policy displays that policy's details page. This page allows you to manage attached file systems and attached audit policy log targets.

The **Syslog Log Targets** panel lists the names of the syslog servers, their URIs, and service types. Clicking an audit policy log target displays that log target's details page.

Creating a File System Audit Policy

Prior to creating a file system audit policy, a syslog log target must have been previously created. Multiple syslog log targets are not supported.

 **Note:** Multiple syslog servers with *data-audit* services and the same *IP/hostname* might lead to a scenario where audit logs will not be captured at the correct syslog server. To avoid this issue, do not create multiple syslog servers with same *IP/hostname* and same services. Removing duplicate syslog servers will also resolve this issue.

To create a file system audit policy, perform the following:

- 1 From the **Policies > File Systems Audit** page, click the **Create (+)** button in the heading of the **File System Audit Policies** panel. The **Create File System Audit 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.

 **Note:** Multiple log targets are not supported in a single audit policy.

- 4 Click the **Add** button in the **Log Targets** field. The **Add Log Targets** pop-up window appears.
- 5 Select a log target from the **Available Log Targets** panel and click **Add**.
- 6 Click **Add** to complete adding the new audit policy.

Once created, the policy will appear in the **File System Audit Policies** panel.

Editing a File System Audit Policy

To edit an existing file system audit policy, perform the following:

- 1 From the **Policies > File Systems Audit** page, click **More Options > Edit** in the same row as the audit policy you wish to edit in the **File System Audit Policies** panel. The **Edit File System Audit Policy** pop-up window appears.

- 2 Enable or disable the policy.

 **Note:** The file system audit policy must have a log target attached. Multiple log targets are not supported in a single audit policy.

- 3 Add or remove log targets to the file system audit policy.
 - a To remove a log target, click the **Remove (X)** button next to the log target you wish to remove.

- b To add a log target, click the **Add** button in the **Log Targets** field. The **Add Log Targets** pop-up window appears.
 - c Select a log target from the **Available Log Targets** panel and click **Add**.
- 4 Click **Save** to complete editing the audit policy.

Attaching a File System Audit Policy to a File System

 **Note:** File systems with NFS/HTTP protocol enabled cannot be attached to audit policies.

File systems can be attached to policies from either the **Policies > File System Audit** page or from the **Storage > File Systems** page.

- 1 Attach an audit policy to a file system from the **Policies > File System Audit** page.
 - a From the **Policies > File Systems Audit** page, click the audit policy you wish to attach a file system. The audit policy's details page appears.
 - b Click **More Options > Add File Systems** in the heading of the **Attached File Systems** panel. The **Add File Systems** pop-up window appears.
 - c Select a file system from the **Available File Systems** panel to attach to the policy.
 - d Click **Add**.
- 2 Attach an audit policy to a file system from the **Storage > File Systems** page.
 - a From the **Storage > File Systems** page, click the file system to which you wish to attach an audit policy.
 - b From the **Attached Policies** panel, select **More Options > Set Audit Policy**. The **Set Audit Policy** pop-up window appears.
 - c Select the policy you wish to attach from the **Selected Policy** list.
 - d Click **Save**.

Removing a File System Audit Policy from a File System

When an audit policy is removed from a file system, auditing for that file system stops.

File systems can be removed from policies from either the **Policies > File System Audit** page or from the **Storage > File Systems** page.

To remove a single file system from an audit policy, perform the following:

- 1 Remove a file system from an audit policy from the **Policies > File System Audit** page.
 - a From the **Policies > File Systems Audit** page, click the audit policy from which you wish to remove a file system. The audit policy's details page appears.
 - b In the **Attached File Systems** panel, click the **Remove (X)** button in the row of the file system you wish to remove. The **Remove File System from Policy** pop-up window appears.
 - c Click **Remove**.
- 2 Remove a file system from an audit policy from the **Storage > File Systems** page.
 - a From the **Storage > File Systems** page, click the file system from which you wish to remove an audit policy.
 - b From the **Attached Policies** panel, click the **Remove (X)** button in the same row of the audit policy you wish to remove. The **Remove Policy from File System** pop-up window appears.
 - c Click **Remove**.

Enabling and Disabling a File System Audit Policy

When a file system audit policy is created, by default it is set to enabled. You can disable and enable an audit policy to stop and start file system auditing with that policy.

To disable an audit policy, perform the following:

- 1 From the **Policies > File Systems Audit** page, click the file system audit policy you wish to disable. The audit policy's details page appears.
- 2 In the **Details** panel, click the **More Options > Disable**. The **Disable Policy** pop-up window appears.
- 3 Click **Disable**.

To re-enable the policy, follow the same procedure and click **More Options > Enable** and then **Enable** to confirm.

Deleting a File System Audit Policy

File system audit policies cannot be deleted while they are in use. Prior to deleting an audit policy, all file systems to which the policy is attached must be removed.

To delete an audit policy, perform the following:

- 1 From the **Policies > File Systems Audit** page, click **More Options > Delete** in the same row as the audit policy you wish to delete in the **File System Audit Policy** panel. The **Delete File System Audit Policy** pop-up window appears.

- 2 Click **Delete**.

Deleting a Syslog Log Target

File system audit policy syslog log targets cannot be deleted while they are used as an audit logging destination. Prior to deleting an audit policy log target, the audit policy using that syslog server as an audit log target must first be deleted.

To delete a file system audit policy syslog log target, perform the following:

- 1 From the **Policies > File Systems > Audit** page, click the **Delete** button in the same row as the syslog log target you wish to delete in the **Syslog Log Targets** panel. The **Delete Log Target** pop-up window appears.
- 2 Click **Delete**.

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.

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

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

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

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

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

- **Policy:** The policy to which a rule is attached.
- **Client:** The IP address, host name, subnet in CIDR notation, netgroup, or anonymous (*).

- **IP address:** An IPv6 or IPv4 IP address.
- **host name:** A host name in short form (without dots) or in FQDN form, that can include the wildcards * or ?.
 - *: Matches any number of characters.
 - ?: Matches a single character in a specific position.
- **subnet in CIDR notation:** An IPv6 or IPv4 IP address and the length of the netmask separated with a slash (/).
- **netgroup:** An Active Directory, LDAP, or NIS netgroup.
- *: The asterisk (*) character specifies no filtering is applied.
- **Permission:** If the file system is read-only (**ro**) or read-write (**rw**).
- **Security:** The rule's security protocol. Values can be **krb5**, **krb5i**, **krb5p**, or **sys**. Note that multiple security protocols can be specified during rule creation and will be displayed.
- **Access:** The NFS export access rule. Values can be **no-squash**, **root-squash**, or **all-squash**.
- **Anon UID:** Anonymous user IDs.
- **Anon GID:** Anonymous group IDs.
- **Transport Security:** The transport security type. Values can be **none**, **tls**, or **mutual-tls**.
- **Secure:** If client ports are restricted to secure ports (**True**) or not (**False**).
- **fileid-32bit:** If 32-bit inodes are used (**True**) or 64-bit inodes are used (**False**).
- **atime:** If the timestamp is updated when a file is accessed (**True**) or not (**False**).
- **Index:** The index of the rule as ordered in its policy.

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

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

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

For additional information about export rules, see [NFS Export Rules](#).

NFS Export Rules

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

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

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

Valid export rules are listed below:

Table 8.2 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.

Table 8.2 Export Rules (continued)

Export Rules	Rules Description
<code>sec=krb5i</code>	Enables Kerberos security authentication and integrity checking of NFS operations.
<code>sec=krb5p</code>	Enables Kerberos security authentication, integrity checking of NFS operations, and NFS traffic encryption.
<code>fileid_32bit</code>	Allows 32-bit inode support for clients.
<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.

Table 8.2 Export Rules (continued)

Export Rules	Rules Description
atime	After a read operation has occurred, the inode access time is updated only if any of the following conditions is true: the previous access time is less than the inode modify time, the previous access time is less than the inode change time, or the previous access time is more than 24 hours ago.
noatime	Disables the update of inode access times after read operations.
secure	Prevents NFS access to client connections coming from non-reserved ports. Applies to NFSv3, NFSv4.1, and auxiliary protocols MOUNT and 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 8.2](#).

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 8.2](#).

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.

Export Rules	Rules Description
* ()	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.
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 Export** settings page, click the **Add (+)** button in the heading of the **NFS Export Policies** panel. The **Step 1: Create NFS Export Policy** pop-up window appears.
- 2 Enter the policy name.
- 3 (Optional) By default, the policy will be enabled upon creation. To disable the policy, click the **Enabled** checkbox.
- 4 Click **Create**. The **Step 2: Add NFS Export Rule** pop-up window appears. You can continue to configure the export rule or click **Skip** to configure the rule at later time.
- 5 In the **Client** field, enter the client as an IP address (IPv6 or IPv4), host name in short or FQDN form (including wildcards), subnet in CIDR notation, netgroup (Active Directory, LDAP, or NIS), or anonymous (*).
- 6 Set the access permission to **Read-Only** or **Read-Write**.
- 7 Select one or more security protocol. The security options `krb5`, `krb5i`, and `krb5p` are not supported by NFSv3.
- 8 Select one of the following NFS export access rules: **no-squash**, **root-squash**, or **all-squash**.
- 9 (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous UID** field enter one or more UIDs.

- 10 (Optional) If you selected **root-squash** or **all-squash**, in the **Anonymous GID** field, enter one or more GIDs.
- 11 Set the data transport security to **none**, **tls**, or **mutual-tls**.
- 12 (Optional) Select to restrict client ports to secure ports only.
- 13 (Optional) Select to use 32-bit inodes instead of 64-bit inodes.
- 14 (Optional) Select to update the timestamp when accessed.
- 15 Click **Save**.

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

Creating File System Exports

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

To create a file system export, perform the following:

- 1 From the **Policies > NFS Export > NFS Export Policies** panel, click the NFS export policy to which you wish to create a file system export. The detail page for that policy appears.
- 2 Click the Add (+) button in the heading of the **File System Exports** panel. The **Create File System Export** pop-up window appears.
- 3 From the **Server** field, select the server to which you are exporting the file system.
- 4 In the **File System** field, select or enter the name of the file system you wish to export.
- 5 In the **Export Name** field, enter a name for the exported file system, which is the name the clients will use to connect to the file system.
- 6 In the **Type** field, **NFS** is pre-selected (cannot be changed), which is the type of policy the file system is attached.
- 7 The **NFS Export Policy** field is pre-populated with the name of the NFS export policy which you selected to create the file system export.
- 8 Click **Create**.

Editing File System Exports

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

To edit a file system export, perform the following:

- 1 From the **Policies > NFS Export > NFS Export Policies** panel, click the NFS export policy with the file system export you wish to edit. The detail page for that policy appears.
- 2 In the **File System Exports** panel, click the Edit button in the same row as the file system export you wish to edit. The **Edit File System Export** pop-up window appears.
- 3 In the **Export Name** field, enter a new name for the exported file system, which is the name the clients will use to connect to the file system.
- 4 Click **Save**.

Deleting File System Exports

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

- 1 To delete a single file system export from an NFS export policy:
 - a From the **Policies > NFS Export > NFS Export Policies** panel, click the NFS export policy from which you wish to delete a file system export. The detail page for that policy appears.
 - b In the **File System Exports** panel, click the Delete button on the same row as the file system export you wish to delete. The **Delete File System Export Configuration** pop-up window appears.
 - c Click **Delete**.
- 2 To delete multiple file system exports from an NFS export policy:
 - a From the **Policies > NFS > NFS Export Policies** panel, click the NFS export policy from which you wish to delete multiple file systems. The detail page for that policy appears.
 - b Click **More Options > Delete** in the heading of the **File System Exports** panel. The **Delete File System Exports** pop-up window appears.

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

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

Adding Additional NFS Export Rules to NFS Export Policies

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

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

Editing NFS Export Rule Clients

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

To edit NFS export rule clients, perform the following:

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

- 2 In the **Existing Rules** panel, select the rules for which you wish to change the client. Each selected rule appears in the **Selected Rules** panel.
- 3 In the **Client** field, enter the new client as an IP address (IPv6 or IPv4), host name in short or FQDN form (including wildcards), subnet in CIDR notation, netgroup (Active Directory, LDAP, or NIS), or anonymous (*).
- 4 Click **Save**.

Editing NFS Export Rules

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

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

Reordering NFS Export Rules

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

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

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

To move an NFS export rule, perform the following:

- 1 From the **Policies > NFS Export > NFS Export Policies** panel, click the NFS export policy whose rule you wish to move. The detail page for that policy appears.
- 2 In the **NFS Export Rules** panel, click **More Options > Move** at the end of the row in which the rule you wish to move appears. The **Move NFS Export Rule** pop-up window appears.
- 3 In the **Before** list, select the new order of the rule.
- 4 Click **Move**.

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

Removing NFS Export Rules

- 1 To remove a single NFS export rule:
 - a From the **Policies > NFS Export > NFS Export Policies** panel, click the policy from which you wish to remove a rule. The detail page for that policy appears.
 - b In the **NFS Export Rules** panel, click **More Options > Remove** on the row of the rule you wish to remove. The **Remove Rule** dialog box appears.
 - c Click **Remove**.
- 2 To remove multiple rules from an object store access policy:
 - a From the **Policies > NFS > NFS Export Policies** panel, click the policy from which you wish to remove multiple rules. The detail page for that policy appears.
 - b Click **More Options > Remove** from the heading of the **NFS Export Rules** panel. The **Remove NFS Export Rules** pop-up window appears.

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

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

Renaming NFS Export Policies

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

To rename an NFS export policy, perform the following:

- 1 From the **Policies > NFS Export > NFS Export Policies** panel, click the NFS export policy you wish to rename. The detail page for that policy appears.
- 2 Click **More Options > Rename** in the heading of the **Details** panel. The **Rename NFS Export Policy** pop-up window appears.
- 3 Enter the policy's new name and click **Rename**.

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

Enabling and Disabling NFS Export Policies

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

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

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

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

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

- 3 To disable the policy:
 - a Click **More Options > Disable** in the **Details** panel. The **Disable Policy** pop-up window appears.
 - b Click **Disable**.

Deleting NFS Export Policies

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

 **Note:** By default, the policies displayed reflect filtering that was applied at the parent table. Selecting **Show all items** removes any previous filtering and all items as shown by default in the parent table are displayed.
 - b From the **Available Policies** panel, select one or more policies from the list. The selected policies appears in the **Selected Policies** panel. To select all policies, select the checkbox directly below the **Available Policies** heading.
 - c Click **Delete**.

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.

The **Object Store Access Policies** panel 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 8.4 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

Table 8.4 Pure-managed Policies (continued)

Access Policy	Policy Definition	Grant*
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
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:ListMultipartUploadParts

Table 8.4 Pure-managed Policies (continued)

Access Policy	Policy Definition	Grant*
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 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 8.5 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.

Table 8.5 S3 Actions (continued)

S3 Action	Permission
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.
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.

Table 8.5 S3 Actions (continued)

S3 Action	Permission
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.
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)

Table 8.5 S3 Actions (continued)

S3 Action	Permission
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 **Object Store Access Policies** panel, click a policy name to open a detail page about that policy.

Figure 8.1 Object Store Policy Details



Rules					1-1 of 1
Name	Effect	Actions	Resources	Conditions	
bucketconfigure1	allow	s3:PutBucketVersioning s3:PutLifecycleConfiguration s3:PutObjectLockConfiguration	*	-	⋮

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 **Object Store Access Policies** panel, click the **Add (+)** button. 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 **Object Store Access Policies** panel, 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 **Object Store Access Policies** panel, 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 **Object Store Access Policies** panel, 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 **Object Store Access Policies** panel, 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 **Object Store Access Policies** panel, 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 **Object Store Access Policies** panel, 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 **Object Store Access Policies** panel, 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 **Object Store Access Policies** panel, 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**.

Object Store Audit

Similar to file system audit policies, object store audit policies are comprised of log targets. A log target works as an endpoint where the audit logs for object store operations will get stored. Attach an audit policy to any bucket and audit logs for operations on that bucket will be stored in the customer's designated log target.

Events will be logged into an object. The log format is similar to AWS Cloudtrail Logging format. For example:

```
{
  "eventId": 27021597764222976,
  "eventTime": "2025-08-28T11:46:46.649Z",
  "eventVersion": "1.0",
  "eventName": "PutObject",
  "readOnly": false,
  "eventCategory": "Data",
  "managementEvent": false,
  "eventSource": "onPrem",
  "sourceIPAddress": "10.201.18.10",
  "bucketName": "mybucket1",
  "recipientAccountId": "dev",
  "resources": [
    {
      "type": "Pure::S3::Object",
      "ARN": "prn::s3::mybucket1/myfile"
    },
    {
      "accountId": "dev",
      "type": "Pure::S3::Bucket",
      "ARN": "prn::s3::mybucket1"
    }
  ],
  "userIdentity": {
    "type": "IAMUser",
    "arn": "prn::iam:array-id/local:obj-account-id/23:user/username",
    "accountId": "dev",
    "accessKeyId": "PSFBSAZRFCDLCCNBJPkPHGKPIMNHOOIGJKJKCBKIDE",
    "userName": "username"
  },
  "userAgent": "",
  "requestParameters": {
    "bucketName": "mybucket1",
    "key": "myfile",
    "Host": "mybucket1.s3.amazonaws.com"
  },
  "responseElements": {
    "x-amz-version-id":
    "AAAAAAAAAPOkdJZusrO7QBgAAAAAAUrAAAAAAAEAYAAAAAEpwAAAAAAAHAAAAAAABcAZm15
    ZmlaAA..",
    "ETag": "\"12345bb861061d1a052c592e2dc6b383\""
  }
}
```

FlashBlade allows buckets to use the same policy, thereby facilitating scaling of buckets by allowing a single change in an audit policy to affect many buckets.

 **Note:** Currently only one audit policy attachment is supported on a bucket. To change the audit policy on a bucket, you must first detach an audit policy and then attach the new audit policy. During this detach and attach policy operation, some audit events might get lost on the source bucket. To avoid this, log targets on the attached audit policy can be edited to achieve the same requirement.

 **Note:** You must configure the log targets before creating an audit policy with it.

The following limits are supported:

Type	Limit
Number of log_targets	100
Number of object audit policies	100
Number of audit prefixes per bucket	5
Number of buckets that can be audited	30,000
Number of audit prefixes across all buckets in the system	60,000

The following operations are supported for object store audit:

Table 8.7 Supported Operations

Operation	Endpoint
Abort	AbortMultipartUpload
Bucket	BucketOptions
Complete	CompleteMultipartUpload
Copy	CopyObject
Delete	- DeleteBucketLifecycle - DeleteMultipleObjects - DeleteObject

Table 8.7 Supported Operations (continued)

Operation	Endpoint
Get	• GetBucketAcl • GetBucketCors • GetBucketLifecycleConfiguration • GetBucketLocation • GetBucketPolicy • GetBucketPolicyStatus • GetBucketVersioning • GetObjectAcl • GetObject • GetObjectLegalHold • GetObjectLockConfiguration • GetObjectRetention • GetObjectTagging • GetPublicAccessBlock
Head	• HeadBucket • HeadObject
Initiate	InitiateMultipartUpload
List	• ListBucketObjectVersions • ListObjects • ListObjectsV2 • ListMultipartUploads
Post	PostObject

Table 8.7 Supported Operations (continued)

Operation	Endpoint
Put	- PutBucketLifecycleConfiguration - PutBucketVersioning - PutObject - PutObjectLegalHold - PutObjectLockConfiguration - PutObjectLockRetention
Touch	TouchObject
Upload	- UploadPart - UploadPartCopy

From the **Policies > Object Store Audit** page, the **Audit Policies (Object Store)** and **Audit Policy Log Targets (Bucket)** panels allow you to view and manage object store audit polices and targets.

The **Audit Policies (Buckets)** panel lists the names of all created audit policies and if those policies are enabled or disabled. Clicking an audit policy displays that policy's details page. This page allows you to manage attached buckets and attached audit policy log targets.

The **Audit Policy Log Targets (Bucket)** panel lists the names of the policies, their target buckets, their log rotation durations, and log name prefixes. Clicking an audit policy log target displays that log target's details page.

Best Practices

Pure Storage recommends the following best practices for object audit log configuration and scaling:

- User access policies should be used to grant appropriate permissions and to restrict the access of audit objects.
- Object Lock/SafeMode should be used in order to make the audit object immutable.
- The cost of auditing small object sizes may be high.

Creating an Object Store Log Target

To create an object store log target, perform the following:

- 1 From the **Policies > Object Store Audit** page, click the **Create (+)** button in the heading of the **Audit Policy Log Targets (Bucket)** panel. The **Create Log Target** pop-up window appears.
- 2 Enter the log target name.
- 3 Select the target bucket.
- 4 Enter the log rotation duration. The log rotation is the duration after which new objects are created where the audit logs are written. The default is 15 minutes.
- 5 Enter the log name prefix. The default prefix is audit_logs.
- 6 Click **Create** to complete adding the new log target.

Once created, the audit policy log target will appear in the **Audit Policy Log Targets (Bucket)** panel.

Editing an Object Store Log Target

To edit an object store log target, perform the following:

- 1 From the **Policies > Object Store Audit** page, click the **Create (+)** button in the heading of the **Audit Policy Log Targets (Bucket)** panel. The **Create Log Target** pop-up window appears.
- 2 Select the target bucket.
- 3 Enter the log rotation duration. The log rotation is the duration after which new objects are created where the audit logs are written. The default is 15 minutes.
- 4 Enter the log name prefix. The default prefix is audit_logs.
- 5 Click **Save** to complete adding the new log target

Creating an Object Store Audit Policy

To create an object store audit policy, perform the following:

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

- 4 Select the log target. Click the **Add (+)** button. The **Add Log Targets** pop-up window appears. Select one log target from the **Available Log Targets** pane and click **Add**.
- 5 Click **Create**.

Once created, the audit policy will appear in the **Audit Policies (Object Store)** panel.

Editing an Object Store Audit Policy

To edit an object store audit policy, perform the following:

- 1 Navigate to the **Policies > Object Store Audit** page.
- 2 In the **Audit Policies (Object Store)** panel, click **More Options > Edit** in the same row as the policy you wish to edit. The **Edit Object Store Audit Policy** pop-up window appears.
- 3 Enable or disable the policy by selecting or deselecting the **Enabled** checkbox.
- 4 Remove and add a new log target.
 - a To remove a log target, click the **Remove (X)** button next to the log target name.
 - b To add a new log target, click the **Add (+)** button. The **Add Log Targets** pop-up window appears. Select one log target from the **Available Log Targets** pane and click **Add**.
- 5 Click **Save**.

Once saved, the updated audit policy will appear in the **Audit Policies (Object Store)** panel.

Adding Buckets to an Object Store Audit Policy

To add buckets to an object store audit policy, perform the following:

- 1 Navigate to the **Policies > Object Store Audit** page.
- 2 In the **Audit Policies (Object Store)** panel, click the policy to which you wish to add buckets. That policy's details page opens.
- 3 In the **Members** panel, click **More Options > Add Members**. The **Add Buckets** pop-up window appears.
- 4 From the **Available Buckets** pane, select one or more buckets to add to the object store audit policy and then click **Add**.

The **Members** panel is updated displaying the added bucket(s).

Removing Buckets from an Object Store Audit Policy

To remove buckets from an object store audit policy, perform the following:

- 1 Navigate to the **Policies > Object Store Audit** page.
- 2 In the **Audit Policies (Object Store)** panel, click the policy from which you wish to remove buckets. That policy's details page opens.
- 3 In the **Members** panel, click **More Options > Remove Members**. The **Remove Buckets** pop-up window appears.
- 4 From the **Current Buckets** pane, select one or more buckets to remove from the object store audit policy and then click **Remove**.
- 5 Alternatively, if you wish to remove only a single bucket, in the **Members** panel, click the **Remove (X)** button in the same row as the bucket you wish to remove. The **Remove Member from Policy** pop-up window appears. Click **Remove**.

Enabling and Disabling an Object Store Audit Policy

When an object store audit policy is created, by default it is set to enabled. You can disable and enable an audit policy to stop and start object store auditing with that policy.

To disable an audit policy, perform the following:

- 1 From the **Policies > Object Store Audit > Audit Policies (Object Store)** page, click the object store audit policy you wish to disable. The audit policy's details page appears.
- 2 In the **Details** panel, click the **More Options > Disable**. The **Disable Object Store Audit Policy** pop-up window appears.
- 3 Click **Disable**.

To re-enable the policy, follow the same procedure and click **More Options > Enable** and then **Enable** to confirm.

Renaming an Object Store Audit Policy

To rename an object store audit policy, perform the following:

- 1 From the **Policies > Object Store Audit > Audit Policies (Object Store)** panel, click the object store audit policy you wish to rename. The audit policy's details page appears.
- 2 In the **Details** panel, click the **More Options > Rename**. The **Rename Object Store Audit Policy** pop-up window appears.

- 3 Enter a new name.
- 4 Click **Rename**.

Deleting an Object Store Audit Policy

Object store audit policies cannot be deleted while they are in use. Prior to deleting an audit policy, all members of the policy must be removed.

To delete an audit policy, perform the following:

- 1 From the **Policies > File Systems Audit > Audit Policies (Object Store)** panel, click **More Options > Delete** in the same row as the audit policy you wish to delete. The **Delete Object Store Audit Policy** pop-up window appears.
- 2 Click **Delete**.

Deleting an Object Store Log Target

Object store log targets cannot be deleted while they are used as an audit logging destination. Prior to deleting an audit policy log target, the audit policy using that audit log target must first be deleted.

To delete an object store audit policy log target, perform the following:

- 1 From the **Policies > Object Store Audit > Audit Policy Log Targets (Bucket)** panel, click the **Delete** button in the same row as the log target you wish to delete. The **Delete Log Target** pop-up window appears.
- 2 Click **Delete**.

File System Quality of Service Policies

FlashBlade allows you to set a ceiling on user requests (rate limiting) for file systems using Quality of Service (QoS) policies. With QoS policies, you can define limits on the amount of resources a file system workload can consume. Ceilings are specified in IOPS/bandwidth and are applied on a per-file system basis by attaching a QoS policy.

With Purity 4.6.3 and later, Pure Storage recommends separating object storage and file system workloads into distinct realms for secure multitenancy if Realm QoS ceiling is configured. This decision is recommended because realm-level Quality of Service (QoS) ceiling is not currently supported for object

storage. This separation ensures proper QoS enforcement for file system workloads within a realm. Realm level QoS ceiling with Object storage will be supported in a future release.

Navigate to the **QoS Policies** panel by clicking **Policies > File Systems > QoS**.

The **QoS Policies** panel displays all configured QoS policies. The following information is displayed:

- **Name:** The QoS policy name.
- **Enabled:** Whether the policy is enabled (True) or disabled (False).
- **Max Total bytes/sec:** The maximum allowed bytes per second totaled across all clients for the policy.
- **Max Total ops/sec:** The maximum allowed operations per second totaled across all clients for the policy.

Clicking a policy name opens the details page for that policy. The policy's details page provides a **Details** panel and an **Attached File Systems** panel.

- The **Details** panel provides summary details for the policy. The panel also provides functionality to rename, edit, and disable/enable the policy.
- The **Attached File Systems** panel displays all file systems to which the policy is attached. The panel also provides functionality to add and remove file systems.

Creating a File System QoS Policy

To create a file system QoS policy, perform the following:

- 1 Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click Add (+). The **Create QoS Policy** pop-up window appears.
- 2 Enter the policy's name.
- 3 Click the **Enabled** checkbox to enable or disable the policy upon creation. By default, the policy will be enabled.
- 4 In the **Max Total byte/sec** field, enter the maximum allowed bytes per second totaled across all clients. The minimum allowed limit is 1M bytes/sec. The maximum allowed limit is 512G bytes/sec. If no value is specified, there will be no limit.
- 5 In the **Max Total ops/sec** field, enter the maximum allowed operations per second totaled across all clients. The minimum allowed limit is 100 OPs/sec. The maximum allowed limit is 100M OPS/sec. If no value is specified, there will be no limit.
- 6 Click **Create**.

Editing a File System QoS Policy

To edit a file system QoS policy, perform the following:

- 1 Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the Edit button in the same row as the policy you wish to edit. The **Edit QoS Policy** pop-up window appears.
- 2 Click the **Enabled** checkbox to enable or disable the policy.
- 3 In the **Max Total byte/sec** field, enter the maximum allowed bytes per second totaled across all clients. The minimum allowed limit is 1M bytes/sec. The maximum allowed limit is 512G bytes/sec. If no value is specified, there will be no limit.
- 4 In the **Max Total ops/sec** field, enter the maximum allowed operations per second totaled across all clients. The minimum allowed limit is 100 OPs/sec. The maximum allowed limit is 100M OPS/sec. If no value is specified, there will be no limit.
- 5 Click **Save**.

Enabling a File System QoS Policy

To enable a file system QoS policy, perform the following:

- 1 Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the policy you wish to enable. The details page for that policy appears.
- 2 In the **Details** panel, click **More Options > Enable**. The **Enable Policy** pop-up window appears.
- 3 Click **Enable**.

Disabling a File System QoS Policy

To disable a file system QoS policy, perform the following:

- 1 Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the policy you wish to disable. The details page for that policy appears.
- 2 In the **Details** panel, click **More Options > Disable**. The **Disable Policy** pop-up window appears.
- 3 Click **Disable**.

Renaming a File System QoS Policy

To rename a file system QoS policy, perform the following:

- 1 Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the policy you wish to rename. The details page for that policy appears.
- 2 In the **Details** panel, click **More Options > Rename**. The **Rename QoS Policy** pop-up window appears.
- 3 Enter a new name for the policy.
- 4 Click **Rename**.

Attaching a QoS Policy to File Systems from the File Systems Details Page

To attach a file system QoS policy to a file system, perform the following:

- 1 Navigate to **Storage > File Systems**. In the **File Systems** panel, click the file system to which you wish to attach a QoS policy. The details page for the file system appears.
- 2 In the **Attached Policies** panel, click **More Options > Set QoS Policy**. The **Set QoS Policy** pop-up window appears.
- 3 From the **Selected Policy** list, select the QoS policy you wish to attach to the file system.
- 4 Click **Save**.

Attaching a QoS Policy to File Systems from the QoS Policy Details Page

To attach a file system QoS policy to a file system, perform the following:

- 1 Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the policy to which you wish to attach a file system. The details page for the policy appears.
- 2 In the **Members** panel, click **More Options > Add File Systems**. The **Add File Systems** pop-up window appears.
- 3 From the **Available File Systems** panel, select one or more file systems to add to the policy.
- 4 Click **Add**.

Removing a Single QoS Policy from a File System from the File Systems Details Page

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

- 1 Navigate to **Storage > File Systems**. In the **File Systems** panel, click the file system from which you wish to remove a QoS policy. The details page for the file system appears.
- 2 In the **Attached Policies** panel, click the **Remove** button in the same row as the policy you wish to remove from the file system. The **Remove Policy from File System** pop-up window appears.
- 3 Click **Remove**.

Removing QoS Policies from File Systems from the QoS Policy Details Page

- 1 To remove a single file system from a QoS policy, perform the following:
 - a Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the policy from which you wish to remove a file system. The details page for the policy appears.
 - b In the **Attached File Systems** panel, click the **Remove** button in the same row of the file system you wish to remove from the policy. The **Remove File System from Policy** pop-up window appears.
 - c Click **Remove**.
- 2 To remove multiple file systems from a QoS policy, perform the following:
 - a Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the policy from which you wish to remove file systems. The details page for the policy appears.
 - b In the **Attached File Systems** panel, click **More Options > Remove File Systems**. The **Remove File Systems** pop-up window appears.
 - c Select the file systems you wish to remove from the policy from the **Current File Systems** list.
 - d Click **Remove**.

Deleting a File System QoS Policy

- 1 To delete a single QoS policy, perform the following:
 - a Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click the **Delete** button in the same row as the policy you wish to delete. The **Delete QoS Policy** pop-up window appears.
 - b Click **Delete**.
- 2 To delete multiple QoS policies, perform the following:
 - a Navigate to **Policies > File Systems > QoS**. In the **QoS Policies** panel, click **More Options > Delete Policies**. The **Delete QoS Policies** pop-up window appears.

- b** Select the policies you wish to delete from the **Available Policies** list.
- c** Click **Delete**.

Object Store Quality of Service Policies

FlashBlade allows you to set a ceiling on user requests (rate limiting) for object store buckets using Quality of Service (QoS) policies. With QoS policies, you can define throughput and IOPs limits for individual buckets, file systems, and realms. Ceilings are specified in IOPS/bandwidth and are applied on a per-bucket basis by attaching a QoS policy.

Navigate to the **QoS Policies** panel by clicking **Policies > Object Store > QoS**.

The **QoS Policies** panel displays all configured QoS policies. The following information is displayed:

- **Name:** The QoS policy name.
- **Enabled:** Whether the policy is enabled (True) or disabled (False).
- **Max Total bytes/sec:** The maximum allowed bytes per second totaled across all clients for the policy.
- **Max Total ops/sec:** The maximum allowed operations per second totaled across all clients for the policy.

Clicking a policy name opens the details page for that policy. The policy's details page provides a **Details** panel and a **Members** panel.

- The **Details** panel provides summary details for the policy. The panel also provides functionality to rename, edit, and disable/enable the policy.
- The **Members** panel displays all buckets to which the policy is attached. The panel also provides functionality to add and remove buckets.

Creating an Object Store QoS Policy

To create an object store QoS policy, perform the following:

- 1** Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click Add (+). The **Create QoS Policy** pop-up window appears.
- 2** Enter the policy's name.
- 3** Click the **Enabled** checkbox to enable or disable the policy upon creation. By default, the policy will be enabled.

- 4 In the **Max Total byte/sec** field, enter the maximum allowed bytes per second totaled across all clients. The minimum allowed limit is 1M bytes/sec. The maximum allowed limit is 512G bytes/sec. If no value is specified, there will be no limit.
- 5 In the **Max Total ops/sec** field, enter the maximum allowed operations per second totaled across all clients. The minimum allowed limit is 100 OPs/sec. The maximum allowed limit is 100M OPS/sec. If no value is specified, there will be no limit.
- 6 Click **Create**.

Editing an Object Store QoS Policy

To edit an object store QoS policy, perform the following:

- 1 Navigate to **Policies > Object Store > QoS**. In the QoS Policies panel, click the Edit button in the same row as the policy you wish to edit. The **Edit QoS Policy** pop-up window appears.
- 2 Click the **Enabled** checkbox to enable or disable the policy.
- 3 In the **Max Total byte/sec** field, enter the maximum allowed bytes per second totaled across all clients. The minimum allowed limit is 1M bytes/sec. The maximum allowed limit is 512G bytes/sec. If no value is specified, there will be no limit.
- 4 In the **Max Total ops/sec** field, enter the maximum allowed operations per second totaled across all clients. The minimum allowed limit is 100 OPs/sec. The maximum allowed limit is 100M OPS/sec. If no value is specified, there will be no limit.
- 5 Click **Save**.

Enabling an Object Store QoS Policy

To enable an object store QoS policy, perform the following:

- 1 Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click the policy you wish to enable. The details page for that policy appears.
- 2 In the **Details** panel, click **More Options > Enable**. The **Enable Policy** pop-up window appears.
- 3 Click **Enable**.

Disabling an Object Store QoS Policy

To disable an object store QoS policy, perform the following:

- 1 Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click the policy you wish to disable. The details page for that policy appears.

- 2 In the **Details** panel, click **More Options > Disable**. The **Disable Policy** pop-up window appears.
- 3 Click **Disable**.

Renaming an Object Store QoS Policy

To rename an object store QoS policy, perform the following:

- 1 Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click the policy you wish to rename. The details page for that policy appears.
- 2 In the **Details** panel, click **More Options > Rename**. The **Rename QoS Policy** pop-up window appears.
- 3 Enter a new name for the policy.
- 4 Click **Rename**.

Attaching an Object Store QoS Policy to a Bucket from the Bucket Details Page

To attach an object store QoS policy to a bucket, perform the following:

- 1 Navigate to **Storage > Object Store**. In the **Buckets** panel, click the bucket to which you wish to attach an object store QoS policy. The details page for the bucket appears.
- 2 In the **Policies** panel, click **More Options > Set QoS Policy**. The **Set QoS Policy** pop-up window appears.
- 3 From the **Selected Policy** list, select the QoS policy you wish to attach to the bucket.
- 4 Click **Save**.

Attaching an Object Store QoS Policy to Buckets from the QoS Policy Details Page

To attach an object store QoS policy to a bucket, perform the following:

- 1 Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click the policy to which you wish to attach a bucket. The details page for the policy appears.
- 2 In the **Members** panel, click **More Options > Add Buckets**. The **Add Buckets** pop-up window appears.
- 3 From the **Available Buckets** panel, select one or more buckets to add to the policy.

- 4 Click **Add**.

Removing a Single Object Store QoS Policy from a Bucket from the Bucket Details Page

To remove a single object store QoS policy from a bucket, perform the following:

- 1 Navigate to **Storage > Object Store**. In the **Buckets** panel, click the bucket from which you wish to remove a QoS policy. The details page for the bucket appears.
- 2 In the **Policies** panel, click the Remove button in the same row as the policy you wish to remove from the bucket. The **Remove Bucket QoS Policy** pop-up window appears.
- 3 Click **Remove**.

Removing Buckets from an Object Store QoS Policy from the QoS Policy Details Page

- 1 To remove a single bucket from a QoS policy, perform the following:
 - a Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click the policy from which you wish to remove a bucket. The details page for the policy appears.
 - b In the **Members** panel, click the Remove button in the same row of the bucket you wish to remove from the policy. The **Remove Bucket from Policy** pop-up window appears.
 - c Click **Remove**.
- 2 To remove multiple buckets from a QoS policy, perform the following:
 - a Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click the policy from which you wish to remove buckets. The details page for the policy appears.
 - b In the **Members** panel, click **More Options > Remove**. The **Remove Members** pop-up window appears.
 - c Select the buckets you wish to remove from the policy from the **Current Members** list.
 - d Click **Remove**.

Deleting an Object Store QoS Policy

- 1 To delete a single QoS policy, perform the following:
 - a Navigate to **Policies > Object Store > QoS**. In the **QoS Policies** panel, click the Delete button in the same row as the policy you wish to delete. The **Delete QoS Policy** pop-up window appears.

- ## Adding Members to a QoS Policy

- 1 Navigate to **Policies > QoS**. In the **QoS Policies** panel, click the policy you want to add members to. The policy details page will appear.
- 2 In the **Members** panel, click **More Options** and select the type of member you want to add to the policy: **Add Buckets**, **Add File Systems**, **Add Realms**. The selected member dialogue window will appear.
- 3 Select one or more members from the left side of the window to add to the right side of the window.
- 4 Click **Add**.

Removing Members from a QoS Policy

- 1 Navigate to **Policies > QoS**. In the **QoS Policies** panel, click the policy you want to remove members from. The policy details page will appear.
- 2 In the **Members** panel, click **More Options > Remove**, the **Remove Members from QoS Policy** dialogue window will appear.
- 3 Select one or more members from the left side of the window to remove to add to the right side of the window.
- 4 Check the acknowledgement statement that you understand the consequences of removing the selected member.
- 5 Click **Remove**.

S3 Export Policies

The S3 export policy is a configuration that controls which S3 buckets are accessible through specific servers for particular accounts in Pure Storage's Object Store. The intent is to restrict or allow bucket access on a per-server and per-account basis.

S3 export policies are attached to "Account Exports," which link an account to a server. Only buckets explicitly included in the export policy for that account are visible/accessible through the server. Policy rules specifying "Allow" or "Deny" actions for sets of resources (currently limited to buckets). Rules define the actions (`pure:S3Access`), effect (allow/deny), and a list of buckets it applies to. Explicit "Deny" takes precedence over "Allow."

There is a default policy called `_s3_export_allow_everything`, which permits access to all buckets. This default policy cannot be modified or deleted.

Creating an S3 Export Policy

To create an S3 export policy for a realm, perform the following:

- 1 Navigate to **Policies > S3 Export**. In the **S3 Export Policies** panel, click Add (+). The **Create S3 Export Policy** dialogue window appears.
- 2 In the **Scope Type** field, select **array** or **realm**.
- 3 In the **Scope Name** field, for an array it defaults to your current FlashBlade system. For a realm select from the drop-down list.
- 4 In the **Name** field, enter a name for the S3 export policy.
- 5 In the **Enabled** field, check the box to enable the policy, or enable it at a later time.
- 6 Click **Create** and a **Step 2: Add Rule to Policy (Policy Name)** dialogue window appears.

 **Note:** Step 2 is optional. Click **Skip** to add a rule at a later time.

- 7 In the **NAME** field, add a name for the policy.
- 8 In the **Effect** field, select the **Allow** or **Deny** option. This allows or denys S3 requests that match the parameters you will select for in the Action and Resources fields. Any explicit "Deny" rule takes precedence over all "Allow" rules.
- 9 In the **Action** field, only `pure:S3Access` is supported at this time. This is the action granted by this rule. Each action may restrict other properties of the rule.
- 10 In the **Resources** field, add a comma-separated list of buckets you want this policy to apply to.

- 11 Click **Add**.

Adding a Rule to an S3 Export Policy

To add a rule to an S3 export policy, perform the following:

- 1 Navigate to **Policies > S3 Export**. In the **S3 Export Policies** panel, click the policy you want to add a rule to.
- 2 In the **S3 Export Rules** panel click Add (+). The **Add Rule to Policy (Policy Name)** dialogue window appears.
- 3 In the **NAME** field, add a name for the rule.
- 4 In the **Effect** field, select the **Allow** or **Deny** option. This allows or denys S3 requests that match the parameters you will select for in the Action and Resources fields. Any explicit “Deny” rule takes precedence over all “Allow” rules.
- 5 In the **Action** field, only pure:S3Access is supported at this time.
- 6 In the **Resources** field, add a comma-separated list of buckets you want this policy to apply to.
- 7 Click **Add**.

Renaming an S3 Export Policy

To rename an S3 export policy, perform the following:

- 1 Navigate to **Policies > S3 Export**. In the **S3 Export Policies** panel, in the same row as the policy you want to rename, click **More Options > Remove**. The **Rename S3 Export Policy** dialogue window appears.
- 2 In the **Existing Name** field, it is greyed out and displays the name of the policy you are editing.
- 3 In the **New Name** field, enter a new name for the policy.
- 4 Click **Rename**.

Disabling an S3 Export Policy

To disable an S3 Export Policy, perform the following:

- 1 Navigate to **Policies > S3 Export**. In the **S3 Export Policies** panel, in the same row as the policy you want to disable, click **More Options > Disable**. The **Disable Policy** dialogue window appears.
- 2 Click **Disable**.

Deleting an S3 Export Policy

To delete an S3 export policy, perform the following:

- 1 Navigate to **Policies > S3 Export**. In the **S3 Export Policies** panel, in the same row as the policy you want to delete, click **More Options > Delete**. The **Delete Policy** dialogue window appears.
- 2 Click **Delete**.

SMB Client and Share Policies

SMB settings 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.

SMB settings 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**).
- **Access-based Enumeration:** Whether the policy has Access-based Enumeration restrictions 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. Additionally, you can enable Access-based Enumeration restrictions when creating SMB client policies to hide files and folders from users who do not have generic read permissions (list folder/read data, read attributes, read extended attributes, read permissions) on a network SMB share.

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 (*).

SMB shares can be made hidden by appending a "\$" to the end of the SMB export name. Hidden shares are typically not listed when Windows obtains a list of shares. The share may still be listed with administrative tools such as Microsoft Management Console. The filtering happens on the client side and as such this is not a feature that should be used for security. To access a hidden share, one needs to know the name of the share, including the "\$" at the end of the name.

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** settings, clicking a policy in the **SMB Client Policies** panel will open the details page for SMB client rules associated with that SMB client policy.

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 and if Access-based Enumeration is enabled. It also allows you to rename and edit 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 **Policies > SMB** settings, 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 By default, the policy will be enabled upon creation. To disable the policy, click the **Enabled** checkbox.
- 4 (Optional) Click to enable **Access-based Enumeration** if you wish to hide files and folders from users who do not have generic read permissions on a network SMB share.
- 5 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.
- 6 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.
- 7 Set rule permission as **Read-Only** or **Read-Write**.
- 8 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.
- 9 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 that file system 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 will 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 a 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.

Editing SMB Client Policies

SMB client policies can be edited at any time after their creation. Editing allows you to enable or disable the SMB client policy as well as enable or disable Access-based Enumeration.

To edit an SMB client policy, perform the following:

- 1 From the **Policies > SMB > SMB Client Policies** panel, click the SMB client policy you wish to edit. The detail page for that policy appears.
- 2 Click **More Options > Edit** in the heading of the **Details** panel. The **Edit SMB Client Policy** pop-up window appears.
- 3 Enable or disable the policy. Note that if the policy is disabled, attached files will no longer be accessible to clients specified by the policy.
- 4 Enable or disable **Access-based Enumeration**. Enabling Access-based Enumeration allows you to hide files and folders from users who do not have generic read permissions on a network SMB share.

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

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.

SMB shares can be made hidden by appending a "\$" to the end of the SMB export name. Hidden shares are typically not listed when Windows obtains a list of shares. The share may still be listed with administrative tools such as Microsoft Management Console. The filtering happens on the client side and as such this is not a feature that should be used for security. To access a hidden share, one needs to know the name of the share, including the "\$" at the end of the name."

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** settings, clicking a policy in the **SMB Share Policies** panel will open the details page for rules associated with that policy.

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 **Policies > SMB** settings, 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 be 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. **Snapshot** settings allows you to manage the array's snapshot policies.

Snapshot settings contains the **Snapshot Policies** panel, which 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 **Snapshot Policies** panel, click a policy name to open a detail page about that snapshot policy.

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 **Policies > Snapshot** settings, 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 **Snapshot Policies** panel, 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][amlpm], 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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.

 **Note:** The Snapshot Rules Edit process provides a warning as this process deletes the existing rule and adds the new rule. When the rule is deleted in the edit process, all the snapshots that were created by this rule are destroyed, but are automatically restored that follows the updated rule. If the new snapshot rule is stricter (with longer “keep-for” [Retention] or shorter “every” [Frequency]), all the snapshots that belonged to the previous rule will be restored.

- 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 **Snapshot Policies** poanel, 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, click the policy from which you wish to remove a snapshot. The policy's detail page appears.

- ## Removing a Policy from a Snapshot

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

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.
- 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 **Snapshot Policies** panel, 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 **Snapshot Policies** panel, 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**.

Storage Class Tiering Policies

FlashBlade Zero Move Tiering allows you to commit entire file system workloads to one of two storage classes: S500X-S for speed (hot data) and S500X-A for archive (cold data). For S500X-S class files that may need to switch class from hot to cold, and vice versa, FlashBlade supports the creation of storage class tiering policies. Storage class tiering policies allow you to set the criteria by which files switch storage class as well as the duration of the storage class switch. For example, you may create a policy that specifies that a file switches from cold to hot storage upon writes to the file and archives to cold storage based on its latest modification after 3 days. Note that attaching storage class tiering policies to S500X-A class file systems is not allowed.

To view and manage storage class tiering policies, navigate to the **Storage Class Tiering Policies** panel by clicking **Policies > Storage Class Tiering**.

The **Storage Class Tiering Policies** panel displays all configured Storage Class Tiering policies. The following information is displayed:

- **Name:** The policy name.
- **Type:** The policy type.
- **Enabled:** Whether the policy is enabled (True) or disabled (False).
- **Archival Rule:** The criteria by which data is moved from hot to cold storage. Move criteria can either be **LATEST_ACCESS** (reads and writes) or **LATEST_MODIFICATION** (writes only). Note that metadata operations will not change the storage class directly, but may be reflected later if you change the modify time of the inode directly
- **Archival Time:** The duration for when data is moved from hot to cold storage.
- **Retrieval Rule:** The criteria by which data should be moved from cold to hot storage. Move criteria can either be **ON_READ** or **ONLY_ON_WRITE**.

Clicking a policy name opens the details page for that policy. The policy's details page provides a **Details** panel and an **Attached File Systems** panel.

- The **Details** panel provides summary details for the policy. The panel also provides functionality to edit and disable/enable the policy.
- The **Attached File Systems** panel displays all file systems to which the policy is attached. The panel also provides functionality to add and remove file systems.

Creating a Storage Class Tiering Policy

To create a storage class tiering policy, perform the following:

- 1 From the **Policies > File Systems > Storage Class Tiering** page, click the **Add (+)** button in the heading of the **Storage Class Tiering Policies** panel. The **Create Storage Class Tiering Policy** pop-up window appears.
- 2 Enter the policy name.
- 3 By default the created policy will be enabled. Click the **Enabled** checkbox to disable the policy upon creation.
- 4 In the **Archival Rules** field, select the criteria by which data should be moved from hot to cold storage. Move criteria can either be by **LATEST_ACCESS** or **LATEST_MODIFICATION**.
- 5 In the **Archival Time** field, enter the duration for when data is moved from hot to cold storage. The duration can be specified in seconds (s), minutes (m), hours (h), or days (d), and must be longer than 2 days. Must be specified in exact days, for example 2d, 48h, etc. Values such as 25h, 3.5d, etc are not allowed.
- 6 In the **Retrieval Rules** field, select the criteria by which data should be moved from cold to hot storage. Move criteria can either be by **ON_READ** or **ONLY_ON_WRITE**.
- 7 Click **Create** to create the policy.

Editing a Storage Class Tiering Policy

To edit a storage class tiering policy, perform the following:

- 1 From the **Policies > File Systems > Storage Class Tiering** page, click the Edit button in the same row as the storage class tiering policy you wish to edit in the **Storage Class Tiering Policies** panel. The **Edit Storage Class Tiering Policy** pop-up window appears.
- 2 Click the **Enabled** checkbox to enable or disable the policy.
- 3 In the **Archival Rules** field, select the criteria by which data should be moved from hot to cold storage. Move criteria can either be by **LATEST_ACCESS** or **LATEST_MODIFICATION**.
- 4 In the **Archival Time** field, enter the duration for when data is moved from hot to cold storage. The duration can be specified in seconds (s), minutes (m), hours (h), or days (d), and must be longer than 2 days. Must be specified in exact days, for example 2d, 48h, etc. Values such as 25h, 3.5d, etc are not allowed.
- 5 In the **Retrieval Rules** field, select the criteria by which data should be moved from cold to hot storage. Move criteria can either be by **ON_READ** or **ONLY_ON_WRITE**.
- 6 Click **Save** to save your changes.

Enabling and Disabling Storage Class Tiering Policies

When storage class tiering policies are created, they are enabled by default. You can disable and enable storage class policies at any time.

- 1 From the **Policies > File Systems > Storage Class Tiering** page, click the storage class tiering policy you wish to enable or disable in the **Storage Class Tiering Policies** panel. The details page for that policy appears.
- 2 In the **Details** panel, click **Enable Policy** or **Disable Policy**. The **Enable Policy** or **Disable Policy** pop-up window appears.
- 3 Click **Enable** or **Disable**.

Deleting a Storage Class Tiering Policy

Storage class tiering policies can be deleted.

Note that policy detection will fail if the policy you wish to delete is currently attached to a file system.

- 1 To delete a single storage class tiering store policy:
 - a From the **Policies > File Systems > Storage Class Tiering** page, click the Delete button in the same row as the storage class tiering policy you wish to delete in the **Storage Class Tiering Policies** panel. The **Delete Storage Class Tiering Policy** pop-up window appears.
 - b Click **Delete**.
- 2 To delete multiple object store policies:
 - a From the **Policies > File Systems > Storage Class Tiering** page, click **More Options > Delete Policies** in the header of the **Storage Class Tiering** panel. The **Delete Storage Class Tiering Policies** pop-up window appears.
 - b Select the policies you wish to delete from the **Available Policies** panel.
 - c Click **Delete**.

Attaching File Systems to Storage Class Tiering Policies

Once you have created storage class tiering policies, you can attach file systems to the policies to manage when file systems data is transitioned from cold to hot, and hot to cold.

To attach file systems to a storage class tiering policy, perform the following:

- 1 From the **Policies > File Systems > Storage Class Tiering** page, click the storage class tiering policy to which you wish to attach file systems in the **Storage Class Tiering Policies** panel. The details page for that storage class tiering policy appears.
- 2 From the header of the **Attached File Systems** panel, click **More Options > Add File Systems**. The **Add File Systems to Tiering Policy** pop-up window appears.
- 3 From the **Available File Systems** panel, select one or more file systems to attach to the policy.
- 4 Click **Add**.

Removing File Systems from Storage Class Tiering Policies

File systems can be removed from storage class tiering policies at any time. Note that when a file system is removed from a policy, the file system will revert to its original storage class.

- 1 To remove a single file system from a storage class tiering store policy:
 - a From the **Policies > File Systems > Storage Class Tiering** page, click the storage class tiering policy from which you wish to remove a file system in the **Storage Class Tiering Policies** panel. The details page for that storage class tiering policy appears.
 - b In the **Attached File Systems** panel, click the Remove button in the same row as the file system you wish to remove from the policy. The **Remove File System from Tiering Policy** pop-up window appears.
 - c Click **Remove**.
- 2 To remove multiple file systems from a storage class tiering store policy:
 - a From the **Policies > File Systems > Storage Class Tiering** page, click the storage class tiering policy from which you wish to remove file systems in the **Storage Class Tiering Policies** panel. The details page for that storage class tiering policy appears.
 - b From the header of the **Attached File Systems** panel, click **More Options > Remove File Systems**. The **Remove File Systems from Tiering Policy** pop-up window appears.
 - c From the **Current File Systems** panel, select the file systems you wish to remove from the policy.
 - d Select the confirmation checkbox.
 - e Click **Remove**.

File System WORM Policies

File system Write Once Read Many (WORM) policies are used to govern WORM file systems. When a WORM policy is attached to a WORM file system, the policy governs file system's retention periods, retention lock settings, and retention mode. Once a file system WORM policy is enabled and attached to a file system, and a file has been committed into WORM state, data can be read, but not changed until the retention periods expire. Once the retention period expires, the file's WORM status can be extended or the file can be deleted.

When a policy is locked, only Pure Technical Services can unlock the policy. When a policy is unlocked, the policy retention attributes can be modified, however the changed retention attributes will only apply to newly committed files.


WORM Policies panel displays a list of all file system WORM policies and the following information for each:

- **Name:** The policy name.
- **Enabled:** Whether the policy is enabled (**True**) or disabled (**False**).
- **Min Retention:** The minimum retention period of the WORM file system. Between 1s and 100 years, in the format: 3h, 2d, 1w. Cannot be larger than the maximum retention period.
- **Max Retention:** The maximum retention period of the WORM file system. Between 1s and 100 years, in the format: 3h, 2d, 1w. Cannot be smaller than the minimum retention period.
- **Default Retention:** The retention period used for committing files to WORM status, between the minimum and maximum retention periods, in the format: 3h, 2d, 1w. The default retention period will be applied if no access time is provided, or the access time is less than the current server time.
- **Mode:** The retention mode. Currently **compliance** retention mode, where files cannot be overwritten or deleted by any user for the duration of the retention period, is the only retention mode supported.
- **Retention Lock:** Displays whether a policy can be modified (**unlocked**) or is read only (**locked**). Locked policies can only be unlocked by Pure Technical Services.

Viewing File System WORM Policy Details

To view a detailed summary about each file system WORM policy in the array, from **WORM Policies** panel, click a policy name to open a detail page about that WORM policy.

Figure 8.2 Viewing WORM Policy Details



The **Details** panel displays the policy's name, status, and retention settings. The **Attached File Systems** displays a list all WORM file systems that have been attached to the policy.

Creating File System WORM Policies

To create a file system WORM policy, perform the following:

- 1 From the **Policies** page, click **WORM** under **File Systems**. Click the Add +) button in the heading of the **WORM Policies** panel. The **Create WORM Policy** pop-up window appears.
- 2 In the **Name** field enter the name of the WORM 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 In the **Min retention** field, enter a minimum retention period. The retention period can be a value from 1 second to 100 years, specified in seconds, hours, days, months, and years, and should be should be equal to, or less than, the maximum retention period.
- 5 In the **Max retention** field, enter the maximum retention period. The retention period can be a value from 1 second to 100 years, specified in seconds, hours, days, months, and years.
- 6 In the **Default retention** field, enter the retention period used for committing files to WORM status. The retention period can be a value from 1 second to 100 years, specified in seconds, hours, days, months, and years. The default retention period should be equal to, or larger than, the minimum retention period, and should be equal to, or less than, the maximum retention period. This default retention period will be applied if, on the client side, no access time is provided, or the access time is less than the current server time.
- 7 By default, the **Mode** is set to **compliance** retention mode, where files cannot be overwritten or deleted by any user for the duration of the retention period. This is currently the only supported retention mode.

- 8 In the **Retention lock** field, select either **unlocked** (the policy can be modified) or **locked** (the policy cannot be modified). Note that this can only be set by the system administrator. If the policy is set to locked, it can only be unlocked by Pure Technical Services.
- 9 Click **Create**.

The **WORM Policies** panel is updated with the newly created policy.

Editing File System WORM Policies

Locked policies can only be unlocked by Pure Technical Services. When a policy is unlocked, the policy retention attributes can be modified, however the changed retention attributes will only apply to newly committed files.

To edit a file system WORM policy, perform the following:

- 1 From the **Policies** page, click **WORM** under **File Systems**. Click the **Edit** button in the same row of the policy you wish to edit from the **WORM Policies** panel. The **Edit WORM Policy** pop-up window appears.
- 2 Select to enable or disable the policy, by clicking the **Enabled** checkbox.
- 3 In the **Min retention** field, enter a minimum retention period. The retention period can be a value from 1 second to 100 years, specified in seconds, hours, days, months, and years, and should be equal to, or less than, the maximum retention period.
- 4 In the **Max retention** field, enter the maximum retention period. The retention period can be a value from 1 second to 100 years, specified in seconds, hours, days, months, and years.
- 5 In the **Default retention** field, enter the maximum retention period. The retention period can be a value from 1 second to 100 years, specified in seconds, hours, days, months, and years. The default retention period should be equal to, or larger than, the minimum retention period, and should be equal to, or less than, the maximum retention period.
- 6 By default, the **Mode** is set to **compliance** retention mode, where files cannot be overwritten or deleted by any user for the duration of the retention period. This is currently the only supported retention mode.
- 7 In the **Retention lock** field, select either **unlocked** (the policy can be modified) or **locked** (the policy cannot be modified). Note that this can only be set by the system administrator.
- 8 Click **Save**.

The **WORM Policies** panel is updated with the edited policy.

Deleting File System WORM Policies

 **Note: WORM policies that are in use cannot be deleted.**

To delete file system WORM policies perform the following:

1 To delete a single file system WORM policy:

a From the **Policies** page, click **WORM** under **File Systems**. Click the **Delete** button in the same row of the policy you wish to delete from the **WORM Policies** panel. The **Delete WORM Policy** pop-up window appears.

b Click **Delete**.

The policy is removed from the **WORM Policies** panel.

2 To delete multiple file system WORM policies:

a From the **Policies** page, click **WORM** under **File Systems**. From the **WORM Policies** panel header, click **More Options > Delete**. The **Delete WORM Policies** pop-up window appears.

b From the **Available Policies** panel, select the policies you wish to delete. Each selected policy appears in the **Selected Policies** panel.

c Click **Delete**.

The policies are removed from the **WORM Policies** panel.

Management Access Policies

The management access policy provides a framework to manage permissions within a Pure ecosystem. It is designed to enhance security and fine-tune user access within Pure storage systems, such as between systems using Fusion and realms in a multi-tenancy environment. This policy transitions from a role-based access control (RBAC) model to a more granular attribute-based access control (ABAC), which is a policy-based access control.

Management access policies replace the concept of roles, which in previous releases defined user permissions. In addition to roles, access policies also define both the scope of the permissions, meaning to which realm or array the permissions apply, and an aggregation strategy, which determines the effect of multiple access policies being attached to a user or group of users. The aggregation strategy can be set either to apply all permissions of the multiple access policies or to apply a permission only if it is granted by every access policy that is assigned to the user or group.

To use management access policies, first you define the access policies you require, then you attach those policies to users or groups as appropriate. An access policy can be attached to multiple users or groups as needed and detached as needed.

Four access policies corresponding to the previous roles are built into Purity//FB, as shown in the following table. These policies are scoped to the local array.

Table 8.8 Build-in Management Access Policies and Their Roles

Built-in Management Access Policy	Current Role	Equivalent Role in Previous Releases
array_admin	admin	array_admin
storage_admin	storage	storage_admin
ops_admin	support	ops_admin
readonly	viewer	readonly

Creating a Management Access Policy

To create a management access policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Access**.
- 2 In the **Management Access Policies** panel header, click the **Add** interface button (+). The **Create Management Access Policy** appears.
- 3 In the **Name** field, enter a name for the policy.
- 4 Select the **Enable** checkbox to ensure that the policy is enabled once you are finished editing.
- 5 In the **Role** list, select the access role of the policy.
- 6 In the **Scope Type** list, select the resource type that the policy will preside over.
- 7 In the **Scope Name** field, enter the name of an array or realm.
- 8 In the **Aggregation Strategy** field, select **all-permissions** or **least-common-permission**.
- 9 Click **Create**.

Deleting a Management Access Policy

To delete a management access policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Access**.

- 2 In the **Management Access Policies** panel, click the **More Options** button and select **Delete**. The **Delete Management Access Policies** pop-up window appears.
- 3 Select one or more policies from the **Available Management Access Policies** panel.
- 4 Verify the selected policies listed in the **Selected Management Access Policies** panel to be deleted.
- 5 Click **Delete**.

Enabling and Disabling a Management Access Policy

To enable or disable a management access policy, perform the following:

- 1  **CAUTION:** Disabling a policy may result in users and groups losing access to permissions granted by the policy.

From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, select **Management Access**.

- 2 In the **Management Access Policies** panel, click the **More Options** button located at the end of the row of the policy you would like to disable then select **Disable**. The **Disable Management Access Policy** pop-up appears.
- 3 Click **Disable**.

Renaming a Management Access Policy

To rename a management access policy, perform the following:

- 1 From the **Policies** page navigation sidebar, click **Management Access** under **Administrative Access**.
- 2 In the **Management Access Policies** panel, find the policy you want to rename and click the **More Options** menu to the right of the policy, then select **Rename**. The **Rename Management Access Policy** pop-up window appears.
- 3 In the New Name field, enter the new name you want to apply to the management access policy.
- 4 Click **Rename**.

Management Authentication Policies

FlashBlade management authentication policies provide granular control over how management users authenticate to the FlashBlade. Management authentication policies allow administrators to define the security factors (password, SSH key, SSH certificate, FIDO2 key) that are required or permitted for management access.

FlashBlade includes a default management authentication policy, where the local array is its member. Management authentication behaves as if all interfaces are configured to allow "default" authentication methods if there is no active management authentication policy at the array level. The default management authentication policy can be modified or replaced to change the baseline policy for all users who do not have a different policy attached to them.

Additional management authentication policies can be created and attached to individual users or to the entire array. Only one policy can be attached to a user or array at a time. The default management authentication policy is applied if no explicit policy is attached to a user.

Management authentication policies define authentication settings for each interface (currently SSH only). Policies specify either:

- **required methods** - all specified methods must be used for authentication.
- **allowed methods** - any of the specified methods can be used for authentication.

Supported authentication methods include:

- **password** - users must enter a password.
- **key** - users can use an SSH key (including FIDO2 keys) or SSH certificate.
- **default** - a placeholder where users can use whichever authentication methods supported by FlashBlade for this interface and version. Note that default can only be used for "allowed" methods.

The default management authentication policy allows SSH key, FIDO2 key, SSH certificate, or password for authentication.

Management authentication policies are created, displayed, and managed from the **Management Authentication Policies** panel. To navigate to the **Management Authentication Policies** panel, click **Policies > Management Authentication**.

The **Management Authentication Policies** panel displays all configured management authentication policies. The following information is displayed:

- **Name:** The policy name.
- **Enabled:** If the policy is enabled (True) or disabled (False).
- **Interface:** The interface type. Currently only SSH is supported.
- **Allowed Methods:** The policy's allowed authentication methods.

- **Required Methods:** The policy's required authentication methods.

Clicking a policy name in the **Management Authentication Policies** panel opens the details page for that policy. The policy's details page includes a **Details** panel and a **Members** panel.

- The **Details** panel provides summary information about the policy. It also provides functionality to edit, rename, enable, and disable that policy.
- The **Members** panel displays the members to which the policy is attached. It also provides functionality to add arrays and users to the policy, and remove members from the policy.

Creating a Management Authentication Policy

To create a management authentication policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Authentication**.
- 2 In the **Management Authentication Policies** panel, click the Create button (+). The **Create Management Authentication Policy** pop-up window appears.
- 3 In the **Name** field, enter a name for the policy.
- 4 Select the **Enable** checkbox to ensure that the policy is enabled once it is created.

 **Note:** You can configure either **SSH Allowed Methods** or **SSH Required Methods** for authentication, but not both. Perform either step 5 or step 6.
- 5 In the **SSH Allowed Methods** field, click the Add (+) button. The **Add Allowed Methods** pop-up window appears.
 - a Select the authentication method(s) from the **Available Methods** panel.
 - b Click **Add**.
- 6 In the **SSH Required Methods** field, click the Add (+) button. The **Add Required Methods** pop-up window appears.
 - a Select the authentication method(s) from the **Available Methods** panel.
 - b Click **Add**.
- 7 Click **Create**.

Adding Members to a Management Authentication Policy

Arrays and users can be added (attached) to a management authentication policy.

To add members to a management authentication policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Authentication**.
- 2 From the **Management Authentication Policies** panel, click the name of the policy to which you wish to add a member. The details page for the policy appears.
- 3 In the **Members** panel header, click **More Options > Add Array** or **More Options > Add User**. The **Add Array** or **Add User** pop-up window appears.
- 4 Enter the array name or the user name you wish to attach to the management authentication policy.
- 5 Click **Add**.

Removing Members from a Management Authentication Policy

Arrays and users can be removed from a management authentication policy.

To remove members from a management authentication policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Access**.
- 2 From the **Management Authentication Policies** panel, click the name of the policy from which you wish to remove a member. The details page for the policy appears.
- 3 In the **Members** panel header, click **More Options > Remove Members**. The **Remove Members** pop-up window appears.
- 4 From the **Available Members** panel, select one or more members you wish to remove from the policy.
- 5 Click **Remove**.

Editing a Management Authentication Policy

To edit a management authentication policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Authentication**.
- 2 In the **Management Authentication Policies** panel, click **More Options > Edit** in the same row as the policy you wish to edit. The **Edit Management Authentication Policy** pop-up window appears.
- 3 Select or deselect the **Enable** checkbox to enable or disable the policy once you have completed editing.

 **Note:** You can configure either **SSH Allowed Methods** or **SSH Required Methods** for authentication, but not both. Perform either step 4 or step 5.

- 4 In the **SSH Allowed Methods** field, perform the following:
 - a To remove an allowed authentication method, click the Remove (x) button next to the listed authentication method.
 - b To add allowed authentication methods, click the Add (+) button. The **Add Allowed Methods** pop-up window appears. Select the authentication method(s) from the **Available Methods** panel and click **Add**.
- 5 In the **SSH Required Methods** field, perform the following:
 - a To remove a required authentication method, click the Remove (x) button next to the listed authentication method.
 - b To add required authentication methods, click the Add (+) button. The **Add Required Methods** pop-up window appears. Select the authentication method(s) from the **Available Methods** panel and click **Add**.
- 6 Click **Save**.

Renaming a Management Authentication Policy

To rename a management authentication policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Authentication**.
- 2 In the **Management Authentication Policies** panel, click **More Options > Rename** in the same row as the policy you wish to rename. The **Rename Management Authentication Policy** pop-up window appears.
- 3 Enter the policy's new name.
- 4 Click **Rename**.

Enabling and Disabling Management Authentication Policies

By default, the enable option is selected during management authentication policy creation.

When a management authentication policy is disabled, all attached members will use the default management authentication policy to log into the array.

To enable or disable a management authentication policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Authentication**.

- 2 In the **Management Authentication Policies** panel, click **More Options > Enable** or **More Options > Disable** in the same row as the policy you wish to enable or disable. The **Enable Management Authentication Policy** or **Disable Management Authentication Policy** pop-up window appears.
- 3 Click **Enable** or **Disable**.

Deleting a Management Authentication Policy

To delete a management authentication policy, perform the following:

- 1 From the **Policies > ADMINISTRATIVE ACCESS** sub-panel, click **Management Authentication**.
- 2 In the **Management Authentication Policies** panel header, click **More Options > Delete**. The **Delete Management Authentication Policies** pop-up window appears.
- 3 Select one or more policies from the **Available Management Authentication Policies** panel.
- 4 Verify the selected policies listed in the **Selected Management Authentication Policies** panel to be deleted.
- 5 Click **Delete**.

Password Policies

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

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

The **Details** panel displays the following:

- **Name:** The name of the password policy.
- **Enabled:** Whether the policy is enabled (**True**) or disabled (**False**).
- **Minimum Password Length:** The minimum allowable password length for the array, from 1 to 100 characters.
- **Maximum Login Attempts:** The maximum number of login attempts before lockout, from 1 to 100. If 0 is displayed, this rule is disabled.

- **Lockout Duration:** The lockout duration after the maximum number of login attempts, from 1 second to 90 days in seconds, minutes, hours, or days. If 0 is displayed, this rule is disabled.
- **Password History:** The number of passwords tracked to prevent password reuse. If 0 is displayed, this rule is disabled.
- **Minimum Password Age:** The minimum duration before the password can be changed from its last change.
- **Maximum Password Age:** The maximum duration before the password must be changed.
- **Minimum Character Groups:** The minimum number of character groups required in the password. There are four character groups:
 - Group 1: a-z
 - Group 2: A-Z
 - Group 3: 0-9
 - Group 4: Special characters
- **Minimum Characters Per Group:** The minimum number of characters per group to count the group as present.
- **Enforce Username Check:** Checks whether the password contains the username is enabled (**True**) or disabled (**False**).
- **Enforce Dictionary Check:** Checks whether the password is from a dictionary of known compromised passwords is enabled (**True**) or disabled (**False**).

Editing Password Policies

To edit a password policy, perform the following:

- 1 From the **Policies > Details** panel, click **More Options** from the panel header and then select **Edit Policy**. The **Edit Password Policy** pop-up window appears.
- 2 In the **Name** field, optionally enter a new policy name. By default, the policy name is **management**.
- 3 Select the **Enable** checkbox to ensure that the policy is enabled once you are finished editing.
- 4 In the **Minimum Password** field, enter the minimum character length for the password, from 1 to 100.

- 5 In the **Maximum Login Attempts** field, enter the maximum allowable number of login attempts before lockout occurs, from 1 to 100. If 0 is entered, the maximum login attempts rule is disabled.
- 6 In the **Lockout Duration** field, enter the lockout time once the maximum number of login attempts is reached, from 1 second to 90 days. Values can be entered in seconds, minutes, hours, and days. If 0 is entered, the lockout duration rule is disabled.
- 7 In the **Password History** field, enter the number of passwords tracked to prevent reuse of old passwords, from 1 to 64. If 0 is entered, the password history rule is disabled.
- 8 In the **Minimum Password Age** field, enter the minimum duration before a password can be changed from its last change, from 0 hours to 1 week. Values can be entered in hours and days. For example, 1d.
- 9 In the **Maximum Password Age** field, enter the maximum age of a password before it must be changed, from 1 day to 99,999 days. Values can be entered in hours, days, week, and years. For example, 3w.
- 10 In the **Minimum Character Groups** field, enter the minimum number of character groups required to be present in the password. There are four character groups:
 - 1: a-z character group. Entering **1** indicates the password must contain one or more members of the a-z character group.
 - 2: A-Z character group. Entering **2** indicates the password must contain one or more members of the A-Z and a-z character groups.
 - 3: 0-9 character group. Entering **3** indicates the password must contain one or more members of the 0-9, A-Z, and a-z character groups.
 - 4: special character group. Entering **4** indicates the password must contain one or more members of the special characters, 0-9, A-Z, and a-z character groups.
- 11 In the **Minimum Characters Per Group** field, enter the minimum number of characters allowed per character group to count the group as present. For example, if **1** is specified in the **Minimum Character Groups** field, entering **2** in the **Minimum Characters Per Group** field indicates that two characters from character group 1 (a-z) must be used in the password.
- 12 Select the **Enforce Username Check** checkbox to check if the password contains the username.
- 13 Select the **Enforce Dictionary Check** checkbox to check if the password appears in a dictionary of known compromised passwords.
- 14 Click **Save**.

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

Disabling Password Policies

Password policies can be disabled at any time.

To disable a password policy, perform the following:

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

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

Enabling Password Policies

To enable a password policy, perform the following:

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

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

SSH Certificate Authority Policies

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

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

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

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

Adding a Public Key

To add a public key, perform the following:

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

Creating an SSH Certificate Authority Policy

To create a policy, perform the following:

- 1 From the **Policies > SSH Certificate Authority** tab, click the **Add (+)** button in the heading of the **SSH Certificate Authority** list. The **Create SSH Certificate Authority Policy** pop-up window appears.
- 2 In the **Name** field enter the name of the SSH certificate policy.
- 3 The **Enabled** check box is enabled by default. Uncheck the box to create a disabled policy. You can enable it at a later time.
- 4 The **Static Authorized Principals** field is optional. If a name is entered then any time the user logs in with the certificate, the principal in the certificate must also match the entered name. If no static authorized principals are set, then successful login will require that the SSH certificate has the username being logged in to as one of its principals. Note that you can add multiple name/principals in the static authorized principal list.
- 5 In the **CA Certificate** field select the CA certificate you wish to use for the policy or import a new CA certificate by clicking **Import CA Certificate** beneath the field. Leave this field blank if this policy is intended to be used with a CA public key
- 6 In the **CA Public Key** field select the CA public key you wish to use for the policy or import a new CA certificate by clicking **Add Public Key** beneath the field. Leave this field blank if this policy is intended to be used with a CA certificate.
- 7 Click **Create**.

Adding Users to an SSH Certificate Authority Policy

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

- 1 From the **Policies > SSH Certificate Authority** tab, click the policy which you want the user added to.
- 2 In the **Members** panel, click **Add User...** from the menu tree. The **Add User** pop-up window will appear.
- 3 In the **Name** field enter a username and click **Add**. The username will now be listed under the **Members** panel.

Adding an Array to an SSH Certificate Authority Policy

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

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

- 1 From the **Policies > SSH Certificate Authority** tab, click the policy which you want the array added to.
- 2 In the **Members** panel, click **Add Array...** from the menu tree. The **Add Array** pop-up window will appear.
- 3 In the **Name** field enter a username and click **Add**. The array will now be listed under the **Members** panel. The current name of the array is listed by default.


Network Access Policies

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

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

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

Figure 8.3 Network Access Policies Panel



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

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

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

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

Figure 8.4 Viewing Network Access Policy Rules

Network Access Rules				1-8 of 8	+	⋮
Effect	Client	Interfaces	Index			
All						
deny	213.145.2.15	local-network-superuser-password-access, management-rest-api, management-ssh, management-web-ui, snmp	1			⋮
deny	192.120.0.0/16	snmp	2			⋮
allow	10.0.0.0/8	local-network-superuser-password-access, management-rest-api, management-ssh, management-web-ui, snmp	3			⋮
allow	*	local-network-superuser-password-access	4			⋮
allow	*	management-ssh	5			⋮
allow	*	management-rest-api	6			⋮
allow	*	management-web-ui	7			⋮
allow	*	snmp	8			⋮
Details						⋮
Name default-network-access-policy						
Type network-access						
Enabled True						

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

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

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

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

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

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

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

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

Renaming a Network Access Policy

To rename a network access policy, perform the following:

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

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

Adding a Network Access Rule to a Network Access Policy

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

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

- 5 At the **Interfaces** field, click the **Add (+)** button. The **Add Interfaces** pop-up window appears.
 - a Select the interface(s) to which access is being allowed or denied from the **Available Interfaces** column. Each selected interface appears in the **Selected Interface** column.
 - b Click **Add** to add the selected interface(s) and return to the **Add Network Access Rule to Policy** pop-up window.
- 6 Click **Add**.

Editing a Network Access Rule

To edit a network access rule, perform the following:

- 1 Select the policy whose rules you wish to edit in the **Network Access Policies** panel. The details page for that policy appears.
- 2 From the **Network Access Rules** panel, locate the rule you wish to modify and select **More Options** > **Edit** at the end of row in which the rule appears. The **Edit Network Access Rule** pop-up window appears.
- 3 In the **Client** field, enter the client as an IP address (IPv6 or IPv4), subnet in CIDR notation, or anonymous (*). Note that ordering of network access rules is such that deny rules will be ordered before allow rules and rules with explicit IP addresses will be ordered before anonymous (*)
- 4 Select to **Allow** or **Deny** access to the specified client. Note that deny rules are ordered before allow rules as explicit deny rules take precedent over allow rules. For example, if there are two allow rules (rule 1 and rule 2), rule 2 cannot be changed to a deny rule.
- 5 At the **Interfaces** field, perform the following to remove or add interfaces:
 - a Remove any interfaces by clicking the **Remove (x)** button next to the interface name.
 - b Add any new interfaces by clicking the **Add (+)** button and then select the interface(s) to which access is being allowed or denied from the **Available Interfaces** column. Click **Add**.
- 6 Click **Save**.

Reordering Network Access Rules

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

Note the following network access rule ordering restrictions:

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

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

To reorder a network access rule, perform the following:

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

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

Removing a Network Access Rule from a Network Access Policy

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

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

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

TLS Policies

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

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

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


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

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

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

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

Figure 8.5 Viewing TLS Policies



Policies

Search Policies

FILE SYSTEMS

- Audit
- NFS Export
- QoS
- SMB
- Snapshot
- WORM

OBJECT STORE

- Object Store Access

ADMINISTRATIVE ACCESS

- Password
- SSH Certificate Authority

NETWORK SECURITY

- Network Access
- TLS**

TLS Policies 11 of 1


Name	Appliance Certificate	Minimum TLS Version	Enabled Ciphers	Disabled Ciphers
default-tls-policy	global	default	default	-

TLS Defaults

Appliance Certificate	Minimum TLS Version	Enabled Ciphers	Disabled Ciphers
-	1.2	ECDHE-ECDSA-AES256-GCM-SHA384 ECDHE-RSA-AES256-GCM-SHA384 DHE-RSA-AES256-GCM-SHA384 ECDHE-ECDSA-CHACHA20-POLY1305 ECDHE-RSA-CHACHA20-POLY1305 DHE-RSA-CHACHA20-POLY1305 ECDHE-ECDSA-AES256-CCM8 ECDHE-ECDSA-AES256-CCM DHE-RSA-AES256-CCM8 DHE-RSA-AES256-CCM ECDHE-ECDSA-ARIA256-GCM-SHA384 ECDHE-ARIA256-GCM-SHA384 DHE-RSA-ARIA256-GCM-SHA384 ECDHE-ECDSA-AES128-GCM-SHA256 ECDHE-RSA-AES128-GCM-SHA256 DHE-RSA-AES128-GCM-SHA256 ECDHE-ECDSA-AES128-CCM8 ECDHE-ECDSA-AES128-CCM DHE-RSA-AES128-CCM8 DHE-RSA-AES128-CCM	-

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

Figure 8.6 Viewing TLS Policy Summary Details



Details

Name	Appliance Certificate	Minimum TLS Version	Enabled TLS Ciphers	Disabled TLS Ciphers
tls_policy36	vipcert36	1.3	default	-

Mutual TLS Settings

Client Certificate Required	Verify Client Certificate Trust	Trusted Client CA(s)	Client CA Type
False	False	-	-

Members 11 of 1

Name	Type
data-10-64-240-36	network-interfaces

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

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

Creating a TLS Policy

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

The following restrictions apply:

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

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

To create a custom TLS Policy, perform the following:

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

c Perform one of the following:

- From the **CA Certificate** list, select the CA certificate used to verify client authenticity.
- From the **CA Certificate Group** list, select the CA certificate group used to verify client authenticity.

8 Click **Create**.

Editing a TLS Policy

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

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

The following restrictions apply:

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

To edit a TLS Policy, perform the following:

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

Deleting a TLS Policy

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

- 1 To delete a single TLS policy:

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

Viewing a TLS Policy's Effective Values

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

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

Setting the Array's Default TLS Policy

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

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

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

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

Attaching TLS Policies to Network Interfaces

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

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

Removing TLS Policies from Network Interfaces

- 1 To remove a TLS policy from a single network interface:
 - a Navigate to **Policies > Network Security > TLS**. In the **TLS Policies** panel, click TLS policy you wish to remove from a network interface. The details screen for that TLS policy appears.
 - b In the **Members** panel, click the delete button in the same row of the interface from which you wish to remove the TLS policy. The **Remove Member from TLS Policy** pop-up window appears.
 - c Click **Remove**.
- 2 To remove a TLS policy from multiple network interfaces:
 - a Navigate to **Policies > Network Security > TLS**. In the **TLS Policies** panel, click TLS policy you wish to remove from multiple network interfaces. The details screen for that TLS policy appears.
 - b In the header of the **Members** panel, click **More Options > Remove**. The **Remove Members from TLS Policy** pop-up window appears.
 - c Select the interfaces from which the TLS policy will be removed from the **Current Members** panel.
 - d Click **Remove**.

Telemetry Policies

Telemetry is the automated process of collecting, transmitting, and analyzing data from systems to monitor health, diagnose issues, and optimize performance. Telemetry provides early warnings of

potential problems before they impact operations, leading to improved cluster stability, enhanced performance, and faster resolution of support cases. Telemetry does not include sensitive customer or personally identifiable data, and all transmissions are encrypted and compliant with industry regulations.

A telemetry policy controls how system telemetry data is collected, managed, and transmitted to targets. It defines which metrics are gathered, how frequently, and the endpoints to which they are sent for analysis. Telemetry metrics and targets are configurable, so you can specify what information to export and where. Supported transports for metric export include secure protocols like OpenTelemetry Protocol (OTLP) over HTTP, with options for TLS encryption. Data is sent at configurable intervals or on demand. Customers can enable/disable telemetry, and disabling may affect proactive support capabilities. Telemetry excludes sensitive and personally identifiable info.

A telemetry target is a designated endpoint to which system telemetry data—such as performance metrics, alerts, diagnostic traces, and operating system details—is periodically exported for analysis, troubleshooting, and support. You configure telemetry targets to determine where telemetry information is sent and how it is transmitted. These are the configurations for a target:

- **Name** - The name you designate for a target.
- **Transport** - The OpenTelemetry Protocol (OTLP) over HTTP with data encoded in protobuf format is the protocol used for the transport of telemetry metrics. The **otlp-http-protobuf** option encrypts and secures transported data using TLS and server authenticity validated by a CA certificate or certificate groups. The **otlp-http-protobuf-without-tls** option transports unencrypted data in plain text.
- **Vendor** - The observability or monitoring platform that receives and processes telemetry data from the FlashBlade. Supported vendors are **generic-otlp**, **datadog**, **dynatrace**, or **splunk**.
- **Endpoint** - The IP address and port of the receiver of the telemetry metrics.
- **API Token** - Used to authenticate the target the metrics are being sent to. This is only required if the target need a token authentication.
- **Source** - The type of metrics to gather. Selecting **virO** will export metrics from the Fabric Module (FM) and selecting **data-vip** will export metrics from the blades.
- **CA Certificates** - Used for TLS server authenticity validation. This is only required if using TLS for the transport.
- **CA Certificates Group** - Used for TLS server authenticity validation. This is only required if using TLS for transport.

Updating an Attached Metrics Target Policy

To update a telemetry metrics target policy, perform the following:

- 1 Navigate to **Policies > Observability > Telemetry Metrics**.
- 2 In the **Telemetry Metrics Policy** panel, click **More Options > Update Attached Targets**. The **Edit Telemetry Metrics Policy** dialogue window appears.
- 3 In the **Name** field, the name is greyed out.
- 4 In the **Enabled** field, it is enabled by default.
- 5 In the **Targets** field, click the plus (+) icon and the **Add Targets** dialogue window appears.
 - a In the **Available Targets** pane on the left, select targets to add to the **Selected Targets** pane on the right.
 - b Click **Add**. The **Add Targets** dialogue window closes.
- 6 Click **Save**.

Editing a Telemetry Metrics Policy

To edit a telemetry metrics policy, perform the following:

- 1 Navigate to **Policies > Observability > Telemetry Metrics**.
- 2 In the **Attached Targets** panel, click **More Options > Update Attached Targets**. The **Edit Telemetry Metrics Policy** dialogue window appears.
- 3 In the **Name** field,
- 4 In the **Targets** field, click the plus (+) icon and the **Add Targets** dialogue window appears.
 - a In the **Available Targets** pane on the left, select targets to add to the **Selected Targets** pane on the right.
 - b Click **Add**. The **Add Targets** dialogue window closes.
- 5 Click **Save**.

Creating a Telemetry Target

To create a telemetry metrics target, perform the following:

- 1 Navigate to **Policies > Observability > Telemetry Metrics**.

- 2 In the **Telemetry Targets** panel, click the plus (+) icon. The **Create Telemetry Target** dialogue window appears.
- 3 In the **Name** field, enter a name for the target.
- 4 In the **Enabled** field, check the box to enable now or uncheck and enable at a later time.
- 5 In the **Transport** field, select **otlp-http-protobuf** or **otlp-http-protobuf-without-tls**.
- 6 In the **Vendor** field, select the vendor (**generic-otlp**, **datadog**, **dynatrace**, or **splunk**) that will receive and process telemetry data.
- 7 In the **Endpoint** field, enter an valid IP address for the target.
- 8 In the **API Token** field, enter the API token used to authenticate the target.
- 9 In the **Source** field, select the datavip source for sending telemetry data from the drop-down list. The vir0 source exports metrics from the FM. The data-vip source exports metrics from the blades.
- 10 In the **CA Certificate** field, select the CA Certificate from the drop-down list.
- 11 In the **CA Certificate Group** field, select the CA Certificate Group from the drop-down list.
- 12 Click **Create**.

Renaming a Telemetry Target

To rename a telemetry target, perform the following:

- 1 Navigate to **Policies > Observability > Telemetry Metrics**.
- 2 In the **Telemetry Targets** pane, find the target to rename and to the right click **More Options > Rename**. The **Rename Target** dialogue window appears.
- 3 In the **New Name** field, enter a new name for the target.
- 4 Click **Rename**.

Disabling a Telemetry Target

To disable a telemetry target, perform the following:

- 1 Navigate to **Policies > Observability > Telemetry Metrics**.
- 2 In the **Telemetry Targets** pane, find the target to disable and to the right click **More Options > Disable**. The **Disable Target** dialogue window appears.
- 3 Click **Disable**.

Editing a Telemetry Target

To edit a telemetry target, perform the following:

- 1 Navigate to **Policies > Observability > Telemetry Metrics**.
- 2 In the **Telemetry Targets** pane, find the target to edit and to the right click **More Options > Edit**. The **Edit Target** dialogue window appears.
- 3 In the **Enabled** field, check the box to enable now or uncheck and enable at a later time.
- 4 In the **Transport** field, select **otlp-http-protobuf** or **otlp-http-protobuf-without-tls**.
- 5 In the **Vendor** field, select the vendor (**generic-otlp**, **datadog**, **dynatrace**, or **splunk**) that will receive and process telemetry data.
- 6 In the **Endpoint** field, enter an valid IP address for the target.
- 7 In the **API Token** field, enter the API token used to authenticate the target.
- 8 In the **Source** field, select the datavip source for sending telemetry data from the drop-down list. The vir0 source exports metrics from the FM. The data-vip source exports metrics from the blades.
- 9 In the **CA Certificate** field, select the CA Certificate from the drop-down list.
- 10 In the **CA Certificate Group** field, select the CA Certificate Group from the drop-down list.
- 11 Click **Save**.

Deleting a Telemetry Target

To delete a telemetry target, perform the following:

- 1 Navigate to **Policies > Observability > Telemetry Metrics**.
- 2 In the **Telemetry Targets** pane, find the target to delete and to the right click **More Options > Delete**. The **Delete Target** dialogue window appears.
- 3 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.


Additionally, tabs are also provided to view performance statistics for NFS clients, groups, and users.

Figure 9.2 Viewing the Performance Page



The Performance Chart

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



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

Latency

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

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

IOPS

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

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

Bandwidth

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

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

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

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

Displaying Protocol-Specific Metrics

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

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

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

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

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

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

Displaying File System Performance


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

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

Figure 9.3 Displaying File System Performance (NFS)



Identifying the Most Active File Systems

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


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

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

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

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

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



Note about the Performance Charts

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

In the **Dashboard** page:

- The performance charts are displayed for the array.

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

In the **Analysis** page:

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

Displaying Client, Group, and User Performance Statistics

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

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

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

Viewing NFS Client Performance Statistics

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

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

- **us/op (read)** - The average time, measured in microseconds, to process a read request.
- **us/op (write)** - The average time, measured in microseconds, to process a 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 the field below each column header, as well as sort by clicking the column header.

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

Viewing S3 Client Performance Statistics

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

To view general statistics for S3 clients:

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

The following information is displayed:

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

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

To view protocol specific S3 client performance:

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

The following information is displayed:

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

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

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

Viewing Group and User Performance Statistics

To view statistics for groups and users:

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

The following information is displayed:

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

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

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


Capacity

By default, the **Capacity** graph displays information for all 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).
- **Shared:** Physical space occupied by deduplicated data, meaning that the space is shared with other filesystems/buckets and snapshots as a result of data deduplication.
- **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) and/or deduplication for all file systems on the array when viewing the the File systems capacity graph, or for all buckets on the array when viewing the Object Store capacity graph.
- **s**
- **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.

Purity DeepReduce™

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

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 9.4 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 9.5 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 9.6 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 9.7 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 9.8 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 9.9 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 9.10 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 9.11 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 is displayed in gibibyte (GiB), gigabyte (GB), tebibyte (TiB), terabyte (TB), pebibyte (PiB), 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, FlashBlade//S, or FlashBlade//S with Expansion chassis) 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.

You can also turn on/off identity LEDs on each blade, DirectFlash Module, and fabric module from the array'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 and //E System Resiliency

FlashBlade//S and //E storage systems provide a high level of resiliency and cluster high-availability during unforeseen operational disruptions.

Single-chassis FlashBlade//S models can tolerate a minimum of 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) per chassis
- A single XFM

The minimum configuration of FlashBlade //E can tolerate a minimum of the following failures:

- Any single blade per chassis

As the cluster grows, by adding additional chassis, blades, and drives, more devices can be lost while still automatically providing a high level of durability and cluster high-availability.

For additional information about system resiliency, contact Pure Technical Services.

Alerts

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

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

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

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

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

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

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

Each alert in the list output includes the following information:

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

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

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

Informational (Info) alerts inform users of a general behavior change and require no action. For example, Purity//FB triggers an informational alert if the NFS service is unhealthy.

- **Code:** Pure Storage alert code number that identifies the type of alert event.
- **State:** Current state of the alert. Possible states include: `open`, `closed`, `closing`, and `waiting to downgrade`.

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

Certain alerts are generated due to intermittent issues. Examples include temperature sensors reporting values outside of normal operating range and space usage exceeding certain capacity thresholds. These types of alerts can change from `open` state to `closing` or `waiting to downgrade` states.

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

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

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

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

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

Flagging Alert Messages

To flag alert messages:

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

Unflagging Alert Messages

To unflag alert messages:

- 1 Select **Health > Alerts**.
- 2 Locate the alert you want to unflag.
- 3 Click the flag icon. When the icon becomes gray, the alert is unflagged.


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

Chapter 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

nova

ID

8a58a4b7-920b-49f5-87ed-bb45a8932d9e

Security Update

-

Time

2024-02-13 14:48:55 GMT-08:00 (PST)

Time Zone

America/Los_Angeles

Login Banner

528 characters, 13 lines

Idle Timeout

-

Network Access Policy

[default-network-access-policy](#)

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.

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

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

Working with System Settings

The following **System** settings are available:

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

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

Alert Watchers

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

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

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

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

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

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

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

Adding an Alert Watcher

Add an alert watcher to receive email alert notifications.

To add an alert watcher, perform the following:

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

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

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

6 Click **Add**.

Enabling and Disabling an Alert Watcher

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

- 1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.
- 2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher for whom you wish to enable or disable and select **Edit**. The **Edit Alert Watchers** pop-up window appears.
- 3 Click the **Enabled** toggle switch to set the alert watcher to enabled (blue) or disabled (gray).
- 4 Click **Save**.

Setting the Severity Level for Alert Watcher Notifications

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

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

- 1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.
- 2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher for whom you wish to set a notification severity level and select **Edit**. The **Edit Alert Watchers** pop-up window appears.
- 3 Select a severity level (**Info+**, **Warning+**, or **Critical**) from the **Notification Severity** menu.
 - **Info+**: If selected, the alert watcher will receive all alert notifications for all severity levels.
 - **Warning+**: If selected, the alert watcher will receive any alert with a severity of Warning or more severe. Info-level alert notifications will not be received.
 - **Critical**: If selected, the alert watcher will only receive the most severe Critical-level alert notifications.

- 4 Click **Save**.

Sending a Test Email

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

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

To send an test email, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.
- 2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher for whom you wish to send a test email and select **Send Test Email**.

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

Deleting an Alert Watcher

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

- 1 From the **Settings** page navigation sidebar, click **Alerts and Notifications** under **System**.
- 2 In the **Alert Watchers** panel, click the **More Options** button in the row of the alert watcher you wish to delete and select **Delete**. The **Delete Alert Watcher** pop-up window appears.
- 3 Click **Delete**.

Alert Routing

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

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

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

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

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

Configuring Alert Routing

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

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

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

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

For IPv6, specify the IP address in the form `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`). If a port number is also specified, enclose the entire address in square brackets (`[]`) and append the port number to the end of the address. For example, `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]:PORT`, where `PORT` represents the port number.

- 5 Select the encryption mode, **None** or **starttls**. If **None** is selected, encryption is Opportunistic TLS. If **starttls** is selected, encryption is enforced. Note that the external SMTP relay server must support TLS 1.2 or above.

- 6 Click **Save**.

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 11.3 Viewing the Array Details Panel

Array Details	
Name nova	
ID 8a58a4b7-920b-49f5-87ed-bb45a8932d9e	
Security Update -	
Time 2024-02-13 14:48:55 GMT-08:00 (PST)	
Time Zone America/Los_Angeles	
Login Banner 528 characters, 13 lines	
Idle Timeout -	
Network Access Policy default-network-access-policy	
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.
- **Network Access Policy** - The current network access policy governing access to the FlashBlade interfaces. Clicking the policy opens that policy's detail page where its network access rules can be viewed and edited. See [Network Access Policies](#) for additional information.
- **Data at Rest Encryption** - Indicates whether data at rest encryption is enabled or disabled. If enabled, inactive data is encrypted preventing unauthorized access.

Items displayed with an edit icon are configurable.

Renaming the Array

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

Array Time

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

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

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

Changing the Time Zone

To change the time zone of the array:

- 1 Navigate to the **Array Details** panel on the **Settings** page.
- 2 Click the edit icon next to **Time Zone**. The **Edit Array Time Zone** pop-up window appears.
- 3 Select the new time zone.
- 4 Click **Save**.

Creating and Editing a Login Banner

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

To create a login banner, perform the following:

- 1 Navigate to the **Array Details** panel on the **Settings** page.
- 2 Click the edit icon next to **Login Banner**. The **Edit Login Banner** pop-up window appears.
- 3 Enter the login banner text. The limit is 8,000 characters.
- 4 Click **Save**.

Setting the Idle Timeout Period

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

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

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

- 1 Navigate to the **Array Details** panel on the **Settings** page.
- 2 Click the edit icon next to **Idle Timeout**. The **Edit Idle Timeout** pop-up window appears.
- 3 Enter the timeout value. The timeout value can be set from five minutes to three hours and must be specified in seconds, minutes, or hours. For example, 300s, 15m, 2h, etc.
- 4 Click **Save**.

Attached Policies

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

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

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

SNMP

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

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

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

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

Table 11.1 System Variables

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

Table 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)

Table 11.2 Array Performance Variables (continued)

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

Table 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	Hardware component name
pureHardwareStatus	iso.3.6.1.4.1.40482.5.1.2	Hardware component status

Table 11.3 Hardware Variables (continued)

Variable	OID	Description
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)
pureDNSName	iso.3.6.1.4.1.40482.6.1.1	DNS name
pureDNSDomain	iso.3.6.1.4.1.40482.6.1.2	DNS domain
pureDNSServers	iso.3.6.1.4.1.40482.6.1.3	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	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

Table 11.5 LAG Variables (continued)

Variable	OID	Description
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

Table 11.8 Network Interface Variables (continued)

Variable	OID	Description
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
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

Table 11.10 Network Performance Variables (continued)

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

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

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

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

Uploading the MIB on the manager side will translate numeric OIDs to variable names when using GET-type SNMP requests. A query to retrieve the pureHardwareName value for CH1 component could then use the path, `pureHardwareName.'`CH1`'`, as opposed to the aforementioned numbered OID path, `.1.3.6.1.4.1.40482.5.1.1.1.'`CH1`'`.

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

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

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

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

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

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

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

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

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

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

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

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

To configure authPriv security, in the **Auth Protocol** field, set the authentication protocol to MD5 or SHA, and in the **Auth Passphrase** field, enter an authentication passphrase. Also, in the **Privacy Protocol** field, set the privacy protocol to AES5 or DES, and in the **Privacy Passphrase** field, enter a privacy passphrase.

Note that privacy cannot be configured without authentication

Once an SNMP manager object is created on the array, the FlashBlade immediately starts transmitting SNMP messages and alerts to the manager.

The **SNMP** panel displays configured SNMP managers. The **SNMP** panel also allows you to configure and manage SNMP managers and agents. From the panel you can also download the SNMP MIB.

To access the **SNMP** panel, from the **Settings** page navigation sidebar, click **Monitoring** under **System**.

Figure 11.4 Viewing the SNMP Panel

SNMP		
No SNMP managers found.		
Syslog Server Connections		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 11.5 Viewing the Syslog Server Connections Panel

SNMP		
No SNMP managers found.		
Syslog Server Connections		1-1 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 Select the service, **data-audit** or **management**.

6 (Optional) Click the arrow in the **Source** field and select the network interface for the syslog server. By default the management network vip is used for syslog server traffic. If configuring for data-audit service, it is recommended to select a data network interface which can be used for routing audit traffic to reduce network resource contention with the management network interface.

7 Click **Create**.

The **Syslog Server Connections** panel is updated with the new syslog server

When establishing TLS communication, each remote syslog servers' certificate must have been signed by the configured CA certificate, or by one of the CA certificates in the configured CA certificate group. See [Selecting CA Certificates for Syslog Server Connections](#). Additionally, your firewall must be set to allow in-bound and out-bound traffic.

Selecting CA Certificates for Syslog Server Connections

For TLS connections, you must select a CA certificate for the remote syslog server.

If you are setting up TLS connections for 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 Select the service, **data-audit** or **management**.
- 5 (Optional) Click the arrow in the **Source** field and select the network interface for the syslog server. By default the management network vip is used for syslog server traffic. If configuring for data-audit service, it is recommended to select a data network interface which can be used for routing audit traffic to reduce network resource contention with the management network interface.
- 6 Click **Save**.

The **Syslog Server Connections** panel is updated with your changes.

Deleting a Syslog Server Connection

To delete a remote syslog server connection, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Monitoring** under **System**.
- 2 In the **Syslog Server Connections** panel click the **Delete** button to the right of the remote syslog server whose connection you wish to delete. The **Delete Syslog Server Connection** pop-up window appears.
- 3 Click **Delete**.

The **Syslog Server Connections** panel is updated with your changes.

NTP Servers

The **NTP Servers** panel displays the hostnames or IP addresses of the Network Time Protocol (NTP) servers that are currently being used by the array to maintain reference time. The installation technician sets the proper time zone for an array when it is installed. During operation, arrays maintain time synchronization by interacting with the NTP server.

To access the **NTP Servers** panel, from the **Settings** page navigation sidebar, click **NTP Servers** under **System**.

Figure 11.6 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.

- 1 From the **Settings** page navigation sidebar, click **NTP Servers** under **System**.
- 2 In the **NTP Servers** panel:
 - To add an NTP server, in the header of the **NTP Servers** panel, click the **Add (+)** button. The **Add NTP Server** pop-up window appears. Type the hostname or IP address of the NTP server, and then Enter multiple servers as comma-separated values.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from 0 to 255 representing a group of 8 bits. For IPv6, specify the IP address in the form `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where `xxxx` is a hexadecimal number representing a group of 16 bits. When specifying an IPv6 address, consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).
 - To remove an NTP server, click the **Delete** button in the same row as the NTP server you wish to remove.

Support

The **Support** panel displays and manages the features used to communicate with Pure Storage Support, the **Challenge Response Verification Key** panel, which allows you to view and update the FlashBlade's signed verification key, and the Maintenance Windows panel, which allows you to schedule a maintenance window for your array.

To access the panels, from the **Settings** page navigation sidebar, click **Support** under **System**.

Figure 11.7 Viewing the Support Panels

Pure1 Support

Phone Home

Enabled

Remote Assist

Disconnected

Proxy Server

-

Support Logs

Download

Challenge Response Verification Key

Name	Key ID
access	1

Maintenance Windows

No maintenance windows found.

Note that if Remote Assist (RA) is enabled and a custom duration has been set for the RA session, the session expiration date and time will be displayed.

Phone Home

The phone home facility allows the array to transmit log and diagnostic information to Pure Storage Support via a secure network connection.

The phone home facility can be enabled and disabled at any time. The current phone home status appears just below the **Phone Home** label.

Enabling and Disabling Phone Home

Enabling phone home functionality allows the array to transmit log and diagnostic information to Pure Storage Support. It is highly recommended that you enable phone home.

To enable (or disable) phone home, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.
- 3 Select the **Phone Home Enabled** checkbox to enable.
- 4 Click **Save**.

Remote Assist

In some cases, the most efficient way for Pure Storage Support to service a FlashBlade array or diagnose problems is through direct access to the array. A remote assistance (RA) session grants Pure Storage Support direct and secure access to the array through a reverse tunnel which you, the administrator, open. This is a two-way communication.

Opening an RA session gives Pure Storage Support the ability to log into the array, effectively establishing an administrative session. Once the RA session is successfully established, the array returns connection details, including the date and time when the session was opened, the date and time when the session expires, and the proxy status (true, if configured).

After the Pure Storage Support team has performed all of the necessary diagnostic or maintenance functions, close the RA session to terminate the connection.

RA sessions can be opened and closed at any time. The current status of the remote assistance session appears just below the Remote Assist label.

An open RA session automatically terminates (closes) after two days have lapsed.

Opening and Closing a Remote Assistance (RA) Session

A remote assistance (RA) session grants Pure Storage Support direct and secure access to the array through a reverse tunnel which you, the administrator, open. This is a two-way communication. By default an RA session lasts 24 hours.

To open or close an RA session, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.
- 3 Select the **Remote Assist Connected** checkbox to enable an RA session. Opening an RA session gives Pure Storage Support direct and secure access to the array.
- 4 Click **Save**.

Configuring a Custom Duration for a Remote Assist Session

A remote assistance (RA) session grants Pure Storage Support direct and secure access to the array through a reverse tunnel which you, the administrator, open. This is a two-way communication. By default, an RA session lasts 24 hours. You can optionally set a custom duration for the RA session. The maximum duration is 14 days.

The option to configure a custom duration is only available if RA session is enabled. Additionally, if an RA session is currently running, you cannot set a custom RA duration (the option to set a custom duration will not be visible).

To set a custom duration for an RA session, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.
- 3 In the **Duration** field, enter the desired RA session duration.
- 4 Click **Save**.

Proxy Server

The proxy hostname, if set, represents the server to be used as the HTTP or HTTPS proxy. The format for the proxy host name is `http(s)://hostname:port`, where `hostname` is the name of the proxy host, and `port` is the TCP/IP port number used by the proxy host.

Configuring the Proxy Host Name

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.

- 3 In the **Proxy Server** field type the proxy host name.

The format for the host name is `http(s)://hostname:port`, where `hostname` is the name of the proxy host, and `port` is the TCP/IP port number used by the proxy host.

- 4 Click **Save**.

Deleting the Proxy Host Name

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the header of the **Pure 1 Support** panel click the **Edit** button. The **Edit Pure1 Support Settings** pop-up window appears.
- 3 In the **Proxy Server** field, delete the proxy host name.
- 4 Click **Save**.

Support Logs

Purity//FB continuously logs a variety of array activity, including performance metrics, hardware and software operations, and administrative actions. The logs contain the activity of both fabric modules and all blades. The time stamp of each log is based on the array time.

If Phone Home is enabled, the logs are periodically transmitted to Pure Storage. The logs are also saved to the array with up to 30 days' worth of activity available for manual download.

When downloading support logs, specify an event start and end date and time. The array generates a file containing a chronological list of array activity that occurred leading up to, during, and a little after the specified event time. Log files are password encrypted and can only be opened by Pure Storage Support.

Downloading Support Logs

Purity//FB continuously logs a variety of array activity. These logs contain the activity of both fabric modules and all blades. These logs are used by Pure Storage Support when assessing the array's performance and for diagnosing and troubleshooting any events or issues.

To download logs, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the **Support Logs** section of the **Pure1 Support** panel, click **Download**. The **Download Support Logs** pop-up window appears.

- 3 Specify the event date and time, representing the approximate array time the activity of interest occurred, and the duration of time for which you wish log information to be collected (1 to 6 hours). Note that the time is based on your local browser's time zone.
- 4 Optionally select the component(s) for which you wish to prepare support logs.

Figure 11.8 Downloading Support Logs for Single-Chassis Systems

The dialog box is titled "Download Support Logs" and contains the following fields and controls:

- Start time:** 2024/03/04 at 05 : 33 PM GMT-08:00 (PST).
- End time:** 2024/03/04 at 06 : 33 PM GMT-08:00 (PST).
- Requested time range in UTC:** 2024-03-05 01:33 - 2024-03-05 02:33
- Components:** A dropdown menu showing "CH1".
- Buttons:** "Prepare all logs" (blue) and "Close" (grey).

Figure 11.9 Downloading Support Logs for Multi-Chassis Systems

The dialog box is titled "Download Support Logs" and contains the following fields and controls:

- Start time:** 12/31/2023 at 11 : 25 AM CET.
- End time:** 12/31/2023 at 12 : 00 PM CET.
- Requested time range in UTC:** 12/31/2023 10:25 - 12/31/2023 11:00
- Components:**
 - ☐ XFM1 ☐ XFM2
 - ☒ CH1 (expanded):
 - ☒ CH1.FM1 ☐ CH1.FM2
 - ☐ CH1.FB1 ☐ CH1.FB2 ☐ CH1.FB3 ☐ CH1.FB4 ☐ CH1.FB5 ☐ CH1.FB6
 - ☐ CH1.FB7 ☐ CH1.FB8 ☐ CH1.FB9 ☐ CH1.FB10
 - ☐ CH2
 - ☒ CH3 (expanded):
 - ☒ CH3.FM1 ☒ CH3.FM2
 - ☒ CH3.FB1 ☒ CH3.FB2 ☒ CH3.FB3
- Buttons:** "Prepare all logs" (blue) and "Close" (grey).

- 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**.

Maintenance Windows

The **Maintenance Windows** panel displays the start and end times of any in-progress maintenance window. The **Name** column displays the maintenance type; currently `array` maintenance is supported. The panel also provides the options to modify the duration of an in-progress maintenance window and to delete an in-progress maintenance window.

Figure 11.10 Viewing the Maintenance Windows Panel

Maintenance Windows				+
Name	Started	Expires		
array	2024-08-29 16:35:12	2024-08-29 20:35:12		

Array users with the roles of `array_admin` and `ops_admin` can create and delete maintenance windows for the array. Array users with the roles of `storage_admin` and `ops_admin` can view/list maintenance windows for the array.

Scheduling a Maintenance Window

A maintenance window is initiated upon completing and submitting the scheduling request. Maintenance windows cannot be scheduled for a future time.

To schedule and initiate a maintenance window for your array, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the header of the **Maintenance Windows** panel click the add button (+). The **Schedule Maintenance Window** pop-up window appears.
- 3 In the **Duration** field, enter the desired maintenance window, from 1 to 48 hours. By default, the maintenance window is 4 hours.
- 4 Click **Schedule**.

The **Maintenance Windows** panel is updated with this window's start and end times, and the maintenance window begins.

Modifying a Maintenance Window

An in-progress maintenance window can be modified at any time before it expires. When modifying an in-progress maintenance window, a new maintenance window duration is applied from the maintenance window's current time, effectively updating the maintenance window. This is described in the following example:

- A 5 hour maintenance window was is at 13:45:45. Its expiration time is 18:45:45.
- At 14:45:45 the maintenance window is modified with a new duration of 1 hour.
- The new expiration time of the maintenance window is 15:45:45.

To modify an in-progress maintenance window, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the **Maintenance Windows** panel click the **Edit** button in the same row of the listed maintenance window. The **Update Maintenance Window** pop-up window appears.
- 3 In the **Duration** field, enter the desired duration that you wish to update the original maintenance window, from 1 to 48 hours. By default, the maintenance window is 4 hours.
- 4 Click **Update**.

The in-progress maintenance window's duration is updated and the **Maintenance Windows** panel is updated to reflect the time the maintenance window was modified and the new expiration time.

Deleting a Maintenance Window

An in-progress maintenance window can be stopped and deleted at any time before it expires.

To stop and delete an in-progress maintenance window, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Support** under **System**.
- 2 In the **Maintenance Windows** panel click the **Delete** button in the same row of the listed maintenance window. The **Delete Maintenance Window** pop-up window appears.
- 3 Click **Delete**.

The in-progress maintenance window is stopped and deleted. The **Maintenance Windows** panel is cleared.

Working with Hardware Settings

The following **Hardware** settings are available:

- **Blades** - Displays information about all blades in the array.
- **Drives** - Displays information about all blades in the array.
- **Hardware** - Displays information about all hardware components in the array.

Viewing Blade Information

The **Blades** panel displays information about each blade used in the array, including status and capacity.

To access the **Blades** panel, from the **Settings** page navigation sidebar click **Blades** under **Hardware**.

Viewing Drive Information

The **Drives** panel displays information about each blade used in the array, including status and capacity.

To access the **Drives** panel, from the **Settings** page navigation sidebar click **Drives** under **Hardware**.

Viewing Array Hardware Information

The **Hardware** panel displays information about all hardware components used in the array. By default, general information about each component is displayed, which includes information such as the hardware type, status, and location. Clicking **Specifications** at the top of the **Hardware** panel displays each component's model number, part number, and serial number.

To access the **Hardware** panel, from the **Settings** page navigation sidebar click **Hardware** under **Hardware**.

Working with Network Settings

The following **Network** settings are available:

- **Connectors** - Displays and manages array network interfaces.
- **Diagnostics** - Allows you to conduct network pings and traceroutes.
- **DNS** - Displays and manages DNS settings.
- **Object Store Virtual Hosts** - Displays and manages virtual hosts for object stores.
- **Link Aggregation Groups** - Displays and manages subnet link aggregation groups (LAGs).
- **Subnets & Interfaces** - Displays and manages all subnets and interfaces used for file exports.

Connectors

The **Connectors** panel displays and allows you to edit the array's connection interfaces.

To access the **Connectors** panel, from the **Settings** page navigation sidebar, click **Connectors** under **Network**. The **Connectors** panel displays the following:

- **Name** - The connector name in the format CHx.FMx.ETHx, where CHx is the chassis number, FMx is the fabric module number, and ETHx is the ethernet port number.

Note that the Out-of-Band management ports for FlashBlade are as follows (these ports are not present in in-band clusters):

- Single-chassis first-generation FlashBlade: CH1.FM1.ETH5, CH1.FM2.ETH5

- Single-chassis FlashBlade//S: CH1.FM1.ETH9, CH1.FM2.ETH9
- Multi-chassis FlashBlade: XFM1.ETH33, XFM2.ETH33
- **Connector Type** - The connector type. Possible values include **QSFP**, **RJ-45**, and **SFP**.
- **Port Count** - The number of ports for the connector.
- **Port Speed** - The speed (in Gb/s) at which the component is operating.
- **Transceiver Type** - The transceiver type.
- **External** - Whether the connector is external (**True**) or not (**False**).

Diagnostics

Click **Diagnostics** under **Network** to access the **Diagnostics** page. The **Diagnostics** page provides two panels: **Network Diagnostics** and **Connection Statistics**.

Use the **Network Diagnostics** panel to perform network pings and traceroutes to determine the reachability of components and troubleshoot network issues.

Use the **Connection Statistics** panel to monitor TCP connections.

Monitoring TCP

TCP (Transmission Control Protocol) monitoring provides information on the state of the connection. The state of the local and remote ports and IP addresses are given so any issues can be diagnosed and debugged.

The impact of TCP states on network performance depends on factors like network congestion, packet loss, and the efficiency of handling retransmissions and acknowledgments.

The states of TCP connections can significantly impact network performance in various ways. For example, during the SYN-SENT and SYN-RECEIVED states, the initial handshake process may introduce delays, especially if there is network congestion or packet loss, resulting in retransmissions and increased latency. This can degrade performance by slowing down the establishment of connections.

In the ESTABLISHED state, TCP connections can handle data transmission efficiently. However, if the network is congested and experiences frequent packet loss, the retransmissions can affect overall throughput and latency.

In the FIN-WAIT-1 and FIN-WAIT-2 states, as connections are being terminated, delays can occur if the termination requests or acknowledgments are lost or delayed. Time spent in the TIME-WAIT state ensures proper closure and prevents issues with stray packets from old connections, but it can tie up resources and potentially impact performance if many connections are in this state simultaneously.

General network performance issues, such as high latency and low throughput, can be exacerbated by TCP states like CLOSE-WAIT and LAST-ACK, where connection termination is delayed and can affect the speed at which new connections can be established.

Use the **Connection Statistics** panel to debug connection issues. The following fields can be optionally populated to TCP connections of interest. If not populated (default) all TCP connections for all configured vips will be displayed.

- **Local Address:** The array configured data IP address. The IP address can be IPv4 or IPv6 in classful IP or CIDR format.
- **Local Port:** The array data service port number, such as for NFS or replication. Port ranges are 0-65535.
- **Remote Address:** The remote (client) IP address. The IP address can be IPv4 or IPv6 in classful IP or CIDR format.
- **Remote Port:** The remote (client) TCP port number. Port ranges are 0-65535.
- **State:** The TCP connection state. The TCP connection state can be one of the following:
 - CLOSE-WAIT: The remote side has closed the connection, and the local end needs to close it.
 - CLOSED: The connection is closed and no further communication is possible.
 - ESTABLISHED: The connection is established and communication is open.
 - ≠ESTABLISHED: All states, except ESTABLISHED, are selected.
 - FIN-WAIT-1: The socket is closing, and it has sent a connection termination request (FIN).
 - FIN-WAIT-2: The socket is waiting for a connection termination request from the remote side.
 - LAST-ACK: The socket is waiting for a final acknowledgment from the remote side.
 - LISTEN: The socket is listening for incoming connections.
 - SYN-RECEIVED: A connection request (SYN) received, and a response (SYN-ACK) is sent back.

- SYN-SENT: A connection request (SYN) is sent and waiting for a matching connection request.
- TIME-WAIT: The socket is waiting to ensure the remote TCP received the acknowledgment of its connection termination request.

Once populated, click **Update** to view the connection status.

Performing Pings and Traceroutes

Network pings and traceroutes are performed in the **Network Diagnostics** panel.

- 1 Enter the name of the blade or controller you wish to ping or run a traceroute in the **Components** field.
- 2 Enter the IP address or fully qualified domain name (FQDN) of the ping or traceroute target in the **Target** field.
- 3 (Optional) Enter the subnet, interface name, or IP address used to run the test in the **Source** field.
- 4 For pings, enter the number of times you wish to ping the target in the **Count** field and then click **Ping**.
- 5 (Optional) For traceroutes, select the **Discover MTU** checkbox and then click **Traceroute**.

The window below the ping and trace parameters displays the results of each ping and traceroute.

Figure 11.11 Viewing Results

```
CH1.FB1:
PING 123.123.123 (123.123.0.123) 56(84) bytes of data.

--- 123.123.123 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1028ms

CH1.FM1:
PING 123.123.123 (123.123.0.123) 56(84) bytes of data.

--- 123.123.123 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1001ms
```

DNS Configurations

The **DNS Configurations** panel displays and manages the DNS attributes for an array's administrative network. The array has one default DNS configuration for management services and the option to add additional configurations for multi-tenancy.

The **DNS Configurations** panel can be accessed from the **Settings** page navigation sidebar, click **DNS** under **Network**.

Figure 11.12 Viewing the DNS Configuration Panel

DNS Configurations					1 of 1 +
Name ▲	Domain	Nameservers	Services	Source	
		Any ▼	Any ▼		
management	alliance.local	10.58.40.152	management, data	-	 

- **Name** - The administrative name for this DNS configuration.
- **Domain** - The DNS domain represents the domain suffix to be appended by the array when performing DNS lookups. For example, `mydomain.com`.
- **Nameservers** - The IP address(es) of the servers to be used when resolved names for the specified domain. Up to three can be entered and they will be queried in the order entered.
- **Services** - The type of service the DNS will handle. The services can be Data and/or Management.
- **Source** - The network interface name for outbound traffic for the DNS server(s) using that particular DNS entry. The Management service is the default DNS configuration and the Source is not selectable.

Creating a DNS Configuration

To create a new DNS, perform the following:

- 1 From the **Settings** page navigation sidebar, click **DNS** under **Network**.
- 2 In the **DNS Configuration** panel, click the **Add** interface button (+). The **Create DNS Configuration** pop-up window appears.
- 3 In the **Name** field, enter the DNS name.
- 4 In the **Domain** field, enter the domain name.
- 5 In the **Nameserver 1** field, enter the name server IP address.

- 6 (Optional) In the **Nameserver 2** field, enter the server IP address.
- 7 (Optional) In the **Nameserver 3** field, enter the server IP address.
- 8 In the **Selected Source** list, select the data service network interface.
- 9 Click **Create**.

Editing DNS Configurations

To edit DNS configurations, perform the following:

- 1 From the **Settings** page navigation sidebar, click **DNS** under **Network**.
- 2 Click the edit button in the header of the **DNS Settings** panel. The **Edit DNS Configuration** pop-up window appears.
- 3 (Optional) In the **Name** field, type a new name for the DNS server.
- 4 (Optional) In the **Domain** field, type the domain suffix to be appended by the array when doing DNS lookups. For example, `mydomain.com`.
- 5 (Optional) In the **Nameserver** fields, specify up to three DNS server IP addresses for Purity//FB to use to resolve hostnames to IP addresses. Enter one IP address in each Nameserver field. Purity//FB queries the DNS servers in the order that the IP addresses are listed.
- 6 (Optional) In the **Selected Source** field, select the network interface name for outbound traffic for the server. Note that the source cannot be changed for the default DNS configuration.
- 7 Click **Save**.

Link Aggregation Groups

A link aggregation group (LAG) is a collection of Ethernet ports on the array. This concept is also known as a port channel, or a bond group. Multiple LAGs can be used to connect to separate physical networks. By default, the array is preconfigured with all ports in a single LAG named **uplink**, which cannot be deleted. Ports can be removed from the uplink LAG and added to another LAG.

The **Link Aggregation Groups** panel displays and allows management of LAGs on the array.

To access the **Link Aggregation Groups** panel, from the **Settings** page navigation sidebar, click **Link Aggregation Groups** under **Network**.

Figure 11.13 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 11.14 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.

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

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

The **Subnets** panel displays the subnets on the array. Attached to each subnet are the data virtual IP addresses (data VIPs) that are used to export file systems, which are displayed in the **Interfaces** panel.

To access the **Subnets** and **Interfaces** panels, from the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.

Figure 11.15 Viewing the Subnets & Interfaces Panel

Subnets

1-4 of 4

Name	Enabled	Prefix	VLAN	Gateway	MTU	LAG	Services	
							Any	
net1	True	10.255.8.0/24	8	10.255.8.1	1500	uplink	data, management, support	
net2	True	10.255.9.0/24	9	10.255.9.1	1500	uplink		
net3	True	2001:8::/64	8	2001:8::1	1500	uplink	management, support	
net4	True	2001:9::/64	9	2001:9::1	1500	uplink		

Interfaces

1-7 of 7

Name	Enabled	Server	Subnet	Address	VLAN	Mask	Gateway	MTU	Service	
									Any	
fm1.admin0	True	-	net1	10.255.8.21	8	255.255.255.0	10.255.8.1	1500	support	
fm1.admin0.1	True	-	net3	2001:8::21	8	ffff:ffff:ffff:ffff::	2001:8::1	1500	support	
fm2.admin0	True	-	net1	10.255.8.22	8	255.255.255.0	10.255.8.1	1500	support	
fm2.admin0.1	True	-	net3	2001:8::22	8	ffff:ffff:ffff:ffff::	2001:8::1	1500	support	
if2	True	server1	net1	10.255.8.55	8	255.255.255.0	10.255.8.1	1500	data	
vir0	True	-	net1	10.255.8.20	8	255.255.255.0	10.255.8.1	1500	management	
vir1	True	-	net3	2001:8::20	8	ffff:ffff:ffff:ffff::	2001:8::1	1500	management	

The **Subnets** panel includes the following information:

- **Name** - Subnet name.
- **Enabled** - Subnet status. When a subnet is enabled (**True**), the data VIPs within the subnet that are enabled are automatically available and ready to connect. Newly created subnets are automatically enabled.
- **Prefix** - The network prefix and prefix length.
- **VLAN** - VLAN ID number.
- **Gateway** - IP address of the gateway through which the data vip communicates with devices outside the given subnet.
- **MTU** - Maximum transmission unit (MTU) for the data vips, in bytes.
- **LAG** - Link aggregation group (LAG) for the data vips.

The **Interfaces** panel includes the following information:

- **Name** - The data VIP name.
- **Enabled** - The interface status. If the subnet to which the interface is attached is enabled, the interface status will display as enabled (**True**).
- **Server** - The server which the data VIP is attached.
- **Subnet** - The subnet to which the data VIP is attached.
- **Addresses** - IP address associated with the data VIP.
- **VLAN** - VLAN ID number inherited from the subnet.
- **Mask** - The subnet mask.
- **Gateway** - IP address of the gateway through which the specified interface is to communicate with devices outside the subnet.
- **MTU** - Maximum transmission unit (MTU) for the data VIPs, in bytes.

Create a new subnet if the desired one does not appear in the **Subnets** panel.

Creating a subnet requires a prefix and VLAN interface. Specify the IP address of the network prefix and prefix length in the form `ddd.ddd.ddd.ddd/dd` for IPv4, or

xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx for IPv6. Specify the VLAN ID to which the subnet is configured. Valid VLAN ID numbers are between 1 and 4094.

Optionally specify the gateway address in the form ddd.ddd.ddd.ddd for IPv4, or

xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx for IPv6.

Optionally specify the maximum transmission unit (MTU), in bytes, of the data VIP. If the MTU is not specified during subnet creation, the value defaults to 1500.

When a data VIP is attached to a subnet, it inherits the gateway, MTU and VLAN ID of the subnet. Likewise, the subnet inherits the `data` service type from the data VIP.

Delete a subnet if it is no longer needed. Note that exporting a file system and file system replication require at least one valid data VIP, which must be attached to a subnet. Therefore, there must always be at least one subnet (for service type `data` or `replication`) on the array.

Replication requires creating a replication VIP. Replication VIPs can be in the same or different subnets as data VIPs, and hence using different physical ports if required. The FlashBlade MTU setting applies to a subnet (VLAN). The MTU configured for replication VLAN must be the minimum MTU along the link between the source array and the target array.

Creating a Subnet

Data VIPs that share the same network prefix and gateway are grouped into the same subnet. Exporting a file system requires a data VIP, which in turn, must be attached to a subnet. Create a subnet if none of the existing ones meet your requirements.

Replication VIPs can be in the same or different subnets as data VIPs, and hence using different physical ports if required. The FlashBlade MTU setting applies to a subnet (VLAN). The MTU configured for the replication VLAN must be the minimum MTU along the link between the source array and the target array. If you notice the arrays are connected, but replication shows little throughput, ensure that the MTU configured for the replication VLAN is the correct size.

It is recommended to create separate VLANs for data and replication so that data and replication can use different MTU settings.

- 1 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 2 Click the **Add (+)** button in the header of the **Subnets** panel. The **Create Subnet** pop-up window appears.
- 3 In the **Name** field, type the name of the subnet.

- 4 In the **Prefix** field, type the IP address of the network prefix and prefix length in the form `ddd.ddd.ddd.ddd/dd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx` for IPv6.
- 5 In the **VLAN** field, specify the VLAN ID to which the subnet is configured. Valid VLAN ID numbers are between 1 and 4094.
- 6 In the **Gateway** field, type the IP address of the gateway through which the data vip communicates with devices outside this network in the form `ddd.ddd.ddd.ddd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx` for IPv6
- 7 In the **MTU** field, specify the maximum transmission unit (MTU) of the data VIP. If the MTU is not specified during subnet creation, the value defaults to 1500.
- 8 Select a Link Aggregation Group from the **LAG** list.
- 9 Click **Create**.

Editing a Subnet

To edit a subnet, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 2 In the **Subnets** panel, locate the subnet you wish to edit and click the **Edit** icon to the right of the subnet name. The **Edit Subnet** pop-up window appears.
- 3 In the **VLAN** field, specify the VLAN ID to which the subnet is configured. Valid VLAN ID numbers are between 1 and 4094.
- 4 In the **Gateway** field, type the IP address of the gateway through which the data VIP communicates with devices outside this network in the form `ddd.ddd.ddd.ddd` for IPv4, or `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx` for IPv6
- 5 In the **MTU** field, specify the maximum transmission unit (MTU) of the data vip. If the MTU is not specified during subnet creation, the value defaults to 1500.
- 6 Select a Link Aggregation Group from the **LAG** list.
- 7 Click **Save**.

Deleting a Subnet

A subnet can only be deleted if it has no interfaces.

To delete a subnet, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.

2 You can delete a single or multiple subnets.

To delete a single subnet:

- a In the **Subnets** panel, locate the subnet you wish to delete and click the **Delete** icon to the right of the subnet name. The **Delete Subnet** pop-up window appears.
- b Click **Delete**.

To delete 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**.

Adding Realm Access to a Subnet

To grant a realm access to a subnet, perform the following:

- 1 Navigate to **Settings > Subnets & Interfaces**. In the **Subnets** panel, click the subnet you want to grant a realm access to. The subnet details page appears.
- 2 In the **Accessible By** panel, click **More Options > Add**. The **Add Realms** dialogue window appears.
- 3 Select one or more realms from the left **Existing Realms** column to add to the right **Selected Realms** column.
- 4 Click **Add**.

Removing Realm Access to a Subnet

To remove a realms access to a subnet, perform the following:

- 1 Navigate to **Settings > Subnets & Interfaces**. In the **Subnets** panel, click the subnet you want to grant a realm access to. The subnet details page appears.
- 2 In the **Accessible By** panel, click **More Options > Remove**. The **Remove Realms** dialogue window appears.
- 3 Select one or more realms from the left **Existing Realms** column to add to the right **Selected Realms** column.
- 4 Click **Remove**.

Creating a Network Interface

Once you have created a subnet you can create and add network interfaces.

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

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

To create a network interface, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 2 In the **Interfaces** panel, click the Add button (+). The **Create Network Interface** pop-up window appears.
- 3 In the **Name** field, enter the interface name.
- 4 Select a subnet from the **Subnet** list.
- 5 In the address field, enter the network address.
- 6 From the **Services** list, select if the interface will be used for **data**, **replication**, or **sts**.
- 7 (Optional) From the **Select Server** list, select the desired server to be used for multi-tenancy.
- 8 Click **Create**.

The **Interfaces** list is updated with the new interface.

Editing a Network Interface

To edit a network interface, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 2 In the **Interfaces** panel, locate the 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**, **replication**, or **sts**.
- 6 Click **Save**.

Deleting a Network Interface

 **Note:** Built-in interfaces cannot be deleted. Additionally, custom interfaces in use cannot be deleted. For example, an interface with replication service cannot be deleted if there are remote FB array/S3 target connections configured.

To delete a network interface, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Subnets & Interfaces** under **Network**.
- 2 You can delete a single or multiple network interfaces.

To delete a single network interface:

- a In the **Interfaces** panel, locate the network interface you wish to delete and click the **Delete** icon to the right of the interface name. The **Delete Network Interface** pop-up window appears.
- b Click **Delete**.

To delete multiple network interfaces:

- a Click the **More Options** button in the **Interfaces** header and select **Delete**. The **Delete Network Interfaces** pop-up window appears.
- b Select one or more 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.
- **Single Sign-On** - Displays and manages single sign-on (SSO) configuration.
- **Users** - Displays and manages local and remote FlashBlade array users.

API Clients

The use of API clients is an alternate method of token authorization that is available on FlashBlade.

The **API Clients** panel displays and provides management of API clients used for authorization with the FlashBlade.

To access the **API Clients** panel, from the **Settings** page navigation sidebar, click **API Clients** under **Security**.

Figure 11.16 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 11.17 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 15:09	pureuser	puretag	delete	test	10.231.213.16	GUI		
1011	2021-01-15 14:59	pureuser	puretag	create	--port 'CHI.FMI.ETH4.CHI.FM2.ETH4' test	10.231.213.16	GUI		
1010	2021-01-15 15:57	pureuser	purenmp	delete	test	10.231.213.16	GUI		
1009	2021-01-15 15:57	pureuser	puretag	delete	test_connection1	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' arggo-snap	10.231.213.148	GUI		
1002	2021-01-15 07:05	ir	purepolicy	rule remove	--keep-for '1d' --destroy-snapshots arggo-snap	10.231.213.148	GUI		
1001	2021-01-15 06:34	ir	pureobj	secure-user access-key create	--user 'tommy@corbio'	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.

⚠ Important! The global certificate is the default certificate after setup of the system for both management GUI/REST and data (ObjectStore/S3) access. It is also the default certificate used by the Default TLS Policy. It can be updated but it cannot be removed.

Purity creates a self-signed certificate and private key when you start the system for the first time. The SSL Certificate panel allows you to view the certificate and import or export certificates and private keys. Intermediate certificates can also be imported to the system.

Imported SSL certificates must be PEM formatted (Base64 encoded), include the "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----" lines.

The **Array Certificates** panel displays imported array and self-signed certificates on the FlashBlade.

To access the **Array Certificates** panel, from the **Settings** page navigation sidebar, click **Certificates** under **Security**.

Figure 11.18 Viewing the Array Certificates Panel

Array Certificates						1-1 of 1 +
Name ▲	Common Name	Organization	Organizational Unit	Valid From	Valid To	
				All ▼	All ▼	
global	Pure Storage	Pure Storage	Pure Storage	2022-02-15	2052-03-29	⋮

The following information is displayed:

- **Name:** The name of the array certificate. Clicking the array certificate name opens the **Certificate Details** panel.
- **Common Name:** The name of the website or entity using the certificate.
- **Organization:** The name of the company or organization.
- **Organizational Unit:** The specific department within the organization.
- **Valid From:** The date from which the certificate is valid.
- **Valid To:** The date the certificate expires.

Array certificates used for signing and decryption of the SAML2 SSO service provider can also be imported from the **Array Certificates** panel.

FlashBlade allows you to create and update self-signed certificate objects, as well as create certificate signing requests based on existing certificate objects.

Self-signed certificates are digital certificates that are signed using their own private key and are a quick and cost-effective option for internal testing and development environments where external trust is not a requirement.

Certificate signing requests (CSRs) are formal requests to a Certificate Authority (CA) to issue a digital certificate containing information about the requesting entity and public key.

Importing Array Certificates

To import array certificates, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Click the Add + button in the **Array Certificates** panel header. The **Import Array Certificate** pop-up window appears.
- 3 Select the scope type.
- 4 In the **Name** field, enter the certificate name.
- 5 Complete or modify the following fields:
 - **Certificate** - Click **Choose File** and select the signed certificate. Verify the certificate is PEM formatted (Base64 encoded), include the "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----" lines.
 - **Private Key** - Click **Choose File** and select the private key.

- **Intermediate Certificate** - (Optional) Click **Choose File** and select the intermediate certificate.
- **Key Phrase** - (Optional) If the private key is encrypted with a passphrase, enter the passphrase.

6 Click **Import**.

Creating a Self-Signed Certificate

To create a self-signed certificate, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Click **More Options** > **Create Self-signed Certificate** in the **Array Certificates** panel header. The **Create Self-signed Certificate** pop-up window appears.
- 3 In the **Name** field, enter the certificate name.
- 4 Complete or modify the following fields:
 - **Key Algorithm** - Select the signature algorithm to generate the key. The default is **rsa**.
 - **Key Size** - Select the key size (bits). The default is **2048**.
 - **Country** - (Optional) Enter the two-letter ISO country code.
 - **State/Province** - (Optional) Enter the state, province, country, or region.
 - **Locality** - (Optional) Enter the city name.
 - **Organization** - (Optional) Enter your full organization name.
 - **Organization Unit** - (Optional) Enter the department or group name using the certificate.
 - **Common Name** - (Optional) Enter the FQDN or management IP address of the current array.
 - **Email** - (Optional) Enter your organization's email address.
 - **Subject Alternative Names** - (Optional) Click the **Add (+)** button and enter a comma-separated list of additional domain names or IP addresses.
 - **Days** - Enter the number of days the self-signed certificate is valid. The default is 3650 days.
- 5 Click **Create**.

Creating a Certificate Signing Request

Certificate signing requests (CSRs) can be created from an existing array certificate.

To create a certificate signing request, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Locate the array certificate for which you wish to create a certificate signing request in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Create Certificate Signing Request**. The **Create Certificate Signing Request** pop-up window appears.
- 3 Complete or modify the following fields:
 - **Country** - (Optional) Enter the two-letter ISO country code.
 - **State/Province** - (Optional) Enter the state, province, country, or region.
 - **Locality** - (Optional) Enter the city name.
 - **Organization** - (Optional) Enter the full organization name.
 - **Organization Unit** - (Optional) Enter the department or group name using the certificate.
 - **Common Name** - (Optional) Enter the FQDN or management IP address of the current array.
 - **Email** - (Optional) Enter your organization's email address.
- 4 Click **Create**.

Viewing Array Certificate Details

To view an array certificate's details, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Click the certificate whose information you wish to view in the **Array Certificates** panel. The **Certificate Details** and **Used By** panels appear.

Exporting Array Certificates

To export array certificates, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Locate the certificate you wish to export in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Export Certificate**. The **Export Certificate** pop-up window appears.
- 3 Click **Download**.

Updating Array Certificates with Importing Options

To update an array certificate with importing operations, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Locate the certificate you wish to update in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Update Certificate with Importing Operations**. The **Update Certificate with Importing Operations** pop-up window appears.
- 3 Complete or modify the following fields:
 - **Certificate** - Click **Choose File** and select the signed certificate. Verify the certificate is PEM formatted (Base64 encoded), include the " -----BEGIN CERTIFICATE----- " and " -----END CERTIFICATE----- " lines.
 - **Private Key** - Click **Choose File** and select the private key.
 - **Intermediate Certificate** - (Optional) Click **Choose File** and select the intermediate certificate.
 - **Key Passphrase** - (Optional) If the private key is encrypted with a passphrase, enter the passphrase.
- 4 Click **Update**.

Updating Array Certificates with Self-Signing Operations

To update an array certificate with self-signing operations, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Locate the certificate you wish to update in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Update Certificate with Self-signing Operations**. The **Update Certificate with Self-signing Operations** pop-up window appears.
- 3 (Optional) Select the **Generate a new key** option.
- 4 Complete or modify the following fields:
 - **Country** - (Optional) Enter the two-letter ISO country code.
 - **State/Province** - (Optional) Enter the state, province, country, or region.
 - **Locality** - (Optional) Enter the city name.
 - **Organization** - (Optional) Enter your full organization name.
 - **Organization Unit** - (Optional) Enter the department or group name using the certificate.
 - **Common Name** - (Optional) Enter the FQDN or management IP address of the current array.

- **Email** - (Optional) Enter your organization's email address.
- **Subject Alternative Names** - (Optional) Click the **Add (+)** button and enter a comma-separated list of additional domain names or IP addresses.
- **Days** - Enter the number of days the certificate is valid. The default is 3650 days.

5 Click **Update**.

Deleting Array Certificates

To delete a non-global array certificate, perform the following:

⚠ Important! The global certificate is the default certificate after setup of the system for both management GUI/REST and data (ObjectStore/S3) access. It is also the default certificate used by the Default TLS Policy. It can be updated but it cannot be removed.

- 1 From the **Settings** page navigation sidebar, click **Certificates** under **Security**.
- 2 Locate the certificate you wish to delete in the **Array Certificates** panel. Click the **More Options** button in the same row and select **Delete**. The **Delete Certificate** pop-up window appears.
- 3 Click **Delete**.


External Certificates

The **External Certificates** panel displays and allows you to import external (CA) certificates for the FlashBlade.

CA certificates are used to provide security authentication between two identities, allowing the FlashBlade to communicate securely with servers and clients.

To access the **External Certificates** panel, from the **Settings** page navigation sidebar, click **Certificates** under **Security**.

Figure 11.19 Viewing the External 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		

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 11.20 Viewing the Certificate Groups Panel

Array Certificates						11 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						11 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 11.21 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 11.22 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 11.23 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 11.24 Viewing Active Directory Account Kerberos Keytabs

Kerberos Keytabs					
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-2.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 11.25 Viewing the Directory Service Panel



Directory Service			
Configuration		Role Mappings	
Enabled	True	Name	Role
URIs	ldap://rockhopper.fbqa.purestorage.com	Group	Group Base
Base DN	DC=fbqa,DC=purestorage,DC=com	array_admin	array_admin
Bind User	pureuser	ops_admin	ops_admin
Bind Password	****	readonly	readonly
CA Certificate	-	storage_admin	storage_admin
CA Certificate Group	-		
User Login Attribute	-		
User Object Class	-		

When the **Directory Service** panel's user roles and directory service protocols are configured and enabled, the FlashBlade leverages a directory service to perform user account and permission level searches.

The FlashBlade supports a single local user, named `pureuser`, with array-wide (`array_admin`) permissions. Users can be added to the array through Lightweight Directory Access Protocol (LDAP) by integrating the array with a directory service such as Active Directory or OpenLDAP and assigning roles for one or more groups. Users can sign into the array only if they are in a group that has been assigned a role.

To integrate the FlashBlade array with a directory service:

- 1 Configure the FlashBlade directory service.

- 2 Test the directory service configuration.
- 3 Enable the directory service.

When configuring the directory service for array management or NFS protocol integration with LDAP, specify the URIs, base DN, and bind username and password of the directory service. If selecting the NIS directory service for NFS, clear any existing LDAP configurations before specifying the NIS servers and NIS domain.

⚠ Important! Anonymous binding is supported for NFS and Management. If configuring anonymous binding for NFS or Management, neither the bind user nor the bind password should be specified.

A "bind" account must be configured to allow the array to read the directory. The bind account should be tied to a service account (rather than an individual), and the account password should never expire or be required to change periodically. The bind account must have read privileges for users and groups and the privileges should be configured at the Organizational Unit.

For Active Directory, users with disabled accounts are not able to access the file systems.

After the directory service has been configured, test the configuration to verify that the URIs can be resolved and the FlashBlade array can bind and query the tree using the bind user credentials.

Enable the directory service after it has passed the directory service test. The directory service can be disabled at any time. Disabling the directory service stops any users in the directory from accessing the file system. In the case of management, disabling directory service prevents users from logging into the array.

For instructions on configuring each directory service, see:

- [Configuring the Array Management Directory Service](#)
- [Configuring the NFS Directory Service with LDAP](#)
- [Configuring the NFS Directory Service with NIS](#)

Kerberos Keytabs

Purity//FB supports Kerberos authentication with directory services and Active Directory (AD) accounts.

Kerberos authentication is automatically enabled when configuring FlashBlade with an Active Directory account (see [Active Directory Overview](#)).

Kerberos network authentication protocols (krb5, krb5i, and krb5p) for NFSv3, NFSv4.1 and SMB clients, in addition to auth_sys, are supported. With Kerberos, NFSv3, NFSv4.1, and SMB clients have to be authenticated by a Key Distribution Center (KDC) to get a session ticket. They have to present this session ticket to the FlashBlade array in order to get access to the file system.

Keytab files can be generated for a FlashBlade array by a KDC, and can then be imported and managed by the FlashBlade. Keytab files contain list of Kerberos principal names and encryption keys to automatically authenticate clients without requiring human interaction or access to password stored in a plain-text file. Keytab file must be manually uploaded with NFS or SMB directory services. With AD accounts, keytab files are automatically generated.

The **Kerberos Keytabs** panel displays and allows management of Kerberos keytabs on the FlashBlade for use with directory services and AD accounts.

To access the **Kerberos Keytabs** panel, from the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.

Figure 11.26 Viewing the Kerberos Keytabs Panel

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 11.27 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 11.28 Viewing the Session Log

Session Log								1-20 of 398 < >
Start Time ▼	End Time	Interface	User	Location	Event	Method	Event Count	Identifying Details
All ▼	All ▼							
2024-06-06 11:47:11	-	REST	pureuser	10.64.146.241	login	API token	1	-
2024-06-06 11:47:11	-	REST	pureuser	10.64.146.241	login	API token	1	-
2024-06-06 11:47:11	-	REST	pureuser	10.64.146.241	login	API token	1	-
2024-06-06 11:47:11	-	REST	pureuser	10.64.146.241	login	API token	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**.
 - **UNKNOWN** is used when GUI and REST access has been denied by a network policy rule.
 - **UNKNOWN** is used when CLI access has been denied by a network policy rule as a result of remote SSH fully disabled for a client source. **ir** is used if only superuser CLI access is denied (governed by the network policy rule, local-network-superuser-password-access).
 - Successful logins will always show the user regardless of what interface was used to login.
- **Interface** - The type of user interface the user logged in from. Interface types are **CLI**, **GUI**, **REST**, and **CONSOLE**.
- **Location** - The location of the user, indicated as an IP address.
- **Event** - The action of the the user's session. Types of events include:
 - **login** - Displayed when the user has logged in and has not yet logged out.
 - **logout** - Displayed when the user has logged out.
 - **user session** - Displayed when the user has logged out. Note that due to memory limitations, when a **Start Time** event has been removed from memory, a **user session** record will change to **logout**.
 - **request without session** - Displayed when a REST request fails.
 - **failed authentication** - Displayed when attempts to access FlashBlade fail.

- **Method** - The access authentication used to log in. Types of authentication are **API token**, **certificate**, **password**, **one-time password**, and **public key**.
 - **one-time password** is only applicable for the user **pureeng**.
 - For network policy rule denial events, the method can be null, but will be **password** for superuser password access.
- **Event Count** - The number of times an action was attempted during a 15 minute time period. Note that the counter increments for unsuccessful events such as **failed authentication** or **request without session** only. The event counter will refresh every 15 minutes.
- **Identifying Details** - The detailed information of the user's SSH certificate and logins with an SSH public key.

Single Sign-On

FlashBlade supports Single Sign-On (SSO) authentication on the management GUI and AWS Security Token Service (STS) on the object store.

FlashBlade supports both the OpenID Connect (OIDC) and SAML (SAML2) authentication protocols to allow identity providers (IdP) and service providers (SPs) (i.e. the FlashBlade) to implement user authentication and access control.

- SAML2 SSO can be configured for both management and object services.
- OIDC SSO can be configured for object services.

SSO Authentication for Object Service

FlashBlade supports OpenID Connect (OIDC) and SAML (SAML2) authentication protocols for SSO to the object store.

OIDC and SAML are protocols used for identity federation in object services allowing the creation of temporary credentials for federated identities, which enables secure access to object storage. This is achieved through the use of roles and trust policies, which define the permissions and trusted entities that can assume these roles.

Purity//FB allows the configuration of OIDC to work with identity providers (IdPs) such as Okta, Google, X/Twitter, etc. The process involves setting up an OIDC identity provider configuration on the FlashBlade array, which includes providing the IdP URL and optionally uploading necessary certificates if the IdP

URL points to a server running with certificates not signed by a well known certificate authority.. Once configured, roles can be created and associated with trust policies that specify the IdP users that are allowed to assume these roles and access the object storage.

Configuring SAML2 SSO for FlashBlade allows federated identities to assume roles and securely access object storage through the use of roles and trust policies, which define the permissions and trusted entities that can assume these roles. Configuration involves setting up the identity provider (IdP) on FlashBlade with the necessary certificates and URLs.

FlashBlade supports up to five SAML and five OIDC configurations for object service.

SSO Authentication for Management Service

⚠ Important! Do not enable Multi-factor Authentication with SAML2 SSO for management service before testing the configuration. If the configuration is enabled with errors, the user will be locked out. Go to Testing the SAML2 SSO Configuration section to perform the test.

Purity//FB supports single sign-on (SSO) integration with the Microsoft® Active Directory Federation Services (AD FS), Okta, Azure Active Directory (Azure AD), and Duo Security via the SAML2 protocol. When SAML2 SSO for management service is configured and enabled, all user logins to the Purity//FB GUI are redirected to the IdP login page for single sign-on.

For example, the AD FS identity provider (IdP) can optionally be configured to enable multi-factor authentication (MFA) and supports X.509 certificates, Microsoft Azure™ authentication, and other authentication methods, in addition to password authentication. This release is verified with X.509 certificate authentication as the multi-factor authentication method.

SSO Limitations

The following considerations apply to this release:

- Only one SAML2 SSO configuration for management service can be created at a time on an array.
- Up to five SAML and five OIDC configurations for object service can be created at a time on an array.
- A single SSO configuration for both object and management services is not allowed.
- SAML2 SSO for management service authentication applies only to GUI logins. SSH logins continue to use their existing password authentication or other authentication mechanism.

Viewing Identity Provider Configurations

ODIC and SAML2 SSO identity provider configurations are displayed and managed in the **ODIC Identity Providers** and **SAML Identity Providers** panels in the **Single Sign-On** screen. To access the **Single Sign-On** screen, from the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.

Figure 11.29 Viewing the Single Sign-On Screen

The screenshot displays the 'Single Sign-On' settings page in the Pure Storage Purity//FB GUI. The left sidebar contains a navigation menu with options like Dashboard, Storage, Protection, Policies, Analysis, Performance, Capacity, Replication, Network, Health, and Settings. The 'Settings' section is expanded, showing sub-panels for SYSTEM, NETWORK, and SECURITY. The 'Single Sign-On' panel is active, showing two sections: 'OIDC Identity Providers' and 'SAML Identity Providers'. The 'OIDC Identity Providers' section has a table with columns: Name, Enabled, Services, Provider URL, CA Certificate, CA Certificate Group, and PRN. One provider named 'test' is listed. The 'SAML Identity Providers' section has tabs for 'General', 'Service Provider', and 'Identity Provider'. It also has a table with columns: Name, Enabled, Services, Binding, Array URL, and PRN. Two providers are listed: 'sso-app' and 'sso-sts-object'. The top right corner shows the user's name 'pureuser' and the time 'GMT-08:00 (PST)'.

Two panels, **OIDC Identity Providers** and **SAML Identity Providers**, provide information about OIDC and SAML identity provider configurations. General (default), Service Provider and Identity Provider-specific information can be displayed in the **SAML Identity Providers** panel.

Clicking the name of an OIDC or SAML identity provider configuration opens a details screen for that configuration.

The **OIDC Identity Provider** details screen provides a summary of the OIDC identity provider configuration information.

Figure 11.30 Viewing OIDC Identity Provider Configuration Details

Details

Enabled	False
Name	test
Services	object
Binding	none
Provider URL	https://idp.providers.com
CA Certificate	-
CA Certificate Group	_default_replication_certs
PRN	prn::iam:array-id/local::oidc-provider/test

The **SAML Identity Provider** details screen provides two panels, **Configuration** and **Role Mappings**.

Figure 11.31 Viewing SAML Identity Provider Configuration Details

SAML Identity Provider
Test

Configuration

GENERAL

Enabled	True
Name	sso-app
Services	management
Binding	http-redirect
Array URL	https://mokujiin.dev.purestorage.com
PRN	prn::iam:array-id/local::saml-provider/sso-app

SERVICE PROVIDER (SP)

SP Entity ID	https://mokujiin.dev.purestorage.com/saml2/service-provider-metadata/sso-app
SP Assertion Consumer URL	https://mokujiin.dev.purestorage.com/login/saml2/sso/sso-app
SP Metadata URL	https://mokujiin.dev.purestorage.com/saml2/service-provider-metadata/sso-app
Sign Request	False
Signing Credential	-
Decryption Credential	-

IDENTITY PROVIDER (IdP)

IdP Entity ID	-
IdP URL	-
IdP Metadata URL	https://dev-16326233.okta.com/app/exkmwchty1cLEjT95d7/sso/saml/metadata
IdP Metadata URL CA Certificate	-
IdP Metadata URL CA Certificate Group	-
Encrypt Assertion	False
Verification Certificate	-

Role Mappings

1-4 of 4

Name	Role	Group	Group Base
array_admin	array_admin	purearrayadmins	OU=puremgmt
ops_admin	ops_admin	pureopsadmins	OU=puremgmt
readonly	readonly	purepeasants	OU=puremgmt
storage_admin	storage_admin	purestorageadmins	OU=puremgmt

Configure group-to-role mappings if your identity provider is set up to send only user group in the SAML response. Use the [Directory Service](#) panel to set up these mappings. Refer to the SAML2 SSO Configuration section of the FlashBlade User Guide for more details.

The **Configuration** panel displays the following:

- **General**

- **Enabled** : If SSO is enabled (**True**) or not (**False**).
- **Name**: The SSO configuration name.
- **Services**: The service type, either **management** or **object**.
- **Binding**: The binding method. For management service, the binding method is **http-direct**. For object service, the binding method is **none**.
- **Array URL**: The FlashBlade array's URL.

- **Service Provider (SP)**

- **SP Entity ID**: The URL of the SP (not used for object service).
- **SP Assertion Consumer URL**: The URL where the identity provider will send its SAML response after authenticating a user (not used for object service).
- **SP Metadata URL**: The URL to the SP metadata file (not used for object service).
- **Sign Request**: If the sign request is enabled (**True**) or not (**False**). When enabled, the matching signing credential must be configured at the identity provider (not used for object service).
- **Signing Credential**: The credential used by the service provider to sign SAML requests (not used for object service).
- **Decryption Credential**: The credential used by the service provider to decrypt encrypted SAML assertions from the identity provider.

- **Identity Provider (IdP)**

- **IdP Entity ID**: The globally unique name for the IdP.
- **IdP URL**: The URL of the IdP.
- **IdP Metadata URL**: The URL to the IdP metadata file.
- **IdP Metadata URL CA Certificate**: The URL to the CA certificate used to validate the authenticity of the metadata file.
- **IdP Metadata URL CA Certificate Group**: The URL to the group of CA certificates used to validate the authenticity of the metadata file.

- **Encrypt Assertion:** If the encryption assertion is enabled (**True**) or not (**False**). When enabled, certificates matching the decryption credential must be configured at the identity provider.
- **Verification Certificate:** The certificate used by the SP to verify the signed SAML response from the IdP.

The **Role Mappings** panel displays current role mapping as configured from the **Directory Service** panel for management service SSO. Note that the **Role Mappings** panel is read only. Any changes to role mapping should be performed from the **Directory Service** panel.

- **Name:** The directory group name.
- **Role:** The role assigned to the group. Values can be array_admin, storage_admin, ops_admin, or read-only (see [Roles](#) for additional information).
- **Group:** The common name of the directory group.
- **Group Base:** The configured group in the directory tree comprised of the OU or CN combined with the base DN attribute.

 **Note:** FlashBlade uses IdP cookies for SSO, and clearing browser cookies/redoing the IdP login may be necessary when IdP authorization settings change.

Configuring SSO Authentication for Object Service

OIDC and SAML2 identity Providers can be configured on the FlashBlade for SSO for object service.

Configuration for Object Service Prerequisites

Before attempting OIDC or SAML SSO configuration for object service, you must request a metadata document from the identity provider (IdP) that contains the following information.

For OIDC configuration, you must have the following:

- CA and CA Group Certificates (not needed if the IdP server certificates are signed by a well known CA).
- The IdP URL, which specifies the URL of the identity provider.

For SAML configuration, you must have the following:

- The identity provider (IdP) entity ID, which specifies the globally unique name for the identity provider.

- The IdP URL, which specifies the URL of the identity provider.
- The IdP verification certificate.

OR

- The IdP metadata URL
- (Optional) CA certificate or CA Certificate Group.

The IdP metadata file contains this information.

 **Note:** The IdP metadata URL is sufficient without the IdP verification certificate.

Configuration for Object Service Overview

The following list summarizes the steps to configure and enable ODIC or SAML2 SSO for object service authentication on a FlashBlade. Configurations are required in both the service provider side (Purity//FB) and on the identity provider side, in order to complete the SSO configuration.

- 1 Create an OIDC or SAML2 identity provider on the FlashBlade.
- 2 Create a role and attach permissions to the role using access policies on the FlashBlade.
- 3 Configure the trust policy for the role to authenticate the IdP.

Configuring an OIDC Identity Provider for Object Service on the FlashBlade

To configure an OIDC identity provider for object service on the FlashBlade array, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
- 2 Click the Add (+) button in the header of the **OIDC Identity Providers** panel. The **Create ODIC Identity Provider** pop-up window appears.
- 3 In the **Name** field, enter the OIDC identity provider name.
- 4 Click **Enabled** to enable the configuration after creation.
- 5 The **Service** and **Binding** fields are automatically populated with **object** and **none**, respectively, and cannot be changed.
- 6 In the **Provider URL** field, enter the identity provider URL.

- 7 (Optional) In the **CA Certificate** field, select the CA certificate that will be used to validate the identity provider's authenticity.
- 8 (Optional) In the **CA Certificate Group** field, select the group of CA certificates that will be used to validate the identity provider's authenticity.
- 9 Click **Create**.

Configuring a SAML2 Identity Provider for Object Service on the FlashBlade

To configure SAML2 identity provider for object service on the FlashBlade array, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
- 2 Click the Add (+) button in the header of the **SAML Identity Providers** panel. The **Create SAML Identity Provider** pop-up window appears.
- 3 In the **Name** field, enter the SAML identity provider name.
- 4 Click **Enabled** to enable the configuration after creation.
- 5 From the **Service** list, select **object**. The **Binding** field is automatically populated with **none**.
- 6 In the **Array URL** field, enter the FlashBlade array's URL.

The **SP Entity ID**, **SP Assertion Consumer URL**, and **SP Metadata URL** fields are populated based on the information entered in the **Name** and **Array URL** fields.

- 7 (Optional) In the **Decryption Credential** field, select the credential from the drop-down list. Note that only array-type certificates are displayed.

 **Note:** Enter either the IdP Metadata URL by itself, or enter the IdP Metadata URL, the optional IdP Metadata CA certificate, and the optional IdP Metadata CA certificate group.

- 8 In the **IdP Entity ID** field, enter the globally unique name for the IdP. Required if the IdP metadata URL is not provided.
- 9 In the **IdP URL** field, enter the IdP's URL. Required if the IdP metadata URL is not provided.
- 10 In the **IdP Metadata URL** field, enter the URL to the IdP metadata file. .
- 11 (Optional) In the **IdP Metadata CA Certificate** field, select the CA certificate to validate the authenticity of the IdP metadata file.
- 12 (Optional) In the **IdP Metadata CA Certificate Group** field, select the group of CA of certificates to validate the authenticity of the IdP metadata file.

13 (Optional) Enable **Encryption Assertion**.

14 In the **IdP Verification Certificate** field, select the certificate used by the SP to verify the signed SAML response from the IdP from the drop-down list. This is required if the IdP metadata URL is not provided.

15 Click **Create**.

Configuring an IdP for SSO using OIDC or SAML

Configuration of an IdP of SSO differs depending on the IdP. The following are general overviews of what is required for configuring an IdP for SSO using OIDC and SAML.

For OIDC, determine the Identity Provider's public URL from documentation or IdP administrative console. When appended with **/.well-known/openid-configuration**, you should see the OpenID Connect Discovery specification-compliant metadata.

For SAML:

- You must access your IdP's administrative console and navigate to the SAML configuration section to create a new SAML application within your IdP.
- Input details such as the SP's entity ID, ACS URL.
- Define which user attributes from your IdP should be sent to the SP.
- Metadata files must be downloaded and then uploaded on the SP for integration. Alternatively, the metadata URL can be passed to the SP.

Creating a Role

Once an OIDC or SAML2 identity provider has been created, you must create a role with access policies. A maximum of 10,000 roles are allowed per FlashBlade.

When creating a role, you can define the allowed session duration and add one or more access policies, which will grant the S3 permissions in the account.

To create a role, perform the following:

- 1** From the **Storage > Object Store** page, click the account name in the **Accounts** panel under which you wish to create a role. The details page for the account appears.
- 2** Click the Add (+) button in the header of the **Roles** panel. The **Step 1: Create Role** pop-up window appears.

- 3 In the **Role Name** field, enter a name for the role.
- 4 In the **Maximum Session Duration** field, enter the maximum session duration from 1 hour to 48 hours. If no value is entered, the default is 1 hour.
- 5 Click **Create**. The **Step 2: Add Access Policies to role <name>** appears.
- 6 From the **Available Policies** panel, select the access policies you wish to assign to the role. At least one policy must be added to grant the role permissions in the account.

If a specific access policy does not exist in the list of available policies, click **Create a new Access Policy** (see [Creating Object Store Access Policies](#) for instructions).

- 7 Click **Add**.

Creating Trust Policy Rules

After creating a role, create a trust policy rule.

To create a trust policy rule, perform the following:

- 1 From the **Storage > Object Store** page, click the account name in the **Accounts** panel under which you created a role. The details page for the account appears.
- 2 Click the created role from the **Roles** panel. The details page for the role appears.
- 3 Click the Add (+) button in the header of the **Trust Policy Rules** panel. The **Create Trust Policy Rule** pop-up window appears.
- 4 In the **Rule Name** field, enter a name for the rule.
- 5 Select whether the action (**Effect**) of the rule will be to **Allow** or **Deny**. Note that explicit Deny rules take precedence over Allow rules.
- 6 Click the Add (+) button in the **Principals** field. The **Add Principals** pop-up window appears.
 - a Select the principals you wish to add from the **Available Principals** panel.
 - b Click **Add**.
- 7 Click the Add (+) button in the **Actions** field. The **Add Actions** pop-up window appears.
 - a Select the actions you wish to add from the **Available Actions** panel.
 - b Click **Add**.
- 8 (Optional) If you wish to add a condition to the rule, click **Add condition**. The condition key must have a format prefix:value. Supported key prefixes are: aws, jwt, saml, sts.
 - a In the Key field, enter the key prefix.

- b** Select **All** or **Any**, the operator, and select **If Exists** under the **Operator** field.
 - c** Click the Add (+) button under the **Values** field. Enter a comma separated value string and then click **OK**.
- 9** Click **Create**.

Uploading a Trust Policy in IAM Format

To upload a trust policy in IAM format, perform the following:

- 1** From the **Storage > Object Store** page, click the account name in the **Accounts** panel with the role that you wish to upload a trust policy in IAM format. The details page for the account appears.
- 2** Click the role from the **Roles** panel. The details page for the role appears.
- 3** Click **More Options > Upload Trust Policy in IAM format** in the header of the **Trust Policy Rules** panel. The **Upload Trust Policy** pop-up window appears.
- 4** In the **Trust Policy** field, copy and paste the trust policy in IAM format. You can optionally load the policy from a file.
- 5** If you did not manually past the trust policy in IAM format in the previous field, you can load the policy from a file. In the **Load from file** field, click **Choose File** and select the file containing the trust policy in IAM format to upload.
- 6** Click **Upload**.

Configuring SSO Authentication with SAML2 SSO for Management Service

 **Important!** Do not enable SSO Authentication with SAML2 SSO for management service before testing the configuration. If the configuration is enabled with errors, the user will be locked out. Go to [Testing the SAML2 SSO Configuration](#) section to perform the test.

SAML2 identity Providers can be configured on the FlashBlade for SSO for management service.

SAML2 SSO Configuration for Management Service Prerequisites

Before attempting SAML SSO configuration for management service, note the following prerequisites:

- The SAML2 SSO configuration for management service steps require either administrator access to the identity provider's management tool, or, for example, coordination with an AD FS administrator to create a relying party trust for the array and to provide the following information:

- The identity provider (IdP) entity ID, which specifies the globally unique name for the identity provider.
- The IdP URL, which specifies the URL of the identity provider.
- The IdP verification certificate.

The IdP metadata file contains this information.

 **Note:** The IdP metadata URL is sufficient without the IdP verification certificate.

- The directory service used with the array must be the same directory service instance as used by the identity provider in the relying party trust configuration for this array. Purity//FB does not support multiple or federated directory services.
- It is common for identity providers to use TLS 1.2 and above for securing communications and enabling services such as SP monitoring in AD FS.

 **Important!** It is highly recommended that before configuring SSO, create a strong password for the **pureuser** and other array administrator accounts, and to save those passwords according to your organization's security policies.

Additionally, note that configurations are required on both the service provider side (Purity//FB) and on the identity provider (IdP) side to complete the SSO configuration.

SAML2 SSO Configuration for Management Service Overview

The following list summarizes the steps to configure and enable SAML2 SSO for management service authentication on a FlashBlade. Configurations are required in both the service provider side (Purity//FB) and on the identity provider side, in order to complete the SSO configuration.

- 1 Configure SAML2 SSO on the FlashBlade.
 - a Obtain IdP information from AD FS, Okta, or an administrator.
 - b Configure the service provider (SP) with the array and IdP information.
- 2 In the identity provider, set up SSO using the SP information from the FlashBlade.
- 3 On the FlashBlade, enable SAML2 SSO Authentication.

For information about Directory Services and Directory Services configuration instructions, see [Directory Services Overview](#) and its related topics.

Use the AD FS Management Tool on AD FS to configure the AD FS IdP for SSO and, optionally, multi-factor authentication (MFA) with the FlashBlade.

Configuration Notes

- The verification certificate (the AD FS primary token-signing certificate) must be an X.509 certificate in PEM format.
- The array name portion of the URL in the browser used to configure the SP must be consistent with the URL entered into the **Array URL** field in the SAML2 SSO panel. Whether the URLs are based on a FQDN (such as pure01.- mycompany.com) or on a hostname (such as pure01), the browser URL and the **Array URL** field configured in the SAML2 SSO panel must be consistent in the way the array name is specified. When there is a mismatch, the browser used to configure SAML SSO will not be able to complete the redirect and authentication process.

Group to Role Mapping

Group to Role Mapping on Purity//FB

Group to role mapping on Purity//FB is only certified with the AD FS server. To configure the group to role mapping, see [Adding Role Mapping](#) .

To configure the identity provider server to send the user group in the SAML response, refer to your identity provider's documentation.

Group to Role Mapping on IdPs

 **Note:** Directory service configuration is not necessary for role mapping on IdPs.

Group to role mapping on IdPs is certified with AD FS and Okta. Each IdP has unique configuration steps to map the user group(s) to the Purity//FB roles (array_admin, storage_admin, ops_admin, and readonly). The specified role is then sent in the SAML response with attribute name purity_roles. Contact Pure Storage Technical Services to create an SSO integration application on the IdP side and to configure group to role mapping for the specified IdP.

Configuring a SAML2 Identity Provider for Management Service on the FlashBlade Array

 **Note:** It is not recommended to enter the service provider (SP) signing and decryption credentials or enable the identity provider (IdP) request or assertion until after the initial configuration passes the end-to-end test.

To configure SAML2 identity provider for management service on the FlashBlade array, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
- 2 Click the Add (+) button in the header of the **SAML Identity Providers** panel. The **Create SAML Identity Provider** pop-up window appears.
- 3 In the **Name** field, enter the SAML identity provider name.
- 4 From the **Service** list, select **management**. The **Binding** field is automatically populated with **http-direct**.
- 5 In the **Array URL** field, enter the FlashBlade array's URL. The URL must use HTTPS.

The **SP Entity ID**, **SP Assertion Consumer URL**, and **SP Metadata URL** fields are populated based on the information entered in the **Name** and **Array URL** fields. The information in these fields is required to create a relaying party for the FlashBlade array in the AD FS identity provider.

- 6 (Optional) Enable **Sign Request**.
- 7 (Optional) In the **Signing Credential** field, select the credential from the drop-down list. Note that only array-type certificates are displayed.
- 8 (Optional) In the **Decryption Credential** field, select the credential from the drop-down list. Note that only array-type certificates are displayed.

 **Note:** Enter either the IdP Metadata URL by itself, or enter the IdP Entity ID and IdP URL fields and the IdP Verification Certificate. If both the IdP Metadata URL and any of the IdP Entity ID, IdP URL, or IdP Verification Certificate are provided, the metadata file is read only for those options that are not provided.

- 9 In the **IdP Entity ID** field, enter the globally unique name for the IdP. For AD FS, this name is found under the AD FS Management Tool under AD FS/Service/Federated Service Properties.
- 10 In the **IdP URL** field, enter the IdP's URL. This value is found under the AD FS Management Tool under AD FS/Service/Endpoints.
- 11 In the **IdP Metadata URL** field, enter the URL to the IdP metadata file. This value is found under the AD FS Management Tool under AD FS/Service/Endpoints.

12 (Optional) Enable **Encryption Assertion**.

13 In the **IdP Verification Certificate** field, select the certificate used by the SP to verify the signed SAML response from the IdP from the drop-down list. For example, this value is found under the AD FS Management Tool under AD FS/Service/Certificates/Token-signing.

14 Click **Create**.

Configuring the IdP

 **Note:** FlashBlade uses IdP cookies for SSO, and clearing browser cookies/redoin the IdP login may be necessary when IdP authorization settings change.

Once SAML2 SSO configuration is complete on the FlashBlade, you must register Purity//FB as a relying party trust on your IdP. Administrator-level access is required to access your IdP's management tool.

Configuration steps will be different depending on the IdP. As an example, the following steps describe how to register Purity//FB as a relying party trust on AD FS.

To register Purity//FB as a relying party trust on AD FS, perform the following:

- 1** On the machine running AD FS, click **Control Panel > System and Security > Administrative Tools** and select **AD FS Management**.
- 2** In the left panel, select **Relying Party Trusts**. In the right panel, select **Add Relying Part Trust**. The **Add Relying Part Trust Wizard** appears.

 **Note:** The Access Control Policy **Permit everyone** corresponds to password authentication. Multi-factor authentication can be configured later if required. Consider waiting until after the SSO configuration passes with password authentication before enabling multi-factor authentication.

Table 11.12 Sample Configuration in the Add Relying Party Trust Wizard

Wizard Page	Action
Welcome	Ensure Claims aware is selected.
Select Data Source	Select Enter data about the relying party manually .
Specify Display Name	Enter a display name for the new relying party (for convenience, this name could match the SP configuration display name, as set in the FlashBlade).
Configure Certificate	No action.

Table 11.12 Sample Configuration in the Add Relying Party Trust Wizard (continued)

Wizard Page	Action
Configure URL	Select Enable support for the SAML 2.0 Web SSO protocol . In the Relying party SAML 2.0 SSO service URL field, enter the Assertion Consumer URL from the Purity SAML2 SSO pane.
Configure Identifiers	In the Relying party trust identifier field, enter the SP entity ID from the Purity SAML2 SSO pane.
Choose Access Control Policy	Select Permit everyone .
Ready to Add Trust	No action.
Finish	Select Configure claims insurance policy for this application .

- 3 Once the new relying party appears in the Relying Party Trusts table, select the relying party and in the right panel select **Edit Claims Insurance Policy**. The **Edit Claims Insurance Policy** page opens.
- 4 Click **Add Rule** at the bottom right of the page. These rules specify the content the IdP returns in an assertion to the SP. Two rules are required: one for a unique identity for the user, the second for group information. The **Add Transform Claim Rule Wizard** opens.
- 5 In the **Claim rule template** field on the **Select Rule Template** page, select **Send LDAP Attributes as Claims** and then click **Next**.
- 6 Click **Choose Rule Type** and follow the instructions as described in the following table:

Table 11.13 Actions in the Add Transform Claim Rule Wizard for Users

Field	Action
Claim rule name	Enter a descriptive name for the rule, such as map name id .
Attribute store	Select Active Directory .
Mapping table LDAP Attribute column	Select SAM-Account-Name .
Outgoing Claim Type column	Select Name ID and then click Finish .

The new rule for name mapping appears in the **Edit Claims Insurance Policy** page **Insurance transform Rules** table.

- 7 Click **Add Rule** for the group information rule. The **Add Transform Claim Rule Wizard** appears.
- 8 In the **Select Rule Template** page, select **Send LDAP Attributes as Claims** from the **Claim rule template** field.

Table 11.14 Actions in the Add Transform Claim Rule Wizard for Groups

Field	Action
Claim rule name	Enter a descriptive name for the rule, such as pass group info .
Attribute store	Select Active Directory .
Mapping table LDAP Attribute column	Select Is-Member-Of-DL .
Outgoing Claim Type column	Select Group and then click Finish .

The new rule for group mapping appears in the **Edit Claims Insurance Policy** page **Insurance Transform Rules** table.

Enabling Multi-factor Authentication

 **Important!** Whenever you modify the SAML2 SSO for management service configuration, it is recommended to re-test to verify the configuration information before enabling.

 **Note:** The types of multi-factor authentication (MFA) available depend on the IdP. The AD FS IdP supports certificate authentication, authentication with Microsoft® Azure™ software, and others.

You can optionally enable MFA as described in the following:

- 1 In the AD FS **Management tool Relying Party Trust** page, select the relying party you created. In the right panel, select **Edit Access Control Policy**.
- 2 On the **Choose an access control policy** page, select **Permit everyone and require MFA** and then click **Next**.
- 3 The next steps, such as configuring a certificate, depends on the type of MFA selected.

Testing the SAML2 SSO for Management Service Configuration

 **Important!** Whenever you modify the SAML2 SSO for management service configuration, you must re-test to verify the configuration information before enabling.

 **Note:** Refer to the [SAML2 SSO Login Error Messages](#) table to assist with troubleshooting end-to-end test error messages. The table also includes actionable tasks to correct the error messages.

Once you have your SAML2 SSO for management service configuration set up, perform the basic test and end-to-end (E2E) test to verify all the data is valid. This test can only be administered by the array_admin role. If SAML2 SSO for management service is not configured then the test option will not be enabled. Use the basic test and E2E test to validate the configuration to be applied before enabling

SSO. If you encounter issues, use the test result page with the SAML response payload from the IdP to correct any configuration errors.

To test the SAML2 SSO for management service configuration, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
- 2 Click the name of the SAML identity provider configured for SSO for management service you wish to test from the **SAML Identity Providers** panel. The details page for the configuration appears.
- 3 Click **Test** in the header of the **SAML Identity Provider** panel. The **Test SAML2 SSO Configuration** pop-up appears.

Test SAML2 SSO Configuration

Basic test results:

Testing array URL correctness: https://flashblade

Testing IdP metadata URL connectivity: https://www.google.com

E2E test result:

End-to-end test has not been started yet.

Close

E2E Test

The IdP metadata URL in the example is only a demonstration that a non-viable site may pass the basic test because it only checks for connectivity and that the site is trusted and supports downloading of the page's content. Use the end-to-end test results to validate the IdP metadata URL.


If the basic test fails, review the error failure messages, close the window and edit your configuration then perform another test.

If you selected to enable your SAML2 SSO for management configuration and click **Save**, without testing, a **Test Reminder** pop-up will appear. It is not recommended to click **Proceed to Save** without testing. You should always test new configurations and modified configurations.

- 4 Click **E2E Test** when your basic test results have passed.
- 5 The ADFS (Active Directory Federation Services) window will appear. Enter your credentials and log in to ADFS and the E2E test will be initiated, as indicated by the **Test SAML2 SSO Configuration** pop-up with the **Check E2E Test Result** button.
- 6 Click the **Check E2E Test Result** button to review the test results when pop-up indicates the test is successful.

The E2E test results may not be ready immediately. If you click the test results before the E2E test completes, you will be notified that the test has not completed.

The SAML2 SSO for management configuration may be incorrect if the test does not complete or is stopped before completion. If the test takes more time than expected, you can close the test and restart it by reopening the **Test SAML SSO Configuration** dialogue box and clicking the **E2E Test** button.



- 7 Review the E2E test results. If your configuration failed, review the error messages and use the [SAML2 SSO Login Error Messages](#) table to troubleshoot your configuration.



End To End Test Result	
Status	Result Details
● Passed	Success

SAML Response ^

```

<samlp:Response ID="_6447c2db-c0d3-491d-be54-a86951a1477c" Version="2.0" IssueInstant="2024-05-02T23:55:53.247Z" Destination="https://fa-saml2-test-array/login/saml2/sso/fa-saml2-test-idp-w-roles"
Consent="urn:oasis:names:tc:SAML:2.0:consent:unspecified" InResponseTo="ARQ61673a3-327c-4f57-af18-a4b88c851edd" xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"><Issuer
xmlns="urn:oasis:names:tc:SAML:2.0:assertion">http://fa-saml2-
test.dev.purestorage.com/adfs/services/trust</Issuer><samlp:Status><samlp:StatusCode
Value="urn:oasis:names:tc:SAML:2.0:status:Success" /></samlp:Status><Assertion ID="_96d4638e-2192-4218-b663-d732a81189de" IssueInstant="2024-05-02T23:55:53.247Z" Version="2.0"
xmlns="urn:oasis:names:tc:SAML:2.0:assertion"><Issuer>http://fa-saml2-
test.dev.purestorage.com/adfs/services/trust</Issuer><ds:Signature
xmlns:ds="http://www.w3.org/2000/09/xmldsig#"><ds:SignedInfo><ds:CanonicalizationMethod
Algorithm="http://www.w3.org/2001/10/xml-exc-c14n#" /><ds:SignatureMethod
Algorithm="http://www.w3.org/2001/04/xmldsig-more#rsa-sha256" /><ds:Reference URI="#_96d4638e-
2192-4218-b663-d732a81189de"><ds:Transforms><ds:Transform
Algorithm="http://www.w3.org/2000/09/xmldsig#enveloped-signature" /><ds:Transform
Algorithm="http://www.w3.org/2001/10/xml-exc-c14n#" /></ds:Transforms><ds:DigestMethod
Algorithm="http://www.w3.org/2001/04/xmldsig#sha256" />
<ds:DigestValue>eFND34vX9jP32BBYMcV5JmCZCMBfOUKGNi0pyNQMbW0=</ds:DigestValue>
</ds:Reference></ds:SignedInfo>
<ds:SignatureValue>wMKIGjn2vOWD2q0UBmvFmWzatW7ltCURwH39WJspNJ/F/tiPBMG9yHUj2u/CeTuK1at
</ds:SignatureValue><KeyInfo xmlns="http://www.w3.org/2000/09/xmldsig#"><ds:X509Data>
<ds:X509Certificate>MIIc/jCCAeagAwIBAgIQJAANoyQudodEpv+dKoxwaDANBgkqhkiG9w0BAQsFADA7M
</ds:X509Certificate></ds:X509Data></KeyInfo></ds:Signature><Subject><NameID>bcaci/NameID>
<SubjectConfirmation Method="urn:oasis:names:tc:SAML:2.0:cm:bearer"><SubjectConfirmationData
InResponseTo="ARQ61673a3-327c-4f57-af18-a4b88c851edd" NotOnOrAfter="2024-05-
03T00:00:53.247Z" Recipient="https://fa-saml2-test-array/login/saml2/sso/fa-saml2-test-idp-w-roles" />
</SubjectConfirmation></Subject><Conditions NotBefore="2024-05-02T23:55:53.231Z"
NotOnOrAfter="2024-05-03T00:55:53.231Z"><AudienceRestriction><Audience>https://fa-saml2-test-
array/saml2/service-provider-metadata/fa-saml2-test-idp-w-roles</Audience></AudienceRestriction>
</Conditions><AttributeStatement><Attribute Name="purity_roles">
<AttributeValue>array_admin</AttributeValue></Attribute><Attribute
Name="http://schemas.xmlsoap.org/claims/Group">
<AttributeValue>CN=pureadmins,OU=PureStorage,OU=SAN,OU=IT,OU=US,DC=dev,DC=purestorage,DC=cc
</AttributeValue></AttributeStatement><AuthnStatement AuthnInstant="2024-05-02T23:55:53.137Z"
SessionIndex="_96d4638e-2192-4218-b663-d732a81189de"><AuthnContext>
<AuthnContextClassRef>urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport</AuthnConte
</AuthnContext></AuthnStatement></Assertion></samlp:Response>
          
```

Close and go back to configurations

- 8 Click the **Close and go back to configurations** button and when ready, enable SAML2 SSO for management service.

Enabling SAML2 SSO for Management Service Authentication

⚠ Important! Whenever you modify the SAML2 SSO for management service configuration, it is recommended to re-test to verify the configuration information before enabling.

After the SAML2 SSO for management service configuration is enabled, SAML users are referred to the identity provider for authentication for all new login attempts to the Purity//FB management GUI. Existing user sessions are not affected.

To enable SAML2 SSO for management service authentication, perform the following:

- 1 (Optional) Log into Purity//FB in a second browser window if login access is interrupted.
- 2 From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
- 3 Click the Edit button in the same row as the SAML identity provider for management service in the **SAML Identity Providers** panel you wish to enable. The **Edit SAML Identity Provider** pop-up window appears.
- 4 Click **Enabled**.
- 5 Click **Save**.

The **SAML Identity Providers** panel displays **True** under the **Enabled** column for the SAML identity provider for management service.

Once enabled, the Purity//FB login screen no longer prompts for a password. Upon a user's first log in, and also after an SSO session expires, the AD FS log in page will open and prompt for the user's AD FS credentials.

Other SAML2 SSO for Management Service Configuration Steps

Enabling IdP Sign Request and Encryption Assertion

⚠ Important! Once the sign request and encryption assertion are enabled, the IdP also needs to have matching certificates imported, if the SP metadata url is not monitored by IdP.

If not enabled during initial SAML2 SSO configuration on the FlashBlade, you can enable IdP Sign Request and Encryption Assertion.

- 1 From the **Settings** page navigation sidebar, click **Single Sign-On** under **Security**.
- 2 Click the **More Options** button in the header of the **Single Sign-On** panel and select **Edit**. The **Edit SAMLs SSO Configuration** pop-up window appears.

- 3 Enable **IdP Sign Request** and select the signing credential in the **SP Signing Credential** field.
- 4 Enable **IdP Encryption Assertion** and select the decryption credential in the **SP Decryption Credential** field.

SSO Session Timeout

When SAML2 SSO is enabled, the Purity//FB management GUI session timeouts are based on the idle timeout configured locally on the FlashBlade. After the GUI session has logged out (due to idle timeout), users can attempt to login to the GUI via SSO. When the IdP session is active, users are prompted to re-authenticate. For example, by default an AD FS SSO session times out after eight hours.

Refer to the IdP documentation for information on customizing the time out setting.

TLS 1.2 or 1.3 Support

The FlashBlade requires TLS 1.2 or 1.3 with strong authentication. The AD FS instance must also support TLS 1.2 or 1.3 with strong authentication in order to test the relying party's federation metadata URL in an AD FS instance. If TLS 1.2 or 1.3 with strong authentication is not configured, an error is seen when clicking the Test URL button when configuring monitoring in the Replying Party Configuration GUI in the AD FS instance.

Refer to Microsoft documentation for instructions on how to enable TLS 1.2/1.3 and strong authentication in AD FS.

SAML2 SSO for Management Service Runtime Notes

- If a user is deactivated in AD FS but is currently logged in, the current login session is not affected. The user is denied access at the next login attempt.
- In case the SAML2 SSO service is temporarily unavailable, users with array-admin role permissions can access the array through the Local Access link on the login page. This link provides emergency administrator access when GUI logins are unavailable.
- SSO session timeouts depend on the IdP. For example, by default AD FS times out after eight hours and Okta times out after 2 hours. A different time out length can be configured in the IdP. For additional information, refer to *SSO Session Timeout* in [Other SAML2 SSO for Management Service Configuration Steps](#).

- When a user logs in through SAML2 SSO, the browser URL field reports its URL based on what is configured in the **Array URL** field, regardless of whether the user entered the array hostname or FQDN when directing the browser to the array.

Understanding SAML2 SSO Login Error Messages

The following table describes error messages that can be returned if login errors occur.

 **Note:** We recommended checking the raw SAML response from the IdP to further understand the error details for errors showing up during the "verification of IdP response" stage. The SAML response XML text is available from the E2E test results page.

Table 11.15 SAML SSO Login Error Messages

Error Message	Cause	Stage	Actions
SSO is not available.	SSO is not enabled or failed to load to the GUI server. The error is displayed when a reserved path is accessed before authentication.	SSO setup.	Enable SSO in the Settings page.
SAML login is not supported for user xyz.	The remote user is trying to login through SAML with same name as a reserved internal user name space, such as ir, root.	Verification of IdP response.	The user is not allowed for SAML.
No user role or group information provided.	No role nor groups information are included in the identity provider's assertion response.	Verification of IdP response.	Check the group information from the identity provider, also use the management DS test function to ensure the group has a mapped role in the FlashBlade.
No valid user role information found.	The role information is provided, but is not a valid role defined in the FlashBlade.	Verification of IdP response.	Check the "purity_roles" definition from the identity provider side.
SSO login fails to get enabled because the identity provider's metadata is not accessible.	The IdP's metadata URL is provided, but is not accessible.	SSO setup.	Check and update the metadata URL of identity provider on the FlashBlade.
Necessary service provider(SP) cookie is not included in the request. Please ensure that cookies are enabled and verify that the accessed URL corresponds to the expected SP's consumer URL.	The SP site cookie is not sent back along with the IdP response, so SP cannot verify which request the response is associated with.	Verification of IdP response.	Check the cookie settings on the browser, also make sure the host name in the login page matches the array URL setting for SSO.
Incomplete SAML response: No subject found in the assertion. Please check the claim setups from the identity provider.	The configuration of the replying party on IdP does not add data such as user name and group/role information to the assertion.	Verification of IdP response.	Ensure the claim setting on the IdP side has added the user name inside the assertion claim.

Table 11.15 SAML SSO Login Error Messages (continued)


Error Message	Cause	Stage	Actions
Failed to decrypt the data received. Check the encryption key in the SAML2 configuration.	The encryption key may not match between SP and IdP.	Verification of IdP response.	Check the encryption credential's expiration and ensure they match with the IdP record.
Assertion signature verification failed. Check the verification certificate in the SAML2 configuration.	The signature verification failed, possibly due to an incorrect verification cert for the SSO configuration.	Verification of IdP response.	Check the verification credential's expiration and ensure they match with the IdP record.
Failed to verify the assertion. Check the verification certificate in the SAML2 configuration.	The general error when assertion is invalid.	Verification of IdP response.	Check the verification certificate.
Failed to authenticate via SSO	Generic error when no explicit error is found.	Verification of IdP response.	n/a
Issuer from the SAML response doesn't match the record. Please check the entity ID details of the SSO config.	Occurs when the issuer of the IdP response does not match the record from the SP side.	Verification of IdP response.	Check and update the SP configuration.
Invalid status [urn:oasis:names:tc:SAML:2.0:status:Responder] for SAML response	Generic error from the IdP. Indicates something is wrong with the IdP configuration.	Verification of IdP response.	Check if the certificate for verifying SAML requests is properly configured when the FlashBlade includes signatures in the SAML request. Check if the rules in the response claim is properly configured.

Users

The **Users** panel manages local and remote FlashBlade array users.

 **Note:** Remote users are listed in the **Users** panel only after API tokens have been explicitly configured for them.

To access the **Users** panel, from the **Settings** page navigation sidebar, click **Users** under **Security**.

Figure 11.32 Users - Users Panel


Users							1-1 of 1
Name ▲	Role	Type	Public Key	API Token	Expires	Lockout Remaining	
pureuser	array_admin	local	-	****	-	-	

The panel displays the following:

- **Name:** The user name.
- **Role:** The role assigned to the user. Values can be `array_admin`, `storage_admin`, `ops_admin`, or `read-only`.
- **Type:** Displays `local` for users local to the array (`pureuser`), or `remote` for LDAP users.
- **Public Key:** The user's public key, if any, displayed as `****`.
- **API Token:** Active API token, if any, displayed as `****`.
- **Expires:** Expiration time, if any, of the API token.

In addition to displaying user attribute details, you can also create and manage local users, user API tokens, and user array authentication in the **Users** panel.

Local Users

Users can be added by configuring local users directly on the array. Alternatively, users can be added via Lightweight Directory Access Protocol (LDAP) by integrating the array with a directory service such as Active Directory or OpenLDAP (refer to the Directory Services Overview for additional information about integrating the array with a directory service).

On the array (locally), users can only be created by array administrators. Local users must have unique names and cannot be the same names as LDAP users. If an LDAP user appears with the same name as a local user, the local users will have priority. Up to 100 local users can be created.

Role-based access control (RBAC) restricts system access and capabilities to each user based on their assigned role in the array. All users are assigned a role on the array, as described in the following table.

Role	Description
readonly	Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array.
storage_admin	Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
array_admin	Array Admin users can perform all FlashBlade operations.

Role	Description
ops_admin	Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.

For local users, the role is set during user creation. For LDAP users, the role is set by configuring groups in the directory that correspond to the FlashBlade user roles. Each local user account on the array is password protected. The password is assigned during user creation and can be modified by array administrators. All local users can manage their own passwords, but only array administrators can manage the passwords of other users. Changing a local user's password requires knowledge of the current password. If the password of a local user is unknown, delete the account and recreate it with the desired password. Note that deleting a local user's account means deleting any public key associated with the user. If the password of the pureuser account is unknown, contact Pure Storage Technical Services to reset the account to the default pureuser password. Passwords of LDAP users are managed in the directory service.

Public keys and API tokens that have been created for users appear masked in the **Users** panel.

API Tokens

The REST API requires permission to run commands. To access the server to run REST API commands, an authentication API token must be generated for the user. Each API token is unique and has an optional expiration setting. When a token expires, it can no longer be used and a new token must be generated to continue running REST API commands. It is a best practice to set an expiration on a token for the duration of time it will be used.

Array Authentication

 **Note:** Setting other users' public keys is available to administrators and users with admin permissions only. The public key should be in the OpenSSH format, which begins with `ssh-rsa`. A maximum of 100 users may have a public key set.

Administrators and users with administrator privileges can set passwords for array users. Additionally, public key authentication can be configured allowing users to SSH to the array.

Creating a Local User

To create a local user, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 Click the **More Options** button in the header of the **Users** panel and select **Create User**. The **Create Local User** pop-up window appears.
- 3 In the **User** field, enter the user name.
- 4 In the **Role** field, select the user's role.
 - **Read-Only:** Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array.
 - **Ops Admin:** Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.
 - **Storage Admin:** Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
 - **Array Admin:** Array Admin users can perform all FlashBlade operations.
- 5 In the **Password** field, enter the user's password.
- 6 In the **Confirm Password** field, re-enter the password.
- 7 Click **Create**.

Editing a Local User

To edit a local user, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 From the **Users** panel, click the **More Options** button in the the same row as the user you wish to edit and select **Edit User**. The **Edit User** pop-up window appears.
- 3 (Optional) In the **Role** field, select the user's role.
 - **Read-Only:** Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array.

- **Ops Admin:** Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.
- **Storage Admin:** Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
- **Array Admin:** Array Admin users can perform all FlashBlade operations.

4 (Optional) In the **Current Password** field, enter the user's password.

5 (Optional) In the **New Password** field, enter the user's new password.

6 (Optional) In the **Confirm Password** field, reenter the new password.

7 Click **Save**.

Deleting a Local User

Local users can be deleted at any time. Once deleted, the user will no longer be able to access the FlashBlade array. Deleted users cannot be restored.

To delete a local user, perform the following:

- 1** From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2** From the **Users** panel, click the **More Options** button in the the same row as the user you wish to delete and select **Delete User**. The **Delete Local User** pop-up window appears.
- 3** Click **Delete**.

Creating an API Token

API tokens allow users to run REST API commands. To access the server to run REST API commands, you must generate an authentication API token for the user. As a best practice set an expiration on a token for the duration of time it will be used.

To create an API token for a user, perform the following:

- 1** From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2** Click the **More Options** button in the header of the **Users** panel and select **Create API Token**. The **Create API Token** pop-up window appears.

- 3 Enter the user for whom the token will be created.
- 4 Enter the token's expiration duration in the format N[m/h/d/w]. For example, **15m**, **1h**, **2d**, **3w**. If not set, the API token will not expire.
- 5 Click **Create**.

The API Token pop-up window appears displaying the token and its creation and expiration dates.

Updating a User's Public Key

A public key can be added or updated for user authentication in order to allow users to SSH to the FlashBlade.

 **Note:** Setting other users' public keys is available to administrators and users with admin permissions only. The public key should be in the OpenSSH format, which begins with `ssh-rsa`. A maximum of 100 users may have a public key set.

To add or update a user's public key, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 Click the **More Options** button in the header of the **API Clients** panel and select **Update Public Key**. The **Update Public Key** pop-up window appears.
- 3 Enter the user for whom you are adding a public key.
- 4 Enter the public key into the **Public Key** field.
- 5 Click **Update**.

Displaying an API Token

If you previously created an API token you can later display and copy the token for use with the REST API.

To display a user's API token for a user, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 In the **Users** panel, locate the user whose API token you wish to display. Click the **More Options** button in the same row, and select **Show API Token**. The **API Token** pop-up window appears.
- 3 Click **Copy** to copy the token if needed.
- 4 Click **Close**.

Recreating an API Token

API tokens can be recreated. Note that recreating an API token invalidates the old token and the recreated token must be authenticated.

To recreate an API token, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 In the **Users** panel, locate the user whose API token you wish to recreate. Click the **More Options** button in the same row, and select **Recreate API Token**. The **Recreate API Token** pop-up window appears.
- 3 Click **Recreate**.

Removing an API Token

To remove an API token from a user, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Users** under **Security**.
- 2 In the **Users** panel, locate the user whose API token you wish to remove. Click the **More Options** button in the same row, and select **Remove API Token**. The **Remove API Token** pop-up window appears.
- 3 Click **Remove**.

Chapter 12

The Purity//FB Fusion

Purity//FB Fusion (Fleet) provide preliminary array federation and remote storage management features, including remote file system and object creation and management, user and group management, and policy management of arrays in a fleet. Fleets are configured and managed from the **Fleet** page of the GUI.

Fusion Management

Fusion allows arrays to be federated into fleets. You can create, manage, and search for objects remotely on any member array from the GUI or CLI of any other array in the fleet.

In the initial view of the **Fleet** page you can create a fleet or join an existing fleet.

On an array that is a member of a fleet, the **Fleet** page displays the member arrays for the current fleet.

Use the **Fleet** page for the following actions:

- Generate a fleet key to allow another array to join the fleet.
- Remove an array from the fleet.
- Search for member arrays.

The following are the possible values for member array **Status**:

- **Joined**: Array has successfully joined the fleet.
- **Joining**: Array is currently joining the fleet.
- **Removing**: Array is currently being removed from the fleet.

 **Note**: Fleet management is enabled by default for Purity//FB 4.5.5 or later release.

Requirements and Restrictions

The following are requirements for arrays in a fleet:

- All arrays in a fleet must be able to connect to each other via management network ports (ex: vif0).
All arrays must be able to communicate with each other on the network over port 443.
- Directory services (Active Directory or LDAP) must be enabled with a consistent configuration on each array in the fleet. Remote provisioning operations are restricted to users authenticated by directory services.
- Arrays in a fleet can run on different versions of Purity. FlashBlades must be on Purity//FB version 4.5.5 or later and FlashArrays must be on Purity//FA version 6.8.1 or later.

Secure Communication in Arrays

In a fleet of arrays, secure communication is established using mutual Transport Layer Security (mTLS). This requires each array to share a common certificate, ensuring that all communications within the fleet are authenticated and encrypted. To maintain security, certificate rotation typically occurs every 45 days, with a Time-To-Live of 90 days. If an array loses connection to the fleet for an extended period and misses the certificate rotation, it will no longer be able to communicate with the fleet until it is reconnected and updated with the current certificate. This approach ensures that only authorized arrays can participate in the fleet, maintaining the integrity and security of the system.

In the case that the array becomes disconnected from the fleet for longer than the standard 45 days and the certificate expires, the array will not be able to connect to the fleet any longer. In order to renew the certificate, run the `purefleet member remove --unreachable <ARRAY>` command on both the array and any other array in the fleet to grant the array a new certificate.

Fleet Keys

A fleet key is an encrypted key containing the information necessary for another array to join an existing fleet.

Because a fleet key can only be created on an array that already is a member of the fleet, the key serves as confirmation that another array is authorized to join the fleet. A valid fleet key is required for an array to join a fleet.

- Each fleet key expires after an hour or on its first use, whichever occurs first.
- Only one fleet key is valid at any point in time. If a second key is created, the first key is deleted, if still present.

Fleet Creation and Membership

Each fleet member array must meet the requirements listed in the *Requirements and Restrictions* section. An array can belong to only one fleet at a time.

 **Note:** The fleet is created with one default member array (the *local array*). Subsequent arrays added to the fleet are *remote arrays*.

To create a fleet:

- 1 On an array that is to be a member of the fleet, click the **Fleet** page in the left navigation panel.
- 2 In the **Create a new Fleet** dialog, enter the name for the new fleet in the **Fleet Name** field.
- 3 Click **Create a Fleet**.

Generating a Fleet Key

The fleet key is necessary for an array to join an existing fleet. As the key is for a single use and expires after an hour, create the key when you are ready to add another array to the fleet and notify that array's Purity administrator of the fleet key.

 **Note:** Save the fleet key as there is no way to retrieve the key after you close the pop-up window.

To create a fleet key:

- 1 From the **Fleet** page, click the **More Options** button in the **Fleet Members** panel and click **Create Fleet Key**. The **Create Fleet Key** pop-up window appears.
- 2 On the **Create Fleet Key** pop-up, click **Copy and Close**. The fleet key is copied to the clipboard.

Joining an Existing Fleet

Follow these steps on an array that you wish to add to a fleet.

To join an existing fleet:

- 1 On the array that is to join the fleet, click **Fleet** in the left navigation panel.
- 2 In the **Fleet Welcome** page, select **Join an Existing Fleet**.
- 3 In the **Join an Existing Fleet** pop-up window, enter the fleet name and the previously generated fleet key.
- 4 Click **Join a Fleet**.

If the fleet join operation times out, the array will be left in fleet-join state. When the array is in fleet-joining state, any attempt to create or join a fleet will fail.

To abort the fleet join operation, go to the CLI and execute the
`purefleet member remove --unreachable <LOCAL-ARRAY>` command.

Removing a Remote Array from the Fleet

In this release, the first array in the fleet cannot be removed from the fleet.

To remove an array from a fleet:

- 1 From the **Fleet > Fleet Members** panel, identify the member array to be removed from the fleet.
- 2 Click the **More Options** button on that member row and click **Remove**.
- 3 In the **Remove Member** pop-up window, confirm that the selected array is the array to be removed.
- 4 Click **Remove**.

Search for Member Arrays

In the **Fleet** page of any member array, begin typing the array name in the **Member Names** field. The **Fleet Members** panel displays the arrays that begin with the same characters.

An asterisk is supported as a wildcard. For example, `pure*db` lists all member arrays that begin with `pure` and also contain `db`.

Summary of Steps to Create a Fleet with Multiple Array Members

The following list summarizes how to create a fleet with multiple member arrays. Detailed instructions are described in the previous sections.

- 1 On the first array:
 - a In the GUI navigation panel, click **Fleet**.
 - b From the fleet **Welcome** page, select **Create a new Fleet**.
 - c From the **Fleet** page, in the fleet name row, click the **More Options** button and select **Generate Fleet Key**.
- 2 On the next array you wish to add to the fleet, from that array's fleet **Welcome** page, select **Join an existing Fleet** and specify the fleet name and its fleet key.
- 3 On any array in the fleet, generate a new fleet key and then repeat step 2 for additional arrays you wish to add to the fleet.

Presets and Workloads

A preset is a reusable template that describes how to provision and configure the Purity objects that make up a workload. When multiple presets are created, it becomes a catalog of customizable workload templates that simplify storage provisioning and enforce configuration standards.

Workloads are provisioned from presets. A workload is a container object that holds references to other Purity objects enabling them to be managed and monitored as a set. Purity objects can be in at-most, one workload.

Creating a Preset

To create a preset, perform the following:

- 1 Navigate to the **Storage > Presets** page and click **Add (+)** in the panel's header. The **Create New Preset** wizard appears.
- 2 Enter the preset details as required in the **Preset Details** section of the wizard, including the workload type and storage class. Workloads have a specific workload pathology that provides a baseline optimization. The storage class defines the type of storage allowed for placing this

workload. Fusion will automatically detect the storage classes in the fleet. When complete, click **Done**.

- 3 Enter the storage resources details such as names and sizes as required in the **Storage Resources** section of the wizard and click **Confirm Storage Resource**.
- 4 (Optional) If you wish to add additional storage resources, click **Add Storage Resources** and repeat step 3.
- 5 Review the preset's configuration and click **Publish**.

Configuring a Workload

Once a preset has been created, you can configure a workload.

To configure a workload, perform the following:

- 1 Navigate to the **Storage > Workloads** page.
- 2 Click the create (+) button in the **Workloads** panel. The workload configuration wizard appears.
- 3 From the **Choose a Preset** screen, select the preset you wish to apply to the workload and click **Next**. The **Configure Workload** screen appears.
- 4 Enter the workload name and any required configuration information required by the preset and click **Next**. The **Workload Placement** screen appears.
- 5 Select the array on which the workload will be deployed and the project period. You can optionally click **Get Recommendations** and Pure1 will provide the recommended arrays to best handle the workload. Click **Next**. The **Review & Deploy** screen appears.
- 6 Review the workload parameters and click **Deploy Workload**.

Fusion creates the storage configuration and the new workload is displayed in the **Workloads** panel.

Portable Presets

 **Note:** Portable presets are available in Purity//FA 6.8.8 and higher and Purity//FB 4.6.7 and higher.

Presets can be copied to other fleets, backed up, shared, and automated.

Presets are editable JSON (text) documents that can be saved off the array and uploaded to any other array. This is performed from the GUI, CLI, or REST API.

Downloading a Workload Preset

To download a work preset as JSON, perform the following:

- 1 Navigate to the **Storage > Presets** page.
- 2 From the **Presets** panel, locate the preset you wish to download and click **More Options > Download** in the same row as the preset. The **Download Preset <Preset Name> as JSON** pop-up window appears.
- 3 Update the file name if necessary.
- 4 Click **Save**.

Uploading a Workload Preset

To upload a work preset from JSON, perform the following:

- 1 Navigate to the **Storage > Presets** page.
- 2 Click **More Options > Upload** from the header of the **Presets** panel. The **Upload Preset (JSON) in <Fleet Name>** pop-up window appears.
- 3 Click **Browse** to navigate to the JSON file you wish to upload from your local machine, or drag and drop the file into the pop-up window.
- 4 Enter a new preset name.
- 5 Click **Save**.

Downloading and Uploading Presets Using the CLI

Presets can be downloaded and uploaded using the CLI.

- 1 Run the **purepreset workload download** command to download a preset.

```
> purepreset workload download <preset_name>
```

This command outputs the JSON to your screen where you can then copy and paste the JSON into the text editor of your choice. If needed, edit the JSON output and copy the contents of the file back to the command line when prompted

- 2 Run the **purepreset workload upload** command to upload a preset.

```
> purepreset workload upload -context <fleet_name> <preset_name>
```


Remote Resource Management

For arrays that are members of the same fleet, any array can view, create, connect to, and resize storage resources on other members of the fleet. Additionally, audit logs show entries for all member arrays of the fleet.

Fleet remote management functionality is available when Fleet View is selected.

On FlashBlades with fleet management enabled, the **Storage > Servers**, **Storage > File Systems**, and **Storage > Object Store** pages have a toggle between local resources and resources on arrays that are members of the fleet. For fleet management features, select **Fleet View**.

Fleet Data Protection

Functionality to manage data protection for all member arrays of a fleet is available for **Protection > Snapshots**, **Protection > File Replication**, and **Protection > Object Replication**.

Fleet View for Snapshots

On the **Protection > Snapshots** page, fleet remote management for snapshots is available when **Fleet View** is selected. In **Fleet View**, all of the snapshots on each member of the fleet are displayed.

To manage fleet snapshots, refer to [Snapshots](#) and [Creating a File System Snapshot](#).

On any member of the fleet, in the **Fleet File System Snapshots** panel in **Fleet View**, the **Context** column displays the location of each snapshot in the fleet.

Fleet View for File Replication

On the **Protection > File Replication** page, fleet remote management for file replication is available when **Fleet View** is selected. In **Fleet View**, all replication links in the fleet are displayed.

To manage fleet file replication, refer to [Replication and Recovery](#).

On any member of the fleet, the **File Replication** page panels in **Fleet View**, the **Context** column displays the source array in the replication pair in the fleet.

Fleet View for Object Replication

On the **Protection > Object Replication** page, fleet remote management for file replication is available when **Fleet View** is selected. In **Fleet View**, all replication links and remote credentials in the fleet are displayed.

To manage fleet object replication, refer to [Replication and Recovery](#).

On any member of the fleet, the **Object Replication** page panels in **Fleet View**, the **Context** column displays the source array in the replication pair in the fleet.

Searching Across a Fleet

In **Fleet View** for the **Storage** and **Protection** options that have fleet functionalities, certain columns (such as Name and Size) have search fields.

In **Fleet View**, these search fields search across the entire fleet.

Viewing Audit Entries for a Remote Array

For members of a fleet, the audit trail shows entries for all arrays in the fleet.

In the **Settings > Security > Audit Trail** page, the **Origin** column shows the array on which each entry occurred.

Chapter 13

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

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

To create an Active Directory account, perform the following:

- 1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.
- 2 Click the Add (+) button. The **Create Active Directory Account** pop-up window appears.

- 3 Select **Create a new account**.
- 4 Enter the name of the new account in the **Name** field.
- 5 In the **Server** list, select the server.
- 6 In the **Computer Name** field, enter the name for the AD computer account that will be created in the AD domain. If a computer name is not entered, the computer name will default to the value entered in the **Name** field. The computer account name should resolve to the data VIP.
- 7 In the **Domain** field, enter the AD domain to which this computer account will be created.
- 8 In the **User Name** field, enter the name of the user who is capable of creating the computer account within the domain.
- 9 Enter the user's password in the **Password** field.

If you wish to enter additional configuration information, click **Additional Settings**.
- 10 In the **Join OU** 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**.
- 11 Select the **Encryption Type**. The default encryption type is **aes256-sha1**. If NTLMv2 is required, **arcfour-hmac** encryption must be enabled.
- 12 Click the Add (+) button to the right of the **Service Principal Names** field and enter a comma-separated list of service principal names (SPNs) or fully qualified domain names (FQDNs) for registering services with the AD domain. If nothing is entered, the value defaults to **Computer Name Domain** as an FQDN.
- 13 Click the Add (+) button to the right of the **Kerberos Servers** field and enter a comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- 14 Click the Add (+) button to the right of the **Directory Servers** field and enter a comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- 15 Click **Create**.

Instructions and troubleshooting on accessing Native SMB file systems using NTLMv2 are available in [this Knowledge Base article](#).

Creating an Active Directory Account using the CLI

Active Directory (AD) accounts can also be created using the **puread account create** CLI command. For full details about the **puread** command, see the *FlashBlade CLI Reference*.

To create an active directory account using the CLI, perform the following:

- 1 Run the **puread account create account --domain DOMAIN ACCOUNT** command, where **ACCOUNT** is the name of the new AD account. Note that the **--domain** argument, which specifies the AD domain to which the computer account will be created, is required. You can also specify the following:
 - **--computer-name:** (Optional) The name for the AD computer account that will be created in the AD domain. If a computer name is not entered, the computer name will default to the name used for the **ACCOUNT-NAME**. The computer account name should resolve to the data VIP.
 - **--join-ou:** (Optional) The organization unit (location) where the computer account will be created. For example, **OU=Arrays**, **OU=Storage**, etc. If an OU is not entered, the value defaults to **CN=Computers**.
 - **--encryption-type:** (Optional) A comma-separated list of encryption types that will be supported for use by clients for Kerberos authentication. Defaults to **AES256-cts-hmac-sha1-96**.
 - **--fqdn:** (Optional) A comma-separated list of fully qualified domain names (FQDNs) used for the creation of keys for Kerberos authentication. If FQDNs are specified, all supported service principals will be registered for them. If neither FQDNs nor SPNs are specified, all supported service principals 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 13.1 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 10.64.242.197 _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 -
```

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 the Add (+) button to the right of the **Kerberos Servers** field and enter a comma-separated list of key distribution servers for Kerberos. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- 8 Click the Add (+) button to the right of the **Directory Servers** field and enter a comma-separated list of directory servers used for authorization. Servers can be listed as IP addresses or DNS names. If nothing is entered, servers are resolved for the domain in DNS.
- 9 Click **Save**.

Instructions and troubleshooting on accessing Native SMB file systems using NTLMv2 are available in [this Knowledge Base article](#).

Editing an Active Directory Account using the CLI

The **puread account setattr** CLI command can also be used to edit the configuration of active directory accounts. For full details about the **puread** command, see the *FlashBlade CLI Reference*.

When running the command, the **ACCOUNT-NAME** must be specified. Similar to creating an AD account, the following can be set/changed:

- **--join-ou:** (Optional) The organization unit (location) where the computer account will be created. For example, **OU=Arrays**, **OU=Storage**, etc. If an OU is not entered, the value defaults to **CN=Computers**.
- **--encryption-type:** (Optional) A comma-separated list of encryption types that will be supported for use by clients for Kerberos authentication. Defaults to **AES256-sha1**.

- `--fqdn:` (Optional) A comma-separated list of fully qualified domain names (FQDNs) used for the creation of keys for Kerberos authentication. If FQDNs are specified, all supported service principals will be registered for them. If neither FQDNs nor SPNs are specified, all supported service principals 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
ZuieCj1U8MmZWOLy016EuQaWLB+GQNWC85MAAABMAAIAClRPVE9STy5ORVQAA25mcwAKT0JBTi1T
TUItMgAAAAFfkquaBQASACCYmferQZVm6J4KOVtwyZlagvI7XoS5BpYsH4ZA1YLzkwAAAFgAAGAK
VE9UT1JPLk5FVAAESE9TVAABv2Jhb1lzbWItMi50b3RvcM8ubmV0AAAAAV+Sq5oFABIAIJIz96tB
lWbongo5VPDjMvQc8jteHLkGliwfhkDVgvOTAAAVwACAAPUT1RPuk8uTkVUAANuZnMAFW9iYW4t
c2liLTludG90b3JvLm5ldAAAAFfkquaBQASACCYmferQZVm6J4KOVtwyZlagvI7XoS5BpYsH4ZA
1YLzkwAAAE0AAgAKVE9UT1JPLk5FVAAESE9TVAABk2Jhb1lzbWItMwAAAAFfkquaBQASACCYmfer
QZVm6J4KOVtwyZlagvI7XoS5BpYsH4ZA1YLzkwAAAEwAAgAKVE9UT1JPLk5FVAADbmZzAApYmFu
LXNtYi0zAAAAAV+Sq5oFABIAIJIz96tB1Wbongo5VPDjMvQc8jteHLkGliwfhkDVgvOTAAAWAAC
AApUT1RPuk8uTkVUAARIT1NUABVvYmFuLXNtYi0zLnRvdG9yby5uZXQAAAABX5KrmgUAEgAgmJn3
...
```

Copy and save the output for use elsewhere.

Deleting a Keytab File

 **Note:** Deleting keytab files can be a disruptive action if there were clients with active sessions that relied on the deleted keytab.

To delete a keytab file, perform the following:

- 1 From the **Settings** page navigation sidebar, click **Kerberos Keytabs** under **Security**.
- 2 To delete one or more keytab files:

- a Click the **More Options** button in the header of the **Kerberos Keytabs** panel and select **Delete**. The **Delete Keytabs** pop-up window appears.
- b Select one or more keytab files from the **Existing Keytabs** column on the left. Each selected keytab file appears under the **Selected Keytabs** column on the right.
- c Click **Delete**.

3 To delete a single keytab:

- a From the **Kerberos Keytabs** panel, click the **More Options** button in the same row as the keytab file you wish to delete and select **Delete**. The **Delete Keytab** pop-up window appears.
- b Click **Delete**

Deleting Keytab Files using the CLI

Use the **purekeytab delete KEYTAB** CLI command to delete one or more keytabs from the FlashBlade. For full details about the **purekeytab** command, see the *FlashBlade CLI Reference*

KEYTAB is the name of the keytab file you wish to delete. Multiple keytab files can be specified in a comma-separated list.

For example:

```
purekeytab delete keytab1,keytab2,keytab3
```

In this example, keytab files `keytab1`, `keytab2`, and `keytab3` are deleted from the AD account.

Deleting an Active Directory Account

Deleting an Active Directory (AD) account deletes the computer account associated with the FlashBlade and all Kerberos keytabs associated with the account.

To delete an Active Directory (AD) account, perform the following:

- 1** From the **Settings** sidebar, click **Directory Service** under **Security**. The **Active Directory Accounts** panel appears.
- 2** Click the **Delete** icon in the row of the account. The **Delete Active Directory Account** pop-up window appears.
- 3** (Optional) If you wish to retain the computer account associated with the FlashBlade in your AD domain, enable (blue) **Local Only**. Note that if enabled, you will need to manually delete Kerberos keytabs.

4 Click **Delete**.

Deleting an Active Directory Account using the CLI

Deleting an Active Directory (AD) account can also be performed using the CLI.

Use the **puread account delete** command to delete an AD account using the CLI. For full details about the **puread** command, see the *FlashBlade CLI Reference*.

To delete an AD account using the CLI, issue the **puread account delete ACCOUNT-NAME** command, where **ACCOUNT-NAME** is the AD account name you wish to delete. For example:

```
puread account delete mynewacct
```

You can optionally specify the `--local-only` argument, which will only delete the AD account on the local array. The computer account will not be deleted in the AD domain.

Directory Services Overview

A directory service allows the configuration of a server that will be used to perform some form of lookup. For example, a management directory server is used to look up users and validate their passwords, as well as to search for groups to which the users belong.

By configuring a directory service, FlashBlade is able to look up relevant information to authenticate users along different data and management paths.

A directory service is defined by the following configuration:

- A list of one or more URIs that specify the directory servers that will be used for lookups by the directory service. These may be specified in the form of IP addresses.
- A Base Distinguished Name (Base DN) as the point from which a server will begin its searches.
- A bind user and bind password (user and password) of a single service account that is used by FlashBlade to search the directory server.

Generally, FlashBlade communicates with directory servers over the Lightweight Directory Access (LDAP) protocol.

FlashBlade has three directory services:

- Array Management (see [Array Management Directory Service](#) for more details)
- NFS (see [NFS Directory Service](#) for more details)
- SMB (see [SMB Directory Service](#) for more details)

If you want to enable Kerberos authentication for FlashBlade's clients when configuring NFS or SMB directory services, you must upload keytab files with keytabs for the appropriate encryption types and service principal names that clients will use. Doing so grants the same client functionality as joining an AD domain.

NFS and SMB Nested User Groups

 **Note:** Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

A nested group is a group which is a member of another group. When using groups to manage permissions, you can create nested groups to allow inheritance of membership permissions from one group to its sub-groups. Similarly, users can belong to multiple group hierarchies and inherit those groups' membership permissions.

Purity//FB supports NFS and native SMB users belonging to a nested hierarchy of group membership (NFS and SMB nested groups). This allows access to users who are not the owners of particular files, but who instead belong to a group that owns the files. A user can belong to a group directly or indirectly (when groups are nested, i.e. when a group belongs to another group). Without NFS or SMB nested groups, Pure's NFS server allows only access to users that are direct members of the group that owns the file. With NFS and SMB nested groups, a user's membership in nested groups, as well as in direct groups, is taken into account for access checks. NFS and SMB interoperability with nested groups is supported; FlashBlade looks up both nested SIDs and GIDs.

The following limits apply to nested groups:

- Maximum number of groups, including direct and indirect groups, per user = 10,000
- Maximum number of nested groups to which a user can be a direct member = 500
- Maximum number of nested groups to which a group can be a direct member = 500
- Maximum number of nested levels = 5

Pure's implementation of NFS nested groups supports the following LDAP-based directory servers:

- OpenLDAP
- Oracle Unified Directory (OUD)
- Active Directory (AD)

In order to use NFS and SMB nested groups with directory services, nesting must be configured according to the RFC2307bis specification on your LDAP server.

Access from multiple domains with two-way trusted domains are supported within the same AD forest. Access across different AD forests is not supported.

Array Management Directory Service

The array management directory service configuration is used to allow users within an LDAP server to log in to a Flashblade (via the CLI, GUI, or REST) and manage it.

In order for a user to log in to the FlashBlade, they must be within a group that is configured to have a role on the array. If a group does not have a role, then its members cannot log in.

When configuring the array management directory service, a bind user and bind password can be optionally specified. If not specified, FlashBlade attempts to bind anonymously (without authenticating) with the configured LDAP servers. Note that if the configured LDAP servers do not grant read permissions on the needed pathways to unauthenticated machines, the anonymous bind will fail.

Roles

Directory services for array management can configure role-based access control (RBAC) for LDAP users by configuring groups in the directory that corresponds to the following permission groups (roles) on the array.

Role	Description
readonly	Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array
storage_admin	Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering file systems and snapshots, as well as object store accounts, users, and access keys. Storage Admin users cannot perform operations that deal with global and system configurations.
array_admin	Array Admin users can perform all FlashBlade operations.

Role	Description
ops_admin	Ops Admin users have all the privileges of Read Only users, plus the ability to run commands related to support operations, such as managing remote assistance sessions, phonehome, and proxy settings. Ops Admin users cannot perform operations that deal with storage or non-support global and system configurations.

When an array management user connects to the FlashBlade, the array confirms the user's identity from the directory service. The response from the directory service includes the user's group, which Purity//FB maps to a role on the array, granting access accordingly.

If users are assigned to more than one role, an alert is issued. For security purposes, a user belonging to more than one FlashBlade role will inherit the permissions of the most restrictive role. For example, a user who has been assigned both the Storage Admin and Read Only roles will inherit the privileges of the more restrictive Read Only role. Note that the user will continue to have privileges from the original role until the directory service role cache is cleared (cache is automatically cleared every 8 hours). To make new permissions take immediate effect, run the **pureadmin refresh** command.

One exception to this rule is users assigned the Ops Admin role—users assigned the Ops Admin role and any other role are denied the ability to log into the FlashBlade. If there is an error with the directory service configuration on the FlashBlade array, contact a user with Array Admin privileges to resolve the issue. If the LDAP server is configured incorrectly, contact your LDAP server administrator. The alert issued in the case where a user is denied access due to being assigned the Ops Admin and another role, provides additional information to help resolve the issue.

As a best practice, after changing a user's role or management directory service configuration settings, use the **pureadmin refresh** command to refresh the user's role or clear directory service role cache to ensure all user privileges are up to date.

For directory service-enabled accounts, user passwords to the FlashBlade are managed through the directory service while API tokens are configured through Purity//FB. The password can be a maximum of 100 characters in length and include any character that is entered from a US keyboard.

In the FlashBlade GUI, buttons and menu items for unauthorized actions are disabled per the user's configured role.

Multiple Directory Service Group Mapping to Role

For an LDAP user to log into the array, the user must belong to a configured group in the LDAP directory, and that group must be mapped to a Purity//FB RBAC role (array_admin, ops_admin, storage_admin, or

readonly) on the array. Once an LDAP group is mapped to a role, all members of the group are assigned that role.

Multiple groups can be mapped to a single role, thereby allowing users from different groups to share the same role and log into the array. This can be useful for large enterprises with complex directory service configurations with many different groups spread across their organization. This feature is enabled by default and allows array_admin users to configure and manage multiple groups to a single role.

- Each Purity//FB role can be mapped to up to 100 LDAP groups.
- Of those groups, different group bases are allowed.
- The permission level of an individual user is cached locally to prevent frequently binding and querying the directory. The cache entries expire after a time limit at which point the directory is queried again and the cache entry is updated. Cache entries are also automatically updated when a new login that requires communication with the configured directory server(s) occurs.

Configuring the Array Management Directory Service

 **Note:** Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The array management directory service enables FlashBlade management access for LDAP users.

To configure the array management directory service using the GUI, perform the following:

- 1 Verify that the directory server is accessible through a management interface.
- 2 Verify that the DNS settings can be used to resolve the directory server domain and server.
- 3 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
- 4 In the **Directory Service** panel, select **Array Management**.
- 5 Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.
- 6 In the **URIs** field, type the comma-separated list of up to 30 URIs of the directory servers.

Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory

service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

If specifying a domain name, it should be resolvable by the configured DNS servers.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from 0 to 255 representing a group of 8 bits.

For IPv6, specify the IP address in the form `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`). Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

Optionally specify a port. Append the port number after the end of the entire address. Default ports are 389 for LDAP, and 636 for LDAPS. Non-standard ports can be specified in the URI if they are in use.

- 7 In the **Base DN** field, type the base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.

 **Important!** Anonymous binding is supported for the array management directory service. If configuring anonymous binding, neither the bind user nor the bind password should be specified.

- 8 In the **Bind User** field, type the username used to bind to and query the directory. For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`. For Active Directory servers, you may instead choose to enter the username - often referred to as the `sAMAccountName` or user login name - of the account that is used to perform directory lookups. The username cannot contain the characters " `[] : ; | = + * ? < > / \`, and cannot exceed 104 characters in length.
- 9 In the **Bind Password** field, type the password for the bind user account.
- 10 (Optional) Select a CA certificate or certificate group to verify the authenticity of configured LDAP servers to establish a TLS (Transport Layer Security) communication for directory services.

- 11 The **User Login Attribute** defaults to the sAMAccountName for Active Directory, or UID for other server types.
- 12 The **User Object Class** defaults to User for AD, posixAccount or shadowAccount for OpenLDAP depending on the server's schema, posixAccount for posix-compliant servers, or person for all other server types.
- 13 Click **Test** to verify the configuration information. The **Test <Directory Service> Configuration** pop-up window appears, displaying the output of the test. During the directory service test, Purity//FB tests the directory service configuration to verify that the URI can be resolved and that the directory service can successfully bind and query the tree using the bind user credentials.

If the test fails, update the configuration information and retest.

Alternatively, you may save the configuration and enable the directory service at a later time.

- 14 If the test passes, select the **Enabled** checkbox to enable the directory service.
- 15 Click **Save**.

Adding Role Mapping

 **Note:** Role mapping can only be performed by array_admin users.

Once the array management directory service is configured, you can add (or map) a group or groups to the FlashBlade roles used with array management directory service. See [Roles](#) for details about FlashBlade roles.

To map a group to a role, perform the following:

- 1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
- 2 In the **Role Mappings** panel, click the **Add (+)** button. The **Add Role Mapping** pop-up window appears.
- 3 Enter a name. The name can be the same as the role.
- 4 Select a role from the **Role** list.
- 5 In the **Group** field, enter the common name (CN) of the directory group. The name should be the Common Name of the group without the "CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, "pureusers,OU=PureStorage", where pureusers is the common name of the directory service group.
- 6 In the **Group Base** field, enter the path to the configured groups in the directory tree. The Group Base consists of one or more path components (OUs and CNs) that, when combined with the base DN attribute and the configured group CNs, complete the full distinguished name of each

group. The group base should specify “OU=” for each OU and “CN=” for each CN. Multiple path components should be separated by commas. The order of the path components should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: “OU=PureGroups,OU=SANManagers”.

7 Click **Add**.

If you wish to map multiple groups to the single role, repeat this procedure for each group you wish to map.

Editing Role Mapping

 **Note:** Role mapping editing can only be performed by array_admin users.

The group and group base of existing role mappings can be edited.

To edit an existing role mapping, perform the following:

- 1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
- 2 In the **Role Mappings** panel, click the **Edit** button at the end of the same row of the role mapping you wish to edit. The **Edit Role Mapping** pop-up window appears.
- 3 In the **Group** field, enter the common name (CN) of the directory group. The name should be the Common Name of the group without the "CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, "pureusers,OU=PureStorage", where pureusers is the common name of the directory service group.
- 4 In the **Group Base** field, enter the path to the configured groups in the directory tree. The Group Base consists of one or more path components (OUs and CNs) that, when combined with the base DN attribute and the configured group CNs, complete the full distinguished name of each group. The group base should specify “OU=” for each OU and “CN=” for each CN. Multiple path components should be separated by commas. The order of the path components should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: “OU=PureGroups,OU=SANManagers”.
- 5 Click **Save**.

Deleting Role Mapping

 **Note:** Role mapping deletion can only be performed by array_admin users.

To delete an existing role mapping, perform the following:

- 1 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.

- 2 In the **Role Mappings** panel, click the **Delete Role Mapping** button at the end of the same row of the role mapping you wish to delete. The **Delete Role Mapping** pop-up window appears.
- 3 Click **Delete**.

Configuring the Array Management Directory Service using the CLI

 **Note:** Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The array management directory service can also be configured using the **pureads setattr management** command. For full details about the **pureads** command, see the *FlashBlade CLI Reference*.

To configure the array management directory service using the CLI, perform the following:

- 1 Issue the **pureads setattr management** command with the following arguments:

- **--uri:** Up to 30 URIs can be specified in a comma-separated list.

Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

If specifying a domain name, it should be resolvable by the configured DNS servers.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from 0 to 255 representing a group of 8 bits.

For IPv6, specify the IP address in the form

`[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`). Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

Optionally specify a port. Append the port number after the end of the entire address. Default ports are 389 for LDAP, and 636 for LDAPS. Non-standard ports can be specified in the URI if they are in use.

- `--base-dn`: The base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.

⚠ Important! Anonymous binding is supported for the array management directory service. If configuring anonymous binding, neither the bind user nor the bind password should be specified.

- `--bind-user`: The username used to bind to and query the directory. For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`. For Active Directory servers, you may instead choose to enter the username - often referred to as the `sAMAccountName` or user login name - of the account that is used to perform directory lookups. The username cannot contain the characters " [] : ; | = + * ? < > / \, and cannot exceed 104 characters in length.
- `--bind-password`: The password for the bind user account.
- `--ca-certificate` or `--ca-certificate-group`: (Optional) A CA certificate or certificate group to verify the authenticity of configured LDAP servers to establish a TLS (Transport Layer Security) communication for directory services.
- `--user-login-attribute`: Defaults to the `sAMAccountName` for Active Directory, or `UID` for other server types.
- `--user-object-class`: Defaults to `User` for AD, `posixAccount` or `shadowAccount` for OpenLDAP depending on the server's schema, `posixAccount` for posix-compliant servers, or `person` for all other server types.

For example:

```
pureds setattr --uri ldaps://[2000:0db1:54a1::ae25:7b1f:0505:1112] --base dn DC=mycompany,DC=com management
```

In this example:

- the URI is set to the IPv6 address `2000:0db1:54a1::ae25:7b1f:0505:1112`
 - the scheme is set to `ldaps://` to enable SSL
 - the base DN is set to be the correct domain components (DC)
 - anonymous binding is used
- 2 Test the configuration using the **pureds test management** command. If the test fails, update the configuration information.

- 3 If the test passes, enable the directory service using the **pureds enable management** command.
- 4 Configure the role(s) you wish to assign the service by issuing the **pureds role setattr --group --group-base ROLE** command, where:
 - **--group**: The common name (CN) of the directory group. The name should be the Common Name of the group without the "CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, "pureusers,OU=PureStorage", where pureusers is the common name of the directory service group.
 - **--group-base**: The organizational unit (OU) to the configured groups in the directory tree. The Group Base consists of OUs that, when combined with the base DN attribute and the configured group CNs, complete the full distinguished name of each group. The group base should specify "OU=" for each OU and multiple OUs should be separated by commas. The order of OUs should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: "OU=PureGroups,OU=SANManagers"
 - **ROLE**: The role assigned to the service. See [Roles](#) for details about FlashBlade roles.

For example:

```
pureds role setattr --group admins --group-base OU=Marketing storage_admin
```

In this example the role `storage_admin` is set as belonging the group `admins` in the group base `Marketing`.

NFS Directory Service

The NFS directory service is used for resolving addresses for netgroups, looking up the extended group membership of NFS data path clients, and looking up name/ID mappings.

Active Directory, OpenLDAP, and Oracle Unified Directory servers (that are RFC-2307 or RFC-2307bis compliant) are supported.

When configuring NFS directory service, you can optionally specify a bind user and bind password. If not specified, FlashBlade attempts to bind anonymously (without authenticating) with the configured LDAP servers. Note that if the configured LDAP servers do not grant read permissions on the needed pathways to unauthenticated machines, the anonymous bind will fail.

If a user belongs to more than 16 groups then you must configure and enable the NFS directory service in order to obtain proper user permissions. If NFS directory service is enabled, FlashBlade attempts to look up user IDs and determine the groups to which they belong.

NIS Domains and NIS Servers

NIS Domains

NIS domains can be used with quotas and normal group lookup.

NIS servers can be configured instead of LDAP servers. When NIS servers are configured, FlashBlade communicates with them over the NIS protocol in order to convert between user IDs and user names, as well as to convert between group IDs and group names.

NIS Servers

NIS servers can also be used with quotas and normal group lookup.

NIS servers can be configured instead of, or in addition to, LDAP servers. The servers can be specified as either host names or IP addresses.

NIS servers are required in order to configure the NFS directory service to use NIS.

Configuring the NFS Directory Service with LDAP

 **Note:** Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The NFS directory service enables file access for users who are defined in a directory server.

Configuring the NFS directory service is a one-time process and is only required if you are integrating the FlashBlade array with a directory service for protocol support.

To export a file system over NFS where a user might belong to more than 16 groups, the FlashBlade array must be integrated with a directory service.

To configure the NFS directory service with LDAP for protocol support, perform the following:

- 1 Verify that the directory server is accessible through a management interface.
- 2 Verify that the DNS settings can be used to resolve the directory server domain and server.
- 3 From the **Settings** sidebar, click **Directory Service** under **Security**. The **Directory Service** panel appears.
- 4 In the **Directory Service** panel, select **NFS**.

- 5 Click the **Edit** icon to the right of **Configuration**. The **Edit Directory Service Configuration** pop-up window appears.
- 6 Under the **Service Name** field, select **LDAP**.
- 7 In the **URIs** field, type the comma-separated list of up to 30 URIs of the directory servers.

Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

If specifying a domain name, it should be resolvable by the configured DNS servers.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from 0 to 255 representing a group of 8 bits.

For IPv6, specify the IP address in the form `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`). Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

Optionally specify a port. Append the port number after the end of the entire address. Default ports are 389 for LDAP, and 636 for LDAPS. Non-standard ports can be specified in the URI if they are in use.

- 8 In the **Base DN** field, type the base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.

 **Important!** Anonymous binding is supported for the NFS directory service. If configuring anonymous binding, neither the bind user nor the bind password should be specified.

- 9 In the **Bind User** field, type the username used to bind to and query the directory. For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`. For Active Directory servers, you may instead choose to enter the username - often referred to as the `sAMAccountName` or user

login name - of the account that is used to perform directory lookups. The username cannot contain the characters " [] : ; | = + * ? < > / \, and cannot exceed 104 characters in length.

- 10 In the **Bind Password** field, type the password for the bind user account.
- 11 (Optional) Select a CA certificate or certificate group to verify the authenticity of configured LDAP servers to establish a TLS (Transport Layer Security) communication for directory services.
- 12 Click **Test** to verify the configuration information. The **Test <Directory Service> Configuration** pop-up window appears, displaying the output of the test. During the directory service test, Purity//FB tests the directory service configuration to verify that the URI can be resolved and that the directory service can successfully bind and query the tree using the bind user credentials.

If the test fails, update the configuration information and retest.

Alternatively, you may save the configuration and enable the directory service at a later time.

- 13 If the test passes, set the **Enabled** toggle to enable (blue) to enable the directory service.
- 14 Click **Save**.

Configuring the NFS Directory Service with LDAP using the CLI

 **Note:** Due to limitations of Microsoft Active Directory software, a user account belonging to more than 1,015 groups may fail to login when FlashBlade's management directory service is configured to communicate with a Windows Server based computer.

The NFS directory service with LDAP can also be configured using the **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 **`pureds setattr nfs`** command. For full details about the **`pureds`** command, see the *FlashBlade CLI Reference*.

To remove the NFS with LDAP configuration using the CLI, perform the following:

- 1 Issue the **`purefs disable nfs`** command to disable NFS.
- 2 Issue the **`purefs 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 **`purefs 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 **puredds setattr nfs** command. For full details about the **puredds** command, see the *FlashBlade CLI Reference*.

To configure the NFS directory service using the CLI, perform the following:

- 1** Issue the **puredds 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:

```
puredds 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 **puredds test nfs** command. If the test fails, update the configuration information.
 - 3** If the test passes, enable the directory service using the **puredds 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.

When joining a domain, a machine account is created with the AD domain. FlashBlade then issues a keytab to that machine account, which is used for future communication with that domain. The default Organizational Unit (OU) where the machine account is created within the AD domain is "Computers".

Because SMB authentication is inherently a Kerberos authentication setup, Kerberos ports must be open for the SMB directory service configuration to function properly.

The following issues can result in loss of SMB functionality and file access:

- If domain controllers delete or otherwise invalidate FlashBlade's keytab or machine account, SMB will stop functioning.
- SMB clients do not have IDs; they have names. FlashBlade resolves the client's name to an ID using the array's SMB mode and the configured SMB directory service. If the SMB directory service is unhealthy or disabled, SMB clients will lose all file access because the name/ID cannot be resolved.



Pure Storage, Inc.

X: @purestorage

2555 Augustine Drive, 1st Floor

Santa Clara, CA 95054

800-379-7873

Sales: sales@purestorage.com

Support: support@purestorage.com

Media: pr@purestorage.com

General: info@purestorage.com