



Volume and LUN administration

Cloud Volumes ONTAP

NetApp
February 27, 2026

Table of Contents

- Volume and LUN administration 1
 - Create a FlexVol volume on a Cloud Volumes ONTAP system 1
 - Before you begin 1
 - Create a volume 1
 - Create a volume on the second node in an HA configuration 6
 - After you create a volume 6
 - Manage volumes on Cloud Volumes ONTAP systems 7
 - Manage volumes 7
 - Resize a volume 9
 - Modify the CIFS server 10
 - Move a volume 11
 - Move a volume when Console displays an Action Required message 12
 - Reasons why a volume move might perform slowly 15
 - View FlexGroup Volumes 15
- Tier inactive Cloud Volumes ONTAP data to a low-cost object storage 16
 - Configurations that support data tiering 17
 - Requirements 17
 - Ensure that tiering is enabled on aggregates 18
 - Tier data from read-write volumes 19
 - Tier data from data protection volumes 20
 - Change the storage class for tiered data 21
 - Change the free space ratio for data tiering 21
 - Change the cooling period for the auto tiering policy 23
 - Remove an S3 bucket on decommissioning a system 23
- Connect to a LUN on Cloud Volumes ONTAP from your host system 23
- Accelerate data access with FlexCache volumes on a Cloud Volumes ONTAP system 24
 - Work with FlexCache when the origin is encrypted 25

Volume and LUN administration

Create a FlexVol volume on a Cloud Volumes ONTAP system

If you need more storage after you launch your initial Cloud Volumes ONTAP system, you can create new FlexVol volumes for NFS, CIFS, or iSCSI from the NetApp Console.

You have several ways to create a new volume:

- Specify details for a new volume and let the Console handle the underlying data aggregates for you. [Learn more](#)
- Create a volume on a data aggregate of your choice. [Learn more](#)
- Create a volume on the second node in an HA configuration. [Learn more](#)

Before you begin

A few notes about volume provisioning:

- When you create an iSCSI volume, the Console automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, [use the IQN to connect to the LUN from your hosts](#).
- You can create additional LUNs from ONTAP System Manager or the ONTAP CLI.
- If you want to use CIFS in AWS, you must have set up DNS and Active Directory. For details, refer to [Networking requirements for Cloud Volumes ONTAP for AWS](#).
- If your Cloud Volumes ONTAP configuration supports the Amazon EBS Elastic Volumes feature, you might want to [learn more about what happens when you create a volume](#).

Create a volume

The most common way to create a volume is to specify the type of volume that you need and then let the Console handle the disk allocation for you. But you also have the option to choose the specific aggregate on which you want to create the volume.

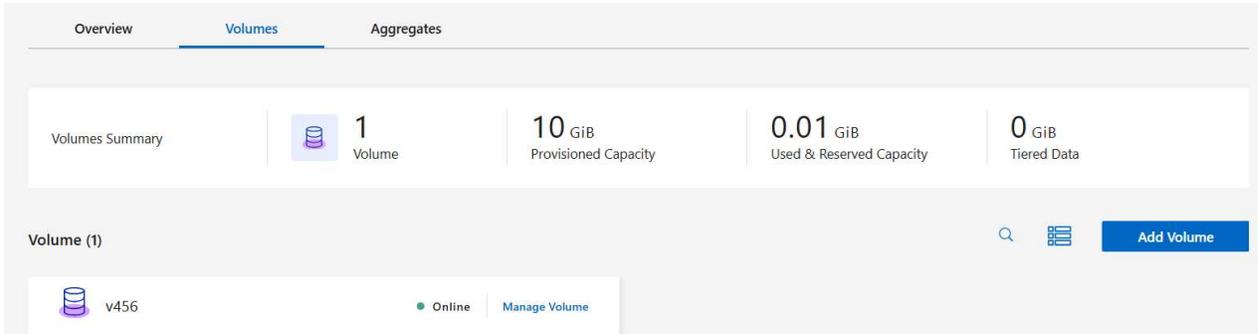
Steps

1. From the left navigation menu, select **Storage > Management**.
2. On the **Systems** page, double-click the name of the Cloud Volumes ONTAP system on which you want to provision a FlexVol volume.

You can create a volume by letting the Console handle the disk allocation for you, or choose a specific aggregate for the volume. Choosing a specific aggregate is recommended only if you have a good understanding of the data aggregates on your Cloud Volumes ONTAP system.

Any aggregate

Select the **Volumes** tab, and click **Add Volume**.



Overview **Volumes** Aggregates

Volumes Summary

1 Volume

10 GiB Provisioned Capacity

0.01 GiB Used & Reserved Capacity

0 GiB Tiered Data

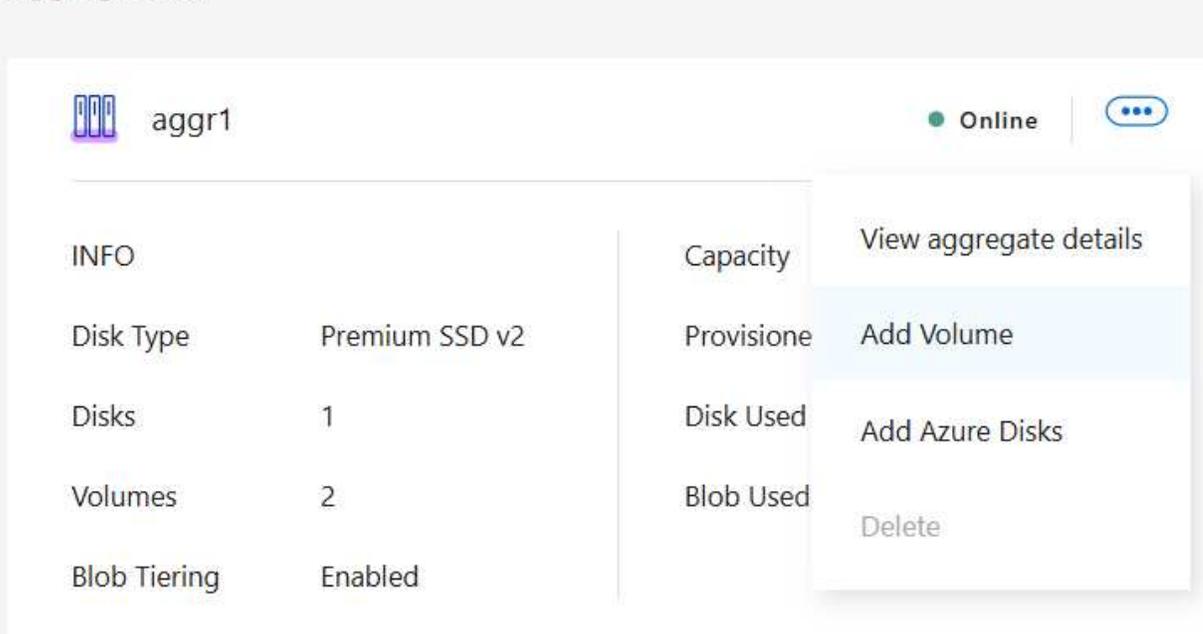
Volume (1) 🔍 ☰ [Add Volume](#)

 v456 ● Online [Manage Volume](#)

Specific aggregate

1. On the **Aggregates** tab, go to the required the aggregate and click the **...** icon.
2. Select **Add Volume**.

Aggregate (1)



 aggr1 ● Online ⋮

| | | | |
|--------------|----------------|-------------|------------------------|
| INFO | | Capacity | View aggregate details |
| Disk Type | Premium SSD v2 | Provisioned | Add Volume |
| Disks | 1 | Disk Used | Add Azure Disks |
| Volumes | 2 | Blob Used | Delete |
| Blob Tiering | Enabled | | |

3. Follow the steps in the wizard to create the volume.

- a. **Details, Protection, and Tags:** Enter basic details about the volume and select a Snapshot policy.

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

| Field | Description |
|-------------|---|
| Volume Name | The identifiable name you can enter for the new volume. |
| Volume 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. |

| Field | Description |
|------------------|--|
| Storage VM (SVM) | A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an SVM or a vserver. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs. You can specify the Storage VM for the new 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. |

- b. **Protocol:** Choose a protocol for the volume (NFS, CIFS, or iSCSI) and then provide the required information.

If you select CIFS and a server isn't set up, the Console prompts you to set up CIFS connectivity after you click **Next**.

[Learn about supported client protocols and versions.](#)

The following sections describe fields for which you might need guidance. The descriptions are organized by protocol.

NFS

Access control

Choose a custom export policy to make the volume available to clients.

Export policy

Defines the clients in the subnet that can access the volume. By default, the Console enters a value that provides access to all instances in the subnet.

CIFS

Permissions and users/groups

Enables you to control the level of access to an SMB 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.

DNS Primary and Secondary IP Address

The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.

If you're configuring Google Managed Active Directory, AD can be accessed by default with the 169.254.169.254 IP address.

Active Directory Domain to join

The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.

Credentials authorized to join the domain

The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.

CIFS server NetBIOS name

A CIFS server name that is unique in the AD domain.

Organizational Unit

The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.

- To configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter **OU=Computers,OU=corp** in this field.
- To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, enter **OU=AADDC Computers** or **OU=AADDC Users** in this field.
[Azure Documentation: Create an Organizational Unit \(OU\) in an Azure AD Domain Services managed domain](#)
- To configure Google Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter **OU=Computers,OU=Cloud** in this field.
[Google Cloud Documentation: Organizational Units in Google Managed Microsoft AD](#)

DNS Domain

The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.

NTP Server

Select **Use Active Directory Domain** to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. For information, refer to the [NetApp Console automation docs](#).

Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.

iSCSI

LUN

iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices. When you create an iSCSI volume, the Console automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, [use the IQN to connect to the LUN from your hosts](#).

Initiator group

Initiator groups (igroups) specify which hosts can access specified LUNs on the storage system

Host initiator (IQN)

iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).

- c. **Disk Type:** Choose an underlying disk type for the volume based on your performance needs and cost requirements.
- [Sizing your system in AWS](#)
 - [Sizing your system in Azure](#)
 - [Sizing your system in Google Cloud](#)
- d. **Usage Profile & Tiering Policy:** Choose whether to enable or disable storage efficiency features on the volume and then select a [volume tiering policy](#).

ONTAP includes several storage efficiency features that can reduce the total amount of storage that you need. NetApp storage efficiency features provide the following benefits:

Thin provisioning

Presents more logical storage to hosts or users than you actually have in your physical storage pool. Instead of preallocating storage space, storage space is allocated dynamically to each volume as data is written.

Deduplication

Improves efficiency by locating identical blocks of data and replacing them with references to a single shared block. This technique reduces storage capacity requirements by eliminating redundant blocks of data that reside in the same volume.

Compression

Reduces the physical capacity required to store data by compressing data within a volume on primary, secondary, and archive storage.

- e. **Review:** Review details about the volume and then click **Add**.

Result

The Console creates the volume on the Cloud Volumes ONTAP system.

Create a volume on the second node in an HA configuration

By default, the Console creates volumes on the first node in an HA configuration. If you need an active-active configuration, in which both nodes serve data to clients, you must create aggregates and volumes on the second node.

Steps

1. From the left navigation menu, select **Storage > Management**.
2. On the **Systems** page, double-click the name of the Cloud Volumes ONTAP system on which you want to manage aggregates.
3. On the Aggregates tab, click **Add Aggregate**, and create the aggregate.

The screenshot shows the 'Aggregates Summary' section with four metrics: Total Aggregates (1), Aggregates with Tiering (1), Aggregates without Tiering (0), and Allocated Disks (1). Below this is a table for 'Aggregate (1)' with one entry, 'aggr1', which is 'Online'. A detailed view of 'aggr1' is shown below, with two columns: 'INFO' and 'Capacity'.

| INFO | | Capacity | |
|--------------|----------------|------------------|------------|
| Disk Type | Premium SSD v2 | Provisioned size | 907.18 GiB |
| Disks | 1 | Disk Used | 1.15 GiB |
| Volumes | 2 | Blob Used | 0 GiB |
| Blob Tiering | Enabled | | |

4. For Home Node, choose the second node in the HA pair.
5. After the Console creates the aggregate, select it and then click **Create volume**.
6. Enter details for the new volume, and then click **Create**.

Result

The Console creates the volume on the second node in the HA pair.



For HA pairs deployed in multiple AWS Availability Zones, you must mount the volume to clients by using the floating IP address of the node on which the volume resides.

After you create a volume

If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.

If you want to apply quotas to volumes, you must use ONTAP System Manager or the ONTAP CLI. Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

Manage volumes on Cloud Volumes ONTAP systems

You can manage volumes and CIFS servers in the NetApp Console. You can also move volumes to avoid capacity issues.

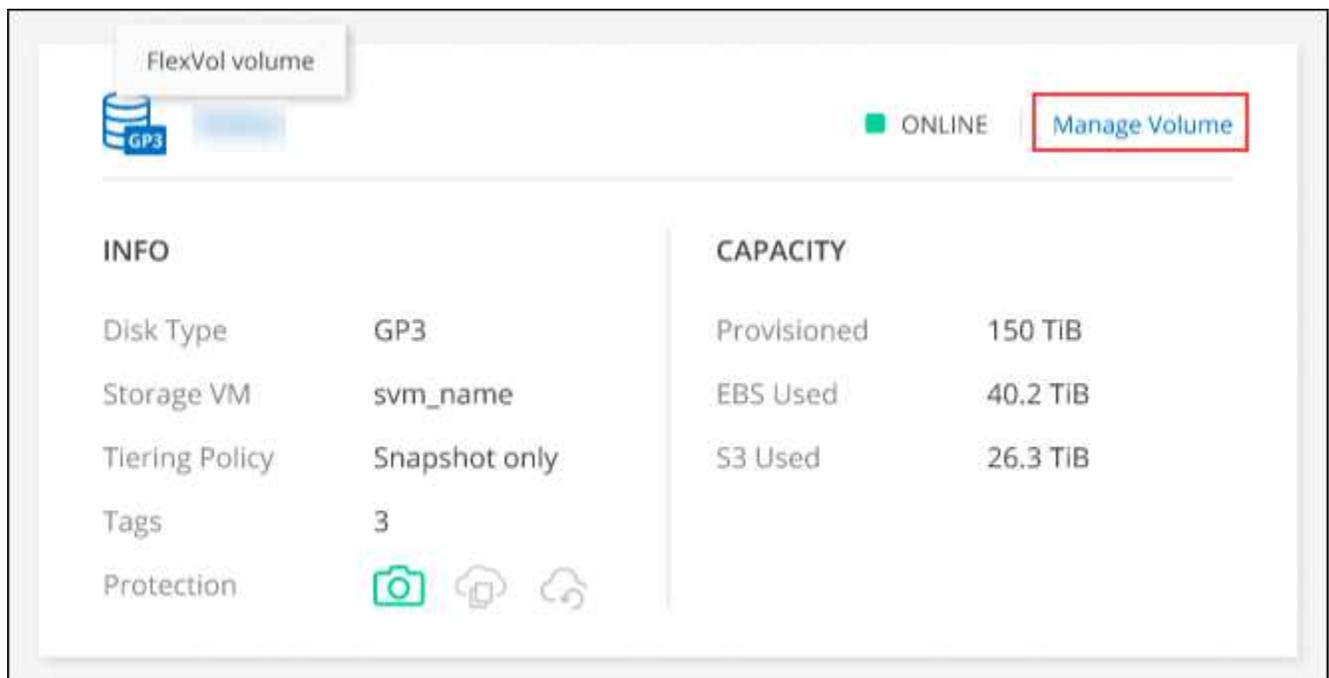
You can manage volumes in the NetApp Console Standard View or through ONTAP System Manager that is included within the Console for advanced volume management. The Standard View provides a limited set of options to modify your volumes. System Manager provides advanced level of management, such as cloning, resizing, changing settings for anti-ransomware, analytics, protection, and activity tracking, and moving volumes across tiers. For information, refer to [Administer Cloud Volumes ONTAP using System Manager](#).

Manage volumes

By using the Standard View of the Console, you can manage volumes according to your storage needs. You can view, edit, clone, restore, and delete volumes.

Steps

1. From the left navigation menu, select **Storage > Management**.
2. On the **Systems** page, double-click the Cloud Volumes ONTAP system on which you want to manage volumes.
3. Select the **Volumes** tab.



4. On the required volume tile, click **Manage volume**.

| Task | Action |
|---------------------------------|--|
| View information about a volume | Under Volume Actions in the Manage volumes panel, click View volume details . |

| Task | Action |
|---|--|
| Get the NFS mount command | <ol style="list-style-type: none"> Under Volume Actions in the Manage volumes panel, click Mount Command. Click Copy. |
| Clone a volume | <ol style="list-style-type: none"> Under Volume Actions in the Manage volumes panel, click Clone the volume. Modify the clone name as needed, and then click Clone. <p>This process creates a FlexClone volume. A FlexClone volume is a writable, point-in-time copy that is space-efficient because it uses a small amount of space for metadata, and then only consumes additional space as data is changed or added.</p> <p>To learn more about FlexClone volumes, refer to the ONTAP 9 Logical Storage Management Guide.</p> |
| Edit a volume (read-write volumes only) | <ol style="list-style-type: none"> Under Volume Actions in the Manage volumes panel, click Edit volume settings Modify the volume's Snapshot policy, NFS protocol version, NFS access control list (export policy), or share permissions, and then click Apply. <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  If you need custom Snapshot policies, you can create them by using ONTAP System Manager. </div> |
| Delete a volume | <ol style="list-style-type: none"> Under Volume Actions in the Manage volumes panel, click Delete the volume. Under the Delete Volume window, enter the name of the volume you want to delete. Click Delete again to confirm. |
| Create a Snapshot copy on demand | <ol style="list-style-type: none"> Under Protection Actions in the Manage Volumes panel, click Create a Snapshot copy. Change the name, if needed, and then click Create. |
| Restore data from a Snapshot copy to a new volume | <ol style="list-style-type: none"> Under Protection Actions in the Manage Volumes panel, click Restore from Snapshot copy. Select a Snapshot copy, enter a name for the new volume, and then click Restore. |

| Task | Action |
|---------------------------------|---|
| Change the underlying disk type | <ol style="list-style-type: none"> Under Advanced Actions in the Manage Volumes panel, click Change Disk Type. Select the disk type, and then click Change. <p> The Console moves the volume to an existing aggregate that uses the selected disk type or it creates a new aggregate for the volume.</p> |
| Change the tiering policy | <ol style="list-style-type: none"> Under Advanced Actions in the Manage Volumes panel, click Change Tiering Policy. Select a different policy and click Change. <p> The Console moves the volume to an existing aggregate that uses the selected disk type with tiering, or it creates a new aggregate for the volume.</p> |
| Delete a volume | <ol style="list-style-type: none"> Select a volume, and then click Delete. Type the name of the volume in the dialog. Click Delete again to confirm. |

Resize a volume

By default, a volume automatically grows to a maximum size when it's out of space. The default value is 1,000, which means the volume can grow to 11 times its size. This value is configurable in the Console agent's settings.

If you need to resize your volume, you can do it from ONTAP System Manager in the Console.

Steps

- Click the System Manager view to resize a volume through ONTAP System Manager. Refer to [How to get started](#).
- From the left navigation menu, select **Storage > Volumes**.
- From the list of volumes, identify the one that you should resize.
- Click the options icon .
- Select **Resize**.
- On the **Resize Volume** screen, edit the capacity and Snapshot reserve percentage as required. You can compare the existing, available space with the modified capacity.
- Click **Save**.

Resize volume ✕

CAPACITY

25
⇅

GiB
▼

SNAPSHOT RESERVE %

1
⇅

| Existing | New |
|------------------|------------------|
| DATA SPACE | DATA SPACE |
| 20 GiB | 24.75 GiB |
| SNAPSHOT RESERVE | SNAPSHOT RESERVE |
| 0 Bytes | 256 MiB |

Cancel
Save

Be sure to take your system's capacity limits into consideration as you resize volumes. Go to the [Cloud Volumes ONTAP Release Notes](#) for more information.

Modify the CIFS server

If you change your DNS servers or Active Directory domain, you need to modify the CIFS server in Cloud Volumes ONTAP so that it can continue to serve storage to clients.

Steps

1. From the **Overview** tab of the Cloud Volumes ONTAP system, click the **Feature** tab under the right-side panel.
2. Under the CIFS Setup field, click the **pencil icon** to display the CIFS Setup window.
3. Specify settings for the CIFS server:

| Task | Action |
|---------------------------------|---|
| Select Storage VM (SVM) | Selecting the Cloud Volume ONTAP storage virtual machine (SVM) displays it's configured CIFS information. |
| Active Directory Domain to join | The FQDN of the Active Directory (AD) domain that you want the CIFS server to join. |

| Task | Action |
|---|---|
| Credentials authorized to join the domain | The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain. |
| DNS Primary and Secondary IP Address | <p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p> <p>If you're configuring Google Managed Active Directory, AD can be accessed by default with the 169.254.169.254 IP address.</p> |
| DNS Domain | The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain. |
| CIFS server NetBIOS name | A CIFS server name that is unique in the AD domain. |
| Organizational Unit | <p>The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.</p> <ul style="list-style-type: none"> • To configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter OU=Computers,OU=corp in this field. • To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, enter OU=AADDC Computers or OU=AADDC Users in this field. Azure Documentation: Create an Organizational Unit (OU) in an Azure AD Domain Services managed domain • To configure Google Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter OU=Computers,OU=Cloud in this field. Google Cloud Documentation: Organizational Units in Google Managed Microsoft AD |

4. Click **Set**.

Result

Cloud Volumes ONTAP updates the CIFS server with the changes.

Move a volume

Move volumes for capacity utilization, improved performance, and to satisfy service-level agreements.

You can move a volume in ONTAP System Manager by selecting a volume and the destination aggregate, starting the volume move operation, and optionally monitoring the volume move job. When using System Manager, a volume move operation finishes automatically.

Steps

1. Use ONTAP System Manager or the ONTAP CLI to move the volumes to the aggregate.

In most situations, you can use System Manager to move volumes.

For instructions, refer to the [ONTAP 9 Volume Move Express Guide](#).

Move a volume when Console displays an Action Required message

The Console might display an Action Required message that says moving a volume is necessary to avoid capacity issues, but that you need to correct the issue yourself. If this happens, you need to identify how to correct the issue and then move one or more volumes.



The Console displays these Action Required messages when an aggregate has reached 90% used capacity. If data tiering is enabled, the messages display when an aggregate has reached 80% used capacity. By default, 10% free space is reserved for data tiering. [Learn more about the free space ratio for data tiering.](#)

Steps

1. [Identify how to correct capacity issues.](#)
2. Based on your analysis, move volumes to avoid capacity issues:
 - [Move volumes to another system to avoid capacity issues.](#)
 - [Move volumes to another aggregate to avoid capacity issues.](#)

Identify how to correct capacity issues

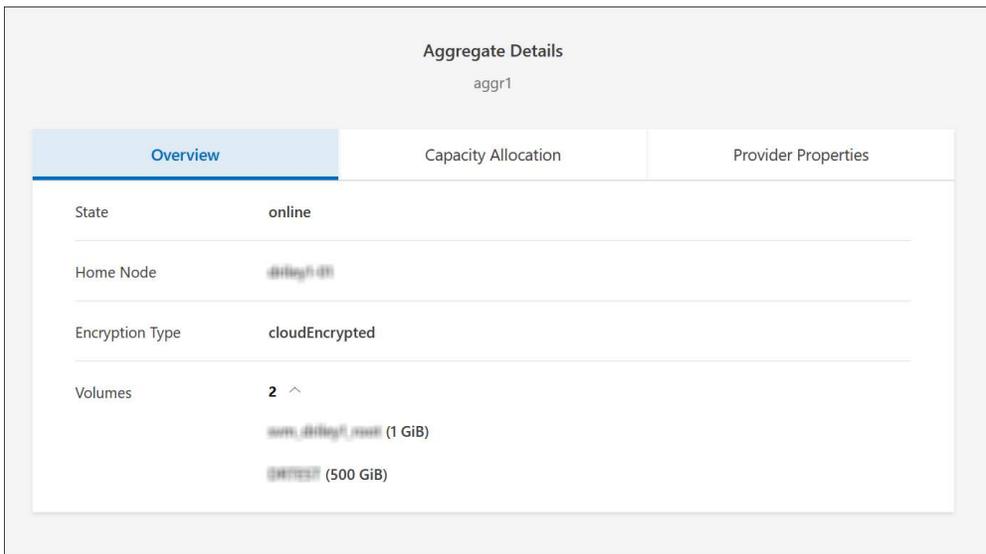
If the Console can't provide recommendations for moving a volume to avoid capacity issues, you must identify the volumes that you need to move and whether you should move them to another aggregate on the same system or to another system.

Steps

1. View the advanced information in the Action Required message to identify the aggregate that has reached its capacity limit.

For example, the advanced information should say something similar to the following: Aggregate aggr1 has reached its capacity limit.
2. Identify one or more volumes to move out of the aggregate:
 - a. In the Cloud Volumes ONTAP system, click the **Aggregates tab**.
 - b. On the aggregate tile, click the **...** icon and then click **View aggregate details**.
 - c. Under the **Overview** tab of the **Aggregate Details** screen, review the size of each volume and choose one or more volumes to move out of the aggregate.

You should choose volumes that are large enough to free space in the aggregate so that you avoid additional capacity issues in the future.



3. If the system has not reached the disk limit, you should move the volumes to an existing aggregate or a new aggregate on the same system.

For information, refer to [Move volumes to another aggregate to avoid capacity issues](#).

4. If the system has reached the disk limit, do any of the following:
 - a. Delete any unused volumes.
 - b. Rearrange volumes to free space on an aggregate.

For information, refer to [Move volumes to another aggregate to avoid capacity issues](#).

- c. Move two or more volumes to another system that has space.

For information, refer to [Move volumes to another aggregate to avoid capacity issues](#).

Move volumes to another system to avoid capacity issues

You can move one or more volumes to another Cloud Volumes ONTAP system to avoid capacity issues. You might need to do this if the system reached its disk limit.

About this task

You can follow the steps in this task to correct the following Action Required message:

Moving a volume is necessary to avoid capacity issues; however, the Console cannot perform this action for you because the system has reached the disk limit.

Steps

1. Identify a Cloud Volumes ONTAP system that has available capacity, or deploy a new system.
2. Drag and drop the source system to the target system to perform a one-time data replication of the volume.

For information, refer to [Replicating data between systems](#).

3. Go to the Replication Status page, and then break the SnapMirror relationship to convert the replicated volume from a data protection volume to a read/write volume.

For information, refer to [Managing data replication schedules and relationships](#).

4. Configure the volume for data access.

For information about configuring a destination volume for data access, refer to the [ONTAP 9 Volume Disaster Recovery Express Guide](#).

5. Delete the original volume.

For information, refer to [Manage volumes](#).

Move volumes to another aggregate to avoid capacity issues

You can move one or more volumes to another aggregate to avoid capacity issues.

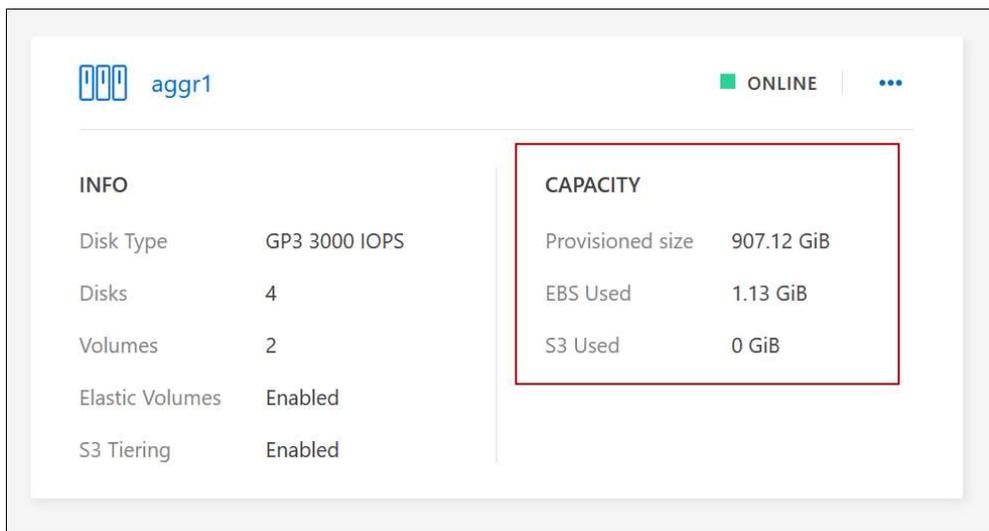
About this task

You can follow the steps in this task to correct the following Action Required message:

Moving two or more volumes is necessary to avoid capacity issues; however, the Console cannot perform this action for you.

Steps

1. Verify whether an existing aggregate has available capacity for the volumes that you need to move:
 - a. On Cloud Volumes ONTAP system, click the **Aggregates** tab.
 - b. On the required aggregate tile, click the **...** icon and then **View aggregate details** to view the available capacity (provisioned size minus used aggregate capacity).



2. If needed, add disks to an existing aggregate:
 - a. Select the aggregate, then click the **...** icon > **Add Disks**.
 - b. Select the number of disks to add, and then click **Add**.
3. If no aggregates have available capacity, create a new aggregate.

For information, refer to [Creating aggregates](#).

4. Use ONTAP System Manager or the ONTAP CLI to move the volumes to the aggregate.
5. In most situations, you can use System Manager to move volumes.

For instructions, refer to the [ONTAP 9 Volume Move Express Guide](#).

Reasons why a volume move might perform slowly

Moving a volume might take longer than you expect if any of the following conditions are true for Cloud Volumes ONTAP:

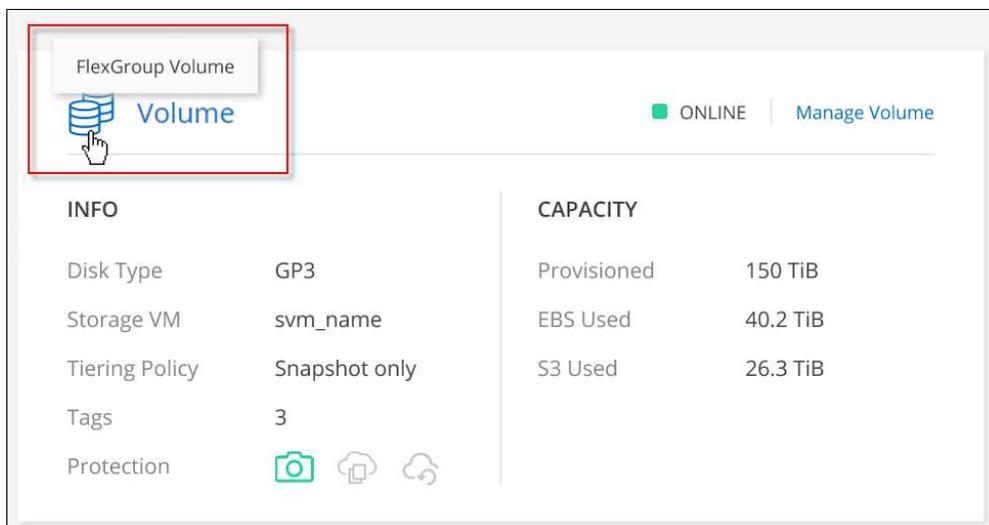
- The volume is a clone.
- The volume is a parent of a clone.
- The source or destination aggregate has a single Throughput Optimized HDD (st1) disk.
- One of the aggregates uses an older naming scheme for objects. Both aggregates have to use the same name format.

An older naming scheme is used if data tiering was enabled on an aggregate in the 9.4 release or earlier.

- The encryption settings don't match on the source and destination aggregates, or a rekey is in progress.
- The *-tiering-policy* option was specified on the volume move to change the tiering policy.
- The *-generate-destination-key* option was specified on the volume move.

View FlexGroup Volumes

You can view FlexGroup volumes created through ONTAP System Manager or the ONTAP CLI directly through the Volumes tab in the Console. You can see detailed information for the FlexGroup volumes through a dedicated **Volumes** tile, where you can identify each FlexGroup volume group through the icon's hover text. Additionally, you can identify and sort FlexGroup volumes under the volumes list view through the Volume Style column.



| INFO | | CAPACITY | |
|----------------|---------------|-------------|----------|
| Disk Type | GP3 | Provisioned | 150 TiB |
| Storage VM | svm_name | EBS Used | 40.2 TiB |
| Tiering Policy | Snapshot only | S3 Used | 26.3 TiB |
| Tags | 3 | | |
| Protection | | | |



Currently, you can only view existing FlexGroup volumes under the Console. You can't create FlexGroup volumes in the Console.

Tier inactive Cloud Volumes ONTAP data to a low-cost object storage

You can reduce storage costs for Cloud Volumes ONTAP by combining an SSD or HDD performance tier for hot data with an object storage capacity tier for inactive data. Data tiering is powered by FabricPool technology. For a high-level overview, refer to [Data tiering overview](#).

To set up data tiering, you need to do the following:

1

Choose a supported configuration

Most configurations are supported. If you have a Cloud Volumes ONTAP system running the most recent version, then you are good to go. [Learn more](#).

2

Ensure connectivity between Cloud Volumes ONTAP and object storage

- For AWS, you'll need a VPC Endpoint to Amazon Simple Storage Service (Amazon S3). [Learn more](#).
- For Azure, you won't need to do anything as long as the NetApp Console has the required permissions. [Learn more](#).
- For Google Cloud, you need to configure the subnet for Private Google Access and set up a service account. [Learn more](#).

3

Ensure that you have an aggregate with tiering enabled

Data tiering should be enabled on an aggregate to enable it on a volume. You should be aware of the requirements for new volumes and for existing volumes. [Learn more](#).

4

Choose a tiering policy when creating, modifying, or replicating a volume

The NetApp Console prompts you to choose a tiering policy when you create, modify, or replicate a volume.

- [Tier data from read-write volumes](#)
- [Tier data from data protection volumes](#)

What's not required for data tiering?

- You don't need to install a feature license to enable data tiering.
- You don't need to create an object store for the capacity tier. The Console does that for you.
- You don't need to enable data tiering at the system level.



The Console creates an object store for cold data when it creates the system, [as long as there are no connectivity or permissions issues](#). After that, you just need to enable data tiering on volumes (and in some cases, [on aggregates](#)).

Configurations that support data tiering

You can enable data tiering when using specific configurations and features.

Support in AWS

- Data tiering is supported in AWS beginning with Cloud Volumes ONTAP 9.2.
- The performance tier can be General Purpose SSDs (gp3 or gp2) or Provisioned IOPS SSDs (io1).



We do not recommend tiering data to object storage when using Throughput Optimized HDDs (st1).

- The inactive data is tiered to Amazon S3 buckets. Tiering to other providers is not supported.

Support in Azure

- Data tiering is supported in Azure as follows:
 - Version 9.4 in with single-node systems
 - Version 9.6 in with HA pairs
- The performance tier can be Premium SSD managed disks, Standard SSD managed disks, or Standard HDD managed disks.
- The inactive data is tiered to Microsoft Azure Blob. Tiering to other providers is not supported.

Support in Google Cloud

- Data tiering is supported in Google Cloud beginning with Cloud Volumes ONTAP 9.6.
- The performance tier can be either SSD persistent disks, balanced persistent disks, or standard persistent disks.
- The inactive data is tiered to Google Cloud Storage. Tiering to other providers is not supported.

Feature interoperability

- Data tiering is supported with encryption technologies.
- Thin provisioning must be enabled on volumes.

Requirements

Depending on your cloud provider, certain connections and permissions must be set up so that Cloud Volumes ONTAP can tier cold data to object storage.

Requirements to tier cold data to Amazon S3

Ensure that Cloud Volumes ONTAP has a connection to Amazon S3. The best way to provide that connection is by creating a VPC Endpoint to the S3 service. For instructions, refer to the [AWS Documentation: Creating a Gateway Endpoint](#).

When you create the VPC Endpoint, be sure to select the region, VPC, and route table that corresponds to the Cloud Volumes ONTAP instance. You must also modify the security group to add an outbound HTTPS rule that enables traffic to the S3 endpoint. Otherwise, Cloud Volumes ONTAP cannot connect to the S3 service.

If you experience any issues, refer to [AWS Support Knowledge Center: Why can't I connect to an S3 bucket using a gateway VPC endpoint?](#).

Requirements to tier cold data to Azure Blob storage

You don't need to set up a connection between the performance tier and the capacity tier as long as the Console has the required permissions. The Console enables a VNet service endpoint for you if the custom role for the Console agent has these permissions:

```
"Microsoft.Network/virtualNetworks/subnets/write",  
"Microsoft.Network/routeTables/join/action",
```

The custom role includes the permissions by default. [View Azure permission for the Console agent](#)

Requirements to tier cold data to a Google Cloud Storage bucket

- The subnet in which Cloud Volumes ONTAP resides must be configured for Private Google Access. For instructions, refer to [Google Cloud Documentation: Configuring Private Google Access](#).
- A service account must be attached to Cloud Volumes ONTAP.

[Learn how to set up this service account.](#)

You're prompted to select this service account when you create a Cloud Volumes ONTAP system.

If you don't select a service account during deployment, you'll need to shut down Cloud Volumes ONTAP, go to the Google Cloud Console, and then attach the service account to the Cloud Volumes ONTAP instances. You can then enable data tiering as described in the next section.

- To encrypt the bucket with customer-managed encryption keys, enable the Google Cloud storage bucket to use the key.

[Learn how to use customer-managed encryption keys with Cloud Volumes ONTAP.](#)

Enable data tiering after implementing the requirements

The Console creates an object store for cold data when the system is created, as long as there are no connectivity or permissions issues. If you didn't implement the requirements listed above until after you created the system, then you'll need to manually enable tiering through the API or ONTAP System Manager, which creates the object store.



The ability to enable tiering through the Console will be available in a future Cloud Volumes ONTAP release.

Ensure that tiering is enabled on aggregates

Data tiering must be enabled on an aggregate in order to enable data tiering on a volume. You should be aware of the requirements for new volumes and for existing volumes.

- **New volumes**

If you're enabling data tiering on a new volume, then you don't need to worry about enabling data tiering on

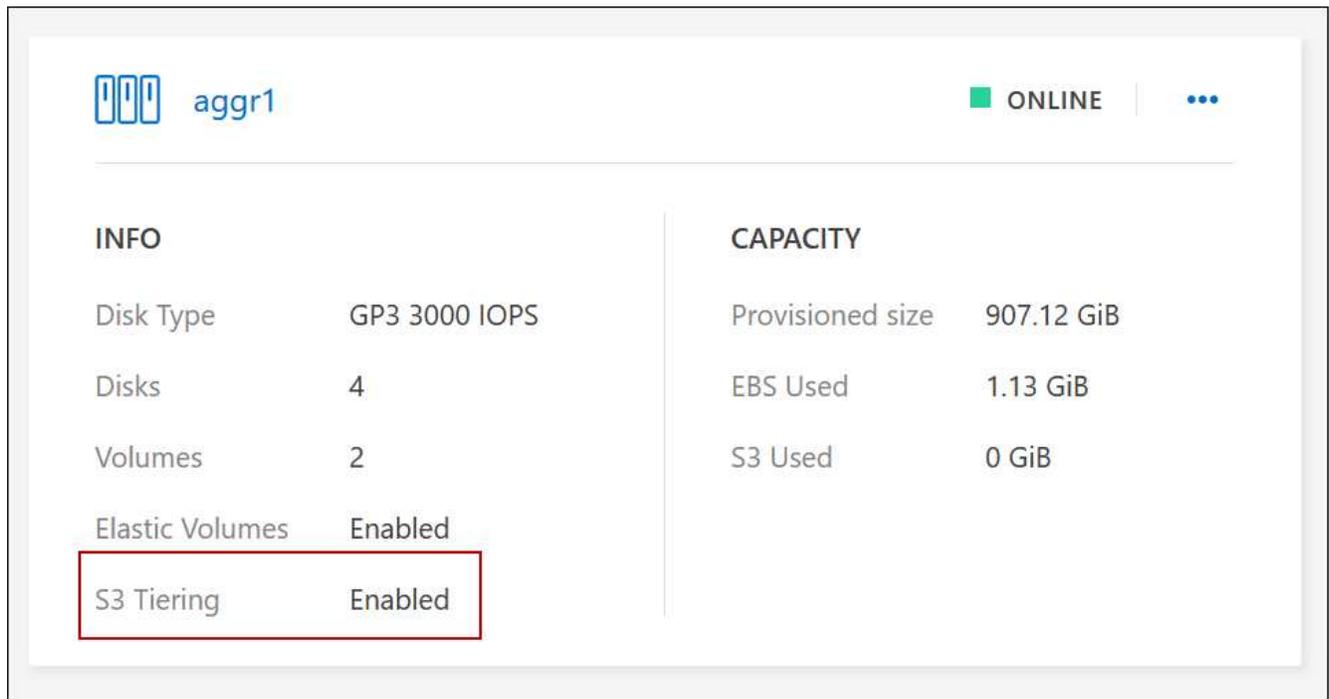
an aggregate. The Console creates the volume on an existing aggregate that has tiering enabled, or it creates a new aggregate for the volume if a data tiering-enabled aggregate doesn't already exist.

- **Existing volumes**

To enable data tiering on an existing volume, ensure it is enabled on the underlying aggregate. If data tiering isn't enabled on the existing aggregate, then you'll need to use ONTAP System Manager to attach an existing aggregate to the object store.

Steps to confirm whether tiering is enabled on an aggregate

1. From the left navigation menu, select **Storage > Management**.
2. Open the Cloud Volumes ONTAP system.
3. Select the **Aggregates** tab and check if tiering is enabled or disabled on the aggregate.



Steps to enable tiering on an aggregate

1. In ONTAP System Manager, click **Storage > Tiers**.
2. Click the action menu for the aggregate and select **Attach Cloud Tiers**.
3. Select the cloud tier to attach and click **Save**.

What's next?

You can now enable data tiering on new and existing volumes, as explained in the next section.

Tier data from read-write volumes

Cloud Volumes ONTAP can tier inactive data on read-write volumes to cost-effective object storage, freeing up the performance tier for hot data.

Steps

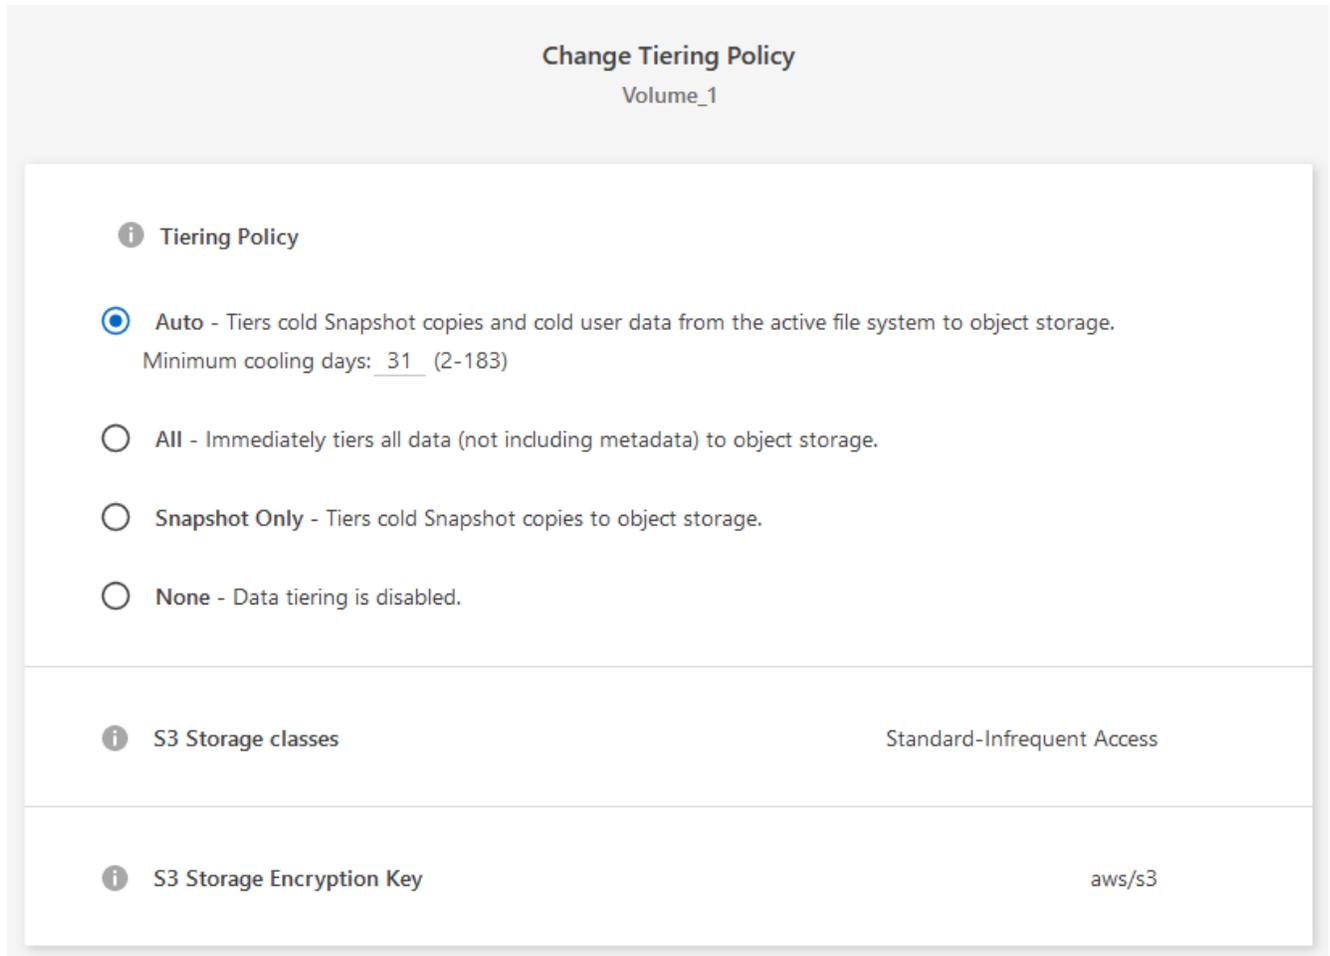
1. In the **Volumes** tab under the system, create a new volume or change the tier of an existing volume:

| Task | Action |
|---------------------------|--|
| Create a new volume | Click Add New Volume . |
| Modify an existing volume | Select the desired volume tile, click Manage volume to access the Manage Volumes right-side panel, and then click Advanced actions and Change tiering policy under the right panel. |

2. Select a tiering policy.

For a description of these policies, refer to [Data tiering overview](#).

Example



The Console creates a new aggregate for the volume if a data tiering-enabled aggregate does not already exist.

Tier data from data protection volumes

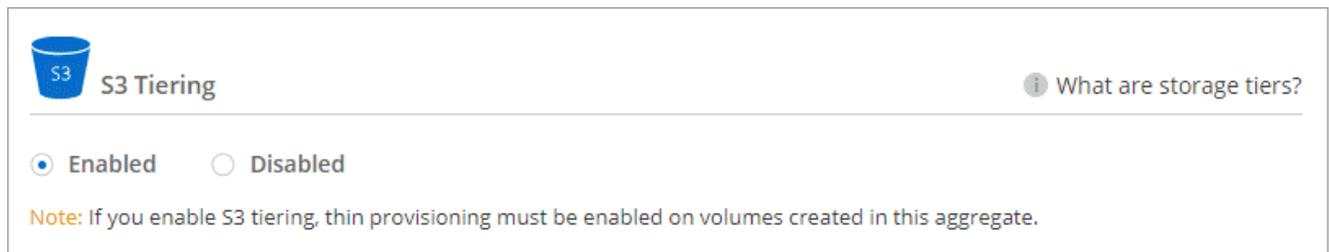
Cloud Volumes ONTAP can tier data from a data protection volume to a capacity tier. If you activate the destination volume, the data gradually moves to the performance tier as it is read.

Steps

1. From the left navigation menu, select **Storage > Management**.

2. On the **Systems** page, select the Cloud Volumes ONTAP system that contains the source volume, and then drag it to the system to which you want to replicate the volume.
3. Follow the prompts until you reach the tiering page and enable data tiering to object storage.

Example



For help with replicating data, refer to [Replicating data to and from the cloud](#).

Change the storage class for tiered data

After you deploy Cloud Volumes ONTAP, you can reduce your storage costs by changing the storage class for inactive data that hasn't been accessed for 30 days. The access costs are higher if you do access the data, so you must take that into consideration before you change the storage class.

The storage class for tiered data is system wide—it's not per volume.

For information about supported storage classes, refer to [Data tiering overview](#).

Steps

1. On the Cloud Volumes ONTAP system, click the menu icon and then click **Storage Classes** or **Blob Storage Tiering**.
2. Choose a storage class and then click **Save**.

Change the free space ratio for data tiering

The free space ratio for data tiering defines how much free space is required on Cloud Volumes ONTAP SSDs/HDDs when tiering data to object storage. The default setting is 10% free space, but you can tweak the setting based on your requirements.

For example, you might choose less than 10% free space to ensure that you are utilizing the purchased capacity. The Console can then purchase additional disks for you when additional capacity is required (up until you reach the disk limit for the aggregate).

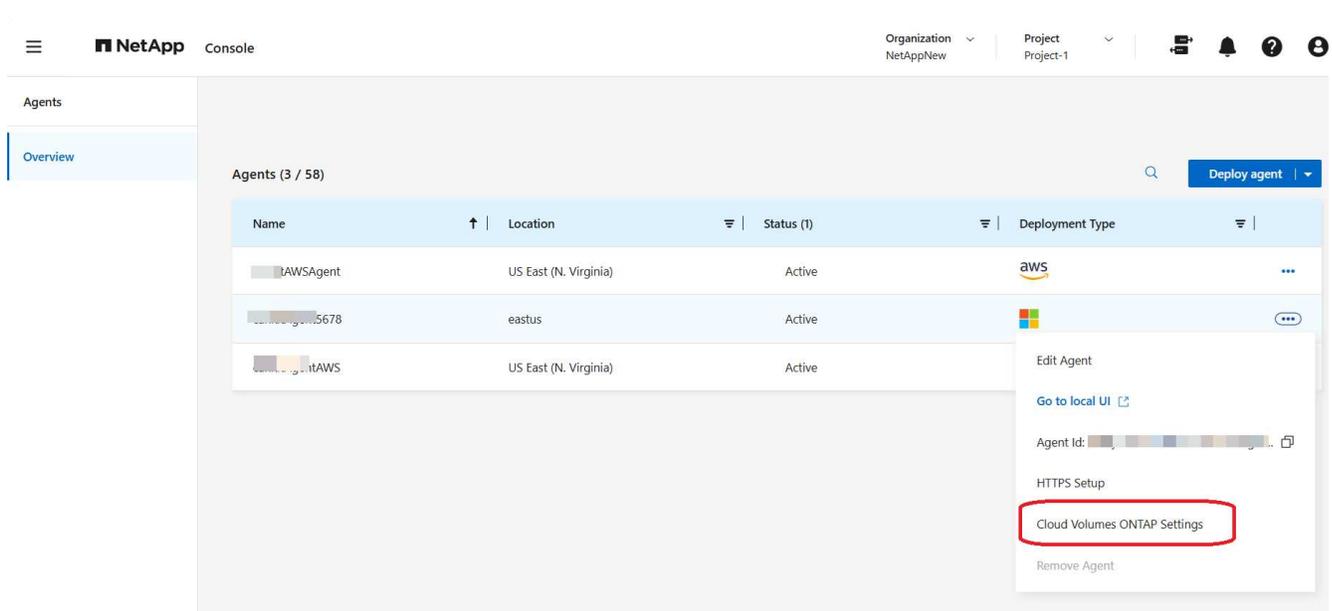


If there isn't sufficient space, then Cloud Volumes ONTAP can't move the data and you might experience performance degradation. Any change should be done with caution. If you're unsure, reach out to NetApp Support for guidance.

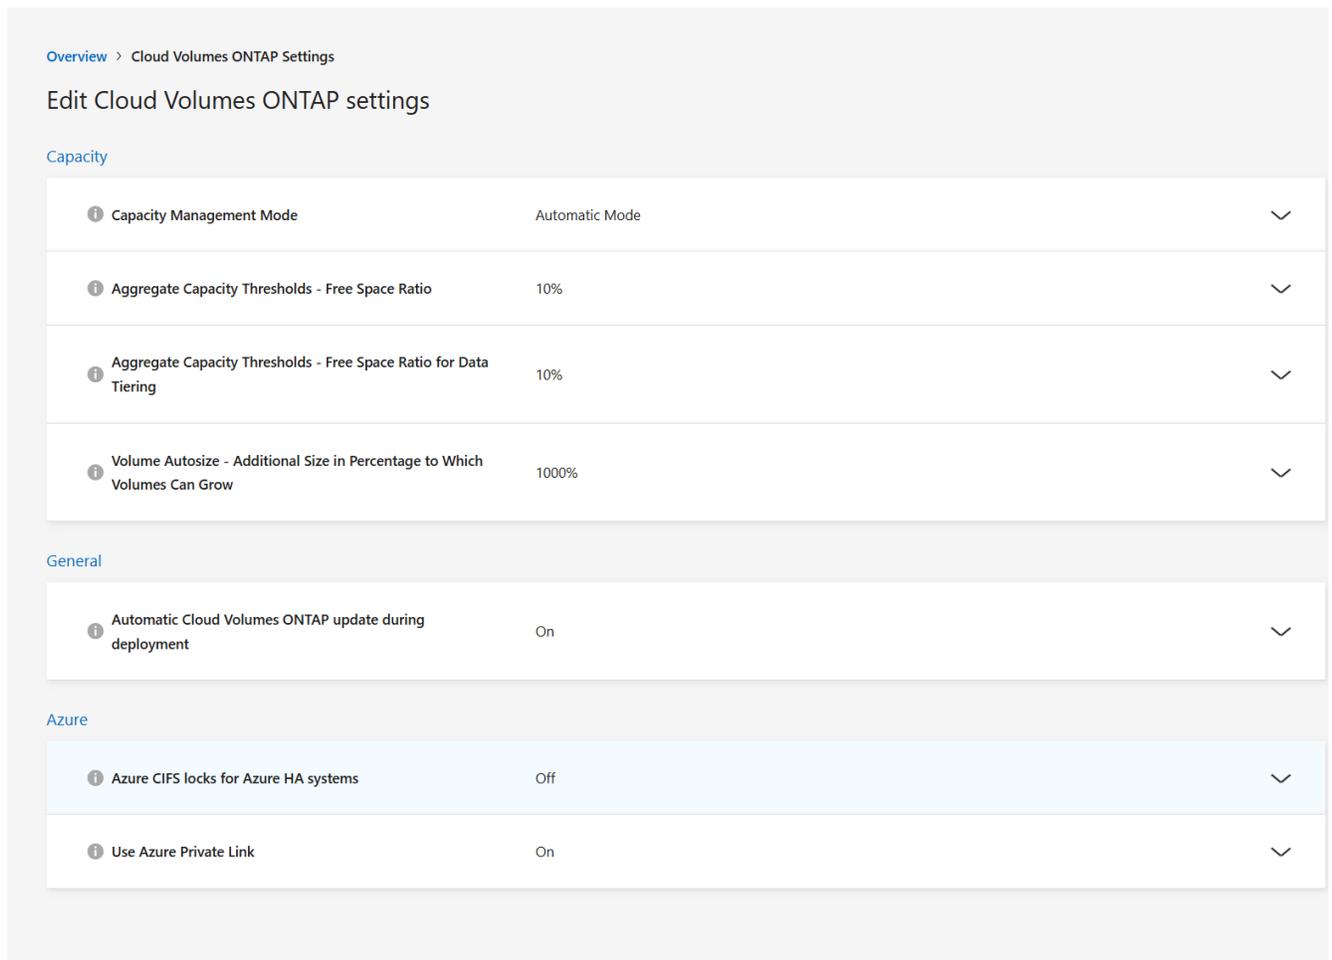
The ratio is important for disaster recovery scenarios because as data is read from the object store, Cloud Volumes ONTAP moves the data to SSDs/HDDs to provide better performance. If there isn't sufficient space, then Cloud Volumes ONTAP can't move the data. Take this into consideration when changing the ratio so that you can meet your business requirements.

Steps

1. From the left navigation pane, go to **Administration > Agents**.
2. Click the **...** icon for the Console agent that manages your Cloud Volumes ONTAP system.
3. Select **Cloud Volumes ONTAP Settings**.



4. Under **Capacity**, click **Aggregate Capacity Thresholds - Free Space Ratio for Data Tiering**.



5. Change the free space ratio based on your requirements and click **Save**.

Change the cooling period for the auto tiering policy

If you enabled data tiering on a Cloud Volumes ONTAP volume using the *auto* tiering policy, you can adjust the default cooling period based on your business needs. This action is supported using ONTAP CLI and API only.

The cooling period is the number of days that user data in a volume must remain inactive before it is considered "cold" and moved to object storage.

The default cooling period for the auto tiering policy is 31 days. You can change the cooling period as follows:

- 9.8 or later: 2 days to 183 days
- 9.7 or earlier: 2 days to 63 days

Step

1. Use the *minimumCoolingDays* parameter with your API request when creating a volume or modifying an existing volume.

Remove an S3 bucket on decommissioning a system

You can delete an S3 bucket with the data tiered from a Cloud Volumes ONTAP system when you decommission the environment.

You can delete the S3 bucket only if:

- The Cloud Volume ONTAP system is deleted from the Console.
- All objects are deleted from the bucket and the S3 bucket is empty.

When you decommission a Cloud Volumes ONTAP system, the S3 bucket that was created for the environment is not deleted automatically. Instead, it remains in an orphaned state to prevent any accidental data loss. You can delete the objects in the bucket, then remove the S3 bucket itself, or keep it for later use. Refer to [ONTAP CLI: vservers object-store-server bucket delete](#).

Connect to a LUN on Cloud Volumes ONTAP from your host system

When you create an iSCSI volume, the NetApp Console automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, use the IQN to connect to the LUN from your hosts.

Note the following:

- The Console's automatic capacity management doesn't apply to LUNs. When it creates a LUN, it disables the autogrow feature.
- You can create additional LUNs from ONTAP System Manager or the ONTAP CLI.

Steps

1. From the left navigation menu, select **Storage > Management**.

2. On the **Systems** page, double-click the Cloud Volumes ONTAP system on which you want to manage volumes.
3. In the system, select the **Volumes** tab.
4. Go to the required volume tile and then select **Manage volume** to access the Manage Volumes panel on the right.
5. Click **Target iQN**.
6. Click **Copy** to copy the iQN name.
7. Set up an iSCSI connection from the host to the LUN.
 - [ONTAP 9 iSCSI express configuration for Red Hat Enterprise Linux: Starting the iSCSI sessions with the target](#)
 - [ONTAP 9 iSCSI express configuration for Windows: Starting iSCSI sessions with the target](#)
 - [ONTAP SAN host configuration](#)

Accelerate data access with FlexCache volumes on a Cloud Volumes ONTAP system

A FlexCache volume is a storage volume that caches SMB and NFS read data from an origin (or source) volume. Subsequent reads to the cached data result in faster access to that data.

You can use FlexCache volumes to speed up access to data or to offload traffic from heavily accessed volumes. FlexCache volumes help improve performance, especially when clients need to access the same data repeatedly, because the data can be served directly without having to access the origin volume. FlexCache volumes work well for system workloads that are read-intensive.

NetApp Console provides management of FlexCache volumes with the [NetApp Volume Caching](#).

You can also use the ONTAP CLI or ONTAP System Manager to create and manage FlexCache volumes:

- [FlexCache Volumes for Faster Data Access Power Guide](#)
- [Creating FlexCache volumes in System Manager](#)



Work with FlexCache when the origin is encrypted

When configuring FlexCache on a Cloud Volumes ONTAP system where the origin volume is encrypted, additional steps are required, to ensure that the FlexCache volume can properly access and cache the encrypted data.

Before you begin

- 1. Encryption setup:** Ensure that the source volume is fully encrypted and operational. For Cloud Volumes ONTAP systems, this involves integrating with cloud-specific key management services. For AWS, this typically means using AWS Key Management Service (KMS). For information, refer to [Manage keys with AWS Key Management Service](#). For Azure, you need to set up Azure Key Vault for NetApp Volume Encryption (NVE). For information, refer to [Manage keys with Azure Key Vault](#). For Google Cloud, it is Google Cloud Key Management Service. For information, refer to [Manage keys with Google's Cloud Key Management Service](#).
- 2. Key management services:** Before creating a FlexCache volume, verify that the key management services are configured correctly on the Cloud Volumes ONTAP system. This configuration is essential for the FlexCache volume to decrypt the data from the origin volume.
- 3. Licensing:** Confirm that a valid FlexCache license is available and activated on the Cloud Volumes ONTAP system.
- 4. ONTAP version:** Ensure that the ONTAP version of your Cloud Volumes ONTAP system supports FlexCache with encrypted volumes. Refer to the latest [ONTAP release notes](#) or compatibility matrix for more information.
- 5. Network Configuration:** Ensure that the network configuration allows for seamless communication between the origin volume and the FlexCache volume. This includes proper routing and DNS resolution in a cloud environment.

Steps

Create a FlexCache volume on your Cloud Volumes ONTAP system with an encrypted source volume. For detailed steps and additional considerations, refer to the following sections:

- [FlexCache Volumes for Faster Data Access Power Guide](#)
- [Creating FlexCache volumes in System Manager](#)

Copyright information

Copyright © 2026 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.