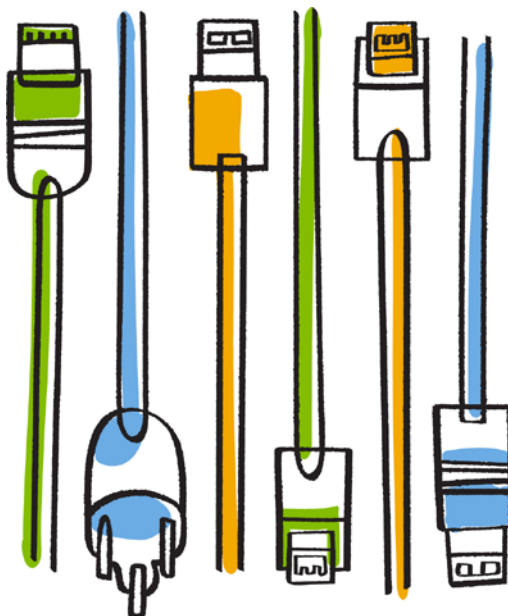


NetApp® AltaVault® Cloud Integrated Storage 4.0

Deployment Guide



© 2015 NetApp, Inc. All rights reserved.

No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. 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, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and Whitewater are trademarks or registered trademarks of NetApp, Inc. and its affiliated entities in the United States and/or other countries. Riverbed is a trademark of Riverbed Technology used pursuant to license. Any other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of certain of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.asp>.

Contents

Contents	iii
Preface.....	1
About This Guide	1
Audience.....	1
Document Conventions	2
Documentation and Release Notes	2
How to Send Your Comments	3
Chapter 1 - Overview of AltaVault Appliance	5
What is AltaVault Appliance?	5
About AltaVault Appliance	5
Chapter 2 - Overview of AltaVault Appliance Deployment	9
AltaVault Deployment Guidelines	9
Deployment Checklist	10
Deployment Diagram	11
Deployment Steps.....	11
Chapter 3 - Disaster Recovery	13
Preparing for Disaster Recovery.....	13
Exporting the Configuration File.....	14
Preparing for Disaster Recovery Testing.....	15
Recovering an AltaVault Configuration	16
Restoring Data for Disaster Recovery	19
Chapter 4 - Data Prepopulation	21
Prepopulating Data	21
Using Prepopulation with Amazon Glacier Cloud Storage.....	28
Automatic Prepopulation	28
Time-Based Automatic Prepopulation	28
Chapter 5 - Configuring Remote Management	29
Configuring Remote Management	29

Remote Management Configuration Examples.....30

Chapter 6 - Cloud Agility.....33

Cloud Agility Overview33

Migration Process34

Chapter 7 - Monitoring Peer Appliances35

Benefits of Monitoring Peer Appliances35

Configuring Appliance Monitoring35

Configuring REST API Access36

Specifying the API Access Code.....37

Index39

Preface

Welcome to the *NetApp AltaVault Cloud Integrated Storage Deployment Guide*. Read this preface for an overview of the information provided in this guide and the documentation conventions used throughout, hardware and software dependencies, and contact information. This preface includes the following sections:

- [“About This Guide” on page 1](#)
- [“Documentation and Release Notes” on page 2](#)
- [“How to Send Your Comments” on page 3](#)

About This Guide

The *NetApp AltaVault Cloud Integrated Storage Deployment Guide* serves as a design guide that helps you deploy and troubleshoot the NetApp AltaVault® Cloud Integrated Storage appliance (AltaVault).

This guide assumes that you are familiar with using the AltaVault command-line interface (CLI) as described in the *NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Manual*.

Audience

This guide is written for storage and backup administrators familiar with Storage Area Network (SAN), Network Attached Storage (NAS), and cloud storage. NetApp assumes that you are already familiar with AltaVault and how it functions.

You must also be familiar with:

- Installing and configuring AltaVault. For details, see *NetApp AltaVault Cloud Integrated Storage User's Guide*.
- Connecting to the AltaVault command-line interface. For details, see the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances*.

Document Conventions

This guide uses the following standard set of typographical conventions.

Convention	Meaning
<i>italics</i>	Within text, new terms, emphasized words, and REST API URIs appear in <i>italic</i> typeface.
boldface	Within text, CLI commands, CLI parameters, and REST API properties appear in bold typeface.
Courier	Code examples appears in Courier font: <pre>amnesiac > enable amnesiac # configure terminal</pre>
< >	Values that you specify appear in angle brackets: interface <ipaddress>
[]	Optional keywords or variables appear in brackets: ntp peer <addr> [version <number>]
{ }	Required keywords or variables appear in braces: {delete <filename>}
	The pipe symbol represents a choice to select one keyword or variable to the left or right of the symbol. The keyword or variable can be either optional or required: {delete <filename> upload <filename>}

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>.

The following documents are provided in this release:

- *NetApp AltaVault Cloud Integrated Storage Getting Started Guide*
- *NetApp AltaVault Cloud Integrated Storage Deployment Guide*
- *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances*
- *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances*
- *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Cloud Appliances*
- *NetApp AltaVault Cloud Integrated Storage User's Guide*
- *NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Manual*

The following guide is referenced for instructions on safely installing the AVA10S shelves: *SAS Disk Shelves Installation and Service Guide for DS4243, DS2246, DS4486, and DS4246*. The AltaVault AVA10S shelf is identical to the DS4246 disk shelf.

If you need more information, see the NetApp Knowledge Base for any known issues, how-to documents, system requirements, and common error messages. You can browse titles or search for keywords and strings. For more information, see the NetApp Support site at <https://mysupport.netapp.com>.

Each software release includes release notes. The release notes identify new features in the software as well as known and fixed problems. To obtain the most current version of the release notes, go to the Software section under Downloads on the NetApp Support site at <https://mysupport.netapp.com>.

Examine the release notes before you begin the installation and configuration process.

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 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

CHAPTER 1 Overview of AltaVault Appliance

This chapter provides an overview of AltaVault cloud integrated storage. It includes the following sections:

- [“What is AltaVault Appliance?” on page 5](#)
- [“About AltaVault Appliance” on page 5](#)

What is AltaVault Appliance?

AltaVault appliance is a disk-to-disk data backup and archive storage optimization system with unique cloud storage integration. AltaVault integrates seamlessly with your existing backup and archive technologies and cloud storage provider APIs to provide rapid replication of data to the cloud for offsite storage and rapid retrieval.

AltaVault is a replacement for tape, virtual tape library (VTL), and disk-to-disk technology. AltaVault becomes the backup target for the enterprise. Rather than writing to tape, disk-to-disk, or VTL, a backup server writes its backups to AltaVault.

AltaVault is an inexpensive solution to storing large numbers of backups, without the cost and maintenance of a secondary data center. It is like having a tape library, a vaulting system, a large number of backups or archives in an offsite storage facility, and a secondary data center in one appliance.

About AltaVault Appliance

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 Appliance** - AltaVault appliance is available in the AVA400 model.
- **Virtual Appliance** - AltaVault® virtual appliance is a virtual machine hosted package.

You can use VMware ESX or ESXi servers, or Microsoft® Hyper-V, to create a virtual machine (VM) and install the AltaVault-v software on the VM.

AltaVault-v is available in the following models for both, VMware or Microsoft Hyper-V:

- AVA-v8
- AVA-v16

- AVA-v32
- **Cloud-Based Virtual Appliance** - AltaVault® cloud based appliances include Amazon Machine Images (AMI) and a single Azure Virtual Machine (AVM).

The model numbers for AMI are:

- AVA-c4

Note: The AVA-c4 has been designed and optimized for deployment on the DS3 instance. Therefore, it is strongly recommended to deploy the AVA-c4 with DS3 instance only.

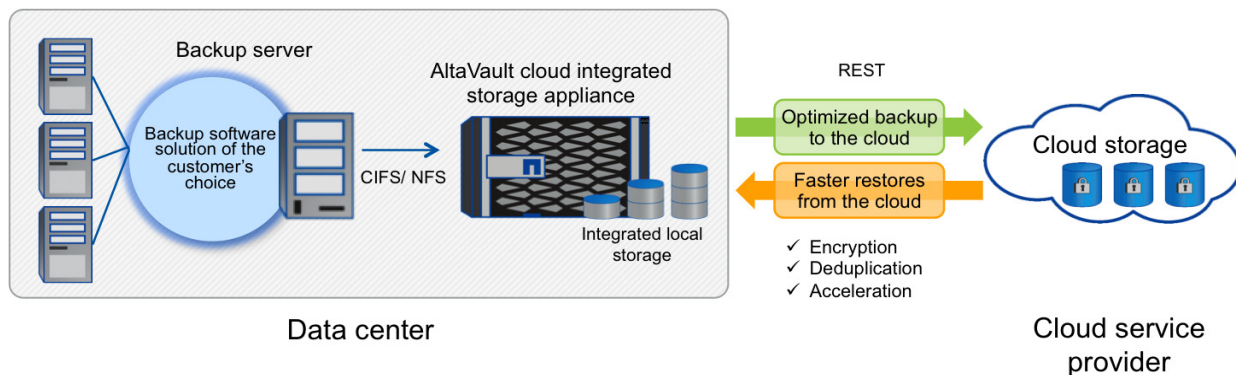
- AVA-c8
- AVA-c16

The model number for AVM is AVA-c4.

Note: Downgrades are not supported in this release.

Figure 1-1 shows an overview of the physical AltaVault deployment.

Figure 1-1. Physical AltaVault Deployment



In Figure 1-1, the application servers, email servers, and file servers connect to the backup server, typically through their backup agents. AltaVault easily integrates into your existing backup infrastructure. You can use your existing backup software, such as Symantec, NetBackup, Symantec Backup Exec, or IBM Spectrum Protect. AltaVault acts as a storage target for your existing infrastructure. AltaVault appears to the backup server as a shared disk, using CIFS (Common Internet File System) or Network File System (NFS) protocols. AltaVault supports CIFS, SMB2, and NFS protocols.

When it is time for a backup, the backup server contacts the backup client. Next, the backup server contacts the backup media (in this case, AltaVault) and starts writing an image of the clients or objects it is backing up.

When you back up to AltaVault, it performs inline (real-time) deduplication of the backup data and replicates data into the cloud. AltaVault uses the local disk to store enough data for recovery of most recent backups. Such a mechanism provides LAN performance for the most likely restores. This deduplication process uses variable segment length inline deduplication plus compression, which is superior to other techniques such as fixed block. AltaVault deduplication level typically ranges between 10 and 30x. Deduplication performance depends on the incoming data type so turn off encryption and compression in the backup applications. Use the native encryption and deduplication in AltaVault to get higher data reduction rates than other typical software products.

AltaVault writes a copy of the data into the cloud storage provider. After AltaVault fills the capacity of its cache, it removes the least recently used data and replaces it with new incoming data. This process is called *eviction*. Evicted data can be recalled from the cloud transparently without user interaction in typical configurations, with most clouds. Amazon Glacier, which is also supported, changes the workflow to be less transparent.

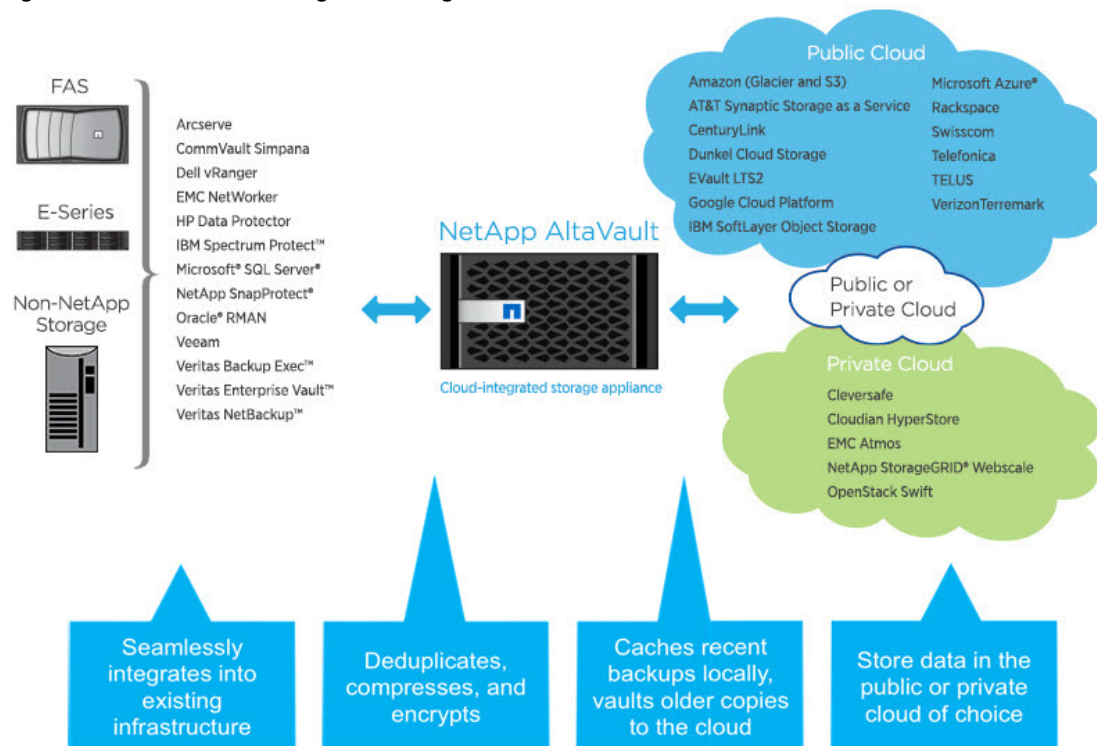
AltaVault also optimizes restores from the cloud because it recalls only deduplicated data (which is not in the local cache) from the cloud. So if the customer is getting 10x deduplication, for example, and he or she needs to restore 10 TB of data, AltaVault needs only about 1 TB to restore. Over a 100-Mb line, this results in a time saving of days.

Data moves from the backup client to the backup server, to AltaVault, and then to the cloud. When you restore data, data moves from the cache in AltaVault, in which it is expanded to its original size to the backup server and to the backup client. If the data is not local, it moves from the cloud to AltaVault, to the backup server, and to the backup client.

Note: Supported Backup Applications: The 7.1.2 release in April'15 IBM TSM data protection solution is re-branded or also known as IBM Spectrum Protect solution.

Figure 1-2 shows the backup applications and cloud providers that AltaVault supports.

Figure 1-2. AltaVault Cloud Integrated Storage



Overview of AltaVault Appliance Deployment

This chapter provides the guidelines for physical and virtual AltaVault deployments. It includes the following section:

- [“AltaVault Deployment Guidelines” on page 9.](#)

AltaVault Deployment Guidelines

Use the following guidelines to deploy a physical AltaVault appliance:

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

Refer to the Interoperability Matrix Tool (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.

- Use the following table to make a comparison of using AltaVault in Backup 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.• Provides the highest ingest performance achieved from backup applications.• 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.	Limits the amount of cloud capacity managed (up to 960TB in the cloud).
Cold Storage Mode	<p>Allows access to far greater cloud capacity (10PB of storage, based on 1.333 billion files of 100MB average file size).</p> <p>Provides expansive long term storage in just one head controller unit.</p> <p>Reduces computing requirements of AltaVault because limited deduplication and compression is performed.</p>	<p>Limits network and WAN performance, dependent on average, file size of objects sent to AltaVault.</p> <p>No expansion capability with shelves.</p> <p>Restores are always from the cloud provider.</p>

- 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 critical folder on one AltaVault data connection, while noncritical backups might be directed by a backup application to point to a non-critical 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 so that generations of the same data arrive at 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, because similar data might be stored in each location.

- If you are choosing to move from one provider to another, you can use the cloud agility feature.

Using a few CLI commands, you designate the new cloud bucket and data is systematically copied from the old provider to the new one. For more information about cloud agility, see [Chapter 6, “Cloud Agility.”](#)

- 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.

- You can deploy each AltaVault to only one cloud storage provider at a time.

If an AltaVault must back up to a different cloud storage provider than the one 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 remains, and you can recover it using AltaVault virtual appliance if necessary.

Deployment Checklist

For guidelines on hardware requirements supported for AltaVault appliance, see the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances*.

Collect the following information to assist you in your deployment as described in this checklist table.

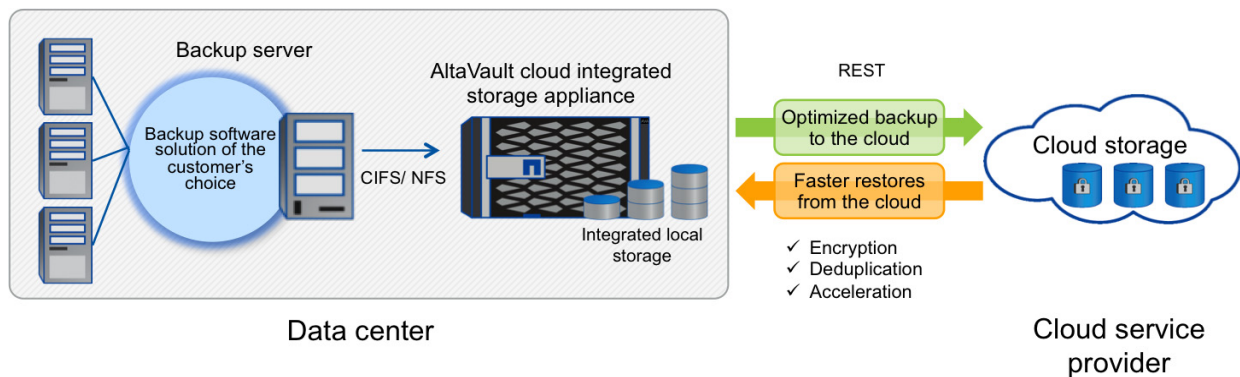
Deployment Information	Required Information that You Must Provide for the Deployment
AltaVault Hostname	Specify the planned hostname.
Primary NIC IP Address	Specify the planned primary NIC IP Address.
Primary NIC Subnet Mask	Specify the planned NIC Subnet Mask.
Primary NIC Gateway	Specify the planned Primary NIC Gateway.
DNS server(s)	Specify the planned DNS server(s).
AltaVault Location (Time Zone)	Specify the planned AltaVault Location (Time Zone).
Domain name	Specify the planned Domain name.
Data Interface NIC IP Address	Specify the planned Data Interface NIC IP Address.

Deployment Information	Required Information that You Must Provide for the Deployment
Data Interface Subnet Mask	Specify the planned Data Interface Subnet Mask.
Data Interface Gateway	Specify the planned Data Interface Gateway.
Cloud Provider	Specify the planned Cloud Provider.
Cloud Provider credentials available (Y/N)	Specify whether Cloud Provider credentials available (Y/N).

Deployment Diagram

Figure 2-1 provides diagram information for the backup application, network, AltaVault, and Cloud Storage Provider for your deployment needs.

Figure 2-1. Physical AltaVault Deployment



Provide deployment information as described in the following table.

Information	Description
Backup Application	Specify the planned Backup Application.
Network	Specify the planned Network.
AltaVault	Specify the planned AltaVault.
Cloud Storage Provider Details	Specify the planned Cloud Storage Provider Details.

Deployment Steps

Use the following table to guide your AltaVault deployment:

Deployment Steps	Reference
Preparing Your Site	<i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances, Chapter 3 - Preparing Your Site.</i>
Installing AltaVault (hardware)	<i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances, Chapter 4 - Installing the AltaVault Controller Chassis.</i>
Connecting the Network (hardware)	<i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances, Chapter 5 - Connecting Networking, Mini-SAS, and Power Cables.</i>

Deployment Steps	Reference
Installing and Configuring an AltaVault physical appliance, AltaVault virtual appliance, or AltaVault cloud appliance	<p>Choose an installation from the following list for your deployment:</p> <ul style="list-style-type: none"> • <i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances.</i> • <i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances, Chapter 1 - Installing and Configuring AltaVault Virtual Appliance on Microsoft Hyper-V.</i> • <i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Virtual Appliances, Chapter 2 - Installing and Configuring AltaVault Virtual Appliance on VMware ESXi.</i> • <i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Cloud Appliances, Chapter 1 - Deploying an Amazon Machine Image.</i> • <i>NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Cloud Appliances, Chapter 2 - Deploying Virtual AltaVault on Microsoft Azure.</i>
Configuring the AltaVault Management Interface	<i>NetApp AltaVault Cloud Integrated Storage Command-Line Interface Reference Manual, Chapter 1 - Configuring the AltaVault Management Interface using the CLI Wizard.</i>
Running the appliance GUI Wizard	<i>NetApp AltaVault Cloud Integrated Storage User's Guide, Chapter 2 Using the AltaVault Configuration Wizard.</i>
Configuring data interface	<i>NetApp AltaVault Cloud Integrated Storage User's Guide, Chapter 4 - Modifying Data Interfaces.</i>
Configure CIFS, SMB2 or NFS shares	<p><i>NetApp AltaVault Cloud Integrated Storage User's Guide, Chapter 3:</i></p> <ul style="list-style-type: none"> • <i>Configuring CIFS</i> • <i>Configuring NFS</i>
Configure email settings	<i>NetApp AltaVault Cloud Integrated Storage User's Guide, Chapter 6 - Configuring Email Settings</i>
Export configuration	<i>NetApp AltaVault Cloud Integrated Storage User's Guide, Chapter 2 - Using the AltaVault Configuration Wizard, see the Using the Export Configuration Wizard section</i>

CHAPTER 3 Disaster Recovery

This chapter describes how to perform disaster recovery using AltaVault. 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:

- “Preparing for Disaster Recovery” on page 13
- “Preparing for Disaster Recovery Testing” on page 15

Preparing for Disaster Recovery

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 AltaVault-v at the recovery site.

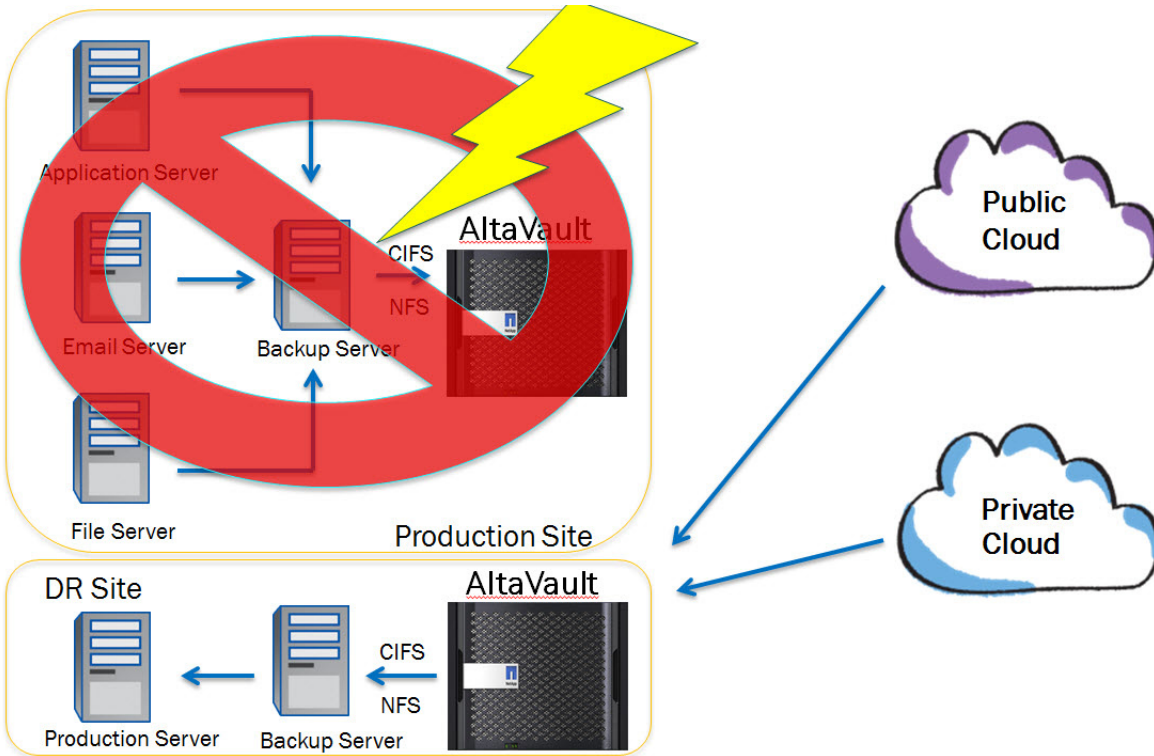
Note: You do not need a license to restore data in read-only mode in AltaVault. You can download AltaVault-v for free from the NetApp Support site at <https://mysupport.netapp.com> and use it to recover your data.

For example, consider a data center with AltaVault located at Site A (shown in [Figure 3-1](#)). The backup site is 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.

For details about AltaVault sizes, see the *NetApp AltaVault Cloud Integrated Storage Installation and Service Guide for Physical Appliances*.

Figure 3-1. Disaster Recovery Process



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.

To export your configuration file

1. Choose Settings > Setup Wizard to display the Wizard Dashboard page.

2. Click **Export Configuration** in the AltaVault wizard dashboard to display the Export Configuration Wizard page.

Figure 3-2. Export Configuration Wizard Page



3. Type the password for the encryption key in the password field. The password field appears only if you specified a password for your encryption key when you generated it in the cloud settings wizard page.
4. Click **Export Configuration** to download the current AltaVault configuration file `AltaVault_config_(HOSTNAME)_(DATETIME).tgz`.
5. Click **Exit** to close the Export Configuration Wizard page and go back to the dashboard.
6. Click **Exit** to close the dashboard.

Preparing for 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.

To disable AltaVault replication (for disaster recovery testing only)

1. Log in to the AltaVault Management Console.

2. Choose Storage > Cloud Settings to display the Cloud Settings page.

Figure 3-3. Cloud Settings Page - Replication

Configure > Storage > Cloud Settings ?

Cloud

Encryption

Replication

Bandwidth

Replication Settings

Pause Replication at: HH:MM:SS

Resume Replication at: HH:MM:SS

Bytes pending replication alert limit: GiB

Bytes pending replication clear limit: GiB

☒ Suspend Replication

Apply

3. Click the Replication tab.
4. Select the **Suspend Replication** check box to pause replication until you resume it again.

Recovering an AltaVault Configuration

The following instructions apply to both DR testing as well as a real DR event.

To recover your configuration to an AltaVault at site B

1. At Site B, plug a serial cable into the console port and a terminal.
2. Configure the AltaVault network information through the serial console:
 - Using telnet, enter the following <terminal server name> <port > to access the AltaVault serial console.
 - If AltaVault appliance is on, the <LOADER> prompt displays.
3. Connect to the AltaVault Management Console.
4. Choose Settings > Setup Wizard to display the Wizard Dashboard page.

5. Click **Import Configuration** in the wizard dashboard to display the Import Configuration page. Import to AltaVault in Site B, the configuration exported from the appliance in Site A. Ensure that the new appliance in Site B uses the same cloud provider credentials, bucket name, and encryption key that Site A uses.

Figure 3-4. Import Configuration Wizard Page



6. Select Local File and click **Choose a File** to select a local configuration file from your computer.
7. Leave the Import Shared Data Only check box selected (by default) to import only the following common settings (the system does not automatically copy the other settings):
 - Cloud settings
 - Email settings
 - Logging
 - NTP settings
 - SNMP settings
 - Statistics or Alarms settings
 - Time zone settings
 - Web and CLI preferences
 - CIFS and NFS 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
8. Select the Password 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.
 9. Click **Import Configuration**.

Caution: After this process completes, the system displays a prompt to restart the storage optimization service. Do not click the restart service button to restart the storage optimization service.

10. Connect to the AltaVault command-line interface using SSH.

11. To perform disaster recovery after a lost primary site, enter the following 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
```

12. To test disaster recovery from a secondary site while the primary site is still alive, 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 data store 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.

Restoring Data for Disaster Recovery

The process for restoring data after a disaster is almost the same as the process used for testing.

To restore data after a disaster

Follow the steps in the section, [“Recovering an AltaVault Configuration” on page 16](#), through Step 11. This completes the process for restoring data after a disaster.

CHAPTER 4 Data Prepopulation

This chapter describes the data population process for AltaVault appliance. It includes the following sections:

- [“Prepopulating Data” on page 21](#)
- [“Using Prepopulation with Amazon Glacier Cloud Storage” on page 28](#)
- [“Automatic Prepopulation” on page 28](#)

Prepopulating Data

You can retrieve the backup data from the cloud and populate AltaVault with it locally so that AltaVault has a local copy of the target data (which improves file access performance) either using the Management Console or using the command-line interface.

NetApp recommends that you use the AltaVault prepopulation process because it is a more efficient way of restoring data from the cloud than using the backup application directly. Although, it might seem longer (because this is a step that occurs before AltaVault restores data through the backup application), the prepopulation process improves restore times. It eliminates sporadic read operations for restore and uses sequential reads, thereby warming the AltaVault cache much more quickly.

To prepopulate data using the AltaVault Management Console

- 1. Choose Storage > Prepopulation to display the Prepopulation page.

Figure 4-1. Prepopulation Page

HOMESTORAGEREPORTSSETTINGSHELP

Prepopulation ⓘ

SaveRestart

Select FileStatus

File Name	Size	Percent Locally Cached
▲ <input type="checkbox"/> cifs		
▶ <input type="checkbox"/> app-setup		
▶ <input type="checkbox"/> foo-dir		
<input type="checkbox"/> scsps0034173001_1429118374_C1_F1.1429118374.img	2.32 GB	
<input type="checkbox"/> scsps0034173001_1429118374_C1_F1.1429118374.info	1.02 kB	
<input type="checkbox"/> scsps0034173001_1429118374_C1_HDR.1429118374.img	512.00 B	
<input type="checkbox"/> scsps0034173001_1429118374_C1_HDR.1429118374.info	1.02 kB	
▲ <input type="checkbox"/> nfs		
▶ <input type="checkbox"/> util		
<input type="checkbox"/> Cisco_WebEx_Add-On.dmg	0.00 B	
<input type="checkbox"/> file1	159.25 kB	
<input type="checkbox"/> foo	0.00 B	
<input type="checkbox"/> fuse-tutorial.tgz	0.00 B	
▶ <input type="checkbox"/> nfsv4		

List of files to be prepopulated

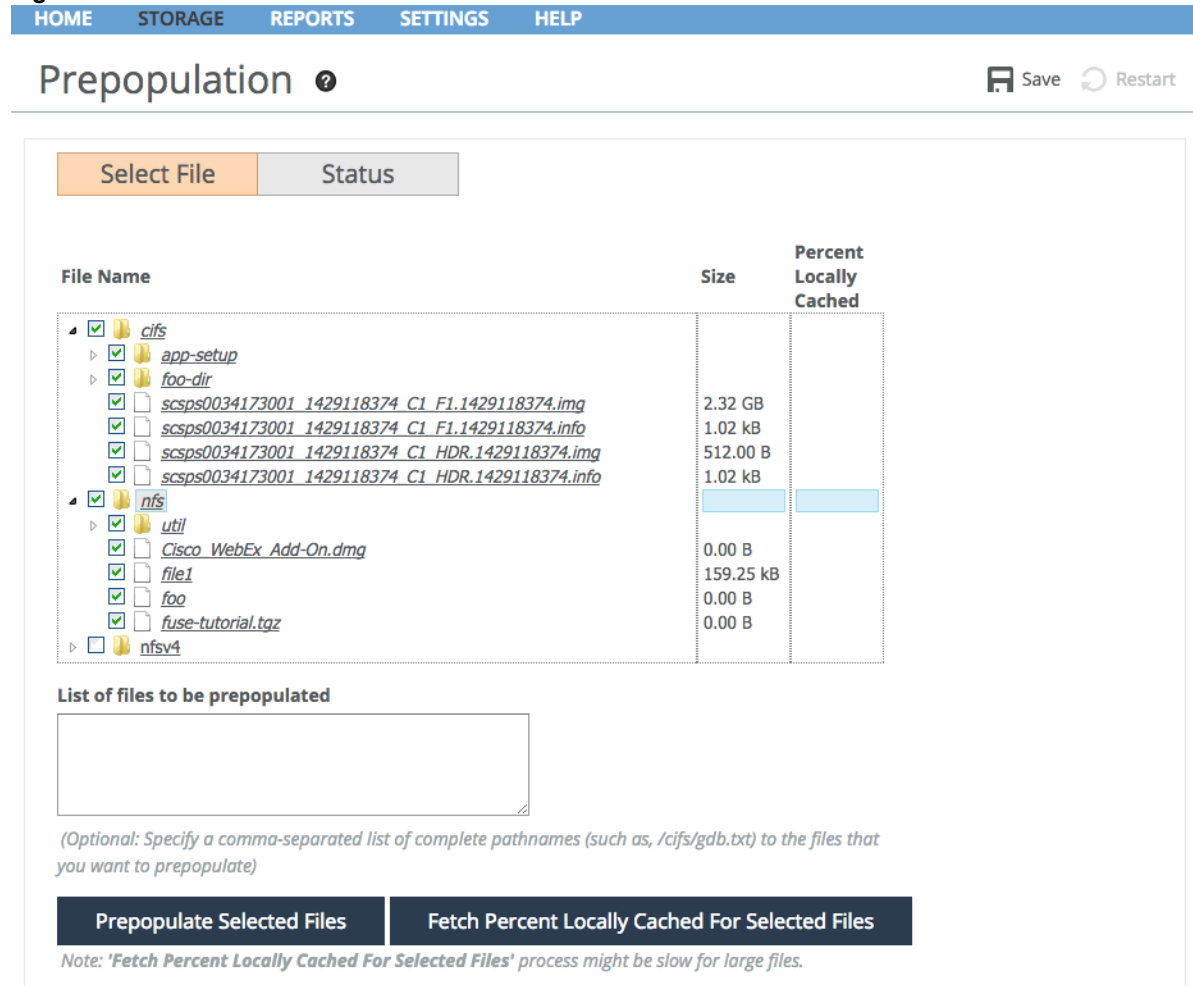
(Optional: Specify a comma-separated list of complete pathnames (such as, /cifs/gdb.txt) to the files that you want to prepopulate)

Prepopulate Selected FilesFetch Percent Locally Cached For Selected Files

Note: 'Fetch Percent Locally Cached For Selected Files' process might be slow for large files.

- Click **Select File** to display the Prepopulation File Browser that contains a list of files that can be prepopulated. You can also display this list by selecting Storage > Prepopulation.

Figure 4-2. Browser List



HOME STORAGE REPORTS SETTINGS HELP

Prepopulation ?

Save Restart

Select File Status

File Name	Size	Percent Locally Cached
<input checked="" type="checkbox"/> cifs ▶ <input checked="" type="checkbox"/> app-setup ▶ <input checked="" type="checkbox"/> foo-dir <input checked="" type="checkbox"/> scsps0034173001_1429118374_C1_F1.1429118374.img <input checked="" type="checkbox"/> scsps0034173001_1429118374_C1_F1.1429118374.info <input checked="" type="checkbox"/> scsps0034173001_1429118374_C1_HDR.1429118374.img <input checked="" type="checkbox"/> scsps0034173001_1429118374_C1_HDR.1429118374.info ▶ <input checked="" type="checkbox"/> nfs ▶ <input checked="" type="checkbox"/> util <input checked="" type="checkbox"/> Cisco_WebEx_Add-On.dmg <input checked="" type="checkbox"/> file1 <input checked="" type="checkbox"/> foo <input checked="" type="checkbox"/> fuse-tutorial.tgz <input type="checkbox"/> nfsv4	2.32 GB 1.02 kB 512.00 B 1.02 kB 0.00 B 159.25 kB 0.00 B 0.00 B	

List of files to be prepopulated

(Optional: Specify a comma-separated list of complete pathnames (such as, /cifs/gdb.txt) to the files that you want to prepopulate)

Prepopulate Selected Files Fetch Percent Locally Cached For Selected Files

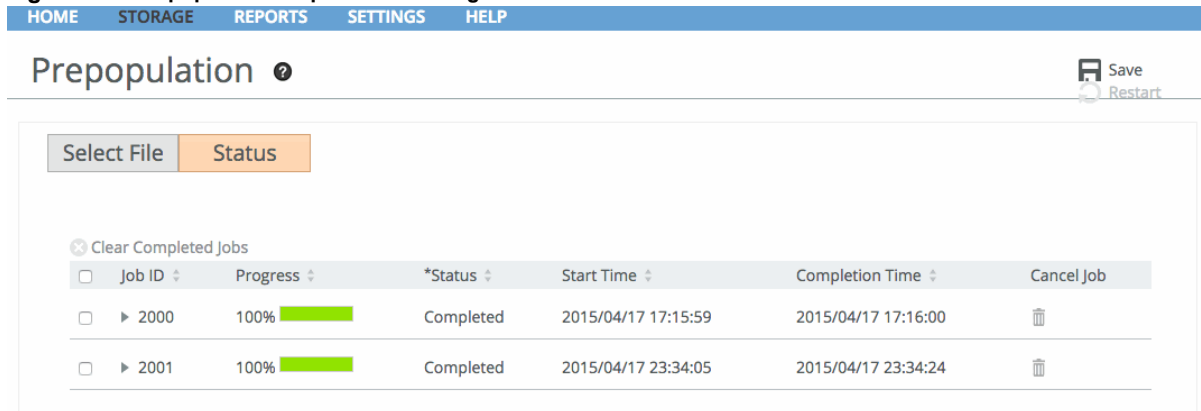
Note: 'Fetch Percent Locally Cached For Selected Files' process might be slow for large files.

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. 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 can be slow for large files.

- Select the check box next to the file you want to prepopulate.

- Click **Prepopulate Selected Files** to prepopulate the files that you selected and display the Prepopulation Report Status page.

Figure 4-3. Prepopulation Report - Status Page



The Prepopulation Report Status page displays the status of the prepopulation task. The following table summarizes the various states.

Status	Description
Enqueued	The prepopulation task has just been recorded. AltaVault has not started processing it. You do not usually see this status (unless there is a large number of prepopulation tasks) because the prepopulation process is very fast and it quickly moves to the next step in the process.
Processing	AltaVault is identifying data that must be restored from the cloud.
Requested	The system has requested all of the data required for the prepopulation request from the cloud.
Downloading	The system has started downloading the data for the prepopulation request. When the cloud provider is Amazon Glacier, it usually takes about five hours for this state to appear.
Completed	This state indicates that the prepopulation task is complete. The completion time also appears in a separate column.
Failed	This state indicates that AltaVault did not restore all of the data and the prepopulation task failed. Check the logs to determine the reason for failure.

- Optionally, click **Clear Completed Jobs** to delete the completed prepopulation tasks (status is Completed).

After a prepopulation job is complete, the system sends an email notification 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 prepopulate data using the command-line interface

1. Connect to the AltaVault command-line interface using SSH.
2. Enter the following command:

```
amnesiac (config) # datastore prepop { [num-days <number-of-days>] | [start-date *] [end-date *]
| [pattern <pattern>] }
```

The following table shows the parameter options.

Parameter	Description
num-days <number-of-days>	Specify 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>	Specify the date from which the data retrieval should start. The system prepopulates the files modified on or before this date.
end-date <end-date>	Specify the date on which the data retrieval should end. Stop prepopulating files on or after this date.
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.

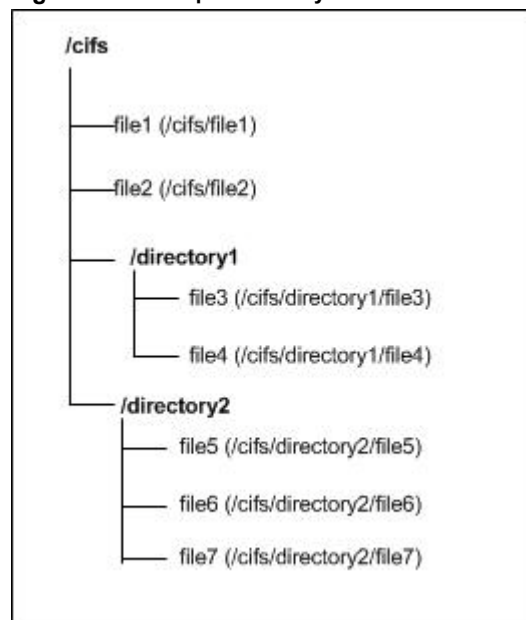
To view the current status of prepopulation, enter the following command:

```
amnesiac (config) # show datastore prepop
```

Example 1 Pattern-based Data Store Prepopulation

This example explains pattern-based data store prepopulation. Consider the directory structure shown in [Figure 4-4](#).

Figure 4-4. Example Directory Structure

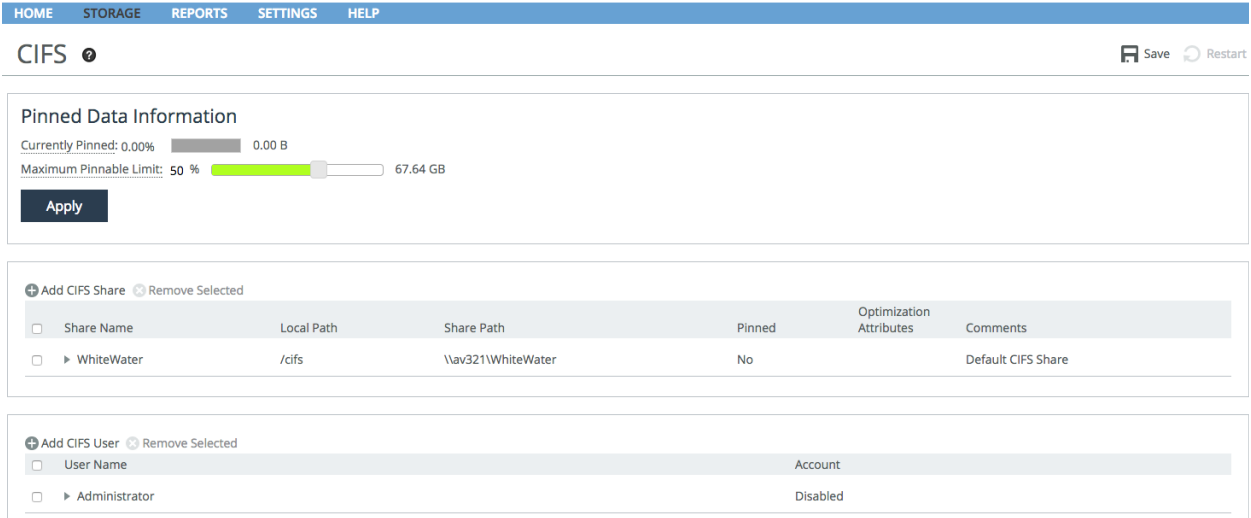


The following table shows different examples of the **datastore prepop** command for this directory structure.

Command	Description
datastore prepop pattern cifs/*	Populates only file1 and file2.
daastore prepop pattern cifs/* recursive	Populates all of the files (file1 through file7) with directory1 and directory2.
datastore prepop pattern cifs/directory1/*	Populates only file3 and file4.

The **datastore prepop** command operates from the local pathname for each CIFS share created as shown in [Figure 4-5](#).

Figure 4-5. Creating CIFS Shares

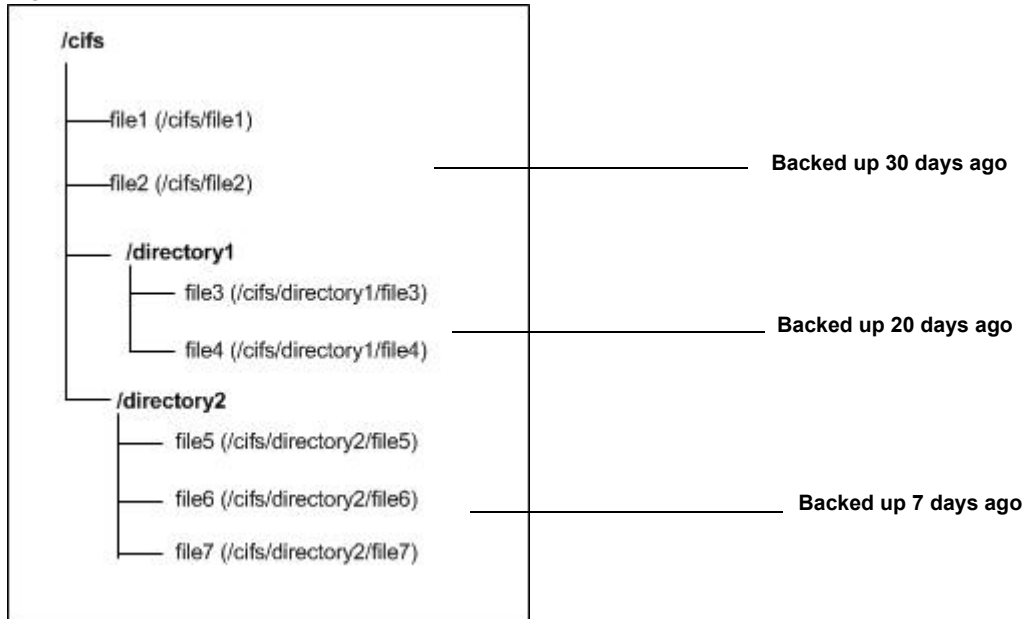


Example 2 Time-based Data Store Prepopulation

This example explains time-based data store prepopulation.

Consider the directory structure shown in [Figure 4-6](#).

Figure 4-6. Example Backup Times



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.

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 for a 24-hour duration starting on 2014-01-01 (YYYY-MM-DD), enter the following command:

```
amnesiac (config) # datastore prepop pattern fulls/hostA/*.img start-date 2014-01-01 end-date 2014-01-02 num-threads 64
```

To prepopulate all backups for Host A that occurred in the past 30 days (from the current time), enter the following command:

```
amnesiac (config) # datastore prepop num-days 30 pattern fulls/hostA/*.img num-threads 64
```

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.

Using Prepopulation with Amazon Glacier Cloud Storage

When you use Amazon Glacier as the cloud storage provider, it takes approximately four to five hours for data to be 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 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.

Automatic Prepopulation

You can also 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.

To enable automatic prepopulation

1. Connect to the AltaVault CLI.
2. Type the following commands:

```
amnesiac(config)# rfsctl exec "-w decoder.file_local_check=true"
amnesiac(config)# rfsctl exec "-w autoprepop.enable=true"
```

Note: The status of the prepopulation job created by the auto-prepopulation functionality can be tracked from the Status tab in the Prepopulation page of the UI. Enabling autoprepop settings can trigger the prepopulation of entire files from the cloud upon read failures, therefore, autoprepop settings must be enabled only after proper consideration.

Time-Based Automatic Prepopulation

When you enable automatic prepopulation, you can also include files that have modification times in a given range relative to the file that is read for prepopulation.

For example, assume that a folder in an AltaVault share contains a file called testfile with a modification time stamp **X**. Assume you enabled automatic prepopulation of the file testfile. If the AltaVault cache does not contain the complete contents of the file testfile, then reading the file triggers a prepopulation job for the file testfile. For details, see [“Automatic Prepopulation” on page 28](#).

Also, if you want to prepopulate all files in the same folder that have modification time stamps in the range (X-delta1, X+delta2), type the following commands on the AltaVault CLI:

```
amnesiac(config)# rfsctl exec "-w autoprepop.time_based.enable=true"
amnesiac(config)# rfsctl exec "-w autoprepop.time_based.pre_delta=<delta1>"
amnesiac(config)# rfsctl exec "-w autoprepop.time_based.post_delta=<delta2>"
```

The units for delta1 and delta2 are in seconds.

CHAPTER 5 **Configuring Remote Management**

This chapter describes how to configure remote management from the Serial Console.

This chapter includes the following sections:

- [“Configuring Remote Management” on page 29](#)
- [“Remote Management Configuration Examples” on page 30](#)

Configuring Remote Management

Access to AltaVault through remote management requires setting up the Service Processor, logging in and entering commands. 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 Telnet access to AltaVault appliance.

To setup the Service Processor

1. Using telnet, enter the following <terminal server name> <port > to access the AltaVault Serial Console.
If AltaVault appliance is on, the <LOADER> prompt displays.

2. Enter the command, `sp setup` to setup the IP addresses for the Service Processor.

3. 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:          172.16.32.1
IPv6:             Disabled

```

- If you do not want to enable DHCP, enter the following from the command line:

```

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]: 172.22.100.4
Please enter the netmask for the SP [unknown]: 255.255.255.0
Please enter the IP address for the SP gateway [unknown]: 172.22.100.1
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:         YES
IP Address:          172.22.100.4
Netmask:             255.255.255.0
Gateway:             172.22.100.1
IPv6:               Disabled

```

4. To verify that the IP addresses are set, enter the command, `sp status`.
5. In a different terminal window, enter the command, `ssh naroot@<IP Address>` to take you to the Service Processor CLI.
6. Enter the autoboot command to restart AltaVault appliance.
7. Set the service processor password using the command:

```
sp password set
```

Remote Management Configuration Examples

The following commands show an example of how to configure the remote management port.

1. Connect to an AltaVault using the command:

```
ssh naroot@172.16.33.155
SP>
```

2. Obtain the IP address 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:             yes
IPv4 configuration:
IP Address:              172.16.33.155
Netmask:                 255.255.252.0
Gateway:                 172.16.32.1
IPv6 configuration: Disabled

```

3. Obtain the IP address of the SP:

```
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:            yes
IPv4 configuration:
  IP Address:           172.16.33.155
  Netmask:              255.255.252.0
  Gateway:              172.16.32.1
IPv6 configuration: Disabled
```

4. Enter the following IPMI command to test the feature:

```
SP> system sensors
```

Sample output:

Sensor Name	Current	Unit	Status	LCR	INC	UNC	UCR
CPU0_Temp_Margin	-66.000	degrees C	ok	na	na	-5.000	0.000
CPU1_Temp_Margin	-68.000	degrees C	ok	na	na	-5.000	0.000
In_Flow_Temp	20.000	degrees C	ok	0.000	10.000	53.000	63.000
Out_Flow_Temp	33.000	degrees C	ok	0.000	10.000	61.000	71.000
Smart_Bat_Temp	30.000	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	1.794	Volts	ok	1.630	1.659	1.950	1.969

5. To login to AltaVault using the system console, enter the following commands.

```
SP> system console
```

6. To exit, type Ctrl-D.

```
AltaVault
amnesiac login: admin
Password:
Last login: Thu Apr 12 13:32:56 from 10.18.5.230
```

7. To terminate the IPMI session, enter the tilde (~) command:

```
SP>
```


CHAPTER 6 Cloud Agility

This chapter describes how to migrate data to a new cloud. It includes the following sections:

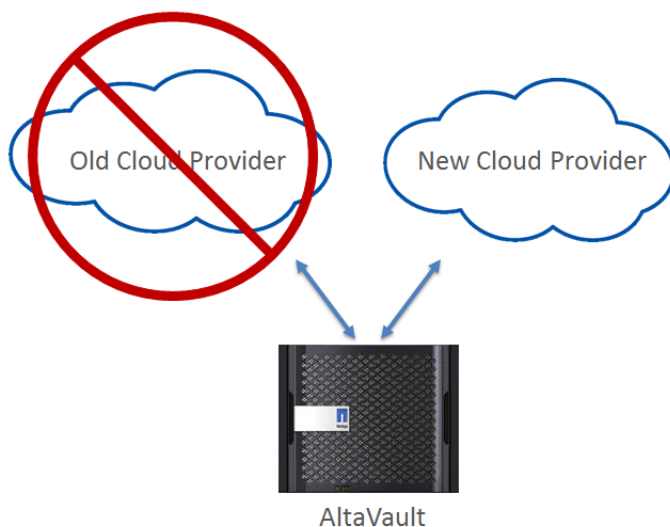
- [“Cloud Agility Overview” on page 33](#)
- [“Migration Process” on page 34](#)

Cloud Agility Overview

AltaVault writes the deduplicated, compressed, and encrypted data to a private or public cloud storage provider. When business requirements dictate that data be migrated to a new cloud (for example, migrating from public cloud to public cloud, private cloud to public cloud, or public cloud to private cloud), the data that resides in that cloud storage must be relocated to a new cloud storage through data migration.

AltaVault implements the cloud migration feature called Cloud Agility to address this requirement. See [“Migration Process” on page 34](#) to see an example of the three commands used for cloud migration. The first two commands provide the cloud credentials and storage target of the new cloud storage service. The third command begins the migration. AltaVault copies all of the data from the first cloud storage to the second cloud storage.

Figure 6-1. Cloud Migration Process



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.

When the data migration process completes, AltaVault automatically updates the cloud storage provider credentials to the new provider and resumes replication of any pending data that was queued during the migration process.

Migration Process

Use the following commands to set up the cloud credentials and perform the migration. The command to set up the authentication type may be different depending on the provider you use.

To set cloud credentials:

```
replication migrate-to provider type <provider-name> bucket-name <bucket-name> hostname <host-name>
port <port-value>
replication migrate-to auth type <authentication-type> acc-key-id <access-key> secret-acc-key
<secret-key>
```

To start cloud migration:

```
replication migrate-to enable [num-threads <value-1-to-128>]
```

To monitor cloud migration:

```
show replication migrate-to estimate
```

Note: Cloud Agility is not supported from Amazon Glacier to any other cloud storage provider.

If migrating from Amazon S3 to Amazon Glacier, use the command:

```
replication s3-to-glacier
```

CHAPTER 7 Monitoring Peer Appliances

This chapter describes how to configure the peer monitoring feature in AltaVault. It includes the following sections:

- [“Benefits of Monitoring Peer Appliances” on page 35](#)
- [“Configuring Appliance Monitoring” on page 35](#)
- [“Configuring REST API Access” on page 36](#)
- [“Specifying the API Access Code” on page 37](#)

Benefits of Monitoring Peer Appliances

The benefits of monitoring peer appliances are:

- Centralizes management
- Improves storage visibility in large or multi-office configurations

Configuring Appliance Monitoring

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

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

To configure appliance monitoring

1. Enable REST-based access on the monitored appliance.
2. Generate the API access code on the monitored appliance.
3. Enter the API access code on the monitoring appliance.

Configuring REST API Access

AltaVault uses REST APIs that you can access to set up peer appliance monitoring.

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

To configure REST API Access

1. Log in to the monitored AltaVault.
2. Choose Settings > REST API Access to display the REST API Access page.

Figure 7-1. REST API Access Page

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

Control	Description
Add Access Code	Displays the controls to add 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 appliance, and a description.
Generate New Access Code	Select to create a new REST API 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 AltaVault.
Remove Selected	Deletes the selected REST API access code.

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

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

Specifying the API Access Code

After you generate the REST API access code on the monitored appliance, you must enter the code in the monitoring appliance to authenticate the monitored appliance.

To specify the API access code in the monitoring appliance

1. Log in to the monitoring appliance.
2. Choose Reports > Appliance Monitoring to display the Appliance Monitoring page.

Figure 7-2. Appliance Monitoring Page

[HOME](#)
[STORAGE](#)
[REPORTS](#)
[SETTINGS](#)
[HELP](#)

Appliance Monitoring ?

Monitored Appliances

▼ Add Monitored Appliance
✕ Remove Selected Appliances

Hostname or IP address:

API Access Code:

eyJhdWQiOiAiaHR0cHM6Ly9hdjMyMS9hcGkvY29tbW9uLzEuMC90b2t1bilsilCjpc3MiOiAiaHR0cHM6Ly9hdjMyMSIsIjwcm4iOiAiYWRTaW4iLCIAianRpljogIjZhNzUzNDhILTBlkZGYtNDdmOC1hY2I5LWRIMzY2MzljMTY2ZCIs

Add

Hostname	Appliance Health	Service State	Expanded Data	Deduplicated Data	Deduplication Factor	Cloud Storage Used	Cloud Storage Available
No AltaVault Appliances added.							

Summary

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:

3. Complete the configuration as described in this table.

Control	Description
Add Monitored Peer	Displays the controls to add a monitored appliance.
Hostname/IP Address	Specify a valid hostname or IP address for the monitored appliance.
API Access Key	Specify the API access code that you obtained from the monitored appliance to access the monitored appliance. To obtain the API access code, see “Configuring REST API Access” on page 36 .
Remove Selected Peers	Select the check box next to the name and click Remove Selected Peers to delete the monitored appliance from the system.

Index

A

AltaVault

- appliance deployments 5
- cloud appliance 6
- definition 5
- overview 5
- physical appliance 5
- physical, deployment guidelines 9

Amazon S3

- migrating to Amazon Glacier 34

AMI 6

- AVA400 5
- AVA-c16 6
- AVA-c4 6
- AVA-c8 6
- AVA-v16 5
- AVA-v32 6
- AVA-v8 5
- AVM 6
- Azure Virtual Machine 6

B

- Backup applications supported 7

C

cloud

- migration 34
- migration process 33
- set credentials 34

Cloud agility 33

cloud migration 33

Cloud providers supported 7

Configuration file, exporting 14

D

data migration 33

Data restoring 15

deployment

- steps 11

Deployment guidelines

- physical AltaVault 9

Disaster recovery

- definition 13
- performing 13
- preparing for 14

Document conventions, overview of 2

downgrades 6

E

ESXi 5

Eviction 7

Exporting, configuration file 14

G

Glacier 34

K

Known issues 2

M

Microsoft Hyper-V 5

Migrate 33

migration

- cloud process 33
- data 33
- monitor cloud 34

models

- virtual 5

O

Online documentation 2

R

Related reading 2

REST API 36

Restoring data 15

S

serial console

- access 16

V

VMware 5

