



NetApp® AltaVault™ Cloud Integrated Storage 4.4.1

Administration Guide

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: + 1 (408) 822-4501
Support telephone: +1(888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-13004_A0
March 2018

Contents

Chapter 1 - Introduction of NetApp AltaVault Cloud Integrated Storage	11
Overview of AltaVault	11
Supported backup applications and cloud destinations	11
AutoSupport	11
System requirements and specifications	11
Documentation and release notes	12
Chapter 2 - Deploying the AltaVault appliance	13
Deployment guidelines	13
Basic configuration	15
Advanced configuration	16
Configuration recovery	16
Chapter 3 - Using the AltaVault configuration wizards	17
Using the AltaVault appliance CLI configuration wizard	17
Using the Management Console	18
Connecting to the Management Console	18
Home page	19
Navigating in the Management Console	19
Getting help	20
Using the Wizard Dashboard	20
Accessing the wizard dashboard	21
Using the System Settings wizard	21
Using the Cloud Settings wizard	22
Using the import configuration wizard	34
Using the export configuration wizard	35
Chapter 4 - Configuring storage settings	37
Configuring cloud settings	37
Configuring cloud provider settings	37
Configuring encryption	38
Configuring replication	38
Configuring bandwidth limits	38
Configuring SMB	39
Configuring an Active Directory domain	40
Adding SMB shares	41

Adding users to access the share	42
Adding SMB local users.....	42
Editing multichannel settings	43
Troubleshooting SMB	43
Configuring NFS	44
Before you begin	44
Configuration tasks.....	44
Editing an NFS configuration.....	46
Troubleshooting NFS	47
Configuring OST	48
Configuring OST shares	48
Adding an OST user to access the share	49
Editing OST user information	50
Configuring SnapMirror.....	50
Enabling SnapMirror service.....	50
Monitoring and deleting SnapMirror shares and Snapshots on AltaVault	51
Enabling long-term retention.....	52
Enabling SnapCenter access.....	53
Chapter 5 - Modifying networking settings	55
Modifying general host settings	55
Modifying management interfaces	57
Modifying data interfaces	58
Modifying virtual interfaces (VIFs)	60
Modifying VLANs	61
Chapter 6 - Configuring system administrator settings	63
Setting announcements	63
Configuring alarm settings	63
Configuring date and time	69
Configuring SNMP basic settings	71
Configuring SNMP v3	73
SNMP authentication and access control	75
Configuring email settings	78
Configuring log settings	78
Chapter 7 - Configuring security settings.....	83
Configuring general security settings	83
Managing user permissions	85
Configuring permissions for user roles	87
Unlocking an account.....	88
Configuring password policy settings	88

Configuring management login from Active Directory domain	89
Configuring login from AD.....	89
Login behavior using AD	90
Setting RADIUS servers.....	91
Configuring TACACS+ access.....	92
Unlocking the secure vault	93
Configuring Web settings	94
Managing web SSL certificates.....	94
Configuring KMIP.....	97
Using the Management Console to configure KMIP	97
Using CLI to configure KMIP.....	99
Troubleshooting KMIP.....	99
Configuring appliance monitoring.....	101
Configuring a management ACL	103
Configuring SSH access and chained authentication	104
Configuring SSO for AltaVault	106
Before you begin	107
Configuring SSO	107
Enabling SSO Service	108
Troubleshooting SSO	110
Chapter 8 - Configuring AltaVault appliances for FIPS-compliant cryptography.....	111
What is FIPS?	111
Understanding FIPS on AltaVault	111
NetApp Cryptographic Security Module.....	112
Compliant FIPS cryptography features	112
Noncompliant FIPS cryptography features	112
Configuring AltaVault for FIPS compliance	113
Configuring AltaVault appliances for FIPS-compliant cryptography	113
Enabling FIPS mode.....	114
Verifying that your system uses FIPS-compliant encryption	114
Working with features to maintain FIPS compliance.....	115
Account passwords.....	115
Cipher requirements	116
Key size requirements	116
NTP.....	117
RADIUS and TACACS+.....	117
SNMP	117
SSH.....	117
Telnet server	118
Web proxy	118
Disabling FIPS mode.....	119
Verifying FIPS mode in system logs	119

Verifying that file transfers operate in FIPS mode	119
Verifying that NTP operates in FIPS mode	120
Verifying that secure vault operates in FIPS mode	120
Verifying that SNMP operates in FIPS mode	120
Verifying that the web interface operates in FIPS mode	120
FIPS CLI	120
Chapter 9 - Managing the AltaVault appliance	121
Starting and stopping the AltaVault appliance	121
Configuring scheduled jobs	122
Managing licenses	123
Managing unlicensed AltaVault appliances	124
Managing licenses using the command-line	124
Managing licenses using the Management Console	125
License limits	125
Model upgrades on the virtual AltaVault appliances	125
Activating support for AltaVault cloud-based appliances	125
Upgrading your software	126
Rebooting and shutting down AltaVault appliance	127
Viewing the current user settings	127
Managing configuration files	128
Chapter 10 - Viewing reports and logs	131
About reports	132
Viewing the storage optimization report	135
Viewing the front-end throughput report	136
Viewing the back-end throughput report	137
Viewing the eviction report	138
Viewing the replication report	139
Viewing the cloud operations report	140
Viewing schedule reports	141
Viewing per share utilization reports	142
Viewing the alarm status report	143
Viewing the CPU utilization report	146
Viewing the memory paging report	147
Viewing the interface counters report	148
Viewing the disk throughput report	149
Viewing the disk IOPS report	150
Viewing the disk utilization report	151

Viewing logs	152
Viewing system logs	152
Viewing user logs	153
Downloading log files	154
Generating system dumps	155
Viewing process dumps	156
Capturing and uploading TCP dumps	156
Viewing a TCP dump	161
Viewing the appliance monitoring report	162
Viewing the shelf details	164
Viewing the storage RAID group	165
Viewing offline file system check page	165
Viewing online file system check page	166
Viewing the verify tool diagnostics	167
Chapter 11 - Transferring data to the cloud using Amazon Snowball	169
Prerequisites	169
Guidelines for using Snowball with AltaVault	169
Seeding data using Snowball	170
Creating a Snowball job in AWS	170
Transferring data from AltaVault to Snowball	171
Managing data transfers on AltaVault	172
Verifying and completing data transfer	173
Chapter 12 - Migrating data to a new cloud	175
Cloud migration overview	175
Cloud-to-cloud migration	176
Canceling cloud-to-cloud migration	177
Amazon S3 or S3-IA to Glacier migration	177
Amazon S3 to S3-IA or Amazon S3-IA to S3 migration	178
Chapter 13 - Migrating data between appliances	179
Data migration overview	179
Data migration connection diagrams	180
Data migration process	182
Prerequisites	182
Prerequisites for the source appliance	183
Prerequisites for the target appliance	184
Performing appliance data migration	184

Post-data migration procedure.....	185
Chapter 14 - Disaster recovery.....	187
Disaster recovery preparations	187
Exporting the configuration file	188
Disaster recovery testing	188
Suspending replication at the production site.....	188
Enabling AltaVault for a disaster recovery test	188
Data restoration for disaster recovery testing	189
Performing post-DR testing activities	190
Disaster recovery	190
Enabling AltaVault for disaster recovery	191
Data restoration for disaster recovery.....	191
Chapter 15 - System components AVA400, AVA800.....	193
AltaVault appliance components	193
System chassis specifications	194
What you need to know about expansion shelves.....	194
Using LEDs to check the status of the system.....	195
Field replaceable units	197
Slot numbering and associated components.....	198
Fan modules and their LEDs	198
Fan redundancy policy	200
Power supplies and their LEDs	201
Power supply LED behaviors.....	201
Controller components and their LEDs	203
Controller LED behaviors	203
Internal FRUs	206
Chapter 16 - System maintenance AVA400, AVA800	209
Remote management port setup	210
Setting the Service Processor password	210
Configuring the remote management port.....	210
Validating the remote management port.....	211
Shutting down the AltaVault controller	212
Replacing controllers	213
Installing a controller in a chassis.....	216
Replacing a controller chassis	218
Hot-swapping controller fan modules	219
Hot-swapping controller power supplies	222
Adding disk shelves.....	225

Changing the shelf ID for a disk shelf	225
Adding an additional RAID group to a configured appliance	227
Replacing a faulty hard disk drive on an AltaVault AVA400 or AVA800 appliance	228
Replacing internal FRUs	228
Replacing a boot device in a controller	229
Replacing system DIMMs	233
Replacing RAID controllers	237
Replacing the RTC clock coin battery	237
Replacing disk shelf power supplies and other FRUs	240
Returning failed parts	240
Disposing of batteries	240
Appendix A - Administrator's configuration worksheet	241
Configuration worksheet	241
Appendix B - AltaVault appliance MIB	245
Accessing AltaVault appliance MIB	245
SNMP traps	245
Appendix C - Amazon AWS IAM and S3 bucket policies	257
Typical AltaVault setup	257
IAM policies for AltaVault	257
Sample of IAM policy	258
Bucket policies for AltaVault	259
Sample of bucket policy	259
Appendix D - Best practices for restoring data from archive	261
Optimizing data movement	261
Protecting data	261
Recovering data from archive	262
Restoring data from the cloud using the prepopulation page	263
Restoring data from the cloud using the command-line interface	264
Automatic prepopulation	268
AltaVault appliance best practices for EMC NetWorker	268
AltaVault appliance best practices for IBM Spectrum Protect	270
AltaVault appliance best practices for Veritas NetBackup	271
AltaVault appliance best practices for Veritas Backup Exec	272
AltaVault appliance best practices for Veeam backup and replication	273

Copyright Information	275
Trademark Information.....	277
How to Send Your Comments	279
Index	281

CHAPTER 1 Introduction of NetApp AltaVault Cloud Integrated Storage

Overview of AltaVault

AltaVault appliance is a disk-to-disk data storage optimization system with unique cloud storage integration. There are three types of AltaVault deployments:

- Physical hardware appliances, available in AVA400 and AVA800 models.
- Virtual appliance, available in AVA-v2, AVA-v8, AVA-v16, and AVA-v32 models.
- Cloud-based virtual appliance:
 - Amazon Machine Images (AMI), available in AVA-c4, AVA-c8, and AVA-c16 models.
 - Microsoft Azure Virtual Machine (AVM), available in the AVA-c4 model.

Supported backup applications and cloud destinations

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the product and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

AutoSupport

AltaVault supports user-triggered and daily AutoSupports (ASUPs) as well as certain event-based triggers. ASUP functionality is supported on all AltaVault models. For event-based triggers, see “[Viewing the alarm status report](#)” on [page 143](#).

For more information on ASUP CLI commands, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Guide*.

System requirements and specifications

This section specifies the hardware and software requirements.

For system requirements for virtual appliances, see the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances*.

For system requirements for cloud, see the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Cloud Appliances*.

For system requirements for physical appliances, see [Chapter 15, “System components AVA400, AVA800.”](#)

Documentation and release notes

To obtain the most current version of all NetApp documentation, go to the NetApp Support site at <https://mysupport.netapp.com>.

CHAPTER 2 Deploying the AltaVault appliance

This chapter includes the following sections:

- [“Deployment guidelines” on page 13](#)
- [“Basic configuration” on page 15](#)
- [“Advanced configuration” on page 16](#)
- [“Configuration recovery” on page 16](#)

Deployment guidelines

- AltaVault is supported with the backup applications and cloud storage providers identified by the IMT (interoperability matrix tool).

Refer to the [IMT](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

- An AltaVault can only be pointed to one cloud storage provider at a time.

If an existing AltaVault needs to be pointed to a different cloud storage provider than the one currently configured, you must clear the AltaVault cache before reconfiguring the new cloud storage provider credentials. All existing data associated with the previous cloud storage provider will remain.

- AltaVault can be deployed in one of two modes: Backup mode or Cold Storage mode. Once deployed, you cannot change the mode. Use the following table to make a comparison of using AltaVault in backup mode versus cold storage mode:

Modes	Pros	Cons
Backup mode	<ul style="list-style-type: none"> • Allows access to the most recent backups on cache. • Allows global deduplication of all data received by AltaVault, leading to higher deduplication rates. • Maximizes data movement efficiency of the WAN through deduplication of data. • Cache expansion capability via add on shelves allows for growth as needed by the business. • Higher ingest performance than Cold Storage mode. 	<ul style="list-style-type: none"> • Cloud capacity managed limited to a maximum of up to 5 times the usable space on the AltaVault's disk cache.
Cold storage mode	<ul style="list-style-type: none"> • Protects archive workloads for long periods of time, typically to cool or cold cloud storage tiers. • Allows access to far greater cloud capacity (Up to 10PB of storage, based on 1.333 billion files of 100MB average file size). • Provides expansive long term storage in just one head controller unit. 	<ul style="list-style-type: none"> • Minimal deduplication compared to backup mode. • Limited network and WAN performance, dependent on average, file size of objects sent to AltaVault. • Only available on AVA400 48 TB and virtual models. • No expansion capability with shelves. • Restores are always from the cloud provider.

- You can configure AltaVault folder shares to help describe a policy target.

For example, you can configure a backup application to direct critical system backups to point to a specific folder on one AltaVault data connection, while noncritical backups might be directed by a backup application to point to another folder on another AltaVault data connection. This method helps balance priorities of data over the network and organize data for recovery in case of a disaster. If possible, organize your backup policies to drive similar data to the same AltaVault unit.

For example, if you are backing up a Windows server farm to multiple AltaVault appliances, operating system backups are likely to have the best deduplication rates when grouped together to the same AltaVault. File and application server backups obtain better deduplication when grouped together as well.

- AltaVault exports its configuration to a file called `altavault_config_(HOSTNAME)_(DATETIME).tgz`.

NetApp recommends that you store the configuration file in different physical locations. The configuration file contains information about the configuration, including the encryption key. Alternatively, you can just export the encryption key alone.

Note: To access the encrypted data, you need an encryption key. If you lose the encryption key, AltaVault cannot reconstitute the encrypted data.

Basic configuration

This procedure assumes that you have already installed your AltaVault appliance as described in the respective installation guide:

AltaVault Model	Installation Guide
AltaVault physical appliance (AVA400, AVA800)	<i>AltaVault System Installation and Setup Instructions</i> (poster)
AltaVault virtual appliance (Microsoft Hyper-V, VMware ESXi, or Linux KVM)	<i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances</i>
AltaVault cloud appliance (Amazon Machine Image or Microsoft Azure Virtual Machine)	<i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Cloud Appliances</i>

After installing the appliance, use the following table to guide your initial AltaVault setup and deployment:

Step	Configuration	Reference
1	Gather configuration information	Appendix A, “Administrator’s configuration worksheet.”
2	Provide the initial system configuration using the CLI Wizard	“Using the AltaVault appliance CLI configuration wizard” on page 17
3	Set the Service Processor password (AVA400, AVA800 models only)	“Setting the Service Processor password” on page 210
4	Configure the remote management port (AVA400, AVA800 models only)	“Configuring the remote management port” on page 210
5	Connect to the Management Console and log in	“Using the Management Console” on page 18
6	Configure the system settings from the System Setup Wizard	“Using the System Settings wizard” on page 21
7	Configure cloud service provider settings using the Cloud Setup Wizard	“Using the Cloud Settings wizard” on page 22
8	Add the license information (virtual appliance models only)	“Managing licenses using the Management Console” on page 125
9	Optionally, activate support (cloud appliance models only)	“Activating support for AltaVault cloud-based appliances” on page 125
10	Configure data interfaces	“Modifying data interfaces” on page 58
11	Optionally, configure virtual interfaces (VIFS)	“Modifying virtual interfaces (VIFS)” on page 60
12	Optionally, configure VLAN interfaces	“Modifying VLANs” on page 61
13	Optionally, join the domain (for SMB)	“Configuring SMB” on page 39
14	Configure storage (select SMB, NFS, OST, or SnapMirror)	“Configuring SMB” on page 39 “Configuring NFS” on page 44 “Configuring OST” on page 48 “Configuring SnapMirror” on page 50
15	Save your configuration to a safe location using the Export Wizard	“Using the export configuration wizard” on page 35

Advanced configuration

The following table summarizes AltaVault's advanced configuration options.

Configuration option	Setting	Reference
Storage settings	Advanced storage settings for SMB, NFS, OST, and SnapMirror	Chapter 4, “Configuring storage settings”
	Configure data prepopulation	“Restoring data from the cloud using the prepopulation page” on page 263
Security settings	Set authentication method, Active Directory (AD) administration, role-based permissions for users, Secure Vault, web settings, REST API access, key management (KMIP), management ACLs, SSH access and chained authentication, Single Sign-On (SSO)	Chapter 7, “Configuring security settings”
	Configure FIPS compliance	Chapter 8, “Configuring AltaVault appliances for FIPS-compliant cryptography”
System administration settings	Set announcements, alarms, date and time, SNMP, email notifications, log settings	Chapter 6, “Configuring system administrator settings”
System monitoring	Viewing reports and logs	Chapter 10, “Viewing reports and logs”
	System monitoring	
	• Schedule jobs	“Configuring scheduled jobs” on page 122
	• Schedule reports	“Viewing schedule reports” on page 141
	• LEDs (AVA400, AVA800 only)	“Using LEDs to check the status of the system” on page 195
	Peer monitoring	“Configuring appliance monitoring” on page 101

Configuration recovery

In the event of a catastrophic event, it might be necessary to recover your configuration if previously saved to another location using the Export Wizard. To recover a saved configuration, see [Chapter 14, “Disaster recovery.”](#)

CHAPTER 3 Using the AltaVault configuration wizards

This chapter includes the following sections:

- [“Using the AltaVault appliance CLI configuration wizard” on page 17](#)
- [“Using the Management Console” on page 18](#)
- [“Using the Wizard Dashboard” on page 20](#)

Using the AltaVault appliance CLI configuration wizard

After installing the AltaVault appliance and logging in for the first time, you are prompted to enter initial system information using command-line interface (CLI).

To run the AltaVault appliance CLI configuration wizard

1. Complete the configuration wizard steps on the client side and server side.

Wizard prompt	Description	Example
Step 1: Admin password?	NetApp requires that you change the default administrator password (<code>password</code>) at this time. The new password must be a minimum of eight ASCII characters and cannot be the word <code>password</code> .	Step 1: Admin password? xxxxyyyy
Step 2: Host name?	Enter the host name for the AltaVault appliance.	Step 2: Hostname? amnesiac
Step 3: Use DHCP on the primary interface?	For AltaVault virtual and physical appliances, DHCP is not recommended. For AltaVault cloud-based virtual appliances, DHCP is required.	Step 3: Use DHCP? yes
Step 4: Primary IP address?	Enter the IP address.	Step 4: Primary IP address? 10.10.10.6
Step 5: Netmask?	Enter the netmask address.	Step 5: Netmask? 255.255.0.0
Step 6: Default gateway?	Enter the default gateway.	Step 6: Default gateway? 10.0.0.1

Wizard prompt	Description	Example
Step 7: Primary DNS server?	Enter the primary DNS server IP address. If you do not specify a valid DNS server, the system does not start.	Step 7: Primary DNS server? 10.0.0.2
Step 8: Domain name?	Enter the domain name for the network that the appliance is connected to. If you set a domain name, you can enter host names in the system without the domain name.	Step 8: Domain name? example.com

- To change an answer, enter the step number to return to. Otherwise press <enter> to save changes and exit. The AltaVault appliance configuration wizard automatically saves your configuration settings. The CLI prompt appears:

```
hostname >
```

- If you are not using DHCP skip to Step 4. If you selected to use DHCP, get the IP address of the appliance by running the following commands:

```
hostname > enable
hostname # configure terminal
hostname (config)# show interfaces primary
```

- To log out of the system, enter `exit` at each of the command-level prompts.

You can now log in to the appliance using a web-based client to access the Management Console (user interface) and Wizard Dashboard for configuring system and cloud service provider (CSP) settings.

Using the Management Console

This section includes the following information:

- [“Connecting to the Management Console” on page 18](#)
- [“Home page” on page 19](#)
- [“Navigating in the Management Console” on page 19](#)

Connecting to the Management Console

To connect to the AltaVault Management Console

- Enter the URL for the Management Console in the location box of your Web browser:

```
https://<host>.<domain>
```

When you connect using HTTPS, you are prompted to inspect and verify the SSL certificate. The SSL certificate is a self-signed certificate used to provide encrypted Web connections to the Management Console. It is re-created when the appliance hostname is changed and when the certificate has expired.

The `<host>` variable is the hostname you assigned to the AltaVault primary interface in the configuration wizard. If your DNS server maps that IP address to a name, you can specify the DNS name.

The `<domain>` variable is the full domain name for the appliance.

You can also specify the IP address instead of the host and domain name.

2. In the Username text box, specify the user login. The default login is admin.
3. In the Password text box, specify the password you assigned in the CLI configuration wizard of the AltaVault. The password cannot be “password.” To change your password, see [“Viewing the current user settings” on page 127.](#)
4. Click **Sign In** to display the AltaVault configuration wizard (when you log in for the first time) or the Home page (for subsequent logins).

Home page

The Home page displays the following parameters:

- Cloud and Disk Storage Allocation - The outer circle represents the cloud storage and the inner circle represents the local AltaVault cache storage. This section also lists the used storage, free storage, and total storage on the cloud and the disk.
- Optimization Service - Specifies whether the Storage Optimization Service is running or has stopped and the status of the service:

Status	Description
Ready	Storage Optimization Service is ready to ingest and replicate data to the cloud.
Not ready	Storage Optimization Service is unavailable. No data will be ingested or replicated.
Replaying	Storage Optimization Service has been terminated during backup replication, either due to loss of power or a crash. During this replay process, the AltaVault verifies data consistency from its transaction logs. The amount of time it takes to replay process to complete will depend on the amount of data in flight at the time the AltaVault appliance was abnormally stopped.
Upgrading	Storage Optimization Service is unavailable due to an in-progress upgrade. No data will be ingested or replicated.

- Cloud Storage Reclamation - Provides the completion percentage of the cloud storage reclamation service (garbage collection). This service runs automatically when needed to clean up fragmented disk and cloud space.
- Alarms Triggered - Displays the appliance health status and software update. To view the alarms triggered, choose Reports > Alarm Status.
- System Status - Displays details such as the AltaVault time, system up time, and optimization service up time.
- Appliance Information - Provides the appliance hostname and its model number.
- Replicated Data - Displays the status of the process of copying data and metadata from the AltaVault to the cloud.
- Storage Optimization - Displays the expanded data, deduplicated data, and deduplication factor. Expanded data is the data that has been backed up locally by the AltaVault. Deduplicated data reflects data that has been optimized through the use of deduplication and compression. Deduplication factor is the ratio of the expanded data and total optimized data. The total optimized data includes both deduplication and compression savings.
- Cloud Information - Displays the status of the cloud connection that the appliance is configured to communicate with.

Navigating in the Management Console

You navigate to the tools and reports available to you in the Management Console using cascading menus.

Saving your configuration

As you apply configuration settings, the values are applied to the running configuration. Most Management Console configuration pages include an **Apply** button for you to commit your changes. When you click **Apply**, the Management Console updates the running configuration.

NetApp recommends that you export your configuration after every change.

A red asterisk next to a control indicates that the field is required. You must specify a valid entry for all of the required controls on a page before saving the changes.

Restarting AltaVault appliance service

Some configuration settings require a restart the services in order for the changes to take effect.

To restart the service, click **Restart** to display the Service page or go to Storage Optimization Service page and restart the service.

Printing pages and reports

You can print Management Console pages and reports using the print option on your Web browser.

To print pages and reports

- Choose File > Print in your Web browser to open the Print dialog box.

Getting help

The Help page provides the following options:

- Online Help - View browser-based online help.
- Technical Support - View links and contact information for NetApp Support.
- Appliance Details - View appliance information such as the model number, hardware revision type, serial number, and software version currently installed on the appliance.

Displaying online help

The Management Console provides page-level help for the appliance.

To display online help

- Click the question mark icon next to the page title. The help for the page appears in a new browser window.

Logging out

In the menu bar, click **Sign out** to end your session.

Using the Wizard Dashboard

The AltaVault configuration wizard appears only after you log in to the appliance for the first time. It enables you to access other configuration wizards, so that you can configure your own system settings, configure cloud settings, and import and export settings.

This section includes the following topics:

- [“Accessing the wizard dashboard” on page 21](#)
- [“Using the System Settings wizard” on page 21](#)
- [“Using the Cloud Settings wizard” on page 22](#)
- [“Using the import configuration wizard” on page 34](#)
- [“Using the export configuration wizard” on page 35](#)

Accessing the wizard dashboard

1. From a web browser, enter the AltaVault IP address to log in to the Management Console.
2. If you are logging in to the Management Console for the first time, the wizard appears, displaying the Welcome page. For subsequent logins, log in to AltaVault and choose **Configure > Setup Wizard**.

Based on your configuration requirements, you can use different wizards from the dashboard.

Task	Reference
Configure networking settings and time zone, use the System Settings wizard.	“Using the System Settings wizard” on page 21
Configure cloud settings, licenses, and encryption key, use the Cloud Settings wizard.	“Using the Cloud Settings wizard” on page 22
Import a previously saved configuration, use the Import Configuration wizard.	“Using the import configuration wizard” on page 34
Export the current configuration from the system, use the Export Configuration wizard.	“Using the export configuration wizard” on page 35

Using the System Settings wizard

Use the System Settings wizard to configure networking settings and time zone.

To use the System Settings wizard

1. From the management console, choose **Configure > Setup Wizard**.
2. Select **System Settings** in the Wizard Dashboard.
The System Settings wizard displays the hostname and DNS server IP address for the AltaVault.
3. In the System Settings wizard, complete the configuration as described in this table.

Control	Description
Obtain IPv4 Address Automatically	Specify this option to automatically obtain the IPv4 address from a valid DHCP server.
Enable IPv4 Dynamic DNS	Enable IPv4 Dynamic DNS - Select this option to enable IPv4 dynamic DNS on the primary interface.

Control	Description
Specify IPv4 Address Manually	Specify this option if you do not use a DHCP server to set the IP address. Specify the following settings: <ul style="list-style-type: none"> IPv4 Address - Specify an IPv4 address. IPv4 Subnet Mask - Specify an IPv4 subnet mask. Default IPv4 Gateway - Specify the default primary gateway IPv4 address. The primary gateway must be in the same network as the primary interface. You must set the primary gateway for interface configurations.
Time Zone	Specify the country and time zone in which the AltaVault is located.
Enable Analytics	Enabling this feature will send daily informational AutoSupport (ASUP) messages to NetApp.

4. Click **Next** to display the Confirmation page.
5. Click **Save and Apply** to display the System Settings Wizard Finish page.
6. Click **Exit** to close the System Settings Wizard and go back to the dashboard.

Using the Cloud Settings wizard

Use the Cloud Settings wizard to configure cloud settings, licensing, and the encryption key.

To use the Cloud Settings wizard

1. From the management console, choose Configure > Setup Wizard.
2. Select Cloud Settings in the Wizard Dashboard.

Note: Check that the datastore is empty. If the datastore is not empty, you cannot change the cloud provider, region, hostname, and bucket name.

3. Under Provider, select and configure your preferred cloud service provider from the drop-down list:

Note: If you are configuring a private cloud, see [“Customizing a private cloud” on page 34](#).

- Alibaba Cloud Object Storage Service - see [“Configuring Alibaba Cloud Object Storage Service \(OSS\)” on page 24](#)
- Amazon Glacier - see [“Configuring Amazon Glacier storage” on page 25](#)
- Amazon S3 - see [“Configuring Amazon S3 storage” on page 27](#)
- AT&T Synaptic Storage - see [“Configuring Atmos-based storage” on page 28](#)
- Cleversafe Cloud Storage - see [“Configuring S3-based storage” on page 31](#)
- Cloudian Cloud Storage - see [“Configuring S3-based storage” on page 31](#)
- Cloudwatt Object Storage - see [“Configuring SWIFT-based storage” on page 32](#)

- EMC Atmos- see [“Configuring Atmos-based storage” on page 28](#)
 - Google Cloud Storage - see [“Configuring Google Cloud Storage” on page 29](#)
 - HGST Storage - see [“Configuring S3-based storage” on page 31](#)
 - Internet Initiative Japan (IIJ) - see [“Configuring S3-based storage” on page 31](#)
 - Microsoft Windows Azure Storage - see [“Configuring Microsoft Windows Azure storage” on page 30](#)
 - NetApp StorageGRID Webscale - see [“Configuring S3-based storage” on page 31](#)
 - OpenStack Object Storage (Swift) - see [“Configuring SWIFT-based storage” on page 32](#)
 - Oracle Storage Cloud Service - Object Storage - see [“Configuring SWIFT-based storage” on page 32](#)
 - Outscale On-Demand Storage - see [“Configuring S3-based storage” on page 31](#)
 - Rackspace Cloud Files - see [“Configuring SWIFT-based storage with region-selection” on page 33](#)
 - S3 Compliant Connector - see [“Configuring S3-based storage” on page 31](#)
 - Scalify RING - see [“Configuring S3-based storage” on page 31](#)
 - SoftLayer Object Storage (Swift) - see [“Configuring SWIFT-based storage with region-selection” on page 33](#)
 - Swisscom Dynamic Storage- see [“Configuring Atmos-based storage” on page 28](#)
 - Verizon Cloud Storage - see [“Configuring S3-based storage” on page 31](#)
4. Configure Encryption Settings in the Wizard Dashboard. This page is only available to users with Read-Only Security Settings permissions or Read and Write Security Settings permissions. Specify the following items:

Control	Description
Create New Datastore Encryption Key	<p>Select this option to establish a new AES-256 bit encryption key that AltaVault uses to secure data.</p> <p>Set Key Passphrase - Optionally, specify a passphrase that will be used to secure the encryption key on AltaVault. This passphrase will be required when importing the encryption key or AltaVault configuration onto a new AltaVault appliance. The passphrase is not stored within a configuration archive and must be kept in a secure location.</p> <p>Confirm Key Passphrase - Confirm the passphrase.</p>
Import Key from File	Select this option to import the key from a file. Select the file to import it onto the appliance. The key must be the key that was generated by an AltaVault appliance.
Import Key from Text	Select this option to import the key from text. The key must be the key that was generated or exported from an AltaVault appliance.
Use Key from KMIP server	Select this option to use the key from the KMIP server. Select a key from the drop-down list.

5. On the Confirmation page verify the information, and click **Save and Apply**.

Note: It is recommended to use a firewall to prevent unauthorized connections to the AltaVault.

Configuring Alibaba Cloud Object Storage Service (OSS)

1. Select **Yes** or **No** to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. Specify a storage class from the drop-down list:
 - Standard
 - Infrequent Access
 - Archive
3. Specify the **Access Key**.
4. Specify the **Secret Key** for your cloud service provider account.
5. Enter the **Hostname** of the cloud provider on which AltaVault stores the replicated data.
6. Specify the **Bucket Name** associated with your cloud service provider account. If the bucket name does not exist, the bucket is created during initial AltaVault replication. Bucket names must be a series of one or more labels separated by a period. If Archive is selected as Storage Class, the bucket should not be created through Alibaba Storage Console. For Archive storage class, AltaVault creates an additional bucket in Standard tier with the name <bucketname>-avastd, which includes the metadata files required for disaster recovery. The archive bucket will automatically be created by AltaVault.

For Alibaba, the bucket names must be a valid DNS name, conforming to the following naming rules:

 - Container names must start with a letter or number and can contain only letters, numbers, and hyphens.
 - Every hyphen must be immediately preceded and followed by a letter or number. You cannot use consecutive hyphens.
 - All letters in a bucket names must be lowercase.
 - Container names must be from 3 to 63 characters (maximum 63 bytes in UTF-8).
7. Specify the port through which replication occurs. Port 443 is recommended.
8. Select **Enable Archiving** if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
9. Optionally, select **Enable Cloud Deduplication**. Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended for Alibaba Archive to improve recovery of recently written data from cache, but can decrease deduplication rates and increase cloud storage costs.
10. Optionally, select **Enable Cloud CA Certificate** to specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.

11. Select **Enable Proxy** to enable proxy server settings. A proxy server acts as an intermediary for requests from clients seeking resources from other servers.

After you select the check box, specify the hostname or the IP address, port number, for access, username and password.

12. Click **Apply**.

Configuring Amazon Glacier storage

1. Select yes or no to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. Specify the Region. You can choose to store your data in the Amazon Glacier Region that meets your regulatory, throughput, and geographic redundancy criteria.

When specifying US East (N. Virginia) or us-east-1 as the region, use US Standard.
3. Custom Region - Optionally, specify the custom region for your cloud service provider account.
4. Authentication Method - Specify one of the following:
 - Standard - Specify selections for the [“Standard authentication” on page 25](#).
 - STS - Specify selections for the [“STS authentication” on page 26](#).

Note: If user files are not cached on AltaVault, they should be pre-populated before reads are performed. This is because restores from Amazon Glacier have a latency of up to 12 hours depending on the retrieval option. For more information, see AWS documentation.

Standard authentication

Note: When S3 or Glacier is configured and Storage Optimization Service fails to start, the logs may contain the error “BucketAlreadyExists: The requested bucket name is not available. The bucket namespace is shared by all users of the system. Please select a different name and try again.” This indicates that the chosen bucket name is not available. You can resolve this by selecting a different bucket name. This error may also be encountered after a cloud migration or changing of the cloud settings. One possible reason may be that the cloud credentials do not belong to the account that owns the bucket. Double-check the credentials and ensure that the correct credentials are entered on the Cloud Settings page.

For the Standard authentication type, make selections for the following:

1. Access Key - Specify the access key for your Amazon S3 (AWS) account.
2. Secret Key - Specify the secret key for your cloud service provider account.
3. Hostname - Verify the hostname of the cloud provider on which AltaVault stores the replicated data.

4. **Bucket Name** - Specify the bucket name associated with your cloud service provider account. If the bucket name does not exist, the bucket is created during initial AltaVault replication. Bucket names must be a series of one or more labels separated by a period
5. **Port** - Specify the port through which replication occurs. Ports 80 or 443 are available.
6. **Enable Archiving** - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
7. **Enable Cloud Deduplication** - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended for Amazon Glacier to improve recovery of recently written data from cache, but can decrease deduplication rates and increase cloud storage costs.
8. **Enable Cloud CA Certificate** - Optionally, specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.
9. **Enable Proxy** - Select to enable proxy server settings. A proxy server acts as an intermediary for requests from clients seeking resources from other servers.

After you select the check box, specify the following settings:

- **Hostname/IP address** - Specify the hostname or the IP address
- **Port** - Specify the port numbers for access
- **Username** - Specify the name of the user for access
- **Password** - Specify the user's password.

STS authentication

1. **Identity Provider URL**: Specify the URL of the provider.

The identity provider is a server that performs two roles: 1) authenticating users and machines wishing to access Amazon AWS services, and 2) providing temporary security credentials with which to access those services. AltaVault makes a call to the identity provider, which in turn makes a call to Amazon STS using the AssumeRole API call to generate temporary security credentials, and then passes these credentials back to AltaVault.

2. **Parameters** - Specify the parameters that the provider expects to authenticate the AltaVault appliance.
3. **Response Type** - JSON is the default.
4. **Method** - Select GET or POST.
5. **CA Certificate** - Specify the certificate that will be used to validate the server certificate of the identity provider. Ensure that the file has the required .pem extension.
6. **Select the Web Settings page link**.
 - **Select the Replace tab**.
 - **Certificate** - Upload the client certificate.
 - **Separate Private Key** - Upload the Private Key.
 - **To replace the certificate and private key**, click Import Certificate and Key.
7. **Hostname** - Verify the hostname of the cloud provider on which AltaVault stores the replicated data.

8. **Bucket Name** - Specify the bucket name associated with your cloud service provider account. If the bucket name does not exist, the bucket is created during initial AltaVault replication. Bucket names must be a series of one or more labels separated by a period.
9. **Port** - Specify the port through which replication occurs. Ports 80 or 443 are available.
10. **Enable Archiving** - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
11. **Enable Cloud Deduplication** - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended for Amazon Glacier to improve recovery of recently written data from cache, but can decrease deduplication rates and increase cloud storage costs.
12. **Enable Cloud CA Certificate** - Optionally, specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.
13. Select the **Enable Proxy** check box to enable proxy server settings and specify:
 - **Hostname/IP address** - Specify the hostname or the IP address
 - **Ports** - Specify the port numbers for access
 - **Username** - Specify the name of the user for access
 - **Password** - Specify the user’s password

Configuring Amazon S3 storage

1. Select yes or no to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. Specify the Region. You can choose an Amazon S3 region to optimize for latency, minimize costs, or address regulatory requirements.
3. **Custom Region** - Optionally, specify the custom region for your cloud service provider account.
4. **Storage Class** - Specify a storage class from the drop-down list:
 - **Standard** (Standard storage class)
 - **Standard-IA** (Standard Infrequent Access)
 - **RRS** (Reduced Redundancy Service)
5. **Authentication Method**- Specify one of the following types:
 - **Standard** - Specify selections for the [“Standard authentication” on page 25](#).
 - **STS** - Specify selections for the [“STS authentication” on page 26](#).

Configuring Atmos-based storage

1. Select yes or no to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. If your provider is AT&T, specify the following settings:

Storage Policy - Select one of the following storage policies from the drop-down list:

 - Local Replication - Stores data stored in one location and protects it using erasure coding.
 - Remote Replication - Stores data in two locations maintains a copy in one data center and replicates it to a geographically remote data center.
3. Specify the following settings:
 - Subtenant ID - Specify the subtenant ID that EMC Atmos uses to authenticate each request.
 - UID - Specify the user ID that EMC Atmos uses to authenticate each request.
 - Shared Secret - Specify the shared secret that EMC Atmos uses to authenticate each request. When the client application builds a Web service request, EMC Atmos uses the shared secret to create a signature entry as a part of the request. The shared secret must be associated with the subtenant ID and application ID created by the EMC Atmos-based storage provider.
4. Specify the hostname.
5. Specify the bucket name associated with your cloud service provider account. You can use buckets to organize your data and control access to your data, but they cannot be nested. If the bucket name does not exist, the bucket is created during initial AltaVault replication.
6. Specify the port number.
7. Enable Archiving - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
8. Enable Cloud Deduplication - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended only for Alibaba Archive, Amazon Glacier, and Azure Archive.
9. Enable Cloud CA Certificate - Optionally, specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.
10. Select the Enable Proxy check box to enable proxy server settings and specify:
 - Hostname/IP address - Specify the hostname or the IP address.
 - Ports - Specify the port numbers for access.
 - Username - Specify the name of the user for access.
 - Password - Specify the user's password.

Configuring Google Cloud Storage

1. Select yes or no to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. Specify the Location from the drop-down list.
3. Storage Class - Specify the storage class from the drop-down list:
 - Standard (Standard storage class)
 - Nearline
4. Project ID - Specify the unique project ID associated with the bucket.
5. Client email - Specify the service account email address value from the API Manager > Credentials page of the Google developers console.
6. Private Key - Select Browse to specify the private key for your Google Cloud Storage service provider account.

Google provides the private key in JSON and PKCS12 format. The AVA cloud credentials page requires a private key with a required extension of .pem or .json. You can read the client email and project ID from the .json file.

Note: When connecting to Google Cloud storage with FIPS enabled, AltaVault requires all imported and generated keys sizes for RSA-based and DSA-based certificates to be 2048 bits or higher. Connections to using 1024-bit certificates will not complete. It is recommended to generate a new private key (2048-bit or higher) for Google Cloud Storage, save it in a .json file, and upload that file when configuring AltaVault with Google Cloud Storage.

7. Specify the hostname.
8. Specify the bucket name associated with your cloud service provider account. If Nearline is selected as Storage Class, the bucket should not be created through Google Developers Console. The Nearline bucket will automatically be created by AltaVault. For Nearline storage class, AltaVault creates an additional bucket in Standard tier with the name <bucketname>-avastd, which includes the metadata files required for disaster recovery. You can use buckets to organize your data and control access to your data, but bucket cannot be nested.

For more information on bucket name restrictions, see [Google documentation](#).
9. Specify the port number. Port 443 is recommended.
10. Enable Archiving - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
11. Enable Cloud Deduplication - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended only for Alibaba Archive, Amazon Glacier, and Azure Archive.

12. Enable Cloud CA Certificate - Specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be a .pem or .json extension file.
13. Select the Enable Proxy check box to enable proxy server settings and specify:
 - Hostname/IP address - Specify the hostname or the IP address.
 - Ports - Specify the port numbers for access.
 - Username - Specify the name of the user for access.
 - Password - Specify the user's password.

Configuring Microsoft Windows Azure storage

1. Select yes or no to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. Specify the following settings:
 - Cloud Type - Select your option from the drop-down list. The options are Azure Government or Azure Public. Use a storage account to access the Cool or Hot access tier.
 - Storage Class - Specify Standard or Archive from the drop-down list.
 - Storage Account - Specify the Microsoft Azure Storage account name. For Standard storage class, AltaVault supports accounts belonging to either the Hot or Cool access tier.
 - Primary or Secondary Access Key - Specify the primary or secondary Microsoft Azure Storage access key that you generated when you created your Microsoft Azure Storage account. The secondary key provides the same access as the primary key and is used for backup purposes.
3. Specify the hostname.
4. Bucket Name - Specify the container name associated with your cloud service provider account. You can use containers to organize your data and control access to your data, but they cannot be nested. If the container name does not exist, the container is created during initial AltaVault replication.

For Azure, the bucket names must be a valid DNS name, conforming to the following naming rules:

 - Container names must start with a letter or number and can contain only letters, numbers, and hyphens.
 - Every hyphen must be immediately preceded and followed by a letter or number. You cannot use consecutive hyphens.
 - All letters in a bucket names must be lowercase.
 - Container names must be from 3 to 63 characters (maximum 63 bytes in UTF-8).
5. Specify the port number.
6. Enable Archiving - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).

7. Enable Cloud Deduplication - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended for Azure Archive to improve recovery of recently written data from cache, but can decrease deduplication rates and increase cloud storage costs.
8. Enable Cloud CA Certificate - Optionally, specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.
9. Select the Enable Proxy check box to enable proxy server settings and specify:
 - Hostname/IP address - Specify the hostname or the IP address.
 - Ports - Specify the port numbers for access.
 - Username - Specify the name of the user for access.
 - Password - Specify the user's password.

Configuring S3-based storage

1. Select yes or no to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. Specify the following settings:
 - Access Key - Specify the access key (same as the username).
 - Secret Key - Specify the secret key (password).
3. Specify the web protocol: HTTP or HTTPS.
4. Specify the hostname.
5. Specify the bucket name associated with your cloud service provider account. You can use buckets to organize your data and control access to your data, but they cannot be nested. If the bucket name does not exist, the bucket is created during initial AltaVault replication.
6. Specify the port number.
7. Enable Archiving - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
8. Enable Cloud Deduplication - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended only for Alibaba Archive, Amazon Glacier, and Azure Archive.
9. Enable Cloud CA Certificate - Optionally, specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.
10. Select the Enable Proxy check box to enable proxy server settings and specify:

- Hostname/IP address - Specify the hostname or the IP address.
- Ports - Specify the port numbers for access.
- Username - Specify the name of the user for access.
- Password - Specify the user's password.

Configuring SWIFT-based storage

1. Select yes or no to use keys from KMIP server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for each of the authentication fields.
2. If your cloud service provider is Oracle Storage Cloud Service - Object Storage, specify the following settings:
 - Storage Class - By default, Storage Class is set to Standard.
3. Specify the following settings:
 - Authentication - Specify the methods that is used to authenticate each request:
 - Access Key ID/Secret Key- Specify the access key ID, secret key, and tenant ID
 - Username/Password - Specify the username, password, and tenant ID
 - Username/API Access Key - Specify the username and the API Access key
 - Authentication URL Path - Specify the cloud server API URL for Cloudwatt Object Storage to authenticate the request. For example, /auth/v1.0 or /auth/v2.0.
 - Web Protocol - Specify whether to use HTTP or HTTPS.
4. Specify the hostname.
5. Specify the bucket name associated with your cloud service provider account. You can use buckets to organize your data and control access to your data, but they cannot be nested. If the bucket name does not exist, the bucket is created during initial AltaVault replication.
6. Specify the port number.
7. Enable Archiving - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
8. Enable Cloud Deduplication - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended only for Alibaba Archive, Amazon Glacier, and Azure Archive.
9. Enable Cloud CA Certificate - Optionally, specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.
10. Select the Enable Proxy check box to enable proxy server settings and specify:

- Hostname/IP address - Specify the hostname or the IP address,
- Ports - Specify the port numbers for access.
- Username - Specify the name of the user for access.
- Password - Specify the user's password.

11. Click **Apply** to apply your changes to the running configuration.

Configuring SWIFT-based storage with region-selection

1. Select yes or no to Use Keys from KMIP Server from the drop-down list. When configuring the KMIP server, you must:
 - Use the same username and password as created in KMS.
 - Upload the same certificate as downloaded from KMS after signing it.
 - Add a symmetric key (KMIP key) as the encryption key.
 - Add a secret data key (KMIP key) for the authentication fields.
2. Specify the following settings:
 - Region - Select the region from the drop-down list:
 - Username - Specify the username to authenticate each request.
 - API Access Key - Specify the API access key.
3. Specify the hostname.
4. Specify the bucket name associated with your cloud service provider account. You can use buckets to organize your data and control access to your data, but they cannot be nested. If the bucket name does not exist, the bucket is created during initial AltaVault replication.
5. Specify the port number.
6. Enable Archiving - Enable this option if you are using the AltaVault for cold storage mode. For more information about cold storage mode, see [“Deployment guidelines” on page 13](#).
7. Enable Cloud Deduplication - Enabling this option may improve deduplication rates for repetitive backup datasets, lowering cloud storage costs. Disabling this option is recommended only for Alibaba Archive, Amazon Glacier, and Azure Archive.
8. Enable Cloud CA Certificate - Optionally, specify a cloud CA certificate that will be used to validate the server certificate of cloud provider. This must be .pem extension file.
9. Select the Enable Proxy check box to enable proxy server settings and specify:
 - Hostname/IP address - Specify the hostname or the IP address.
 - Ports - Specify the port numbers for access.
 - Username - Specify the name of the user for access.
 - Password - Specify the user's password.

Customizing a private cloud

You need to contact NetApp [technical support](#) to configure a private cloud. After you configure a private cloud, the cloud appears as the cloud provider in the cloud settings page.

To customize a private cloud

1. Contact NetApp Support to convert a cloud to a private cloud.
2. After you configure a private cloud using the CLI, it appears in the Cloud Settings page and the dashboard in the Cloud Information section as the Provider. For more information on CLI, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Guide*.
3. Choose Configure > Cloud Settings.
4. Select the Cloud tab.
5. Under Cloud Provider Settings, complete the configuration as necessary. Refer to your private cloud configuration for the required authentication credentials needed to communicate with this cloud.

Using the import configuration wizard

Use the Import Configuration wizard to import a previously saved configuration into the AltaVault. The Import Configuration Wizard will fail if the AltaVault already has an encryption key set.

It is recommended to set the time zone to the AltaVault prior to uploading the configuration.

To use the import configuration wizard

1. From the management console, choose Configure > Setup Wizard.
2. Select Import Configuration in the Wizard Dashboard.
3. Select one of the following options:
 - Select Local File and click **Browse** to select a local configuration file from your computer.
 - or-
 - Select URL and specify the URL of an appliance whose configuration you want to import.
4. Leave the Import Shared Data Only check box selected to import only the following common settings:
 - Cloud settings
 - Email settings
 - Logging
 - NTP settings
 - SNMP settings
 - Statistics or Alarms settings
 - Time zone settings
 - Web and CLI preferences
 - SMB, NFS, OST, SnapMirror configuration

When you select the Import Shared Data Only check box, the following settings are not imported:

- General Security Settings
 - Static host configuration
 - Appliance licenses
 - Interface configuration, IP configuration, static routes, and virtual interfaces.
 - Radius protocol settings
 - Name server settings and domains
 - Scheduled Jobs
 - SSH server settings and public or private keys
 - Hostname, Message of the Day (MOTD), and Fully Qualified Domain Name (FQDN)
 - TACACS protocol settings
 - Telnet server settings
5. Select the Key Phrase protect the Encryption Key check box to specify a password for the encryption key. If you select this option, you must enter the same password when you import or export the encryption key.
 6. Click **Import Configuration**.
 7. Click **Exit**.

Using the export configuration wizard

To use the Export Configuration wizard

1. From the management console, choose Configure > Setup Wizard.
2. Select Export Configuration.
3. Click **Export Configuration** to download the current AltaVault configuration file `AltaVault_config_(HOSTNAME)_(DATETIME).tgz`.

If an encryption key passphrase is configured on AltaVault at the time you export the configuration, the configuration file will require this passphrase when imported to another AltaVault appliance. For more information about the encryption key passphrase, go to [“Configuring encryption” on page 38](#).

CHAPTER 4 **Configuring storage settings**

This chapter includes the following sections:

- [“Configuring cloud settings” on page 37](#)
- [“Configuring SMB” on page 39](#)
- [“Configuring NFS” on page 44](#)
- [“Configuring OST” on page 48](#)
- [“Configuring SnapMirror” on page 50](#)

Configuring cloud settings

You can specify cloud settings in the [Configure > Cloud Settings](#) page.

Before you configure cloud settings, you must configure DNS settings to access the cloud service provider host machine on the [Configure > Host Settings](#) page.

This section includes the following topics:

- [“Configuring cloud provider settings” on page 37](#)
- [“Configuring encryption” on page 38](#)
- [“Configuring replication” on page 38](#)
- [“Configuring bandwidth limits” on page 38](#)

To transition cloud credentials and the encryption key from the AltaVault to a Key Management Server (KMS), refer to the section [“Configuring KMIP” on page 97](#).

Configuring cloud provider settings

This setting enables you to access the storage and software through the Internet. For more details on cloud provider settings, see [“Using the Cloud Settings wizard” on page 22](#).

Only users who have Read-Only Replication Settings permission or Read and Write Replication Settings permission can access and configure the Cloud Settings Page.

Configuring encryption

The new datastore encryption key can be generated or imported from an existing one.

To secure the encryption key, protect it using a key passphrase. This passphrase will be used to encrypt the datastore encryption key and must be provided whenever importing this datastore encryption key, such as for disaster recovery. It is not stored within a configuration archive and must be kept in a secure location.

For more information on encryption, see [“Using the Cloud Settings wizard” on page 22](#).

Configuring replication

Replication is the process of copying data and metadata from the AltaVault to the cloud. The AltaVault replicates data to the cloud asynchronously.

Only users who have Read-Only Replication Settings permission or Read and Write Replication Settings permission can access and configure the Replication Settings Page.

To configure replication

1. Choose **Configure > Cloud Settings**.
2. Select the **Replication** tab.
3. Under **Replication Settings**, complete the configuration as described in this table.

Control	Description
Pause Replication at	Specify the time (in HH:MM:SS format) at which you want replication to pause.
Resume Replication at	Specify the time (in HH:MM:SS format) at which you want replication to resume.
Bytes pending replication alert limit	Displays an alarm if the number of bytes pending replication to the cloud exceeds the value you specify. The default value is 500 GiB.
Bytes pending replication clear limit	Specify the lower limit at which the bytes pending replication alert limit notification is cleared. The default value is 450 GiB
Suspend Replication	Click to temporarily stop replication.

4. Click **Apply** to complete your changes.

Configuring bandwidth limits

You can limit the bandwidth that the AltaVault uses to replicate data and restore data in the bandwidth limit settings page.

Only users who have Read-Only Replication Settings permission or Read and Write Replication Settings permission can access and configure the Bandwidth Limit Settings Page.

To configure bandwidth limits

1. Choose **Configure > Cloud Settings**.

2. Select the Bandwidth tab and specify:

Control	Description
Cloud Replication Interface	<p>Select a data interface to use for sending data to and restoring data from the cloud. Select the interface in the drop-down list and then specify the bandwidth limits and scheduling. You must first configure the data interfaces before they appear in the drop-down list.</p> <p>Setting the replication interface to Primary/Default is not recommended as this is the management interface for the appliance.</p>
Replication Limit Rate	Specify a rate to limit the data transmitted to the cloud storage provider in kilobits per seconds (kbps).
Restore Limit Rate	Specify a rate to limit the data restored in kilobits per second (kbps).
Enable Bandwidth Limit Scheduling	<p>Before you select this option, you must specify the replication/restore options above. Select the check box and specify:</p> <ul style="list-style-type: none"> • Start Time - the time at which the bandwidth limit should start. • End Time - the time at which the bandwidth limit should end. • Replication Limit Rate - the replication rate during the defined schedule. The bandwidth reverts to the normal replication limit rate outside the scheduled times. • Restore Limit Rate - the restore rate during the defined schedule. • Include Weekends - apply schedule to weekdays and weekends.

3. Click **Apply** to apply your changes to the running configuration.

After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports.

Configuring SMB

SMB is currently enabled in two versions: SMBv2 and SMBv3. SMBv2 is the default protocol that is used with Windows 2008 systems, and SMBv3 is the default protocol that is used with Windows 2012 systems. AltaVault supports SMB2 and SMB3. You can configure SMB access for Microsoft Windows based clients to the AltaVault in the Configure > SMB page.

Note: If you are upgrading to AVA4.2 or later releases, migration of your CIFS configuration from earlier AVA releases to AVA SMB is supported. For complete information about SMB changes associated with AltaVault software upgrades, see *AltaVault Cloud Integrated Storage Release Notes*.

When configuring SMB, you perform the following tasks:

- [“Configuring an Active Directory domain” on page 40](#)
- [“Adding SMB shares” on page 41](#)
- [“Adding users to access the share” on page 42](#)
- [“Adding SMB local users” on page 42](#)
- [“Editing multichannel settings” on page 43](#)

Configuring an Active Directory domain

If your network has an Active Directory (AD) domain, you can add the AltaVault to the domain and enable domain users to access AltaVault SMB shares. You can add the AltaVault only to one domain. Ensure that you have permissions to join appliances to the domain.

1. Choose **Configure > SMB**.

The SMB page does not appear until the Storage Optimization Service is started. If needed, choose **Maintenance > Service** and click **Start** to initialize the service.

2. Optionally, you can specify up to three preferred domain controllers. Under Preferred Domain Controllers, enter a fully qualified domain name or IPv4 address for each controller. AltaVault accesses preferred controllers in order, starting with Domain Controller 1.

If no controllers are specified, AltaVault uses DNS to discover domain controllers.

3. Click **Apply**.

4. To join the AltaVault to an AD domain, go to the Domain section and specify these settings:.

Control	Description
Domain Name	Specify the fully qualified domain name of the AD that the AltaVault will join.
Username	Specify the username of a user who has appropriate permissions to add computers to the domain.
Password	Specify the user's domain password.

5. Click **Show Advanced Settings** to display Advanced Settings to (optionally) configure the domain:.

Control	Description
Hostname	Optionally, specify the hostname that the AltaVault will use as part of the domain. AltaVault then appears as the specified hostname in the AD.

6. Click **Join Domain**. AltaVault attempts to join the AD domain.

7. After you join a domain, the Domain section of the SMB page changes to reflect the domain that the AltaVault has joined.

When you leave a domain, specify:.

Control	Description
Username	Optionally, specify the username of a user which has appropriate permissions to add computers to the domain.
Password	Optionally, specify the user's domain password.
Leave Domain	Attempt to remove AltaVault from the domain.

Reboot all client machines that were used to connect to the AltaVault to delete cached domain credentials.

Adding SMB shares

1. The SMB page does not appear until the Storage Optimization Service is started. If needed, choose Maintenance > Service and click Start to initialize the service.
2. Optionally, under Pinned Data Information, slide the indicator along the bar to select the bytes allowed for share pinning. Share pinning instructs the share to always retain data on AltaVault locally without fetching it from the cloud.

Note that the maximum number of bytes selected for share pinning is the total available to all protocols (SMB, NFS, and OST) and is cumulative on the AltaVault. For example, if you are configuring SMB shares and set the pinning percentage to 20 percent, that much of the local cache is available for share pinning across NFS, SMB and OST.

You can increase or decrease the percentage of cache available for share pinning at any time; however, once you establish pinning and begin consuming cache space for pinned data, you cannot decrease the percentage below the total number of bytes used by pinned shares across all protocols (SMB, NFS, and OST). If the total of bytes used by pinned shares exceeds the percentage of cache available for pinning, writes to pinned shares fail.

3. To add an SMB share, complete the configuration as described in this table and click Add.

Control	Description
Add SMB Share	Displays the controls to add a new SMB share.
Share Name	Specify the name of the share. Note: AltaVault does not support having two shares with the same name.
Pin Share	Optionally, enable data pinning on the share. Select Yes or No from the drop-down list to specify whether the SMB share should be pinned. Share pinning enables the share to always contain data that is available to the AltaVault locally without fetching it from the cloud. You can pin SMB shares only at the time of share creation. Existing unpinned shares cannot be pinned. Once a share is pinned, unpinning of that share can be performed via CLI and requires optimization service to be offline. Unpinning a share can be a time-consuming operation. Unpinning a share does not result in erasing the previously pinned data. After unpinning, the previously pinned data becomes available for eviction. You cannot remove a pinned share if it contains data.
Early Eviction	Specify whether or not data from this share should be assigned a higher priority for eviction. If you select yes, data written to this share is eligible for eviction earlier than other data.
Disable Dedupe	Specify whether or not data written to this share should be checked for duplication. If you select yes, then the AltaVault will not perform duplication checks on data written to the share.
Disable Compression	Specify whether or not data written to this share should be compressed. Select yes if your data set is already in a compressed format and will not benefit from further compression attempts.
Local Path	Specify the internal pathname on the AltaVault to which this SMB share writes data. Note: AltaVault does not support having two shares with the same local path. Do not create two shares with the same local path. Additionally, nesting shares (local path of a share is part of the local path of another share) is not allowed.
Comment	Enter a comment about the share. You can use supported Unicode characters, underscores, hyphens, and spaces.

Control	Description
Allow Everyone Access	Enable global access to the SMB share.
Add Share	Adds the SMB share to the AltaVault.
Remove Selected	Deletes the selected SMB share.

The share you configured appears in the list of shares on the page along with the option to add access control entries for domain and local users.

4. Optionally, to edit share information, extend the share name to edit the configuration, and click **Apply**.

Adding users to access the share

SMB share security and access can be administered for Active Directory users and groups. If AltaVault is not within a domain, use a local user account to gain access to a share.

1. To add a user to access the share that you created, expand the share name to complete the configuration as described in the following table.

Note: A local SMB user must first be created as described in the [“Adding SMB local users” on page 42](#) before you can add the access control entry for that user to a share.

Control	Description
Add a User	Displays the controls to add access control entries to a share ACL.
User	Enter either a domain user or group name, or local user name. Domain user or group names must include the domain. For example, DOMAIN\user1.
Access	Select one of the following options from the drop-down list: Full Control - Allows the user read, write, delete, update ACL, and ownership access to the share. Write - Allows the user read and write access to the share. Read - Allows the user read-only access to the share. Deny - Denies the user access the share.
Remove Selected	Deletes the selected user from the SMB server.
Add User	Adds the SMB user.

2. Optionally, to edit user permissions after adding a user to a share, extend the user name to edit the configuration, and click **Apply Changes**.

Adding SMB local users

You can configure AltaVault with a set of local users for access to SMB shares.

1. To add local users to access the share that you created, complete the configuration as described in this table.

Control	Description
Add SMB User	Displays the controls to add a user to the SMB share.
User Name	Specify the user name of a local user to access the SMB share. The user name is case-insensitive.
Password	Specify the password for the new user.
Password Confirm	Re-enter the new password for the new user.
Admin	Select one of the following options from the drop-down list: <ul style="list-style-type: none"> • Yes - Provides Administrator privileges to user • No - Disables Administrator privileges to user
Account	Select one of the following options from the drop-down list: <ul style="list-style-type: none"> • Enabled - Enables local user account for accessing SMB share • Disabled - Disables local user account from accessing SMB share
Remove Selected	Removes local SMB user configuration.
Add	Adds local SMB user.

2. Optionally, to edit local user information, extend the user name to edit the configuration and click **Apply**.

Editing multichannel settings

SMB multichannel is a feature that enables SMB3 clients to establish multiple channels for SMB sessions. '

1. Multichannel support is enabled for all interfaces by default. To disable an interface (e0a, e0b, e0c), select the checkbox for the interface and click **Disable**.

Troubleshooting SMB

Use the following table to help resolve SMB issues.

Symptom	Description
Naming conflicts with SMB share or user names.	AltaVault does not support having two shares or users with the same name. Select different names. In AltaVault 4.4 and later releases, share and user names are case insensitive. Upgrading from earlier releases can result in naming conflicts. For more information on resolving conflicts, refer to the AltaVault release notes.
Shares cannot be mapped and are marked invalid.	AltaVault 4.4.1 does not support nested share paths. Upgrading from earlier releases can result in path conflicts resulting in shares not being accessible. For more information on resolving conflicts, refer to the AltaVault release notes.

Configuring NFS

You can configure Network File System (NFS) for Unix and Linux based clients in the Configure > NFS page. Before you configure NFS, choose Maintenance > Service and click **Stop** to stop the Storage Optimization Service.

Kerberos authentication is optional with NFS and works only with NFSv4 exports. If you are not using Kerberos, AltaVault does not use any other means of authentication for NFSv4 exports.

This section includes the following topics:

- [“Before you begin” on page 44](#)
- [“Configuration tasks” on page 44](#)
- [“Editing an NFS configuration” on page 46](#)
- [“Troubleshooting NFS” on page 47](#)

Before you begin

For NFS configurations using Kerberos authentication, confirm the following:

- The hosts file on the Kerberos Key Distribution Center (KDC) and NFS client has been modified to include the AltaVault. Additionally, the hosts file on AltaVault must include the IP address and host name for the KDC and client.
- On the KDC, the principals “host/...” and “nfs/...” exist for AltaVault and client, and you have run `ktadd` to create the keytab file for them.

Configuration tasks

You can configure NFS on the Configure > NFS page.

To configure NFS protocol

1. The NFS page does not appear until the Storage Optimization Service is started. If needed, choose Maintenance > Service and click Start to initialize the service.
2. Choose Configure > NFS.
3. Optionally, under Pinned Data Information, slide the indicator along the bar to select the maximum bytes allowed for share pinning. Share pinning instructs the share to always contain data that is available to the AltaVault locally without fetching it from the cloud.

Note that the maximum number of bytes selected for share pinning is the total available to all protocols (SMB, NFS, and OST) and is cumulative on the AltaVault. For example, if you are configuring SMB shares and set the pinning percentage to 20 percent, that much of the local cache is available for share pinning across NFS, SMB and OST.

You can increase or decrease the percentage of cache available for share pinning at any time; however, once you establish pinning and begin consuming cache space for pinned data, you cannot decrease the percentage below the total number of bytes used by pinned shares across all protocols (SMB, NFS, and OST). If the total of bytes used by pinned shares exceeds the percentage of cache available for pinning, writes to pinned shares fail.

4. If using Kerberos, upload the Kerberos keytab file (krb5.keytab), then upload a valid Kerberos configuration file (krb5.conf).

The keytab file is an encrypted, local, on-disk copy of the host's key. The configuration file contains Kerberos configuration information, including the locations of KDCs and administration servers for the Kerberos realms, default parameters for the current realm and for Kerberos applications, and mappings of host names onto Kerberos realms.

Note: Share access can be lost if Kerberos is enabled without uploading the krb5.keytab and krb5.conf files.

5. If using Kerberos, choose Configure > Host Settings > Hosts and add information for the following hosts: KDC, NFS client(s), and AltaVault.

6. Under Add an Export, complete the configuration as described in this table:

Control	Description
Add an Export	Displays the controls to export an NFS share.
Name	Specify the name of the export share.
Export as NFSv4	Specify the type of NFS export. If you select yes, the export will be configured as NFSv4 export. If you select no, the export will be configured as NFSv3 export.
Kerberos Authentication	Kerberos authentication works only with NFSv4 exports. It is optional. If you are not using Kerberos, AltaVault does not use any other means of authentication for NFSv4 exports.
Pin Export	<p>Optionally, enable data pinning on the share. Select Yes or No from the drop-down list to specify whether the NFS export should be pinned. Share pinning enables the share to always contain data that is available to the AltaVault locally without fetching it from the cloud. You can pin NFS exports only at the time of share creation. Existing unpinned shares cannot be pinned.</p> <p>Once a share is pinned, unpinning of that share can be performed via CLI and requires optimization service to be offline. Unpinning a share can be a time-consuming operation. Unpinning a share does not result in erasing the previously pinned data. After unpinning, the previously pinned data becomes available for eviction.</p> <p>You cannot change this option after the NFS export is created.</p>
Early Eviction	<p>Specify whether or not data from this share should be assigned a higher priority for eviction.</p> <p>If you select yes, data written to this share is eligible for eviction earlier than other data.</p>
Disable Dedupe	Specify whether or not data written to this share should be checked for de-duplication. If you select yes, then the AltaVault will not perform duplication checks on data written to the share.
Disable Compression	<p>Specify whether or not data written to this share should be compressed.</p> <p>Select yes if your data set is already in a compressed format and will not benefit from further compression attempts.</p>
Local Path	Specify the internal pathname on the AltaVault to which this share writes data.
Comment	Enter a comment about the NFS share. You use only alphanumeric characters, underscores, hyphens, and spaces in this field.

Control	Description
Export Asynchronously	<p>Select the check box to export the NFS share asynchronously. Click the icon for the following information:</p> <p>Exporting NFS asynchronously forces the server to drop all fsync requests from the client. It is required to obtain good performance with NFS clients that issue frequent NFS COMMIT operations, which might degrade the AltaVault performance significantly. Many UNIX clients often execute NFS COMMIT operations when low on memory. To understand the circumstances that cause this behavior and to detect and prevent it, contact your client operating system vendor. The AltaVault automatically synchronizes any file that is idle for a configurable amount of time. The default value is 10 seconds. Although there is a window of time (after the server responds with success for an fsync request, and before the data is written to disk), this window is small and performance benefits are large. NetApp recommends exporting NFS asynchronously.</p>
Allow Specified Clients	<p>Specify which clients can connect to the NFS share.</p> <p>To limit access, specify the client's IP address and subnet mask.</p> <p>By default, all clients can access the share, until the first client is entered. To revert to full access after adding a client, specify 0.0.0.0/0 in the Client IP/Network field IP/Network field.</p>
Allow All Clients	<p>Enables all clients connected to the AltaVault system to access the NFS share.</p> <p>WARNING: Enabling all clients to access the NFS share is not recommended.</p>
Add	Adds the export path and client IP address to the AltaVault NFS server.
Remove Selected	Select the check box next to the name and click Remove Selected .

The share you configure and its parameters appear in the list of shares on the page.

- Click **Add** to apply your changes to the running configuration.

Editing an NFS configuration

To edit an existing configuration

- Choose Configure > NFS and click the share name at the bottom of the page.
- Select the NFS share name and specify:

Control	Description
Edit Export	Select tab to edit the exported NFS share.
Local Path	Change the export file pathname, which starts with a forward slash (/).
Comment	Specify or change the comment about the NFS share.
Export as NFSv4	Not available for editing.
Kerberos Authentication	Available for editing only when Export as NFSv4 is selected.
Pinned	Not available for editing.
Early Eviction	<p>Select yes or no from the drop-down list to specify whether or not data from this share should be assigned a higher priority for eviction.</p> <p>If you select yes, data written to this share is eligible for eviction earlier than other data. If you select no, data written to this share is evicted using the default method.</p>

Control	Description
Disable Dedupe	Specify whether or not data written to this share should be checked for de-duplication. If you select yes, then the AltaVault will not perform duplication checks on data written to the share.
Disable Compression	Specify whether or not data written to this share should be compressed. Select yes if your data set is already in a compressed format and will not benefit from further compression attempts.
Export Asynchronously	Select the check box to export the NFS share asynchronously. Click the icon for the following information: Exporting NFS asynchronously forces the server to drop all fsync requests from the client. This is a feature of the NFS protocol. It is required to obtain good performance with NFS clients that issue frequent NFS COMMIT operations, which might degrade the AltaVault performance significantly. Many UNIX clients often execute NFS COMMIT operations when low on memory. To understand the circumstances that cause this behavior and to detect and prevent it, contact your client operating system vendor. The AltaVault automatically synchronizes any file that is idle for a configurable amount of time. The default value is 10 seconds. Although there is a window of time (after the server responds with success for an fsync request and before the data is written to disk), this window is small and performance benefits are large. NetApp recommends exporting NFS asynchronously.
Allow All Clients	Enables all clients connected to the AltaVault system to access the NFS share.
Allow Specified Clients	Enables only the clients that you specify to connect to the AltaVault system to access the NFS share. If you select this option, you must specify the client's IP address and subnet mask in the Client IP/Network text field below it. To enable all clients to access the NFS share, specify 0.0.0.0/0 in the Client IP/Network field.
Mount Commands	Select this tab to display the Linux and UNIX NFS mount commands. You configure the mount commands through the command-line. These commands are for your reference. If the AltaVault is a secondary appliance, the mount commands enable only read permissions and not write permissions.

3. Click **Apply** to apply your changes to the running configuration.

Troubleshooting NFS

Use the following table to help resolve NFS issues.

Symptom	Description
User attempts to map an NFS share fail; users are unable to connect to a share after a client or AltaVault reboot	Certain services, such as NFSv3, rely on RPC to assign a port number to services. Reboots of clients or AltaVault can cause a port re-negotiation, which is expected and normal for TCP/IP and UDP protocols. Your firewall must be configured to allow ports or the connection can be denied. In AltaVault, the NFS daemon listens on port 2049, lockd is bound to port 4001, and mountd is bound to port 32767. Check your firewall configuration and update access policies as necessary.

Configuring OST

OpenStorage Technology (OST) is a proprietary protocol created by Veritas for ingesting backup data streams to (third-party) disk-like storage devices. OST is implemented as a plug-in (shared object/DLL) running in NetBackup media server process address space and streaming data to the OST server running on the AltaVault.

You can perform the following tasks:

- [“Configuring OST shares” on page 48](#)
- [“Adding an OST user to access the share” on page 49](#)
- [“Editing OST user information” on page 50](#)

For information on configuring up the AltaVault OST Plug-in for communication with AltaVault, see the *NetApp AltaVault OST Plug-in Deployment Guide*.

Configuring OST shares

1. The OST page does not appear until the Storage Optimization Service is started. If needed, choose Maintenance > Service and click Start to initialize the service
2. Optionally, under Pinned Data Information, slide the indicator along the bar to select the maximum bytes allowed for share pinning. Share pinning instructs the share to always retain data locally on the AltaVault without fetching it from the cloud.

Note that the maximum number of bytes selected for share pinning is the total available to all protocols (SMB, NFS, and OST) and is cumulative on the AltaVault. For example, if you are configuring SMB shares and set the pinning percentage to 20 percent, that much of the local cache is available for share pinning across NFS, SMB and OST.

You can increase or decrease the percentage of cache available for share pinning at any time; however, once you establish pinning and begin consuming cache space for pinned data, you cannot decrease the percentage below the total number of bytes used by pinned shares across all protocols (SMB, NFS, and OST). If the total of bytes used by pinned shares exceeds the percentage of cache available for pinning, writes to pinned shares fail.

3. Click **Apply** to apply your changes.
4. To add an OST share, click **Add OST Share** and specify:.

Control	Description
Share Name	Specify the name of the share.
Type	Select regular or cloud. Regular shares treat incoming data by AltaVault the same as traditional SMB shares or NFS exports by writing the data on cache, and replicating the data to the cloud. Cloud shares are used to create an optimized duplicate of data in regular shares replicated in the cloud that are managed via NetBackup storage lifecycle policies (SLP).

Control	Description
Pin Share	<p>Optionally, enable data pinning on the share. Select Yes or No from the drop-down list to specify whether the OST share should be pinned. Share pinning enables the share to always contain data that is available to the AltaVault locally without fetching it from the cloud. You can pin OST shares only at the time of share creation. Existing unpinned shares cannot be pinned.</p> <p>Once a share is pinned, unpinning of that share can be performed via CLI and requires optimization service to be offline. Unpinning a share can be a time-consuming operation. Unpinning a share does not result in erasing the previously pinned data. After unpinning, the previously pinned data becomes available for eviction.</p> <p>You cannot remove a pinned share if it contains data.</p>
Early Eviction	<p>Specify whether or not data from this share should be assigned a higher priority for eviction.</p> <p>If you select yes, data written to this share is eligible for eviction earlier than other data.</p>
Disable Dedupe	Specify whether or not data written to this share should be checked for duplication. If you select yes, then the AltaVault will not perform duplication on data written to the share.
Disable Compression	<p>Specify whether or not data written to this share should be compressed.</p> <p>Select yes if your data set is already in a compressed format and will not benefit from further compression attempts.</p>
Add Share	Adds the OST share to the AltaVault.

- Optionally, to remove an empty OST share, select the OST share from the table and click **Remove Selected**
- Optionally, to enable SSL communication between the AltaVault OST Plug-in and the AltaVault, select the checkbox, **Enable SSL**, in the Global OST Settings section, and click **Apply**. AltaVault will communicate with the AltaVault OST Plug-in using secured port 8085.

Adding an OST user to access the share

- To add a user to access the share that you created, select the share and specify:

Note: OST shares must have an associated user to be used by NetBackup storage server. Multiple users per OST share are allowed.

Control	Description
Add OST User	Displays the controls to add a user to the OST share.
User Name	Type the user name that you would use for authenticating the share from NetBackup.
Password	Specify the new password for the user.
Password confirm	Re-enter the password.
Add User	Adds the OST user.

Editing OST user information

1. Extend the user name to complete the configuration as described in this table.

Control	Description
Change Password	Select the check box to change the password.
Password	Specify the new password.
Password Confirm	Re-enter the new password.
Account	Specify: <ul style="list-style-type: none"> • Enabled - Enables local user account • Disabled - Disables local user account
Apply	Applies the changes to the OST share users.

Configuring SnapMirror

AltaVault supports backup and restore operations for ONTAP FlexVol volumes using the SnapMirror service. Backup relationships are created and managed from ONTAP using ONTAP commands or SnapCenter software. SnapMirror support is available on AltaVault physical and virtual appliance models. For more information about AltaVault operation with ONTAP, see [NetApp Technical Report 4624: Solution Deployment: AltaVault with ONTAP Using SnapMirror](#).

This section includes the following topics

- “Enabling SnapMirror service” on page 50
- “Monitoring and deleting SnapMirror shares and Snapshots on AltaVault” on page 51
- “Enabling long-term retention” on page 52
- “Enabling SnapCenter access” on page 53

Enabling SnapMirror service

To enable SnapMirror service

1. Choose **Configure > SnapMirror** in the Management Console.
2. Under SnapMirror Service, click **Enable**.
3. If the “Service restart required” prompt appears, click the **Restart** button that becomes enabled in the upper right portion of the AltaVault Management Console.
4. Under Whitelist IP, click **Add Whitelist IP**.
The Whitelist specifies which addresses are authorized to communicate with AltaVault.
5. Enter the IP addresses of ONTAP intercluster LIFs from which AltaVault will accept connections for backup and restore operations, and click **Add**.

The list of authorized IP addresses must be populated prior to initiating a connection from the ONTAP system or the connection will be rejected.

To remove an IP address, select the IP Address and click **Remove Selected**. Removing an IP address from the whitelist disables access to the AltaVault from that IP address.

Note: If you are using SnapCenter to manage backups, SnapCenter automatically creates the whitelist of approved IP addresses when you initiate a backup from SnapCenter. In this case, there is no need to create the IP whitelist.

To disable SnapMirror service

1. Choose Configure > SnapMirror in the Management Console.

2. Under SnapMirror Service, click **Disable**.

When SnapMirror service is disabled, the shares and Snapshots that exist on AltaVault are not deleted and are kept intact. Snapshots are not accessible while service is disabled. Snapshots can be restored only when SnapMirror service is enabled.

3. If the Service restart required prompt appears, click the Restart button that becomes enabled in the upper right of the console.

Monitoring and deleting SnapMirror shares and Snapshots on AltaVault

A SnapMirror share is created automatically when the SnapMirror relationship with the AltaVault is created in ONTAP or in SnapCenter. Based on SnapMirror policies, Snapshot copies of ONTAP volumes are backed up to the associated AltaVault share. AltaVault provides global deduplication on all Snapshot backup streams prior to replication to the cloud.

Snapshots backed up to AltaVault shares are read-only copies and can only be restored back to ONTAP using ONTAP commands or SnapCenter.

To view SnapMirror shares and Snapshots on AltaVault

1. Under SnapMirror Shares, review the information fields associated with a share:

Field	Description
Name	Specifies the name of the share created in ONTAP using the ONTAP CLI or by using SnapCenter software. When the ONTAP administrator creates a SnapMirror relationship with AltaVault, a share is automatically created in AltaVault. Each share is associated with one ONTAP FlexVol volume. AltaVault supports up to 500 SnapMirror shares. To view a list of Snapshots associated with each share, select a share name.
Peer Path	Identifies the source volume in ONTAP that is being backed up to AltaVault.
UUID	Lists the unique identifier associated with each SnapMirror share. The UUID value is generated by AltaVault.
Size	Specifies the size of the SnapMirror share. The size can grow or shrink as Snapshots are backed up or deleted from the share. Shares on the AltaVault have no size limitation but are bound by the AltaVault's cache capacity. The size of source volume, change rate, and number of Snapshots will impact the number and size of SnapMirror shares on the AltaVault.

- To view the list of Snapshots for a share, select a share and review the Snapshot information:

Field	Description
Name	List Snapshot for a share. Snapshot backups can be triggered in ONTAP through SnapMirror policies, by explicitly running the ONTAP update command, or through SnapCenter software.
UUID	Lists the unique identifier associated with each Snapshot copy. The UUID value is generated by ONTAP.
Created	Displays the date and time when the Snapshot was created in ONTAP.
Size	Specifies the size of a Snapshot. During the lifetime of a share, there is only one baseline Snapshot. Any Snapshot after the baseline is always incremental. Baseline transfer can take a long time to complete depending on the size of the Snapshot. During incremental Snapshot backups, only the changed blocks between two Snapshots are transferred.
Status	Identifies the status of Snapshot replication to the cloud. Replication status can be either Completed or Pending.

To delete SnapMirror shares and Snapshots on AltaVault

- To remove a share or Snapshot, select the share or Snapshot and click **Remove Selected**.

Snapshots can be deleted on the AltaVault through ONTAP SnapMirror policies or SnapCenter policies, or by manual deletion on AltaVault. When a share is deleted, Snapshots belonging to that share are also deleted. AltaVault reclamation will recover the space occupied by the deleted Snapshot or Share asynchronously, and Share size may not immediately reflect available space from the operation.

Note: You cannot delete the latest Snapshot. Also, a Snapshot cannot be deleted while a restore is in progress.

Enabling long-term retention

AltaVault supports up to 500 SnapMirror shares in one of two modes: short-term retention (default) or long-term retention. For short-term retention, each share supports up to 251 Snapshots, and Snapshot retention is dependent upon the retention policy set up in ONTAP. For example, suppose a share has a two-tier retention policy supporting 50 hourly and 100 daily Snapshots. In this case, when the count of hourly Snapshots exceeds 50 or the daily count exceeds 100, the oldest snapshot of the respective tier is deleted.

For long-term retention, each share supports up to a maximum of 3700 Snapshots, which is equivalent to 10 years worth of daily Snapshots. Long-term retention allows AltaVault to continue storing Snapshots until it reaches the maximum. If a share exceeds 3700 Snapshots, AltaVault begins deleting the oldest Snapshot copies to make room for new ones.

When long-term retention is turned off (disabled), AltaVault reverts to using the retention policy set up in ONTAP, which supports a maximum 251 Snapshots per share. If there are large numbers of Snapshots (more than 251) when long-term retention is turned off, the number of snapshots will be reduced to match the count set in the retention policy.

The retention method used for Snapshot retention applies to all SnapMirror shares created on the AltaVault.

To enable long-term retention mode

- Under Long Term Retention, click **Enable**.

To disable long-term retention, click **Disable**.

Important: If SnapCenter is being used to manage backups, long-term retention will be enabled or disabled from SnapCenter. Do not disable or enable long-term retention on the AltaVault appliance explicitly while SnapCenter is managing backups.

Enabling SnapCenter access

SnapCenter can be used to back up and delete Snapshots, and to perform single file restores. If you are using SnapCenter to manage backups, you must enable SnapCenter access on AltaVault.

Additionally, before you can use SnapCenter to manage backups on AltaVault, you must configure a role-based account on AltaVault for SnapCenter administrator access. This account must have the read/write permissions for the following user roles: General, Replication, Storage.

To create a role-based user account for SnapCenter on AltaVault

1. Choose **Configure > User permissions** in the Management Console.
2. Under role-based accounts, select **Add a New User**.
3. Enter an account name and password, and check **Enable Account**.
4. Select **Read/Write** permission for the following roles: General Settings, Replication Settings, Storage Settings.
5. Click **Add**.

To enable SnapCenter access to AltaVault

1. Under SnapCenter Access, click **Enable**.
To disable SnapCenter access, click **Disable**.

CHAPTER 5 **Modifying networking settings**

This chapter includes the following sections:

- [“Modifying general host settings” on page 55](#)
- [“Modifying management interfaces” on page 57](#)
- [“Modifying data interfaces” on page 58](#)
- [“Modifying virtual interfaces \(VIFs\)” on page 60](#)
- [“Modifying VLANs” on page 61](#)

Modifying general host settings

You can view and modify general host settings in the Configure > Host Settings page. Use the following groups of controls on this page only if modifications or additional configuration is required:

- **Name** - Modify the hostname.
- **DNS Settings** - NetApp recommends that you use DNS resolution.
- **Hosts** - If you do not use DNS resolution, or if the host does not have a DNS entry, you can assign a host-IP address resolution map.
- **Web/FTP Proxy** - Configure proxy addresses for Web or FTP proxy access from the AltaVault. The proxy settings do not affect cloud connections or AutoSupport uploads originating from the AltaVault. Web proxy settings are used for uploading system, process, or TCP dumps to NetApp support.

To view general host settings

- Choose Configure > Host Settings.

To change the hostname

1. Choose Configure > Host Settings.
2. Under Name, modify the value in the Hostname field.
3. Click **Apply** to apply your changes to the running configuration.

To specify DNS settings

1. Choose **Configure > Host Settings**.
2. Under **DNS Settings**, complete the configuration as described in this table.

Control	Description
Primary DNS Server	Specify the IP address for the primary name server.
Secondary DNS Server	Optionally, specify the IP address for the secondary name server.
Tertiary DNS Server	Optionally, specify the IP address for the tertiary name server.
DNS Domain List	Specify an ordered list of domain names. If you specify domains, the system automatically finds the appropriate domain for each of the hosts that you specify in the system.

3. Click **Apply** to apply your changes to the running configuration.

To add a new host

1. Choose **Configure > Host Settings**.
2. Under **Hosts**, complete the configuration as described in this table.

Control	Description
Add a New Host	Displays the controls for adding a new host.
IP Address	Specify the IP address for the host.
Hostname	Specify a hostname.
Add	Adds the host.
Remove Selected	Select the check box next to the name and click Remove Selected .

3. Click **Apply** to apply your changes to the running configuration.

To set a Web/FTP proxy

1. Choose **Configure > Host Settings**.
2. Under **Web/FTP Proxy**, complete the configuration as described in this table.

Control	Description
Enable Web Proxy	Enables the appliance to use a Web proxy to contact NetApp. Web proxy access is disabled by default.
Web/FTP Proxy	Specify the IP address for the Web or FTP proxy.

Control	Description
Port	Optionally, specify the port for the Web or FTP proxy. The default port is 1080.
Enable Authentication	<p>Optionally, select to require user credentials for use with Web or FTP proxy traffic. Specify the following settings to authenticate the users:</p> <ul style="list-style-type: none"> • User Name - Specify a username. • Password - Specify a password. • Authentication Type - Select an authentication method from the drop-down list: <ul style="list-style-type: none"> – Basic - Authenticates user credentials by requesting a valid username and password. This is the default setting. – NTLM - Authenticates user credentials based on an authentication challenge and response. – Digest - Provides the same functionality as Basic authentication; however, Digest authentication improves security because the system sends the user credentials across the network as a Message Digest 5 (MD5) hash.

3. Click **Apply** to apply your changes to the running configuration.

The proxy settings do not affect cloud connections originating from the AltaVault.

Modifying management interfaces

You can view and modify settings for the appliance interfaces in the Management Interfaces page. Use the following groups of controls on this page only if you require modifications or additional configuration:

- Primary Interface - The primary interface is the interface used to manage the device. It is the interface utilized to get to the Management Console and command-line interface (CLI). This is also the default port used for replication if no other interface is set up for replication traffic as described in [“Configuring bandwidth limits” on page 38](#).
- Main IPv4 Routing Table - Displays a summary of the main routing table for the appliance. You can add static routes that might be required for some subnets.

To display and modify the configuration for management interfaces

1. Choose **Configure > Management Interfaces**.
2. Under Primary Interface, complete the configuration as described in this table.

Control	Description
Enable Primary Interface	<p>Enables a primary interface for the AltaVault.</p> <p>If only one interface is set up, both appliance management and replication traffic will traverse it.</p>
Obtain IPv4 Address Automatically	<p>Automatically obtain an IPv4 address from a DHCP server.</p> <ul style="list-style-type: none"> • Enable IPv4 Dynamic DNS - Select this option to enable IPv4 dynamic DNS on the primary interface.

Control	Description
Specify IPv4 Address Manually	Specify this option to set a static IP address. <ul style="list-style-type: none"> IPv4 Address - Specify an IPv4 address. IPv4 Subnet Mask - Specify an IPv4 subnet mask. Default IPv4 Gateway - Specify the default primary gateway IPv4 address. The primary gateway must be in the same network as the primary interface.
MTU	Specify the Maximum Transmission Unit (MTU) value. The default value is 1500.

- Click **Apply** to apply your changes to the running configuration.

To modify main IPv4 routing table

- Choose **Configure > Management Interfaces**.
- Under **Main IPv4 Routing Table**, complete the configuration as described in following table.

Control	Description
Add a New Route	Displays the controls for adding a new route.
Destination IPv4 Address	Specify the destination IPv4 address for the appliance.
IPv4 Subnet Mask	Specify the IPv4 subnet mask.
Gateway IPv4 Address	Specify the IPv4 address for the gateway.
Interface	Select the interface from the drop-down list.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click Remove Selected .

- Click **Apply**.

You can verify whether changes have had the desired effect by reviewing related reports.

Modifying data interfaces

You can view and modify settings for the data interfaces in the **Configure > Data Interfaces** page.

To display and modify the configuration for data interfaces

- Choose **Configure > Data Interfaces**.

- Under Physical Interface, click the arrow next to the name of the interface and complete the configuration as described in this table.

Control	Description
Enable Data Interface	Select the check box to enable the data interface and specify the following settings: <ul style="list-style-type: none"> IPv4 Address - Specify an IPv4 address. IPv4 Subnet Mask - Specify a subnet mask. IPv4 Gateway - Specify the gateway IP address. MTU - Specify the MTU value. The default value is 1500.

If a physical interface is a member of a virtual interface, it is owned by the virtual interface and you can only enable it by editing the virtual interface.

- Under Routing Table for <physical interface>, you can configure static routes if your network requires them. You can add or remove routes from the table as described in following table.

Control	Description
Add a New Route	Displays the controls for adding a new route.
Destination IP Address	Specify the destination IP address for the appliance.
Subnet Mask	Specify the subnet mask.
Gateway IP Address	Specify the IP address for the gateway.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click Remove Selected .

- Under Virtual Interface, click the arrow next to the name of the interface to enable and configure the VIF networking configuration. Create virtual interfaces from the Configure > VIFs page.

Control	Description
Virtual Interface	Displays the controls to add a virtual network interface.
IP Configuration	Displays the IP address of the network interface.
Enabled	Displays the state of the interface.
Members	Specify a comma-separated list of the data interfaces that are members of this VIF.
Enable Virtual Interface	Select this check box to enable the data interface and specify the following settings: <ul style="list-style-type: none"> IPv4 Address - Specify an IPv4 address. IPv4 Subnet Mask - Specify a subnet mask. IPv4 Gateway - Specify the gateway IP address. MTU - Specify the MTU value. The default value is 1500.

- Under VLAN Interface, click the arrow next to the name of the interface to complete the configuration. Create virtual interfaces from the Configure > VLANs page.

Control	Description
IP Configuration	Displays the IP address of the network interface.
Enabled	Displays the state of the interface.
Enable Interface	Select the check box to enable the data interface and specify the following settings: <ul style="list-style-type: none"> IPv4 Address - Specify an IPv4 address. IPv4 Subnet Mask - Specify a subnet mask. IPv4 Gateway - Specify the gateway IP address. MTU - Specify the MTU value. The default value is 1500.

If an interface is a member of a virtual interface, you can only enable it by editing the virtual interface.

- Under Routing Table for <VLAN interface>, you can configure static routes if your network requires them. You can add or remove routes from the table as described in following table.

Control	Description
Add a New Route	Displays the controls for adding a new route.
Destination IP Address	Specify the destination IP address for the appliance.
Subnet Mask	Specify the subnet mask.
Gateway IP Address	Specify the IP address for the gateway. The gateway must be in the same network as the network interface you are configuring.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click Remove Selected .

- Click **Apply** to save your changes.

Modifying virtual interfaces (VIFs)

You can view, add and modify virtual interfaces (VIFs) in the Configure > VIFs page. A VIF is a logical bonded interface created by aggregating multiple physical interfaces.

To display, add, and modify the VIF configuration

- Choose Configure > VIFs.
- Click Add a Virtual Interface and complete the configuration as described in this table.

Control	Description
Enable VIF	Enables VIF feature.
Virtual Interface Name	Specify a name for the virtual interface.
Member Interfaces	Specify a comma-separated list of the data interfaces that are members of this VIF.

Control	Description
Mode	<p>Select one of the following modes:</p> <ul style="list-style-type: none"> • 802.3ad - Enables IEEE 802.3ad Dynamic Link Aggregation. This mode enables you to bundle or aggregate multiple physical interfaces into a single VIF and enables load balancing between the interfaces. • Transmit/Receive Load Balance - Provides both transmit and receive load balancing. • Transmit Load Balance - Provides adaptive-transmit load balancing. The AltaVault distributes the outgoing traffic based on the current load on each member interface. One of the member interfaces of the VIF receives the incoming traffic.
Monitoring interval	Specifies the Media Independent Interface (MII) link monitoring frequency in milliseconds. This determines how often the link state of each slave is inspected for link failures. A value of zero disables MII link monitoring. A value of 50 is a good starting point.
Add	Adds the VIF to your configuration.
Remove Selected	Select the check box next to the existing VIF to remove, and click Remove Selected.

3. Choose Maintenance > Service and click **Restart** for the configuration changes to take effect.

Modifying VLANs

VLAN tagging enables AltaVault to direct network packets to specific virtual local area networks (VLANs) in order to segment data traffic.

To display, add, or modify a VLAN configuration

1. Stop the Storage Optimization Service before adding or removing a VLAN. If needed, choose Maintenance > Service and click **Stop** to terminate the service.
2. Choose Configure > VLANs.
3. Click Add a VLAN Interface and complete the configuration as described in this table.

Control	Description
VLAN ID	Specify the VLAN tag identifier. This can be an integer from 2 to 4094. AltaVault supports up to 16 VLAN tags per interface.
Interface Type	Select from Data Interface or Virtual Interface (VIF).
Data Interfaces	Select from the drop-down list.
VIFs	Select from the drop-down list.
Add	Adds the VLAN interface to your configuration.
Remove Selected	<p>Select the check box next to the name and click Remove Selected.</p> <p>Note: A restart of AltaVault is required before performing any further networking changes.</p>

4. Choose Maintenance > Service and click **Start** for the changes to take effect.

CHAPTER 6 **Configuring system administrator settings**

This chapter includes the following sections:

- [“Setting announcements” on page 63](#)
- [“Configuring alarm settings” on page 63](#)
- [“Configuring date and time” on page 69](#)
- [“Configuring SNMP basic settings” on page 71](#)
- [“Configuring email settings” on page 78](#)
- [“Configuring log settings” on page 78](#)

Setting announcements

You can create or modify a login message to be displayed in the Management Console Login page. You can also post a message of the day to appears in the Home page and when you first log in to the CLI.

To set an announcement

1. Choose **Configure > Announcements**.
2. Use the controls to complete the configuration as described in this table.

Control	Description
Login Message	Type a message in the text box to appear on the Login page.
MOTD	Type a message in the text box to appear on the Home page as the message of the day.

3. Click **Apply** to view the message before saving.

Configuring alarm settings

You can set alarms in **Configure > Alarms** page. Enabling alarms is optional.

AltaVault uses hierarchical alarms. The system groups certain alarms into top-level categories, such as the Link Duplex Alarm. When an alarm triggers, its parent expands to provide more information. As an example, the Disk Full top-level parent alarm aggregates over multiple partitions. If a specific partition is full, the Disk Full parent alarm triggers, and the System Status report displays more information regarding which partition caused the alarm to trigger.

Disabling a parent alarm disables its children. You can enable a parent alarm and disable any of its child alarms. You cannot enable a child alarm without first enabling its parent.

When the top-level parent box is unchecked, the Alarm category is disabled, and grayed out. All the child alarms are disabled. When the top-level parent box is checked, the Alarm category is enabled, and no longer grayed out. All the child alarms will be enabled, but can be disabled individually if needed.

The child alarm of a disabled parent appears on the System Status report with a suppressed status. Disabled children alarm of an enabled parent appears on the System Status report with a disabled status.

To set alarm parameters

1. Choose **Configure > Alarms**.

2. Under Enable Alarms, complete the configuration as described in this table.

Alarm	Description
Admission Control	<p>Enables an alarm if the AltaVault reaches the maximum number of connections that can be made to the AltaVault.</p> <p>By default, this alarm is enabled.</p>
Cloud Bucket Consistency	<p>Enables an alarm if there is data in the cloud, but the AltaVault data store is empty. To clear this alarm, enable replication and recovery to ensure that the cloud storage is synchronized with the data store.</p> <p>This alarm occurs when you perform disaster recovery without specifying the correct parameters.</p>
Cloud Bucket Disparity	<p>Enables an alarm when the cloud bucket that the AltaVault is trying to connect to is being used by another AltaVault appliance. This alarm prevents corruption of the files in the cloud.</p>
Cloud Bucket Over Capacity	<p>Enables an alarm when the cloud bucket that the AltaVault connects to has exceeded the licensed cloud capacity.</p>
CPU Utilization	<p>Enables an alarm if the average and peak thresholds for the CPU utilization are exceeded. When an alarm reaches the rising threshold, it is activated; when it reaches the lowest or reset threshold, it is reset. After an alarm is triggered, it is not triggered again until it has fallen below the reset threshold.</p> <p>If the CPU utilization alarm triggers when the AltaVault is under a heavy load, you can ignore it.</p> <p>By default, this alarm is enabled.</p> <p>Rising Threshold - Specify the rising threshold. When an alarm reaches the rising threshold, it is activated. The default value is 95%.</p> <p>Reset Threshold - Specify the reset threshold. When an alarm reaches the lowest or reset threshold, it is reset. After an alarm is triggered, it is not triggered again until it has fallen below the reset threshold. The default value is 70%.</p>
Data Integrity Error	<p>Enables an alarm when inconsistency in the data stored on the disk is detected.</p>
Datastore Eviction	<p>Indicates that the system has detected an issue with datastore eviction.</p> <p>The alarm triggers when the appliance starts evicting data from the local disk cache and the age of the evicted data is relatively young. If disk space runs low, the appliance starts evicting cached data that has not been used recently, keeping only the most recent data.</p> <p>The AltaVault keeps statistics about how old the evicted data is (this is the average evicted age). Usually, only old data is evicted. However, the appliance might be experiencing a large workload where more recent data needs to be evicted from the appliance to make space for incoming data. This causes the average evicted age to decrease, and when it goes below a certain threshold, the average evicted age alarm triggers. This alarm is an anomalous event, signaling that the appliance is handling a much larger workload than expected.</p> <p>This alarm is useful in detecting whether the appliance is undersized relative to your normal workload. If the alarm is constantly triggered, then you should consider increasing AltaVault's disk cache.</p>
Datastore Low Space	<p>Indicates that the local data store is running out of space and the eviction process on the AltaVault is unable to run at a sufficient pace to create space on the disk cache.</p> <p>This alarm might also trigger when replication is too slow.</p> <p>View the Eviction Optimization report (choose Reports > Eviction) to determine how much disk cache is available.</p>

Alarm	Description
Disk Full	<p>Enables an alarm if the system partitions (not the AltaVault data store) are full or almost full. For example, AltaVault monitors the available space used to hold logs, statistics, system dumps, and TCP dumps.</p> <p>By default, this alarm is enabled.</p> <p>This alarm monitors the following system partitions:</p> <ul style="list-style-type: none"> • /boot Full • /bootmgr Full • /config Full • /tmp Full • /var Full
Hardware	<p>Fan Error - Enables an alarm when an appliance fan error is detected (the fan is either missing or running at a low speed).</p> <p>Battery Backup Unit - Enables an alarm when battery backup unit is detected.</p> <p>IPMI - Indicates that there has been a physical security intrusion, triggering an Intelligent Platform Management Interface (IPMI) error. The following events trigger the IPMI alarm:</p> <ul style="list-style-type: none"> • Chassis intrusion (physical opening and closing of the appliance case) • Memory errors (ECC memory errors that can or cannot be corrected) • Hard drive faults or predictive failures • Power supply status or predictive failures <p>The option to reset the alarm appears only after the service triggers the IPMI alarm. To reset the alarm, click Clear the IPMI alarm now.</p> <p>Memory Error - Enables an alarm when there is a memory error in one or more memory modules. Unplug the power cords from the power supply and try reseating the memory.</p> <p>Power Supply - Enables an alarm when an inserted power supply cord does not have power, as opposed to a power supply slot with no power supply cord inserted. By default, this alarm is enabled.</p> <p>RAID - Indicates that the system has encountered RAID errors.</p> <p>For drive rebuilds, if a drive is removed and then reinserted, the alarm continues to be triggered until the rebuild is complete.</p> <p>Important: Rebuilding a disk drive can take 12 hours or longer.</p> <p>By default, this alarm is enabled.</p> <p>You can enable or disable the alarm for a specific RAID disk. To enable or disable an alarm, choose Settings > Alarms and select or clear the check box next to the RAID disk name. This alarm monitors and displays the status of the RAID disks.</p> <p>RAID Integrity Check - Enables an alarm when RAID integrity check is needed.</p> <p>Shelf Power Supply - Enables an alarm when shelf power supply is needed.</p>
Inconsistent Cloud Connectivity	Enables an alarm when the connection to the cloud provider is inconsistent.
Inconsistent Cloud Data	Enables an alarm when inconsistency in the data stored in the cloud is detected.

Alarm	Description
Licensing	<p>Enables an alarm and sends an email notification if a license on the AltaVault is removed, is about to expire, has expired, or is invalid.</p> <p>The licenses expiring and licenses expired alarms are triggered per feature. For example, if you install two license keys for a feature, AVA-FOO-xxx (expired) and AVA-FOO-yyy (not expired), the alarms do not trigger, because the feature has one valid license.</p> <p>By default, this alarm is enabled.</p>
Link Duplex	<p>Enables an alarm and sends an email notification when an interface is not configured for half-duplex negotiation but has negotiated half-duplex mode. Half-duplex significantly limits the optimization service results.</p> <p>The alarm displays which interface is triggering the duplex alarm.</p> <p>By default, this alarm is enabled.</p>
Link I/O Errors	<p>Enables an alarm and sends an email notification when the link error rate exceeds 0.1% while either sending or receiving packets. The alarm clears when the rate drops below 0.05%.</p> <p>You can change the default alarm thresholds by entering the alarm link_errors err-threshold xxxxx CLI command at the system prompt. For details, see the <i>NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Guide</i>.</p> <p>By default, this alarm is enabled.</p> <p>You can enable or disable the alarm for a specific interface. For example, you can disable the alarm for a link where you have decided to tolerate the errors. To enable or disable an alarm, choose Settings > Alarms and select or clear the check box next to one or more of the link names.</p>
Link State	<p>Enables an alarm and sends an email notification if an Ethernet link is lost.</p> <p>By default, this alarm is disabled.</p> <p>You can enable or disable the alarm for a specific interface. To enable or disable an alarm, choose Settings > Alarms and select or clear the check box next to one or more link names.</p>
Low Memory	<p>Enables an alarm when there is not enough memory in the system to start the Storage Optimization Service.</p>
Max inodes limit	<p>Enables an alarm when the maximum number of files that can be stored has been reached.</p>
Max Pinnable Limit	<p>Enables an alarm when the share has reached the maximum pinnable limit. If you configure a share to be pinned, it always has data available locally in the AltaVault; data need not be fetched from the cloud.</p>
Memory Paging	<p>Enables an alarm when the system has reached the memory paging threshold. If the AltaVault is exceeding 100 pages are swapped approximately every two hours, then reboot the AltaVault from the Maintenance > Reboot/Shutdown page to clear this alarm.</p> <p>If the memory paging alarm triggers when the AltaVault is under a heavy load, you can ignore it.</p>
Metadata Space Full	<p>Enables an alarm when the data reserved for storing system metadata has filled up and leading to reduced deduplication.</p>
Process Dump Creation Error	<p>Enables an alarm and sends an email notification if the system detects an error while trying to create a process dump. When the alarm is raised, the directory is blacklisted.</p> <p>By default, this alarm is enabled.</p>
Secure Vault	<p>Enables an alarm and sends an email notification if the system encounters a problem with the secure vault:</p> <ul style="list-style-type: none"> Secure Vault Locked - Indicates that the secure vault is locked. To optimize SSL connections or to use data store encryption, the secure vault must be unlocked. Go to Configure > Secure Vault and unlock the secure vault.

Alarm	Description
SMB	<p>Enables an alarm when AltaVault detects the Domain Controller is not reachable.</p> <ul style="list-style-type: none"> Domain Controller Network Status - Indicates the Domain Controller is unreachable. The alarm is cleared when network connectivity to the Domain Controller is restored. If the alarm is not cleared after the network connectivity is restored, you can clear the alarm manually using alarm smb_alarms clear command.
Software update available	Enables an alarm when a new version of the software is available.
Shelf Error	<p>Shelf Missing - This alarm is applicable only to the AltaVault models. The AltaVault Expansion Shelf is missing or cannot be accessed.</p> <hr/> <p>Shelf <shelf name></p> <ul style="list-style-type: none"> Shelf Inconsistent - The AltaVault Expansion Shelf is not consistent with the stored configuration. Shelf Not Empty - You have added a new AltaVault Expansion Shelf that is not empty. A new AltaVault Expansion Shelf must be empty before you add it to AltaVault appliance. Shelf Not Valid - The AltaVault Expansion Shelf is not a valid shelf. For details, choose Reports > Storage RAID Groups and click the serial number of the shelf.
Storage Optimization Service	<ul style="list-style-type: none"> Storage Optimization Service Down - Enables an alarm and sends an email notification if the Storage Optimization Service encounters a service condition. By default, this alarm is enabled. The message indicates the reason for the condition. The following conditions trigger this alarm: <ul style="list-style-type: none"> Configuration errors: examples include no encryption key set, incorrect appliance time, or incorrect cloud credentials. An AltaVault appliance reboot for example, during an appliance software update. A system crash due to a power failure A Storage Optimization Service restart due to a cloud storage provider change. A user enters the CLI command no service enable or shuts down the Storage Optimization Service from the Management Console A user restarts the optimization service from either the Management Console or CLI Storage Optimization Service Error - Enables an alarm and sends an email notification if the Storage Optimization Service encounters a condition that might degrade optimization performance. By default, this alarm is enabled. Go to the Maintenance > Service page and restart the optimization service.
Storage Optimization Service Replication	<ul style="list-style-type: none"> Replication Error - Enables an alarm when the replication to the cloud encounters an error. Displays an error message that indicates the type of error such as, a file cannot be replicated to the cloud. Replication Paused - Enables an alarm when the replication to the cloud pauses, because there is a cloud connection error, or you entered the CLI command no replication enable, or because you are using replication scheduling (nonbandwidth limit type). This alarm warns you that the AltaVault is not replicating data in the cloud. <p>By default, this alarm is enabled.</p>
System Reserved Space Full	Indicates that the space used for internal data structures is full. De-duplication performance is impacted while the appliance is in this state.

Alarm	Description
Temperature	<ul style="list-style-type: none"> • Critical Temperature - Enables an alarm and sends an email notification if the CPU temperature exceeds the rising threshold. When the CPU returns to the reset threshold, the critical alarm is cleared. The default value for the rising threshold temperature is 80° C; the default reset threshold temperature is 67° C. • Warning Temperature - Enables an alarm and sends an email notification if the CPU temperature approaches the rising threshold. When the CPU returns to the reset threshold, the warning alarm is cleared. • Rising Threshold - Specifies the rising threshold. The alarm activates when the temperature exceeds the rising threshold. The default value is 80° C. • Reset Threshold - Specifies the reset threshold. The alarm clears when the temperature falls below the reset threshold. The default is 67° C. <p>After the alarm triggers, it cannot trigger again until after the temperature falls below the reset threshold and then exceeds the rising threshold again.</p>
Upgrade Status	Indicates the status of the upgrade. By default, this alarm is enabled.

3. Click **Apply** to apply your changes to the running configuration.

Configuring date and time

You set the system date and time in the Configure > Date and Time page.

You can either set the system date and time by entering it manually, or by assigning an NTP server to the AltaVault. By default, the appliance uses the NetApp-provided NTP server.

To set the date and time manually

1. Choose Configure > Date and Time.
2. Complete the configuration as described in this table.

Control	Description
Time Zone	<p>Select a time zone from the drop-down list.</p> <p>If you change the time zone, log messages retain the previous time zone until you reboot the AltaVault.</p>
Set Time Manually	<p>Change Date - Specify the date in this format: YYYY/MM/DD.</p> <p>Change Time - Specify military time in this format: HH:MM:SS.</p>

3. Click **Apply** to apply your changes to the running configuration.

To use Network Time Protocol (NTP) time synchronization

1. Choose Configure > Date and Time.
2. Under Date and Time, select Use NTP Time Synchronization.
3. As a best practice, configure your own internal NTP servers.

Current NTP status

Brief status information appears just below the Use NTP Time Synchronization button. The label Current NTP server is followed by either a server name or nothing if no NTP server is active.

This information appears after an NTP server name:

- Authentication information; “unauthenticated” appears after the server name when it is not using authentication.
- When the system has no NTP information about the current server, nothing appears.

When you configure an NTP server pool, the current NTP server that appears after the label Current NTP server never matches the hostname of the server pool.

NTP MD5-based authentication

NTP authentication verifies the identity of the NTP server sending timing information to the AltaVault. The system supports MD5-based Message-Digest Algorithm symmetric keys for NTP authentication. MD5 is a widely used cryptographic hash function that produces a 128-bit (16-byte) hash value.

NTP authentication is *optional*.

Configuring NTP authentication involves these steps that you can perform in any order:

- Configure a key ID and a secret pair.
- Configure the NTP server with the key ID.

NTP servers

NetApp recommends synchronizing the AltaVault to an NTP server of your choice.

To add an NTP server

1. Choose Configure > Date and Time.
2. Under Requested NTP Servers, complete the configuration as described in this table.

Control	Description
Add a New NTP Server	Displays the controls to add a server.
Hostname or IP Address	Specify the hostname or IP address for the NTP server.
Version	Select the NTP server version from the drop-down list: 3 or 4.
Enabled/Disabled	Select Enabled from the drop-down list to connect to the NTP server. Select Disabled from the drop-down list to disconnect from the NTP server.
Key ID	Specify the MD5 key identifier to use to authenticate the NTP server. The valid range is 1 to 65534. The key ID must appear on the trusted keys list.
Add	Adds the NTP server to the server list.
Remove Selected	Select the check box next to the name and click Remove Selected .

NTP authentication keys

NTP authentication uses a key and a shared secret to verify the identity of the NTP server sending timing information to the AltaVault. The system encrypts the shared secret text using MD5, and uses the authentication key to access the secret.

To add an NTP authentication key

1. Choose **Configure > Date and Time**.
2. Under **NTP Authentication Keys**, complete the configuration as described in this table.

Control	Description
Add a New NTP Authentication Key	Displays the controls to add an authentication key to the key list. Both trusted and untrusted keys appear on the list.
Key ID	Optionally, specify the secret MD5 key identifier for the NTP server. The valid range is 1 to 65534.
Key Type	Select MD5 or SHA1 option.
Secret (Text)	Specify the shared secret. You must configure the same shared secret for both the NTP server and the NTP client to use MD5-based cryptography. The shared secret: <ul style="list-style-type: none"> • is limited to 16 ASCII characters or fewer • cannot include white space or #s • cannot be empty • is case sensitive The secret appears in the key list as its MD5 hash value.
Add	Adds the authentication key to the trusted keys list.
Remove Selected	Select the check box next to the name and click Remove Selected .

NTP key information

NTP keys appear in a list that includes the key ID, type, secret (displays as the MD5 hash value), and whether the system trusts the key for authentication.

Configuring SNMP basic settings

You configure Simple Network Management Protocol (SNMP) contact and trap receiver settings to enable event reporting to an SNMP entity in the **Configure > SNMP Basic** page.

Traps are messages sent by an SNMP entity that indicate the occurrence of an event. The default system configuration does not include SNMP traps.

AltaVault supports the following SNMP Basic settings:

- SNMP Version 1
- SNMP Version 2c
- SNMP Version 3, which provides authentication through the User-based Security Model (USM)
- View-Based Access Control Mechanism (VACM), which provides richer access control

- Enterprise Management Information Base (MIB)
- ACLs (Access Control Lists) for users (v1 and v2c only)

To set general SNMP basic parameters

1. Choose **Configure > SNMP Basic**.
2. Under **SNMP Server Settings**, complete the configuration as described in this table.

Control	Description
Enable SNMP Traps	Enables event reporting to an SNMP entity.
System Contact	Specify the username for the SNMP contact.
System Location	Specify the physical location of the SNMP system.
Read-Only Community String	Specify a password-like string to identify the read-only community. For example, public. This community string overrides any VACM settings. Community strings do not allow non-printable ASCII characters, except for spaces. Also, the community strings cannot begin with '#' and '-'.

3. Click **Apply** to apply your changes to the running configuration.

To add or remove a trap receiver

1. Under **Trap Receivers**, complete the configuration as described in this table.

Control	Description
Add a New Trap Receiver	Displays the controls to add a new trap receiver.
Receiver	Specify the destination IP address or hostname for the SNMP trap.
Destination Port	Specify the destination port.
Receiver Type	Select SNMP version v1, v2c, or v3 (user-based security model). Note: SNMP v1 and v2c are less secure, v3 is recommended.
Remote User	(Appears only when you select v3). Specify a remote username.
Authentication	(Appears only when you select v3). Optionally, select either Supply a Password or Supply a Key to use while authenticating users.
Authentication Protocol	(Appears only when you select v3). Select an authentication method from the drop-down list: <ul style="list-style-type: none"> • MD5 - Specifies the Message-Digest 5 algorithm, a widely used cryptographic hash function with a 128-bit hash value. This is the default value. • SHA - Specifies the Secure Hash Algorithm, a set of related cryptographic hash functions. SHA is considered to be the successor to MD5.
Password/Password Confirm	(Appears only when you select v3 and Supply a Password). Specify a password. The password must have a minimum of eight ASCII characters. Confirm the password in the Password Confirm text box.

Control	Description
Security Level	<p>(Appears only when you select v3).Determines whether a single atomic message exchange is authenticated. Select one of the following settings from the drop-down list:</p> <ul style="list-style-type: none"> No Auth - Does not authenticate packets and does not use privacy. This is the default setting. Auth - Authenticates packets but does not use privacy. AuthPriv - Authenticates packets using AES 128 and DES to encrypt messages for privacy. <p>A security level applies to a group, not to an individual user.</p>
Community	For v1 or v2 trap receivers, specify the SNMP community name. For example, public or private v3 trap receivers need a remote user with an authentication protocol, a password, and a security level.
Enable Receiver	Select to enable the new trap receiver. Clear to disable the receiver.
Add	Adds a new trap receiver to the list.
Remove Selected	Select the check box next to the name and click Remove Selected .

After upgrade, all previous traps and community string intact are visible.

To test an SNMP trap

1. Choose Configure > SNMP Basic.
2. Under SNMP Trap Test, click **Run**.

Configuring SNMP v3

SNMP v3 provides additional authentication and access control for message security. For example, you can verify the identity of the SNMP entity (manager or agent) sending the message.

Using SNMP v3 is more secure than SNMP v1 or v2; however, it requires more configuration steps to provide the additional security features.

Basic steps

1. Create the SNMP-server users. Users can be authenticated using either a password or a key.
2. Configure SNMP-server views to define which part of the SNMP MIB tree are visible.
3. Configure SNMP-server groups, which map users to views, allowing you to control who can view what SNMP information.
4. Configure the SNMP-server access policies that contain a set of rules defining access rights. Based on these rules, the entity decides how to process a given request.

To create users for SNMP v3

1. Choose Configure > SNMP v3.

2. Under Users, complete the configuration as described in this table.

Control	Description
Add a New User	Displays the controls to add a new user.
User Name	Specify the username.
Authentication Protocol	<p>Select an authentication method from the drop-down list:</p> <ul style="list-style-type: none"> MD5 - Specifies the Message-Digest 5 algorithm, a widely used cryptographic hash function with a 128-bit hash value. This is the default value. SHA - Specifies the Secure Hash Algorithm, a set of related cryptographic hash functions. SHA is considered to be the successor to MD5.
Authentication	Optionally, select either Supply a Password or Supply a Key to use while authenticating users.
Password/Password Confirm	<p>Specify a password. The password must have a minimum of eight ASCII characters. Confirm the password in the Password Confirm text box.</p> <p>The password cannot be "password."</p>
MD5 Key	(Appears only when you select Supply A Key). Specify a unique authentication key. The key is a MD5 or SHA-1 digest created using md5sum or sha1sum.
Privacy MD5/SHA Key	(Appears only when you select v3 and Privacy as Supply a Key). Specify the privacy authentication key. The key is either a 32-hexadecimal digit MD5 or a 40-hexadecimal digit SHA digest created using md5sum or sha1sum.
Use Privacy Option	<p>Privacy Protocol - Select the privacy protocol from the drop-down list. Choose AES or DES.</p> <p>Privacy - Select the privacy option from the drop-down list. Choose Same as Authentication, Supply a Password, or Supply a Key.</p>
Add	Adds the user.
Remove Selected	Select the check box next to the name and click Remove Selected .

3. Click **Apply** to apply your changes to the running configuration.

SNMP authentication and access control

The features on this page apply to SNMP v1, v2c, and v3 unless noted otherwise:

- Security Names - Identify an individual user (v1 or v2c only).
- Secure Groups - Identify a security-name, security model by a group, and referred to by a group-name.
- Secure Views - Create a custom view using the View-based Access Control Model (VACM) that controls who can access which MIB objects under agent management by including or excluding specific Object Identifiers (OIDs). For example, some users have access to critical read-write control data, while some users have access only to read-only data.
- Security Models - A security model identifies the SNMP version associated with a user for the group in which the user resides.
- Secure Access Policies - Defines who gets access to which type of information. An access policy contains <group-name, security-model, security-level, read-view-name>:
 - read-view-name is a preconfigured view that applies to read requests by this security-name.
 - write-view-name is a preconfigured view that applies to write requests by this security-name.
 - notify-view-name is a preconfigured view that applies to write requests to this security-name.

An access policy is the configurable set of rules, based on which the entity decides how to process a given request.

To set secure usernames

1. Choose Configure > SNMP ACLs.

2. Under Security Names, complete the configuration as described in this table.

Control	Description
Add a New Security Name	Displays the controls to add a security name.
Security Name	<p>Specify a name to identify a requestor allowed to issue gets and sets (v1 and v2c only). The specified requestor can make changes to the view-based access-control model (VACM) security name configuration.</p> <p>Community strings do not allow printable ASCII characters, except for spaces.</p> <p>Also, community strings cannot begin with '#' or '-' (hash or hyphen).</p> <p>This control does not apply to SNMPv3 queries. To restrict v3 USM users from polling a particular subnet, use the Management ACL feature.</p> <p>Traps for v1 and v2c are independent of the security name.</p>
Community String	<p>Specify the password-like community string to control access using a combination of uppercase, lowercase, and numerical characters to reduce the chance of unauthorized access to the AltaVault.</p> <p>Community strings do not allow printable ASCII characters, except for spaces. Also, the community strings cannot begin with '#' and '-'.</p> <p>If you specify a read-only community string (located in the SNMP Basic page under SNMP Server Settings), it takes precedence over this community name and allows users to access the entire MIB tree from any source host. If this is not desired, delete the read-only community string.</p> <p>To create multiple SNMP community strings on a AltaVault, leave the default public community string and then create a second read-only community string with a different security name. Or, you can delete the default public string and create two new SNMP access control lists (ACLs) with unique names.</p>
Source IP Address and Mask Bits	Specify the host IP address and mask bits to which you permit access using the security name and community string.
Add	Adds the security name.
Remove Selected	Select the check box next to the name and click Remove Selected .

3. Click **Apply** to apply your changes to the running configuration.

To set secure groups

1. Choose Configure > SNMP ACLs.

2. Under Groups, complete the configuration as described in this table.

Control	Description
Add a New Group	Displays the controls to add a new group
Group Name	Specify a group name.
Security Model and Name Pairs	<p>Click the + button and select a security model from the drop-down list:</p> <ul style="list-style-type: none"> v1 or v2c - displays another drop-down list; select a security name. v3 (usm) - displays another drop-down list, select a user. <p>To add another Security Model and Name pair, click the plus sign (+).</p>
Add	Adds the group name and security model and name pairs.
Remove Selected	Select the check box next to the name and click Remove Selected .

3. Click **Apply** to apply your changes to the running configuration.

To set secure views

1. Choose Configure > SNMP ACLs.
2. Under Views, complete the configuration as described in this table.

Control	Description
Add a New View	Displays the controls to add a new view.
View Name	Specify a descriptive view name to facilitate administration.
Includes	Specify the object identifiers (OIDs) to include in the view, separated by commas. For example, .1.3.6.1.4.1. By default, the view excludes all OIDs. You can specify .iso or any subtree or subtree branch. You can specify an OID number or use its string form. For example, .iso.org.dod.internet.private.enterprises.xxx.products.AltaVault.system.model
Excludes	Specify the OIDs to exclude in the view, separated by commas. By default, the view excludes all OIDs.
Add	Adds the view.
Remove Selected	Select the check box next to the name and click Remove Selected .

3. Click **Apply** to apply your changes to the running configuration.

To add an access policy

1. Choose Configure > SNMP ACLs.
2. Under Access Policies, complete the configuration as described in this table.

Control	Description
Add a New Access Policy	Displays the controls to add a new access policy.
Group Name	Select a group name from the drop-down list.
Security Level	Determines whether a single atomic message exchange is authenticated. Select one of the following from the drop-down list: <ul style="list-style-type: none"> • No Auth - Does not authenticate packets and does not use privacy. This is the default setting. • Auth - Authenticates packets but does not use privacy. • AuthPriv - Authenticates packets using AES or DES to encrypt messages for privacy. A security level applies to a group, not to an individual user.
Read View	Select a view from the drop-down list.
Add	Adds the policy to the policy list.
Remove Selected	Select the check box next to the name and click Remove Selected .

3. Click **Apply** to apply your changes to the running configuration.

Configuring email settings

You can set email notification parameters for events and failures in the **Configure > Email** page.

By default, email addresses are not specified for event and failure notification.

To set event and failure email notification

1. Choose **Configure > Email**.
2. Under **Email Notification**, complete the configuration as described in this table.

Control	Description
SMTP Server	Specify the SMTP server. You must have external DNS and external access for SMTP traffic for this feature to function. Make sure you provide a valid SMTP server to ensure that the users you specify receive email notifications for events and failures.
SMTP Port	Specify the port number for the SMTP server.
Report Events via Email	Specify this option to report events through email. Specify a list of email addresses to receive the notification messages. Separate addresses by spaces, semicolons, commas, or vertical bars.
Report Failures via Email	Specify this option to report failures through email. Specify a list of email addresses to receive the notification messages. Separate addresses by spaces, semicolons, commas, or vertical bars.
Override Default Sender's Address	Select this option to configure the SMTP protocol for outgoing server messages for errors or events. Specify a list of email addresses to receive the notification messages. Separate addresses by commas. You can also configure the outgoing email address sent to the client recipients. The default outgoing address is <code>do-not-reply@hostname.domain</code> . If you do not specify a domain the default outgoing email is <code>do-not-reply@hostname</code> .

3. Click **Apply** to apply your changes to the running configuration.

Configuring log settings

You set up local and remote logging in the **Configure > Logging** page.

By default, the system rotates each log file every 24-hours or if the file size reaches one Gigabyte uncompressed. You can change this to rotate every week or month and you can rotate the files based on file size.

The automatic rotation of system logs deletes your oldest log file, labeled as **Archived log #10**, pushes the current log to **Archived log # 1**, and starts a new current-day log file.

To set up logging

1. Choose **Configure > Logging**.
2. To rotate the logs immediately, under **Log Actions** at the bottom of the page, click **Rotate Logs**. After the logs are rotated, the following message appears:

```
logs have been successfully rotated
```

You can also schedule a log rotation based on time or the amount of disk space the log uses, described next.

3. Under Logging Configuration, complete the configuration as described in this table.

Control	Description
Minimum Severity	<p>Select the minimum severity level for the system log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list:</p> <ul style="list-style-type: none"> Emergency - Emergency, the system is unusable. Alert - Action must be taken immediately. Critical - Conditions that affect the functionality of the AltaVault. Error - Conditions that probably affect the functionality of the AltaVault. Warning - Conditions that could affect the functionality of the AltaVault, such as authentication failures. Notice - Normal but significant conditions, such as a configuration change. Info - Informational messages that provide general information about system operations. This is the default setting. <p>This control applies to the system log only. It does not apply to the user log.</p>
Maximum Number of Log Files	Specify the maximum number of logs to store. The default value is 10.
Lines Per Log Page	Specify the number of lines displayed per page when viewing the logs. The default value is 100.
Rotate Based On	<p>Specifies the rotation option:</p> <ul style="list-style-type: none"> Time - Select Day, Week, or Month from the drop-down list. The default setting is Day. Disk Space - Specify how much disk space, in megabytes, the log uses before it rotates. The default value is 16 MB. <p>The log file size is checked at 10-minute intervals. If there is an unusually large amount of logging activity, it is possible for a log file to grow larger than the set disk space limit in that period of time.</p>

4. Click **Apply** to apply your changes to the running configuration.

To add or remove a log server

1. Under Remote Log Servers, complete the configuration as described in this table.

Control	Description
Add a New Log Server	Displays the controls for configuring new log servers.
Server IP	Specify the server IP address.

Control	Description
Minimum Severity	<p>Select the minimum severity level for the log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list:</p> <ul style="list-style-type: none"> Emergency - Emergency, the system is unusable. Alert - Action must be taken immediately. Critical - Conditions that affect the functionality of the AltaVault. Error - Conditions that probably affect the functionality of the AltaVault. Warning - Conditions that could affect the functionality of the AltaVault, such as authentication failures. Notice - Normal but significant conditions, such as a configuration change. This is the default setting. Info - Informational messages that provide general information about system operations.
Add	Adds the server to the list.
Remove Selected	Select the check box next to the name and click Remove Selected .

- Click **Apply** to apply your changes to the running configuration.

Filtering logs by application or process

You can filter a log by one or more applications or one or more processes. This is particularly useful when capturing data at a lower severity level at which the AltaVault might not be able to sustain the flow of logging data that the service is committing to disk.

Log filters enable you to specify the logging level of individual processes independently.

To filter a log

- Choose **Configure > Logging**.

2. Under Per-Process Logging, complete the configuration as described in this table.

Control	Description
Add a New Process Logging Filter	Displays the controls for adding a process-level logging filter.
Process	<p>Select a process to include in the log from the drop-down list:</p> <ul style="list-style-type: none"> alarmd - Alarm Manager. cli - Command Line Interface. hald - Hardware abstraction daemon, which handles access to the hardware. Isiraid - LSI raid daemon. mgmtd - Device control and management, which directs the entire device management system. It handles message passing between various management daemons, managing system configuration and general application of system configuration on the hardware underneath through the hardware abstraction layer daemon (HALD). pm - Process Manager, which handles launching of internal system daemons and keeps them running. sched - Process Scheduler that handles one-time scheduled events. statsd - Statistics Collector that handles the statistics. wdt - Watchdog Timer, the motherboard watchdog daemon. webasd - Web Application Process, which handles the Web user interface.
Minimum Severity	<p>Select the minimum severity level for the log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list:</p> <ul style="list-style-type: none"> Emergency - Emergency, the system is unusable. This is the default setting. Alert - Action must be taken immediately. Critical - Conditions that affect the functionality of the AltaVault. Error - Conditions that probably affect the functionality of the AltaVault. Warning - Conditions that could affect the functionality of the AltaVault, such authentication failures. Notice - Normal but significant conditions, such as a configuration change. Info - Informational messages that provide general information about system operations.
Add	Adds the filter to the list, after which it logs at the selected severity and higher.
Remove Selected	Select the check box next to the name and click Remove Selected to remove the filter.

3. Click **Apply** to apply your changes to the running configuration.

CHAPTER 7 **Configuring security settings**

This chapter includes the following sections:

- [“Configuring general security settings” on page 83](#)
- [“Managing user permissions” on page 85](#)
- [“Configuring management login from Active Directory domain” on page 89](#)
- [“Setting RADIUS servers” on page 91](#)
- [“Configuring TACACS+ access” on page 92](#)
- [“Unlocking the secure vault” on page 93](#)
- [“Configuring Web settings” on page 94](#)
- [“Configuring KMIP” on page 97](#)
- [“Configuring appliance monitoring” on page 101](#)
- [“Configuring a management ACL” on page 103](#)
- [“Configuring SSH access and chained authentication” on page 104](#)
- [“Configuring SSO for AltaVault” on page 106](#)

Configuring general security settings

You can prioritize local, RADIUS, and TACACS+ authentication methods for the system and set the authorization policy and default user for RADIUS and TACACS+ authorization systems in the **Configure > General Settings** page.

Make sure to put the authentication methods in the order in which you want authentication to occur. If authorization fails on the first method, the next method is attempted until all of the methods have been attempted.

To set TACACS+ authorization levels (admin or read-only) to allow certain members of a group to log in, add the following attribute to users on the TACACS+ server:

```
service = rbt-exec {  
    local-user-name = "monitor"  
}
```

Replace `monitor` with `admin` for write access.

To set general security settings

1. Choose **Configure > General Settings**.
2. Under **Authentication Methods**, complete the configuration as described in this table.

Control	Description
Authentication Methods	<p>Select an authentication method from the drop-down list. The methods are listed in the order in which they occur. If authorization fails on the first method, the next method is attempted until all of the methods have been attempted.</p> <hr/> <p>Note: Prior to selecting the Kerberos/AD Only method, the AltaVault must have joined the AD domain. Additionally, AltaVault requires at least one local user account configured having an Admin role with Read/Write privileges. A local admin user account is required for system access should it become necessary to fall back to Local authentication.</p> <hr/>
For RADIUS/TACACS+, fallback only when servers are unavailable	Specifies that the AltaVault uses a RADIUS or TACACS+ server only when all other servers do not respond. Enabled is the default setting.
Authorization Policy	<p>Appears only for some Authentication Methods. Optionally, select one of the following policies from the drop-down list:</p> <ul style="list-style-type: none"> • Remote First - Checks for an authentication policy on the remote server first and only checks locally if the remote server does not have a policy set. • Remote Only- Only check the remote server. This is the default. • Local Only - Checks only the local server. All remote users are mapped to the user specified. Any vendor attributes received by an authentication server are ignored.
Default User	Optionally, select Admin or Monitor from the drop-down list to define the default authentication policy.

3. Click **Apply** to apply your changes to the running configuration.

Managing user permissions

You can change the administrator or monitor passwords and define role-based users in the **Configure > User Permissions** page.

There are two types of accounts:

- [“Capability-based accounts” on page 85](#)
- [“Role-based accounts” on page 85](#)

Capability-based accounts

The system has two built-in accounts, based on what actions you can take:

- **Admin** - The administrator user has full privileges. For example, as an administrator you can set and modify configuration settings, add and delete users, restart the AltaVault service, reboot the AltaVault, and create and view performance and system reports.
- **Monitor** - Monitor users can view reports and user logs and change their own password. A monitor user cannot make configuration changes.

Role-based accounts

Use the role-based management feature of AltaVault to specify what roles a user is assigned to, and what actions a user is permitted to perform on the appliance in each of those roles. You can specify role-based accounts for admin settings, general settings, prepopulation (prepop) settings, replication settings, report settings, security settings, and storage settings in the AltaVault.

A role-based account cannot modify another role-based or capability-based account. Only the Admin account and accounts with the admin settings role can create and modify role-based accounts.

This section describes the roles that you can assign for specific features.

Admin settings

You can assign users permissions to perform administrator activities, including creating and deleting other users. Users with the Admin role always have read/write permission for all other roles, even if those other roles explicitly indicate Deny for the user.

General settings

You can assign users permissions to configure the following General Settings:

- Software upgrades
- Licenses
- Email, SNMP settings, and Web settings.
- Hardware RAID settings
- Shelf settings
- Starting and stopping the Storage Optimization Service
- All Networking Settings
- All Maintenance Settings

- All System Settings
- System logs
- Accessing system dumps and process dumps
- Debugging commands such as the alarm command
- Tcpdumps

Prepop settings

You can assign users permissions to start a new prepopulation task and to view an existing prepopulation task.

Replication settings

You can assign users permissions to configure the following Replication Settings:

- Cloud configuration
- Replication settings
- Starting and stopping the Storage Optimization Service

Report settings

You can assign users permissions to configure the following read-only Report Settings:

- Interface statistics
- Alarm Status
- View report graphs and statistics

Security settings

You can assign users permissions to configure the following Security Settings:

- Kerberos/AD
- RADIUS
- TACACS
- FIPS
- Secure vault
- Import, export, generate, and reset encryption key
- Import
- Export
- Reboot/Shutdown
- Appliance Monitoring
- Alarm
- User Log

Storage settings

You can assign users permissions to configure the following Storage Settings:

- SMB
- NFS
- OST
- SnapMirror

Configuring permissions for user roles

You can specify the following permissions for each role:

- Deny - You cannot view settings or make configuration changes for a feature.
- Read-Only - You can view current configuration settings but not change them.
- Read/Write - You can view settings and make configuration changes for a feature.

To configure user permissions

1. Choose Configure > User Permissions.
2. Under Capability-Based Accounts, complete the configuration as described in this table.

Control	Description
admin/monitor	Click the magnifying glass icon to change the administrator or monitor password.
	Enable Account - Click the check box to enable or disable the administrator or monitor account.
	Change Password - Select the check box to change password protection.
	<ul style="list-style-type: none"> • New Password - Specify a password in the text box. The password cannot be “password” or any case combination of “password” for any user including admin and root. You will be prompted with the following message: Password “password” and its case combinations are not allowed. The password must be at least 6 ASCII characters long. • New Password Confirm - Confirm the new administrator password.

3. Under Role-Based Accounts, complete the configuration as described in this table.

Control	Description
Add a New User	Click to display the controls for creating a new role-based account.
Account Name	Specify a name for the role-based account.
	<p>Note: If you are creating a user role for management login from the Active Directory domain, the name you enter must be the same as the user name in the Active Directory.</p>
Password	Specify the new password. The password cannot be “password” or any case combination of “password” for any user including admin and root and must be at least 6 ASCII characters long.
	This password can be different from the AD password.
New Password Confirm	Confirm the new password.

Control	Description
External Authentication Only	If this option is selected, then this user can only be authenticated via external authentication methods. If Kerberos/AD authentication is enabled, the local password originally configured for a user is no longer retained by AltaVault. If you disable external authentication, you will need to create a new password.
Enable Account	Select the check box to enable the new role-based account.
Roles and Permissions	For the account being created, specify the desired permissions for each role. Click Select All to choose the given access level for all feature settings.
Add	Adds your settings to the system.
Remove Selected Users	Select the check box next to the name and click Remove Selected .

Unlocking an account

AltaVault temporarily locks out an account after a user exceeds the configured number of login attempts. Account lockout information appears on the [Configure > User Permissions](#) page.

When an account is locked out, the lockout ends after:

- The configured lockout time elapses.
—or—
- The administrator unlocks the account. AltaVault never locks out the capability-based admin account.

To unlock an account

1. Log in as admin or any role-based user with read/write permission for the admin role.
2. Choose [Configure > User Permissions](#).
3. Select the user to display Edit User section.
4. Click **Clear Login Failure Details** to unlock the user account.

When you log in to your account successfully, AltaVault resets the login failure count.

Configuring password policy settings

You configure password complexity and lockout requirements for local management logins using Password Policy settings.

To configure password policy

1. Choose [Configure > User Permissions](#).
2. Click **Password Policy** at the bottom of the page.
3. Select **Enable Account Control**.
4. Optionally, you can choose to populate the password settings with a predetermined set of values.
To see these values, move your cursor over each of the template options: **Strong Security Template** or **Basic Security Template**. The default values appear next to each field. Click on a template to select it.

For new installations, the password settings are prepopulated with basic security values.

5. Specify values for each of the following settings (default values shown):

- Login attempts before logout (no limit)
- Timeout for user login after logout (seconds) (300)
- Days before password expires (no limit)
- Days to warn user of an expiring password (no limit) - takes effect after setting Days before password expires
- Days to keep account active after password expires (no limit)
- Days between password changes (no limit)
- Minimum Interval for password reuse (0)
- Minimum password length (6)
- Minimum uppercase characters (0)
- Minimum lowercase characters (0)
- Minimum numerical characters (0)
- Minimum special characters (0)
- Minimum character difference between passwords (0)
- Maximum consecutively repeating characters (no limit)
- Choose whether to prevent dictionary words (yes)

6. Click **Apply** to save your settings.

Configuring management login from Active Directory domain

AltaVault supports management login from either the Management Console (UI) or command-line interface (CLI) for domain users using their Active Directory (AD) credentials.

Note: The built-in AltaVault admin and monitor user accounts cannot be used for AD login. After AD login is enabled, you will not be able to log in using the built-in admin or monitor account. Management login from the AD domain requires you to add user accounts with the read/write permission for the Admin settings role.

This section covers the following information:

- [“Configuring login from AD” on page 89](#)
- [“Login behavior using AD” on page 90](#)

Configuring login from AD

To configure management login via Active Directory

1. From the Management Console, choose **Configure > Host Settings**.

2. In the DNS settings area, specify the DNS servers that can contact the domain controllers used by AltaVault. The preferred domain controllers AltaVault can use are specified in the next steps.
3. From the Management Console, choose **Configure > SMB**.
4. If not already configured, select **Domain** and complete the domain configuration as described in [“Configuring an Active Directory domain,”](#) then click **Join Domain**.
For Username, you can enter any user that has administrator privileges to join the domain.
5. From the Management Console, choose **Configure > User Permissions**.
6. Under **Role-based Accounts**, select **Add a New User** and enter a user name and password. The user name must map to that of an existing user in the AD domain. Do not qualify the user name with a domain name. For example, “user” is acceptable, but DOMAIN\user or user@DOMAIN is not.
7. Under **Roles and Permissions**, select the roles and permissions provided to the user.
To enable AD login, you must assign this user with the **Admin** role and **read** and **write** permissions. This user will then have privileges to add, delete or change permissions for other users.
8. Click **Add** to save user roles and permissions.
9. Repeat steps 6 through 8 to add additional users.
10. From the Management Console, choose **Configure > General Settings**.
11. Under **Authentication Methods**, select **Kerberos/AD Only** from the drop down menu and click **Apply** to save your settings and enable management login from AD.

Note: You must have joined the AD domain and have created an admin user account prior to setting the authentication method.

12. Optionally, if your security policy requires that user passwords cannot be stored locally, choose **Configure > User permissions** from the Management Console. Select the user you wish to edit, and check the box **External Authentication Only**.

When this box is checked, the local password for this user is deleted from AltaVault and you must log in using AD credentials.

13. Optionally, to further limit AltaVault logins to use AD credentials only, disable SSH public key authentication in the CLI:

```
no ssh server pub-key-auth
```

Login behavior using AD

After enabling Kerberos for Active Directory login, accessing AltaVault has the following behaviors:

- Password authentication will be checked against Active Directory credentials, not local passwords.
- If the user password is changed in Active Directory, that user must log in using the new Active Directory password.

- If user is disabled or deleted in Active Directory, that user will not be able to log in to the AltaVault. To avoid losing access to the AltaVault, it is recommended that you configure more than one Admin user account for Active Directory access.
- AltaVault supports only individual Active Directory user accounts.

Setting RADIUS servers

You can optionally configure Remote Authentication Dial-in User Server (RADIUS) server authentication in the **Configure > RADIUS** page.

RADIUS is an access control protocol that uses a challenge and response method for authenticating users.

To configure RADIUS server authentication

1. Choose **Configure > RADIUS**.
2. Under **Default RADIUS Settings**, complete the configuration as described in this table.

Control	Description
Set a Global Default Key	Enables a global server key for the RADIUS server.
Global Key	Specify the global server key.
Confirm Global Key	Confirm the global server key.
Timeout (seconds)	Specify the time-out period in seconds (1 to 60). The default value is 3.
Retries	Specify the number of times that you want to allow the user to retry authentication. The default value is 1.

3. Click **Apply** to apply your changes to the running configuration.
4. To add a new RADIUS server, complete the configuration as described in this table.

Control	Description
Add a RADIUS Server	Displays the controls for defining a new RADIUS server.
Hostname or IP Address	Specify the hostname or IP address.
Authentication Port	Specify the port for the server.
Authentication Type	Select one of these authentication types: <ul style="list-style-type: none"> • PAP - Password authentication protocol (PAP), which validates users before allowing them access to the RADIUS server resources. PAP is the most flexible protocol but is less secure than CHAP. • CHAP - Challenge-Handshake Authentication Protocol (CHAP), which provides better security than PAP. CHAP validates the identity of remote clients by periodically verifying the identity of the client using a three-way handshake. This happens at the time of establishing the initial link and might happen again at any time afterwards. CHAP bases verification on a user password and transmits an MD5 sum of the password from the client to the server.

Control	Description
Override the Global Default Key	Select this check box to override the global server key for the server and specify the following: <ul style="list-style-type: none"> Server Key - Specify the override server key. Confirm Server Key - Confirm the override server key.
Timeout (seconds)	Specify the time-out period in seconds (1 to 60). The default value is 3.
Retries	Specify the number of times that you want to allow the user to retry authentication. Valid values are 0 to 5. The default value is 1.
Enabled	Select the check box to enable the new server.
Add	Adds the RADIUS server to the list.
Remove Selected	Select the check box next to the name and click Remove Selected .

If you add a new server to your network and you do not specify these fields at that time, the global settings are applied automatically.

Configuring TACACS+ access

You can optionally set up TACACS+ (Terminal Access Controller Access-Control System) server authentication in the **Configure > TACACS+** page.

TACACS+ is an authentication protocol that allows a remote access server to forward a login password for a user to an authentication server to determine whether access is allowed to a given system.

You can prioritize local, RADIUS, and TACACS+ authentication methods for the system and set the authorization policy and default user for RADIUS and TACACS+ authorization systems in the **General Settings** page.

To configure a TACACS+ server

1. Choose **Configure > TACACS+**.
2. Under **Default TACACS+ Settings**, complete the configuration as described in this table.

Control	Description
Set a Global Default Key	Enables a global server key for the server.
Global Key	Specify the global server key.
Confirm Global Key	Confirms the global server key.
Timeout (seconds)	Specify the time-out period in seconds (1 to 60). The default value is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are 0 to 5. The default is 1.

3. Click **Apply** to apply your changes to the running configuration.

4. To add or remove a TACACS+ server, complete the configuration as described in this table.

Control	Description
Add a TACACS+ Server	Displays the controls for defining a new TACACS+ server.
Hostname or IP Address	Specify the hostname or server IP address.
Authentication Port	Specify the port for the server. The default value is 49.
Authentication Type	Select either PAP or ASCII as the authentication type. The default value is PAP.
Override the Global Default Key	Specify this option to override the global server key for the server.
Server Key	Specify the override server key.
Confirm Server Key	Confirm the override server key.
Timeout (seconds)	Specify the time-out period in seconds (1 to 60). The default is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are 0 to 5. The default is 1.
Enabled	Enables the new server.
Add	Adds the TACACS+ server to the list.
Remove Selected	Select the check box next to the name and click Remove Selected .
If you add a new server to your network and you do not specify these fields, the system automatically applies the default settings.	

Unlocking the secure vault

The secure vault contains sensitive information from your AltaVault configuration, including the encryption key. These configuration settings are encrypted on the disk at all times, using 256-bit AES encryption.

You can unlock and change the password for the secure vault in the Secure Vault page.

Initially, the secure vault is keyed with a default password known only to the AltaVault software. This allows the AltaVault to automatically unlock the vault during system startup. You can change the password, but the secure vault does not automatically unlock on startup. If not using the default password, the user will need to provide the password to unlock secure vault. To use encryption, the secure vault must be unlocked.

If a password policy is enabled, the number of retries allowed for unlocking the secure vault is the same as the number of retries for locking out a user. The lockout duration is also the same as set in the password policy. To change the password policy, choose **Configure > User Permissions** and select **Password Policy** as the bottom of the page.

To unlock or change the password of the secure vault

1. Choose **Configure > Secure Vault**.

- Under Unlock Secure Vault, complete the configuration as described in this table.

Control	Description
Password	Type a password and click Unlock Secure Vault. Initially, the secure vault is keyed with a default password known only to the AltaVault software. This allows the system to automatically unlock the vault during system startup. You can change the password, but the secure vault does not automatically unlock on startup.
Unlock Secure Vault	Unlocks the vault.

- Under Change Password, complete the configuration as described in this table.

Control	Description
Current Password	Specify the current password. If you are changing the password that ships with the product, leave the text box blank.
New Password	Specify a new password for the secure vault.
New Password Confirm	Confirm the new password for the secure vault.
Change Password	Changes the password for the secure vault.

Configuring Web settings

You can modify Management Console Web user interface settings in the Configure > Web Settings page. For information on managing Web SSL certificates, see [“Managing web SSL certificates” on page 94](#).

To modify web settings

- Choose Configure > Web Settings.
- Under Web Settings, complete the configuration as described in this table.

Control	Description
Default Web Login ID	Specify the username that appears in the authentication page. The default value is admin.
Web Inactivity Timeout (minutes)	Specify the number of idle minutes before time-out. The default value is 15. A value of 0 disables time-out.
Allow Session Timeouts When Viewing Auto-Refreshing Pages	By default, session time-out is enabled. Clear the Allow box to disable the session time-out and remain logged-in indefinitely. Disabling this feature is not recommended and can pose a security risk.

- Click **Apply** to apply your changes to the running configuration.

Managing web SSL certificates

The AltaVault provides the following additional security features to manage SSL certificates used by the AltaVault Management Console Web user interface using HTTPS.

- Generate the certificate and key pairs on the AltaVault. This overwrites the existing certificate and key pair regardless of whether the previous certificate and key pair was self-signed or user added. The new self-signed certificate lasts for one year (365 days).
- Create certificate signing requests from the private key.
- Replace a signed certificate with one created by an administrator or generated by a third-party certificate authority.

To modify web SSL certificates

1. Choose Configure > Web Settings.
2. Under Web Certificate, select the Details tab.

The AltaVault identity certificate details appear, as described in this table.

Control	Description
Issued To/Issued By	Common Name - Specifies the common name of the certificate authority.
	Email - Specifies the email address of the contact person.
	Organization - Specifies the organization name (for example, the company).
	Locality - Specifies the city.
	State - Specifies the state.
	Country - Specifies the country.
Validity	Issued On - Specifies the date the certificate was issued.
	Expires On - Specifies the date the certificate expires.
Fingerprint	SHA1 - Specifies the SSL fingerprint.
Key	Type - Specifies the key type.
	Size - Specifies the size in bits.

3. To import certificate and private key, under Web Certificate, select the Replace tab and complete the configuration as described in this table.

Control	Description
Import Certificate and Private Key	<p>Select this option to import certificate and private key.</p> <p>Upload (PKCS-12, PEM or DER formats) - Select this option to upload the CA-signed certificate file. The page displays a CA-Signed Public Certificate control for browsing to the key and certificate files or a text box for copying and pasting the key and certificate.</p> <p>Paste it here (PEM only) - Select this option to paste the CA-signed certificate.</p> <p>Private Key - Select an option from the following:</p> <ul style="list-style-type: none"> • This private key is in a separate file (below) • This file includes the certificate and private key • The private key for this certificate was created with a CSR generated on this appliance
Separate Private Key	<p>Upload (PEM or DER formats) - Select this option to upload the private key file. The page displays a Private Key control for browsing to the key or a text box for copying and pasting the key. Click Browse to navigate to the file.</p> <p>Paste it here (PEM only) - Select this option to paste the private key.</p> <p>Decryption password - Specify the decryption password. It is required for PKCS-12 files.</p>
Import Certificate and Key	Imports the new private key and certificate.

4. To generate self-signed certificate and new private key, under Web Certificates, select the Replace tab and complete the configuration as described in this table.

Control	Description
Organization Name	Specify the organization name (for example, the company).
Organization Unit Name	Specify the organization unit name (for example, the section or department).
Locality	Specify the city.
State	Specify the state.
Country	Specify the country (2-letter code only).
Email Address	Specify the email address of the contact person.
Validity Period	Specify the validity period. You can select from 60 to 3650 days.
Cipher	Default is RSA.
Cipher Bits	Select private key size from the drop-down list
Generate Certificate and Key	Generates the Certificate and private key.

5. To generate a CSR, under Web Certificate, select the Generate CSR tab and complete the configuration as described in this table.

Control	Description
Common Name	Specify the common name.
Organization Name	Specify the organization name (for example, the company).
Organization Unit Name	Specify the organization unit name (for example, the section or department).

Control	Description
Locality	Specify the city.
State	Specify the state.
Country	Specify the country (2-letter code only).
Email Address	Specify the email address of the contact person.
Generate CSR	Generates the Certificate Signing Request.

- To view the certificate in PEM format, under Web Certificate, select the PEM tab.

Configuring KMIP

Key Management Interoperability Protocol (KMIP) is a standard describing communication between key management servers and their clients. AltaVault manages several important pieces of information that must be kept secure. These pieces include the datastore encryption key that encrypts user data and cloud credentials (which allow AltaVault to authenticate itself to the cloud provider). Without KMIP, these pieces of information are stored on a disk in an encrypted partition of AltaVault called the Secure Vault. They can also be exported in configuration archives. It is up to the user to keep these archives secure.

A user's environment may be running multiple AltaVault's as well as other appliances or services which also require own encryption keys and other sensitive information. The need for centralized key management has led to development of key management servers (KMS), which operates as the KMIP server.

During setup, the administrator specifies an external KMS to manage AltaVault's keys and cloud authentication parameters. The datastore encryption key and/or cloud authentication parameters will then be managed by the KMS. If AltaVault uses KMIP, the KMS must be running nominally in order for AltaVault to be accessible.

AltaVault implements the following KMIP functionality:

- Registering keys with a KMS
- Fetching previously registered keys from a KMS

Note: Keys retrieved from a key server are never stored on a disk, only in memory. You cannot export fetched keys from a key server.

This section includes:

- [“Using the Management Console to configure KMIP”](#)
- [“Using CLI to configure KMIP”](#)
- [“Troubleshooting KMIP”](#)

Using the Management Console to configure KMIP

This section includes the following information:

- [“To add a KMIP server” on page 98](#)
- [“To add KMIP keys” on page 98](#)

- [“To configure cloud settings” on page 99](#)
- [“To configure the encryption key” on page 99](#)

To add a KMIP server

Before you add a KMIP server, check Web Settings page to verify that you have a certificate under the PEM tab.

1. Choose Configure > KMIP.
2. Under KMIP Servers, select Add a New Server and complete as described in the table.

Control	Description
Key Server Name	Specify the key server name.
Hostname	Specify the hostname of the server.
Port	Specify the port number.
Protocol Version	Select the protocol version from the drop-down list.
Username	Specify the username.
Password	Specify the password.
Upload CA Certificate	Select Browse to navigate to the CA certificate. The certificate must be a .pem file.
Add	Adds the KMIP server to the AltaVault. The KMIP server displays in the table below.
Remove Selected	Select a KMIP server and click Remove Selected to delete. This will result in AltaVault not using the key any longer. But the key will remain on the KMS. Deleting the key from the KMS has to be done through the UI provided by the KMS

To add KMIP keys

1. Under KMIP Keys, select Add a New Key and complete as described in the table.

Control	Description
Key Server Name	Select the key server name that was added earlier from the drop-down. If the server is not available, you must add the KMIP server.
Key Name	Specify the key name of the server.
Type	Select the type from the drop-down. Secret Data - Select this option to manage cloud authentication. Symmetric Key - Select this option to manage datastore encryption key. The selected key must be an AES-256 key.
Register Key	Select yes or no from the drop-down list. Note: Select yes only if this key does not exist on the KMIP server. Select no if the key already exists on the KMIP server.
Key Data	Specify the cloud authentication parameters. This field displays only when the Register Key is set to Yes, and the Type is set to Secret Data.
UUID	Specify the UUID from your server. The UUID field displays only if the Register Key is set to No.

Control	Description
Add	Adds the KMIP keys to the AltaVault. The KMIP key displays in the table below.
Remove Selected	Select a KMIP key and click Remove Selected to delete.

To configure cloud settings

1. Choose Configure > Cloud Settings.
2. Select the Cloud tab.
3. Select your cloud provider.
4. Select Yes from the Use Keys from KMIP Server drop-down list.
5. Select the correct secret data object names for each cloud authentication parameter (Access Key and Secret Key).
6. Click **Apply** to save your settings.
7. Choose Maintenance > Service, and select **Start** to start the Storage Optimization Service.

To configure the encryption key

1. Choose Configure > Cloud Settings.
2. Select the **Encryption** tab.
3. Select **Use Key from KMIP server**.
4. Select a key from the drop-down list and click **Apply**.

Using CLI to configure KMIP

You can use CLI to configure KMIP. For more information, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Reference Guide* available on the NetApp Support at <https://mysupport.netapp.com> under the Documentation tab.

Troubleshooting KMIP

KMIP commands are normally used by the AltaVault to the KMIP server in two situations: when the service comes up (most common) and when an object is registered with the server. Activity from these actions are recorded by AltaVault in the Maintenance > System Logs page.

Example of a successful command

```
[mgmtd.INFO]: Executing KMIP command
[mgmtd.INFO]: KMIP request:
[mgmtd.INFO]: Tag: REQUEST_MESSAGE (0x420078), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:   Tag: REQUEST_HEADER (0x420077), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:     Tag: PROTOCOL_VERSION (0x420069), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:       Tag: PROTOCOL_VERSION_MAJOR (0x42006a), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]:       Tag: PROTOCOL_VERSION_MINOR (0x42006b), Type: INTEGER (0x02), Data: 0x00000000
[mgmtd.INFO]:     Tag: AUTHENTICATION (0x42000c), Type: STRUCTURE (0x01), Data: null
[mgmtd.INFO]:     Tag: CREDENTIAL (0x420023), Type: STRUCTURE (0x01), Data: null
[mgmtd.INFO]:       Tag: CREDENTIAL_TYPE (0x420024), Type: ENUMERATION (0x05), Data: 0x00000001 (USERNAME_AND_PASSWD
[mgmtd.INFO]:       Tag: CREDENTIAL_VALUE (0x420025), Type: STRUCTURE (0x01), Data: null
[mgmtd.INFO]:         Tag: USERNAME (0x420099), Type: TEXT_STRING (0x07), Data: testuser
[mgmtd.INFO]:         Tag: PASSWORD (0x4200a1), Type: TEXT_STRING (0x07), Data: *****
[mgmtd.INFO]:       Tag: BATCH_COUNT (0x42000d), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]:     Tag: BATCH_ITEM (0x42000f), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:       Tag: OPERATION (0x42005c), Type: ENUMERATION (0x05), Data: 0x0000000a (GET)
[mgmtd.INFO]:     Tag: REQUEST_PAYLOAD (0x420079), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:       Tag: UNIQUE_IDENTIFIER (0x420094), Type: TEXT_STRING (0x07), Data: E39C6A6BF0E74E52DCA3CE1ABFBCD7C
4B2153F24041439E4A0C4EA5A0A
[mgmtd.INFO]:       Tag: KEY_FORMAT_TYPE (0x420042), Type: ENUMERATION (0x05), Data: 0x00000002 (OPAQUE)
[mgmtd.INFO]: KMIP response:
[mgmtd.INFO]: Tag: RESPONSE_MESSAGE (0x42007b), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:   Tag: RESPONSE_HEADER (0x42007a), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:     Tag: PROTOCOL_VERSION (0x420069), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:       Tag: PROTOCOL_VERSION_MAJOR (0x42006a), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]:       Tag: PROTOCOL_VERSION_MINOR (0x42006b), Type: INTEGER (0x02), Data: 0x00000000
[mgmtd.INFO]:     Tag: TIME_STAMP (0x420092), Type: DATE_TIME (0x09), Data: 0x000000000560b08d1 Tue Sep 29 21:55:29 201
[mgmtd.INFO]:     Tag: BATCH_COUNT (0x42000d), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]:     Tag: BATCH_ITEM (0x42000f), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:       Tag: OPERATION (0x42005c), Type: ENUMERATION (0x05), Data: 0x0000000a (GET)
[mgmtd.INFO]:       Tag: RESULT_STATUS (0x42007f), Type: ENUMERATION (0x05), Data: 0x00000000 (SUCCESS)
[mgmtd.INFO]:     Tag: RESPONSE_PAYLOAD (0x42007c), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]:       Tag: OBJECT_TYPE (0x420057), Type: ENUMERATION (0x05), Data: 0x00000007 (SECRET_DATA)
[mgmtd.INFO]:       Tag: UNIQUE_IDENTIFIER (0x420094), Type: TEXT_STRING (0x07), Data: E39C6A6BF0E74E52DCA3CE1ABFBCD7C
```

Example of an unsuccessful command

```
[mgmtd.INFO]: KMIP request:
[mgmtd.INFO]: Tag: REQUEST_MESSAGE (0x420078), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: REQUEST_HEADER (0x420077), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: PROTOCOL_VERSION (0x420069), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: PROTOCOL_VERSION_MAJOR (0x42006a), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]: Tag: PROTOCOL_VERSION_MINOR (0x42006b), Type: INTEGER (0x02), Data: 0x00000000
[mgmtd.INFO]: Tag: AUTHENTICATION (0x42000c), Type: STRUCTURE (0x01), Data: null
[mgmtd.INFO]: Tag: CREDENTIAL (0x420023), Type: STRUCTURE (0x01), Data: null
[mgmtd.INFO]: Tag: CREDENTIAL_TYPE (0x420024), Type: ENUMERATION (0x05), Data: 0x00000001 (USERNAME_AND_PASSW
[mgmtd.INFO]: Tag: CREDENTIAL_VALUE (0x420025), Type: STRUCTURE (0x01), Data: null
[mgmtd.INFO]: Tag: USERNAME (0x420099), Type: TEXT_STRING (0x07), Data: testuser
[mgmtd.INFO]: Tag: PASSWORD (0x4200a1), Type: TEXT_STRING (0x07), Data: *****
[mgmtd.INFO]: Tag: BATCH_COUNT (0x42000d), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]: Tag: BATCH_ITEM (0x42000f), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: OPERATION (0x42005c), Type: ENUMERATION (0x05), Data: 0x0000000a (GET)
[mgmtd.INFO]: Tag: REQUEST_PAYLOAD (0x420079), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: UNIQUE_IDENTIFIER (0x420094), Type: TEXT_STRING (0x07), Data: asdf2
[mgmtd.INFO]: Tag: KEY_FORMAT_TYPE (0x420042), Type: ENUMERATION (0x05), Data: 0x00000002 (OPAQUE)
[mgmtd.INFO]: KMIP response:
[mgmtd.INFO]: Tag: RESPONSE_MESSAGE (0x42007b), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: RESPONSE_HEADER (0x42007a), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: PROTOCOL_VERSION (0x420069), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: PROTOCOL_VERSION_MAJOR (0x42006a), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]: Tag: PROTOCOL_VERSION_MINOR (0x42006b), Type: INTEGER (0x02), Data: 0x00000000
[mgmtd.INFO]: Tag: TIME_STAMP (0x420092), Type: DATE_TIME (0x09), Data: 0x00000000561d66a2 Tue Oct 13 13:16:34 201
[mgmtd.INFO]: Tag: BATCH_COUNT (0x42000d), Type: INTEGER (0x02), Data: 0x00000001
[mgmtd.INFO]: Tag: BATCH_ITEM (0x42000f), Type: STRUCTURE (0x01), Data:
[mgmtd.INFO]: Tag: OPERATION (0x42005c), Type: ENUMERATION (0x05), Data: 0x00000018 (QUERY)
[mgmtd.INFO]: Tag: RESULT_STATUS (0x42007f), Type: ENUMERATION (0x05), Data: 0x00000001 (OPERATION FAILED)
[mgmtd.INFO]: Tag: RESULT_REASON (0x42007e), Type: ENUMERATION (0x05), Data: 0x00000003 (AUTHENTICATION_NOT_SUCCES
[mgmtd.INFO]: Tag: RESULT_MESSAGE (0x42007d), Type: TEXT_STRING (0x07), Data: Authentication failed
[mgmtd.WARNING]: KMIP command returned OPERATION_FAILED, reason: AUTHENTICATION_NOT_SUCCESSFUL, message: Authenticatio
```

Example of an unsuccessful command (failure to connect)

```
[mgmtd.INFO]: [SSL_connect:before/connect initialization]
[mgmtd.INFO]: [SSL_connect:error in SSLv2/v3 write client hello B]
[mgmtd.INFO]: 140013611939528:error:2006A066:BIO routines:BIO_get_host_ip:bad hostname lookup:b_sock.c:146:host=null
[mgmtd.WARNING]: KMIP command returned IO, message: I/O Error
```

Common Errors

An authentication error could be caused the following:

- Incorrect username or password
- Incorrect client certificate.
- Misconfigured certificates.

Configuring appliance monitoring

You can set up any AltaVault as the monitoring master appliance that monitors peer AltaVaults. The AltaVault uses REST APIs that you can access to set up peer appliance monitoring.

After you configure REST API access and add the API access code for the monitored appliance, the Appliance Monitoring report enables you to view the health status, disk space, and cloud service status of the AltaVault.

The monitoring appliance probes the monitored peer appliances every 60 seconds by default.

To configure REST API Access

When you add an appliance to be monitored by the AltaVault, you must generate an API access code to enable authenticated communication between the monitoring master appliance and the monitored peer appliance.

1. Log in to the monitored AltaVault appliance.

2. Choose **Configure > REST API Access**.
3. To enable access to the REST APIs, under REST API Access Settings, select the **Enable REST API Access** check box.
4. Click **Apply**.
5. Complete the configuration as described in the table.

Control	Description
Add Access Code	Displays the controls to generate the API access code.
Description of Use	Specify a clear description of the monitoring appliance such as the hostname or IP address of the monitoring master appliance and a description such as “monitoring appliance.”
Generate New Access Code	Generates the new access code.
Use Existing Access Code	Select to use an existing REST API access code. When you are monitoring multiple appliances, you can use the same access code instead of creating a new one for each appliance.
Add	Adds the API access code to the AltaVault.
Remove Selected	Select an access code description from the table below and click Remove Selected to delete the selected REST API access code.

The added access code description appears in the Access Code Description table, along with the name of the user who created it.

6. Click the **Access Code Description**.
7. Copy the Access Code from the text field into a text editor, such as Notepad.

To specify the API access code in the monitoring appliance

1. Log in to the monitoring AltaVault appliance.
2. Choose **Reports > Appliance Monitoring**.
3. Complete the configuration as described in this table.

Control	Description
Add Monitored Appliance	Displays the controls to add a monitored appliance.
Hostname or IP address	Specify a valid hostname or IP address for the monitored appliance.
API Access Code	Specify the API access code that you obtained from the monitored appliance as specified in “To configure REST API Access” on page 101 .
Add	Adds the API access code to the AltaVault appliance.
Remove Selected Appliances	Select an access code from the table below and click Remove Selected Appliances to delete the selected REST API access code.

Configuring a management ACL

You can secure access to the AltaVault using an internal management Access Control List (ACL) in the **Configure > Management ACL** page. For information on the ACL rules, see [“ACL Management Rules” on page 103](#).

Using an internal management ACL, you can:

- restrict access to certain interfaces or protocols of an appliance.
- restrict inbound IP access to the AltaVault, protecting it from access by hosts that do not have permission.
- specify which hosts or groups of hosts can access and manage the AltaVault by IP address.

The Management ACL provides the following safeguards to prevent accidental disconnection from the AltaVault:

- It detects the IP address you are connecting from and displays a warning if you add a rule that denies connections to that address.
- It converts well-known port and protocol combinations such as SSH, Telnet, HTTP, HTTPS, SNMP, and SOAP into their default management service and protects these services from disconnection. For example, if you specify protocol 6 (TCP) and port 22, the management ACL converts this port and protocol combination into SSH and protects it from denial.
- It tracks changes to default service ports and automatically updates any references to changed ports in the access rules.

To set up a management ACL

1. Choose **Configure > Management ACL**.
2. Under **Management ACL Settings**, complete the configuration as described in this table.

Control	Description
Enable Management ACL	Select the check box to secure access to a AltaVault using a management ACL.

3. Click **Apply** to apply your changes to the running configuration.

If you add, delete, or modify a rule that could disconnect connections to the AltaVault, a warning message appears. Click **Confirm** to override the warning and allow the rule definition anyway. Use caution when overriding a disconnect warning.

ACL Management Rules

The management ACL contains rules that define a match condition for an inbound IP packet. You set a rule to allow or deny access to a matching inbound IP packet. When you add a rule on a AltaVault, the destination specifies the AltaVault itself, and the source specifies a remote host.

To add an ACL management rule

1. Choose **Configure > Management ACL**.

2. Under Add a new rule, complete the configuration as described in this table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule.
Action	<p>Select one of the following rule types from the drop-down list:</p> <ul style="list-style-type: none"> Allow - Allows access when packets match the specified criteria. This is the default action. Deny - Denies access when packets match the specified criteria.
Service	Optionally, select Specify Protocol, or HTTP, HTTPS, SOAP, SNMP, SSH, Telnet. When specified, the Destination Port is dimmed and unavailable.
Protocol	(Appears only when Service is set to Specify Protocol.) Optionally, select All, TCP, UDP, or ICMP from the drop-down list. The default setting is All. When set to All or ICMP, the Service and Destination Ports are dimmed and unavailable.
Source Network	Optionally, specify the source subnet of the inbound packet. For example, 1.2.3.0/24.
Destination Port	Optionally, specify the destination port of the inbound packet, either a single port value or a port range of port1-port2, where port1 must be less than port2. Leave it blank to specify all ports.
Interface	Optionally, select an interface name from the drop-down list. Select All to specify all interfaces.
Description	Optionally, describe the rule to facilitate administration.
Rule Number	<p>Optionally, select a rule number from the drop-down list. By default, the rule goes to the end of the table (just above the default rule).</p> <p>AltaVaults evaluate rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied, and the system moves on to the next packet. If the conditions set in the rule do not match, the system consults the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted.</p> <p>The default rule, Allow, which allows all remaining traffic from everywhere that has not been selected by another rule, cannot be removed and is always listed last.</p>
Log Packets	Tracks denied packets in the log. By default, packet logging is enabled.
Add	Adds the rule to the list. The Management Console redisplay the Rules table and applies your modifications to the running configuration, which is stored in memory.
Remove Selected	Select the check box next to the name and click Remove Selected .
Move Selected	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

Configuring SSH access and chained authentication

AltaVault supports SSH access to the management port of the appliance. SSH access can be done using the client public key, user credentials (username/password), or chained authentication using both public key and user credentials (multifactor authentication).

Authentication using user credentials is provided by default when accessing the AltaVault using SSH. This section describes how to use AltaVault CLI commands to configure SSH access using public keys and chained authentication.

This section covers the following topics:

- [“To configure SSH access via public key” on page 105](#)
- [“To enable SSH access via chained authentication” on page 106](#)
- [“To disable SSH service” on page 106](#)

To configure SSH access via public key

SSH access using public keys is enabled by default. Use this procedure to set up user public keys.

1. Log in to the AltaVault using the credentials for a user with security settings role.
2. At the user prompt, enter configuration mode:

```
hostname> enable
hostname# configure terminal
hostname (config)#
```

3. At the configuration mode prompt, enter the following SSH command, providing the public SSH key information in the <authorized-key> field.

```
hostname (config)# ssh server user <username> authorized-key <authorized-key>
```

For example:

```
hostname (config) # ssh server user admin authorized-key "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQ7ZCATt6tD3t5JmS276WzIJpVoPn+0ReCbRThpyh/2GlsV346
XV3OdZ1954gd2n1kpeMckt7iSv6EgF2oGWMPE1h9tiCY5PKumZsT7bwQ94Y8IML+ZsldggVDOqRXRyXsInAm0hLCOFp3Ux
g4SUxjcTwJM82ZP6jTSMVLVjxWJhZKqJLzQpsUgv0BuAaQWdeS6vyNmgxfm+Fpv4Ov2o376sPSmnePodkyGnXTnn1JoQPH0
ICrrwt8of6IxObKH9HEBUaO94qZ+XLT+7SM6s9j4uR53KON8DnHNkntpGFDmR9hL6Krg9KWVCOb7Z0amNDk1p4y4bOkcMk
AXMm+v6ldT user@example.com"
```

To delete an authorized public key (disabling SSH access for a given user), enter the index number of the public key. For example, the first key added would have index value 1.

```
hostname (config) # no ssh server user admin authorized-key ?
hostname (config) # no ssh server user admin authorized-key 1
```

Note: With the public-key deleted, SSH login is accomplished with user credentials.

4. Enter the following command to display the SSH public key settings and to verify public key authentication is enabled.

```
hostname (config)# show ssh server
SSH server enabled: yes
SSH server listen enabled: no
SSH password authentication enabled: yes
SSH public key authentication enabled: yes
SSH chained authentication enabled: no
SSH port: 22
SSH max auth tries: 6
SSH v2 only: yes
```

5. If public-key authentication is disabled, enable it:

```
hostname (config)# ssh server pub-key-auth
```

To disable authentication, use this command:

```
hostname (config)# no ssh server pub-key-auth
```

To enable SSH access via chained authentication

With chained authentication, the public key is checked first. If the key information doesn't match, access is denied. If the key information matches, the user is prompted for a password. Password authentication can be local, Kerberos/AD, RADIUS or TACACS.

Prior to enabling chained authentication, password and public-key authentication must be enabled. See [“To configure SSH access via public key” on page 105](#). Additionally, at least one role-based admin account requires being configured with a public key. See [“Managing user permissions” on page 85](#).

1. Log in to the AltaVault with the login name and password.
2. At the user prompt, enter configuration mode: command:

```
hostname> enable
hostname# configure terminal
hostname (config)#
```

3. Enable chained authentication using the following command:

```
hostname (config) # ssh server chained-auth
```

Enabling chained-authentication applies to all SSH users regardless of role. To disable chained authentication, enter the `no ssh server chained-auth` command.

4. Enter the following command to display the SSH public key settings and to verify chained authentication is enabled.

```
hostname (config)# show ssh server
SSH server enabled: yes
SSH server listen enabled: no
SSH password authentication enabled: yes
SSH public key authentication enabled: yes
SSH chained authentication enabled: yes
SSH port: 22
SSH max auth tries: 6
SSH v2 only: yes
```

To disable SSH service

SSH service is enabled by default on the AltaVault. To disable SSH service you must be connected to the AltaVault serial console. You cannot disable SSH service while connected using SSH. To disable SSH service, enter the following command:

```
hostname (config) # no ssh server enable
```

Configuring SSO for AltaVault

Single Sign-On (SSO) enables web access to AltaVault using an external identity provider (IdP). AltaVault supports SSO using the standards-based Security Assertion Markup Language (SAML) 2.0. With SAML, the AltaVault acts as a service provider (SP), redirecting log in requests to IdP for authentication. The IdP and AltaVault exchange metadata files, as part of the configuration, to establish trust between them.

By redirecting AltaVault login requests to the IdP, users are subject to the authentication policies of the IdP, including additional credentials that might be required for multifactor authentication (MFA), such as hardware tokens, PINs, or mobile passcodes.

After user authentication with the IdP, the IdP issues a SAML response message to the AltaVault. The message includes assertions about authentication along with the attributes of the user. The IdP then redirects the user back to AltaVault. As long as the token is valid, users can access the AltaVault without further authentication.

SSO enables AltaVault users to login to any AltaVault that has been configured to use the same IdP. However, the user's access on each AltaVault will be dependent on the user role-based permissions assigned to that user on each AltaVault. That means the same IdP authenticated user may have different roles and permissions on each AltaVault.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) for a list of supported IdPs.

This section includes the following topics:

- [“Before you begin” on page 107](#)
- [“Configuring SSO” on page 107](#)
- [“Enabling SSO Service” on page 108](#)
- [“Troubleshooting SSO” on page 110](#)

Before you begin

- Identify the IdP metadata URL or file for uploading to the AltaVault.
- Check that local users with the appropriate privileges exist on the AltaVault. See [“Managing user permissions” on page 85](#).
- Verify the IdP and AltaVault are synchronized to a known good time source. See [“Configuring date and time” on page 69](#).
- Establish an SSH or console connection with the AltaVault and log in. Providing access to the AltaVault command line interface is best practice when configuring SSO and can help to recover system access in the event of a configuration issue or error.

Configuring SSO

1. From the Management Console, choose **Configure > SSO**.
2. Under Identity Provider Metadata, complete the configuration as described in this table:

Control	Description
Upload IdP Metadata file	<p>Select this option to specify the file from which to upload IdP metadata to AltaVault. This information is required by AltaVault to establish a trust relationship with the IdP.</p> <p>To select a file, click Browse.</p>

Control	Description
Specify URL	<p>Select this option to specify a URL from which IdP metadata can be fetched by AltaVault. This information is required by AltaVault to establish a trust relationship with the IdP.</p> <p>Disable Certificate Check - HTTPS automatically checks for certificates from a trusted certificate authority (CA) when retrieving metadata. Check this box if you are using a self-signed or internal root CA certificate.</p>
Username Mapping	<p>Edit these setting to map SSO login users to local user accounts:</p> <p>Attribute Name - The Attribute Name identifies the incoming username provided by the IdP to the AltaVault (SP). Enter the attribute name exactly as it appears in the IdP. The value of this attribute must match the local user being logged in on AltaVault.</p> <p>Attribute Name Format - The Attribute Name Format defines how the name information is presented in the security assertion. By default the Attribute Name Format is set to URI reference as: urn:oasis:names:tc:SAML:2.0:attrname-format:uri". If the format is different than this for the attribute that maps to the authenticated username, you must reset it as required.</p> <p>Note: An incorrect Attribute Name or Attribute Name Format will cause errors and prevent user log in. The AltaVault default admin capability-based account can log in using SSO if there is a user account on the IdP mapped to "admin."</p>

3. Click **Apply**.

AltaVault displays the SSO state information, SP entity ID, and IdP entity ID. Verify the SP entity ID includes the current AltaVault hostname and the IdP entity ID reflects the IdP provider. If necessary, repeat steps 2 and 3.

4. Under Service Provider Actions, complete the configuration as described in this table:

Control	Description
Reset SP Entity ID	Click this field to update the service provider (SP) metadata information on the AltaVault with the current hostname.
Download SP Metadata	<p>After resetting the SP Entity ID, click this field to download the information required by the IdP to establish a trust relationship with the AltaVault.</p> <p>Caution: You must download the SP metadata prior to enabling SSO service. Enabling SSO without configuring trust on IdP by uploading SP metadata will result in a lock out.</p>

5. Configure your IdP with AltaVault (SP) metadata and name mapping information.

For example, in ADFS, add a Relying Party Trust using the downloaded SP metadata file. When prompted, create a rule for mapping the LDAP attribute SAM-Account-Name to the outgoing claim type Common Name. For a successful login, AltaVault must have local users with the same username as the SAM-Account-Name for users on the Active Directory. Refer to your IdP for detailed configuration information.

Enabling SSO Service

Prior to enabling SSO Service, complete the procedure for configuring SSO on the AltaVault.

1. Under Single Sign-On Service, click **Enable** to start the service. A message appears asking for confirmation to enable SSO, reminding to download SP metadata and advising that you will be logged out of the current web session.

Note: AltaVault establishes a trust relationship with the IdP using DNS names. To access the AltaVault after enabling SSO, use the hostname instead of the IP address. After enabling SSO, all existing Web user sessions with AltaVault will be required to re-authenticate.

To disable the SSO service, click **Disable**. After disabling SSO, users log in using previously established authentication methods.

2. Click **Yes**. The system redirects you to the IdP login page.
3. Log in using your IdP/SSO credentials.

If login with the IdP is successful, you will be logged in to the AltaVault with the privileges provided for the local user account mapped to your IdP credentials. You must have admin or security user privileges to view the SSO page on the AltaVault.

Troubleshooting SSO

Issue	Problem and Resolution
Unable to log in using SSO (incorrect user ID or password)	<p>Problem: The IdP provider could be down or the IdP account information for the user is incorrect.</p> <p>Resolution: Verify the IdP status and user account information in the database and update as necessary.</p>
Unable to access AltaVault using IP address	<p>Problem: AltaVault establishes a trust relationship with the IdP using DNS names. When accessing AltaVault using the IP address, the IdP attempts to identify the trust relationship but is unable find a match.</p> <p>Resolution: Use hostname instead of IP address to access AltaVault. If your environment does not have DNS, then update the hosts file to add an entry for AltaVault.</p>
Unable to access the AltaVault after SSO login (unauthorized user)	<p>Problem: AltaVault local user account name does not match the value of IdP name attribute or the attribute mapping for username is incorrect.</p> <p>Resolution: Perform SSO login again using a different user account that has admin role privileges. If access to AltaVault is successful, verify and update the local user account information.</p> <p>If login is unsuccessful using a different user account, log in to the CLI and verify the user account information. Enter the show users and show rbm users to determine user privileges and roles. Update user information as necessary using the username and rbm user commands and trying logging in again.</p> <p>If the error persists, enter the show sso idp attribute map username command to verify the name mapping is properly configured. To change the name mapping attribute or format, use the sso idp attribute-map and sso idp metadata fetch commands, respectively.</p>
Unable to access the AltaVault after SSO login (SAML exception error)	<p>Problem: SSO service may have been enabled on the AltaVault prior to downloading the SP metadata file, or the SP metadata file uploaded to the IDP is invalid.</p> <p>Resolution: Enter the show sso sp-metadata command from the CLI and print out the SP metadata information. Copy that content to an .xml file and upload that to IdP.</p> <p>Alternatively, enter the no sso enable command to disable SSO. Log in using the Management Console and configure SSO again.</p>
Unable to access the AltaVault after SSO login (SAML exception error, message expired)	<p>Problem: A clock synchronization issue may exist between the IdP and the SP. Assertions messages exchanged between the IdP and SP include attributes that define how long an assertion is valid.</p> <p>Resolution: Ensure IdP and SP clocks are synchronized to a known good time source.</p>
Cannot change hostname	<p>Problem: It is not possible to change the AltaVault hostname after enabling SSO. This is the intended behavior.</p> <p>Resolution: If it is necessary to change the hostname, first disable SSO and then change the hostname. After changing the hostname, reset the SP entity ID, download SP metadata, and replace the SP metadata file on the IdP.</p>

CHAPTER 8 **Configuring AltaVault appliances for FIPS-compliant cryptography**

This chapter includes the following sections:

- [“What is FIPS?” on page 111](#)
- [“Understanding FIPS on AltaVault” on page 111](#)
- [“Configuring AltaVault for FIPS compliance” on page 113](#)
- [“Configuring AltaVault appliances for FIPS-compliant cryptography” on page 113](#)
- [“Disabling FIPS mode” on page 119](#)
- [“Verifying FIPS mode in system logs” on page 119](#)
- [“FIPS CLI” on page 120](#)

You configure and manage FIPS mode through the Command-Line Interface (CLI). For detailed information about the CLI, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Reference Guide*.

What is FIPS?

Federal Information Processing Standard (FIPS) is a publicly announced set of validation standards developed by the United States National Institute of Standards and Technology (NIST) for use by government agencies and by government contractors.

FIPS 140-2 details the U.S. and Canadian Government requirements for cryptographic modules. Protection of a cryptographic module within a security system is necessary to maintain the confidentiality and integrity of the information protected by the module.

This standard specifies the security requirements satisfied by a cryptographic module utilized within a security system protecting sensitive but unclassified data.

Understanding FIPS on AltaVault

This section describes the NetApp Cryptographic Security Module as well as the system features that are FIPS compliant and those that are not.

AltaVault offers end-to-end security for data at rest and in flight using FIPS 140-2 level 1-compliant encryption with the NetApp Cryptographic Security Module.

AltaVault requires all imported and generated keys sizes for RSA-based and DSA-based certificates to be 2048 bits or higher.

NetApp Cryptographic Security Module

The NetApp Cryptographic Security Module is the part of AltaVault software that separates the cryptography that is FIPS compliant from the rest of the AltaVault.

The NetApp Cryptographic Security Module is compatible with FIPS 140-2 Level 1 requirements.

The NetApp Cryptographic Security Module appears as the validated cryptographic module on the NIST vendor page instead of a specific AltaVault. The NIST vendor page is available at this URL:

<http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140val-all.htm>

Note: Throughout this guide, *FIPS-mode* and *FIPS-compliance* refers to use of the NetApp Cryptographic Security Module.

Compliant FIPS cryptography features

The following features use FIPS-compliant cryptography:

- Web interface (Apache Web server)
- Local user passwords and local authentication using SHA256-based or SHA512-based hash
- Image integrity checks for AltaVault OS
- File transfers
- NTP with SHA authentication
- Secure vault
- SNMP except if SNMP user passwords are configured with MD5 or DES protocols
- SSH with approved ciphers
- SSL optimization
- AltaVault Storage Optimization Service
- AltaVault data replication
- Domain-join feature in the AltaVault

Noncompliant FIPS cryptography features

The following features are not FIPS compliant. The system does not prevent you from using these features, but it does warn you that they are not FIPS compliant. You need to ensure that the system is configured in FIPS mode and uses only FIPS-compliant features to achieve full compliance.

Features Depending on NTLM or Kerberos Domain Authentication

- SMB signing

Features That Use Cryptographic Libraries Outside the NetApp Cryptographic Security Module Boundary

- Kerberos
- SSH with unapproved ciphers

Features with Protocol Specifications That Can Use Noncompliant Hash Algorithms

- Some cloud providers (that you use with the AltaVault) are not fully supported in FIPS mode
 - AT&T Synaptic and EMC Atmos clouds are supported in FIPS mode only in service versions 2.1 and later. Older versions use noncompliant hash algorithms.
 - Not all cloud providers use FIPS 140-2 validated cryptography. It is your responsibility to ensure that the configured cloud provider meets regulatory requirements.
- Local user passwords and local authentication with MD5-based hash
- NTP with MD5 authentication
- RADIUS
- SNMP with users configured with MD5 or DES protocols
- TACACS+

Configuring AltaVault for FIPS compliance

To achieve FIPS compliance on an AltaVault, configure the system to run in FIPS operation mode and adjust the configuration of any features that are not FIPS compliant.

With FIPS mode enabled, the system monitors configuration changes and provides warnings if you configure a feature to be noncompliant with FIPS. These warning messages appear when you try to change a configuration setting to an unsupported option. You can also view these warnings using the `show fips status` command.

Configuring AltaVault appliances for FIPS-compliant cryptography

This section includes the following information:

- [“Enabling FIPS mode” on page 114](#)
- [“Verifying that your system uses FIPS-compliant encryption” on page 114](#)
- [“Working with features to maintain FIPS compliance” on page 115](#)
- [“Account passwords” on page 115](#)
- [“Cipher requirements” on page 116](#)
- [“Key size requirements” on page 116](#)
- [“NTP” on page 117](#)
- [“RADIUS and TACACS+” on page 117](#)
- [“SNMP” on page 117](#)

- [“SSH” on page 117](#)
- [“Telnet server” on page 118](#)
- [“Web proxy” on page 118](#)
- [“Verifying that file transfers operate in FIPS mode” on page 119](#)
- [“Verifying that NTP operates in FIPS mode” on page 120](#)
- [“Verifying that secure vault operates in FIPS mode” on page 120](#)
- [“Verifying that SNMP operates in FIPS mode” on page 120](#)
- [“Verifying that the web interface operates in FIPS mode” on page 120](#)

Enabling FIPS mode

FIPS mode ensures the system uses only FIPS-compliant encryption algorithms.

Note: Before you can enable FIPS mode, you need to ensure passwords for user accounts use FIPS-compliant encryption. For details, see [“Account passwords” on page 115](#).

Note: Prior to enabling FIPS mode on the AltaVault appliance, remove previously configured local SMB users. FIPS mode does not support local SMB users.

To enable FIPS mode

1. Connect to the CLI.
2. Enter configuration mode, enable FIPS mode, and restart the system.

At the system prompt, enter the following set of commands:

```
amnesiac > enable
amnesiac # configure terminal
amnesiac (config) # fips enable
You must save the configuration and reload the system to enable FIPS mode.
amnesiac (config) # write memory
amnesiac (config) # reload
Rebooting...
amnesiac (config) #
```

Restarting the system applies the FIPS-specific settings on supported features and data.

3. Run the `show fips status` command to view compliance status and make any required configuration changes.

Verifying that your system uses FIPS-compliant encryption

To verify that your system is FIPS compliant

1. Connect to the CLI.

2. Enter the following set of commands:

```
amnesiac > enable
amnesiac # configure terminal
amnesiac (config) # show fips status

FIPS Mode: Enabled
```

The output indicates if FIPS mode is enabled and displays any warnings for features that affect FIPS compliance. If no warnings appear and FIPS mode is enabled, your system is FIPS compliant. If warnings appear, you need to make configuration changes to achieve full compliance.

You cannot review FIPS compliance from the Management Console; however, if you attempt to configure features that affect FIPS compliance through the Management Console when in FIPS mode, the Web interface produces an error message warning you of the conflict.

Working with features to maintain FIPS compliance

It is the responsibility of the system administrator of AltaVault appliance to ensure the system is FIPS compliant. Not all features can be operated in a FIPS-compliant manner and they need to be disabled when in FIPS mode. Some features can be operated in a FIPS-compliant manner by following noted guidance and other features prevent you from entering into FIPS mode.

The system generates a warning message if you configure non-compliant features. You can view the warning messages with the `show fips status` command.

The following sections describe configurable features that can affect FIPS compliance and describe how to resolve the warnings. The system does not prevent you from using noncompliant features in FIPS mode, but the system does warn you that they are not FIPS compliant. (The exception is account passwords; you cannot enable FIPS mode if all account passwords do not use compliant encryption).

Note: The commands in the following sections are configuration commands. You need to run these from configuration mode in the CLI. For detailed information, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Reference Guide*.

Account passwords

FIPS compliance requires that passwords for user accounts are encrypted using an SHA256-based or SHA512-based hash.

In systems with RCSM, SHA512 is the default hash when creating and updating a user password. However, previous releases used MD5 encryption. So, when you upgrade to a software release supporting FIPS mode from a release with MD5-based passwords, the MD5 passwords remain in the configuration.

If you attempt to enter FIPS mode on a system with accounts that have MD5 passwords, you see the following error:

```
amnesiac (config) # fips enable
% User admin has a password hashed using a non-FIPS-allowed hash.
The password(s) must be changed before FIPS mode can be enabled.
```

The error message identifies the user accounts that need to be updated; in this example, the admin account. You must update the noncompliant passwords or delete the accounts before you can enable FIPS mode. From the CLI, enter the `username <username> password <password>` command to change passwords.

Cipher requirements

You need to use the following cipher string when running in FIPS mode: `TLSv1.2:kRSA:!eNULL:!aNULL`

This requirement impacts SSL optimization, secure peering, and the Web interface security settings.

Note: It is advisable to allow TLS 1.1 or TLS 1.2.

To configure the cipher

- Enter the command `web ssl cipher`.

The format of the command is:

```
web ssl cipher TLSv1.2:kRSA:!eNULL:!aNULL
```

If you do not configure the required cipher string, the following message appears after enabling FIPS mode or with the `show fips status` command:

```
Web SSL ciphers must include the elements in TLSv1.2:kRSA:!eNULL:!aNULL and may optionally
delete ciphers.
```

This message also appears if you make any changes to the Web SSL cipher.

Key size requirements

FIPS specifies three techniques for the generation and verification of digital signatures for the protection of data: the Digital Signature Algorithm (DSA), the Elliptic Curve Digital Signature Algorithm (ECDSA), and the Rivest-Shamir-Adleman (RSA) Algorithm.

FIPS includes key size requirements when running in FIPS mode. All imported and generated keys need to be the following sizes:

- RSA-based and DSA-based certificates:
 - 2048 bits
 - 3072 bits
 - 4096 bits
- ECDSA certificates:
 - 224 bits and higher

These requirements apply to SSL optimization, SSL secure peering, and the Web interface.

Web user interface

You need to ensure imported and generated certificates for the Web interface adhere to FIPS size requirements and use only 2048-bit or higher key sizes.

You manage Web interface certificate keys using the `web ssl cert generate key-size *` command in the CLI and the Configure > Web Settings page in the Management Console. These methods always generate RSA based self-signed certificates.

In addition to self-signed certificates, you can import certificates using the `web ssl cert import-cert *` and `web ssl cert import-cert-key *` commands or the Configure > Web Settings page in the Management Console.

If you specify a key size that is not 2048-bit or higher with FIPS mode enabled, the system blocks the key generation and warns that the key size is not supported in FIPS mode.

NTP

NTP using either SHA authentication keys or no authentication keys is FIPS compliant. NTP using MD5 keys is not FIPS compliant.

If you configure an MD5 key for NTP using the following command, the system generates a warning message and the system will not be FIPS compliant:

```
amnesiac (config) # ntp authentication key <id> type MD5 secret <secret password>
```

To verify that NTP is running in FIPS mode, examine the system log when NTPD starts (this occurs whenever the NTP configuration is modified) and ensure that the NTPD entry sets FIPS mode:

```
Mar 18 15:49:57 amnesiac pm[4989]: [pm.NOTICE]: Launched ntpd with pid 27617
Mar 18 15:49:57 amnesiac ntpd[27617]: ntpd 4.2.6p4@1.2324-o Thu May 17 21:31:11 UTC 2012 (1)
...
Mar 18 15:49:57 amnesiac ntpd[27617]: FIPS_mode_set(1)
```

For more information about system logs, see [“Viewing system logs” on page 152](#).

RADIUS and TACACS+

The RADIUS and TACACS+ protocols are not FIPS compliant. These protocols use noncompliant hash algorithms. The system displays a warning message if you configure these features in FIPS mode.

The following commands generate a configuration warning in FIPS mode:

```
aaa accounting per-command default tacacs+
aaa authentication [console-login | login] default [radius | tacacs+]
aaa authorization per-command default tacacs+
```

SNMP

SNMP is FIPS compliant except if SNMP user passwords are configured with noncompliant hash algorithms. If you configure an SNMP user password with MD5 or DES protocols using the following command, the system generates a warning message and the system will not be FIPS compliant:

```
snmp-server user <username> password plain-text <password> [auth-protocol MD5 priv-protocol DES
priv-key plain-text <password>]
```

To verify that SNMP runs in FIPS mode, look for entries similar to the following in the system log when SNMP starts (this occurs whenever the SNMP configuration changes) and ensure that FIPS mode is set:

```
Mar 18 16:05:10 amnesiac pm[4989]: [pm.NOTICE]: Launched snmpd with pid 31709
Mar 18 16:05:10 amnesiac snmpd[31709]: FIPS_mode_set(1)
...
Mar 18 16:05:10 amnesiac snmpd[31709]: NET-SNMP version 5.3.1
```

For more information about system logs, see [“Viewing system logs” on page 152](#).

SSH

SSH requires the use one of the following ciphers to run in FIPS mode:

- aes128-cbc
- aes192-cbc
- aes256-cbc
- aes128-ctr
- aes192-ctr
- aes256-ctr

Configuring any other ciphers displays a warning message and the system will not be FIPS compliant.

Note: The default ciphers for SSH are aes128-cbc, aes192-cbc, and aes256-cbc. These ciphers are FIPS compliant.

- You can configure SSH ciphers with the following command:

```
amnesiac (config) # ssh server allowed-ciphers aes128-cbc,aes192-cbc,aes256-cbc,aes128-ctr,aes192-ctr,aes256-ctr
amnesiac (config) # write memory
```

- To verify your SSH settings, enter the following command:

```
amnesiac (config) # show ssh server allowed-ciphers
SSH server allowed ciphers:
-----
aes128-cbc
aes192-cbc
aes256-cbc
aes128-ctr
aes192-ctr
aes256-ctr
```

- To verify that SSH is running in FIPS mode, look for entries similar to the following in the syslog when a user logs in:

```
Mar 18 15:00:30 amnesiac sshd: FIPS_mode_set(1)
Mar 18 15:00:30 amnesiac sshd[14594]: FIPS mode initialized
```

Telnet server

Telnet functionality is not FIPS compliant. Enabling this feature triggers a configuration warning in FIPS mode.

Telnet must be disabled. If Telnet is enabled, an error message appears if you try to enable FIPS mode. If FIPS mode is enabled, the system prevents you from enabling Telnet and provides an error message.

- To disable this feature, use the following commands:

```
amnesiac (config) # no telnet-server enable
amnesiac (config) # no telnet-server permit-admin
amnesiac (config) # write memory
```

- To verify your settings, enter the following command:

```
amnesiac (config) # show telnet-server
Telnet server enabled: no
```

Web proxy

Web proxy functionality for licensing is not FIPS compliant.

- To disable this feature, enter the following commands:

```
amnesiac (config) # no web proxy host <ip-address>
amnesiac (config) # write memory
```

Enabling this feature triggers a configuration warning in FIPS mode.

Disabling FIPS mode

If you no longer want to use FIPS mode, you can turn off this feature.

To disable FIPS mode

1. Connect to the CLI.
2. Enter configuration mode, disable FIPS mode, and restart the system.

At the system prompt, enter the following set of commands:

```
amnesiac > enable
amnesiac # configure terminal
amnesiac (config) # no fips enable
You must save the configuration and reload the system to disable FIPS mode.
amnesiac (config) # write memory
amnesiac (config) # reload
Rebooting...
amnesiac (config) #
```

When you disable FIPS mode, the system is less restrictive and FIPS compliance configuration warnings no longer appear. Any configuration changes that you made while in FIPS mode (such as disabling certain features or setting specific ciphers) are not modified.

Verifying FIPS mode in system logs

You can review the system logs to ensure that features use FIPS mode. Features that run in FIPS mode have entries in the system log that include **FIPS_mode_set(1)**.

The following sections show several examples.

For more information about system logs, see [“Viewing logs” on page 152](#).

Verifying that file transfers operate in FIPS mode

File transfers, such as configuration fetch, run in FIPS mode and are FIPS compliant. To verify, look for file transfer entries in the syslog when initiating a file download. Ensure these entries have **FIPS_mode_set(1)**.

For example:

```
Mar 18 16:28:34 amnesiac curl: FIPS_mode_set(1)
```

Verifying that NTP operates in FIPS mode

To verify that NTP is running in FIPS mode, examine the system log when NTPD starts (this occurs whenever the NTP configuration is modified) and ensure that the NTPD entry sets FIPS mode:

```
Mar 18 15:49:57 amnesiac pm[4989]: [pm.NOTICE]: Launched ntpd with pid 27617
Mar 18 15:49:57 amnesiac ntpd[27617]: ntpd 4.2.6p4@1.2324-o Thu May 17 21:31:11 UTC 2012 (1)
...
Mar 18 15:49:57 amnesiac ntpd[27617]: FIPS_mode_set(1)
```

Verifying that secure vault operates in FIPS mode

The secure vault contains sensitive information from your AltaVault appliance configuration, including SSL private keys and the data store encryption key. These configuration settings are encrypted on the disk using AES 256-bit encryption.

The secure vault always runs in FIPS mode. To verify, look for the following in the system log at startup:

```
Mar 11 18:28:06 amnesiac encfs: FIPS_mode_set(1)
```

Verifying that SNMP operates in FIPS mode

To verify that SNMP is running in FIPS mode, look for entries similar to the following in the system log when SNMP starts (this occurs whenever the SNMP configuration changes) and ensure that FIPS mode is set:

```
Mar 18 16:05:10 amnesiac pm[4989]: [pm.NOTICE]: Launched snmpd with pid 31709
Mar 18 16:05:10 amnesiac snmpd[31709]: FIPS_mode_set(1)
...
Mar 18 16:05:10 amnesiac snmpd[31709]: NET-SNMP version 5.3.1
```

Verifying that the web interface operates in FIPS mode

The Apache web server for the AltaVault appliance always runs in FIPS mode.

To verify that the web server is in FIPS mode, look for entries similar to the following in the system log:

```
Mar 18 16:22:11 amnesiac httpd: FIPS_mode_set(1)
Mar 18 16:22:11 amnesiac httpd: [Mon Mar 18 16:22:11 2013] [notice] Operating in SSL FIPS mode
Mar 18 16:22:11 amnesiac httpd: [Mon Mar 18 16:22:11 2013] [notice] Init: Skipping generating
temporary 512 bit RSA private key in FIPS mode
Mar 18 16:22:11 amnesiac httpd: [Mon Mar 18 16:22:11 2013] [notice] Init: Skipping generating
temporary 512 bit DH parameters in FIPS mode
Mar 18 16:22:11 amnesiac httpd: [Mon Mar 18 16:22:11 2013] [notice] Init: Skipping generating
temporary 512 bit RSA private key in FIPS mode
Mar 18 16:22:11 amnesiac httpd: [Mon Mar 18 16:22:11 2013] [notice] Init: Skipping generating
temporary 512 bit DH parameters in FIPS mode
Mar 18 16:22:11 amnesiac httpd: [Mon Mar 18 16:22:11 2013] [notice] Apache/2.2.23 (Unix) mod_ssl/
2.2.23 OpenSSL/1.0.1c-fips configured -- resuming normal operations
```

FIPS CLI

For information about FIPS CLI commands, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Reference Guide*.

CHAPTER 9 Managing the AltaVault appliance

This chapter includes the following sections:

- [“Starting and stopping the AltaVault appliance” on page 121](#)
- [“Configuring scheduled jobs” on page 122](#)
- [“Managing licenses” on page 123](#)
- [“Activating support for AltaVault cloud-based appliances” on page 125](#)
- [“Upgrading your software” on page 126](#)
- [“Rebooting and shutting down AltaVault appliance” on page 127](#)
- [“Viewing the current user settings” on page 127](#)
- [“Managing configuration files” on page 128](#)

Starting and stopping the AltaVault appliance

The AltaVault Storage Optimization Service is a daemon that executes in the background, performing operations when required.

You can start, stop, and restart AltaVault Storage Optimization Service in the Maintenance > Service page. You can also use this page to reset the service alarm after it has been triggered.

Restarting AltaVault service disrupts ingest sessions established with the AltaVault. This has the following impact:

- All shares (SMB, NFS, OST, and SnapMirror) will be unavailable and cloud connections will fail until the service initializes again.
- Active backup jobs targeted at the appliance will fail and will need to be restarted after the Storage Optimization Service is running again.

Note: Prior to stopping or restarting Storage Optimization Service, quiesce all backup applications.

To start, stop, or restart the Storage Optimization Service

1. Choose Maintenance > Service.

2. Click **Stop**, **Start**, or **Restart**.

If Status displays **Replaying**, the Storage Optimization Service has been terminated during backup replication. During this replay process, the AltaVault is verifying data consistency from its transaction logs and, if needed, restoring data slabs (file chunks and references) from the cloud to the local AltaVault. This might take a long time depending on the available WAN bandwidth and data to be restored.

If the replay process is very slow or seems to be stuck, contact NetApp Support at <https://mysupport.netapp.com>.

3. To verify that the service is running again, go to Home > Optimization Service. You should see that Service is running and Status is ready.

To check the service status from the command line interface, run the `show service` command.

Configuring scheduled jobs

Jobs are commands that are scheduled to execute at a time you specify.

You can view completed, pending, and inactive jobs, as well as jobs that were not completed because of an error, in the Scheduled Jobs page. You can also delete a job, change its status, or modify its properties.

You can use the Management Console to:

- Schedule a software upgrade.
- Schedule multiple TCP trace dumps.

To schedule all other jobs, you must use the NetApp CLI.

For details about scheduling jobs using the CLI, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Reference Guide*.

To view scheduled jobs

1. Choose Reports > View Scheduled Jobs.
2. Click the Job ID number to display details about the job.

3. Under Details for Job <#>, complete the configuration as described in this table.

Control	Description
ID	Specify the job ID number.
Name	Specify a name for the job.
Comment	Add any comments.
Interval	Specify the recurrence interval in seconds
Executes On	Specify the date on which the job runs.
Created	Specify the date and time at which the job was created.
Last Run	Specify the last date and time at which the job was executed.
Enable/Disable Job	Enables the job.
Apply Changes	Applies the changes to the current configuration.
This Job	s the job.
Execute Now	Runs the job.
Remove Selected Jobs	Select the check box next to the name and click Remove Selected Jobs .

4. Click **Apply Changes** to save your settings permanently.

Managing licenses

There are four types of license keys (AVA-v2, AVA-v8, AVA-v16, AVA-v32) and are only required for the AltaVault virtual models. CPU, memory, disk, and cloud capacity enforcement is based on the license applied.

Note: You can install only one license key at a time.

To add or remove a license requires the optimization service to be shutdown in order to allow it to perform capacity and usage checks at the time of starting up. Choose Maintenance > Service to shut down the service.

If an attempt is made to configure license while service is running the following message will appear in the CLI:

```
Cannot configure license while optimization service is running
```

Managing licenses includes the following sections:

- [“Managing unlicensed AltaVault appliances” on page 124](#)
- [“Managing licenses using the command-line” on page 124](#)
- [“Managing licenses using the Management Console” on page 125](#)
- [“License limits” on page 125](#)
- [“Model upgrades on the virtual AltaVault appliances” on page 125](#)

Managing unlicensed AltaVault appliances

When an AltaVault virtual appliance is deployed and brought up for the first time, it enters a 90-day evaluation period and the appliance runs unlicensed. During this evaluation period, the full functionality of AltaVault is available for use for proof-of-concept trial and evaluation. Once the 90-day evaluation expires, the virtual machine will only allow read-only access to the data ingested into the appliance and new data will not be ingested. Once a valid license is installed, the full functionality of the virtual machine returns.

By default, AltaVault appliances are deployed as AVA-v2 models, but can be modified to any other model (AVA-v8, AVA-v16, AVA-v32) by using the command:

```
CLI (config)> license virtual-model <model type>
```

For example, to modify the default AVA-v2 to an AVA-v8

```
CLI (config)> license virtual-model v8
```

Reboot the appliance to change the model. The requirements for the new model will be checked and enforced after the reboot occurs.

Note: Upgrading a virtual AltaVault 4.1 or earlier model that is unlicensed will result in the model being reset to an AVA-v2 model. Use the above command to change the model back to its previous configuration.

Managing licenses using the command-line

You can install, delete, or view an existing licenses using the command-line.

To install a license

- Use the following command to install a license:

```
CLI> license install <license-key>
```

If a license already exists, the `license install <license-key>` command fails and it warns the user of any unintended license installs of a lower capacity.

To delete a license

- Use the following command to delete a license that already exists:

```
CLI> no license install <license-key>
```

This enable users to switch a license that is already installed on a system with a different capacity.

To view existing licenses

- Use the following command to view existing licenses:

```
CLI> show licenses
Key: <license-key>
Capacity: <capacity in TB>
Version: <software version>
Type: Capacity"
Installation: <date/time>
Expiration: <date>
```

This command does not require the optimization service to be shutdown.

If there is no license, the following command displays:

```
No license found
```

Managing licenses using the Management Console

You can add or remove a license using the Management Console.

1. Choose **Maintenance > Licenses**.
2. Select **Add a New License**.
3. Enter or paste the license into the text box, and click **Add**.
4. To remove a license, check the license you want to remove and click **Remove Selected**.

License limits

The AltaVault license limit controls the amount of data you can back up to the cloud from the AltaVault. The license limit varies for each AltaVault model. For details, see the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances*.

The AltaVault lets you exceed the license limit by 10 percent before triggering the over_capacity alarm. When you exceed 110 percent of the limit, the AltaVault pauses all replication activity and does not upload any additional data to the cloud until you increase the licensed capacity or reclaim the space by deleting files.

Restore operations are not affected by the over_capacity alarm and all files, whether stored locally or in the cloud. Any new files written to the AltaVault are queued pending replication to the cloud during the over_capacity alarm.

You can find the pending data to be replicated in the Replication Optimization report by choosing **Reports > Replication**.

Model upgrades on the virtual AltaVault appliances

Upgrading an existing licensed AltaVault virtual appliance to a higher model requires removing the old license key and installing a new license key. Only one license key can be present on the appliance at any time.

Refer to the sections “[Managing licenses using the command-line](#)” on page 124 and “[Managing licenses using the Management Console](#)” on page 125 on applying the new license.

After the new license is applied, halt the AltaVault virtual appliance (using the **Maintenance > Reboot/Shutdown** page) and provision additional CPU, memory, and disk as required by the model corresponding to the new license. Refer to the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances* for details on model requirements.

Activating support for AltaVault cloud-based appliances

NetApp technical support is available for AltaVault cloud appliances starting with version 4.4.1 and higher. You must register the serial number with NetApp for each appliance instance to activate support entitlement.

If you are a new customer to NetApp and do not have a NetApp Support Site (NSS) account, or you are having issues with support activation through this product, go to <http://register.netapp.com> to activate support for the serial number listed in the Support Activation page. Otherwise, follow the steps below to automatically activate support for your AltaVault instance.

1. Select **Maintenance > Support Activation**.

The system displays the activation page along with the Account ID and Instance ID from Amazon or Azure. AltaVault generates a serial number for each instance based on system information.

2. Start the serial number registration process with NetApp by providing your NetApp Support Site (NSS) SSO username and password in the UI. This NSS account must be a customer-level account and not temp or guest status.
3. Click **Apply** to complete the activation of the appliance instance with NetApp.

The AltaVault appliance will validate your NSS account credentials, create the serial number in NetApp support business systems and associate it to the account tied to the NSS account used. The activation process usually takes only a few seconds.

The AltaVault UI will display the Support Activation Status page indicating the status of the request. It will auto-refresh the page every few seconds. In the event the status does not show activated after a few minutes, there is an option to try again.

Once activation is successful, the Support Summary page appears indicating that support has been activated for the appliance with the specified serial number. The owner of that account will receive an email with the status of the support registration.

Deactivating support

When the decision is made to terminate the instance from the AWS or Azure console, please deactivate support with the following steps:

1. Select **Maintenance > Support Activation**.

The system displays a summary page indicating the support status.

2. Click **Deactivate**.

Note: Once a serial number is deactivated, the appliance will not let you reactivate it. Contact NetApp technical support for assistance reactivating the serial number in the event it was done unintentionally.

Upgrading your software

You can upgrade your AltaVault software in the Maintenance > Software Upgrade page. The bottom of the page displays the AltaVault software version history, which includes the version number and the software installation date.

Review the following installation notes before proceeding with the upgrade:

- Refer to the release notes for release specific upgrade information prior to upgrading software.
- Quiesce all backup applications prior to upgrading AltaVault software.
- AltaVault software upgrades cannot be downgraded to an older software version once applied.

To upgrade software versions

1. Choose Maintenance > Service and click **Stop** to shut down the Storage Optimization Service.
2. Choose Maintenance > Software Upgrade.

- Under Install Upgrade, complete the configuration as described in this table.

Control	Description
From URL	Select this option and type the URL. If you specify a URL in the URL text box, the image is uploaded and installed and the system is rebooted at the time you specify. Use the image.img file for upgrades, found on the NetApp support site.
From Local File	Select this option and type the path, or click Browse to go to the local file directory. If you specify a file to upload in the From Local File text box, the image is uploaded immediately; however, the image is installed and the system is rebooted at the time you specify.
Schedule Upgrade for Later	Schedules the upgrade process. Specify the date and time to run the upgrade: <ul style="list-style-type: none">• Date and Time - Use the following format: YYYY/MM/DD, HH:MM:SS.
Install	Installs the software upgrade on your system.

- Choose Maintenance > Reboot/Shutdown, and click Reboot.

After successfully upgrading the AltaVault software, the appliance does not support start up from the alternate boot partition containing the previous version of software. The alternate boot partition exists to applied software upgrades without affecting the current running AltaVault software version. If the software upgrade process fails or experiences an error, then the current running version of AltaVault software is not affected.

Rebooting and shutting down AltaVault appliance

You can reboot or shut down the system in the Maintenance > Reboot/Shutdown page. The reboot and shutdown processes can take a few minutes and you will need to remount all NFS shares.

To restart the system, you must manually turn on the AltaVault.

Rebooting the AltaVault physical appliances does not reboot the attached add-on shelves. You do not need to power off the add-on shelves even if the AltaVault is rebooting. The add-on shelves should always be powered on before the AltaVault boots and starts running the Storage Optimization Service.

To reboot or shut down the system

- Choose Maintenance > Reboot/Shutdown.
- Click **Reboot**. After you click Reboot, you are logged out of the system and it is rebooted.
- Click **Shutdown** to shut down the system. After you click Shutdown, the system is turned off. To restart the system, you must manually turn on the AltaVault.

Viewing the current user settings

You can change your login password, reset your user preferences, and view your role-based permissions in the My Account page.

To display system permissions

1. Choose **Configure > My Account**.
2. To change the password of the currently logged in user, complete the configuration as described in this table.

Control	Description
Change Password	Select the check box to change the password for the currently logged in user.
New Password/Confirm New Password	Type a password in the text box. Retype the password in the Confirm New Password text box. The password cannot be "password."
Apply	Click to confirm the password change.

To restore user preferences

1. To restore the default user preferences for the currently logged in user, click the **Restore Defaults** button.

The user preferences are used to remember the state of the Management Console across sessions on a per-user basis. They do not affect the configuration of the appliance.

Managing configuration files

The admin account, and users with admin role privileges, can save and activate configurations in the **Configure > Configurations** page.

Each AltaVault has an active, running configuration and a written, saved configuration.

The default configurations are as follows:

Configuration	Description
initial	This is the initial configuration.
initial.bak	This is the backup of the initial configuration.

To manage configurations

1. Choose **Configure > Configurations**.
2. Under **Current Configuration: <filename>**, view or save the configuration as described in this table.

Control	Description
Current Configuration: <configuration name>	View Running Config - Displays the running configuration settings in a new browser window. Save - Saves settings that have been applied to the running configuration. Revert - Reverts your settings to the running configuration.
Save Current Configuration	Specify a new filename to save settings that have been applied to the running configuration as a new file, and then click Save As .

3. Click **Activate** to change the active configuration to the configuration you select.

4. Restart the AltaVault service on the Maintenance > Service page.

Click the configuration name to display the configuration settings in a new browser window.

CHAPTER 10 Viewing reports and logs

This chapter includes the following sections:

- [“About reports” on page 132](#)
- [“Viewing the storage optimization report” on page 135](#)
- [“Viewing the front-end throughput report” on page 136](#)
- [“Viewing the back-end throughput report” on page 137](#)
- [“Viewing the eviction report” on page 138](#)
- [“Viewing the replication report” on page 139](#)
- [“Viewing the cloud operations report” on page 140](#)
- [“Viewing schedule reports” on page 141](#)
- [“Viewing per share utilization reports” on page 142](#)
- [“Viewing the alarm status report” on page 143](#)
- [“Viewing the CPU utilization report” on page 146](#)
- [“Viewing the memory paging report” on page 147](#)
- [“Viewing the interface counters report” on page 148](#)
- [“Viewing the disk throughput report” on page 149](#)
- [“Viewing the disk IOPS report” on page 150](#)
- [“Viewing the disk utilization report” on page 151](#)
- [“Viewing logs” on page 152](#)
- [“Downloading log files” on page 154](#)
- [“Generating system dumps” on page 155](#)
- [“Viewing process dumps” on page 156](#)
- [“Capturing and uploading TCP dumps” on page 156](#)
- [“Viewing the appliance monitoring report” on page 162](#)
- [“Viewing the shelf details” on page 164](#)
- [“Viewing the storage RAID group” on page 165](#)
- [“Viewing offline file system check page” on page 165](#)

- “Viewing online file system check page” on page 166
- “Viewing the verify tool diagnostics” on page 167

About reports

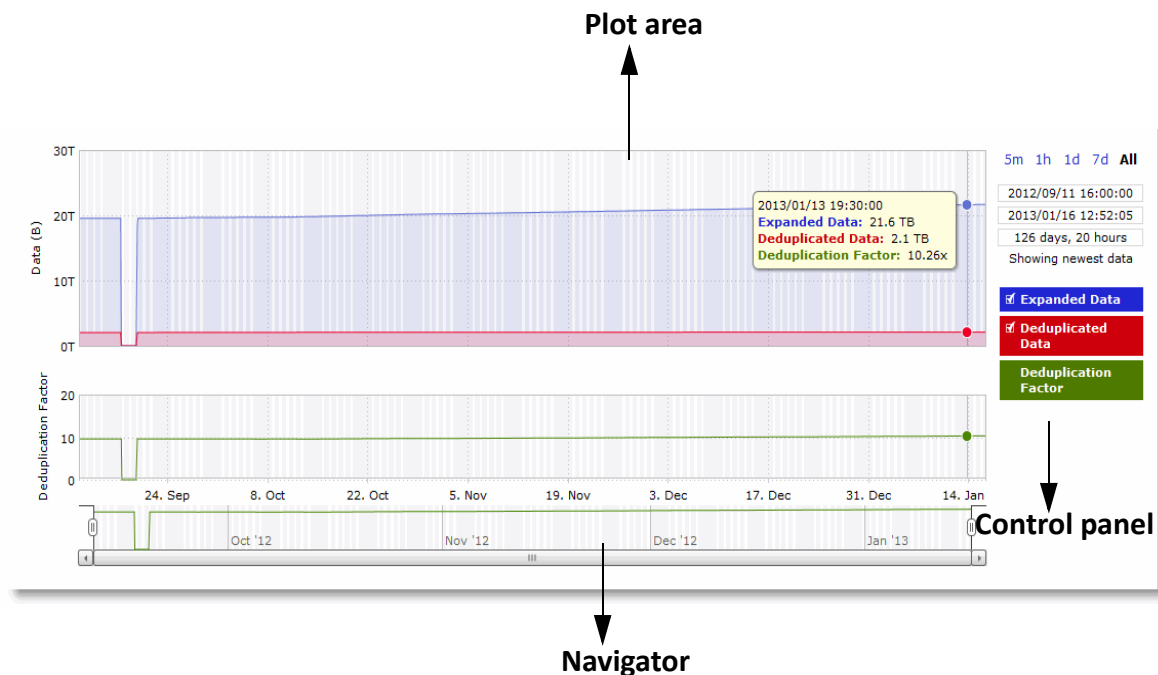
All of the time-series reports in AltaVault have a look and feel that is clear, interactive, and easy to navigate. This section describes the new report format in detail.

Navigating the report layout

The AltaVault report format not only makes data easily accessible, but also enhances your ability to explore data in context. An example of a typical report is shown in [Figure 10-1](#), with the key areas labeled. For details about individual reports, see the report description.

Statistics used in generating reports are retained for one year.

Figure 10-1. Report layout



The report sections, going counter-clockwise from the top-left of the report window shown in [Figure 10-1](#), are:

- Plot area
- Navigator
- Control panel

The summary tables, which appeared under the graphs in the legacy report style, have been removed. The summary statistics appear in the control panel legend.

Plot area

The plot area is where data visualization occurs. Reports can display either a single-pane or dual-pane layout. In a dual-pane layout, both panes remain synchronized with respect to the x-axis. Each pane is capable of having two y-axes (a primary on the left and a secondary on the right).

The reports present the majority of data series as simple line series graphs, but some reports display area series graphs where appropriate. The types of area series graphs are:

- **Layered series**, which appear on top of each other. These are identified by transparent colors.
- **Stacked area series**, which appear on top of each other in the y direction. The AltaVault uses stacked area graphs to depict an aggregate broken down by its constituent parts. In this type of graph, each series is a mutually exclusive partition of some aggregate data set, identified by opaque colors. A stacked series is appropriate when the sum of all the series is meaningful.

The report format enables you to hover over individual data points to view the detailed information.

To view the timestamp and value of each data series at that time

- Place the mouse pointer over the plot area.

A tool tip displays the time stamp and the value of each data series at that time. The plot area colors the series names appropriately and the data values have their associated units.

The plot area also displays subtle shading to denote work hours (white background) and nonwork hours (gray background). The AltaVault defines work hours as 8:00 AM to 5:00 PM on weekdays are not configurable.

To pan the plot area

1. Place the mouse pointer over the plot area, and then click and hold the left mouse button.
2. Move the mouse left or right to slide the window of visible data left or right.

The navigator reflects the changing chart window as do the associated controls in the control panel.

Navigator

Directly above the scroll bar is the navigator, which shows a much smaller and simpler display of the data in the plot area. The navigator displays only one data series.

Use the navigator to navigate the entire range of chart data. The scroll bar at the bottom displays which portion of the total data range is currently displayed in the plot area.

The navigator display can appear very different than the plot area display when an interesting or eye-catching series in the plot area is not the series in the navigator.

To resize the current chart window

- Move the handles on either side of the chart window in the navigator.

The charts have a minimum chart window size (currently five minutes), so if you try to resize the chart window to something smaller, the chart window springs back to the minimum size.

You can also click the data display portion of the navigator (not the scroll bar) and the chart window moves to wherever you clicked.

Control panel

Use the control panel to control how much data the chart displays, to see chart properties, and to view or hide the summary statistics.

To change the chart interval

- Click a link: 5m (five minutes), 1h (one hour), 1d (one day), 1w (seven days), or All (last 30 days).

If the current size of the chart window matches any of the links, the link appears in bold black text; the system ignores any clicks on that link. If the time duration represented by any of the links is greater than the total data range of the chart, those links are dimmed and unavailable.

- **Chart window controls** - More window-related controls appear below the chart window interval links. These controls offer more precise control of the window and also display various window properties. From top to bottom, the controls are:

- Text field containing the left edge (starting time) of the chart window.
- Text field containing the right edge (ending time) of the chart window.
- Text field containing the chart window interval. The chart window interval in this text field is not always exactly correct, but it is correct to two units (with the units being days, hours, minutes, and seconds). For example, if the chart window interval is exactly two days, three hours, four minutes, and five seconds, this text field displays 2 days, 3 hours.
- Link or static text that represents the chart window state of *attachment* to the end of the chart. When the chart window is attached, the report replaces the link with the static text **Showing newest data**. When the chart is showing newest data, you can see new data points as the system adds them automatically to the chart every 10 seconds. This can be very powerful when you launch a new configuration and need to analyze its impact quickly. You cannot change the 10-second default.

When the chart window is not attached to the end of the chart, the report replaces the static text with a link that displays **Show newest data**. Click this link to slide the chart window to the end of the chart range of data and attach the window.

With all three text fields, if the focus leaves the field (either because you click outside the field or press the Tab key), the chart window updates immediately with the new value. Pressing Enter while in one of these fields removes the focus.

Custom controls

Below the chart window controls is an optional section of custom, report-specific controls. The custom controls vary for each report. In [Figure 10-1](#), the Bandwidth Optimization report displays Port and Direction drop-down lists.

When you change the value of a custom control, the system sends a new request for data to the server. During this time, the control panel is unavailable and an updating message appears on the chart. When the report receives a response, the system replaces the chart, populates it with the new data, and makes the control panel available again.

Chart legend

The chart legend correlates the data series names with line colors and contains a few other features.

You can hide and show individual data series. When a white check box icon appears next to the data series name, you can hide the series from the plot area.

To hide individual series from the plot area

- Clear the check box next to the data series name.

To display individual series in the plot area

- Select the check box next to the data series name.

You cannot toggle the visibility of all series. In [Figure 10-1](#), you can hide the LAN Throughput and WAN Throughput series, but you cannot hide the Data Reduction series.

The legend also displays statistics. Each report defines any number of statistics for any of the data series in the chart. The system bases the statistics computation on the subset of each data series that is visible in the current chart window. The statistics display changes as the chart window changes. The reports also support nonseries statistics (for example, composite statistics that incorporate the data from multiple data series); these statistics (not pictured in [Figure 10-1](#)) appear at the very bottom of the legend, below all the series.

Setting user preferences

You can change report default settings and chart windows to match your preferred style. The system saves the setting on the server on a per-user basis. A message appears at the top of each page when multiple user are logged in explaining the user preferences might be overwritten.

Viewing the storage optimization report

The Storage Optimization report is the same report that appears on the home page of the AltaVault Management Console. It summarizes the percentage of the data storage optimized within the time period specified. It includes the following statistics that describe the storage optimization activity for the time period you specify.

Field	Description
Expanded Data	Data that has been backed up locally by the AltaVault.
Deduplicated Data	Current data de-duplicated by the AltaVault and sent to the cloud.
Deduplication Factor	Ratio of the expanded data and total optimized data. The total optimized data includes both deduplication and compression savings.

The Storage Optimization report answers the following questions:

- What is the latest user data stored?
- What is the latest amount of deduplicated data stored?
- What is the deduplication factor (ratio of user data and deduplicated data)?

To view the Storage Optimization report

1. Choose Reports > Storage Optimization.
2. Use the controls to customize the report as described in this table.

Control	Description
Period	Select the time period from the drop-down list. For Custom, enter the Start Time and End Time and click Go . Use the following format: YYYY/MM/DD HH:MM:SS.
Refresh	Select a refresh rate from the drop-down list. To turn refresh off, click Off .
Go	Displays the report.

Viewing the front-end throughput report

The Front-End Throughput report summarizes the front-end (SMB/NFS/OST/SnapMirror) data read into and written out of the AltaVault within the time period specified. The front-end data is sent from your local servers or backup software to AltaVault.

This report includes the following statistics that describe the storage optimization activity for the time period you specify.

Field	Description
Front-End In	
Average	Average front-end (SMB/NFS/OST/SnapMirror) data that the backup server writes to the AltaVault.
Total Data	Total data transferred from the backup server to the AltaVault
Front-End Out	
Average	Average front-end (SMB/NFS/OST/SnapMirror) data read by the backup server.
Total Data	Total data transferred from the front-end SMB, NFS, OST, or SnapMirror system to the backup server.

The Front-End Throughput report answers the following questions:

- What is the average front-end (SMB/NFS/OST/SnapMirror) data that the backup server writes to the AltaVault?
- What is the total amount of data transferred from the backup server to the AltaVault?
- What is the average front-end (SMB/NFS/OST/SnapMirror) data read by the backup server?
- What was the peak amount of data transmitted?

To view the Front-End Throughput Optimization report

1. Choose Reports > Front-End Throughput.

2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.

Viewing the back-end throughput report

The Back-End Throughput report summarizes the back-end cloud data read by and written out of the AltaVault within the time period specified. The back-end data is data that the AltaVault sends to the cloud.

This report includes the following statistics that describe the storage optimization activity for the time period you specify.

Field	Description
Back-End In	
Average	Average back-end data read from the cloud by the AltaVault
Total Data	Total data transferred from the cloud to the AltaVault
Back-End Out	
Average	Average back-end data that the AltaVault writes to the cloud
Total Data	Total data transferred from the AltaVault to the cloud

The Back-End Throughput report answers the following questions:

- What is the average back-end data read from the cloud by the AltaVault?
- What is the total amount of back-end data transferred from the cloud to the AltaVault?
- What is the average back-end data that the AltaVault writes to the cloud?
- What was the peak amount of data transmitted?

To view the Back-End Throughput Optimization report

1. Choose Reports > Back-End Throughput.

2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing the eviction report

The Eviction report summarizes the evicted bytes and the age of the data leaving the AltaVault local storage disk within the time period specified. The report first displays the available data cache which is the amount of free space available locally on the AltaVault data store. It also displays the amount of data in bytes that has been evicted from the local AltaVault data store to make space for the new deduplicated data. The second graph shows the average age of data evicted from the AltaVault.

The Eviction reports includes the following statistics that describe the eviction optimization activity for the time period you specify.

Field	Description
Evicted Data	Evicted bytes are the amount of data forced out from the local storage disk because the AltaVault needs to store new deduplicated data on the disk and disk space is insufficient.
Eviction Age	Average value of the data (calculated using the date and time it was created) leaving the AltaVault local storage disk.

The Eviction report answers the following questions:

- How much data is forced out from the AltaVault local disk?
- What is the average of the most recent data (calculated using the date and time it was created) leaving the AltaVault local storage disk?

To view the Eviction report

1. Choose Reports > Eviction.

2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing the replication report

When you back up data to the AltaVault, it writes data on to its disks while also replicating the data in the cloud. The Replication report shows how much data is currently waiting to be replicated to the cloud and the time and date the cloud was last synchronized.

The Replication report displays the following statistics.

Field	Description
Data Pending Replication	Amount of data that is currently waiting to be replicated to the cloud.
Time to complete replication	Estimated time for replication to complete.
Cloud Synchronized Until	Data ingested up to this time has been replicated to the cloud.

The Replication report answers the following questions:

- How many bytes to be replicated in the cloud are pending?
- What is the estimated time for replication to complete?
- How much data ingested until the specified time is replicated to the cloud?

To view the Replication report

1. Choose Reports > Replication.

2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing the cloud operations report

The Cloud Operations report provides statistics about the total put, head, post, get, and delete operations to and from the cloud.

The Cloud Operations report provides only the main operation statistics and does not cover all statistics. Do not use the Cloud Operations report to measure the number of cloud operations issued. Use the cloud provider statistics for cloud cost computations.

The Cloud Operations report displays the following statistics.

Field	Description
Put Operations	The number of HTTP operations to upload a file in the cloud.
Head Operations	The number of HTTP operations that return only the meta-information in the HTTP headers; they do not return a message-body in the response.
Post Operations	The number of HTTP post operations that have been communicated to the cloud.
Get Operations	The number of HTTP get operations performed to download files from the cloud.
Delete Operations	The number of files deleted in the cloud.

The Cloud Operations report answers the following questions:

- What are the total put operations?
- What are the total head operations?
- What are the total post operations?
- What are the total get operations?
- What are the total delete operations?

To view the Cloud Operations report

1. Choose Reports > Cloud Operations.
2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing schedule reports

The AltaVault emails you scheduled reports that contain information such as the hostname, model number, serial number, version, uptime, and statistics related to storage optimization, replicated data, disk storage allocation, and cloud storage allocation that display in the Schedule Reports page.

Use the Schedule Reports page to configure the frequency of report distribution.

The Schedule Reports answer the following questions:

- What is the hostname of the AltaVault?
- What is the model name of the AltaVault?
- What is the AltaVault serial number?
- What is the AltaVault software version?
- What is the date on which the report was sent?
- What is the AltaVault up time?
- What is the expanded data in the AltaVault?
- What is the AltaVault deduplicated data?
- What is the deduplication factor (ratio of the expanded data and the deduplicated data)?
- How much data is replicated to the cloud?
- What is the time to complete replication?
- What are the pending replication bytes?
- How much disk storage is used?
- How much disk storage is free?

- What is the total disk storage?
- How much cloud storage has been used?
- What is the total cloud storage?

To configure Schedule Reports

1. Choose **Configure > Schedule Reports**.
2. Under **Schedule AltaVault Reports via Email**, use the controls in the control panel to customize the report as described in this table.

Control	Description
Enable Scheduling	Select this check box for the AltaVault to email scheduled reports.
Weekly	Select this check box to schedule weekly reports and then specify the day of the week on which the AltaVault must email the report.
Monthly	Select this check box to schedule monthly reports and then specify the day of the month and time at which the AltaVault must email the report.

If you do not specify a weekly or monthly schedule, the reports will be sent daily at the selected time.

3. Click **Apply** to apply your changes to the running configuration.

Viewing per share utilization reports

The AltaVault emails you Share Utilization reports that contain information such as the hostname, model number, serial number, version, date, uptime, and evicted percentage.

The Share Utilization reports answer the following questions:

- What is the hostname of the AltaVault?
- What is the model name of the AltaVault?
- What is the AltaVault serial number?
- What is the AltaVault software version?
- What is the date on which the report was sent?
- What is the AltaVault up time?
- What is the evicted percentage?

To configure and view Share Utilization report

1. Choose **Configure > Email**.

The system sends an email to all email addresses on the **Configure > Email** page appearing under **Report Events via Email**. The email notifies users that the trigger has fired.

2. To generate share statistics, use the following CLI command:

```
share-stats generate
```

Viewing the alarm status report

The Alarm Status report takes a current snapshot of the system to provide a single report you can use to check for any issues with the AltaVault. The report examines the status of key system components. Use this report to gather preliminary system information before calling NetApp Support to troubleshoot an issue.

Alarm	Description
Admission Control	Indicates that AltaVault has entered the admission control state. Admission control limits the number of connections made to the AltaVault, so that you do not over-consume resources on your system.
Cloud Bucket Consistency	Indicates that there is data in the cloud although the AltaVault datastore is empty. Enable replication and recovery to ensure that the cloud storage is synchronized with the datastore.
Cloud Bucket Disparity	Indicates that the cloud bucket that the AltaVault is trying to connect to might be in use by another AltaVault. This prevents corruption of the files in the cloud.
Cloud Bucket Over Capacity	Indicates that the cloud bucket size is greater than the licensed amount of cloud capacity that you are using. You must upgrade to a higher-capacity license. Generates an AutoSupport message when AutoSupport is enabled.
CPU Utilization	Indicates that the system has reached the CPU threshold for one or more of the CPUs in the AltaVault. If the system has reached the CPU threshold, check and adjust your settings as appropriate.
Data Integrity Error	Indicates that there was an inconsistency in the data stored on disk. Generates an AutoSupport message when AutoSupport is enabled.
Datastore Eviction	Indicates that the system has detected an issue with datastore eviction. The alarm triggers when the appliance has started evicting data from the local disk cache when the age of the evicted data is relatively young. An AltaVault has disk space much smaller than the total addressable space on the cloud, and if disk space runs low, the appliance starts evicting data from disk that has not been used recently. This keeps only fresh and frequently accessed data in cache. The AltaVault keeps statistics about how old the evicting data is (this is the average evicted age). Usually, only old data is evicted. This behavior is generally not a problem and does not trigger an alarm. However, the appliance might be experiencing such a huge workload that more and more recent data needs to be evicted from the appliance to make space for incoming data. This causes the average evicted age to decrease, and when it goes below a certain threshold, the average evicted age alarm triggers. This alarm is an anomalous event because it signals that the appliance is handling a much larger workload than expected. The alarm is useful in detecting whether the appliance is undersized relative to the your normal workload. If the alarm is constantly triggered, then you should consider increasing the AltaVault's cache size with either expansion shelves or a larger virtual capacity.
Datastore Low Space	Indicates that the local data store (disk cache on the AltaVault) is running out of space and the eviction process on the AltaVault is unable to run at a sufficient pace to create space on the disk cache. This alarm might trigger when replication is too slow. View the Eviction Optimization report to determine how much disk cache is available.

Alarm	Description
Disk Full	<p>Enables an alarm if the system partitions (not the AltaVault data store) are full or almost full. For example, AltaVault monitors the available space on the directory that contains logs, statistics, system dumps, and TCP dumps.</p> <p>Generates an AutoSupport message when AutoSupport is enabled.</p> <p>By default, this alarm is enabled.</p> <p>This alarm monitors the following system partitions:</p> <ul style="list-style-type: none"> • Directory "/boot" free space • Directory "/bootmgr" free space • Directory "/config" free space • Directory "/tmpFull" • Directory "/var Full"
Hardware	<p>Fan Error - Indicates that the system has detected a problem with the fans. Fans for many systems can be replaced. Contact NetApp Support at https://mysupport.netapp.com and file a trouble ticket to order a replacement fan.</p> <p>IPMI - Indicates that there has been a physical security intrusion, triggering an Intelligent Platform Management Interface (IPMI) error. The following events trigger the IPMI alarm:</p> <ul style="list-style-type: none"> • Chassis intrusion (physical opening and closing of the appliance case) • Memory errors (ECC memory errors that can or cannot be corrected) • Hard drive faults or predictive failures • Power supply status or predictive failures <p>The option to reset the alarm appears only after the service triggers the IPMI alarm. To reset the alarm, click Clear the IPMI alarm now.</p>
Inconsistent Cloud Connectivity	Indicates that the connection to the cloud provider is inconsistent, leading to a large number of connection errors.
Inconsistent Cloud Data	Indicates that an inconsistency in the data stored in the cloud was detected.
Licensing	<p>Indicates that your appliance does not have a valid license.</p> <ul style="list-style-type: none"> • Evaluation Mode Expire - Evaluation mode license expired. • License Expired - A license expired. • License Expiring - A license will expire. • License Missing - A license is missing.
Link Duplex	<p>Indicates that an interface is not configured for half-duplex negotiation but has negotiated half-duplex mode. Half-duplex significantly limits the optimization service results.</p> <p>The alarm displays which interface is triggering the duplex error.</p> <p>Choose Configure > Data Interfaces and examine the AltaVault link configuration. Next, examine the peer switch user interface to check its link configuration. If the configuration on one side is different from the other, traffic is sent at different rates on each side, causing many collisions.</p> <p>To troubleshoot, change both interfaces to automatic duplex negotiation. If the interfaces do not support automatic duplex, configure both ends for full duplex.</p> <p>You can enable or disable the alarm for a specific interface. To disable an alarm, choose Configure > Alarms and select or clear the check box next to the link alarm.</p>

Alarm	Description
Link I/O Errors	<p>Indicates that the error rate on an interface has exceeded 0.1% while either sending or receiving packets. This threshold is based on the observation that even a small link error rate reduces TCP throughput significantly. A properly configured connection experiences very few errors. The alarm clears when the error rate drops below 0.05%.</p> <p>You can change the default alarm thresholds by entering the alarm link_errors err-threshold xxxxx CLI command at the system prompt. For details, see the <i>NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Guide</i>.</p> <p>To troubleshoot, try a new cable and a different switch port. Another possible cause is electromagnetic noise nearby.</p> <p>You can enable or disable the alarm for a specific interface. For example, you can disable the alarm for a link after deciding to tolerate the errors. To enable or disable an alarm, choose Configure > Alarms and select or clear the check box next to the link name.</p>
Link State	<p>Indicates that the system has detected a link that is down. You are notified through SNMP traps, email, and alarm status.</p> <p>The system has lost one of its Ethernet links due to an unplugged cable or dead switch port. Check the physical connectivity between the AltaVault and its neighbor device. Investigate this alarm as soon as possible. Depending on what link is down, the system might no longer be optimizing and a network outage could occur.</p> <p>You can enable or disable the alarm for a specific interface. To enable or disable the alarm, choose Configure > Alarms and select or clear the check box next to the link name.</p>
Low Memory	Indicates that low memory exists.
Max inodes limit	Indicates that the maximum number of files that can be stored has been reached.
Max Pinnable Limit	Indicates that the maximum pinnable limit has been reached.
Memory Paging	<p>Indicates that the system has reached the memory paging threshold. If 100 pages are swapped approximately every two hours, the AltaVault is functioning properly. If thousands of pages are swapped every few minutes, then reboot the AltaVault. If rebooting does not solve the problem, contact NetApp Support at https://mysupport.netapp.com.</p> <p>If the memory paging alarm triggers when the AltaVault is under a heavy load, you can ignore it.</p>
Metadata Space Full	Indicates that the data reserved for storing system metadata has filled up, leading to reduced deduplication.
Process Dump Creation Error	Indicates that the system has detected an error while trying to create a process dump. Contact NetApp Support at https://mysupport.netapp.com to correct the issue.
Secure Vault	<p>Enables an alarm and sends an email notification if the system encounters a problem with the secure vault:</p> <ul style="list-style-type: none"> Secure Vault Locked - Indicates that the secure vault is locked. To optimize SSL connections or to use data store encryption, the secure vault must be unlocked. Go to Configure > Secure Vault and unlock the secure vault.
SMB	Domain Controller Network Status - Indicates a Domain Controller is unreachable. The alarm is cleared when network connectivity to a Domain Controller is restored. If the alarm is not cleared after the network connectivity is restored, you can clear the alarm manually using alarm smb_alarms clear command.
Software Update Available	<p>Indicates that a newer version of the software is available.</p> <ul style="list-style-type: none"> To disable the alarm enter the following CLI command: <code>software_upgrade_available_clear</code>

Alarm	Description
Storage Optimization Service	<ul style="list-style-type: none"> Storage Optimization Service Down - Enables an alarm and sends an email notification if the Storage Optimization Service encounters a service condition. By default, this alarm is enabled. The message indicates the reason for the condition. The following conditions trigger this alarm: <ul style="list-style-type: none"> Configuration errors: examples include no encryption key set, incorrect appliance time, or incorrect cloud credentials. An AltaVault appliance reboot for example, during an appliance software update. A system crash for example, due to a power failure A Storage Optimization Service restart for example, due to a cloud storage provider change. A user enters the CLI command no service enable or shuts down the Storage Optimization Service from the Management Console A user restarts the optimization service from either the Management Console or CLI Storage Optimization Service Error - Enables an alarm and sends an email notification if the Storage Optimization Service encounters a condition that might degrade optimization performance. By default, this alarm is enabled.
Storage Optimization Service Replication	<ul style="list-style-type: none"> Replication Error - Enables an alarm when the replication to the cloud encounters an error. Displays an error message that indicates the type of error such as, a file cannot be replicated to the cloud. Replication Paused - Enables an alarm when the replication to the cloud pauses, because there is a cloud connection error, or you entered the CLI command no replication enable, or because you are using replication scheduling (nonbandwidth limit type). This alarm warns you that the AltaVault is not replicating data in the cloud. <p>By default, this alarm is enabled.</p>
System Reserved Space Full	Indicates that the space used for internal data structures is full. If you write more data to the appliance, it reduces deduplication performance.
Upgrade Status	Indicates the current AltaVault status.

The Alarm Status report answers the following question:

- What is the current status of the AltaVault?

To view the Alarm Status report

- Choose Reports > Alarm Status. Alternately, you can click the current alarm status that appears in the status box in the upper-right corner of each screen (Healthy, Degraded, or Critical).

Viewing the CPU utilization report

The CPU Utilization report summarizes the percentage of the CPU used within the time period specified.

Typically, the AltaVault operates on approximately 30 to 40 percent CPU capacity during nonpeak hours and approximately 60 to 70 percent capacity during peak hours. No single AltaVault CPU usage should exceed 95 percent.

The CPU Utilization report answers the following questions:

- How much of the CPU is being used?
- What is the average and peak percentage of the CPU being used?

To view the CPU Utilization report

1. Choose Reports > CPU Utilization.
2. Use the controls to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

To generate share level expanded bytes report

- To generate the share level expanded bytes report using the CLI command:

```
share-stats generate [num-threads]
```

This command scans all the files in the various shares and exports created and accumulates the statistics. The report contains the following statistics for all the shares in the appliance:

- Size of the share
- Number of files in the share
- Average file size
- Evicted percentage

The size of the share is the expanded bytes that is the actual size of the data being written to the appliance.

When the command is executed, an email containing the report is sent to the administrator. The report will not be display in the Web.

Viewing the memory paging report

The Memory Paging report provides the total number of memory pages, per second, utilized in the time period specified. It includes the following data that describe memory paging activity for the time period you specify.

Field	Description
Page Swap Out Rate	<p>Specifies the rate at which pages swapped. If 100 pages are swapped every couple of hours, the AltaVault is functioning properly. If thousands of pages are swapped every few minutes, contact NetApp Support at https://mysupport.netapp.com.</p>

The Memory Paging report answers the following questions:

- How much memory is being used?
- What is the average and peak number of memory pages swapped?

To view the Memory Paging report

1. Choose Reports > Memory Paging.
2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing the interface counters report

The Interface Statistics report displays information for each AltaVault interface, such as the interface name, its IP address, the Ethernet MAC address settings, the link settings, and the number of packets received and transmitted by the interface.

The Interface Statistics report includes the following statistics.

Field	Description
Interface	Name of the AltaVault interface.
IP	IP address of the network interface.
Ethernet	The MAC address information for the interface. The speed can be Auto, 1000, 100, or 10. The default value is Auto.
Link	Parameter that indicates whether there is an active connection plugged into the interface. Displays true if the interface is connected; otherwise, it displays false.
Receive Packets	Number of packets received by the interface. It also provides details such as the number of packets discarded, packet errors, packet overruns, packet frames, and packets multicast.
Transmit Packets	Number of packets transmitted by the interface. It also provides details such as the number of packets discarded, packet errors, packet overruns, packet frames, and packets multicast.

The Interface Statistics report answers the following questions:

- What is the name, IP address, Ethernet information, and link state of each AltaVault interface?
- What is the number of packets received by the interface? How many of these packets were discarded, and how many had errors, overruns, frames, and multicast?
- What is the number of packets transmitted by the interface? How many of these packets were discarded, and how many had errors, overruns, frames, and multicast?

To view the Interface Counters report

- Choose Reports > Interface Counters.

Viewing the disk throughput report

The Disk Throughput report displays information for each AltaVault disk, such as the average read throughput, and average write throughput.

The Disk Throughput report includes the following statistics.

Field	Description
Read	Average data read from the AltaVault internal disk subsystem.
Write	Average data written to the AltaVault internal disk subsystem.

The Disk Throughput report answers the following questions:

- What is the average read throughput that the AltaVault is experiencing at the specified time?
- What is the average write throughput that the AltaVault is experiencing at the specified time?

The average disk throughput is a measure of performance that help you to discover disk-based bottlenecks.

The rate at which the AltaVault can read from the internal disk subsystem characterizes the performance by which the AltaVault can serve read requests to client machines and transfer data to the cloud storage provider.

The rate at which the AltaVault can write to the internal disk subsystem characterizes the performance by which the AltaVault can serve write requests from the client machines and retrieve data from the cloud storage provider.

To view the Disk Throughput report

1. Choose Reports > Disk Throughput.

2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing the disk IOPS report

The Disk IOPS report displays the average Input/Output Operations Per Second (IOPS) information for each AltaVault disk.

The Disk IOPS report includes the following statistics.

Field	Description
Read	Average value of input/output operations per second for read operations from the AltaVault internal disk subsystem.
Write	Average value of input/output operations per second for write operations to the AltaVault internal disk subsystem.

The Disk IOPS report answers the following questions:

- What is the average read IOPS that the AltaVault is experiencing at the specified time?
- What is the average write IOPS that the AltaVault is experiencing at the specified time?

The average disk IOPS is a measure of performance that help you to discover disk-based bottlenecks.

To view the Disk IOPS report

1. Choose Reports > Disk IOPS.

2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing the disk utilization report

The Disk Utilization report displays information for each AltaVault disk, such as the average percentage utilization, cumulative average queue size, cumulative average wait, and cumulative average service time.

The Disk Utilization report includes the following statistics.

Field	Description
Utilization	The average percentage of the aggregate measurement of access time and transfers per second.
Queue Size	A running average of the number of requests either in service or waiting for service on the AltaVault internal disk subsystem.
Wait Time	A running average of the wait time per I/O operation experienced by the AltaVault internal disk subsystem.
Service Time	A running average service time per I/O operation experienced by the AltaVault internal disk subsystem.

The Disk Utilization report answers the following questions:

- What is the estimated service time for an I/O request on the AltaVault internal disk subsystem?
- What is the availability of AltaVault services to process I/O requests?
- How well are the AltaVault services using the internal disk subsystem?

To view the Disk Utilization report

1. Choose Reports > Disk Utilization.

2. Use the controls in the control panel to customize the report as described in this table.

Control	Description
5m 1h 1d 1w All	<p>Select one of the following report time intervals to filter the display:</p> <ul style="list-style-type: none"> • 5m - Displays five minutes of data. • 1h - Displays one hour of data. • 1d - Displays one day of data. • 1w - Displays one week of data. • All - Displays all data available for the past 30 days. <p>To type a custom time interval, enter the start time and end time (using the format YYYY/MM/DD HH:MM:SS) in the text field.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically.</p>
Show newest data	<p>Click this option to display only the latest data available. The latest data available depends on the time interval that you specify. For example, if 1h is the specified time interval, then clicking Show newest data displays the last 60 minutes of data.</p>

Viewing logs

The AltaVault log reports provide a high-level view of network activity. You can view both user and system logs.

- [“Viewing system logs” on page 152](#)
- [“Viewing user logs” on page 153](#)

Viewing system logs

You can view system logs in the Maintenance > System Logs page. View System logs to monitor system activity and to troubleshoot problems. The most recent log events are listed first.

To view system logs

1. Choose Maintenance > System Logs.
2. Use the controls to customize the report as described in this table.

Control	Description
Show	Select one of the archived logs or Current Log from the drop-down list.
Lines per page	Specify the number of lines that you want to display in the page.
Jump to	<p>Select one of the following options from the drop-down list:</p> <ul style="list-style-type: none"> • Page - Specify the number of the page that you want to display. • Time - Specify the time for the log that you want to display.

Control	Description
Filter	Select one of the following filtering options from the drop-down list: <ul style="list-style-type: none"> • Regular expression - Specify a regular expression on which to filter the log. • Error or higher - Displays error-level logs or higher. • Warning or higher - Displays warning-level logs or higher. • Notice or higher - Displays notice-level logs or higher. • Info or higher - Displays informational-level logs or higher.
Go	Displays the report.

To view a continuous log

1. Choose Maintenance > System Logs.
2. Customize the log as described in [“To view system logs” on page 152](#).
3. Click **Launch Continuous Log** in the upper-right corner of the page.

You can continuously display new lines as the log expands.

Viewing user logs

You can view user logs in the Maintenance > User Logs page. The user log filters messages from the system log to display messages that are of immediate use to the system administrator.

View user logs to monitor system activity and to troubleshoot problems. For example, you can monitor who logged in, who logged out, and who entered particular CLI commands, and you can monitor alarms and errors. The most recent log events are listed first.

To view and customize user logs

1. Choose Maintenance > User Logs.
2. Use the controls to customize the log as described in this table.

Control	Description
Show	Select one of the archived logs or Current Log from the drop-down list.
Lines per Page	Specify the number of lines that you want to display in the page.
Jump to	Select one of the following options from the drop-down list: <ul style="list-style-type: none"> • Page - Specify the number of the page that you want to display. • Time - Specify the time for the log that you want to display.

Control	Description
Filter	<p>Select one of the following filtering options from the drop-down list:</p> <ul style="list-style-type: none"> • Regular expression - Specify a regular expression on which to filter the log. • Error or higher - Displays error-level logs or higher. • Warning or higher - Displays warning-level logs or higher. • Notice or higher - Displays notice-level logs or higher. • Info or higher - Displays informational-level logs or higher.
Go	Displays the report.

3. Click the **Launch Continuous Log** in the upper-right corner of the page.

You can continuously display new lines as the log expands.

To view a continuous log

1. Choose Maintenance > User Logs.
2. Customize the log as described in [“To view and customize user logs” on page 153](#).
3. Click **Launch Continuous Log** in the upper-right corner of the page.

Downloading log files

This section describes how to download user and system log files.

You can download both user and system logs.

- [“Downloading user log files” on page 154](#)
- [“Downloading system log files” on page 155](#)

The download page displays up to 10 archived log files plus the current day log file. By default, the system rotates each file every 24 hours or if the file size reaches one gigabyte uncompressed size. You can change this to rotate every week or month. Additionally, you can rotate the files based on file size.

The automatic rotation of system logs deletes your oldest log file, labeled as Archived log #10, pushes the current log to Archived log # 1, and starts a new current-day log file.

Downloading user log files

You can download user logs in the Maintenance > User Logs Download page. Download user logs to monitor system activity and to troubleshoot problems.

To download user logs

1. Choose Maintenance > User Logs Download.
2. Click the name of the log to display the dialog box to display or save the log to disk.

3. Click **Rotate Logs** to archive the current log to a numbered archived log file and then clear the log so that it is empty again.

Downloading system log files

You can download system logs in the Maintenance > System Logs Download page. Download system logs to monitor system activity and to troubleshoot problems.

To download system logs

1. Choose Maintenance > System Logs Download.
2. Click the name of the log to display the dialog box to display or save the log to disk.
3. Click **Rotate Logs** to archive the current log to a numbered archived log file and then clear the log so that it is empty again.

Generating system dumps

You can generate and download system dumps in the Maintenance > System Dumps page. A system dump contains a copy of the kernel data on the system. System dump files can help you diagnose problems in the system.

To generate system dump files

1. Choose Maintenance > System Dumps.
2. Click the name of the system dump file and then click **Download System Dump** to download it.
3. Optionally, upload a file to NetApp support.

Enter your case number and click Upload to upload the system dump to the NetApp Support site. If you do not have a case number, visit support.netapp.com to open a support case.

If a proxy is required at your site for outgoing FTP, select Configure > Host Settings and configure the Web Proxy settings.
4. Under Generate System Dump, select the type of information to include in the report:
 - **Include Statistics** - Select to collect and include CPU, memory, and other statistics in the system dump (this option is enabled by default).
 - **Include All Logs** - Removes the 50 MB limit for compressed log files, to include all logs in the system dump.

5. Click **Generate System Dump**.

A spinner will be displayed during the system dump creation. When the system dump is complete, it appears in the list of links to download.

To view system dump files

1. Choose Maintenance > System Dumps.

2. Click **Download Link** to view a previously saved system dump.
3. Select the filename to open a file or save the file to disk.
4. To remove an entry, check the box next to the name and click **Remove Selected**.

Viewing process dumps

You can display and download process dumps in the Maintenance > Process Dumps page. A process dump is a saved copy of memory, including the contents of all memory, bytes, hardware registers, and status indicators. It is periodically taken to restore the system in the event of failure. Process dump files can help you diagnose problems in the system.

To view process dump files

1. Choose Maintenance > Process Dumps.
2. Click the filename to open a file or save the file to disk.
3. To remove an entry, select the box next to the name and click **Remove Selected**.

Capturing and uploading TCP dumps

You can create, download, and upload TCP dumps in the Maintenance > TCP Dumps page.

TCP dumps contain summary information for every Internet packet received or transmitted on the interface to help diagnose problems in the system.

The AltaVault provides an easy way to capture and retrieve multiple capture files from the Management Console. You can create TCP dumps from multiple interfaces at the same time, limit the size of the TCP dump, and schedule a specific date and time to capture TCP information. Scheduling and limiting a TCP capture file by time or size allows unattended captures.

The top of the TCP Dumps page displays a list of existing TCP dumps and the bottom of the page displays controls to create a capture file. The bottom of the page also includes the capture files that are currently running, and controls to create a trigger that stops a capture when a specific event occurs. The Running Capture Name list includes captures running at a particular time. It includes captures started manually and also any captures that were scheduled previously and are now running.

To capture TCP trace dumps

1. Choose Maintenance > TCP Dumps.
2. Under TCP Dumps Currently Running, complete the configuration as described in this table.

Control	Description
Add a New TCP Dump	Displays the controls for creating a capture file.

Control	Description
Capture Name	<p>Specify the name of the capture file. Use a unique filename to prevent overwriting an existing capture file. The default filename uses this format:</p> <p><i>hostname_interface_timestamp.cap</i></p> <p><i>hostname</i> is the hostname of the AltaVault, <i>interface</i> is the name of the interface selected for the trace (for example, lan0_0, wan0_0), and <i>timestamp</i> is in the YYYY-MM-DD-HH-MM-SS format.</p> <p>If this capture file relates to an open NetApp Support case, specify the capture filename case_ <i>number</i> where <i>number</i> is your NetApp Support case number. For example, case_12345.</p> <p>The .cap file extension is not included with the filename when it appears in the capture queue.</p>
Endpoints	<p>Specify IP addresses and port numbers to capture packets between them:</p> <p>IPs - Specify IP addresses of endpoints on <i>one side</i>. Separate multiple IP addresses using commas. You can enter IPv6 addresses separated by commas. The default setting is all IP addresses.</p> <p>Ports - Specify ports on <i>one side</i>. Separate multiple ports using commas. The default setting is all ports.</p> <p>—<i>and</i>—</p> <p>IPs - Specify IP addresses of endpoints on the <i>other side</i>. Separate multiple IP addresses using commas. You can enter IPv6 addresses separated by commas. The default setting is all IP addresses.</p> <p>Ports - Specify ports on the <i>other side</i>. Separate multiple ports using commas. The default setting is all ports.</p> <p>To capture traffic flowing in only one direction or to enter a custom command, use the CLI tcpdump command. For details, see the <i>NetApp AltaVault Cloud Integrated Storage Command-Line Reference Guide</i>.</p>
Capture Interfaces	<p>Captures packet traces on the selected interfaces. You can select all interfaces or a base, in-path, or RSP interface. The default setting is none. You must specify a capture interface.</p> <p>If you select several interfaces at a time, the data is automatically placed into separate capture files.</p> <p>When path selection is enabled, NetApp recommends that you collect packet traces on all LAN and WAN interfaces.</p>

Control	Description
Capture Parameters	<p>These parameters let you capture information about dot1q VLAN traffic. You can match traffic based on VLAN-tagged or untagged packets, or both. You can also filter by port number or host IP address and include or exclude ARP packets. Select one of these parameters for capturing VLAN packets:</p> <ul style="list-style-type: none"> • Capture Untagged Traffic Only - Select this option for the following captures: <ul style="list-style-type: none"> – All untagged VLAN traffic. – Untagged 7850 traffic and ARP packets. – Only untagged ARP packets. • Capture VLAN-Tagged Traffic Only - Select this option for the following captures: <ul style="list-style-type: none"> – Only VLAN-tagged traffic. – VLAN-tagged packets with host 10.11.0.6 traffic and ARP packets. You must also specify 10.11.0.6 in the IPs field, and specify or arp in the custom flags field on this page. – VLAN-tagged ARP packets only. You must also specify and arp in the custom flags field on this page. • Capture both VLAN and Untagged Traffic - Select this option for the following captures: <ul style="list-style-type: none"> – All VLAN traffic. – Both tagged and untagged 7850 traffic and ARP packets. You must also specify the following in the custom flags field on this page: (port 7850 or arp) or (vlan and (port 7850 or arp)) – Both tagged and untagged 7850 traffic only. You must also specify 7850 in one of the port fields on this page. No custom flags are required. – Both tagged and untagged ARP packets. You must also specify the following options in the custom flags field on this page: (arp) or (vlan and arp)
Capture Duration (Seconds)	<p>Specify a positive integer to set how long the capture runs, in seconds. The default value is 30. Specify 0 or continuous to initiate a continuous trace.</p> <p>For continuous capture, NetApp recommends specifying a maximum capture size and a nonzero rotate file number to limit the size of the TCP dump.</p>
Maximum Capture Size	<p>Specify the maximum capture file size, in megabytes. The default value is 100. After the file reaches the maximum capture size, TCP dump starts writing capture data into the next file, limited by the Number of Files to Rotate field.</p> <p>NetApp recommends a maximum capture file size of 1024 MB (1 GB).</p>
Buffer Size	<p>Optionally, specify the maximum amount of data, in kilobytes, allowed to queue while awaiting processing by the capture file. The default value is 154 KB.</p>
Snap Length	<p>Optionally, specify the snap length value for the capture file, which equals the number of bytes captured for each packet. Having a snap length smaller than the maximum packet size on the network enables you to store more packets, but you might not be able to inspect the full packet content. Specify 0 for a full packet capture (recommended for SMB, MAPI, and SSL captures). The default value is 1518 bytes.</p>
Number of Files to Rotate	<p>Specify how many capture files to keep for each interface before overwriting the oldest file. To stop file rotation, you can specify 0; however, NetApp recommends rotating files, because stopping the rotation can fill the disk partition.</p> <p>This limits the number of files created to the specified number, and begins overwriting files from the beginning, thus creating a rotating buffer.</p> <p>The default value is five files per interface. The maximum value is a 32-bit integer.</p>

Control	Description
Custom Flags	<p>Specify custom flags as additional statements within the filter expression. Custom flags are added to the end of the expression created from the Endpoints fields and the Capture Parameters radio buttons (pertaining to VLANs).</p> <p>If you require an “and” statement between the expression created from other fields and the expression that you are entering in the custom flags field, you must include the “and” statement at the start of the custom flags field.</p> <p>Do not use host, src, or dst statements in the custom flags field. Although it is possible in trivial cases to get these to start without a syntax error, they do not capture GRE-encapsulated packets that some modes of AltaVault communications use, such as WCCP deployments or Interceptor connection-setup traffic. NetApp recommends using bidirectional filters by specifying endpoints.</p> <p>For complete control of your filter expression, use the CLI tcpdump command. For details, see the <i>NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Guide</i>.</p>
Schedule Dump	<p>Schedules the capture to run at a later date and time.</p> <p>Start Date - Specify a date to initiate the capture, in this format: YYYY/MM/DD.</p> <p>Start Time - Specify a time to initiate the capture, in this format: HH:MM:SS.</p>
Add	Adds the capture request to the capture queue.

Troubleshooting

If your command results in a syntax error with an immediate or scheduled TCP dump, this message appears:

“Error in tcpdump command. See System Log for details.”

Review the system log to see the full tcpdump command attempt. Check the expression for issues such as a missing “and,” as well as contradictory instructions such as looking for VLAN-tagged traffic AND non-tagged traffic.

Custom flag use examples

The examples in this table focus on the custom flag entry but rely on other fields to create a complete filter.

Filter Purpose	Custom Flag
To capture all traffic on VLAN 10 between two specified endpoints: 1.1.1.1 and 2.2.2.2	and vlan 10
To capture any packet with a SYN or an ACK	tcp[tcpflags] & (tcp-syn tcp-ack) != 0
To capture any packet with a SYN	tcp[tcpflags] & (tcp-syn) != 0 —or— tcp[13] & 2 == 2
To capture any SYN to or from host 1.1.1.1	and (tcp[tcpflags] & (tcp-syn) != 0) —or— and (tcp[13] & 2 == 2)

Stopping a TCP dump after an event occurs

TCP dumps offer visibility into intermittent network issues, but the amount of traffic they capture can be overwhelming. Be aware that automated log rotation on the AltaVault can overwrite debugging information specific to the event.

To limit the amount of data collected, you can use a trigger to stop a continuous TCP dump after a specific log event occurs. The stop trigger continuously scans the system logs for a search pattern. The result is a smaller file to help pinpoint what makes the event happen.

The stop trigger continuously scans the system logs for a search pattern. When it finds a match, it stops all running captures.

To stop a TCP dump after a specific log event

1. Choose Maintenance > TCP Dumps.
2. Add a new TCP dump and schedule it to run. For more details about adding a TCP Dump, see [“To capture TCP trace dumps” on page 156](#).
3. Scroll down to the TCP Dump Stop Trigger section on the TCP Dumps page.
4. In the Pattern text box, enter a Perl regular expression (regex) to find in a log. The system compares the Perl regex against each new line in the system logs and the trigger stops if it finds a match.

The simplest regex is a word or a string of characters. For example, if you set the pattern to “Limit,” the trigger matches the line “Connection Limit Reached.”

Notes:

- Perl regular expressions are case-sensitive.
- Perl treats the space character “\space “ like any other character in a regex.
- Perl reserves some characters, called metacharacters, for use in regex notation. The metacharacters are:
`{ } [] () ^ $. | * + ? \`

You can match a metacharacter by putting a backslash before it. For example, to search for a backslash in the logs, you must enter two backslashes (\\) as the pattern.

- The pattern follows Perl regular expression syntax. For details, go to:
<http://perldoc.perl.org/perlre.html>
 - You cannot change the pattern while a scan is running. You must stop the scan before changing a pattern.
 - You do not need to wrap the pattern with the metacharacters to match the beginning or end of a line (^ \$) or with the wildcard character (*).
5. Specify the amount of time to pause before stopping all running captures when the system finds a match. This gives the system some time to log more data without abruptly cutting off the capture. The default is 30 seconds. Specify 0 for no delay; the capture stops immediately.

After a trigger has fired, the capture can stop by itself before the delay expires. For example, the capture duration can expire.

6. Click **Start Scan**.

When the scan stops, the system sends an email to all email addresses on the Configure > Email page appearing under Report Events via Email. The email notifies users that the trigger has fired.

The page indicates “Last Triggered: Never” if a TCP Dump stop trigger has never triggered on the AltaVault. After the delay duration of the stop trigger, the system displays the last triggered time.

Before changing the Perl regular expression or amount of delay, you must first stop the process.

To stop a running scan

- Click **Stop Scan** to halt the background process that monitors the system logs. The system dims this button when the stop trigger is idling.

Stop trigger limitations

These limitations apply to the trigger:

- You cannot create a trigger to stop a specific capture; the trigger affects all running captures.
- If the search pattern contains a typo, the trigger might never find a match.
- Only one instance of a trigger can run at one time.

Viewing a TCP dump

The top of the TCP Dumps page displays a list of existing TCP trace dumps.

To view a capture file

1. Choose Maintenance > TCP Dumps.
2. Under Stored TCP Dumps, select the trace dump name to open the file.
3. Click **Download** to view a previously saved capture file.
4. To remove a capture file, select the check box next to the name and click **Remove Selected**.

To print a capture file

1. Choose Maintenance > TCP Dumps.
2. Under Download Link, select the capture filename to open the file.
3. When the file opens, choose File > Print in your Web browser to open the Print dialog box.

To stop a running a capture

1. Choose Maintenance > TCP Dumps.
2. Select the capture filename in the Running Capture Name list.
3. Click **Stop Selected Captures**.

Uploading a TCP dump

NetApp offers several methods to upload capture files to the support server for sharing with the support team while diagnosing issues.

To upload the capture file to NetApp support

1. In continuous mode, on the TCP Dumps page, select the running capture and click **Stop Selected Captures**.

For timed captures that are complete, skip to Step 2.

The capture appears as a download link in the list of Stored TCP Dumps.

2. Select the capture filename.

3. Optionally, specify a case number that corresponds to the capture. NetApp Support recommends using a case number. For example, 194170.

To specify a URL instead of a case number, you must use the CLI. You can enter the CLI command **file tcpdump upload URL**. When you specify a URL, the capture file goes directly to the URL.

If the URL points to a directory on the upload server, it must have a trailing backslash (/).

For example:

`ftp://ftp.netapp.com/incoming/`

(not `ftp://ftp.netapp.com/incoming`)

The filename as it exists on the appliance will then match the filename on the upload server.

For details, see the *NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Guide*.

4. Click **Upload**.

A progress bar displays the percentage of the total upload completed, the case number (if applicable), and the date and time the upload began. When the capture file finishes uploading, the date, time, and a status of either uploaded (appears in green) or failed (appears in red).

Successful uploads show the status, the case number (if applicable), and the date and time the upload finished.

For uploads that fail, an explanation, the case number (if applicable), and the upload starting date and time appear.

Viewing the appliance monitoring report

The Appliance Monitoring report enables you to view the health status, disk space, and cloud service status of multiple AltaVaults using a monitoring AltaVault.

The Appliance Monitoring report includes the following statistics.

Field	Description
Hostname	Hostname of the monitored appliance.
Appliance Health	<p>Current state of the appliance, which is one of the following options:</p> <ul style="list-style-type: none"> • Healthy - The AltaVault is active and optimizing storage effectively. You do not need to take any action. • Needs Attention - Accompanies a healthy state to indicate management-related issues that does not affect the ability of the AltaVault to optimize storage. Address the management-related issue in a low-priority order. • Degraded - The appliance is optimizing storage but the system has detected an issue. Address the issue that the system has detected. • Admission Control - The appliance is optimizing storage but has entered the admission control state. Admission control limits the number of connections made to the AltaVault, so that you do not over-consume resources on your system. To clear this state, reduce the number of connections to the AltaVault. • Critical - The AltaVault might or might not be optimizing storage. You must address a critical issue immediately.
Service State	Current state of the Storage Optimization Service.
Expanded Data	Data that is backed up locally by the AltaVault.
Deduplicated Data	Data that has been deduplicated. Deduplication is a method of reducing storage space by eliminating redundant data. It stores only one unique instance of the data. It replaces redundant data with a pointer to the unique data copy.
Deduplication Factor	The average deduplication factor, which is the ratio of the expanded data and total optimized data. The total optimized data includes both deduplication and compression savings.
Cloud Storage Used	The amount of cloud storage used by the appliance.
Cloud Storage Available	The amount of storage available in the cloud.
AltaVault Storage Available	The amount of storage available on the monitored AltaVault.
Bytes Pending Replication	The number of bytes in the local disk that are pending replication to the cloud.

The Appliance Monitoring report provides the following statistics for each monitored appliance: hostname, appliance health, service state, expanded data, deduplicated data, deduplication factor, cloud storage used, cloud storage available, AltaVault storage available, and bytes pending replication.

It also displays a summary of the total expanded data, total deduplicated data, average deduplication factor, total cloud storage used, total cloud storage available, total AltaVault storage available, total bytes pending replication to the cloud for all monitored appliances in the network.

To configure appliance monitoring

1. Choose Reports > Appliance Monitoring.

2. Under Monitored Appliances, complete the configuration as described in this table.

Control	Description
Add Monitored Appliance	Displays controls to add a monitored appliance.
Hostname or IP Address	Specify a valid hostname or IP address for the monitored appliance.
API Access Code	Specify the API access key to access the monitored appliance. To obtain the API access key: <ul style="list-style-type: none"> • Open a separate browser. • Log in to the appliance that you want to monitor. • Choose Configure > REST API Access. • Generate the API Access key. • Copy it into Notepad. • Log back in to the monitoring AltaVault. • Paste the API access key in this text box.
Add	Click Add to add the peer appliance to the list of monitored appliances.
Remove Selected Appliances	Select the check box next to the name and click Remove Selected Appliances to delete it from the system.

Viewing the shelf details

The Shelf Details report enables you to view information about the AltaVault hardware RAID. It is applicable only to AltaVault hardware models; not to the virtual or cloud models.

The Shelf Details displays the AltaVault enclosure ID and a visual replica of the hardware disk array that includes the head unit and the expansion shelves.

The shelf serial number can be found on the back of the shelf. The slot number identifies the slot within the shelf.

Hover your mouse over each disk shown on the Shelf Details report to obtain the status of the disk and its serial number.

The top right LED is the power LED. If the disk is online, it is constantly green. If the disk is offline, it is amber.

The Shelf Details report provides the serial number and status of the AltaVault hardware disk array.

To identify physical location of a drive

- Use the following CLI:

```
CLI > show hwraid disk information
```

To view the shelf details

1. Choose Reports > Shelf Details.
2. Click the serial number of a shelf to obtain alarm information for the RAID.

Viewing the storage RAID group

The Storage RAID groups enables you to view the serial number, status, size, used, model number, and product ID. For AVA400, the maximum number of RAID groups supported is five. For AVA800, the maximum number of RAID groups supported is seven.

The Storage RAID groups includes the following statistics.

Field	Description
Serial Number	The serial number of the Storage RAID group.
Status	Current state of the Storage RAID group, which is one of the following options: <ul style="list-style-type: none">Valid - The Storage RAID group is valid and active.Invalid - The Storage RAID group is not valid. Click the serial number of the alarm information for the Storage RAID group.
Size	The size of the Storage RAID group (terabytes).
Used	Current space occupied by the Storage RAID group.
Model	The Storage RAID group model.
Product ID	The Storage RAID group product ID number.

The Storage RAID Group report provides statistics for each group.

To view the Storage RAID Group report

1. Choose Reports > Storage RAID groups.
2. Click the serial number of a group to obtain information for the RAID.

To view diagnostics using the CLI

- Run the following command to view diagnostics:

```
CLI > # show raidgroups
```

Viewing offline file system check page

The Offline File System Check page checks the integrity of the local file system and provides the diagnosis. You cannot run the Offline File System Check if the Storage Optimization Service is running. Before you perform the Offline File System Check, choose Maintenance > Service and click **Stop** to stop the Storage Optimization Service.

To view the offline file system check page

1. Choose Maintenance > Offline File System Check.
2. Click the filename to open an Offline File System Check log file or save the file to disk.
3. To remove an entry, select the box next to the name and click **Remove Selected**.

4. Under Offline File System Check Actions, complete the configuration as described in this table.

Control	Description
Status	Displays the status of the Offline File System Check report such as, running or not running.
Check type	Select the type of diagnosis from the drop-down list: <ul style="list-style-type: none"> Internal Data Consistency - Checks only the internal consistency of the data. This is the fast mode. Complete - Checks the data completely from end to end. This is the slow mode.
Repair errors	Click this option to repair errors from which the system can recover.
Replay log after a crash	Click this option to replay log after a crash.
Start Mfsck	Starts the Offline File System Check tool.
Stop Mfsck	Stops the Offline File System Check tool.

5. Open and check the Offline File System Check log file that you downloaded. If your configuration is accurate, the system displays the following results:

- File verification summary displays: File check result: Volume consistent
- Label verification summary displays: Label check result: Volume consistent.

Otherwise, the log file displays an error message.

Viewing online file system check page

The Online File System Check performs online verification of the data as it is written to the cloud. The Online File System Check report displays the log files that contain the integrity check data, the date and time of the check, and the file size.

You can perform the Online File System Check only if the Storage Optimization Service is running. You can stop the check at any time.

Field	Description
Download Link	Option to download a specific log file.
Integrity Check done up to	Displays the date and time up to which the AltaVault has performed the Online File System Check.
Size	Displays the size of the log file.

To view the Online File System Check page

1. Choose Maintenance > Online File System Check.
2. Under Online File System Check Actions, click **Start Data Integrity Check** to start the check.

-or-

Click **Stop Data Integrity Check** to stop the check.

To set the file system check duration using the CLI

In some cases, file system checking can impact overall utilization of system resources. You can manage system utilization for file system checks by setting the duration in the CLI. For example, setting a longer duration extends the amount of time over which file system checks are done, reducing system utilization for a given time.

1. To change the file system check duration, enter the following command:

```
hostname (config)# datastore integrity check duration <duration>
```

Set the duration from 1 to 60 days. The default duration is 30 days. The setting takes effect in the next check cycle.

Viewing the verify tool diagnostics

The Verify page provides the diagnosis for checking cloud replication consistency. Start the Verify operation only after all pending replication has completed successfully and the value of Replication Bytes Pending is zero.

You cannot run the Verify tool if the storage optimization service is running.

Note: If the cloud provider is Alibaba Archive, Amazon Glacier, or Azure Archive, validate only the checksum using the checksum only action.

To view the Verify tool diagnostics

1. Choose Maintenance > Replication Verify Check.
2. Click the filename to open a Verify log file or save the file to disk.
3. Under Actions, complete the configuration as described in this table.

Control	Description
Status	Displays the status of the Offline File System Check report, such as running or not running.
Only validate checksums	Click this option to verify only the checksums.
Start	Starts the Verify tool.
Stop	Stops the Verify tool.

4. Open and check the Verify log file. If all data in the local disk is replicated in the cloud, the file displays the following result:

```
Verification complete: collection is properly replicated.
```

Otherwise, the log file displays an error message.

CHAPTER 11 Transferring data to the cloud using Amazon Snowball

This chapter describes transferring large amounts of data from your AltaVault to the cloud using Amazon Snowball. It includes the following sections:

- [“Prerequisites” on page 169](#)
- [“Guidelines for using Snowball with AltaVault” on page 169](#)
- [“Seeding data using Snowball” on page 170](#)

Snowball provides a cost-effective and efficient means of moving data to the cloud while taking advantage of the deduplication, compression, and encryption capabilities of AltaVault. A single Snowball appliances can transfer terabytes of information from AltaVault to the Amazon cloud. Seeding your cloud storage using Snowball alleviates the challenges of moving large data sets across the Internet during the initial backup of your site.

Prerequisites

- AltaVault must be configured to use the Amazon S3 cloud. For AltaVault cloud configuration details, see [“To use the Cloud Settings wizard” on page 22](#).
- AltaVault replication must be suspended immediately after configuring AltaVault to use the S3 cloud so no data replication to S3 occurs over the Internet. To suspend replication, see [“Configuring replication” on page 38](#).
- You must use the same AWS account to order Snowball as was used to configure S3 bucket/cloud credentials.
- Snowball Edge is not supported.
- AltaVault virtual appliances must be configured with an additional 8 GB of memory to support Snowball transfers. For example, increase AVA-v8 memory from 24 to 32 GB prior to starting data transfer using Snowball. For AVA-v2, increase memory from 6 to 14 GB. For AVA-v16 and AVA-v32, increase memory from 48 GB to 56 GB and from 96 GB to 104 GB, respectively.

Guidelines for using Snowball with AltaVault

- After initially configuring Amazon S3 as the cloud storage with AltaVault, you must suspend replication on AltaVault. No data should be allowed to replicate from AltaVault to the cloud prior to performing one or more Snowball operations.
- Snowball configuration and management on AltaVault is provided through the command-line interface (CLI) only.

- Seeding data from AltaVault to the cloud is generally a one-time operation that is done after initially configuring a cloud service provider. It assumes no data exists in the cloud bucket.
- The amount of time required for transferring data from AltaVault to Snowball can vary based on the data size and throughput rate for your network connection. Estimates range from hours to days when completely filling a Snowball. The data transfer rate when using a direct connection is approximately 150 MB per second.
- Snowball is available in 50 TB or 80 TB capacities. AltaVault supports writing to only one Snowball at a time.
- Do not resume replication until after the seeded data has been verified in Amazon and the seeding operation is complete.
- Any data transferred to Snowball remains and will be copied to the cloud bucket when the Snowball is returned to Amazon. If a Snowball job is canceled and the Snowball is re-used for a new seeding operation, the existing contents of the Snowball are overwritten with the latest copy.

Seeding data using Snowball

The Snowball appliance is available through Amazon Web Services (AWS). This section covers the following topics:

- [“Creating a Snowball job in AWS” on page 170](#)
- [“Transferring data from AltaVault to Snowball” on page 171](#)
- [“Managing data transfers on AltaVault” on page 172](#)
- [“Verifying and completing data transfer” on page 173](#)

Read this entire procedure before deploying Snowball with AltaVault.

Creating a Snowball job in AWS

1. Log in to the Amazon Web Services at <https://console.aws.amazon.com/>.
2. Select **Snowball** from the list of Storage & Content Delivery services.
3. In the Job dashboard, select **Create a job**.
4. Under Plan your job, click **Import into Amazon S3**
5. Click **Next**.
6. Under Give shipping details, select the shipping address and shipping speed.
7. Click **Next**.
8. Under Give job details, specify the job information:

Job information	Description
Import job name	Specify a name to apply to the job. This can be the same as the bucket name.
Region	Select a destination region from the pull-down menu.

Job information	Description
Bucket name	Select a destination bucket name for the data. This is the existing bucket that was specified when configuring the cloud connection on AltaVault to Amazon S3.
Storage capacity	Select the Snowball storage capacity for data size.

9. Click **Next**.

10. Under Set security, select your settings:

Security Setting	Descriptions
Permission	Click Create/Select IAM role to provide the Identity and Access Management setting.
Encryption	For KMS key, select your setting from the pull-down menu. If you do not have a KMS key, choose “(default) aws/importexport.”

11. Click **Next**.

12. Under Set notifications, optionally, you can choose to select **Create new SNS topic** and enter a topic name and email address if you wish to receive notifications.

Under Job status, select each event for which you would like to receive notification.

13. Click **Next**.

14. Review the summary information for your job. If everything is correct, click **Create job**.

Amazon prepares and ships the Snowball appliance.

Transferring data from AltaVault to Snowball

1. After receiving the Snowball appliance, plug it in, and connect it to the AltaVault using a 10GbE SFP+ Optical (recommended) or a 1GbE Ethernet cable. When using copper cables with direct connection, a cross-over cable is required. For best performance it is recommended to use direct connections between the appliances.
2. Using the E-Ink display (touch screen) on Snowball, select your network interface type, and configure an IP address, netmask, and default gateway.
3. From the AWS Job dashboard, select the job name to view Job status for the Snowball appliance.
4. Click **Get credentials** to access the job manifest and a 25 character unlock code. Both the manifest file and the unlock code are required to start Snowball. A separate unlock code and manifest file is provided for each job.
5. Click **Download manifest** to download a copy of the file to a server that can be accessed by AltaVault.
6. Log in to the AltaVault CLI.
7. Upload a copy of the manifest file to the AltaVault:

```
hostname (config)# seed snowball manifest-file fetch <manifest-file-url>
```

For the `manifest-file-url`, HTTP, HTTPS, SCP, or FTP can be used. For example:

```
hostname (config)# seed snowball manifest-file fetch http://192.168.1.1/tools/DIJ9bfa9234-0256-c451-8abc-62db308017040_manifest.bin
```

8. Enter the following command to verify the manifest file has been downloaded:

```
hostname (config)# show seed snowball manifest-file
```

9. Enter the following command to verify connectivity with Snowball:

```
hostname (config)# ping <snowball ip address>
```

10. Start transferring data to Snowball:

```
hostname (config)# seed snowball start ip <ip> unlock-code <code> secret-key <key>
```

11. Enter the following command to check the seeding status:

```
hostname (config)# show seed
```

When data transfer is done, or the Snowball appliance is full, the seeding status changes to Paused.

12. When the transfer is complete, power off and disconnect the Snowball from AVA. A shipping label will appear on the E-Ink display. Ship the Snowball appliance(s) back to AWS for copying of the data to the cloud bucket.

13. Enter the following command to delete the current manifest file from AltaVault:

```
hostname (config)# seed snowball manifest-file delete
```

Note: Do not start replication service on AltaVault at this time. You must verify the data transfer on Amazon before enabling replication to the cloud.

14. After shipping all of the Snowball(s) used by AltaVault to Amazon, verify the data transfer in AWS and complete the seeding operation as described in [“Verifying and completing data transfer” on page 173](#).

Managing data transfers on AltaVault

AltaVault provides additional commands for monitoring and managing Snowball data transfer on the AltaVault.

To display the seeding progress to a Snowball

1. Enter the following command to display seed status from AltaVault to Snowball using the `show seed` command.

Status	Description
Not started	Snowball copy is not started. AltaVault can also enter this state after successfully stopping or canceling a job.
Copying data	Copying data to Snowball is in progress.

Status	Description
Paused	<p>The Snowball copy process can enter Paused state for a number of reasons:</p> <ul style="list-style-type: none"> • The seeding to Snowball is complete. • The seeding to Snowball exited or paused due to an error. • The seeding to Snowball is paused using the <code>seed snowball pause</code> command. • Verification exited due to an error, including exiting Storage Optimization Service.
Verifying imported data	Data has been imported into the Amazon cloud, and AltaVault is in the process of verifying the data transferred to the cloud.
Completed	The data has been transferred successfully to cloud.

2. To see detailed progress information, enter the following command:

```
hostname (config)# show seed detailed
```

AltaVault displays more detail with regard to each status.

To pause and resume a Snowball operation

1. Enter the following command to manually pause data transfer to Snowball:

```
hostname (config)# seed snowball pause
```

2. Enter the following command to resume data transfer to Snowball:

```
hostname (config)# seed snowball resume
```

Data transfer will pick up where it left off.

To cancel a Snowball operation

1. To cancel a Snowball data transfer job, enter the following command:

```
hostname (config)# seed snowball cancel
```

After canceling a job, AltaVault cleans up related files on the AltaVault appliance, including deletion of the manifest file. Any data transferred to Snowball remains and will be copied to the cloud bucket when the Snowball is returned to Amazon. If a Snowball job is canceled and the Snowball is re-used for a new seeding operation, the existing contents of the Snowball are overwritten with the latest copy.

Verifying and completing data transfer

After the Snowball appliance is returned, Amazon copies the data to the S3 bucket created when you ordered Snowball. If you set up job notifications, Amazon emails you with status at each event. When you are notified that the transfer is finished, you must verify that the transfer completed correctly and complete the Snowball actions on AltaVault.

To verify the data transfer

1. Check that Storage Optimization Service is running:

```
hostname (config)# show service
```

If service is not running, start the service:

```
hostname (config)# service enable
```

2. Check that the data transfer completed successfully using the following command:

```
hostname (config)# seed snowball verify
```

The AltaVault runs a series of checks to verify the information transferred to the cloud bucket matches what was copied to Snowball.

3. Enter the following command to check the verification status:

```
hostname (config)# show seed
```

If the transfer to the cloud was successful, the verification status changes to Completed. For a list of status messages, see [“Managing data transfers on AltaVault” on page 172](#).

4. After the verification is complete, instruct AltaVault to end Snowball actions by entering the following command:

```
hostname (config)# seed snowball stop
```

After stopping the process, AltaVault cleans up related files, including deletion of the manifest file.

5. Start the replication service on AltaVault:

```
hostname (config)# replication service enable
```

Normal backup operations to the cloud resume, including any data that was not copied to the cloud using Snowball.

CHAPTER 12 Migrating data to a new cloud

This chapter describes how to migrate data to from an existing cloud to a new cloud. It includes the following sections:

- [“Cloud migration overview” on page 175](#)
- [“Cloud-to-cloud migration” on page 176](#)
- [“Canceling cloud-to-cloud migration” on page 177](#)
- [“Amazon S3 or S3-IA to Glacier migration” on page 177](#)
- [“Amazon S3 to S3-IA or Amazon S3-IA to S3 migration” on page 178](#)

Cloud migration overview

It might be necessary to move AltaVault protected cloud data from one cloud to another some time after the original deployment. For example, there may be a strategy shift to move from a public cloud to a private cloud in order to keep data on premise rather than online.

AltaVault offers cloud migration to address this requirement.

For general cloud-to-cloud migration, AltaVault acts as a data replicator during the migration. As the data flows from the existing cloud, through AltaVault, and then on to the new cloud, AltaVault does not reprocess the data. Therefore, no data is evicted from the AltaVault cache during this process; the data simply flows through the networking components of the appliance. AltaVault also continues to accept data from backup applications, so no interruption to backup schedules occurs during migration. Upon completion, AltaVault automatically resumes replication of any pending data that was queued during the migration process.

Be aware that replication is suspended during cloud migration. Therefore, if the AltaVault cache becomes full (since data from the cache cannot be evicted if it is not replicated yet), then backup application attempts to write to the AltaVault cache will fail, leading to failure of backup applications.

In addition to general cloud-to-cloud migration, AltaVault can also perform direct cloud migration between different storage classes of Amazon under the same Amazon account and security credentials. This cloud migration occurs directly between Amazon cloud storage classes without sending the data back through AltaVault. As with general cloud-to-cloud migration, replication is suspended during cloud migration. Upon completion, reboot the AltaVault appliance to resume replication of any pending data queued during the migration process.

AltaVault does not support cloud migration in the following cases:

- Migration from Azure Standard to Azure Archive.
- Migration from Alibaba Archive, Amazon Glacier, and Azure Archive to any other cloud storage provider.

- Using user-defined cloud provider life-cycle policies configured in Amazon Web Services for moving data between storage classes. AltaVault has been optimized to maintain data integrity with specific provider life-cycle policies when migrating data between cloud providers or across storage classes (for example, S3 to S3-IA). To transfer data between providers or storage classes, use only the AltaVault migration procedures described in this chapter.

To migrate data to a new cloud, use one of the following procedures:

- [“Cloud-to-cloud migration” on page 176](#)
- [“Amazon S3 or S3-IA to Glacier migration” on page 177](#)
- [“Amazon S3 to S3-IA or Amazon S3-IA to S3 migration” on page 178](#)

Cloud-to-cloud migration

Note: If you are migrating between Amazon storage classes for the same security credentials, see [“Amazon S3 or S3-IA to Glacier migration” on page 177](#) or [“Amazon S3 to S3-IA or Amazon S3-IA to S3 migration” on page 178](#).

Use the following commands to perform the migration between cloud service providers.

1. Connect to the AltaVault command-line interface using SSH.

2. Enter the following commands to access configuration mode:

```
hostname > enable
hostname # configure terminal
hostname (config)#
```

3. Enter the credentials that you will use to connect to the new cloud:

- a. Enter the new cloud provider name, bucket name, cloud provider endpoint hostname, and port to which you are migrating:

```
CLI (config)# replication migrate-to provider type <provider-name> bucket-name <bucket-name>
hostname <host-name> port <port-value>
```

- b. Enter the authentication type and security credentials for the new cloud to which you are migrating:

```
CLI (config)# replication migrate-to auth type <authentication-type> acc-key-id <access-key>
secret-acc-key <secret-key>
```

- c. Optionally, enter the proxy name and credentials for the cloud to which you are migrating and enable the proxy:

```
CLI (config)# replication migrate-to proxy hostname <hostname> [port <port number>] [username
<username>] [password <password>]
```

```
CLI (config)# replication migrate-to proxy enable
```

4. Start cloud migration specifying the number of processes (threads) AltaVault uses to complete the transfer. The default is 128 threads:

```
CLI (config)# replication migrate-to enable [num-threads <value-1-to-128>]
```

5. Monitor cloud migration status. Enter the following command to show the estimated time until migration completes:


```
CLI (config)# show replication migrate-to estimate
```

Upon completion, AltaVault displays a message indicating the migration was successful along with the number of bytes transferred.

6. If migration fails for any reason, you can try continuing the migration from where it left off, or deleting the data in the new cloud and start over.

To restart migration from where it left off, enter the following command:

```
CLI (config)# replication migrate-to skip-existing
Replication will be paused and service will be restarted
Continue? [y/n]
```

To delete the data in the new cloud and start the migration over again, enter the following command:

```
CLI (config)# replication migrate-to format
Data in new cloud will be deleted. Do you want to continue? (y/N)
```

Canceling cloud-to-cloud migration

Use the following commands to cancel a migration that is already in progress.

1. Cancel the migration:

```
CLI (config)# replication migrate-to cancel
```

AltaVault displays a message prompting you to confirm cancellation of the job.

2. Select **Y** or **N** at the prompt. Select **N** to continue the migration. Select **Y** to end an in-progress migration. AltaVault displays messages describing possible next steps:

```
To format the new cloud bucket, use "replication migrate-to format" CLI command
To resume migration, use "replication migrate-to enable skip-existing" CLI command
```

If you decide to continue the migration again, enter the `replication migrate-to enable skip-existing` command to pick up where you left off, skipping data already transferred to the new bucket.

To wipe out any data transferred to the new bucket, enter the `replication migrate-to format` command.

Amazon S3 or S3-IA to Glacier migration

If you are migrating from Amazon S3 or S3-infrequent access storage classes to Amazon Glacier, perform the following steps:

1. Connect to the AltaVault command-line interface using SSH.
2. Enter the following commands to access configuration mode:

```
hostname > enable
hostname # configure terminal
hostname (config)#
```

3. Stop the AltaVault storage optimization service:

```
hostname (config) # no service enable
Terminating optimization service...
....
```

4. Start cloud migration:

```
CLI (config) # replication s3-to-glacier
Cloud based deduplication is currently enabled. Disabling cloud based deduplication could take
a few hours for a large cloud bucket. Disable? (y/N) y
Cloud based deduplication turned off
Successfully switched from S3 to Glacier
S3 prefixes already enabled
```

5. Start the AltaVault Storage Optimization Service:

```
hostname (config) # service enable
```

Amazon S3 to S3-IA or Amazon S3-IA to S3 migration

If you are migrating from Amazon S3 to S3-infrequent access storage classes or vice versa, perform the following steps:

1. Connect to the AltaVault command-line interface using SSH.**2. Enter the following commands to access configuration mode:**

```
hostname > enable
hostname # configure terminal
hostname (config)#
```

3. Stop the AltaVault storage optimization service

```
hostname (config) # no service enable
Terminating optimization service...
....
```

4. If you are migrating from Amazon S3 to S3-IA, enter the following command:

```
CLI (config)# replication provider type s3 storage-class standard-IA
```

If you are migrating from Amazon S3-IA to S3, enter the following command:

```
CLI (config)# replication provider type s3 storage-class standard
```

5. Start the AltaVault Storage Optimization Service:

```
hostname (config) # service enable
```

CHAPTER 13 Migrating data between appliances

This chapter includes the following sections:

- [“Data migration overview” on page 179](#)
- [“Data migration connection diagrams” on page 180](#)
- [“Data migration process” on page 182](#)
- [“Prerequisites” on page 182](#)
- [“Performing appliance data migration” on page 184](#)

Data migration overview

In some cases, it may be necessary migrate your data from one type of AltaVault appliance to another. This can happen when you upgrade hardware, or you want to migrate your data from a virtual appliance to a higher-capacity physical appliance. The following table describes the supported migration options and software requirements:

Source appliance	Target appliance	Software requirements: source/target
SteelStore physical or virtual appliance	AltaVault physical appliance (AVA400/800)	3.2.3/4.1.1 Note: For SteelStore appliances, you must first migrate to an AltaVault running version 4.1.1. After the migration to AltaVault 4.1.1 is complete, you can upgrade to later versions of AltaVault software.
AltaVault virtual appliance	AltaVault physical appliance (AVA400/800)	Source and target appliances must be running the same AltaVault version.
SteelStore virtual appliance	AltaVault virtual appliance	See <i>AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances</i> . Do not follow the instructions in this chapter.

In all cases, the cache size on the target appliance must be equal to or larger than the source. Additionally, for physical appliances the number of RAID groups on the target appliance must be equal to or larger than the source appliance.

Migrating data from one appliance to another involves a data migration tool to transfer data from a physical or virtual appliance to a physical AltaVault appliance. The data is transferred preserving shelf and RAID locality. For example, data from the controller on the source will be migrated to the first RAID group on shelf 1 on the target appliance. Data from the first shelf on the source will be migrated to the second RAID group on shelf 1 on the target appliance, and so on. During data migration both the source and target appliances will be unavailable for backup and recovery.


Data migration connection diagrams

Below are the recommended data migration connectivity diagrams for physical and virtual appliances. The connections you establish will depend on whether you have 1GbE or 10GbE connectivity in your network. Connectivity can include switches, but for best performance it is recommended to use direct connections between appliances using crossover cables. If connections use static routes, all interfaces can be on the same subnet. If connections use dynamic routing, NetApp recommends that interfaces be on separate subnets. The diagrams cover the following connections:


- [“SteelStore physical appliance to AltaVault physical appliance” on page 181](#)
- [“SteelStore or AltaVault virtual appliance to AltaVault physical appliance” on page 181](#)

SteelStore physical appliance to AltaVault physical appliance


10G Connectivity



1G connectivity



SteelStore or AltaVault virtual appliance to AltaVault physical appliance



Data migration process

The data migration process has multiple high-level steps as described in the following table:

Steps	Details
Validate prerequisite steps for source and target appliances	For more information, see “Prerequisites” on page 182 .
Install the new AltaVault target appliance	For information on the procedure, see the current version of the <i>NetApp AltaVault Cloud Integrated Storage Installation Guide for the Virtual Appliances</i> or <i>AltaVault System Installation and Setup Instructions</i> (poster)
Connect the source appliance to the target appliance	For more information, refer to the suggested network diagram configurations as shown in the “Data migration connection diagrams” on page 180 .
Run the migration tool to migrate data on to the target appliance	For more information on how to migrate, see “Performing appliance data migration” on page 184 .
Bring up the new appliance	For details, see “Post-data migration procedure” on page 185 .
Decommission the source appliance	The source appliance needs to be decommissioned according to your company’s practices regarding decommissioned hardware. If the source appliance will be reset and used in a new data protection capacity, contact NetApp support for assistance on steps needed to redeploy the appliance.

Prerequisites

The following describes the prerequisites required of the source and target appliances prior to performing data migration.

- [“Prerequisites for the source appliance” on page 183](#)
- [“Prerequisites for the target appliance” on page 184](#)

Prerequisites for the source appliance

Requirement	Description
RAID groups	<p>If the source is a physical appliance, total number of source storage shelves, including the source controller, must be equal to or less than the number of target RAID groups. For example, a SteelStore 3030 with 3 shelves would require an AltaVault appliance with at least 2 shelves, which in turn supports 4 RAID groups. This does not apply when the source is a virtual appliance.</p> <p>To gather shelf information using the CLI, open a SSH session to the source appliance and issue the following command:</p> <pre>enable configure terminal show shelves</pre>
Time zone	The time zone on the source appliance must match that of the target appliance.
Software version	<p>When SteelStore is the source appliance, it must be running version 3.2.3. To download the required software upgrade, go to https://mysupport.netapp.com.</p> <p>When the source is an AltaVault appliance, both source and target must be running the same software version.</p>
Appliance service	<p>Shut down the storage optimization service by navigating to Maintenance > Service page. To stop the service using the CLI, open an SSH session to the source appliance and issue the following commands:</p> <pre>enable configure terminal no service enable</pre>
Encryption key passphrase	If an encryption key passphrase was created on the source appliance, make sure to record the value. This passphrase is required for the data migration procedure.
Network connectivity	Setup network connectivity to the target appliance. Connection via multiple data network interfaces is recommended for faster data transfer and fault tolerance. For more details, see “Data migration connection diagrams” on page 180 .
Data integrity	<p>(Optional) If the source appliance has recently recovered from a non-fatal appliance error condition, such as an unexpected reboot, run the File System Check and Verify operations to validate that the data integrity has not been compromised prior to performing data migration.</p> <p>Note: File System Check and Verify operations can take several days to run depending on the amount of storage that must be scanned.</p>

Prerequisites for the target appliance

Requirement	Description
RAID groups	<p>The total number of target RAID groups must be equal to or greater than the number of source storage shelves, including the source controller. For example, a SteelStore 3030 with 3 shelves would require an AltaVault appliance with at least 2 shelves, which in turn supports 4 RAID groups. To gather shelf information using the CLI, open a SSH session to the source appliance and issue the following command:</p> <pre>enable configure terminal show shelves</pre>
Time zone	The time zone on the source appliance must match that of the target appliance.
Software version	<p>When the source is a SteelStore appliance, the target must be running AltaVault 4.1.1.</p> <p>When the source is an AltaVault appliance, both source and target must be running the same software version.</p>
Appliance service	<p>Shut down the storage optimization service. To stop the service using the CLI, open an SSH session to the source appliance and issue the following commands:</p> <pre>enable configure terminal no service enable</pre>
Network connectivity	The target appliance must be minimally configured with a management IP address and IP addresses for any data interfaces that the appliance will use to receive data during the data migration process. For details on configuring the management and data interfaces, see “Modifying management interfaces” on page 57 and “Modifying data interfaces” on page 58 .
Data integrity	<p>The target datastore must be empty. Use the following CLI commands to format the local datastore:</p> <pre>enable config terminal no service enable datastore format local</pre>

Performing appliance data migration

1. On the target appliance, start data migration using the following CLI commands. After executing these commands, you will be prompted to enter the encryption key passphrase used to setup the encryption key on the source appliance. This passphrase will be used to import the encryption key on the target appliance. If no passphrase was created, press Enter to continue.

```
<target appliance> > enable
<target appliance> # configure terminal
<target appliance (config)> # datamigration receive
```

2. On the source appliance, check whether cloud deduplication is enabled:

```
<source appliance> > enable
<source appliance> # configure terminal
<source appliance (config)> # no replication cloud-dedupe
```

If cloud deduplication is enabled, AltaVault displays a message asking if you want to disable cloud deduplication: If prompted, enter **N** (no).

The cloud deduplication status is needed in [Step 3](#) for determining whether to use the metadata-only option.

3. Once the target appliance is ready, data migration can be initiated on the source appliance. Data migration can transmit both the metadata and the data or just the metadata. If only metadata is migrated, all data must have been previously replicated to cloud storage by the source appliance. Migrating data and metadata requires additional time, but allows AltaVault to have immediate access to data in cache for quick restore. If the data migration is stopped or interrupted and restarted, then migration will continue from where it was interrupted.

(a). To transfer the data and metadata, use the following CLI commands on the source appliance to start data migration:

```
<source appliance> > enable
<source appliance> # configure terminal
<source appliance (config)> # datamigration send dest-ip <ipaddress1,ipaddress2...>
```

(b). To exclude data and only send metadata to the target appliance, use the following CLI command.

```
<source appliance (config)> # datamigration send dest-ip <ipaddress1,ipaddress2...>
metadata-only
```

With metadata-only option, the data pinned on the source appliance will not be transferred to the target appliance. Newly ingested data to pinned shares on the target appliance will continue to be pinned.

When data migration is performed with metadata-only option, the migration process completes faster than migrating both the data and metadata. After metadata migration is complete, deduplication of new backup and/or archival data to the target appliance will be similar to the source appliance. However, the metadata-only option is not optimized for data restore on the target appliance as it has not transferred the locally cached data.

One or more destination interface IP addresses can be provided as part of the data migration command. If more than one destination interface IP address can be provided in comma-separated form as part of the data migration across these interfaces

4. Once data migration is initiated on the source appliance, the tool will validate the state on the target appliance. This step can take up to 5 minutes per shelf on the source appliance. To display the status of data migration, use the following CLI command:

```
<source appliance (config)> # show datamigration status
```

5. To stop data migration tool on the source appliance, use the following command:

```
<source appliance (config)> # datamigration send stop
```

6. To stop data migration on the target appliance, use the following command:

```
<target appliance (config)> # datamigration receive stop
```

7. To reset and restart data migration from the beginning, issue the following CLI commands on both on the source and target appliances. After entry, repeat step 1 and step 2 to initiate a new data migration.

```
<target appliance (config)> # datamigration receive reset
<source appliance (config)> # datamigration send reset
```

Note: Data pinned on the source appliance will continue to be pinned on the target appliance.

Post-data migration procedure

Once data migration is complete you will need to complete the setup of the target appliance for production use.

To set up the target appliance for production

1. Disconnect the target appliance if it is connected directly to the source appliance, and connect it to the network infrastructure.
2. If required, reconfigure the data interfaces on the target appliance for the new network.

Note: The metadata-only migration option is not supported in cases where the source appliance has cloud deduplication disabled. In this case, the target appliance would receive the metadata for files in the local cache of the source, but not the actual data. The datastore cache contains dedupe hashes of only those data segments that are cached locally. A metadata-only transfer of an appliance that is configured with cloud deduplication turned off might cause data ingested into the new appliance to dedupe against data segments found only in cloud. If metadata-only option is used when cloud deduplication is disabled, run `no replication cloud dedupe force` command to delete hashes for data not cached in the target appliance.

3. Issue the following CLI commands on the target appliance to complete data migration and return the appliance to production use:

```
<target appliance (config)> # megastore guid reset
<target appliance (config)> # service restart
```

4. When migrating from a SteelStore appliance, after the migration to AltaVault 4.1.1 is complete, you can upgrade to later versions of AltaVault software. For AltaVault software upgrade instructions, see [“Upgrading your software” on page 126](#).

CHAPTER 14 Disaster recovery

Disaster recovery is the process of recovering the technology infrastructure critical to an organization after a natural or man-made disaster. AltaVault supports disaster recovery by enabling you to retrieve your data in case of a failure.

This chapter includes the following sections:

- “Disaster recovery preparations” on page 187
- “Disaster recovery testing” on page 188
- “Disaster recovery” on page 190

Disaster recovery preparations

You can enable AltaVault at the disaster recovery site to access backups that originated from an AltaVault at the affected data center. Depending on the data size, you can also use an AltaVault virtual appliance at the recovery site.

Note: You do not need a license to restore data in read-only mode in AltaVault. You can download AltaVault-v for free from the NetApp Support site at <https://mysupport.netapp.com> and use it to recover your data.

For example, consider a data center with AltaVault located at the Production Site (site A). The backup site is the DR (Site B), located in a different physical location (such as different city, country, or continent). If there is a disaster at Site A, the data still resides in the cloud. Site B contains a passive AltaVault that is not powered on.

You can also use AltaVault-v at Site B, depending on the size of the data that you need to restore. AltaVault-v can store data up to 32 TB. NetApp recommends that you use an appliance in the disaster recovery site (Site B) that has the same or greater local storage capacity as the affected AltaVault (in Site A). If the appliances at the two sites do not match, you can still initiate the recovery process; however, it recovers only as much data as the size of the storage on AltaVault at the disaster recovery site. If the recovery process attempts to bring back more data than the disaster recovery AltaVault can handle, then the recovery process might fail.

Exporting the configuration file

To prepare for disaster recovery, export your current configuration file from AltaVault at Site A, `altavault_config_(HOSTNAME)_(DATETIME).tgz`, and store it in a safe place, such as with your business continuity plans. For details on exporting your configuration file, see [“Using the export configuration wizard” on page 35](#).

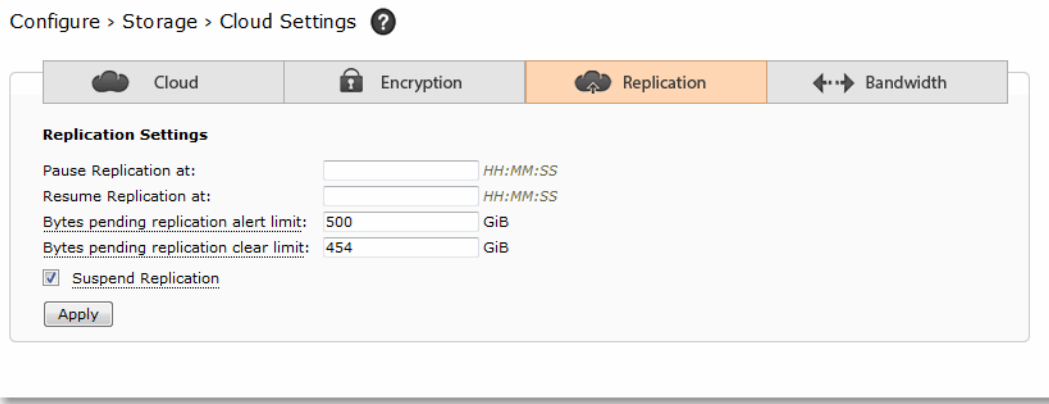
Disaster recovery testing

If you are restoring data for disaster recovery testing, you must first disable replication on AltaVault at site A and then restore your data at site B.

Suspending replication at the production site

To suspend replication

5. On the original production AltaVault appliance, suspend replication by selecting **Configure > Cloud Settings**.



6. Select the Replication tab.
7. To suspend replication, select the check box, Suspend Replication.
8. Click **Apply**.

Enabling AltaVault for a disaster recovery test

To recover your configuration to an AltaVault at site B

1. Log in to the appliance with the admin account using the serial console. If you are using a virtual AltaVault appliance, use the VM console to log in.
2. You can now configure the management IP address using the Configuration Wizard.

3. Use the IP address to connect to the appliance UI using a browser.
4. Log in to the appliance as `admin`. Provide the password configured during step 2.
5. Choose `Configure > Setup Wizard`.
6. Run the System Settings wizard to set the time zone to the same time zone as that of the original appliance.
7. Start the Import Configuration wizard.
8. Import to AltaVault in Site B, the configuration exported from the appliance in Site A:
 Select the Import Shared Data Only checkbox to import only common settings to this secondary AltaVault appliance. If you have configured the encryption key with an encryption key passphrase on the AltaVault in Site A, enter that passphrase into the encryption key passphrase field. After the import occurs, ensure that the new appliance in Site B uses the same cloud provider credentials, bucket name, and encryption key that Site A uses.
 For more information on wizard related steps, see [“Using the AltaVault configuration wizards” on page 17](#).
9. To test disaster recovery from a secondary site while the primary site is still alive, connect to the CLI by logging in as `admin` using SSH, and enter the following commands:

```
amnesiac > enable
amnesiac # configure terminal
amnesiac (config) # no service enable
amnesiac (config) # datastore format local
amnesiac (config) # replication dr-test enable
amnesiac (config) # service enable
amnesiac (config) # show service
```

The recovery process for both testing and an actual recovery can take anywhere from a few seconds to a few hours, depending on the backup(s) being restored. During the recovery process, the system communicates with the cloud provider and recovers all the namespace files that existed before the failure. The duration of this process depends on how many files you stored on AltaVault before the failure. Enter the **show service** command to determine the date and time until which the datastore has been replicated.

After your service restarts, you can browse to your share and see your files. Because the recovery process downloads only the namespace and metadata, initial file access might be slow, because AltaVault downloads all of the data from the cloud.

Data restoration for disaster recovery testing

The process for restoring data from AltaVault for a disaster recovery test is similar to the process for restoring data from AltaVault under normal conditions. However, because none of the data is local after recovering the AltaVault configuration on a new appliance, AltaVault must recover the data back from cloud storage when requested by the backup application or user accessing the files.

Data is recovered serially as designated by read requests of the files from the SMB shares, OST shares, SnapMirror shares, or NFS mounts. To speed up the recovery process, you can prepopulate data to the cache, which allows AltaVault to request recovery of data segments in parallel. For more information, see [“Restoring data from the cloud using the prepopulation page” on page 263](#).

Performing post-DR testing activities

After completing a disaster recovery test, perform the following steps to clean up the disaster recovery environment, and to re-enable operations on the production AltaVault appliance.

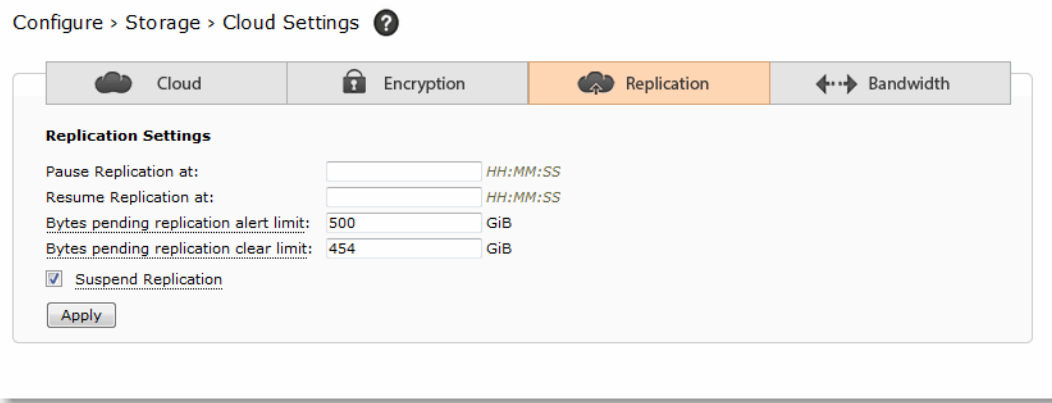
To revert back to original state after the DR test is complete

1. When disaster recovery testing is complete, the AltaVault appliance used for DR testing is no longer needed. To revert the appliance back to an initial, non-configured state for future DR tests and to release control of the cloud bucket, enter the following commands:

```
amnesiac > enable
amnesiac # configure terminal
amnesiac (config) # no service enable
amnesiac (config) # datastore format local
amnesiac (config) # reset factory
```

Note: The AltaVault appliance shuts down as the final action in the reset factory command.

2. On the original production AltaVault appliance, re-enable replication by selecting Configure > Cloud Settings.



3. Select the **Replication** tab.
4. To resume replication, clear the check box, **Suspend Replication**, and click **Apply**.
5. If pending data is awaiting to be replicated to the cloud, review the Reports > Back-End Throughput graph to confirm that data has resumed replication to cloud storage.

Disaster recovery

Disaster recovery actions must be performed as follows to enable a new AltaVault appliance for recovery of data backed up by the original AltaVault appliance lost in the disaster.

Enabling AltaVault for disaster recovery

To recover your configuration to an AltaVault at site B

1. Log in to the appliance using the admin account using the serial console. If you are using a virtual AltaVault appliance, use the VM console.
2. You can now configure the management IP address using the Configuration Wizard.
3. Use the IP address to connect to the appliance UI using a browser.
4. Log in to the appliance as `admin`. Provide the password configured during step 2.
5. Choose **Configure > Setup Wizard**.
6. Run the System Settings wizard to set the time zone to the same time zone as that of the original appliance.
7. Start the Import Configuration wizard.
8. Import to AltaVault in Site B, the configuration exported from the appliance in Site A:
 Use the checkbox, **Import Shared Data Only**, if you are importing settings to a new AltaVault appliance. If you have configured the encryption key with an encryption key passphrase on the AltaVault in Site A, enter that passphrase into the encryption key passphrase field. After the import occurs, ensure that the new appliance in Site B uses the same cloud provider credentials, bucket name, and encryption key that Site A uses.
 For more information on wizard related steps, see [Chapter 3, “Using the AltaVault configuration wizards.”](#)
9. To perform disaster recovery after a lost primary site, connect to the CLI by logging in as `admin` using SSH, and enter the following CLI commands:

```
amnesiac > enable
amnesiac # configure terminal
amnesiac (config) # no service enable
amnesiac (config) # datastore format local
amnesiac (config) # replication recovery enable
amnesiac (config) # service enable
amnesiac (config) # show service
```

Data restoration for disaster recovery

The process for restoring data from AltaVault after a disaster is similar to the process for restoring data from AltaVault under normal conditions. However, because none of the data is local after recovering the AltaVault configuration on a new appliance, AltaVault must recover the data back from cloud storage when requested by the backup application or user accessing the files.

Data is recovered serially as designated by read requests of the files from the SMB shares, OST shares, SnapMirror shares, or NFS mounts. To speed up the recovery process, you can prepopulate data to the cache, which allows AltaVault to request recovery of data segments in parallel. For more information, see [“Restoring data from the cloud using the prepopulation page” on page 263](#).

Cloud data available for recovery reflects the replication state of the AltaVault prior to DR. For example, if AltaVault was in the process of replicating 100 files to the cloud at the time of the disaster but only had time to replicate 75 of those files, the remaining files would not be available in the cloud for DR. Likewise, if AltaVault was pending deletion of 10 files when a disaster occurred, those 10 files would still be available in the cloud for DR operations.

CHAPTER 15 **System components AVA400, AVA800**

This section provides system specifications for the NetApp AltaVault appliance. It includes the following sections:

- “AltaVault appliance components” on page 193
- “Using LEDs to check the status of the system” on page 195
- “Field replaceable units” on page 197
- “Fan modules and their LEDs” on page 198
- “Power supplies and their LEDs” on page 201
- “Controller components and their LEDs” on page 203
- “Internal FRUs” on page 206

For a detailed list of system hardware components and specifications, see [NetApp Hardware Universe](#).

AltaVault appliance components

The AltaVault AVA400 or AVA800 chassis contains the following components:

- AVA400 or AVA800 controller
- Front panel with Power, Warning, and controller Activity LEDs
- Two field-replaceable 1300W, 100-240V AC auto-ranging, plug-in power supply units
- Three field-replaceable high-speed fan modules
- Rack mounting kit for standard racks

The AV10S disk shelf contains:

- Up to 24 hot-swappable 4TB (AVA400) or 6TB (AVA800) disk drives
- A front panel with Power, Warning, Activity, and Shelf ID LEDs
- Two field-replaceable 530W, 100-240V AC auto-ranging, plug-in power supply units
- 6Gb/s SAS IOMs for connectivity to the controller

- A rack mounting kit for standard racks (optional)

Note: Removing a drive out of a RAID group and moving it to a different RAID group is not supported. The disk drives have been formatted to a unique RAID configuration and should not be used for Hot-Swappable drive replacements.

System chassis specifications

The following table summarizes the physical specifications for the system chassis:

Description	Controller
Rack units	6 U
Height	10.2 in. (25.9 cm)
Width	Without mounting flanges: 17.6 in. (44.68 cm) With mounting flanges: 19 in. (48.26 cm)
Depth	24.3 in. (61.72 cm)
Weight	105 lbs (47.6 kg)

What you need to know about expansion shelves

You must understand the following information for a successful installation:

- The system uses RAID6 controllers, twelve drives per RAID group, two RAID groups per shelf (AVA400/AVA800).
- You can hot-swap failed hard drives on an AltaVault expansion shelf, but you must run a CLI command to add it to the RAID array. For more information see, *NetApp AltaVault Cloud Integrated Storage User's Guide*.
- You must add a drive to the RAID controller using the Management Console or the command-line interface. New expansion shelves cannot be added when system is running. Drives are not added to the RAID array automatically. You must shut down the system, add the drives, and restart the system.
- The two power supplies on an AltaVault expansion shelf are field replaceable units (FRUs).
- The power supplies have fans:
 - The PSU fans are for PSU cooling and maintaining chassis thermals.
 - The fans continue to run, even if the power supply unit (PSU) itself fails.
- If the PSU fails, do not remove it from the chassis until a replacement is immediately available. It is recommended to replace the PSU as soon as possible to restore full redundancy.


Using LEDs to check the status of the system

You can check the front of the bezel to verify that the power is turned on, controller is active, system is halted, or whether a fault has occurred in the chassis.

To check the status of the system

1. Locate the LEDs on the front of the bezel to verify the power is turned on, the controller is active, the system is halted, or whether a fault has occurred in the chassis.

The image on the left displays a system bezel. The diagram on the right displays the system LEDs on the bezel.






2. Use the following table to understand the system LEDs:

Component	Description
1	<p>System LEDs</p> <p>When the bezel is in place, the LEDs are arranged horizontally in the following left-to-right order:</p> <ul style="list-style-type: none"> • Power • Fault (System Attention LED) • Controller A activity (controller installed in the top bay only with no secondary controller) <hr/> <p>Note: When the bezel is removed, the system LEDs on the chassis are arranged in the same order as on the bezel, but in a top-to-bottom vertical orientation.</p> <hr/>


3. Use the following table to understand the AltaVault appliance system components on the chassis front:

Component	Description
Bezel	The bezel covers the front of the AltaVault system chassis. From the bottom of the front bezel, you can see four system LEDs (arranged horizontally) that indicate the status of the system.
Fan modules	<p>There are six redundant and hot-swappable fan modules installed in slots A1 through A3 and B1 through B3.</p> <p>The bottom row contains a blank.</p> <p>You have to remove the front bezel to see the fan modules. The Attention LED is set to Amber if the fan module is not working properly.</p>
System LEDs	<p>The chassis has four LEDs (arranged vertically) that indicate the status of the system:</p> <ul style="list-style-type: none"> • Power • Attention • Activity A (Controller) • Activity B (unused in AltaVault) <p>These LEDs are located behind the bezel. Remove the front bezel to see these LEDs.</p> <p>The system LEDs and the LEDs visible when the bezel is installed are the same, except that the bezel LEDs are aligned horizontally and the system LEDs are aligned vertically. Both sets of LEDs provide the same information about the system.</p>
Chassis handles	The chassis has two handles on each side that assist with lifting the system. The handles are integrated into the chassis and fold flush when not in use.

4. Use the following table to understand the behavior of the system LEDs:

LED label	LED name	LED status	Description
	Power	Green	At least one of the two power supplies (PSUs) is delivering power to the system.
		Clear (off)	Neither PSU is delivering power to the system.
	Fault (System Attention LED)	Amber	<p>The system is shut down or a fault occurred in the chassis. The error might be in a PSU, fan, or controller. The LED also is lit when there is an internal FRU failure, or the system is in maintenance mode.</p> <p>Note: You can check the fault LED on the back of each controller to see where the problem occurred.</p>
		Clear (off)	The system is operating normally.
	Activity	Blinking green	The AltaVault OS is running. The length of time that the light remains on is proportional to the controller's activity.
		Clear (off)	There is no activity on the controller.

5. Locate port LEDs on the rear of the controller to verify link and activity status:



6. Use the following table to identify the rear controller components:

Component	Description
10GbE Ports	
Link LED	Displays green when the port is linked and displays white when there is no link.
Activity LED	Displays orange when the port is active and displays white when there is no activity.
PCM	PCM Attention LED is used on the AltaVault controller for debugging purposes.
GbE and management ports	
Link LED	Displays green when the port is linked and displays white when there is no link.
Activity LED	Displays orange when the port is active and displays white when there is no activity.
Note: Private management ports are unused on the AltaVault controller.	

Field replaceable units

Field-replaceable units (FRUs) are components that can be replaced at your site. There are two types of FRUs:

- Hot-swappable FRUs: Component which can be replaced while the system is still powered.
- Non hot-swappable FRUs: Component which cannot be replaced while the system is still powered.

The following are available system FRUs:

- Fan modules: Hot-swappable
- Power supplies: Hot-swappable
- Controller module: Not hot-swappable

Each controller module has internal FRUs which are not hot-swappable. Use the `reload halt` command to ensure a clean shutdown, and then power off the controller to replace these FRUs.

The following are available internal FRUs:

- Boot media device
- RAID controller assemblies
- Real-Time Clock (RTC) coin battery
- System DIMM

Slot numbering and associated components

An AltaVault controller controls three fan modules. Controller A controls the fans in slots A1 through A3. If a controller is removed from the chassis, the fans in all three associated fan modules stop spinning.

Three fan modules are shipped preinstalled on the front of the chassis in the slots labeled A1 through A3. The chassis has six slots on the chassis front and four slots on the chassis rear. The bottom row of fan modules are replaced with a blank because these fan modules are not required.

To locate chassis components in an AltaVault system


1. You must remove the front bezel to see the slots that contain the fan modules in the chassis.
2. Use the information that follows to locate the slots and the components that they accept:

Location	Slot label	Slot component assignments	Description
Chassis front	A1	One fan module in each slot	The fan modules in these slots are controlled by the controller module installed in slot A on the chassis rear. Note: A blank is installed in Slots B1, B2, and B3.
	A2		
	A3		
Chassis rear	A	One controller module in this slot	AltaVault systems support only one controller module in the chassis. A blank is installed over slot B.
	B		
	1	An AC power supply is located in slots 1 and 2.	Both power supplies are preinstalled.
	2		

Fan modules and their LEDs

If the controller does not detect the presence of the associated fan module, the system fails to boot. Each fan module is a hot-swappable field-replaceable unit (FRU) that contains two fans.

Each fan module has a cam handle, a cam handle release latch, and an attention LED that indicates the status of the fan module. The image below displays the fan module component locations.




The following table displays the fan module components and descriptions:

Component	Description
1	Cam handle
2	Fan module
3	Cam handle release latch
4	Fan module attention LED

To locate the fan modules and their LEDs

1. Remove the front bezel.

2. Check the attention LED on each fan module to determine whether the fan module is working properly.



An amber LED light indicates a fan module failure:

Icon	LED	Color	Description	Corrective action
!	Fan module attention LED	Amber	The fan module is not working properly.	Remove the faulty fan module and install a replacement module.
		Clear (off)	No errors.	None

3. Hot swap the fan module as required.

The system fails to start if the controller does not detect the presence of the associated fan module. You must replace a fan module within two minutes of removing it from the chassis to minimize disruption to the system's airflow. If the fan module is not replaced within two minutes, the system shuts down to avoid overheating. For information on replacing fan modules, see [“Hot-swapping controller fan modules” on page 219](#).

Fan redundancy policy

The AltaVault system fan redundancy policy enables the system to continue operating with a single fan failure (one fan failure in any one fan module). Each fan module has two fans.

The following events indicate fan failure:

- The fans speed has fallen below the critical low threshold.
- One or both fans in the fan module have stopped spinning.
- The fan module is not present.
- The fans in the module are malfunctioning.

In all these events, the system issues a warning for fan failure on the system console and boosts the speed of the remaining fans to the maximum. The system behaves differently for single fan failure scenarios versus multiple fan failure scenarios.

If one fan in a fan module fails (of the three fan modules installed per controller), the system remains operational and an alert displays in the UI:

- **Single fan failures:** If one fan in a fan module fails (of the three fan modules installed per controller), the system remains operational and an alert displays in the UI (Reports > Alarms Status).
- **Multiple fan failures:** In the event that two fans within the same fan module or multiple fans across different fan modules fail simultaneously, the system shuts down within two minutes if no action is taken. The AltaVault OS boosts the speed of the remaining fans to maximum to prevent the chassis from overheating. Replace the faulty fan modules within two minutes of the Alarm event email. The alarm can be checked on the UI (Reports -> Alarms Status).

Power supplies and their LEDs

AltaVault systems ship with two AC power supplies preinstalled in the chassis, in the slots labeled 1 and 2. The AC power supplies are installed on the rear of the chassis. The power supplies are fully redundant, hot-swappable field-replaceable units (FRUs). The system remains operational even if one power supply fails.

Note: When you remove a failed power supply, you must replace it promptly to minimize disruption to the controller's airflow. The system continues to function normally. The AltaVault OS logs receive alert messages in the event log file about the degraded power supply until the power supply is replaced.

Each power supply contains the following:

- Two power supply LEDs
- One on/off power switch – Switch to the OFF ((O) position) before connecting or disconnecting power to the power supply.
- One power cord inlet
- One power cord retainer clip – Holds the power cord in place.
- One cam handle – Grip and move to an open position to remove and install the power supply.
- One cam handle release latch – Releases to unseat the power supply.
- Two integrated fans – Cools the power supplies


Power supply LED behaviors

This section describes the AltaVault system LEDs, their behavior, and information about the power supplies used in the system.

You must power off the controller before powering off the AVA10S shelves.

Locations and descriptions of system LEDs

Each AC power supply has two LEDs on its faceplate that display the status of the power supply. The image below displays the location of these LEDs.



The following table displays information about the power supply LEDs:

Component	Description
1	Power LED
2	Attention LED

The following table describes the behavior of the power supply LEDs. These LEDs help monitor the status of the power supply:

Icon	LED	Color	Description	Corrective action
!	Power supply attention LED	Amber	The power supply module is not working properly.	Remove the faulty fan module and install a replacement module.
		Clear (off)	No errors.	None.

The following table summarizes the specifications of the AC power supplies used in AltaVault systems:

Description	Value
AC input power frequency	50 to 60 Hz
AC input power voltage/operating range (auto-ranging)	100 to 120 VAC 200 to 240 VAC
AC output power	1300 W
PSU output power	1050W

Controller components and their LEDs

The controller module is the component of the AltaVault system that runs the AltaVault OS (AVOS) operating system and controls its disk subsystem. The AltaVault system comes with one controller. You install the AltaVault controller on the rear of the chassis, in the slot labeled A, also referred to as *Controller 1* or *Controller A*, the storage controller or the controller. The bottom slot is empty and contains a blank.

Note: AltaVault systems support only one controller module per chassis.

Each AltaVault controller contains the following:

- Components for removing and installing controller
- Ports and LEDs on the AltaVault controller
- Internal FRUs and their LEDs

Use the components in this table to help remove or install a controller in a chassis:


Component	Description
Cam handle	Use to install and remove the controller handle.
Thumb screw	Use to secure the controller module in the chassis.
Release latch stop	Prevents the controller from sliding out of the chassis by stopping the controller mid-way during removal. Press this latch stop to remove the controller from the system.

Controller LED behaviors

The LEDs on the face plate of the controller display the status of its network or disk shelf connections and identifies the controller where a fault has occurred. To aid in understanding, the controller faceplate information is divided into two sections, the left and right side of the controller faceplate.

Controller left side ports and LEDs





This section describes the controller ports and LEDs on the left side of the controller faceplate.



The following table describes ports and LEDs on the left side of the controller:


Component	Description
1	SAS port LEDs (not used)
2	SAS port (not used)
3	10 GbE port LEDs
4	10 GbE Ports
5	CNA port LEDs (not used)
6	CNA ports (not used)
7	NVRAM LED (not used) Note: The NVRAM is physically present, but not used in the AltaVault software.
8	Controller attention LED

The following table describes the ports and the status of the port LEDs:

Port labels	Port descriptions	LED labels	LED status	LED descriptions
e0a through e0d and 	10 GbE ports Each controller has four 10Gb Ethernet (10GbE) ports, identified with labels e0a, e0b, e0c, and e0d.	LNK	Green	A link is established between the port and some upstream device.
			Amber flashing	Traffic is flowing over the connection.
			Off	No traffic is flowing over the connection.
	Controller attention LED		Amber	A problem has occurred in the controller. This in turn has caused the system attention LED on the chassis front to become illuminated.
			Off	The controller is functioning properly.





Controller right side ports and LEDs

This section describes the ports and LEDs on the right side of the controller faceplate.



Component	Description
1	1GbE port LEDs
2	1GbE ports
3	Management Ethernet port (wrench icon) LEDs
4	Management Ethernet port (wrench icon)
5	Private management Ethernet port (wrench lock icon) LEDs (not used)
6	Private management Ethernet port (wrench lock icon) (not used)
7	USB port (not used)
8	Console port

The following table describes the ports in use and the status of port LEDs:

Port label	Port description	LED labels	LED status	Description
e0i through e0l and 	GbE ports Each controller has four 1Gb Ethernet (GbE) ports, identified with labels e0i, e0j, e0k, and e0l.	LNK (Left LED)	Green	A link is established between the port and some upstream device.
			Off	No link is established.
		Activity (Right LED)	Amber flashing	Traffic is flowing over the connection.
			Off	No traffic is flowing over the connection.
 and 	Remote management The remote management port is labeled with a wrench symbol. It is identified as (primary) e0M in the CLI commands and output. This port is used for managing the AltaVault system.	LNK (Left LED)	Green	A link is established between the port and some upstream device.
			Off	No link is established.
		Activity (Right LED)	Amber flashing	Traffic is flowing over the connection.
			Off	No traffic is flowing over the connection.
	Serial console port The RJ-45 console port enables you to communicate with the controller directly even when a network connection is not available or when you set up the system for the first time by using the CLI.	No LED on this port.		

Internal FRUs

The AltaVault controller has many internal FRUs. These FRUs are not hot-swappable and require the system to be powered off before replacing any of these components.

The controller has a FRU map printed on the CPU cover. This FRU map lists all the internal FRUs in the controller and their general location on the board and their corresponding LEDs.

The following table lists the AltaVault controller's internal FRUs:

Internal FRU	Description	Attention LED Label
Boot device	The boot device stores a primary and secondary set of system files (also called the boot image) that the system uses when it starts. The attention LED is next to the boot device on the motherboard.	Not applicable
System DIMM	There are eight 32G DIMMs. All eight DIMMs must be installed for the system to be fully functional.	Not applicable
RAID Controller assemblies	The RAID controller assemblies provide reliability, high performance, and fault-tolerant disk subsystem management. The RAID controller assemblies are installed in PCIe slots 1 and 3.	Not applicable
Real-time clock (RTC) coin battery	The real-time clock (RTC) coin battery provides power to the real-time clock so that time synchronization functions continue to function properly. The attention LED is next to the battery on Riser-R.	Not applicable

For more information on replacing specific internal FRUs, see [“Replacing internal FRUs” on page 228](#).

CHAPTER 16 System maintenance AVA400, AVA800

This section provides information on maintaining AltaVault system components. It includes the following sections:

- “Shutting down the AltaVault controller” on page 212
- “Shutting down the AltaVault controller” on page 212
- “Replacing controllers” on page 213
- “Installing a controller in a chassis” on page 216
- “Replacing a controller chassis” on page 218
- “Hot-swapping controller fan modules” on page 219
- “Hot-swapping controller power supplies” on page 222
- “Adding disk shelves” on page 225
- “Changing the shelf ID for a disk shelf” on page 225
- “Adding an additional RAID group to a configured appliance” on page 227
- “Replacing internal FRUs” on page 228
- “Replacing a boot device in a controller” on page 229
- “Replacing system DIMMs” on page 233
- “Replacing RAID controllers” on page 237
- “Replacing the RTC clock coin battery” on page 237
- “Replacing disk shelf power supplies and other FRUs” on page 240
- “Replacing a faulty hard disk drive on an AltaVault AVA400 or AVA800 appliance” on page 228
- “Returning failed parts” on page 240
- “Disposing of batteries” on page 240

For a list of system hardware components and specifications, see [NetApp Hardware Universe](#).

Note: Best practices calls for allowing the replication queue to drain fully to the cloud prior to engaging in maintenance activities.

Remote management port setup

Setting the Service Processor password

The Service Processor (SP) is a subcomponent of the controller. There is no password on the Service Processor until you set the password. The Service Processor password must be set separately from the AltaVault OS administrator password.

To set the SP password

1. Log in to the CLI using either the serial connection, or an SSH session using the configured primary management interface set up using the CLI configuration wizard steps. For more information, see [“Using the AltaVault appliance CLI configuration wizard” on page 17](#).
2. Enter configuration mode and set the Service Processor password in the controller using the following CLI commands:

```
CLI > enable
CLI # configure terminal
CLI (config) # sp password set
```
3. Use the new password when logging into the Service Processor.

Configuring the remote management port

Access to the Service Processor for remote management is accomplished by connecting through the serial console port or an SSH session using the configured primary management interface.

The Service Processor is responsible for monitoring sensors, managing the physical environment of the system, capturing events, logs and forensics, and sending notifications and alerts. It also provides remote management features for administrators.

Before you begin, ensure that you have serial console access to the AltaVault appliance.

To setup the remote management port

1. Log in to the CLI using the serial connection.
2. At the CLI prompt, reboot the appliance and interrupt the AUTOBOOT process by pressing CTRL + C. This halts the appliance at the LOADER prompt.

```
(config) # reload
Rebooting...
.....
.....
Starting AUTOBOOT press Ctrl-C to abort...
***Hit CTRL+C***
Autoboot of PRIMARY image aborted by user.
LOADER-A>
```

3. At the LOADER prompt, enter the command, `sp setup`, to setup the IP address for the Service Processor.

```
LOADER-A> sp setup
```

4. Based on your setup, choose one of the following:

- To enable DHCP, enter the following:

```
LOADER-A> sp setup
Would you like to configure the SP? [y/n] y
Would you like to enable DHCP on the SP LAN interface? [y/n] y
Do you want to enable IPv6 on the SP? [y/n] n
```

```
Output sample:
Service Processor New Network Configuration
Ethernet Link:      up, full duplex, auto-neg complete
Mgmt MAC Address: 00:A0:98:54:F9:F6
IPv4 Settings
Using DHCP:        YES
IP Address: 172.16.33.154
Netmask: 255.255.252.0
Gateway: 172.16.33.1
IPv6: Disabled
```

Note: Make a note of the IP address. This IP address is used in subsequent steps to validate the Remote Management Port.

- If you do not want to enable DHCP, enter the IP address, netmask, and gateway address from the command line. The IP address of the SP interface must be on the same subnet as the primary interface. The SP interface and primary interface are connected to the wrench port on the AltaVault by means of an internal switch.

```
Would you like to configure the SP? [y/n] y
Would you like to enable DHCP on the SP LAN interface? [y/n] n
Please enter the IP address for the SP [unknown]: 172.16.100.4
Please enter the netmask for the SP [unknown]: 255.255.255.0
Please enter the IP address for the SP gateway [unknown]: 172.16.100.1
Do you want to enable IPv6 on the SP? [y/n] n
```

```
Output sample:
Service Processor New Network Configuration
Ethernet Link:      up, full duplex, auto-neg complete
Mgmt MAC Address: 00:A0:98:5D:34:BC
IPv4 Settings
Using DHCP:NO
IP Address:172.16.100.4
Netmask:255.255.255.0
Gateway:172.16.100.1
IPv6: Disabled
```

5. After the Service Processor is configured, run the command, autoboot, at the LOADER prompt to start the AltaVault appliance.

```
LOADER-A> autoboot
```

Validating the remote management port

The following commands show an example of how to validate the remote management port.

1. Connect to the service processor using one of the following methods:
 - Open a telnet connection to the serial console port using a terminal program to log in to the AltaVault command line prompt. Press Ctrl+G to get into Service Processor mode. To exit service processor mode, press Ctrl+D.

or

- ssh to service processor using the IP address obtained while configuring the remote management port.

2. Obtain the Service Processor IP configuration of the remote appliance:

```
SP> sp status
Firmware Version: 3.0.2
Debug Mode: Enabled
Mgmt MAC Address: 00:A0:98:65:03:24
Ethernet Link: Up, 1000Mb, Full-Duplex, Auto-neg enabled,completed
Using DHCP: no
IPv4 configuration:
IP Address: 172.16.33.154
Netmask: 255.255.240.0
Gateway: 172.16.33.1
IPv6 configuration: Disabled
```

3. Enter the following IPMI command to test the service processor feature:

```
SP> system sensors
```

Sample output (truncated):

Sensor Name	Current	Unit	Status	LCR	LNC	UNC	UCR
CPU0_Temp_Margin	-65.000	degrees C	ok	na	na	-5.000	0.000
CPU1_Temp_Margin	-65.000	degrees C	ok	na	na	-5.000	0.000
In_Flow_Temp	23.000	degrees C	ok	0.000	10.000	53.000	63.000
Out_Flow_Temp	35.000	degrees C	ok	0.000	10.000	61.000	71.000
Smart_Bat_Temp	31.000	degrees C	ok	0.000	10.000	59.000	69.000
CPU0_Error	0x0	discrete	Deasserted	na	na	na	na
CPU0_Therm_Trip	0x0	discrete	Deasserted	na	na	na	na
CPU0_Hot	0x0	discrete	Deasserted	na	na	na	na
Memory0_Hot	0x0	discrete	Deasserted	na	na	na	na
CPU1_Error	0x0	discrete	Deasserted	na	na	na	na
CPU1_Therm_Trip	0x0	discrete	Deasserted	na	na	na	na
CPU1_Hot	0x0	discrete	Deasserted	na	na	na	na
Memory1_Hot	0x0	discrete	Deasserted	na	na	na	na
PCH_Hot	0x0	discrete	Deasserted	na	na	na	na
P5V_STBY	5.026	volts	ok	4.246	4.343	5.661	5.807
P3V3_STBY	3.296	volts	ok	2.960	3.040	3.568	3.664
P1V8_STBY	1.804	volts	ok	1.630	1.659	1.950	1.969
P1V2_STBY	1.193	volts	ok	1.086	1.106	1.300	1.319
P0V9_STBY	0.892	volts	ok	0.805	0.854	0.951	0.999
P5V	5.051	volts	ok	4.246	4.343	5.661	5.807
P3V3	3.280	volts	ok	2.960	3.040	3.568	3.664
PVDDQ_DDR3_AB	1.339	volts	ok	0.010	0.019	2.454	2.464
PVTT_DDR3_AB	0.660	volts	ok	0.010	0.019	2.454	2.464
PVCCP_CPU0	0.980	volts	ok	0.010	0.019	2.454	2.464
PVDDQ_DDR3_CD	1.339	volts	ok	0.010	0.019	2.454	2.464
NVRAM_PG_3.3V	0x0	discrete	Asserted	na	na	na	na
NVRAM_PG_3.0V	0x0	discrete	Asserted	na	na	na	na
NVRAM_PG_1.8V	0x0	discrete	Asserted	na	na	na	na
NVRAM_PG_1.35V	0x0	discrete	Asserted	na	na	na	na
NVRAM_PG_2.5V	0x0	discrete	Asserted	na	na	na	na

Shutting down the AltaVault controller

An AltaVault controller has field-replaceable units (FRUs) that are not hot-swappable. For a list of AltaVault system FRUs, see [“Field replaceable units” on page 197](#).

When you are replacing field-replaceable units (FRUs) that are not hot-swappable or when you are replacing the controller itself, shut down the controller. A clean shutdown of the AltaVault controller ensures that all data has been written to the storage subsystems. You must also disconnect power from the power supplies before these replacement procedures.

Important: You must power off the controller before powering off the shelf.

To shut down a stand-alone AltaVault controller

1. To shut down the AltaVault controller, enter the following CLI command:

```
CLI> reload halt
```

When the shutdown is successful, the box powers off, there are no prompts or output to the console.

2. If *you* are not already grounded, properly ground *yourself*.
3. Turn off the power supplies by turning the on/off switch to the OFF (O) position. If the power source for this power supply has a on/off switch, turn the switch to the OFF (O) position.
4. Unplug the power cords from the power supplies and the power source:
 - Pinch the tab on the locking mechanism of the cable retainer clip, and open the retainer clip.
 - Slide the retainer clip off the cord.
 - Unplug the power cord from the power supply. Then, unplug the cord from the power source.
 - Repeat Step 4 for the second power supply.

Replacing controllers

These instructions describe how to replace an AltaVault controller in the case of an RMA (Return Merchandise Authorization).

Caution: In most cases, replacing an AltaVault controller is performed by NetApp Support. Only experienced technicians should attempt to replace a controller because of the risk of losing data.

If you are replacing your AltaVault Controller because of a faulty appliance, NetApp recommends that you preserve the existing disk cache so that you do not need to download large amounts of data from the cloud, which can be expensive in terms of time, bandwidth, and costs.

When an AltaVault Controller is replaced by an RMA, you must physically swap the RAID cards between the old controller and the new one (provided the fault is not related to the RAID card assembly). Since both models are exactly the same, the total size of the disks remains unchanged. The disks must be installed in order, populating shelf space 1 through 11, they cannot be jumbled.

To preserve your data, ensure that you export the current configuration in your source AltaVault appliance using the Export Configuration Wizard (choose Settings > Setup Wizard from the Management Console). NetApp recommends that you export your configuration and store it in a secure location as a preventative measure against unexpected failures.

Replacing an AltaVault controller includes the following processes:


- [“Shutting down and removing the controller from the chassis” on page 214](#)
- [“Moving working FRUs from the controller” on page 215](#)
- [“Reinstalling and connecting the controller in the chassis” on page 215](#)

Shutting down and removing the controller from the chassis

Before shutting down the controller, you must export the controller's current configuration, shut down the controller, and remove it from the chassis.

To shut down the controller and remove it from the chassis:

1. Ensure that the current configuration from the source appliance has been exported. For details, see the section, *Using the Export Configuration Wizard*, in the *NetApp AltaVault Cloud Integrated Storage User's Guide*. Ensure that the exported configuration from the source appliance is safe and accessible.
2. Shut down the source AltaVault appliance by selecting the following from the Management Console:
 - a. Select Settings > Reboot/Shutdown.
 - b. Click **Shutdown**.
3. If you are not already grounded, properly ground yourself.
4. Unplug the various system cables and SFPs from the controller module. Keep track of where the cables were connected so that when you reinstall the controller, you can reattach the cables in the same configuration.
5. Loosen the thumbscrew on the cam handle.



6. Identify components as described in this table:

component	Description
1	Thumbscrew
2	Cam Handle

7. Pull the cam handle downward to unseat the controller and slide the controller module out of the chassis until it catches. Then, press the release latch on the left side of the controller module and slide the controller module completely out of the chassis, making sure that you support the base of the module with your free hand.
8. Place the controller module on a clean, flat surface.

Moving working FRUs from the controller

Remove all the working field replaceable units (FRUs) from the impaired controller and move them to the new controller module. To reduce the possibility of damage to the components that are being moved, you should minimize the handling of the components by installing them in the replacement controller as soon as they are removed from the impaired controller. Then, install the replacement controller module in the system chassis.

To move all the working FRUs from the impaired controller to the new controller module

1. Place the replacement controller module next to the impaired module and open the CPU cover and the left and right side panels to install system components.

Note: Do not power on your system until you have moved all the internal FRUs to the replacement module.

2. Move the eight system DIMMs to the replacement controller module.
 - For instructions on replacing the system DIMMs, see [“Replacing system DIMMs” on page 233](#).
 - Install the eight DIMMs using the steps described in the procedure, [“Installing system DIMMs” on page 235](#).
3. Move the boot device to the replacement controller module
 - a. For instructions on removing the boot media from the impaired module, see [“Replacing a boot device in a controller” on page 229](#).
 - b. Open the boot device cover in the replacement controller module. Install the boot media in the boot media holder using the following procedure, [“Installing a boot device” on page 231](#).
 - c. If the boot device is damaged or corrupted during the replacement process, contact technical support to assist you in restoring the system to a healthy state.

Reinstalling and connecting the controller in the chassis

When you finish moving the internal FRUs to the replacement control, reinstall the controller module in the chassis.

To reinstall the controller module in the chassis

1. Use the steps described in the procedure, [“Installing a controller in a chassis” on page 216](#) to reinstall the controller module in the chassis.
2. Swap the RAID card assembly from the source AltaVault controller to the target controller. Ensure that you move the RAID cards from the source controller into the corresponding slots on the target controller. Do not jumble the cards as position and order are important. Keep track of which RAID cards you have swapped.
3. Reconnect the Storage Shelves to the target controller in the same configuration as the source controller.
4. Restart the target AltaVault appliance. After swapping the disks, when you restart the target AltaVault appliance, the following message displays:

The secure vault fails to unlock. This is because of the serial number mismatch between the appliances.
You cannot use the data store because you cannot access the data store encryption key in the secure vault.
5. Connect to the target AltaVault appliance through the serial cable.

6. Log in using the default login admin and password.
7. At the command line, enter the following CLI commands:

```
CLI> enable
CLI> configuration terminal
CLI> update controller config
```

The CLI command, `update controller config`, does the following:

- Updates the add-on RAID groups with the new controller configuration.
 - Clears the secure vault. It is safe to run this CLI command because the information in the secure vault is already backed up.
 - Restarts the system.
8. After AltaVault appliance starts, import the previously-saved configuration into the target AltaVault appliance using the Import Configuration Wizard.
 9. Reset the Megastore GUID by entering the following CLI commands at the command line:

```
CLI> enable
CLI> configuration terminal
CLI> megastore guid reset
```

The last command generates a new megastore GUID based on the serial number of the target AltaVault appliance. It is important to perform this step before restarting the storage optimization service.

Installing a controller in a chassis

After internal FRU replacement or movement tasks are complete, reinstall the controller module in the system chassis.

Before you begin

- Ensure that the chassis enclosure is securely installed in the rack or cabinet.
- Ensure the three fan modules associated with the controller module are installed in the system using the steps described in the procedure, [“Installing fan modules” on page 221](#).
- To ensure that the system has rebooted correctly, verify that the System attention LED on the chassis front and the Controller attention LED on the chassis rear are not lit after the system has finished rebooting.


To install a Controller in a Chassis

1. Align the controller module with the opening in the chassis and then gently push the controller module halfway into the system.
2. Re-cable the system:
 - Re-cable the QSFP connections to the disk shelf and the mini-SASHD connections to the controller. Remember to reinstall the media converters (SFPs) for FC cables.
 - Re-cable the management and console port connections.

3. With the cam handle in the open position, firmly push the controller module into the chassis until the controller module meets the mid-plane.

Note: Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors on the rear of the module.

4. Close the cam handle so that the latch clicks into the locked position and the controller module is fully seated in the chassis. Tighten the thumbscrew.



5. Identify components as described in this table:

Component	Description
1	Thumbscrew
2	Cam handle

6. Reconnect the power cables to the power supplies and secure them using the cable retaining clips.
7. Reconnect the power cables to the power source.

Note: Power on the disk shelf attached to the AltaVault system and set the disk shelf ID before powering on the controller. For information on setting the disk shelf ID, see the “*NetApp AltaVault Cloud Integrated Storage Poster*.”

8. Turn the on/off switch on the power source and the power supply to the ON (I) position to start the boot process for the system.
9. If all the FRU replacements were successful, the system should boot normally.
 - To ensure that the system has booted correctly, verify that the System attention LED on the chassis front and Controller attention LED on chassis rear are not lit after the system has finished booting.
 - Verify the health of your system and the system configuration using the command, `show info`.

Regarding the boot device:

- If the boot device is replaced, you must contact technical support to assist in the initialization of the device.

- If the boot device is damaged or corrupted during the replacement process, contact technical support for help in restoring the system to a healthy state.

Replacing a controller chassis

The chassis is the external metal casing that houses the controller module, power supplies and fan modules. The chassis has mounting flanges that are used to install the chassis in a rack or cabinet. The chassis is not a hot-swappable field-replaceable unit (FRU); you must shut down your system before replacing the chassis.

To replace a chassis for an AltaVault system, you must remove the power supplies, fan modules, and controller module from the old chassis, remove the old chassis from the rack or cabinet, install the new chassis, and then reinstall the components in the new chassis.

Before you begin

- You must perform a clean system shutdown to ensure that all data has been written to the storage subsystems. Use the `reload halt` command to shut down the system.
- If you cannot gracefully shut down the system, contact technical support for assistance. Power off the system and disconnect power from it. For detailed instructions on powering off the system and disconnecting the power as described in [“Shutting down the AltaVault controller” on page 212](#).

To remove a chassis

1. If *you* are not properly grounded, properly *ground yourself*.
2. Remove the two power supplies installed in the chassis using the steps described in the procedure, [“Removing power supplies” on page 222](#).
3. Set the power supplies aside. You will re-install them in the replacement chassis.
4. Removing the controller module from the chassis requires assistance from NetApp Support.
5. Set the controller module aside. You will re-install it in the replacement chassis.
6. Remove the fans modules installed on the chassis front using the steps described in the procedure, [“Removing fan modules” on page 220](#).
7. Set the fan module aside. You will re-install them in the replacement chassis.
8. Remove the screws from the chassis mount points.
9. With the help of two or three people, slide the empty chassis off the rails and set it aside.

Installing the chassis

Before you begin, verify that the replacement chassis you are installing is an approved part from NetApp.

To install the chassis

1. With the help of two or three people, slide the replacement chassis into the rack or cabinet by guiding the chassis onto the rails or brackets.

2. Slide the chassis all the way into the rack or cabinet.
3. Secure the front of the chassis to the rack or cabinet, using the screws you removed in the previous procedure.
4. Install the fan modules in the chassis as described in [“Installing fan modules” on page 221](#).

Note: You must install all three fan modules for a controller. If all three fan modules associated with a controller are not detected when the system boots, the system will crash. For details, see [“Slot numbering and associated components” on page 198](#).

5. Install both power supplies in the chassis using the steps described in the procedure, [“Installing power supplies on a controller” on page 224](#).

Note: Do not turn on the power supplies at this time.

6. Reinstall the controller module in the chassis as described in [“Installing a controller in a chassis” on page 216](#). AltaVault systems support only one controller per chassis. Install the blank panel over the second slot.
7. Turn the power switch on both power supplies and the power source to the ON (I) position. The system will start booting.
8. To ensure that the system has booted correctly, verify that the System attention LED on the chassis front and Controller attention LED on chassis rear are not lit after the system has finished booting.

Hot-swapping controller fan modules

Fan modules in an AltaVault system are hot-swappable field-replaceable units (FRUs). You can hot-swap the fan modules while the system is powered on, without disrupting the normal operation of the system.

When a fan module fails, the system logs messages in the event log file and an alarm is raised, indicating which power supply has failed.

Note: When a fan module fails, an alarm is raised. The UI shows a degraded state with the appropriate alarm information, and if configured, an email is sent to the administrator.

Removing fan modules

You must replace the fan module within two minutes of removing it from the chassis to minimize disruption to the system's airflow. System airflow is disrupted and the controller module associated with the failed fan module shuts down after two minutes to avoid overheating. To understand AltaVault system behavior for single and multiple fan failure scenarios, see [“Fan redundancy policy” on page 200](#).


To remove a fan module

1. If *you* are not already grounded, properly *ground yourself*.
2. With two hands, grasp the openings on each side of the bezel and pull it toward you until the bezel releases from the four ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the error messages and looking at the attention LED on each fan module cam handle.

Note: If the fan attention LED is lit solid amber, the fan module has failed.

4. Press down the release latch on the fan module cam handle and pull the cam handle downward to unseat the fan module from the chassis.
5. Pull the fan module straight out from the chassis, as shown below. Make sure that you support the base of the fan module with your free hand, so that it does not fall out of the chassis.

Note: Controller fan modules are short. Always support the bottom of the module with your free hand, so that it does not suddenly drop free from the chassis.



6. Identify fan module components as described in this table:

Component	Description
1	Cam handle
2	Fan module
3	Cam handle release latch
4	Fan module attention LED


7. Set the failed fan module aside.
8. Before removing another fan module from the chassis, install a replacement fan module using the steps described in the procedure, [“Installing fan modules” on page 221](#).

Installing fan modules

Before you begin, verify that the controller fan module that you are installing is supported by your controller model.

To install a fan module

1. If *you* are not already grounded, properly *ground yourself*.
2. Insert the replacement fan module into the chassis by aligning it with the opening and sliding it into the chassis.
3. Push firmly on the fan module housing to ensure that it is seated all the way into the chassis. The cam handle raises slightly when the fan module is completely seated.
4. Swing the cam handle up to its closed position, as shown below. Make sure that the cam handle release latch clicks into the locked position.



5. Identify fan module components as described in this table:

Component	Description
1	Cam handle
2	Fan module
3	Cam handle release latch
4	Fan module attention LED

6. Repeat the procedure for the remaining fan modules, if any.
7. After all fan modules have been replaced, align the bezel with the ball studs on the chassis and gently push it onto the ball studs.

Hot-swapping controller power supplies

Power supply units (PSUs) in a AltaVault system are auto-ranging, redundant, hot-swappable field-replaceable units (FRUs). You can hot-swap the power supplies on a controller while the system is powered on without disrupting the normal operation of the system.

Note: When a power supply fails, an alarm is raised. The UI shows a degraded state with the appropriate alarm information, and if configured, an email is sent to the administrator.

Before you begin

- If you are replacing more than one power supply on a controller, you must do so one at a time to prevent system downtime.
- If you must remove all power supplies, leaving the controller without any power, you must first shut down and then power off the system using the steps described in the procedure, [“Shutting down the AltaVault controller” on page 212](#).


Removing power supplies

This section describes how to remove a power supply.

Note: Replace a failed power supply promptly to minimize disruption to the controller's airflow. The system continues to function normally, but the AltaVault OS logs alert messages in the event log file about the degraded power supply until the power supply is replaced.

To remove a power supply

1. Identify the power supply you want to replace, based on the event log messages or through the amber fault LED on the power supply.
2. If you are not already grounded, properly *ground yourself*.
3. Turn off the target power supply by turning the on/off switch to the OFF (O) position. If the power source for this power supply has a on/off switch, turn the switch to the OFF (O) position.
4. Remove the power cord from the power supply, using the image below as a reference.
 - Pinch the tab on the locking mechanism of the power cord retainer clip, and open the retainer clip.
 - Slide the retainer clip off the cord.
 - Unplug the power cord from the power supply and the power source.
5. Press down the release latch on the power supply cam handle to unseat the power supply.
6. Lower the cam handle to the fully open position, and then slide the power supply out of the chassis, as shown below. Make sure that you support the power supply with your free hand.



7. Identify power supply components as described in this table:

Component	Description
1	Power supply
2	Cam handle release latch
3	Power and fault LEDs
4	Cam handle
5	Power cord locking mechanism


Installing power supplies on a controller

Before you begin, verify that the power supply you are installing is supported by your controller model.

To install a power supply

1. Verify that the on/off switch on the power supply is in the OFF (O) position. If the power source has an on/off switch, ensure that it is also set to the OFF (O) position.
2. If *you* are not already grounded, properly *ground yourself*.
3. With the cam handle in the open position, align the edges of the power supply with the opening in the system chassis and gently push the power supply into the chassis until it is almost flush with the chassis, as shown below.

Important: Do not use excessive force when sliding the power supply into the chassis; you can damage the connectors on the rear of the power supply.



4. Identify power supply components as described in this table:

Component	Description
1	Power supply
2	Cam handle release latch
3	Power and fault LEDs
4	Cam handle
5	Power cord locking mechanism

5. Push on the power supply to seat it all the way into the chassis, and then push the cam handle to the closed position, making sure that the cam handle release latch clicks into its locked position.
6. Reconnect the power cord, and secure it to the power supply using the power cord retainer clip.
7. Reconnect the power supply cable to the power source.

8. Turn the on/off switch on the power source and the power supply to the ON (I) position. Verify that the power supply is working correctly by observing that the power LED is lit green and the fault LED is not lit.
9. Repeat these steps for the remaining power supplies, if any.

Adding disk shelves

You can add a new AVA10S shelf to a configured and running AVA400 or AVA800 appliance. The maximum number of supported shelves is three for the AVA400 and four for the AVA800 models.

AltaVault does not support the addition of an existing AVA-10S shelf that has been previously used by another AltaVault appliance.


For instructions on installing the AVA10S shelves, and for general safety guidelines, see the guide, [SAS Disk Shelves Installation and Service Guide](#) for DS4243, DS2246, DS4486, and DS4246. The AltaVault AVA10S shelf is identical to the DS4246 disk shelf.

Important: For AltaVault specific shelf cabling information, see [AltaVault System Installation and Setup Instructions](#).


Changing the shelf ID for a disk shelf

If a change to a shelf ID is required, such as when expanding existing AltaVault capacity with a new shelf addition, perform the following steps:

1. If running, shut down the AltaVault Controller as described in “[Shutting down the AltaVault controller](#)” on [page 212](#).
2. Attach the power cable to the new shelf. To power on the disk shelf for both power supplies, flip the power switches to the on (I) position, as seen below.



3. a. Remove the left ear cover by using one hand to hold the left ear cover between your thumb and index finger.
- b. Pull either the top or the bottom of the cover until one end is released, then, pull off the left ear cover, as displayed below.



4. Press and hold the shelf ID button until the first digit on the digital display blinks.
5. Press the button to select a number 0 or 9. This number continues to blink.

Important: Valid AltaVault shelf IDs are 00, 01, 10, or 11. The actual setting depends on which shelf is being installed. For detailed setting information see the [AltaVault System Installation and Setup Instructions](#).

6. Repeat Step 4 and Step 5 for the second digit.
7. Press and hold the button until the second number stops blinking.
Both numbers on the digital display should blink and the shelf fault LED illuminates within five seconds. The fault LED stays lit until you power-cycle the shelf.
8. To power-cycle the disk shelf to ensure the new disk shelf ID takes effect, flip the Power Switch to the **off** position, wait several seconds, then flip it back to the **on** position.
9. Replace the left ear cover.
10. Reconnect and power on the system.

Adding an additional RAID group to a configured appliance

Note: Removing a drive out of a RAID group and moving it to a different RAID group is not supported. The disk drives have been formatted to a unique RAID configuration and should not be reused for hot-swappable drive replacements.

Refer to [NetApp Hardware Universe](#) for drive information.

The AltaVault appliance can accept additional packs of disks into added AVA10S shelves of the appliance. You can add an additional 12-pack RAID group to a configured and running AltaVault appliance. For AVA400, the maximum number of additional RAID groups is 5. For AVA800, the maximum number of additional RAID groups is 7. A RAID group ships with an initialized RAID volume. The following steps apply when installing a newly preconfigured 12-pack of disk drives into the shelf.

Before you begin

- Ensure you have an empty AVA10S shelf or an existing shelf with 12 empty disk slots. The empty slots number from 0-11 (bottom half of the shelf) and 12-23 (top half of the shelf).
- If you are installing two sets of 12-pack disk drives, install the entire 12-pack disk drives in either in the top half of the rack or the bottom half of the rack. Do not mix up the disks of a disk pack between the top and bottom halves of the rack as they are preconfigured as a group.
- AltaVault appliance storage configurations support only homogeneous drive sizes; the AVA400 supports 4TB and the AVA800 supports 6TB drives.
- Ensure the controller and the shelf are powered off. To shut down the controller, follow the procedure in the section, [“Shutting down the AltaVault controller” on page 212](#).

To add a 12-pack RAID group to a system

1. Place the 12 drives into the empty slots on the powered off disk shelf. The slots are numbered and grouped 0-11 and 12-23. Any drive in the 12-pack can go into any slot within the selected group.
2. Power on the shelf and then the controller.
3. Login as admin to the CLI using an SSH connection to the management NIC or serial console.
4. Enter the following CLI command in the configuration terminal:

```
CLI> show raidgroups
```

This command lists all the add-on RAID Groups connected to the AltaVault system. For the newly added RAID group, the following message displays: `raidgroup import <vd_id>`. The `vd_id` is an ID generated by the system.

5. Use the following command to import the newly added RAID Group into the AltaVault system and make it available for use.

```
CLI> raidgroup import <vd_id>
```

6. Confirm that the RAID group has been imported by running the following command:

```
show raidgroups
```

For information on replacing a faulty disk drive, see [“Replacing a faulty hard disk drive on an AltaVault AVA400 or AVA800 appliance” on page 228](#).

Replacing a faulty hard disk drive on an AltaVault AVA400 or AVA800 appliance

An alarm is raised when a disk drive fails. The alarm indicates which drive has failed and the state of the RAID being degraded. For information on replacing a faulty hard disk drive on an AltaVault AVA400 or AVA800 appliance, see the [KB Article](#) on the NetApp Support site.

Note: Rebuilding a disk drive can take 12 hours or longer. To view the rebuild progress, use the `show hwraid disk rebuild` CLI command.

For information on installing a new preconfigured 12-pack of disk drives into the shelf, see [“Adding an additional RAID group to a configured appliance” on page 227](#).

Replacing internal FRUs


To replace internal field-replaceable units (FRUs) inside the AltaVault controller module, you must perform a clean shutdown of the system, remove the controller module from the chassis, replace the faulty internal FRU, then, re-install the controller module in the system chassis.

Before you begin, ensure that you replace the impaired controller with a controller module that you received from your provider.

To replace an internal FRU

1. Use the `reload halt` command to perform a clean system shutdown of your system.
2. Power off the system and disconnect the power from the system using the steps described in the procedure, [“Shutting down the AltaVault controller” on page 212](#).
3. If *you* are not already grounded, properly *ground yourself*.

4. Loosen the thumbscrew on the cam handle as shown below.



5. Identify the components as described in this table:

Component	Description
1	Thumbscrew
2	Cam handle

6. Pull the cam handle downward to unseat the controller and slide the controller module out of the chassis until it catches. Then press the release latch on the left side of the controller module and slide the controller module completely out of the chassis, making sure that you support the base of the module with your free hand.
7. Place the controller module on a clean, flat surface.
8. Use one or more of the following procedures for the FRU you are replacing:
 - [“Installing a controller in a chassis” on page 216](#)
 - [“Replacing system DIMMs” on page 233](#)
 - [“Replacing the RTC clock coin battery” on page 237](#)
 - [“Replacing a boot device in a controller” on page 229](#)
 - [“Replacing a controller chassis” on page 218](#)
 - [“Disposing of batteries” on page 240](#)
9. After you have completed replacing the internal FRU, reinstall the controller module in the chassis using the steps described in the procedure, [“Installing a controller in a chassis” on page 216](#).

Replacing a boot device in a controller

The boot device stores a primary and secondary set of system files (also called the boot image) that the system uses when it boots. The boot device is not a hot-swappable field-replaceable unit (FRU).

The following error messages on the system console indicate that the boot media may have failed:

- Unrecoverable fsck error messages while the AltaVault OS is booting.

Example

```
Could not load fat://boot0/AV/image1/vmlinuz:Device not found
ERROR: Error booting OS on: 'boot0' file: fat://boot0/mars/image1/vmlinuz (boot0,fat)
```

```
Autoboot of PRIMARY image failed. Device not found (-6)
```

- Unrecoverable write errors messages
- Alerts from mDir service concerning configuration backup (configbkp) or boot device.

Replacing the boot device involves keeping track of current system image, shutting down the system cleanly, removing the old boot media from the controller, installing the new boot device, copy system files to the new boot device, and rebooting the system.

Note: Re-installing the OS is not a supported procedure in this release. An Return Merchandise Authorization (RMA) is required for replacing a boot device in a controller. Contact technical support to help you determine if you need to replace this field-replaceable unit (FRU) and to assist you during the replacement procedure.

Removing the boot device from the controller


Before you begin, note the details about the image installed on your system before you do a clean shutdown of the system for replacing the boot device. The same system image should be installed after boot device replacement.

Note: You must connect to the console port of the AltaVault controller to carry out the tasks in this procedure.

To remove the boot device from the controller

1. On the system console, type the following commands to note details of the operating system image that is installed on your system. This same OS version and revision should be installed after the boot device replacement.
 - `show info` to note the configuration details of your system
 - `show images` command to view the NetApp Release, Revision, Build date, Location, and Install Date
2. Perform a clean system shutdown using the `reload halt` command.
3. Power down your system and disconnect power from using the steps described in the procedure, [“Shutting down the AltaVault controller” on page 212](#).
4. Remove the controller from the chassis using the steps described in the procedure, [“Replacing internal FRUs” on page 228](#).
5. If *you* are not already grounded, properly *ground yourself*.

6. Locate the boot device holder using the FRU map on the CPU cover in the controller. The attention LED next to the boot device holder is lit. The image below shows the boot device and holder.



7. Identify boot device components as described in this table:

Component	Description
1	Boot device
2	Boot device holder: not removable

8. Open the boot device cover. Hold the boot device by its edges, gently lift it straight upwards to remove it out of the holder. Lifting the boot device at an angle can bend or break the connector pins in the boot device.

Note: Do not remove the boot device holder from the controller; it is not a FRU.

9. Set the boot device aside.

Installing a boot device


After you remove the faulty boot device from the controller, you must copy system files and restore configuration information to the replacement boot device.

Note: Re-installing the OS is not a supported procedure in this release. Contact technical support to help you determine if you need to replace this field-replaceable unit (FRU) and to assist you during the replacement procedure.

Before you begin, you must have the blank replacement boot device that you received from your provider.

To install the boot device

1. If *you* are not already grounded, properly *ground yourself*.
2. Locate the boot device holder in the controller. Use the FRU map on the controller module to help you locate the boot device holder.
3. Open the boot device cover, if applicable.
4. Align the boot device with the boot device socket or connector, and then firmly push the boot device straight down into the socket or connector.



5. Identify boot device components as described in this table:

Component	Description
1	Boot device
2	Boot device holder: not removable

6. Check the boot device to make sure that it is seated squarely and completely in the socket or connector. If necessary, remove the boot device and reseat it into the socket.
7. Close the boot device cover.

8. Reinstall the controller module in the chassis and connect the power as described in [“Installing a controller in a chassis” on page 216](#). In this case, the system begins to boot but stops at the loader prompt.

Note: All subsequent steps in this procedure are done at the loader prompt.

9. For initializing the new boot media, you must contact technical support.

Important: This procedure requires advanced knowledge of AltaVault systems and should not be done without help from technical support. Contact technical support to assist you in copying the system files and OS version to the boot media and bringing your system back up to a healthy state, after the boot media replacement procedure.

Replacing system DIMMs

An AltaVault controller has eight 32G DIMMs, also known as system memory. The DIMMs encounter correctable and uncorrectable errors during the normal operation of the system.

- With a uncorrectable memory errors, the system reboots.
- With correctable memory errors, the system does not reboot; the errors are recoverable and messages are recorded in the log.

Removing system DIMMs

System DIMMs are not hot-swappable FRUs. To remove a system DIMM from the AltaVault controller, shut down the system and remove the controller module from the chassis. Replace the failed DIMM with a replacement DIMM that is supported on your storage system.

Before you begin

- Perform a clean system shutdown using the `reload halt` command.
- Power down the system and disconnect power from it. For detailed instructions on performing a clean system shutdown and disconnection power, see [“Shutting down the AltaVault controller” on page 212](#).
- Remove the controller requires assistance from NetApp support.


To remove a system DIMM

1. If *you* are not already grounded, properly *ground yourself*.
2. Open the CPU cover in the AltaVault controller to access DIMMs 1, 2, 5, and 6.

Loosen the thumbscrew on the appropriate side panel and remove the side panel to access DIMMs 3, 4, 7, and 8; left side panel for DIMMs 3 and 4, and right side panel for DIMMs 7 and 8.

3. Locate the DIMM that needs to be replaced.

Figure 16-1. Locating system DIMMs




4. Identify DIMM components as described in this table:

Component	Description
1	CPU cover
2	DIMM 6, DIMM 5, DIMM 1, DIMM 2 (left to right)
3	DIMM 4, DIMM 3 (left to right)
4	DIMM 7, DIMM 8 (left to right)

5. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
6. Press down simultaneously on the two DIMM ejector tabs on either side of the DIMM to eject the DIMM from its slot, and then carefully lift it out of the slot.

All system DIMMs have white ejector latches, shown in the image below.



Important: Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.


7. Place the DIMM in an anti-static bag.
8. Repeat these steps to remove additional DIMMs as needed.

Installing system DIMMs

The AltaVault controller has eight 16GB system DIMMs installed; all eight DIMMs are required for optimal system performance. Before you begin, verify that the system DIMM you are installing is supported by your controller model.

To install a system DIMM

1. If *you* are not already grounded, properly *ground yourself*.
2. Open the CPU cover in the AltaVault controller to access the slots for DIMMs 1, 2, 5, and 6.
Loosen the thumbscrew on the appropriate side panel and remove the panel to access the slots for DIMMs 3, 4, 7, and 8; left side panel for DIMMs 3 and 4, and right side panel for DIMMs 7 and 8.
3. Locate the slot where you will be installing the new DIMM.
4. Ensure that the latches are in the open position.



5. Identify DIMM components as described in this table

Component	Description
1	CPU cover
2	DIMM 6, DIMM 5, DIMM 1, DIMM 2 (left to right)
3	DIMM 4, DIMM 3 (left to right)
4	DIMM 7, DIMM 8 (left to right)


6. Remove the replacement DIMM from the anti-static shipping bag, hold the DIMM by the corners, and align it over the slot. The notch among the pins on the DIMM should line up with the tab in the socket.

Important: Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

7. Insert the DIMM squarely into the slot. The DIMM fits tightly into the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Important: Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Push carefully, but firmly, on the top edge of the DIMM until the latches snap into place over the notches at the ends of the DIMM, displayed below. An audible click sound indicates the DIMM is securely installed in the slot.



9. Repeat the preceding steps to install additional DIMMs as needed.
10. Close the CPU cover and close and lock the side panel.
11. Reinstall the controller module in the chassis, connect power and boot up the system using the steps described in the procedure, [“Installing a controller in a chassis” on page 216](#).
12. To ensure that the system has booted correctly, verify that the System attention LED on the chassis front and Controller attention LED on chassis rear are not lit after the system has finished booting.

Replacing RAID controllers

RAID Controllers are assigned to slots 1 and 3 in the controller. The BBU cards are located in brackets that are attached to the existing controller shelves behind the installed cards.

For information on replacing a RAID controller, contact NetApp support. The procedure is complex and should only be done by field support.

Note: Perform a clean shutdown before replacing any RAID Controllers.

Replacing the RTC clock coin battery

You must replace a faulty real-time clock (RTC) coin battery in the controller module to ensure that your system's services and applications that depend on accurate time synchronization continue to function properly.

Removing an RTC battery

Removing an RTC battery entails shutting down the system, locating the battery in the controller module, and removing the battery.


Before you begin

- Perform a clean system shutdown using the `reload halt` command.
- Power down the system and disconnect its power using the steps described in the procedure, [“Shutting down the AltaVault controller” on page 212](#).
- Removing the controller from the chassis requires assistance from NetApp Support.

To remove an RTC battery

1. Locate the RTC coin battery in the controller module using the FRU map on the CPU cover. The attention LED next to the battery is lit.

The RTC coin battery in the AltaVault controller is located near the boot device, almost in the center of the controller as shown below.



2. Identify battery components as described in this table:

Component	Description
1	RTC battery and controller

3. Place your thumb or forefinger on the battery, gently push the battery away from the holder, and then lift the battery out of the holder.

Note: The polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder in the correct orientation, when replaced. A plus sign near the holder tells you how the battery should be positioned.

4. Place the battery on an anti-static surface.

Installing an RTC battery

Before you begin, verify that the RTC battery you are installing is supported by your controller model.

To install an RTC battery


1. If you are not already grounded, properly ground yourself.
2. Remove the replacement battery from the anti-static shipping bag.

3. Hold the battery such that the plus sign on the battery is facing you and away from the battery holder. You must install the battery in this orientation for the polarity to be correct.

Note: A plus sign near the battery holder indicates the battery polarity and how the battery should be positioned in the chassis.

4. Locate the empty battery holder in the controller module and insert the battery into the holder by tilting the battery at an angle and gently pushing down.

The image below shows the RTC battery in the controller. The battery should slide easily into the battery holder. If it does not, remove the battery and try again.



5. Identify battery components as described in this table:

Component	Description
1	RTC battery and controller

6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
7. Reinstall the controller module in the chassis, connect power and reboot the system using the steps described in the procedure, [“Installing a controller in a chassis” on page 216](#).
8. To ensure that the system has booted correctly, verify that the System attention LED on the chassis front and Controller attention LED on chassis rear are not lit after the system has finished booting.

Disposing of batteries

Dispose of batteries according to local regulations regarding battery recycling or disposal. If you cannot properly dispose of the battery, return it to NetApp, as described in the RMA instructions shipped with the kit.

Replacing disk shelf power supplies and other FRUs

For replacing disk power supplies, and other disk shelf FRUs, see [SAS Disk Shelves Installation and Service Guide for DS4243, DS2246, DS4486, and DS4246](#).

Returning failed parts

Return failed parts to NetApp as described in the RMA instructions shipped with the kit.

Contact technical support at mysupport.netapp.com, 888-463-8277 (North America/Canada), 00-800- 44-638277 (Europe/EMEA), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Disposing of batteries

Dispose of batteries according to local regulations regarding battery recycling or disposal. If you cannot properly dispose of the battery, return it to NetApp, as described in the RMA instructions shipped with the kit.

APPENDIX A Administrator's configuration worksheet

Configuration worksheet

1	Appliance information	Notes
1.1	Appliance host name: This is the name of the appliance and needs to be configured in DNS as well.	
1.2	IP address for Primary Interface: This IP address is required for configuring the PRI (Primary) / Management Interface when the CLI wizard runs automatically during the initial deployment. It is recommended to use a static IP address.	
1.3	Netmask	
1.4	Primary DNS server IP	
1.5	NTP Server	
1.6	FQN and IP address of SMTP Relay Server: This is required to configure and enable email notifications.	
1.7	Email address or alias for Notification of Events and Failures: There are two groups of notifications - "Events" Group and "Failure" Group.	
1.8	Domain name for the appliance: This needs to be configured in DNS as well to resolve fqdn of the appliance.	
1.10	Time zone in which the appliance will be installed	
2	Cloud provider credentials and storage configuration	Notes
2.1	Name of Preferred Cloud Provider: For example, AWS etc. The information required to configure the cloud provider is dependent on the specific cloud provider you have selected.	
2.2	Region: The information required to configure the cloud provider is dependent on the specific cloud provider you have selected.	
2.3	Credentials to Cloud Object Storage: For example, Amazon S3 access and secret key.	
2.4	Bucket name: The bucket name must be unique (across all of AWS).	

2	Cloud provider credentials and storage configuration	Notes
2.5	Is connectivity to cloud, in place - Yes/No? Connectivity to the cloud is mandatory in order for the appliance to be configured. WAN circuits, firewalls (port 443) etc. must be configured prior to commencing configuration. If not when will it be in place?	
2.6	Bandwidth to cloud.	
2.7	Port value (typically 443)?	
2.8	Cloud CA Certificate (if using private cloud): .pem file for private cloud provider	
2.9	Encryption key: If you import from another AltaVault appliance, have key file or contents and key passphrase available (if required).	
2.10	Replication interface - Replication interface defaults to the Primary interface unless another interface is selected and configured.	
2.11	Number of shares to be created on the AltaVault appliance.	
2.12	SMB Share Name(s) / naming convention to be used: Specify the share name(s) to be configured or whether any naming convention is to be followed.	
2.13	SMB Domains: Will the AltaVault appliance be part of AD Domain? Yes / No. If Yes, Domain administrator credentials will be required to join the AD Domain.	
2.14	Preferred domain controllers - specify up to three domain controllers	
2.15	SMB Username(s) / Groups to be given access to the share,	
2.16	NFS Export Name / naming convention to be used,	
2.17	Is NFSv4 Kerberos required? If yes, Kerberos keytab and conf files are required.	
2.18	OST Share Name(s) / naming convention to be used: Specify the share name(s) to be configured or whether any naming convention is to be followed.	
2.19	OST Username(s) to be given access to the share,	
2.22	SnapMirror Whitelist - IP addresses from which Snapshots will be allowed	
3	Network connectivity	Notes
3.1	Number of 1GbE/10GbE Data Interfaces to be used?	
3.2	Specify speed of Data (Backup) LAN - 1Gbe or 10Gbe? Specify if the Data LAN is be 1gb or 10gb NICs. AltaVault Appliance supports 4x1gb and 4x10gb NICs. IP addresses will be required for each port depending on type of NIC supplied (4x1gb or 4x10gb).	
3.3	Describe the LAN topology for Backup Data to the Appliance: Flat LAN or VLANs.	
3.4	Will Data Interfaces be connected to different subnets / VLANs? AltaVault supports VLAN tagging. Note the VLAN IDs associated with each data interface.	

3	Network connectivity	Notes
3.5	Do you want to configure a virtual interface (802.3ad link aggregation) for the Data Interfaces? If so does the LAN switch support 802.3ad?	
3.6	<ol style="list-style-type: none"> 1. Provide up to 4 x IP address / Netmask / Gateway for each 4x1Gbe port. 2. Provide up to 4 x IP addresses / NetMask/Gateway for each 4x10Gbe port. 	
3.7	Specify the type of SFP for 10Gbe Optical NICs.	
4	Advance features	Notes
4.1	Bandwidth throttling for replication to cloud	
4.2	Alarms, Announcements, Logging, Scheduled Reports	
4.3	SNMP	

APPENDIX B AltaVault appliance MIB

This section provides a reference to the AltaVault Management Information Base (MIB) and SNMP traps.

This appendix includes the following sections:

- [“Accessing AltaVault appliance MIB” on page 245](#)
- [“SNMP traps” on page 245](#)

Accessing AltaVault appliance MIB

AltaVault MIB monitors device status and peers, and provides network statistics for seamless integration into network management systems such as Hewlett-Packard OpenView Network Node Manager, PRTG, and other SNMP browser tools.

For details about configuring and using these network monitoring tools, consult their companies’ Web sites.

The following guidelines describe how to download and access the AltaVault MIB using common MIB browsing utilities:

- You can download the AltaVault MIB (NTAP-MIB.txt or AVA-MIB.txt) from the Management Console and load it into any MIB browser utility.
- Some utilities might expect a file type other than a text file. If this occurs, change the file type to the one expected.
- Some utilities assume that the root is mib-2 by default. If the utility sees a new node, such as enterprises, it might look under mib-2.enterprises. If this occurs, use .iso.org.dod.internet.private.enterprises.rbt as the root.
- Some command-line browsers might not load all MIB files by default. If this occurs, find the appropriate command option to load the NTAP-MIB.txt file. For example, for NET-SNMP browsers, snmpwalk -m all.

SNMP traps

Every AltaVault supports SNMP traps and email alerts for conditions that require attention or intervention. An alarm fires for most, but not every, event and the related trap is sent. For most events, when the condition clears, the system clears the alarm and also sends out a clear trap. The clear traps are useful in determining when an event has been resolved.

This section describes the SNMP traps. It does not list the corresponding clear traps.

AltaVault includes support for SNMP v3.

You can view the AltaVault health at the top of each Management Console page, by entering the CLI `show info` command, and through SNMP (`health`, `systemHealth`).

The AltaVault tracks key hardware and software metrics and alerts you of any potential problems so that you can quickly discover and diagnose issues. Appliance health falls into one of the following states:

- **Healthy** - The AltaVault is functioning and optimizing storage.
- **Needs Attention** - The AltaVault is optimizing storage, but there are management-related issues.
- **Degraded** - The AltaVault is optimizing storage but the system has detected an issue.
- **Critical** - The AltaVault might not be optimizing storage and a critical issue needs to be addressed.

This table summarizes the SNMP traps sent out from the system to configured trap receivers and their effect on AltaVault health state.

Trap and OID	AltaVault Appliance State	Text	Description
procCrash (1.3.6.1.4.1.17163.1.102.4.0.1)		A procCrash trap signifies that a process managed by PM has crashed and left a core file. The variable sent with the notification indicates which process crashed.	A process has crashed and subsequently been restarted by the system. The trap contains the name of the process that crashed. A system snapshot associated with this crash has been created on the appliance and is accessible via the CLI or the Management Console. NetApp Support might need this information to determine the cause of the crash. No other action is required on the appliance because the crashed process is automatically restarted.
procExit (1.3.6.1.4.1.17163.1.102.4.0.2)		A procExit trap signifies that a process managed by PM has exited unexpectedly, but not left a core file. The variable sent with the notification indicates which process exited.	A process has unexpectedly exited and been restarted by the system. The trap contains the name of the process. The process might have exited automatically or due to other process failures on the appliance. Review the release notes for known issues related to this process exit. If none exist, Contact NetApp Support to determine the cause of this event. No other action is required on the appliance because the crashed process is automatically restarted.
configChange ((1.3.6.1.4.1.17163.1.102.4.0.3)		A change has been made to the system's configuration.	A configuration change has been detected. View the log files around the time of this trap to determine what changes were made and whether they were authorized.
cpuUtil (1.3.6.1.4.1.17163.1.102.4.0.4)	Degraded	The average CPU utilization in the past minute has gone above the acceptable threshold.	Average CPU utilization has exceeded an acceptable threshold. If CPU utilization spikes are frequent, it might be because the system is undersized. Sustained CPU load can be symptomatic of more serious issues. Consult the CPU Utilization report to gauge how long the system has been loaded and also monitor the amount of traffic currently going through the appliance. A one-time spike in CPU is normal but NetApp recommends reporting extended high CPU utilization to NetApp Support. No other action is necessary because the alarm clears automatically.
pagingActivity (1.3.6.1.4.1.17163.1.102.4.0.5)	Degraded	The system has been paging excessively (thrashing).	The system is running low on memory and has begun swapping memory pages to disk. This event can be triggered during a software upgrade while the storage optimization service is still running, but there can be other causes. If this event triggers at any other time, generate a debug sysdump report and send it to NetApp Support. No other action is required because the alarm clears automatically.

Trap and OID	AltaVault Appliance State	Text	Description
linkError (1.3.6.1.4.1.17163.1.102.4.0.6)	Degraded	An interface on the appliance has lost its link.	<p>The system has lost one of its Ethernet links due to a network event. Check the physical connectivity between the AltaVault and its neighbor device. Investigate this alarm as soon as possible. Depending on what link is down, the system might no longer be optimizing and a network outage could occur.</p> <p>This is often caused by surrounding devices, like routers or switches transitioning their interfaces. This alarm also accompanies service or system restarts on the AltaVault.</p>
powerSupplyError (1.3.6.1.4.1.17163.1.102.4.0.7)	Degraded	A power supply on the appliance has failed (not supported on all models).	A redundant power supply on the appliance has failed and needs to be replaced. Contact NetApp Support for an RMA replacement soon.
fanError (1.3.6.1.4.1.17163.1.102.4.0.8)	Degraded	A fan has failed on this appliance (not supported on all models).	A fan is failing or has failed and needs to be replaced. Contact NetApp Support for an RMA replacement soon.
memoryError (1.3.6.1.4.1.17163.1.102.4.0.9)	Degraded	A memory error has been detected on the appliance (not supported on all models).	A memory error has been detected. A system memory stick might be failing. Try reseating the memory first. If the problem persists, contact NetApp Support for an RMA replacement soon.
ipmi (1.3.6.1.4.1.17163.1.102.4.0.10)	Degraded	An IPMI event has been detected on the appliance. Check the details in the alarm report on the Web UI (not supported on all models).	<p>An Intelligent Platform Management Interface (IPMI) event has been detected. Check the Alarm Status page for more detail. You can also view the IPMI events on the AltaVault by entering the CLI command:</p> <pre>show hardware error-log all</pre>
localFSFull (1.3.6.1.4.1.17163.1.102.4.0.11)		The appliance local file system is full.	The AltaVault local file system is full. Check the Eviction report and contact NetApp Support.
temperatureCritical (1.3.6.1.4.1.17163.1.102.4.0.12)	Critical	The system temperature has reached a critical stage.	This trap/alarm triggers a critical state on the appliance. This alarm occurs when the appliance temperature reaches 90 degrees Celsius. The temperature value is not user configurable. Reduce the appliance temperature.
temperatureWarning (1.3.6.1.4.1.17163.1.102.4.0.13)	Degraded	The system temperature has exceeded the threshold.	The appliance temperature is a configurable notification. By default, this notification is set to trigger when the appliance reached 70 degrees Celsius. Raise the alarm trigger temperature if it is normal for the AltaVault to get that hot, or reduce the temperature of the AltaVault.
scheduledJobError (1.3.6.1.4.1.17163.1.102.4.0.14)		A scheduled job has failed during execution.	A scheduled job on the system (for example, a software upgrade) has failed. To determine which job failed, use the CLI or the Management Console.

Trap and OID	AltaVault Appliance State	Text	Description
confModeEnter (1.3.6.1.4.1.17163.1.102.4.0.15)		A user has entered configuration mode.	A user on the system has entered a configuration mode from either the CLI or the Management Console. A login to the Management Console by user admin sends this trap as well. This message is for notification purposes only; no other action is necessary.
confModeExit (1.3.6.1.4.1.17163.1.102.4.0.16)		A user has exited configuration mode.	A user on the system has exited configuration mode from either the CLI or the Management Console. A logout of the Management Console by user admin sends this trap as well. This message is for notification purposes only; no other action is necessary.
secureVaultLocked (1.3.6.1.4.1.17163.1.102.4.0.17)	Needs Attention	The secure vault is locked. The secure data cannot be used.	The secure vault is locked. The datastore cannot be encrypted. Check the Alarm Status page for more details. The alarm clears when the secure vault is unlocked.
dirtyCloud (1.3.6.1.4.1.17163.1.102.4.0.19)		The cloud bucket is not empty even though the local datastore is.	There is data in the cloud although the AltaVault datastore is empty. Enable replication and recovery to ensure that the cloud storage is synchronized with the datastore.
license (1.3.6.1.4.1.17163.1.102.4.0.21)		The WWBASE license is missing, expired, or invalid.	A license on AltaVault has been removed, has expired, or is invalid. The alarm clears when a valid license is added or updated.
lowSpace (1.3.6.1.4.1.17163.1.102.4.0.22)		The backup service is running out of space.	The AltaVault storage optimization service is running out of space. Delete extra files and ensure that there is more space.
overCapacity (1.3.6.1.4.1.17163.1.102.4.0.24)		The cloud bucket has reached its licensed capacity.	AltaVault capacity license manages the capacity of the storage within the cloud that the system can address. The cloud storage has reached the licensed capacity limit. You must increase your capacity license or decrease your data storage.
replicationError (1.3.6.1.4.1.17163.1.102.4.0.25)		Replication encountered a non-fatal error.	There was an error in the replication process. The system automatically retries the replication process. Contact your cloud service provider or NetApp Support.
replicationPause (1.3.6.1.4.1.17163.1.102.4.0.26)		Replication was paused.	Replication was paused because you scheduled it to pause at the current time.
serviceError (1.3.6.1.4.1.17163.1.102.4.0.27)	Degraded	The backup service could not initialize properly.	The storage optimization service has encountered a condition that might degrade its performance. See the system log for more information. No other action is necessary.
portalUnreachable (1.3.6.1.4.1.17163.1.102.4.0.28)		The NetApp Cloud Portal is unreachable.	The AltaVault cannot access the NetApp Cloud Portal.

Trap and OID	AltaVault Appliance State	Text	Description
rplctdBytesPending (1.3.6.1.4.1.17163.1.102.4.0.29)		Too many bytes pending to be replicated.	The number of bytes pending replication to the cloud has exceeded the “Bytes pending replication” alert limit. This might be due to a slow replication link or because the appliance is undersized. Increase the amount of replication bandwidth available to the AltaVault or temporarily stop backup operations so that replication can normalize.
cloudDispError (1.3.6.1.4.1.17163.1.102.4.0.30)		The appliance might be connecting to another appliance’s cloud bucket.	This indicates that the cloud bucket that the AltaVault is trying to connect to might be in use by another AltaVault. This prevents corruption of the files in the cloud.
dataEviction (1.3.6.1.4.1.17163.1.102.4.0.31)		The appliance had to evict recent data to prevent running out of space in the datastore.	<p>This indicates that the system has detected an issue with datastore eviction.</p> <p>The alarm triggers when the appliance starts evicting data from the local disk cache and the age of the evicted data is relatively young. An AltaVault has disk space much smaller than the total addressable space on the cloud, and if disk space runs low, the appliance starts evicting data from disk that has not been used recently. This keeps only fresh and frequently accessed data in cache.</p> <p>The AltaVault keeps statistics about how old the evicted data is (this is the average evicted age). Usually, only old data is evicted. This behavior is generally not a problem and does not trigger an alarm. However, the appliance might be experiencing such a huge workload that more and more recent data needs to be evicted from the appliance to make space for incoming data. This causes the average evicted age to decrease, and when it goes below a certain threshold, the average evicted age alarm triggers. This alarm is an anomalous event, signaling that the appliance is handling a much larger workload than expected.</p> <p>This alarm is useful in detecting whether the appliance is undersized relative to your normal workload. If the alarm is constantly triggered, then you should consider moving your data to a AltaVault model with a larger disk cache.</p>

Trap and OID	AltaVault Appliance State	Text	Description
approachingCapacity (1.3.6.1.4.1.17163.1.102.4.0.32)		The cloud bucket is approaching its licensed capacity.	<p>This indicates that the amount of storage that the AltaVault used in the cloud bucket is approaching the licensed cloud capacity amount.</p> <p>Consider the following options:</p> <ul style="list-style-type: none"> • Upgrade the AltaVault to a version that can accommodate more space. • Add more AltaVaults (each of which will account for subsets of your entire data) in your network. • Delete data from the AltaVault and reclaim used space.
blockstoreFull (1.3.6.1.4.1.17163.1.102.4.0.34)		The System reserved space is full.	The amount of reserved space on the AltaVault is full. Contact NetApp Support.
shelfError (1.3.6.1.4.1.17163.1.102.4.0.35)		One or more shelves have errors.	One or more expansion shelves attached to the AltaVault have errors. Contact NetApp Support.
lowMemory (1.3.6.1.4.1.17163.1.102.4.0.36)		The system doesn't have enough memory to start service.	Memory usage on the AltaVault is high. The storage optimization system performance might decrease if the amount of memory available to the AltaVault remains low.
hwraidDiskIndivError (1.3.6.1.4.1.17163.1.102.4.0.38)		The appliance has detected an error with one or more disks.	One or more disks had has detected an error.
shelfPowerSupply (1.3.6.1.4.1.17163.1.102.4.0.39)		One or more shelves have a power supply error.	One or more shelves have a power supply error.
hwraidBbuError (1.3.6.1.4.1.17163.1.102.4.0.40)		The Storage Optimization Service is disabled because the RAID is degraded and the battery backup unit is not sufficiently charged.	The Storage Optimization Service is disabled because the RAID is degraded and the battery backup unit is not sufficiently charged.
bbuIndivError (1.3.6.1.4.1.17163.1.102.4.0.41)		One or more battery backup units have errors.	One or more battery backup units have errors.
hwraidIntegrityCheckError (1.3.6.1.4.1.17163.1.102.4.0.42)		RAID integrity check alarm triggered due to unclean shutdown.	RAID integrity check alarm triggered due to unclean shutdown.
evalModeExpiry (1.3.6.1.4.1.17163.1.102.4.0.43)		Evaluation mode for this virtual appliance is expiring soon.	Evaluation mode for this virtual appliance is expiring soon.
metadataSpaceFull (1.3.6.1.4.1.17163.1.102.4.0.44)		The space reserved for metadata is full.	The space reserved for metadata is full.
inodesThresholdReached (1.3.6.1.4.1.17163.1.102.4.0.45)		The maximum inodes limit has been reached.	The maximum inodes limit has been reached.

Trap and OID	AltaVault Appliance State	Text	Description
dataIntegrityError (1.3.6.1.4.1.17163.1.102.4.0.46)		Data integrity check has reported an inconsistency.	The data integrity check has reported an inconsistency.
replicationInconsistent (1.3.6.1.4.1.17163.1.102.4.0.47)		Replication is experiencing high number of retries.	The replication is experiencing high number of retries.
inconsistentCloudData (1.3.6.1.4.1.17163.1.102.4.0.48)		Possible cloud data corruption or inconsistency.	A possible cloud data corruption or inconsistency has been detected.
softwareUpdateAvailable (1.3.6.1.4.1.17163.1.102.4.0.49)		Software updated available.	A software update is available.
firmwareUpgrade (1.3.6.1.4.1.17163.1.102.4.0.50)		Firmware upgrade available.	A firmware upgrade is available.
controllerPowerSupplyError (1.3.6.1.4.1.17163.1.102.4.0.51)		A power supply on the controller has failed.	A power supply on the controller has failed.
controllerFanError (1.3.6.1.4.1.17163.1.102.4.0.52)		A fan has failed on the controller.	A fan has failed on the controller.
cpuUtilClear (1.3.6.1.4.1.17163.1.102.4.0.10004)		The average CPU utilization has fallen back within the acceptable threshold.	Average CPU utilization exceeded an acceptable threshold, but has now fallen back to an acceptable threshold. If CPU utilization spikes are frequent, it might be because the system is undersized. Sustained CPU load can be symptomatic of more serious issues. Consult the CPU Utilization report to gauge how long the system has been loaded and also monitor the amount of traffic currently going through the appliance. A one-time spike in CPU is normal but NetApp recommends reporting extended high CPU utilization to NetApp Support. No other action is necessary because the alarm clears automatically.
pagingActivityClear (1.3.6.1.4.1.17163.1.102.4.0.10005)	Degraded	The system has stopped paging excessively (thrashing).	The system that was running low on memory and was swapping memory pages to disk has stopped paging excessively, and the alarm clears.
linkErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10006)		An interface on the appliance has regained its link.	The system has regained the Ethernet link it lost due to a network event and clears the alarm.
powerSupplyErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10007)		All power supplies are now functioning normally.	A redundant power supply had failed on the appliance, but this is now corrected and all power supplies are functioning properly.
fanErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10008)		All system fans are now functioning normally.	A fan on the appliance had failed, but now all fans are functioning normally.

Trap and OID	AltaVault Appliance State	Text	Description
memoryErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10009)		A memory error has been rectified on the appliance.	There was an alarm due to a memory error on the appliance, but this is now cleared because the memory error has been corrected.
ipmiClear (1.3.6.1.4.1.17163.1.102.4.0.10010)		An IPMI event has been rectified on the appliance.	An Intelligent Platform Management Interface (IPMI) event was detected, but it is now corrected.
localFSFullClear (1.3.6.1.4.1.17163.1.102.4.0.10011)		The appliance local file system usage is below threshold.	The appliance local file system usage exceeded the permitted threshold, but the usage is now below the threshold.
temperatureNonCritical (1.3.6.1.4.1.17163.1.102.4.0.10012)		The system temperature is no longer in a critical stage.	The alarm that occurred when the appliance temperature reached 90 degrees Celsius is now cleared.
temperatureNormal (1.3.6.1.4.1.17163.1.102.4.0.10013)		The system temperature is back within the threshold.	The temperature of the appliance exceeded the threshold (default value of 70 degrees Celsius), but is now back within the threshold.
secureValutUnlocked (1.3.6.1.4.1.17163.1.102.4.0.10017)		Secure vault is unlocked. The secure data store can be used now.	The secure vault, which was locked, is now unlocked and you can use the secure datastore.
dirtyCloudClear (1.3.6.1.4.1.17163.1.102.4.0.10019)		The cloud consistency issue has been resolved.	Both the cloud bucket and the datastore are synchronized.
licenseClear (1.3.6.1.4.1.17163.1.102.4.0.10021)		The appliance is now properly licensed.	A license on the AltaVault was removed, had expired, or was invalid, but the license has been installed correctly again.
lowSpaceClear (1.3.6.1.4.1.17163.1.102.4.0.10022)		The appliance now has enough space for datastore.	The appliance did not have enough space for datastore earlier. Now, it has enough space and the alarm is cleared.
overCapacityClear (1.3.6.1.4.1.17163.1.102.4.0.10024)		The cloud bucket is no longer over its licensed capacity.	The cloud storage meets the storage limit allocated by its capacity license.
replicationErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10025)		Replication error has been cleared.	There was an error in the replication process, but this is now cleared.
replicationPauseClear (1.3.6.1.4.1.17163.1.102.4.0.10026)		Replication was resumed.	Replication had paused, but it has resumed.
serviceErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10027)		The backup service initialized properly.	There was an error in the storage optimization service, but this is now corrected.
portalUnreachableClear (1.3.6.1.4.1.17163.1.102.4.0.10028)		The NetApp Cloud Portal is now reachable.	The AltaVault could not access the NetApp Cloud Portal, but this is now corrected.
rplctdBytesPendingClear (1.3.6.1.4.1.17163.1.102.4.0.10029)		Number of bytes pending replication is within normal range now.	The number of bytes pending replication had exceeded the bytes pending alert limit, but this is now corrected.
cloudDispErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10030)		The appliance is now connecting to its own cloud bucket.	The cloud bucket that the AltaVault was connecting to was in use by another AltaVault. This issue is now corrected and the AltaVault is connected to its own cloud bucket.

Trap and OID	AltaVault Appliance State	Text	Description
dataEvictionClear (1.3.6.1.4.1.17163.1.102.4.0.10031)		The appliance is no longer evicting data.	The datastore eviction issue is now corrected.
blockstoreFullClear (1.3.6.1.4.1.17163.1.102.4.0.10033)		The System reserved space is now enough.	The amount of reserved space on the AltaVault is now sufficient for the storage optimization service to run efficiently.
shelfErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10034)		All shelves are working properly.	There was an error in one of the expansion shelves attached to the AltaVault, but this is now corrected.
lowMemoryClear ((1.3.6.1.4.1.17163.1.102.4.0.10035)		The system has enough memory.	The amount of memory available to the AltaVault is now sufficient.
hwraidDiskIndivClear (1.3.6.1.4.1.17163.1.102.4.0.10037)		All disks are functioning properly.	All the disks are functioning properly.
shelfPowerSupplyClear (1.3.6.1.4.1.17163.1.102.4.0.10038)		All shelves have proper power supply.	All the shelves have proper power supply.
hwraidBbuErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10039)		The storage optimization service is no longer disabled.	The storage optimization service is no longer disabled.
bbuIndivClear (1.3.6.1.4.1.17163.1.102.4.0.10040)		All battery backup units are functioning properly.	All the battery backup units are functioning properly.
hwraidIntegrityCheckClear (1.3.6.1.4.1.17163.1.102.4.0.10041)		RAID integrity check alarm cleared.	The RAID integrity check alarm cleared.
evalModeExpiryClear (1.3.6.1.4.1.17163.1.102.4.0.10042)		The virtual appliance is no longer running in evaluation mode.	The virtual appliance is no longer running in evaluation mode.
metadataSpaceFullClear (1.3.6.1.4.1.17163.1.102.4.0.10043)		The metadata reserved space is now enough.	The metadata reserved space is sufficient.
inodesThresholdReachedClear (1.3.6.1.4.1.17163.1.102.4.0.10044)		The number of inodes is now below the maximum limit.	The number of inodes is now below the maximum limit.
dataIntegrityErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10045)		The data integrity check alarm cleared.	The data integrity check alarm has cleared.
replicationInconsistentClear (1.3.6.1.4.1.17163.1.102.4.0.10046)		The replication is no longer experiencing retries now.	The replication is no longer experiencing retries.
inconsistentCloudDataClear (1.3.6.1.4.1.17163.1.102.4.0.10047)		Restored cloud data consistency.	The cloud data consistency has been restored.
softwareUpdateAvailableClear (1.3.6.1.4.1.17163.1.102.4.0.10048)		Software update available alarm cleared.	The software update alarm has been cleared.
firmwareUpgradeClear (1.3.6.1.4.1.17163.1.102.4.0.10049)		Firmware upgrade alarm cleared.	The firmware upgrade alarm is cleared.

Trap and OID	AltaVault Appliance State	Text	Description
controllerPowerSupplyErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10050)		All power supplies are now functioning normally.	All the power supplies are functioning.
controllerFanErrorClear (1.3.6.1.4.1.17163.1.102.4.0.10051)		All controller fans are now functioning normally.	All the controller fans are functioning.

APPENDIX C Amazon AWS IAM and S3 bucket policies

Amazon AWS provides the ability to specify Identity and Access Management (IAM) policies and bucket policies to control permissions related to AWS users and S3 cloud buckets. In general, IAM users and buckets should be configured with the minimum permissions required for normal operation. For more details about Amazon's best practices, see

<http://docs.aws.amazon.com/IAM/latest/UserGuide/IAMBestPractices.html>.

For more information on Amazon AWS, see the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Cloud Appliances*.

This appendix includes the following sections:

- “Typical AltaVault setup” on page 257
- “IAM policies for AltaVault” on page 257
- “Bucket policies for AltaVault” on page 259

Typical AltaVault setup

A typical AltaVault setup includes the following AWS configuration:

- One IAM user created exclusively for AltaVault. Access keys are generated for the user and entered into the AltaVault cloud configuration. AltaVault never requires access keys for the root AWS account. It is recommended that access keys are not generated for the root account.
- An IAM group is created with the AltaVault user. A policy is set on the group that allows only the permissions used by AltaVault.
- If a bucket policy is required, then a bucket is created for use by AltaVault, with a policy that allows only the AltaVault user to access it. It is not necessary to create the bucket prior to AltaVault using it if bucket policies are not required.

IAM policies for AltaVault

IAM policies allow access to the Amazon S3 account and its associated cloud buckets via different users with restricted permissions, in contrast to the root account which has unrestricted access to the account and cloud buckets. It is recommended that programmatic access (including access via appliances such as AltaVault) to Amazon AWS and S3 are done via IAM users with the appropriate permissions rather than via the root AWS account.

AltaVault requires the following IAM user permissions:

- On all buckets:
 - ListAllMyBuckets (not required for normal operation, but some features may not work)
- On the configured cloud bucket:
 - CreateBucket (not required if the bucket has been created beforehand)
 - GetBucketLocation
 - ListBucket
 - ListBucketMultipartUploads
 - GetLifecycleConfiguration
 - PutLifecycleConfiguration
- On objects inside the configured cloud bucket:
 - AbortMultipartUpload
 - DeleteObject
 - GetObject
 - ListMultipartUploadParts
 - PutObject
 - RestoreObject

Sample of IAM policy

Below is a sample of the IAM policy implementing the above permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1394143726000",
      "Effect": "Allow",
      "Action": [
        "s3:ListAllMyBuckets"
      ],
      "Resource": [
        "arn:aws:s3:::*"
      ]
    },
    {
      "Sid": "Stmt1394143742000",
      "Effect": "Allow",
      "Action": [
        "s3:CreateBucket",
        "s3:GetBucketLocation",
        "s3:ListBucket",
        "s3:ListBucketMultipartUploads",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration"
      ],
      "Resource": [
        "arn:aws:s3:::bucket_name"
      ]
    }
  ]
}
```

```
{
  "Sid": "Stmt1394143790000",
  "Effect": "Allow",
  "Action": [
    "s3:AbortMultipartUpload",
    "s3:DeleteObject",
    "s3:ListMultipartUploadParts",
    "s3:GetObject",
    "s3:PutObject",
    "s3:RestoreObject"
  ],
  "Resource": [
    "arn:aws:s3:::bucket_name/*"
  ]
}
```

Bucket policies for AltaVault

Amazon S3 bucket policies can be configured to allow only specific users (including users outside the AWS account) to access an S3 cloud bucket, and can be used in conjunction with IAM user policies. AltaVault requires that the cloud bucket (configured in the AltaVault management console under Configure > Cloud Settings) allows access by the IAM user configured for AltaVault. No access by any other user is required.

AltaVault requires a set of permissions in the bucket policy similar to the set of permissions for an IAM policy, with the exception of `s3:ListAllMyBuckets` and `s3:CreateBucket`, which are not relevant at the bucket level.

Sample of bucket policy

Below is a sample of the bucket policy:

```
{
  "Id": "Policy1394662102999",
  "Statement": [
    {
      "Sid": "Stmt1394661890920",
      "Action": [
        "s3:GetBucketLocation",
        "s3:GetLifecycleConfiguration",
        "s3:ListBucket",
        "s3:ListBucketMultipartUploads",
        "s3:PutLifecycleConfiguration"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::bucket_name",
      "Principal": {
        "AWS": [
          "arn:aws:iam::123456789012:user/user_name"
        ]
      }
    },
    {
      "Sid": "Stmt1394661925663",
      "Action": [
        "s3:AbortMultipartUpload",
        "s3:DeleteObject",
        "s3:ListMultipartUploadParts",
        "s3:GetObject",

```

```
        "s3:PutObject",
        "s3:RestoreObject"
    ],
    "Effect": "Allow",
    "Resource": "arn:aws:s3:::bucket_name/*",
    "Principal": {
        "AWS": [
            "arn:aws:iam::123456789012:user/user_name"
        ]
    }
}
]
```

APPENDIX D **Best practices for restoring data from archive**

This section describes best practices for restoring data from Alibaba Archive, Amazon Glacier, or Microsoft Azure Archive. It includes the following sections:

- [“Optimizing data movement” on page 261](#)
- [“Protecting data” on page 261](#)
- [“Recovering data from archive” on page 262](#)
- [“AltaVault appliance best practices for EMC NetWorker” on page 268](#)
- [“AltaVault appliance best practices for IBM Spectrum Protect” on page 270](#)
- [“AltaVault appliance best practices for Veritas NetBackup” on page 271](#)
- [“AltaVault appliance best practices for Veritas Backup Exec” on page 272](#)
- [“AltaVault appliance best practices for Veeam backup and replication” on page 273](#)

Optimizing data movement

AltaVault appliances can be configured in backup or cold storage mode for use with Alibaba Archive, Amazon Glacier, or Azure Archive archive tiers. As most use cases for cloud provider archive tiers are typically for long term storage of inactive data, AltaVault is generally configured in cold storage mode to maximize cloud storage. Be aware that when using AltaVault in cold storage mode, the appliance cache is used for holding a greater amount of deduplicated data in cloud provider archive tiers, which results in a lower amount of cache being available to hold data for immediate recovery needs. For more information on the two appliance modes, see [“Deployment guidelines” on page 13](#).

Protecting data

Protecting backup and critical production servers using an archive tier also adds additional considerations. In most data protection scenarios, backup servers would be protected the same way as any other production server backup data, and stored in the same storage target but under a different set of retention requirements and typically a different target location. During disaster recovery (DR), backup and critical production server recoveries are the first processes that occur, and thus retrieval of the backup server backup is critical in the DR process. If you use a cloud storage target like S3, then this would typically not pose a problem in terms of the DR time frame. However, if you use an archive tier,

then this would most likely incur a heavy delay that most businesses would find unacceptable from an RTO perspective. The appropriate approach for handling backup and critical production server protection if using an archive tier would be to have the backup and critical production server backups go to a separate share on AltaVault which is pinned. AltaVault pinned data is always held on local cache and is never evicted, allowing administrators to have access to these server backups immediately and without any delays associated with accessing data stored in archive tiers.

For Amazon Glacier, AltaVault can also be configured to delay the migration of data from Amazon S3 to Glacier. By default, data is typically migrated within a 24 hour window once it arrives on Amazon S3. Using the AltaVault CLI command, `replication migration-delay`, Amazon S3 can be instructed to maintain the data for a longer period of time (set in days), until the data is moved to Amazon Glacier. However, it is important to understand that while this can extend the retention on Amazon S3 for new data segments, iterations of the same data segments sent to AltaVault (such as during subsequent full backups) will not result in the migration delay being extended for that deduplicated data segment already residing on Amazon S3. It is likely that over time, only a subset of a backup will reside on Amazon S3 awaiting migration, but a majority of the data will reside in Amazon Glacier since it hasn't changed. Refer to the *NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Guide* for more details.

Recovering data from archive

Recovery times can vary depending on the cloud provider.

When you use Amazon Glacier for cloud storage provider, you can select from one of several retrieval modes: Expedited, Standard (default), or Bulk. Data retrieval speeds for these modes are 5 minutes, 5 hours, and 12 hours, respectively. For the data transfer pricing available with each method, see Amazon Web Services. Faster retrieval times will result in higher costs.

Note: For each retrieval option, AltaVault waits for the specified time prior to attempting to restore data from the cloud. Actual retrieval times from Amazon Glacier to Amazon S3 tier are subject to Amazon service level agreements.

When you use Microsoft Azure Archive for cloud storage, the data retrieval speed is set at 15 hours.

When you use Alibaba Archive for cloud storage, the data retrieval speed is several minutes.

The retrieval speed is the time it takes to make the data available for download after you send the initial request to the cloud. Due to this delay, if data is not available on the local cache, it cannot be paged back from the cloud on demand. In such cases, you must first manually restore the files to be read from the cloud to the local cache on AltaVault using either the prepopulation GUI or CLI commands. After the data is restored from the cloud, it can be read from the local cache.

When recovering data from the archive tier, NetApp requires that all the data segments related to a restore or retrieve by the backup application be present on the AltaVault cache first. This avoids error or retry conditions by the backup or archive application while it waits on the cloud service provider to send the data to the AltaVault cache. To recover the data segments associated with the files to the AltaVault cache, use the AltaVault prepopulation feature. This limits the time penalty for data retrieval to the waiting period associated your selected retrieval mode. If several, separate retrieve requests are made for individual data segments (for example, when a backup application restores files sequentially, where each file is not in the AltaVault cache), this can slow down the overall retrieval process from the cloud service provider, and the overall recovery.

Data prepopulation is described in the following sections:

[“Restoring data from the cloud using the prepopulation page” on page 263](#)

[“Restoring data from the cloud using the command-line interface” on page 264](#)

[“Automatic prepopulation” on page 268](#)

Restoring data from the cloud using the prepopulation page

The Prepopulation page provides a granular status of tasks that were started using the Management Console or the command-line interface. For each task, you can view the list of files that are being restored by the task.

To restore data from the cloud using the prepopulation page

1. Choose **Configure > User Permissions**.
2. Select **Role-based Accounts** and edit the user information to include Read/Write permissions for Prepop Settings.
3. Click **Apply**.
4. Choose **Configure > Prepopulation**.
5. Click the **Select File** tab to display the Prepopulation File Browser that contains a list of files that can be prepopulated.

The Prepopulation File Browser enables you to browse the files on the AltaVault shares. For each file, it displays the file size, modification time (appears when you hover the cursor over a specific file), and its estimated size on disk.

Note: In the list of filenames, “lrse” represents SnapMirror shares.

Select a file or a list of files, and click **Fetch Percent Locally Cached for selected files** to obtain the locally cached percent in the AltaVault cache. This process might be slow for large files.

-or-

Optionally, specify a list of complete path names, separated by the pipe (|) character, to the files that you want in the List of files to be prepopulated. Refer to the following examples:

To prepopulate the file /nfs/dir_1/dir_2/test-2.txt, enter the following path name:

/nfs/dir_1/dir_2/test-2.slabs

To prepopulate the files /nfs/dir_1/test-1.txt and /nfsv4/dir_1/dir_2/test-2.txt, enter the following path names:

/nfs/dir_1/test-1.slabs/nfsv4/dir_1/dir_2/test-2.slabs

To prepopulate all the files under /nfs, enter the following path name:

/nfs/*

To prepopulate the files /smb/app_1/test_file.txt and /smb/app_2/test_file.txt, enter the following path names:

/smb/app_1/test_file.txt/smb/app_2/test_file.txt

To prepopulate all the files under /smb/app_1/, enter the following path name:

/smb/app_1/*

6. Select the check box next to the file (displayed in the list) names you want to prepopulate.
7. Click **Prepopulate Selected Files**.

If using Amazon Glacier, select the retrieval method for the files: Expedited, Standard, or Bulk. If no option is selected, the recovery type defaults to Standard.

- Click OK to begin prepopulation of the selected files. AltaVault displays the Prepopulation Report Status page, providing information related to each job:

Field	Description
Job ID	AltaVault assigned job number associated with each retrieval operation.
Progress	Percent of job downloaded from the cloud to the AltaVault local cache.
Status	Retrieval operation status: <ul style="list-style-type: none"> New — No files have been identified for prepopulation. Creating — System is identifying files for prepopulation. Enqueued — The prepopulation task has been recorded. The AltaVault has not started processing it. Processing — AltaVault is identifying data that must be restored from the cloud. Requested — AltaVault has requested all the data from the cloud. Downloading — The system has started downloading data for the prepopulation request. For Amazon Glacier, the time it takes for this state to appear depends on the retrieval type. Completed — Retrieval from the cloud to the AltaVault has completed. Canceled — Data download from the cloud has been canceled. Failed — AltaVault was unable to restore all of the data and the task failed. Check the logs to determine the reason for failure.
Start Time	Time when prepopulation from the cloud to AltaVault was started.
Completion Time	Time when prepopulation from the cloud to the AltaVault completed.
Retrieval Type	Glacier retrieval option: Expedited (5 minutes), Standard (5 hours; default), or Bulk (12 hours).
Cancel job	Cancels a job that has not yet completed.

An email completion notification is sent to the email recipients configured to receive email notifications.

If the prepopulation job is successful, the email notification contains the following information:

For a successful prepop:

Subject: Prepopulation Job Completed

Body: Prepopulation job #[job id] has completed successfully.

If the prepopulation job fails, the email notification contains the following information:

Subject: Prepopulation Job Failed

Body: Prepopulation job #[job id] has failed. Please check the system log for more information.

To configure users who should receive email notifications, go to the [Configure > Email](#) page and configure the email notification settings.

Restoring data from the cloud using the command-line interface

- Connect to the AltaVault command-line interface using SSH.
- Enter the following command in configuration mode:


```
hostname (config) # datastore prepop {[num-days <number-of-days>] | [start-date *] [end-date *]} [pattern <pattern>] [retrieval-type] [dryrun]
```

The `retrieval-type` option applies when using Amazon Glacier.

The following table shows the parameter options:

Parameter	Description
num-days <number-of-days>	Specifies the number of last-modified days to start data retrieval (from the present date to the number of days you specify).
start-date <start-date>	Specifies the date from which the data retrieval should start. The system prepopulates the files modified on or before this date. Enter a date in the format yyyy-mm-dd.
end-date <end-date>	Specifies the date on which the data retrieval should end. Stop prepopulating files on or after this date. Enter a date in the format yyyy-mm-dd.
pattern <pattern>	Filters the data retrieved by the pattern you specify. The pattern specified contains a required internal share name created on AltaVault, one or more optional subfolder names from the external share name visible to the user, and finally a required regular expression describing the file or files to be prepopulated. The asterisk (*) symbol with the regular expression matches all characters.
retrieval-type	Specifies the Glacier retrieval option: Expedited (5 minutes), Standard (5 hours; default), or Bulk (12 hours). If no retrieval-type is specified, the default method is Standard.
dryrun	Estimates the size and status of the prepopulation job. No information is downloaded using this option. The size is shown as two values: Restore Size — Total size of the restore job, which includes information that is already available in the local cache and amount of data required to be downloaded from the cloud. Download Size — Amount of data that would be downloaded from the cloud to prepopulate a file or file set. This data is the data that is not already available on the local cache.

3. To view the current status of prepopulation, enter the `show datastore prepop` command:

```
hostname (config) # show datastore prepop jobs files
Job: Job 4002
  Status:           Failed
  Start Time:       2017/02/10 23:49:53
  Retrieval Type:   Standard
  Files in job:

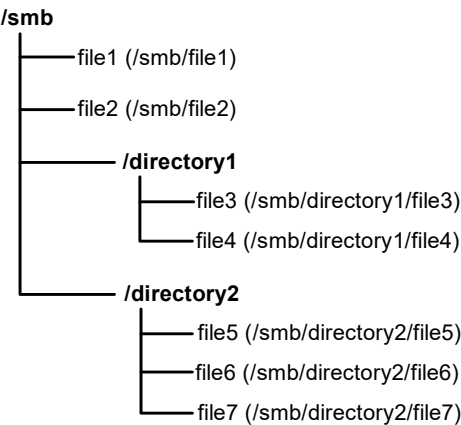
Job: Job 4003
  Status:           Completed
  Start Time:       2017/02/10 23:50:53
  Complete Time:    2017/02/10 23:56:59
  Retrieval Type:   Expedited
  Files in job:
    /nfs/test-1
```

The following output appears

Example 1: Pattern-based datastore prepopulation

This example explains pattern-based datastore prepopulation. Consider the directory structure example shown in the figure below.

Directory structure



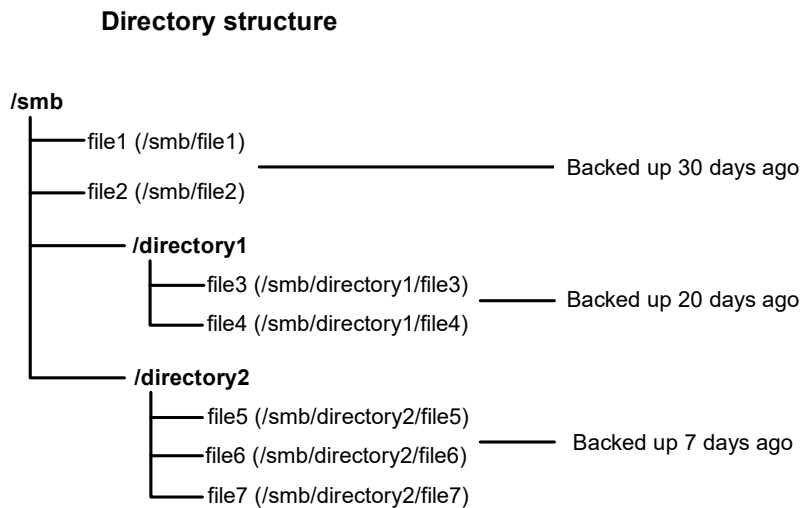
The following table shows different examples of the `datastore prepop` command for this directory structure:

Command	Description
<code>datastore prepop pattern smb/f*</code>	Populates only file1 and file2.
<code>datastore prepop pattern smb/* retrieval-type bulk</code>	Populates all of the files (file1 through file7) with directory1 and directory2 using Bulk retrieval.
<code>datastore prepop pattern smb/ directory1/* retrieval-type standard</code>	Populates only file3 and file4 using Standard retrieval.
<code>datastore prepop pattern smb/ directory1/file3 /smb/directory2/ file7 retrieval-type expedited</code>	Populates only file3 and file7 using Expedited retrieval.

The `datastore prepop` command operates from the local pathname for each SMB share created.

Example 2: Time-based datastore prepopulation

This example explains time-based datastore prepopulation. Consider the directory structure example shown in the figure below.



To obtain the most recent files backed up, enter the following command on the AltaVault command-line interface:

```
datastore prepop num-days 7
```

This command fetches data that is seven days old from the cloud.

To fetch files backed up 20 days ago using the using Amazon Glacier Expedited retrieval-type option, enter the following command:

```
datastore prepop num-days 20 retrieval-type Expedited
```

Example 3: Prepopulating from backups

In this example, assume that:

- All full backups are stored in a directory called fulls.
- All full backups for Host A are stored in a subdirectory called hostA.

To prepopulate all backups for Host A that occurred in the past 30 days (from the current time) using default (standard) retrieval, enter the following command:

```
hostname (config) # datastore prepop num-days 30 pattern fulls/hostA/*.img
```

To prepopulate all backups for Host A that occurred for a 24-hour duration starting on 2017-01-01 (YYYY-MM-DD), using expedited retrieval, enter the following command:

```
hostname (config) # datastore prepop pattern fulls/hostA/*.img start-date 2017-01-01 end-date 2017-01-02 retrieval-type expedited
```

To prepopulate all backups for Host A that occurred in the past 30 days (from the current time) using bulk retrieval, enter the following command:

```
hostname (config) # datastore prepop num-days 30 pattern fulls/hostA/*.img retrieval-type bulk
```

After this process finishes, you can initiate a restore process using the restore feature of the backup application. For details about how to restore your backups, refer to the relevant documentation for your backup application.

Automatic prepopulation

You can use settings in AltaVault to automatically trigger prepopulation of a file when you try to read the file and find that data must be restored from the cloud. For example, a backup application attempt to restore data from files stored in the archive tier will fail but trigger an automatic prepopulation request. For S3, automatic prepopulation uses Standard retrieval method. If the backup application can be configured to retry the restore operation after the retrieval time frame of the archive tier, the retry will succeed without any additional user intervention to perform prepopulation.

To enable automatic prepopulation

1. Connect to the AltaVault command-line interface using SSH.
2. In configuration mode, enter the following command:

```
hostname (config)# datastore prepop auto-enable
```

Note: Enabling automatic prepopulation settings can trigger the prepopulation of entire files from the cloud upon read failures. This can result in restore charges related to recovering the data back from the cloud, as well as possible eviction of backup data in order to place the recovered data on cache. Enable automatic prepopulation only after careful consideration.

To confirm the prepopulation status, enter the following command:

```
hostname (config)# show datastore prepop auto-enable
autoprepop.enable: true
```

To disable automatic prepopulation

1. Connect to the AltaVault command-line interface using SSH.
2. In configuration mode, enter the following command:

```
CLI (config)# no datastore prepop auto-enable
```

To confirm the prepopulation status, enter the following command:

```
CLI (config)# show datastore prepop auto-enable
autoprepop.enable: false
```

The discovery of which files have data which must be retrieved from Amazon Glacier to the AltaVault cache varies by application. Recommendations are highlighted for specific applications in the remainder of this chapter.

AltaVault appliance best practices for EMC NetWorker

The following process allows you to stage a restore operation, identifying the files required for the restore and migrate the files from the archive tier to AltaVault.

1. Stage a *saveset* restore operation, by issuing the following command from the NetWorker command line:

```
Networker User > Operation > Save Set Recover > Source Client (oak-cs.cb2k3r2.com)
Save Set Name:  D:\ (number of versions:  2)
Version Date:
[x] 6/13/2013 11:06 AM    38765393    browsable
[ ] 6/13/2013 11:49 AM    55531031    browsable

Required Volumes:  oak_cs.cb2k3r2.com.001.RO
```

2. Assume that you want the saveset from "6/13/2013 11:06 AM" and go to the Windows command prompt and run the following commands:

```
cd "c:\Program Files\Legato\nsr\bin"

mminfo -otc -v -q name=D:\
volume      type      client      date      time      size  ssid      fl  lvl  name
oak_cs.cb2k3r2.com.001  adv_file oak-cs.cb2k3r2.com 6/13/2013 11:06:53 AM 38 GB 4273605193 cb manual D:\
oak_cs.cb2k3r2.com.001.RO adv_file oak-cs.cb2k3r2.com 6/13/2013 11:06:53 AM 38 GB 4273605193 cb manual D:\
oak_cs.cb2k3r2.com.001  adv_file oak-cs.cb2k3r2.com 6/13/2013 11:49:49 AM 55 GB 4256830541 cb manual D:\
oak_cs.cb2k3r2.com.001.RO adv_file oak-cs.cb2k3r2.com 6/13/2013 11:49:49 AM 55 GB 4256830541 cb manual D:\
```

3. Using the SSID value from the output of the above, find the name of the corresponding file on the disk:

```
mminfo -q "ssid=4273605193" -r "ssid(53)"
8d5688a2-00000006-feba0a49-51ba0a49-00030e00-a8d346b6
```

4. Now that the filename needed for restore or retrieve has been identified, locate the file on the AltaVault appliance share or mount.

Linux:

```
find <AVAmntpointname> -name "8d5688a2-00000006-feba0a49-51ba0a49-0003"
< AVAmntpointname >/networker/66/76/notes/8d5688a2-00000006-feba0a49-51ba0a49-00030e00-a8d346b6
< AVAmntpointname >/networker/66/76/8d5688a2-00000006-feba0a49-51ba0a49-00030e00-a8d346b6
```

Windows:

Map AltaVault share path as Windows mapped network drive (Z: in this example)

```
Z:\>dir "8d5688a2-00000006-feba0a49-51ba0a49-00030e00-a8d346b6" /s
```

Volume in drive Z is LKWZE

Volume Serial Number is 009A-9A03

Directory of Z:\66\76

```
06/13/2013 11:19 AM 39,695,762,716 8d5688a2-00000006-feba0a49-51ba0a49-00030e00-a8d346b6
1 File(s) 39,695,762,716 bytes
```

Directory of Z:\66\76\notes

```
06/13/2013 11:19 AM 236 8d5688a2-00000006-feba0a49-51ba0a49-00030e00-a8d346b6
1 File(s) 236 bytes
```

Total Files Listed:

```
2 File(s) 39,695,762,952 bytes
```

```
0 Dir(s) 499,903,440,459,776 bytes free
```

Note: The mminfo query for the session id (SSID) should list out multiple files if the backup spans multiple media files. You need to list out all levels of the backup through mminfo. To get the SSID, you can also go to Network Administration > Media > Disk Volumes > Show Save Sets or Media > Save Sets and query them.

5. Prepopulate the file or files identified in Step 4 using the Prepopulation GUI as described in [“Restoring data from the cloud using the prepopulation page” on page 263](#), and wait until the files migrate from the archive tier to the AltaVault cache.
6. Initiate your restore from *Networker User* as you normally would to complete the recovery.

AltaVault appliance best practices for IBM Spectrum Protect

The following process allows you to identify AltaVault volumes and initiate a Spectrum Protect client restore.

1. On the Spectrum Protect server administrative command line issue the following SELECT statement to identify the object ID for the file you want to restore:

```
select * from backups where node_name='<tsmclientnodename>' and ll_name='<filenamestore>'
select * from backups where node_name='CLIENT1' and LL_NAME='MYFILE.TXT'
```

```

      NODE_NAME: CLIENT1
    FILESPACE_NAME: \\CLIENT1\s$
      FILESPACE_ID: 1
        STATE: ACTIVE_VERSION
        TYPE: FILE
      HL_NAME: \BACKUP\SET2\
      LL_NAME: MYFILE.TXT
      OBJECT_ID: 1062
    BACKUP_DATE: 2013-04-23 20:02:38.000000
  DEACTIVATE_DATE:
      OWNER:
    CLASS_NAME: DEFAULT

      NODE_NAME: CLIENT1
    FILESPACE_NAME: \\CLIENT1\s$
      FILESPACE_ID: 1
        STATE: ACTIVE_VERSION
        TYPE: FILE
      HL_NAME: \BACKUP\SET3\
      LL_NAME: MYFILE.TXT
      OBJECT_ID: 6786
    BACKUP_DATE: 2013-04-23 20:06:19.000000
  DEACTIVATE_DATE:
      OWNER:
    CLASS_NAME: DEFAULT
```

2. From the list of file versions above, identify the version needed (for this example, the version from 20:06:19 on 4/23/2013), and then issue the following command using the OBJECT_ID value from the output.

Note: Note that if the *bitfile* is part of a super-bitfile, rerun the below command against the super-bitfile OBJECT_ID below.

```
show bfo 6786
Bitfile Object: 6786
Active
**Sub-bitfile 6786 is stored in the following aggregate(s)
Super-bitfile: 6783, Offset: 2000, Length 694, Deduped: F
```

```
show bfo 6783
Bitfile Object: 6783
**Super-bitfile 6783 contains following aggregated bitfiles,
Bitfile Id, offset, length, active state or owner, link bfid
6783          0          671      Active
6784          671         663      Active
6785         1334         666      Active
6786         2000         694      Active
6787         2694        1406      Active
6788         4100         676      Active
```

```

6789          4776          679          Active
.....
**Sub-bitfile 6783 is stored in the following aggregate(s)
   Super-bitfile: 6783, Offset: 0, Length 671, Deduped: F

**Disk Bitfile Entry
   Bitfile Type: PRIMARY   Storage Format: 22
   Logical Size: 25880969   Physical Size: 25886720   Number of Segments: 1,
Deleted: False
   Storage Pool ID: 4   Volume ID: 3   Volume Name: H:\TSMVOLS\AVAAVOL001.BFS

```

3. Prepopulate the volume identified in Step 2, using the Prepopulation GUI as described in [“Restoring data from the cloud using the prepopulation page” on page 263](#), and wait until the files migrate from the archive tier to the AltaVault cache.
4. Initiate your restore from the Spectrum Protect client as you normally would to complete the recovery.

To identify all the volumes related to a AltaVault based storage pool

1. Issue the following Spectrum Protect administrative SELECT command, using the appropriate storage pool name that points to AltaVault:

```

select volume_name from volumes where stgpool_name='<AVASTGPPOOLNAME>'

select volume_name from volumes where stgpool_name='AVACOPYPOOL'

VOLUME_NAME: \\AltaVault-01\TSM\00000002.BFS
VOLUME_NAME: \\AltaVault-01\TSM\00000003.BFS
VOLUME_NAME: \\AltaVault-01\TSM\00000004.BFS

```

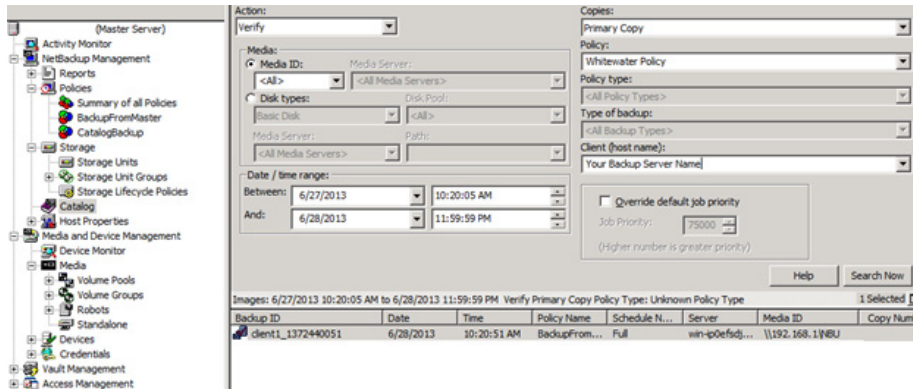
2. Prepopulate the volumes identified using the Prepopulation GUI as described in [“Restoring data from the cloud using the prepopulation page” on page 263](#), and wait until the files migrated from archive to the AltaVault cache.
3. Perform primary Spectrum Protect storage pool recovery as appropriate if the storage pool above is a copy storage pool.
4. Initiate your restore from the Spectrum Protect client as you normally would to complete the recovery.

AltaVault appliance best practices for Veritas NetBackup

The NetBackup Catalog maintains an inventory of the backups and can be used to identify which media volumes are required for restore.

1. Go to the Catalog menu item from the NetBackup GUI and search for the backup from which you want to restore data.

2. Use the filter criteria to select the policy that goes to AltaVault, the client you want to recover data from, and the date and time range.



3. In the results field at the bottom of the page, identify the Backup_ID.
This corresponds to a portion of the file name of the volume created on AltaVault for the backup.
4. Prepopulate the identified volume identified in step 2, using the Prepopulation GUI as described in [“Restoring data from the cloud using the prepopulation page” on page 263](#), and wait until the files migrate from the archive tier to the AltaVault cache.
5. Initiate your restore from the NetBackup GUI as you normally would to complete the recovery.

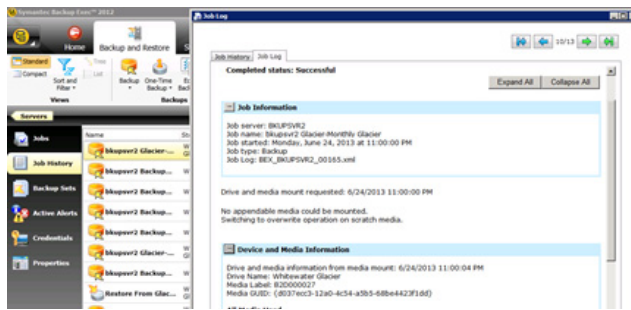
AltaVault appliance best practices for Veritas Backup Exec

The Backup Exec job activity page maintains inventory of the backups, which can be used to identify which media volumes are required for restore. Those volumes can then be located by referring to the Backup-To-Disk-Folder or Storage target and locating the volume from the corresponding AltaVault SMB share.

To identify which medium volumes are required for restore for Backup Exec 2012:

1. Select the Backup and Restore tab
2. Select the server from which you want to restore objects.
3. Select the Job History tab from the left side.
4. Select the backup job you want to restore from the list of backups that display.

5. In the Job Log page that displays, select the Job Log tab and find the media volumes used.



Note: Backup Exec users sending data to archive should run the following AltaVault CLI command to ensure that media recycling of volumes on AltaVault appliances occurs correctly with respect to archive storage: `megastore keep-bkf-local enable`.

6. Prepopulate the file or files identified in Step 5 using the Prepopulation GUI as described in [“Restoring data from the cloud using the prepopulation page” on page 263](#), and wait until the files migrate from the archive tier to the AltaVault cache.

AltaVault appliance best practices for Veeam backup and replication

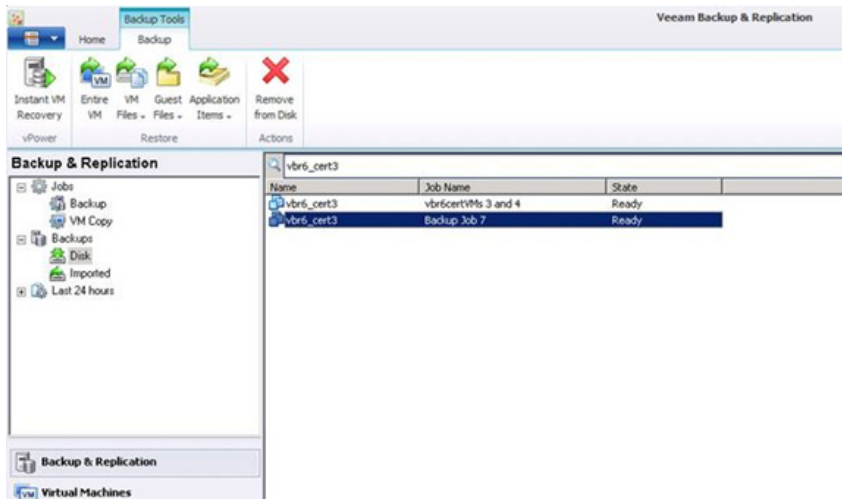
For Veeam Backup & Replication, folders are named after their respective job numbers in Veeam, as are the files.

Note: Note that all of the files and folders are time stamped automatically starting from Veeam 5.

To locate the backup required for prepopulation

1. Locate which backup job to which the VM belongs.

2. In this case, the example is VM is *vbr6_cert3*. In Veeam, you can locate the backup job by searching the backup database. In the figure below, you can see that *vbr6_cert3* has been backed up by a job named *Backup Job 7*.



3. After identifying the job name, prepopulate the most recent *.vbk* file plus all subsequent *.vib* files, and the job metadata file with the *.vbm* extension, using the Prepopulation GUI as described in “[Restoring data from the cloud using the prepopulation page](#)” on page 263, and wait until the files migrated from the archive tier to the AltaVault cache.

Note: There is no need to prepopulate older backup chains. Veeam time stamps the backup job files to make it easy to identify.

The figure below shows an example of how to prepopulate the current Veeam backup files needed to restore the most recent version of the VM named vbr6 cert3.



4. Initiate your restore from the Veeam Backup and Replication GUI as you normally would to complete the recovery.

Copyright Information

Copyright © 1994-2018 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>.

How to Send Your Comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

Numerics

802.3ad 61

A

Access Control List 103

Access Key 25, 27, 33

access policy 75

adding 77

Account Name 87

Accounts

capability-based 85

privileges 85

Action 104

Add 43, 58, 59, 60, 61, 88, 92, 93, 104

Add a New Group 76

Add a New Log Server 79

Add a New NTP Authentication Key 71

Add a New NTP Server 70

Add a New Route 58, 59, 60

Add a New Rule 104

Add a New Security Name 76

Add a New TCP Dump 156

Add a New Trap Receiver 72

Add a New User 74

Add a Radius Server 91

Add a TACACS+ Server 93

Add an Export 45, 46

Add Share 42

Add SMB Share 41

Add SMB User 42, 43

Admin 85

administrator configuration worksheet 241

Administrator password 85

Admission Control 65, 143

advance features 243

Alarm

admission control 143

cloud bucket consistency 143

cloud bucket disparity 143

cloud bucket over capacity 143

CPU utilization 143

data integrity error 143

datastore eviction 143

licensing 144

link state 145

memory paging 145

process dump staging directory
inaccessible 145

secure vault 145

software update available 145

system disk full 68, 146

Alarm status

link duplex error 144

link I/O error 145

Alarm Status report

definition of 143

viewing 146

Alarm thresholds, setting 63

Alert 79, 80, 81

Alibaba

configuring Object Storage Service
(OSS) 24

recovering data from Archive 262

Allow All Clients 46, 47

Allow Everyone Access 42

Allow Session Timeouts When Viewing

Auto-Refreshing Pages 94

Allow Specified Clients 46, 47

AltaVault

bezel 196

chassis components 196

components 193

deployment guidelines 13

AltaVault appliance

deployments 11

Amazon

configuring Glacier storage 25

configuring S3 storage 27

migrating to Amazon Glacier 178

recovering data from Glacier 262

transferring data using Snowball 169

Announcement, setting on home page 63

appliance 173

Appliance Information 19

Appliance monitoring report, viewing 162

Appliance monitoring, configuring 97, 101

Apply Changes 123

Atmos-based storage

configuring 28

Authentication 72, 74

Authentication methods 84

Local 84

Authentication Port 91, 93

Authentication Protocol 72, 74

Authentication Type 91, 93

authentication type

- standard 25
- Authentication, setting general security 83, 129
- Authorization Policy 84
- Autodiscover rules, overview of 104
- Autolicense critical event 67
- Azure
 - configuring storage 30
 - recovering data from Archive 262
- Azure Archive
 - storage class 30
- B**
- Back-End Throughput Optimization report 137
- Bandwidth Limit, setting 39
- battery
 - disposal 239, 240
- BBU card
 - location 237
- boot device 229
 - install 231
 - removal 230
- boot image 229
- boot media device 198
- C**
- Cancel This Job 123
- Capability-based accounts 85
- chained authentication using SSH 104
- Change Password 94
- CHAP 91
- chassis 196
 - depth 194
 - description 196
 - handles 196
 - height 194
 - installation 218
 - removal 218
 - slots 198
 - weight 194
 - width 194
- CIFS users, adding 48
- cli 81
- CLI configuration wizard 17
- cloud
 - migration 169, 175, 176
 - settings 190
- Cloud and Disk Storage Allocation 19
- cloud appliance
 - deactivating support 126
- Cloud Bucket Consistency 65, 143
- Cloud Bucket Disparity 65, 143
- Cloud Bucket Over Capacity 65, 143
- Cloud Information 19
- cloud migration 175
- Cloud provider credentials 241
- Cloud Settings wizard
 - using 22
- cloud storage configuration 241
- Cloud Storage Reclamation 19
- Comment 41, 45, 46, 123
- Common Name 95
- Community 73
- Community String 76
- configuration
 - file export 188
 - recovery 188
- Configuration files, managing 128
- configuration worksheet 241
- Configuration, saving 20
- Confirm Global Key 92
- Confirm Server Key 92, 93
- Continuous log, viewing 154
- controller
 - module installation in chassis 216
 - moving FRUs 215
 - re-installation 215
 - removal 214
 - replacement 213
 - replacing chassis 218
 - shut down 212
- CPU utilization 65, 143
- CPU Utilization report
 - definition of 146
 - viewing 147
- Created 123
- Critical 79, 80, 81
- Current Configuration 128
- Current Password 94
- D**
- data
 - preservation 213
 - restoration 188
 - restoration for disaster recovery 191
- Data Integrity Error 143
- Data interface routing, configuring 59
- Data interfaces, modifying 58
- data recovery
 - Alibaba Archive 262
 - Amazon Glacier 262
 - Microsoft Azure Archive 262
 - restoring data 263
- data restoration
 - disaster recovery testing 189
- data storage
 - optimizing data movement 261
 - protecting 261
- Datastore Eviction 65, 143
- Datastore Low Space 65, 143
- Deduplicated Data 135
- Deduplication 19
- Deduplication factor 135
- Default User 84
- Deny in-path rules, overview of 104
- deployment
 - guidelines 13
 - steps 15
- Description 104
- description
 - rack units 194
 - system chassis 194
- Destination 59, 60
- Destination IPv4 Address 58
- Destination Port 72, 104
- DHCP 18
- DIMM
 - removal 233
- Disable Compression 41, 45, 47
- Disable Dedup 41, 45, 47
- disaster recovery

- definition 187
- preparation 188
- testing 188
- testing activities 190
- Disk Full 66, 144
- Disk Statistics report
 - definition of 149, 150
 - viewing 149, 150
- Disk Utilization report
 - definition of 151
 - viewing 151
- DNS settings, specifying 56
- Dynamic DNS 21, 57
- E**
 - Early Eviction 41, 45, 46
 - Edit Export 46
 - Email notification, setting 78
 - Emergency 79, 80, 81
 - Enable Account 88
 - Enable Data Interface 59, 60
 - Enable Scheduling 142
 - Enable SNMP Traps 72
 - Enable/Disable Job 123
 - Enabled 92, 93
 - Error 79, 80, 81
 - Event and failure notification, setting 78
 - Eviction Optimization report
 - definition of 138
 - viewing 138
 - Execute Now 123
 - Executes On 123
 - Expanded Data 135
 - expansion shelves
 - about 194
 - Export Asynchronously 46, 47
 - Export Configuration wizard
 - using 35
 - Export, adding 45, 46
- F**
 - failed part
 - returns 240
 - Fan Error 66, 144
 - fan module
 - control 198
 - controller 198
 - description 196
 - detection 198
 - failure 219
 - hot swap 219
 - removal 220
 - fan module failure 201
 - Filter logs 154
 - Fingerprint 95
 - FIPS-mode
 - verifying your system is compliant 114
 - From Local File 127
 - From URL 127
 - Front-End Throughput Optimization report 136
 - FRU 194
 - hot-swappable 197
 - replacements 197
 - system 198
 - types 197
- FRUs
 - internal 228
- FTP proxy access 55
- G**
 - Gateway 59, 60
 - Gateway IPv4 Address 58
 - Glacier 175
 - Global Key 91, 92
 - Google Cloud Storage
 - configuring 29
- H**
 - hald 81
 - Home page announcement, setting 63
 - Home page, overview of 19
 - Host settings, modifying 65
 - Hostname 70, 93
 - Hostname, changing 55
 - Hosts 55
 - hot-swap
 - fan module 219
 - PSUs 222
- I**
 - Import Configuration wizard
 - using 34
 - Info 79, 80, 81
 - Install 127
 - install
 - AVA10S Shelves 225
 - chassis 218
 - controller in chassis 216
 - fan module 221
 - Interface 57, 104
 - management 57
 - interface
 - primary 57
 - Interface Statistics report
 - definition of 148
 - viewing 149
 - internal FRU
 - replacement 228
 - Interval 123
 - IP Address 70, 93
 - IP Configuration 59, 60
 - IPMI 66, 144
 - IPv4 address
 - obtaining automatically 21, 57
 - specifying manually 22, 58
 - IPv4 Subnet Mask 58
 - Issued By 95
 - Issued To 95
- J**
 - Jobs
 - scheduling 122
 - viewing details 122
- K**
 - Key 74, 95
 - Key ID 70, 71
- L**
 - Last Run 123
 - Licensing 67, 144

- Lines Per Log Page 79
- Link Duplex 67
- Link duplex alarm status 144
- Link I/O error alarm status 145
- Link I/O Errors 67
- Link state 67, 145
- Local logging 69
- Local logging, setting 78
- Local Only 84
- Local Path 45
- Local path 46
- Log Packets 104
- Log server, adding or removing 79
- Logs
 - downloading 154
 - filtering 80
 - rotating 78
 - system logs, downloading 155
 - system logs, viewing 152
 - user logs, downloading 154
 - user logs, viewing 153
 - viewing 152
 - viewing continuous 153
- long-term retention, SnapMirror 52
- Low Memory 67

M

- Main routing table 57
- Management ACL 103
- Management console
 - connecting to 18
 - navigating 19
- Management interfaces, modifying 57
- Manually specifying
 - primary interface IPv4 address 22, 58
- Max Pinnable Limit 67
- Maximum No. of Log Files 79
- MD5 72, 74
- MD5 key identifier 71
- Media Independent Interface 61
- Megastore GUID
 - reset 216
- Member Interfaces (comma-separated) 59, 60
- Memory Error 66
- Memory Paging 67, 145
- Memory Paging report
 - definition of 147
 - viewing 148
- Message of the day 63
- mgmtd 81
- Microsoft Azure
 - Archive 30
 - Standard 30
- migration
 - monitor cloud 176
- Minimum Severity 79, 80
- Mode 61
- Modifying
 - data interfaces 58
 - host settings 55
 - management interfaces 57
- Monitor 85
- Monitor password 85
- monitored appliance 101
- monitoring appliance 101

- Monitoring interval 61
- MOTD, setting 63
- MOTD. See Message of the day.
- Mount Commands 47
- Move Selected 104
- MTU 58
- MTU value, setting 58, 59, 60
- multifactor authentication using SSH (see chained authentication) 104

N

- Name 45, 55, 123
- network connectivity information 242
- Network File System 44
- New host, adding 56
- New Password 94
- New Password Confirm 94
- NFS
 - configuring 44
- NFS, adding an export 45, 46
- NFS. See Network File System.
- Notice 79, 80, 81

O

- Object Identifiers, viewing through
 - SNMP 77
- Offline File System Check
 - definition of 165
 - viewing report 165
- Online File System Check 166
- OpenStorage Technology (OST) 48
- Optimization Service 19
- OST 48
 - configuring 48
- Override the Global Default Key 92, 93

P

- Pages and reports, printing 20
- PAP 91
- Password 94
- password policy 88
- password set
 - service processor 210
- Path 41
- Perl regular expression 160
- Pin Export 45
- Pin Share 41
- Pinned Data 41
- pinned data 48
- Pinned Data Information 44
- pm 81
- power supplies 194
- Power Supply 66
- power supply
 - alerts and logs 222
 - install 224
 - remove 223
- power supply units
 - hot swap 222
- power supply units (PSUs) 222
- Prepopulation
 - restoring data 263
 - viewing 263
- Primary Interface 57
- Primary interface

- setting 57
 - specifying IPv4 address 22, 58
- private cloud
 - customizing 34
- Private cloud, customizing 34
- Process Dump Creation Error 67
- Process Dump Staging Directory
 - Inaccessible 145
- Process dumps, displaying and downloading 156
- Process dumps, viewing 156
- protecting data 261
- Protocol 104
- Protocols
 - NFS 44
 - OST 48
 - SMB 39
 - SnapMirror 50
- Proxies 55
- Proxy
 - addresses for Web access 55
 - setting an IP for Web/FTP 56
- PSU
 - remove 222

Q

- Queue
 - capture file 157
 - specifying the trace dump size 158

R

- rack unit
 - specifications 194
- RADIUS authentication method, setting 83
- RAID 66
- RAID battery information 19
- RAID controller 237
- RAID controller assemblies 198
- RAID group
 - adding disk 12-pack 227
 - disk replacement 227
 - update 216
- RAID6 194
- Read-Only Community String 72
- real-time clock 198
 - replacement 237
- Reboot 127
- Receiver 72
- Receiver Type 72
- Remote First 84
- remote management port
 - configuration 210
 - validation 211
- Remote Only 84
- Remote User 72
- Remove Selected 42, 43, 46, 58, 59, 60, 92, 93, 104
- Remove Selected Jobs 123
- Remove Selected Users 88
- Replicated Data 19
- Replication 38
 - definition 19, 37
 - scheduling 38
- Replication Error 68, 146
- Replication Optimization report
 - definition of 139

- viewing 139, 141
- Replication Paused 68, 146
- Reports
 - Alarm Status 143
 - Appliance Monitoring 162
 - Back-End Throughput Optimization 137
 - CPU Utilization 146
 - Disk Statistics 149, 150
 - Disk Utilization 151
 - Eviction Optimization 138
 - Front-End Throughput Optimization 136
 - Interface Statistics 148
 - Memory Paging 147
 - Offline File System Check 165
 - Replication Optimization 139
 - Scheduled Reports 141
 - Storage Optimization 135
 - TCP trace dump 156
 - Verify 167
- restoring data 188
- Restrict inbound IP access to SteelStore
 - gateway 103
- Retries 91, 92, 93
- RMA (Return Merchandise Authorization) 213
- Role-based
 - user permissions 85
- Rotate Based On 79
- RTC battery
 - install 238
 - replacement 238
- Rule Number 104

S

- S3-based storage
 - configuring 31
- Save Current Configuration 128
- Scanning system logs 160
- sched 81
- Schedule Reports 141
- Schedule Upgrade for Later 127
- Scheduled Reports
 - purpose 141, 142
- Scheduled reports
 - configuring 142
- Scheduling, replication 38
- Secret (Text) 71
- Secret Key 29, 31
- Secure access by inbound IP address 103
- Secure Access Policies 75
- Secure Groups 75
- Secure groups, setting 76
- Secure Vault 67, 145
 - unlocking and changing the password 93
- Secure Vault Locked 145
- Secure Views 75
 - setting 77
- Security Models 75
- Security Models and Name Pairs 76
- Security Name 76
- Security Names 75
- Security, general 83, 129
- Server IP 79
- Server Key 92, 93
- Service 104

- service processor (SP)
 - set password 210
 - accessing 211
 - Services, starting, stopping, restarting 121
 - Set a Global Default Key 92
 - SHA 72, 74
 - Share Name 41
 - Share pinning 41, 44
 - share pinning 48
 - Shared Secret 28
 - Shelf Details report, viewing 164
 - Shelf Error 68
 - short-term retention, SnapMirror 52
 - Shut down 127
 - Single Sign-On (SSO) 106
 - SMB
 - configuring 39
 - SMB users, adding 40, 42, 43, 49
 - SMTP Port 78
 - SMTP Server 78
 - SnapCenter access 53
 - SnapMirror
 - CLI configuration 53
 - configuring 50
 - long-term retention 52
 - short-term retention 52
 - SnapCenter access 53
 - SNMP
 - access policies 77
 - access policy security 73, 77
 - access policy, adding 77
 - ACLs 72
 - adding groups 76
 - adding or removing trap receivers 72
 - adding views 77
 - configuring v3 73
 - features 75
 - including specific OIDs in a view 77
 - MIB, accessing 209, 245
 - secure groups, setting 76
 - secure views, setting 77
 - setting trap receivers 71
 - supported versions 71
 - testing a trap 73
 - traps, summary of sent 247
 - Snowball
 - guidelines with AltaVault 169
 - managing transfers 172
 - prerequisites 169
 - seeding data 170
 - transferring data to Amazon 169
 - verifying transfer 173
 - Software Update Available 145
 - Software, upgrading 126
 - Source IP Address and Mask Bits 76
 - Source Network 104
 - specifications
 - AltaVault chassis 194
 - SSH
 - chained authentication 104
 - configuring access 104
 - SSL
 - peering list 95
 - SSO (Single Sign-On) 106
 - standard authentication type 25
 - Storage Optimization 19
 - Storage optimization report, viewing 135
 - Storage Optimization Service 68, 146
 - Storage Optimization Service Down 68, 146
 - Storage Optimization Service Error 68, 146
 - Storage Optimization Service Replication 68
 - Sub Tenant ID 28, 30
 - Subnet Mask 59, 60
 - support activation
 - deactivating support 126
 - SWIFT-based storage
 - configuring 32
 - configuring with region-selection 33
 - Synchronizing, AltaVault appliance to NTP server 70
 - system
 - re-cabling 216
 - restart 216
 - system chassis specifications 194
 - System Contact 72
 - system DIMM 198
 - install 235
 - removal 233
 - System Disk Full 68, 146
 - System dump, viewing 155
 - system LEDs
 - behavior 196
 - description 196
 - System Location 72
 - System logs, downloading 155
 - System logs, viewing 152
 - system memory
 - DIMMs 233
 - System Settings wizard
 - using 21
 - System Status 19
 - System, logging out of 20
- T**
- TACACS+ authentication method, setting 83
 - TCP dump 156
 - TCP trace dump 156
 - Temperature 69
 - Tenant ID 32
 - Throughput Optimization report
 - viewing 136, 137
 - Time zone 69
 - Timeout 91, 92, 93
 - Transmit Load Balance 61
 - Transmit/Receive Load Balance 61
 - Trap receivers, adding or removing 72
 - Traps 71
 - Traps, summary of SNMP traps sent 247
- U**
- UID 28, 30
 - Unlock Secure Vault 94
 - User logs
 - downloading 154
 - viewing 153
 - User Name 42, 43
 - User permissions 85
 - using Cloud Settings 22

V

- Validity 95
- Vault, unlocking and changing the password 93
- Verify report
 - definition of 167
 - viewing 167
- Version 70
- View-Based Access Control Mechanism 71
- Virtual Interface 59
- Virtual Interface Name 60
- VLAN 61
- VLAN Interface 60

W

- Warning 79, 80, 81
- wdt 81
- Web Inactivity Timeout 94
- webasd 81
- Wizard dashboard 22
 - accessing 21
 - exporting configurations 35
 - importing configurations 34
 - using 20, 21