



Replicate and protect data

Cloud Manager 3.7

NetApp
June 06, 2022

Table of Contents

- Replicate and protect data 1
 - Discovering and managing ONTAP clusters 1
 - Replicating data between systems 2
 - Backing up data to Amazon S3 9
 - Syncing data to Amazon S3 18

Replicate and protect data

Discovering and managing ONTAP clusters

Cloud Manager can discover the ONTAP clusters in your on-premises environment, in a NetApp Private Storage configuration, and in the IBM Cloud. Discovering these clusters enables you to easily replicate data across your hybrid cloud environment directly from Cloud Manager.

Discovering ONTAP clusters

Discovering an ONTAP cluster in Cloud Manager enables you to provision storage and replicate data across your hybrid cloud.

Before you begin

You must have the cluster management IP address and the password for the admin user account to add the cluster to Cloud Manager.

Cloud Manager discovers ONTAP clusters using HTTPS. If you use custom firewall policies, they must meet the following requirements:

- The Cloud Manager host must allow outbound HTTPS access through port 443.

If Cloud Manager is in AWS, all outbound communication is allowed by the predefined security group.

- The ONTAP cluster must allow inbound HTTPS access through port 443.

The default "mgmt" firewall policy allows inbound HTTPS access from all IP addresses. If you modified this default policy, or if you created your own firewall policy, you must associate the HTTPS protocol with that policy and enable access from the Cloud Manager host.

Steps

1. On the Working Environments page, click **Discover** and select **ONTAP Cluster**.
2. On the **ONTAP Cluster Details** page, enter the cluster management IP address, the password for the admin user account, and the location of the cluster.
3. On the Details page, enter a name and description for the working environment, and then click **Go**.

Result

Cloud Manager discovers the cluster. You can now create volumes, replicate data to and from the cluster, and launch OnCommand System Manager to perform advanced tasks.

Provisioning volumes on ONTAP clusters

Cloud Manager enables you to provision NFS and CIFS volumes on ONTAP clusters.

Before you begin

NFS or CIFS must be set up on the cluster. You can set up NFS and CIFS using System Manager or the CLI.

About this task

You can create volumes on existing aggregates. You cannot create new aggregates from Cloud Manager.

Steps

1. On the Working Environments page, double-click the name of the ONTAP cluster on which you want to provision volumes.
2. Click **Add New Volume**.
3. On the Create New Volume page, enter details for the volume, and then click **Create**.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Usage Profile	Usage profiles define the NetApp storage efficiency features that are enabled for a volume.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.

Replicating data between systems

You can replicate data between working environments by choosing a one-time data replication for data transfer, or a recurring schedule for disaster recovery or long-term retention. For example, you can set up data replication from an on-prem ONTAP system to Cloud Volumes ONTAP for disaster recovery.

Cloud Manager simplifies data replication between volumes on separate systems using SnapMirror and SnapVault technologies. You simply need to identify the source volume and the destination volume, and then choose a replication policy and schedule. Cloud Manager purchases the required disks, configures relationships, applies the replication policy, and then initiates the baseline transfer between volumes.



The baseline transfer includes a full copy of the source data. Subsequent transfers contain differential copies of the source data.

Data replication requirements

Before you can replicate data, you should confirm that specific requirements are met for both Cloud Volumes ONTAP systems and ONTAP clusters.

Version requirements

You should verify that the source and destination volumes are running compatible ONTAP versions before replicating data. For details, see the [Data Protection Power Guide](#).

Requirements specific to Cloud Volumes ONTAP

- The instance's security group must include the required inbound and outbound rules: specifically, rules for ICMP and ports 10000, 11104, and 11105.

These rules are included in the predefined security group.

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).
- To replicate data between a Cloud Volumes ONTAP system in AWS and a system in Azure, you must have a VPN connection between the AWS VPC and the Azure VNet.

Requirements specific to ONTAP clusters

- An active SnapMirror license must be installed.
- If the cluster is on your premises, you should have a connection from your corporate network to AWS or Azure, which is typically a VPN connection.
- ONTAP clusters must meet additional subnet, port, firewall, and cluster requirements.

For details, see the Cluster and SVM Peering Express Guide for your version of ONTAP.

Setting up data replication between systems

You can replicate data between Cloud Volumes ONTAP systems and ONTAP clusters by choosing a one-time data replication, which can help you move data to and from the cloud, or a recurring schedule, which can help with disaster recovery or long-term retention.

About this task

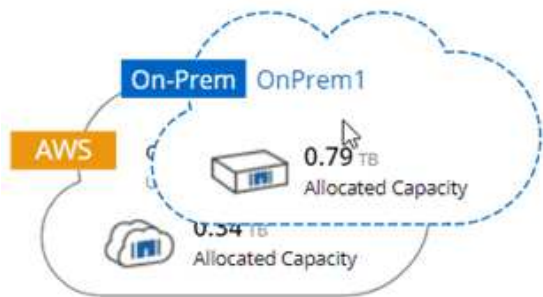
Cloud Manager supports simple, fanout, and cascade data protection configurations:

- In a simple configuration, replication occurs from volume A to volume B.
- In a fanout configuration, replication occurs from volume A to multiple destinations.
- In a cascade configuration, replication occurs from volume A to volume B and from volume B to volume C.

You can configure fanout and cascade configurations in Cloud Manager by setting up multiple data replications between systems. For example, by replicating a volume from system A to system B and then by replicating the same volume from system B to system C.

Steps

1. On the Working Environments page, select the working environment that contains the source volume, and then drag it to the working environment to which you want to replicate the volume:



2. If the Source and Destination Peering Setup pages appear, select all of the intercluster LIFs for the cluster peer relationship.

The intercluster network should be configured so that cluster peers have *pair-wise full-mesh connectivity*, which means that each pair of clusters in a cluster peer relationship has connectivity among all of their intercluster LIFs.

These pages appear if an ONTAP cluster that has multiple LIFs is the source or destination.

3. On the Source Volume Selection page, select the volume that you want to replicate.
4. On the Destination Volume Name and Tiering page, specify the destination volume name, choose an underlying disk type, change any of the advanced options, and then click **Continue**.

If the destination is an ONTAP cluster, you must also specify the destination SVM and aggregate.

5. On the Max Transfer Rate page, specify the maximum rate (in megabytes per second) at which data can be transferred.
6. On the Replication Policy page, choose one of the default policies or click **Additional Policies**, and then select one of the advanced policies.

For help, see [Choosing a replication policy](#).

If you choose a custom backup (SnapVault) policy, the labels associated with the policy must match the labels of the Snapshot copies on the source volume. For more information, see [How backup policies work](#).

7. On the Schedule page, choose a one-time copy or a recurring schedule.

Several default schedules are available. If you want a different schedule, you must create a new schedule on the *destination* cluster using System Manager.

8. On the Review page, review your selections, and then click **Go**.

Result

Cloud Manager starts the data replication process. You can view details about the replication in the Replication Status page.

Managing data replication schedules and relationships

After you set up data replication between two systems, you can manage the data replication schedule and relationship from Cloud Manager.

Steps

1. On the Working Environments page, view the replication status for all working environments in the

workspace or for a specific working environment:

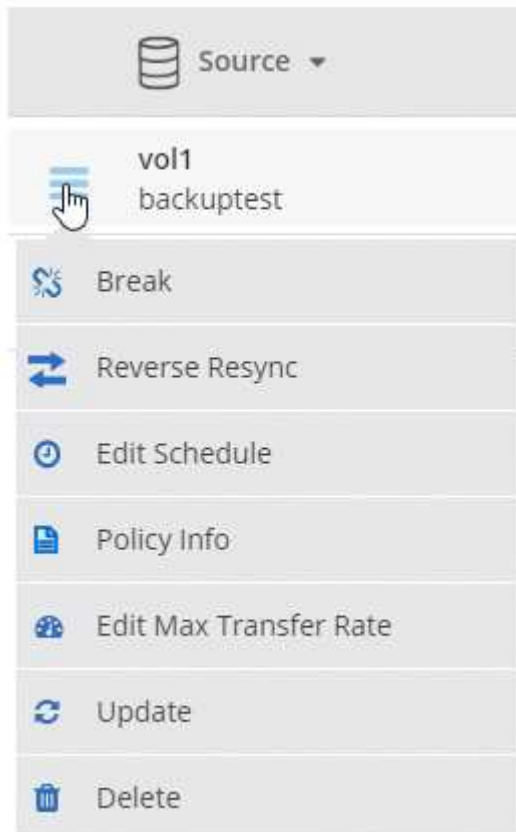
Option	Action
All working environments in the workspace	At the top of Cloud Manager, click Replication Status .
A specific working environment	Open the working environment and click Replications .

2. Review the status of the data replication relationships to verify that they are healthy.




If the Status of a relationship is idle and the Mirror State is uninitialized, you must initialize the relationship from the destination system for the data replication to occur according to the defined schedule. You can initialize the relationship by using System Manager or the command-line interface (CLI). These states can appear when the destination system fails and then comes back online.

3. Select the menu icon next to the source volume, and then choose one of the available actions.



The following table describes the available actions:

Action	Description
Break	<p>Breaks the relationship between the source and destination volumes, and activates the destination volume for data access.</p> <p>This option is typically used when the source volume cannot serve data due to events such as data corruption, accidental deletion, or an offline state.</p> <p>For information about configuring a destination volume for data access and reactivating a source volume, see the ONTAP 9 Volume Disaster Recovery Express Guide.</p>
Resync	<p>Reestablishes a broken relationship between volumes and resumes data replication according to the defined schedule.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  <p>When you resynchronize the volumes, the contents on the destination volume are overwritten by the contents on the source volume.</p> </div> <p>To perform a reverse resync, which resynchronizes the data from the destination volume to the source volume, see the ONTAP 9 Volume Disaster Recovery Express Guide.</p>
Reverse Resync	<p>Reverses the roles of the source and destination volumes. Contents from the original source volume are overwritten by contents of the destination volume. This is helpful when you want to reactivate a source volume that went offline.</p> <p>Any data written to the original source volume between the last data replication and the time that the source volume was disabled is not preserved.</p>
Edit Schedule	Enables you to choose a different schedule for data replication.
Policy Info	Shows you the protection policy assigned to the data replication relationship.
Edit Max Transfer Rate	Enables you to edit the maximum rate (in kilobytes per second) at which data can be transferred.
Update	Starts an incremental transfer to update the destination volume.
Delete	Deletes the data protection relationship between the source and destination volumes, which means that data replication no longer occurs between the volumes. This action does not activate the destination volume for data access. This action also deletes the cluster peer relationship and the storage virtual machine (SVM) peer relationship, if there are no other data protection relationships between the systems.

Result

After you select an action, Cloud Manager updates the relationship or schedule.

Choosing a replication policy

You might need help choosing a replication policy when you set up data replication in Cloud Manager. A replication policy defines how the storage system replicates data from a source volume to a destination volume.

What replication policies do

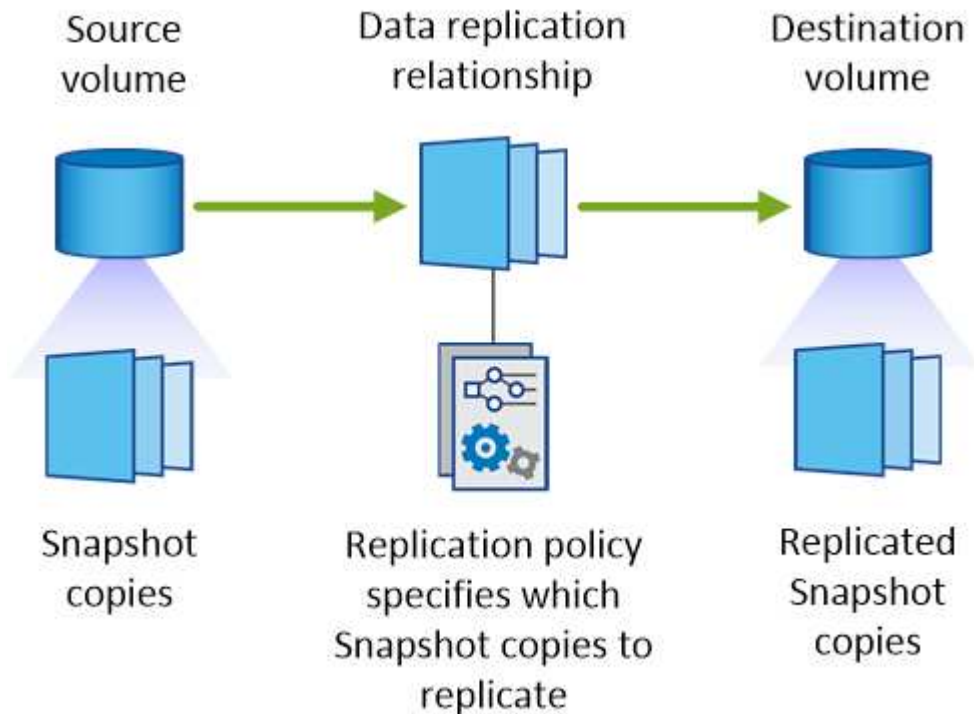
The ONTAP operating system automatically creates backups called Snapshot copies. A Snapshot copy is a read-only image of a volume that captures the state of the file system at a point in time.

When you replicate data between systems, you replicate Snapshot copies from a source volume to a destination volume. A replication policy specifies which Snapshot copies to replicate from the source volume to the destination volume.



Replication policies are also referred to as *protection* policies because they are powered by SnapMirror and SnapVault technologies, which provide disaster recovery protection and disk-to-disk backup and recovery.

The following image shows the relationship between Snapshot copies and replication policies:



Types of replication policies

There are three types of replication policies:

- A *Mirror* policy replicates newly created Snapshot copies to a destination volume.

You can use these Snapshot copies to protect the source volume in preparation for disaster recovery or for one-time data replication. You can activate the destination volume for data access at any time.


- A *Backup* policy replicates specific Snapshot copies to a destination volume and typically retains them for a longer period of time than you would on the source volume.

You can restore data from these Snapshot copies when data is corrupted or lost, and retain them for standards compliance and other governance-related purposes.


- A *Mirror and Backup* policy provides both disaster recovery and long-term retention.

Each system includes a default Mirror and Backup policy, which works well for many situations. If you find that you need custom policies, you can create your own using System Manager.

The following images show the difference between the Mirror and Backup policies. A Mirror policy mirrors the Snapshot copies available on the source volume.



A Backup policy typically retains Snapshot copies longer than they are retained on the source volume:



How Backup policies work

Unlike Mirror policies, Backup (SnapVault) policies replicate specific Snapshot copies to a destination volume. It is important to understand how Backup policies work if you want to use your own policies instead of the default policies.

Understanding the relationship between Snapshot copy labels and Backup policies

A Snapshot policy defines how the system creates Snapshot copies of volumes. The policy specifies when to create the Snapshot copies, how many copies to retain, and how to label them. For example, a system might create one Snapshot copy every day at 12:10 a.m., retain the two most recent copies, and label them "daily".

A Backup policy includes rules that specify which labeled Snapshot copies to replicate to a destination volume and how many copies to retain. The labels defined in a Backup policy must match one or more labels defined in a Snapshot policy. Otherwise, the system cannot replicate any Snapshot copies.

For example, a Backup policy that includes the labels "daily" and "weekly" results in replication of Snapshot

copies that include only those labels. No other Snapshot copies are replicated, as shown in the following image:

Default policies and custom policies

The default Snapshot policy creates hourly, daily, and weekly Snapshot copies, retaining six hourly, two daily, and two weekly Snapshot copies.

You can easily use a default Backup policy with the default Snapshot policy. The default Backup policies replicate daily and weekly Snapshot copies, retaining seven daily and 52 weekly Snapshot copies.

If you create custom policies, the labels defined by those policies must match. You can create custom policies using System Manager.

Backing up data to Amazon S3

Backup to S3 is an add-on feature for Cloud Volumes ONTAP that delivers fully-managed backup and restore capabilities for protection, and long-term archive of your cloud data. Backups are stored in S3 object storage, independent of volume Snapshot copies used for near-term recovery or cloning.

When you enable Backup to S3, the service performs a full backup of your data. All additional backups are incremental, which means that only changed blocks and new blocks are backed up.

[Go to NetApp Cloud Central for pricing details.](#)

Note that you must use Cloud Manager for all backup and restore operations. Any actions taken directly from ONTAP or from Amazon S3 results in an unsupported configuration.

Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



Verify support for your configuration

Verify the following:

- Cloud Volumes ONTAP 9.4 or later is running in a supported AWS region: N. Virginia, Oregon, Ireland, Frankfurt, or Sydney
- You have subscribed to the new [Cloud Manager Marketplace offering](#)
- TCP port 5010 is open for outbound traffic on the security group for Cloud Volumes ONTAP (it's open by default)
- TCP port 8088 is open for outbound traffic on the security group for Cloud Manager (it's open by default)
- The following endpoint is accessible from Cloud Manager:

`https://w86yt021u5.execute-api.us-east-1.amazonaws.com/production/whitelist`

- There's room for Cloud Manager to allocate up to two interface VPC endpoints in the VPC (the AWS limit

per VPC is 20)


- Cloud Manager has permission to use the VPC endpoint permissions listed in the latest [Cloud Manager policy](#):

```
"ec2:DescribeVpcEndpoints",  
"ec2:CreateVpcEndpoint",  
"ec2:ModifyVpcEndpoint",  
"ec2>DeleteVpcEndpoints"
```

2

Enable Backup to S3 on your new or existing system

- New systems: The Backup to S3 feature is enabled by default in the working environment wizard. Be sure to keep the option enabled.
- Existing systems: Open the working environment, click the backup settings icon and enable backups.



3

If needed, modify the backup policy

The default policy backs up volumes every day and retains 30 backup copies of each volume. If needed, you can change the number of backup copies to retain.

Backup to S3

Backup Working Environment Automatically back up all volumes

Policy - Retention & Schedule

Backup every: Number of backups to retain:

4

Restore your data, as needed

At the top of Cloud Manager, click **Backup & Restore**, select a volume, select a backup, and then restore data from the backup to a new volume.

vol1

Select the backup you want to restore

Aug 21, 2019 05:01:34 PM UTC



Requirements

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to S3.

Supported ONTAP versions

Backup to S3 is supported with Cloud Volume ONTAP 9.4 and later.

Supported AWS regions

Backup to S3 is supported with Cloud Volumes ONTAP in the following AWS regions:

- US East (N. Virginia)
- US West (Oregon)
- EU (Ireland)
- EU (Frankfurt)
- Asia Pacific (Sydney)

AWS permissions required

The IAM role that provides Cloud Manager with permissions must include the following:

```
"ec2:DescribeVpcEndpoints",  
"ec2:CreateVpcEndpoint",  
"ec2:ModifyVpcEndpoint",  
"ec2>DeleteVpcEndpoints"
```

AWS subscription requirement

Starting with the 3.7.3 release, a new Cloud Manager subscription is available in the AWS Marketplace. This subscription enables deployments of Cloud Volumes ONTAP 9.6 and later PAYGO systems and the Backup to S3 feature. You need to [subscribe to this new Cloud Manager subscription](#) before you enable Backup to S3. Billing for the Backup to S3 feature is done through this subscription.

Port requirements

- TCP port 5010 must be open for outbound traffic from Cloud Volumes ONTAP to the backup service.
- TCP port 8088 must be open for outbound traffic on the security group for Cloud Manager.

These ports are already open if you used the predefined security groups. But if you used your own, then you'll need to open these ports.

Outbound internet access

Ensure that the following endpoint is accessible from Cloud Manager:

<https://w86yt021u5.execute-api.us-east-1.amazonaws.com/production/whitelist>

Cloud Manager contacts this endpoint to add your AWS account ID to the list of allowed users for Backup to S3.

Interface VPC endpoints

When you enable the Backup to S3 feature, Cloud Manager creates an interface VPC endpoint in the VPC where Cloud Volumes ONTAP is running. This *backup endpoint* connects to the NetApp VPC where Backup to S3 is running. If you restore a volume, Cloud Manager creates an additional interface VPC endpoint—the *restore endpoint*.

Any additional Cloud Volumes ONTAP systems in the VPC use these two VPC endpoints.

[The default limit for interface VPC endpoints is 20 per VPC](#). Make sure that your VPC hasn't reached the limit before you enable the feature.


Enabling backups to S3 on a new system

The Backup to S3 feature is enabled by default in the working environment wizard. Be sure to keep the option enabled.

Steps

1. Click **Create Cloud Volumes ONTAP**.
2. Select Amazon Web Services as the cloud provider and then choose a single node or HA system.
3. Fill out the Details & Credentials page.

4. On the Backup to S3 page, leave the feature enabled and click **Continue**.



5. Complete the pages in the wizard to deploy the system.

Result


The Backup to S3 feature is enabled on the system and backs up volumes every day and retains 30 backup copies. [Learn how to modify backup retention](#).

Enabling backups to S3 on an existing system

You can enable backups to S3 on an existing Cloud Volumes ONTAP system, as long as you are running a supported configuration. For details, see [Requirements](#).

Steps

1. Open the working environment.
2. Click the backup settings icon.



3. Select **Automatically back up all volumes**.
4. Choose your backup retention and then click **Save**.

Backup to S3

Backup Working Environment Automatically back up all volumes

Policy - Retention & Schedule

Backup every: Number of backups to retain:

Result

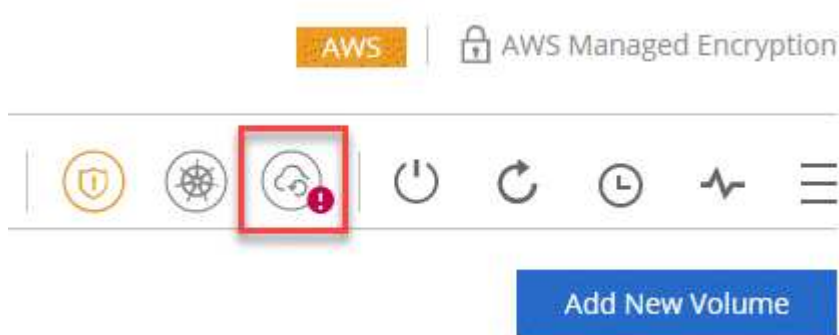
The Backup to S3 feature starts taking the initial backups of each volume.

Changing the backup retention

The default policy backs up volumes every day and retains 30 backup copies of each volume. You can change the number of backup copies to retain.

Steps

1. Open the working environment.
2. Click the backup settings icon.



3. Change the backup retention and then click **Save**.

Backup to S3

Backup Working Environment Automatically back up all volumes

Policy - Retention & Schedule

Backup every: Number of backups to retain:

Restoring a volume

When you restore data from a backup, Cloud Manager performs a full volume restore to a *new* volume. You can restore the data to the same working environment or to a different working environment.

Steps

1. At the top of Cloud Manager, click **Backup & Restore**.
2. Select the volume that you want to restore.

Working Environment	Source Volume	Last Backup	Policy	Retention	Relationship Status	
BackupandRestore (On)	vol1 (Available)	Aug 21, 2019 05:01:34 PM U...	Daily	30	Active (Idle)	View Backup List

3. Find the backup that you want to restore from and click the restore icon.

vol1

Select the backup you want to restore

Aug 21, 2019 05:01:34 PM UTC



4. Select the working environment to which you want to restore the volume.
5. Enter a name for the volume.
6. Click **Restore**.

Deleting backups


All backups are retained in S3 until you delete them from Cloud Manager. Backups are not deleted when you delete a volume or when you delete the Cloud Volumes ONTAP system.

Steps

1. At the top of Cloud Manager, click **Backup & Restore**.
2. Select a volume.
3. Find the backup that you want to delete and click the delete icon.

vol1

Select the backup you want to restore




4. Confirm that you want to delete the backup.

Disabling backups to S3

Disabling backups to S3 disables backups of each volume on the system. Any existing backups will not be deleted.

Steps

1. Open the working environment.
2. Click the backup settings icon.



3. Disable **Automatically back up all volumes** and then click **Save**.

How Backup to S3 works

The following sections provide more information about the Backup to S3 feature.

Where backups reside

Backup copies are stored in a NetApp-owned S3 bucket, in the same region where the Cloud Volumes ONTAP system is located.

Backups are incremental

After the initial full backup of your data, all additional backups are incremental, which means that only changed blocks and new blocks are backed up.

Backups are taken at midnight

Daily backups start just after midnight each day. At this time, you can't schedule backup operations at a user specified time.

Backup copies are associated with your Cloud Central account

Backup copies are associated with the [Cloud Central account](#) in which Cloud Manager resides.

If you have multiple Cloud Manager systems in the same Cloud Central account, each Cloud Manager system will display the same list of backups. That includes the backups associated with Cloud Volumes ONTAP instances from other Cloud Manager systems.

The backup policy is system wide

The number of backups to retain are defined at the system level. You can't set a different policy for each volume on the system.

Security

Backup data is secured with AES-256 bit encryption at-rest and TLS 1.2 HTTPS connections in-flight.

Data travels across secured Direct Connect links to the service, and is protected at rest by AES 256-bit encryption. The encrypted data is then written to cloud using HTTPS TLS 1.2 connections. Data also travels to Amazon S3 only through secure VPC endpoint connections, so no traffic is sent across the internet.

Each user is assigned a tenant key, in addition to an overall encryption key owned by the service. This requirement is similar to needing a pair of keys to open a customer safe in a bank. All keys, as cloud credentials, are stored securely by the service and are restricted to only certain NetApp personnel responsible for maintaining the service.

Limitations

- If you use any of the following instance types, a Cloud Volumes ONTAP system can back up a maximum of 20 volumes to S3:
 - m4.xlarge
 - m5.xlarge
 - r4.xlarge
 - r5.xlarge
- Volumes that you create outside of Cloud Manager aren't automatically backed up to S3.

For example, if you create a volume from the ONTAP CLI, ONTAP API, or System Manager, then the

volume won't be automatically backed up.

If you want to backup these volumes, you would need to disable Backup to S3 and then enable it again.

- When you restore data from a backup, Cloud Manager performs a full volume restore to a *new* volume. This new volume isn't automatically backed up to S3.

If you want to backup volumes created from a restore operation, you would need to disable Backup to S3 and then enable it again.

- You can back up volumes that are 50 TB in size or less.
- Backup to S3 can maintain up to 245 total backups of a volume.
- WORM storage is not supported on a Cloud Volumes ONTAP system when backup to S3 is enabled.


Syncing data to Amazon S3

You can sync data from ONTAP volumes to an Amazon S3 bucket by integrating a working environment with [NetApp Cloud Sync](#). You can then use the synced data as a secondary copy or for data processing using AWS services like EMR and Redshift.

How the sync to S3 feature works

You can integrate a working environment with the Cloud Sync service at any time. When you integrate a working environment, the Cloud Sync service syncs data from the selected volumes to a single S3 bucket. The integration works with Cloud Volumes ONTAP working environments, as well as ONTAP clusters that are on-premises or part of a NetApp Private Storage (NPS) configuration.

To sync the data, the service launches a data broker instance in your VPC. Cloud Sync uses one data broker per working environment to sync data from volumes to an S3 bucket. After the initial sync, the service syncs any changed data once per day at midnight.



If you want to perform advanced Cloud Sync actions, go directly to the Cloud Sync service. From there, you can perform actions such as syncing from S3 to an NFS server, choosing different S3 buckets for volumes, and modifying schedules.

14-day free trial

If you are a new Cloud Sync user, your first 14 days are free. After the free trial ends, you must pay for each *sync relationship* at an hourly rate or by purchasing licenses. Each volume that you sync to an S3 bucket is considered a sync relationship. You can set up both payment options directly from Cloud Sync in the License

Settings page.

How to get help

Use the following options for any support related to the Cloud Manager sync to S3 feature or for Cloud Sync in general:

- General product feedback: ng-cloudsync-contact@netapp.com
- Technical Support options:
 - NetApp Cloud Sync Communities
 - In-product chat (lower-right corner of Cloud Manager)

Integrating a working environment with the Cloud Sync service

If you want to sync volumes to Amazon S3 directly from Cloud Manager, then you must integrate the working environment with the Cloud Sync service.

 | https://img.youtube.com/vi/3hOtLs70_xE/maxresdefault.jpg

Steps

1. Open a working environment and click **Sync to S3**.
2. Click **Sync** and follow the prompts to sync your data to S3.



You cannot sync data protection volumes to S3. The volumes must be writable.

Managing volume sync relationships

After you integrate a working environment with the Cloud Sync service, you can sync additional volumes, stop syncing a volume, and remove the integration with Cloud Sync.

Steps

1. On the Working Environments page, double-click the working environment on which you want to manage sync relationships.
2. If you want to enable or disable sync to S3 for a volume, select the volume and then click **Sync to S3** or **Delete Sync Relationship**.
3. If you want to delete all sync relationships for a working environment, click the **Sync to S3** tab and then click **Delete Sync**.

This action does not delete synced data from the S3 bucket. If the data broker is not being used in any other sync relationships, then the Cloud Sync service deletes the data broker.

Copyright Information

Copyright © 2022 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, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.