



## **Tier on-premises data to the cloud**

### NetApp Cloud Tiering

NetApp  
November 10, 2025

# Table of Contents

Tier on-premises data to the cloud .....	1
Tier data from on-premises ONTAP clusters to Amazon S3 in NetApp Cloud Tiering .....	1
Quick start .....	1
Network diagrams for connection options .....	2
Prepare your Console agent .....	3
Prepare your ONTAP cluster .....	4
Prepare your AWS environment .....	5
Tier inactive data from your first cluster to Amazon S3 .....	8
Tier data from on-premises ONTAP clusters to Azure Blob storage in NetApp Cloud Tiering .....	12
Quick start .....	12
Requirements .....	12
Tier inactive data from your first cluster to Azure Blob storage .....	15
Tier data from on-premises ONTAP clusters to Google Cloud Storage in NetApp Cloud Tiering .....	18
Quick start .....	18
Requirements .....	19
Tier inactive data from your first cluster to Google Cloud Storage .....	21
Tiering data from on-premises ONTAP clusters to StorageGRID in NetApp Cloud Tiering .....	24
Quick start .....	24
Requirements .....	25
Tier inactive data from your first cluster to StorageGRID .....	27
Tier data from on-premises ONTAP clusters to S3 object storage in NetApp Cloud Tiering .....	29
Quick start .....	29
Requirements .....	30
Tiering inactive data from your first cluster to S3-compatible object storage .....	33

# Tier on-premises data to the cloud

## Tier data from on-premises ONTAP clusters to Amazon S3 in NetApp Cloud Tiering

Free space on your on-premises ONTAP clusters by tiering inactive data to Amazon S3 in NetApp Cloud Tiering.

### Quick start

Get started quickly by following these steps. Details for each step are provided in the following sections in this topic.

1

#### Identify the configuration method you'll use

Choose whether you'll connect your on-premises ONTAP cluster directly to AWS S3 over the public internet, or whether you'll use a VPN or AWS Direct Connect and route traffic through a private VPC Endpoint interface to AWS S3.

[See the available connection methods.](#)

2

#### Prepare your Console Agent

If you already have the Console agent deployed in your AWS VPC or on your premises, then you're all set. If not, then you'll need to create the agent to tier ONTAP data to AWS S3 storage. You'll also need to customize network settings for the agent so that it can connect to AWS S3.

[See how to create a agent and how to define required network settings.](#)

3

#### Prepare your on-premises ONTAP cluster

Discover your ONTAP cluster in the NetApp Console, verify that the cluster meets minimum requirements, and customize network settings so the cluster can connect to AWS S3.

[See how to get your on-premises ONTAP cluster ready.](#)

4

#### Prepare Amazon S3 as your tiering target

Set up permissions for the agent to create and manage the S3 bucket. You'll also need to set up permissions for the on-premises ONTAP cluster so it can read and write data to the S3 bucket.

[See how to set up permissions for the agent and for your on-premises cluster.](#)

5

#### Enable Cloud Tiering on the system

Select an on-premises system, select **Enable** for the Cloud Tiering service, and follow the prompts to tier data to Amazon S3.

See how to enable Tiering for your volumes.

6

## Set up licensing

After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP Cloud Tiering BYOL license, or a combination of both:

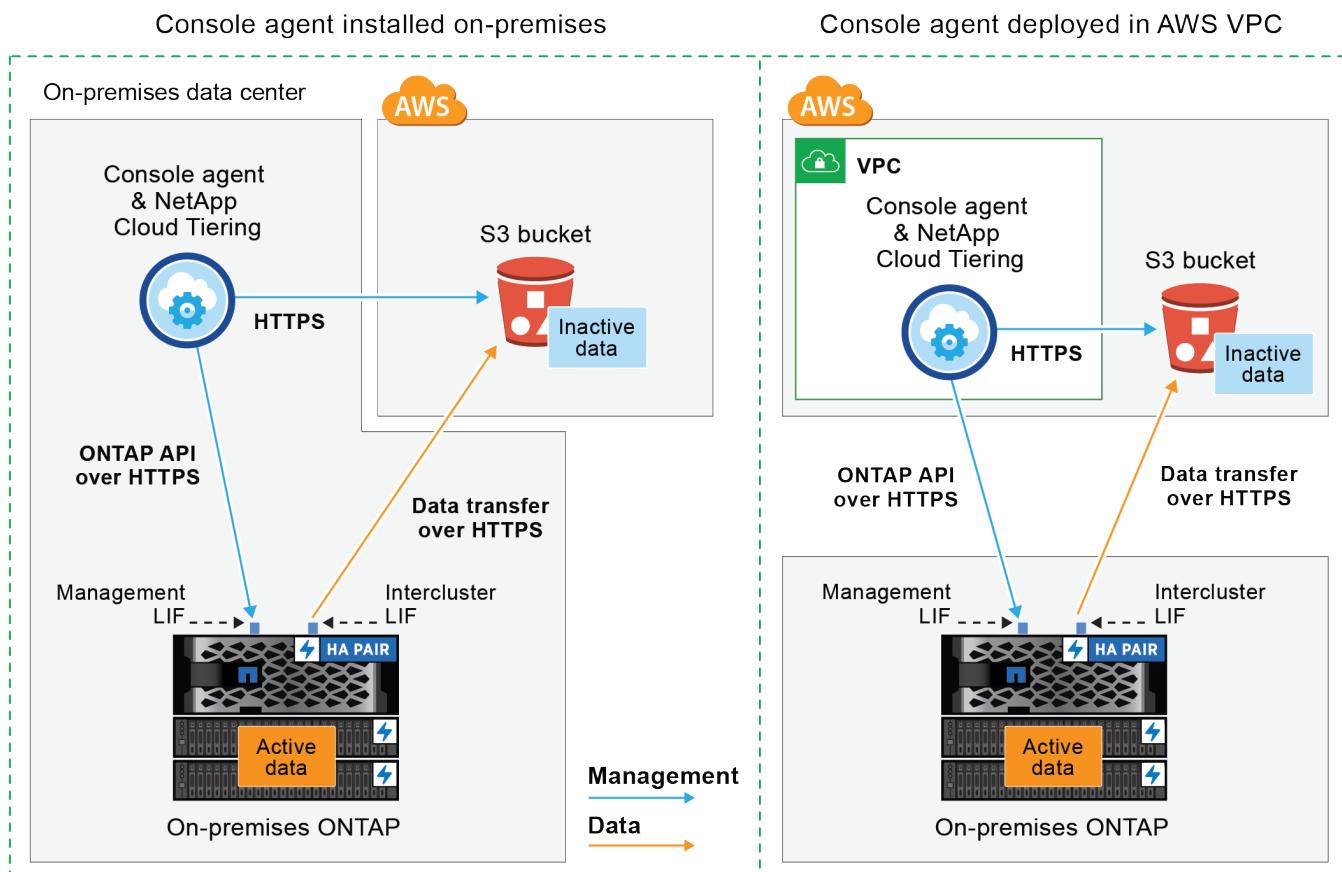
- To subscribe from the AWS Marketplace, [go to the Marketplace offering](#), select **Subscribe**, and then follow the prompts.
- To pay using a Cloud Tiering BYOL license, [contact us if you need to purchase one](#), and then [add it to the NetApp Console](#).

## Network diagrams for connection options

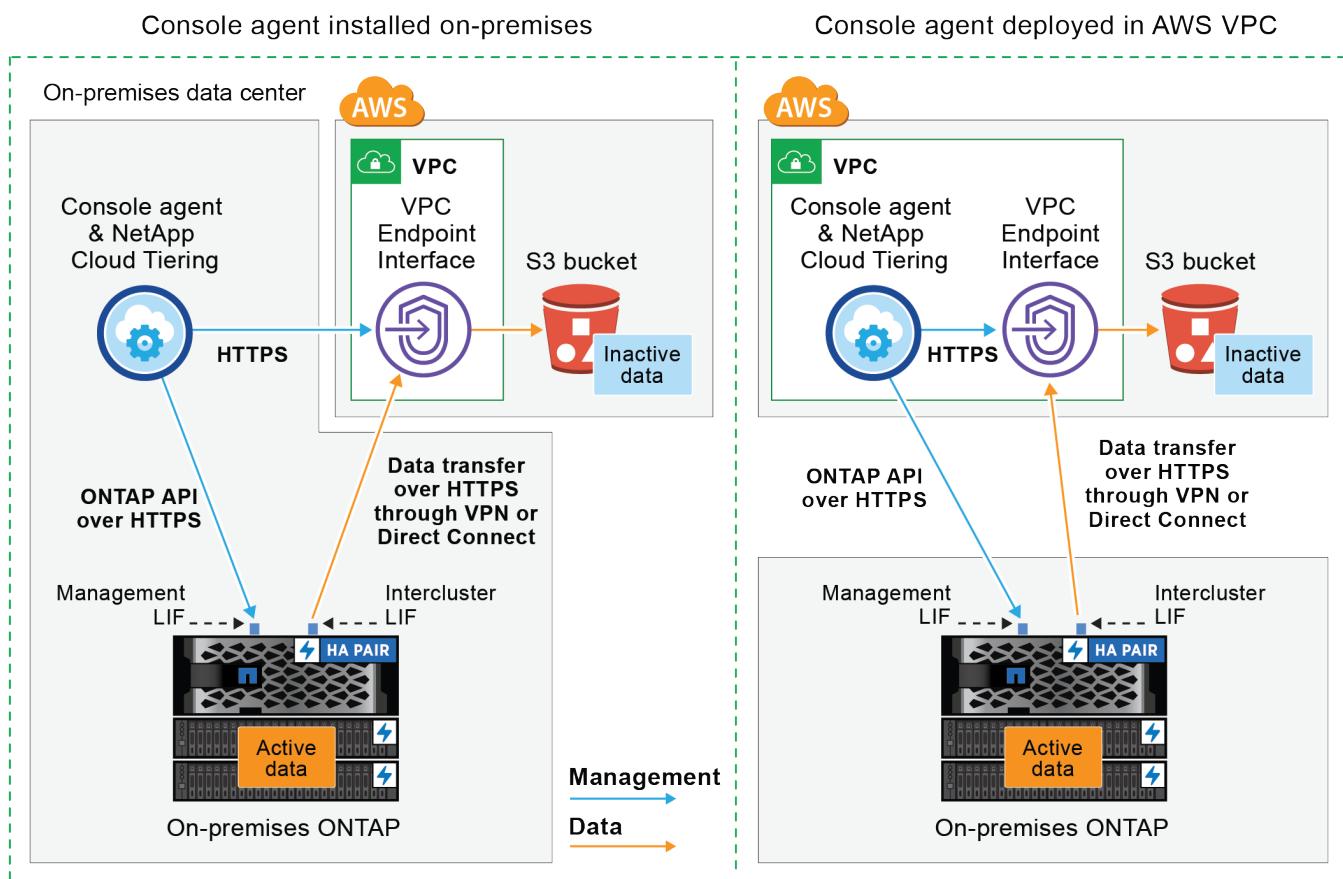
There are two connection methods you can use when configuring tiering from on-premises ONTAP systems to AWS S3.

- Public connection - Directly connect the ONTAP system to AWS S3 using a public S3 endpoint.
- Private connection - Use a VPN or AWS Direct Connect and route traffic through a VPC Endpoint interface that uses a private IP address.

The following diagram shows the **public connection** method and the connections that you need to prepare between the components. You can use the Console agent that you've installed on your premises, or an agent that you've deployed in the AWS VPC.



The following diagram shows the **private connection** method and the connections that you need to prepare between the components. You can use the Console agent that you've installed on your premises, or an agent that you've deployed in the AWS VPC.



Communication between an agent and S3 is for object storage setup only.

## Prepare your Console agent

The agent enables tiering capabilities from the NetApp Console. An agent is required to tier your inactive ONTAP data.

### Create or switch agents

If you already have an agent deployed in your AWS VPC or on your premises, then you're all set. If not, then you'll need to create an agent in either of those locations to tier ONTAP data to AWS S3 storage. You can't use an agent that's deployed in another cloud provider.

- [Learn about Console agents](#)
- [Deploying a agent in AWS](#)
- [Installing an agent on a Linux host](#)

### Agent networking requirements

- Ensure that the network where the agent is installed enables the following connections:
  - An HTTPS connection over port 443 to the Cloud Tiering service and to your S3 object storage ([see](#)

the list of endpoints)

- An HTTPS connection over port 443 to your ONTAP cluster management LIF
- Ensure that the agent has permissions to manage the S3 bucket
- If you have a Direct Connect or VPN connection from your ONTAP cluster to the VPC, and you want communication between the agent and S3 to stay in your AWS internal network (a **private** connection), you'll need to enable a VPC Endpoint interface to S3. [See how to set up a VPC endpoint interface.](#)

## Prepare your ONTAP cluster

Your ONTAP clusters must meet the following requirements when tiering data to Amazon S3.

### ONTAP requirements

#### Supported ONTAP platforms

- When using ONTAP 9.8 and later: You can tier data from AFF systems, or FAS systems with all-SSD aggregates or all-HDD aggregates.
- When using ONTAP 9.7 and earlier: You can tier data from AFF systems, or FAS systems with all-SSD aggregates.

#### Supported ONTAP versions

- ONTAP 9.2 or later
- ONTAP 9.7 or later is required if you plan to use an AWS PrivateLink connection to object storage

#### Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. Refer to the ONTAP documentation for [functionality or features not supported by FabricPool](#).



Cloud Tiering supports FlexGroup volumes starting with ONTAP 9.5. Setup works the same as any other volume.

### Cluster networking requirements

- The cluster requires an inbound HTTPS connection from the Console agent to the cluster management LIF. A connection between the cluster and Cloud Tiering is not required.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to tier. These intercluster LIFs must be able to access the object store.

The cluster initiates an outbound HTTPS connection over port 443 from the intercluster LIFs to Amazon S3 storage for tiering operations. ONTAP reads and writes data to and from object storage — the object storage never initiates, it just responds.

- The intercluster LIFs must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces.](#)

When you set up Cloud Tiering, you are prompted for the IPspace to use. You should choose the IPspace that these LIFs are associated with. That might be the "Default" IPspace or a custom IPspace that you created.

If you are using a different IPspace than "Default", then you might need to create a static route to get access to the object storage.

All intercluster LIFs within the IPspace must have access to the object store. If you can't configure this for the current IPspace, then you'll need to create a dedicated IPspace where all intercluster LIFs have access to the object store.

- If you are using a Private VPC Interface Endpoint in AWS for the S3 connection, then in order for HTTPS/443 to be used, you'll need to load the S3 endpoint certificate into the ONTAP cluster. [See how to set up a VPC endpoint interface and load the S3 certificate.](#)
- [Ensure that your ONTAP cluster has permissions to access the S3 bucket.](#)

## Discover your ONTAP cluster in NetApp Console

You need to discover your on-premises ONTAP cluster in the NetApp Console before you can start tiering cold data to object storage. You'll need to know the cluster management IP address and the password for the admin user account to add the cluster.

[Learn how to discover a cluster.](#)

## Prepare your AWS environment

When you set up data tiering for a new cluster, you're prompted whether you want the service to create an S3 bucket or if you want to select an existing S3 bucket in the AWS account where the agent is set up. The AWS account must have permissions and an access key that you can enter in Cloud Tiering. The ONTAP cluster uses the access key to tier data in and out of S3.

By default, Cloud tiering creates the bucket for you. If you want to use your own bucket, you can create one before you start the tiering activation wizard and then select that bucket in the wizard. [See how to create S3 buckets from the NetApp Console.](#) The bucket must be used exclusively for storing inactive data from your volumes - it cannot be used for any other purpose. The S3 bucket must be in a [region that supports Cloud Tiering.](#)

 If you are planning to configure Cloud Tiering to use a lower cost storage class where your tiered data will transition to after a certain number of days, you must not select any lifecycle rules when setting up the bucket in your AWS account. Cloud Tiering manages the lifecycle transitions.

## Set up S3 permissions

You'll need to configure two sets of permissions:

- Permissions for the agent so it can create and manage the S3 bucket.
- Permissions for the on-premises ONTAP cluster so it can read and write data to the S3 bucket.

## Steps

### 1. Console agent permissions:

- Confirm that [these S3 permissions](#) are part of the IAM role that provides the agent with permissions. They should have been included by default when you first deployed the agent. If not, you'll need to add any missing permissions. See the [AWS Documentation: Editing IAM policies](#) for instructions.
- The default bucket that Cloud Tiering creates has a prefix of "fabric-pool". If you want to use a different prefix for your bucket, you'll need to customize the permissions with the name you want to use. In the

S3 permissions you'll see a line "Resource": ["arn:aws:s3:::fabric-pool\*"]. You'll need to change "fabric-pool" to the prefix that you want to use. For example, if you want to use "tiering-1" as the prefix for your buckets, you'll change this line to "Resource": ["arn:aws:s3:::tiering-1\*"].

If you want to use a different prefix for buckets that you'll use for additional clusters in this same NetApp Console organization, you can add another line with the prefix for other buckets. For example:

```
"Resource": ["arn:aws:s3:::tiering-1*"]
"Resource": ["arn:aws:s3:::tiering-2*"]
```

If you are creating your own bucket and do not use a standard prefix, you should change this line to "Resource": ["arn:aws:s3:::\*"] so that any bucket is recognized. However, this may expose all your buckets instead of those you have designed to hold inactive data from your volumes.

## 2. Cluster permissions:

- When you activate the service, the Tiering wizard will prompt you to enter an access key and secret key. These credentials are passed to the ONTAP cluster so that ONTAP can tier data to the S3 bucket. For that, you'll need to create an IAM user with the following permissions:

```
"s3>ListAllMyBuckets",
"s3>ListBucket",
"s3:GetBucketLocation",
"s3:GetObject",
"s3:PutObject",
"s3:DeleteObject"
```

See the [AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#) for details.

## 3. Create or locate the access key.

Cloud Tiering passes the access key on to the ONTAP cluster. The credentials are not stored in the Cloud Tiering service.

[AWS Documentation: Managing Access Keys for IAM Users](#)

## Configure your system for a private connection using a VPC endpoint interface

If you plan to use a standard public internet connection, then all the permissions are set by the agent and there is nothing else you need to do. This type of connection is shown in the [first diagram above](#).

If you want to have a more secure connection over the internet from your on-premises data center to the VPC, there's an option to select an AWS PrivateLink connection in the Tiering activation wizard. It's required if you plan to use a VPN or AWS Direct Connect to connect your on-premises system through a VPC Endpoint interface that uses a private IP address. This type of connection is shown in the [second diagram above](#).

- Create an Interface endpoint configuration using the Amazon VPC Console or the command line. [See details about using AWS PrivateLink for Amazon S3](#).
- Modify the security group configuration that's associated with the agent. You must change the policy to "Custom" (from "Full Access"), and you must [add the required S3 agent permissions](#) as shown earlier.

Search for services, features, marketplace products, and docs [Option+S]

hendr@xcp-liveng

Select security groups ▾

Filter by tags and attributes or search by keyword

Group ID	Group Name	VPC ID	Description	Owner ID
sg-012e998e...	pvt-link-testO...	vpc-09c1028...	EC2-VPC	NetApp OCCM Instance ... 464

**Policy\***

Full Access - Allow access by any user or service within the VPC using credentials from any AWS accounts to any resources in this AWS service. All policies — IAM user policies, VPC endpoint policies, and AWS service-specific policies (e.g. Amazon S3 bucket policies, any S3 ACL policies) — must grant the necessary permissions for access to succeed.

Custom

Use the [policy creation tool](#) to generate a policy, then paste the generated policy below.

```
"s3:GetObject",
"s3:DeleteObject",
"s3>ListBucket",
"s3>ListAllMyBuckets",
"s3:GetBucketTagging",
"s3:GetBucketLocation",
```

If you're using port 80 (HTTP) for communication to the private endpoint, you're all set. You can enable Cloud Tiering on the cluster now.

If you're using port 443 (HTTPS) for communication to the private endpoint, you must copy the certificate from the VPC S3 endpoint and add it to your ONTAP cluster, as shown in the next 4 steps.

### 3. Obtain the DNS name of the endpoint from the AWS Console.

New VPC Experience Tell us what you think

VIRTUAL PRIVATE CLOUD

Your VPCs

Subnets

Route Tables

Internet Gateways

Egress Only Internet Gateways

Carrier Gateways

DHCP Options Sets

Elastic IPs

Managed Prefix Lists

**Endpoints**

Endpoint Services

NAT Gateways

Peering Connections

SECURITY

Network ACLs

Security Groups

NETWORK ANALYSIS

Reachability Analyzer

Network Access Analyzer

VPC > Endpoints > vpce-0ff5c15df7e00fbab

vpce-0ff5c15df7e00fbab / RN-VPC-Endpoint

**Details**

Endpoint ID: vpce-0ff5c15df7e00fbab Status: Available

VPC ID: vpc-673ae603 [VPC EMEA] Status message:

DNS names:

- \*.vpce-0ff5c15df7e00fbab-yxs7it8v.s3.us-west-2.amazonaws.com (Z1YSA3EXCYUU9Z)
- \*.vpce-0ff5c15df7e00fbab-yxs7it8v-us-west-2.s3.us-west-2.amazonaws.com (Z1YSA3EXCYUU9Z)
- \*.vpce-0ff5c15df7e00fbab-yxs7it8v-us-west-2a.s3.us-west-2.amazonaws.com (Z1YSA3EXCYUU9Z)
- \*.vpce-0ff5c15df7e00fbab-yxs7it8v-us-west-2c.s3.us-west-2.amazonaws.com (Z1YSA3EXCYUU9Z)

Subnets Security Groups Notification Policy Monitoring Tags

### 4. Obtain the certificate from the VPC S3 endpoint. You do this by [logging into the VM that hosts the agent](#) and running the following command. When entering the DNS name of the endpoint, add "bucket" to the beginning, replacing the "\*":

```
[ec2-user@ip-10-160-4-68 ~]$ openssl s_client -connect bucket.vpce-0ff5c15df7e00fbab-yxs7lt8v.s3.us-west-2.vpce.amazonaws.com:443 -showcerts
```

5. From the output of this command, copy the data for the S3 certificate (all data between, and including, the BEGIN / END CERTIFICATE tags):

```
Certificate chain
0 s:/CN=s3.us-west-2.amazonaws.com`  
i:/C=US/O=Amazon/OU=Server CA 1B/CN=Amazon  
-----BEGIN CERTIFICATE-----  
MIIM6zCCC9OgAwIBAgIQA7MGJ4FaDBR8uL0KR3o1tTANBgkqhkiG9w0BAQsFADBG  
...  
...  
GqvbOz/oO2NWLLFCqI+xmkLcMiPrZy+/6Af+HH2mLCM4EsI2b+IpBmPkriWnnxo=  
-----END CERTIFICATE-----
```

6. Log into the ONTAP cluster CLI and apply the certificate you copied using the following command (substitute your own storage VM name):

```
cluster1::> security certificate install -vserver <svm_name> -type  
server-ca  
Please enter Certificate: Press <Enter> when done
```

## Tier inactive data from your first cluster to Amazon S3

After you prepare your AWS environment, start tiering inactive data from your first cluster.

### What you'll need

- [A managed on-premises system in the Console](#).
- An AWS access key for an IAM user who has the required S3 permissions.

### Steps

1. Select the on-premises ONTAP system.
2. Click **Enable** for Cloud Tiering from the right panel.

If the Amazon S3 tiering destination exists as a system on the Systems page, you can drag the cluster onto the system to initiate the setup wizard.

sgtest1  
On

DETAILS  
On-Premises ONTAP

SERVICES

	Replication	■ Off	<span>Enable</span>	
	Tiering	■ Off	<span>Enable</span>	

3. **Define Object Storage Name:** Enter a name for this object storage. It must be unique from any other object storage you may be using with aggregates on this cluster.
4. **Select Provider:** Select **Amazon Web Services** and select **Continue**.

Tiering Setup

- 1 S3 Bucket
- 2 Storage Class Life Cycle
- 3 Credentials
- 4 Networking

Provider Settings

AWS Account ID

awsAccount1

S3 Bucket

The default bucket name prefix is 'fabric-pool'. If you want to use a different prefix to identify buckets for tiered data, you must revise the AWS Connector permissions to include the prefix. [Learn more](#)

Select bucket type  Add a new S3 bucket  Select an existing S3 bucket

The default bucket name prefix is 'fabric-pool'. If you choose to change it, you must add permissions for the new defined prefix of bucket name otherwise actions will fail. [Learn more](#)

Prefix	Bucket name	Bucket region
fabric-pool		Region

5. Complete the sections in the **Tiering Setup** page:
  - a. **S3 Bucket:** Add a new S3 bucket or select an existing S3 bucket, select the bucket region, and select **Continue**.

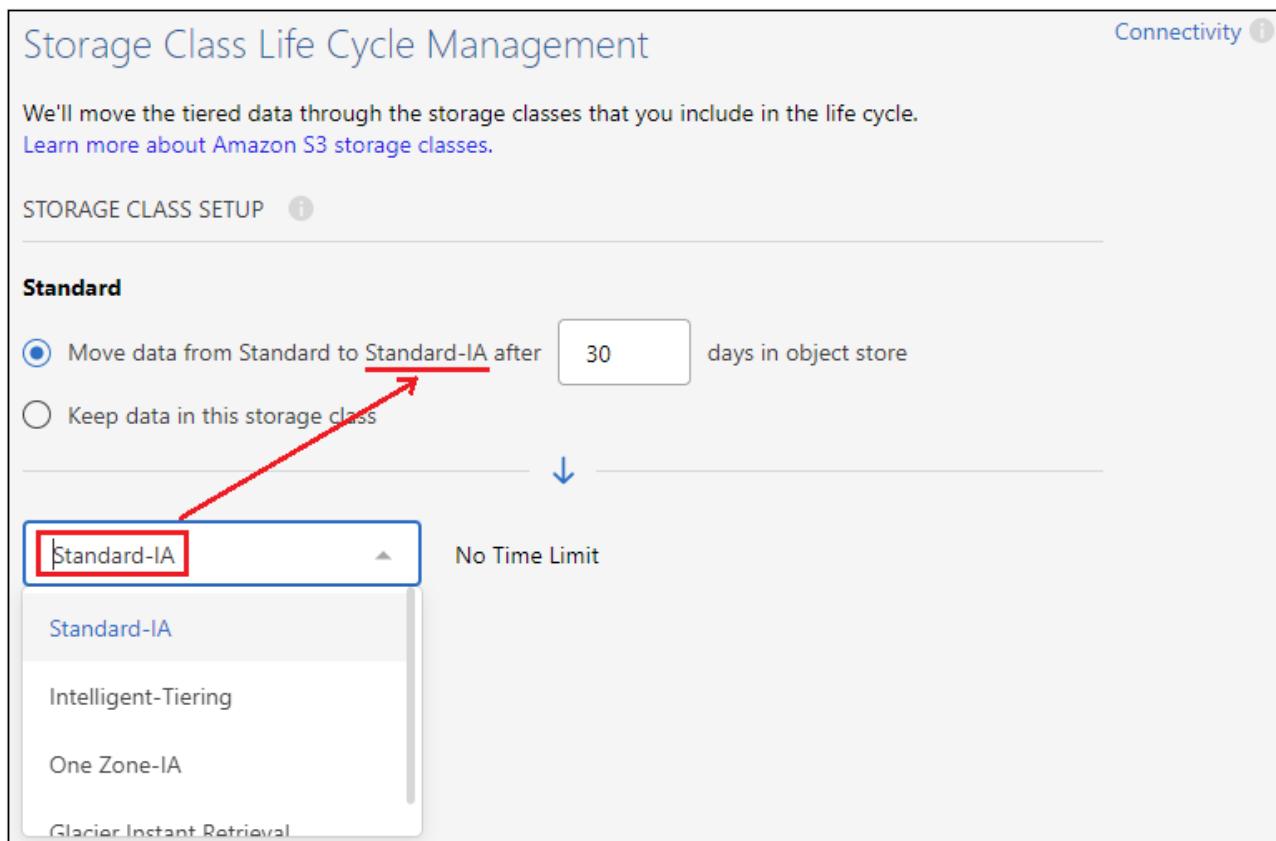
When using an on-premises agent, you must enter the AWS Account ID that provides access to the existing S3 bucket or new S3 bucket that will be created.

The *fabric-pool* prefix is used by default because the IAM policy for the agent enables the instance to perform S3 actions on buckets named with that exact prefix. For example, you could name the S3 bucket *fabric-pool-AFF1*, where AFF1 is the name of the cluster. You can define the prefix for the buckets used for tiering as well. See [setting up S3 permissions](#) to make sure you have AWS permissions that recognize any custom prefix you plan to use.

b. **Storage Class:** Cloud Tiering manages the lifecycle transitions of your tiered data. Data starts in the *Standard* class, but you can create a rule to apply a different storage class to the data after a certain number of days.

Select the S3 storage class that you want to transition the tiered data to and the number of days before the data is assigned to that class, and select **Continue**. For example, the screenshot below shows that tiered data is assigned to the *Standard-IA* class from the *Standard* class after 45 days in object storage.

If you choose **Keep data in this storage class**, then the data remains in the *Standard* storage class and no rules are applied. [See supported storage classes](#).



Storage Class Life Cycle Management

We'll move the tiered data through the storage classes that you include in the life cycle.  
[Learn more about Amazon S3 storage classes](#).

STORAGE CLASS SETUP

**Standard**

Move data from Standard to Standard-IA after  days in object store

Keep data in this storage class

↓

**Standard-IA**

No Time Limit

Standard-IA

Intelligent-Tiering

One Zone-IA

Glacier Instant Retrieval

Note that the lifecycle rule is applied to all objects in the selected bucket.

c. **Credentials:** Enter the access key ID and secret key for an IAM user who has the required S3 permissions, and select **Continue**.

The IAM user must be in the same AWS account as the bucket that you selected or created on the **S3 Bucket** page.

d. **Networking:** Enter the networking details and select **Continue**.

Select the IPspace in the ONTAP cluster where the volumes you want to tier reside. The intercluster LIFs for this IPspace must have outbound internet access so that they can connect to your cloud provider's object storage.

Optionally, choose whether you'll use an AWS PrivateLink that you have previously configured. [See the setup information above](#). A dialog box is displayed to help guide you through the endpoint configuration.

You can also set the network bandwidth available to upload inactive data to object storage by defining the "Maximum transfer rate". Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.

6. On the *Tier Volumes* page, select the volumes that you want to configure tiering for and launch the *Tiering Policy* page:

- To select all volumes, check the box in the title row ( **Volume Name**) and select **Configure volumes**.
- To select multiple volumes, check the box for each volume ( **Volume\_1**) and select **Configure volumes**.
- To select a single volume, select the row (or icon) for the volume.

Tier Volumes								
Volumes (16) <small>(2 selected)</small>	Volume Name	Aggregate/s Name	SVM Name	Node/s Name	Volume Size	Snapshot Size	Cold Data...	Tier Status
<input type="checkbox"/>	volume_1	aggr-1	svm_1	volume_1_node	20 TB	0 B	10 TB   10 %	Not Tiered Volume
<input checked="" type="checkbox"/>	volume_2	aggr-1	svm_1	volume_2_node	15 TB	205 GB	2.8 TB   70 %	Not Tiered Volume
<input checked="" type="checkbox"/>	volume_3	aggr-1	svm_1	volume_3_node	8 TB	0 B	716.8 GB   70 %	Not Tiered Volume
<input type="checkbox"/>	volume_4	aggr-1	svm_1	volume_4_node	3 TB	3 GB	35 GB   70 %	Not Tiered Volume

7. In the *Tiering Policy* dialog, select a tiering policy, optionally adjust the cooling days for the selected volumes, and select **Apply**.

[Learn more about volume tiering policies and cooling days.](#)

Select volume tiering policy

The policy is applied to the volume, or volumes, you selected in the previous page.

Volume_1	5 TiB	512 GiB   25% Cold data <small>i</small>	0 Bytes	10 TiB Used size <small>i</small>
Select tiering policy				
<input type="radio"/> No policy	<input type="radio"/> Cold snapshots <small>i</small>	<input checked="" type="radio"/> Cold user data & snapshots <small>i</small>	<input type="radio"/> All user data <small>i</small>	
Adjust cooling days <small>i</small> 62 Days				

## Result

You've successfully set up data tiering from volumes on the cluster to S3 object storage.

## What's next?

[Be sure to subscribe to the Cloud Tiering service.](#)

You can review information about the active and inactive data on the cluster. [Learn more about managing your tiering settings.](#)

You can also create additional object storage in cases where you may want to tier data from certain aggregates on a cluster to different object stores. Or if you plan to use FabricPool Mirroring where your tiered data is

replicated to an additional object store. [Learn more about managing object stores.](#)

# Tier data from on-premises ONTAP clusters to Azure Blob storage in NetApp Cloud Tiering

Free space on your on-premises ONTAP clusters by tiering inactive data to Azure Blob storage.

## Quick start

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

1

### Prepare to tier data to Azure Blob storage

You need the following:

- A source on-premises ONTAP cluster that's running ONTAP 9.4 or later that you have added to the NetApp Console, and an HTTPS connection to Azure Blob storage. [Learn how to discover a cluster.](#)
- A Console agent installed in an Azure VNet or on your premises.
- Networking for an agent that enables an outbound HTTPS connection to the ONTAP cluster in your data center, to Azure storage, and to the Cloud Tiering service.

2

### Set up tiering

In the NetApp Console, select an on-premises ONTAP system, select **Enable** for the Tiering service, and follow the prompts to tier data to Azure Blob storage.

3

### Set up licensing

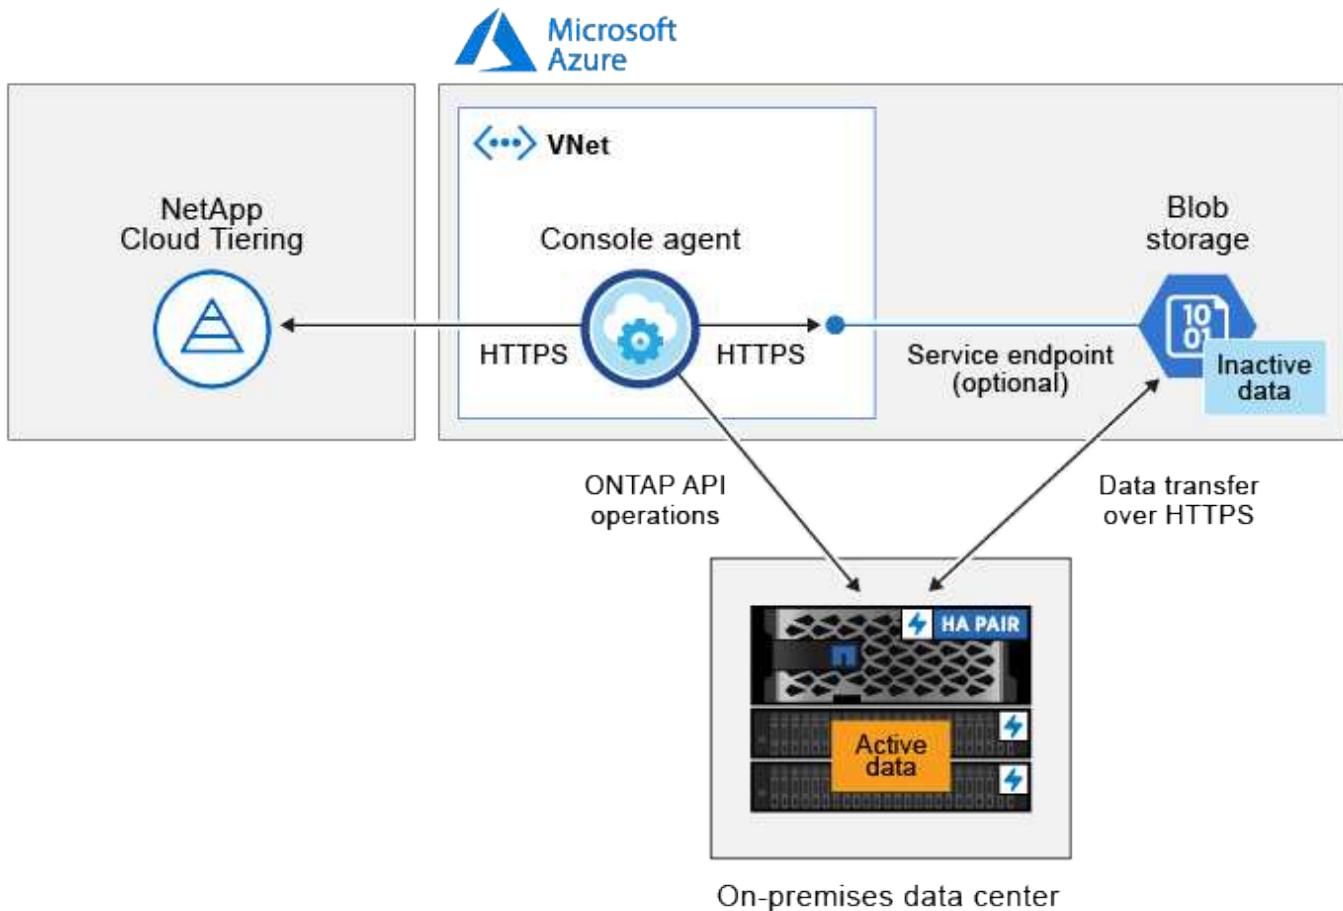
After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP Cloud Tiering BYOL license, or a combination of both:

- To subscribe from the Azure Marketplace, [go to the Marketplace offering](#), select **Subscribe**, and then follow the prompts.
- To pay using a Cloud Tiering BYOL license, [contact us if you need to purchase one](#), and then [add it to the NetApp Console](#).

## Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between the Console agent and Blob storage is for object storage setup only. The agent can reside on your premises, instead of in the cloud.

## Prepare your ONTAP clusters

Your ONTAP clusters must meet the following requirements when tiering data to Azure Blob storage.

### Supported ONTAP platforms

- When using ONTAP 9.8 and later: You can tier data from AFF systems, or FAS systems with all-SSD aggregates or all-HDD aggregates.
- When using ONTAP 9.7 and earlier: You can tier data from AFF systems, or FAS systems with all-SSD aggregates.

### Supported ONTAP version

ONTAP 9.4 or later

### Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 to Azure Blob storage.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

Although ExpressRoute provides better performance and lower data transfer charges, it's not required between the ONTAP cluster and Azure Blob storage. But doing so is the recommended best practice.

- An inbound connection is required from the agent, which can reside in an Azure VNet or on your premises.

A connection between the cluster and the Cloud Tiering service is not required.

- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to tier. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

When you set up data tiering, Cloud Tiering prompts you for the *IPspace* to use. You should choose the *IPspace* that each LIF is associated with. That might be the "Default" *IPspace* or a custom *IPspace* that you created. Learn more about [LIFs](#) and [IPspaces](#).

## Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. Refer to the ONTAP documentation for [functionality or features not supported by FabricPool](#).



Cloud Tiering supports FlexGroup volumes, starting with ONTAP 9.5. Setup works the same as any other volume.

## Discover an ONTAP cluster

You need to add an on-premises ONTAP system to the NetApp Console before you can start tiering cold data.

[Learn how to discover a cluster.](#)

## Create or switch agents

An agent is required to tier data to the cloud. When tiering data to Azure Blob storage, you can use an agent that's in an Azure VNet or in your premises. You'll either need to create a new agent or make sure that the currently selected agent resides in Azure or on-premises.

- [Learn about agents](#)
- [Deploying an agent in Azure](#)
- [Installing an agent on a Linux host](#)

## Verify that you have the necessary agent permissions

If you created the Console agent using version 3.9.25 or greater, then you're all set. The custom role that provides the permissions that an agent needs to manage resources and processes within your Azure network will be set up by default. See the [required custom role permissions](#) and the [specific permissions required for Cloud Tiering](#).

If you created the agent using an earlier version, then you'll need to edit the permission list for the Azure account to add any missing permissions.

## Prepare networking for the Console agent

Ensure that the Console agent has the required networking connections. The agent can be installed on-premises or in Azure.

## Steps

1. Ensure that the network where the agent is installed enables the following connections:
  - An HTTPS connection over port 443 to the Cloud Tiering service and to your Azure Blob object storage ([see the list of endpoints](#))
  - An HTTPS connection over port 443 to your ONTAP cluster management LIF
2. If needed, enable a VNet service endpoint to Azure storage.

A VNet service endpoint to Azure storage is recommended if you have an ExpressRoute or VPN connection from your ONTAP cluster to the VNet and you want communication between the agent and Blob storage to stay in your virtual private network.

## Prepare Azure Blob storage

When you set up tiering, you need to identify the resource group you want to use, and the storage account and Azure container that belong to the resource group. A storage account enables Cloud Tiering to authenticate and access the Blob container used for data tiering.

Cloud Tiering supports tiering to any storage account in any region that can be accessed via the agent.

Cloud Tiering supports only the General Purpose v2 and Premium Block Blob types of storage accounts.



If you are planning to configure Cloud Tiering to use a lower cost access tier where your tiered data will transition to after a certain number of days, you must not select any lifecycle rules when setting up the container in your Azure account. Cloud Tiering manages the lifecycle transitions.

## Tier inactive data from your first cluster to Azure Blob storage

After you prepare your Azure environment, start tiering inactive data from your first cluster.

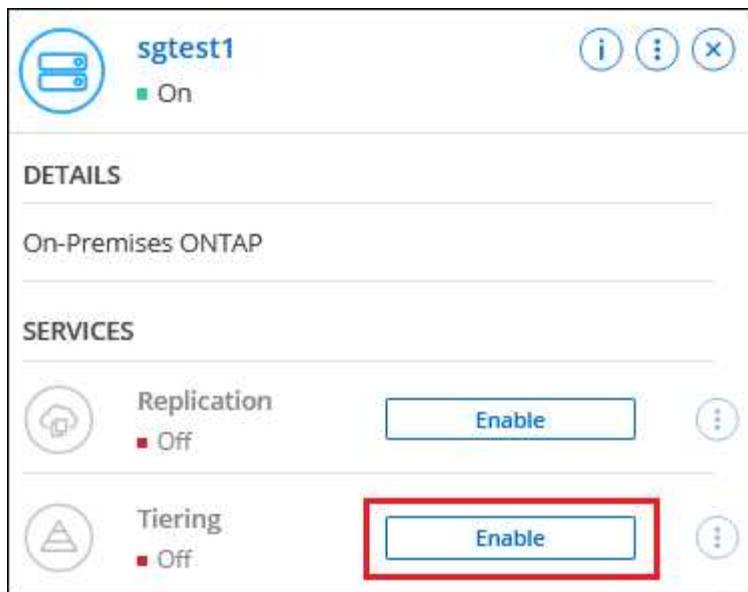
### What you'll need

[An on-premises ONTAP system to the NetApp Console.](#)

### Steps

1. Select the on-premises ONTAP system.
2. Click **Enable** for the Tiering service from the right panel.

If the Azure Blob tiering destination exists as a system on the Systems page, you can drag the cluster onto the Azure Blob system to initiate the setup wizard.



3. **Define Object Storage Name:** Enter a name for this object storage. It must be unique from any other object storage you may be using with aggregates on this cluster.
4. **Select Provider:** Select **Microsoft Azure** and select **Continue**.
5. Complete the steps on the **Create Object Storage** pages:
  - a. **Resource Group:** Select a resource group where an existing container is managed, or where you'd like to create a new container for tiered data, and select **Continue**.

When using an on-premises agent, you must enter the Azure Subscription that provides access to the resource group.
  - b. **Azure Container:** Select the radio button to either add a new Blob container to a storage account or to use an existing container. Then select the storage account and choose the existing container, or enter the name for the new container. Then select **Continue**.

The storage accounts and containers that appear in this step belong to the resource group that you selected in the previous step.
  - c. **Access Tier Lifecycle:** Cloud Tiering manages the lifecycle transitions of your tiered data. Data starts in the *Hot* class, but you can create a rule to apply the *Cool* class to the data after a certain number of days.

Select the access tier that you want to transition the tiered data to and the number of days before the data is assigned to that tier, and select **Continue**. For example, the screenshot below shows that tiered data is assigned to the *Cool* class from the *Hot* class after 45 days in object storage.

If you choose **Keep data in this access tier**, then the data remains in the *Hot* access tier and no rules are applied. [See supported access tiers](#).

## Access Tier Life Cycle Management

We'll move the tiered data through the access tiers that you include in the life cycle. [Learn more about Azure Blob storage access tiers.](#)

### ACCESS TIER SETUP

#### Hot

Move data from Hot to Cool after  days  
 Keep data in this storage class



#### Cool

No Time Limit 

Note that the lifecycle rule is applied to all blob containers in the selected storage account.

d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage, and select **Continue**.

Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.

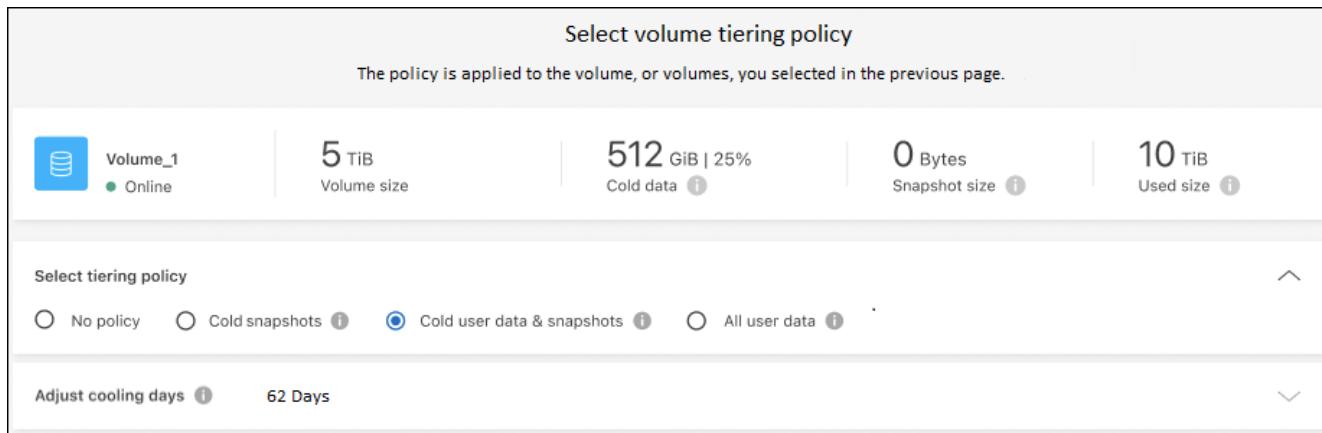
You can also set the network bandwidth available to upload inactive data to object storage by defining the "Maximum transfer rate". Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.

6. On the *Tier Volumes* page, select the volumes that you want to configure tiering for and launch the *Tiering Policy* page:

- To select all volumes, check the box in the title row (  ) and select **Configure volumes**.
- To select multiple volumes, check the box for each volume (  ) and select **Configure volumes**.
- To select a single volume, select the row (or  icon) for the volume.

7. In the *Tiering Policy* dialog, select a tiering policy, optionally adjust the cooling days for the selected volumes, and select **Apply**.

[Learn more about volume tiering policies and cooling days.](#)



Select volume tiering policy

The policy is applied to the volume, or volumes, you selected in the previous page.

Volume_1	5 TiB Volume size	512 GiB   25% Cold data	0 Bytes Snapshot size	10 TiB Used size
----------	----------------------	----------------------------	--------------------------	---------------------

Select tiering policy

No policy  Cold snapshots  Cold user data & snapshots  All user data

Adjust cooling days  62 Days

## Result

You've successfully set up data tiering from volumes on the cluster to Azure Blob object storage.

## What's next?

[Be sure to subscribe to the Cloud Tiering service.](#)

You can review information about the active and inactive data on the cluster. [Learn more about managing your tiering settings.](#)

You can also create additional object storage in cases where you may want to tier data from certain aggregates on a cluster to different object stores. Or if you plan to use FabricPool Mirroring where your tiered data is replicated to an additional object store. [Learn more about managing object stores.](#)

## Tier data from on-premises ONTAP clusters to Google Cloud Storage in NetApp Cloud Tiering

Free space on your on-premises ONTAP clusters by tiering inactive data to Google Cloud Storage in NetApp Cloud Tiering.

### Quick start

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

1

#### Prepare to tier data to Google Cloud Storage

You need the following:

- A source on-premises ONTAP cluster that's running ONTAP 9.6 or later that you have added to the NetApp Console, and a connection over a user-specified port to Google Cloud Storage. [Learn how to discover a cluster.](#)
- A service account that has the predefined Storage Admin role and storage access keys.
- A Console agent installed in a Google Cloud Platform VPC.
- Networking for the agent that enables an outbound HTTPS connection to the ONTAP cluster in your data center, to Google Cloud Storage, and to the Cloud Tiering service.

**2**

## Set up tiering

In the NetApp Console, select an on-premises system select **Enable** for the Tiering service, and follow the prompts to tier data to Google Cloud Storage.

**3**

## Set up licensing

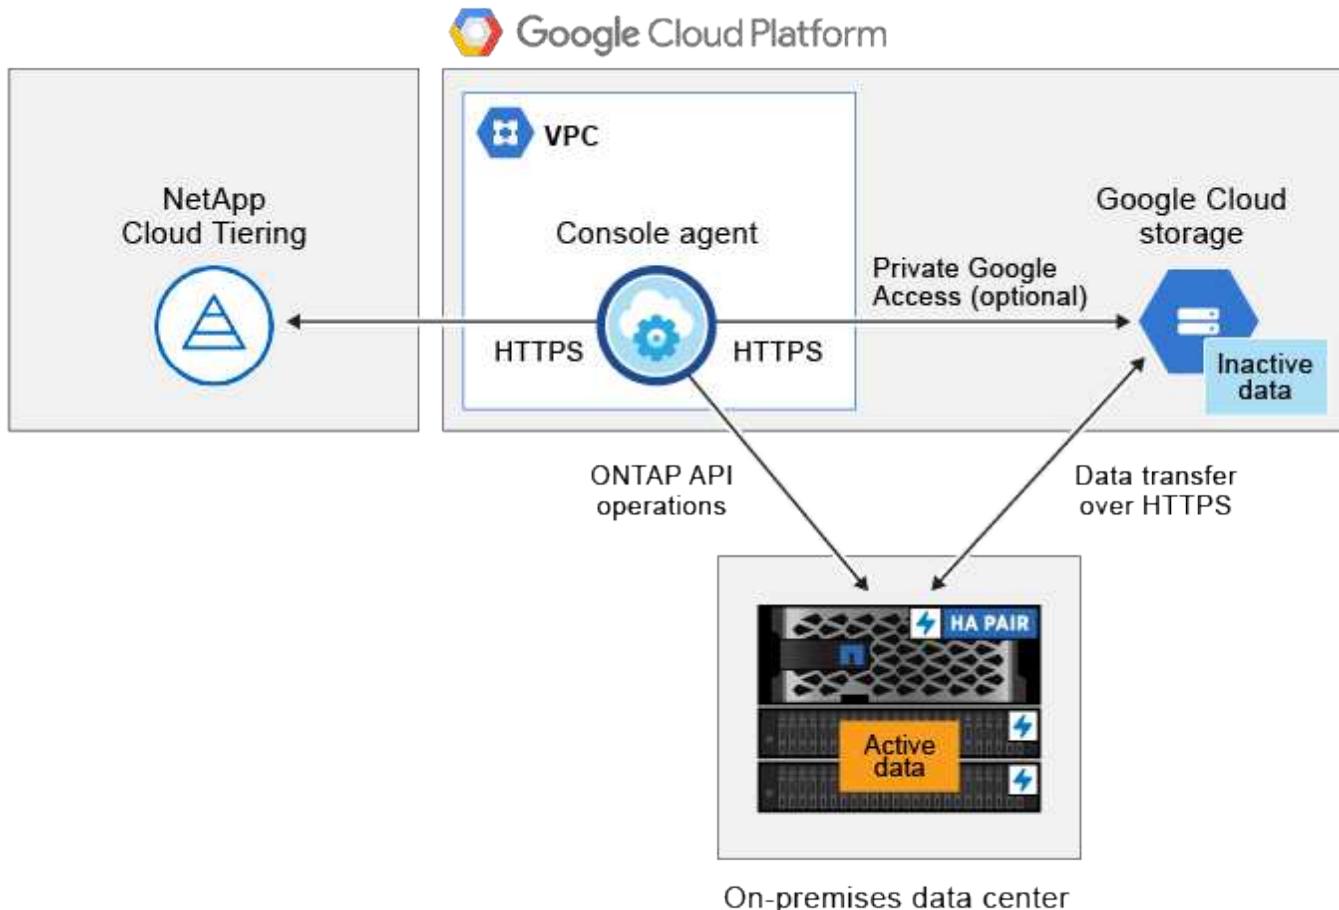
After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP Cloud Tiering BYOL license, or a combination of both:

- To subscribe from the Google Cloud marketplace, go to the [Marketplace offering](#), select **Subscribe**, and then follow the prompts.
- To pay using a Cloud Tiering BYOL license, [contact us if you need to purchase one](#), and then [add it to the NetApp Console](#).

## Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between the agent and Google Cloud Storage is for object storage setup only.

## Prepare your ONTAP clusters

Your ONTAP clusters must meet the following requirements when tiering data to Google Cloud Storage.

### Supported ONTAP platforms

- When using ONTAP 9.8 and later: You can tier data from AFF systems, or FAS systems with all-SSD aggregates or all-HDD aggregates.
- When using ONTAP 9.7 and earlier: You can tier data from AFF systems, or FAS systems with all-SSD aggregates.

### Supported ONTAP versions

ONTAP 9.6 or later

### Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 to Google Cloud Storage.  
ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.  
Although a Google Cloud Interconnect provides better performance and lower data transfer charges, it's not required between the ONTAP cluster and Google Cloud Storage. But doing so is the recommended best practice.
- An inbound connection is required from the agent, which resides in a Google Cloud Platform VPC.  
A connection between the cluster and the Cloud Tiering service is not required.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to tier. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

When you set up data tiering, Cloud Tiering prompts you for the *IPspace* to use. You should choose the *IPspace* that each LIF is associated with. That might be the "Default" *IPspace* or a custom *IPspace* that you created. Learn more about [LIFs](#) and [IPspaces](#).

### Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. Refer to the ONTAP documentation for [functionality or features not supported by FabricPool](#).



Cloud Tiering supports FlexGroup volumes. Setup works the same as any other volume.

### Discover an ONTAP cluster

You need to add your on-premises ONTAP system to the NetApp Console before you can start tiering cold data.

[Learn how to discover a cluster.](#)

### Create or switch Console agents

A Console agent is required to tier data to the cloud. When tiering data to Google Cloud Storage, an agent must be available in a Google Cloud Platform VPC. You'll either need to create a new agent or make sure that the currently selected agent resides in Google Cloud.

- [Learn about agents](#)
- [Deploying an agent in Google Cloud](#)

## Prepare networking for the Console agent

Ensure that the Console agent has the required networking connections.

### Steps

1. Ensure that the VPC where the agent is installed enables the following connections:
  - An HTTPS connection over port 443 to the Cloud Tiering service and to your Google Cloud Storage ([see the list of endpoints](#))
  - An HTTPS connection over port 443 to your ONTAP cluster management LIF
2. Optional: Enable Private Google Access on the subnet where you plan to deploy the agent.

**Private Google Access** is recommended if you have a direct connection from your ONTAP cluster to the VPC and you want communication between the agent and Google Cloud Storage to stay in your virtual private network. Note that Private Google Access works with VM instances that have only internal (private) IP addresses (no external IP addresses).

## Prepare Google Cloud Storage

When you set up tiering, you need to provide storage access keys for a service account that has Storage Admin permissions. A service account enables Cloud Tiering to authenticate and access Cloud Storage buckets used for data tiering. The keys are required so that Google Cloud Storage knows who is making the request.

The Cloud Storage buckets must be in a [region that supports Cloud Tiering](#).

 If you are planning to configure Cloud Tiering to use lower cost storage classes where your tiered data will transition to after a certain number of days, you must not select any lifecycle rules when setting up the bucket in your GCP account. Cloud Tiering manages the lifecycle transitions.

### Steps

1. [Create a service account that has the predefined Storage Admin role](#).
2. Go to [GCP Storage Settings](#) and create access keys for the service account:
  - a. Select a project, and select **Interoperability**. If you haven't already done so, select **Enable interoperability access**.
  - b. Under **Access keys for service accounts**, select **Create a key for a service account**, select the service account that you just created, and select **Create Key**.

You'll need to enter the keys later when you set up Cloud Tiering.

## Tier inactive data from your first cluster to Google Cloud Storage

After you prepare your Google Cloud environment, start tiering inactive data from your first cluster.

### What you'll need

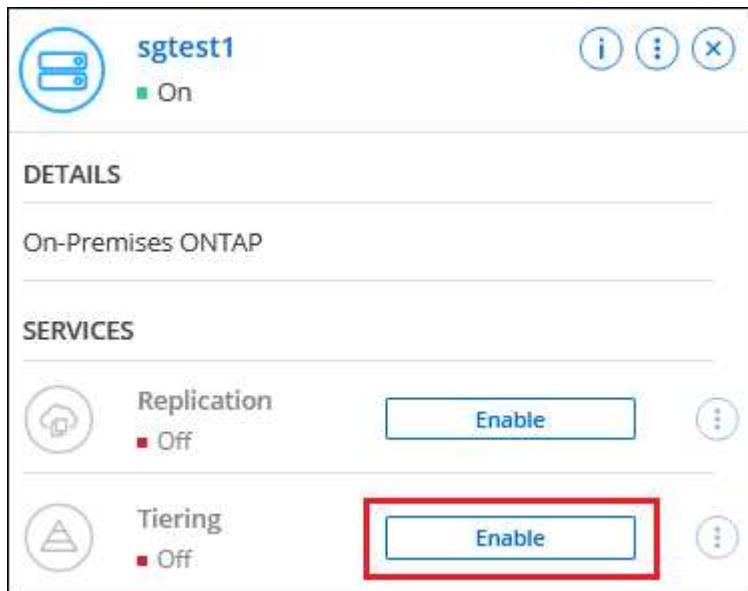
- [An on-premises system added to the NetApp Console](#).

- Storage access keys for a service account that has the Storage Admin role.

## Steps

1. Select the on-premises ONTAP system.
2. Click **Enable** for the Tiering service from the right panel.

If the Google Cloud Storage tiering destination is available on the **Systems** page, you can drag the cluster onto the Google Cloud Storage system to initiate the setup wizard.



3. **Define Object Storage Name:** Enter a name for this object storage. It must be unique from any other object storage you may be using with aggregates on this cluster.
4. **Select Provider:** Select **Google Cloud** and select **Continue**.
5. Complete the steps on the **Create Object Storage** pages:
  - a. **Bucket:** Add a new Google Cloud Storage bucket or select an existing bucket.
  - b. **Storage Class Lifecycle:** Cloud Tiering manages the lifecycle transitions of your tiered data. Data starts in the *Standard* class, but you can create rules to apply different storage classes after a certain number of days.

Select the Google Cloud storage class that you want to transition the tiered data to and the number of days before the data is assigned to that class, and select **Continue**. For example, the screenshot below shows that tiered data is assigned to the *Nearline* class from the *Standard* class after 30 days in object storage, and then to the *Coldline* class after 60 days in object storage.

If you choose **Keep data in this storage class**, then the data remains in the that storage class. [See supported storage classes](#).

## Storage Class Life Cycle Management

We'll move the tiered data through the storage classes that you include in the life cycle. [Learn more about Google Cloud Storage classes.](#)

### STORAGE CLASS SETUP

#### Standard

Move data from Standard to Nearline after  days

Keep data in this storage class



#### Nearline

Move data from Nearline to Coldline after  days

Keep data in this storage class



#### Coldline

Move data from Coldline to Archive after  days

Keep data in this storage class



#### Archive

No Time Limit

Note that the lifecycle rule is applied to all objects in the selected bucket.

- c. **Credentials:** Enter the storage access key and secret key for a service account that has the Storage Admin role.
- d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage.

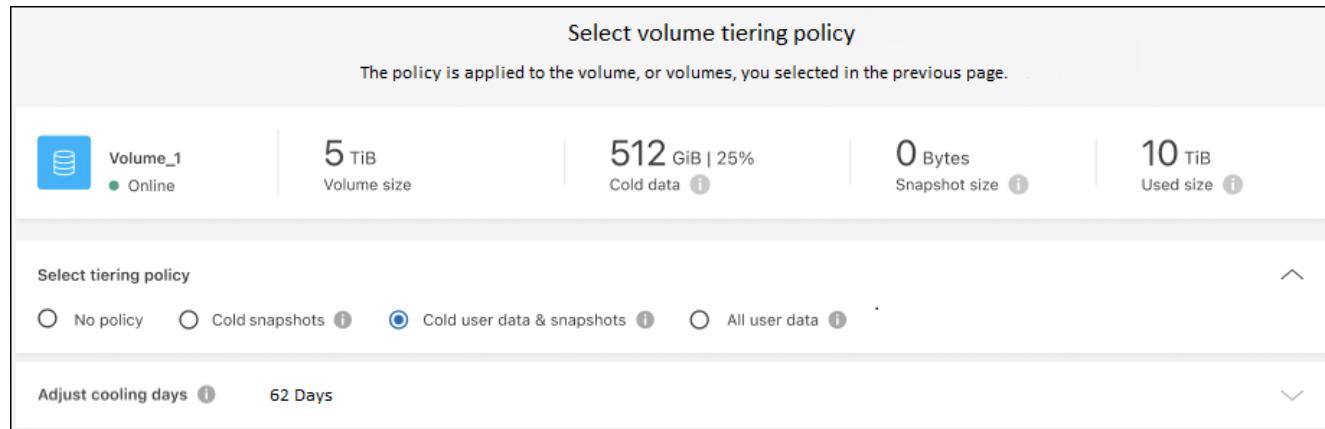
Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.

You can also set the network bandwidth available to upload inactive data to object storage by defining the "Maximum transfer rate". Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.

6. Click **Continue** to select the volumes that you want to tier.
7. On the *Tier Volumes* page, select the volumes that you want to configure tiering for and launch the **Tiering Policy** page:
  - To select all volumes, check the box in the title row ( **Volume Name**) and select **Configure volumes**.
  - To select multiple volumes, check the box for each volume ( **Volume\_1**) and select **Configure volumes**.
  - To select a single volume, select the row (or  icon) for the volume.

8. In the *Tiering Policy* dialog, select a tiering policy, optionally adjust the cooling days for the selected volumes, and select **Apply**.

Learn more about volume tiering policies and cooling days.



## Result

You've successfully set up data tiering from volumes on the cluster to Google Cloud object storage.

## What's next?

Be sure to subscribe to the Cloud Tiering service.

You can review information about the active and inactive data on the cluster. [Learn more about managing your tiering settings](#).

You can also create additional object storage in cases where you may want to tier data from certain aggregates on a cluster to different object stores. Or if you plan to use FabricPool Mirroring where your tiered data is replicated to an additional object store. [Learn more about managing object stores](#).

## Tiering data from on-premises ONTAP clusters to StorageGRID in NetApp Cloud Tiering

Free space on your on-premises ONTAP clusters by tiering inactive data to StorageGRID in NetApp Cloud Tiering.

### Quick start

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

1

#### Prepare to tier data to StorageGRID

You need the following:

- A source on-premises ONTAP cluster that's running ONTAP 9.4 or later that you have added to the NetApp Console, and a connection over a user-specified port to StorageGRID. [Learn how to discover a cluster](#).
- StorageGRID 10.3 or later with AWS access keys that have S3 permissions.

- A Console agent installed on your premises.
- Networking for the agent that enables an outbound HTTPS connection to the ONTAP cluster, to StorageGRID, and to the Cloud Tiering service.

2

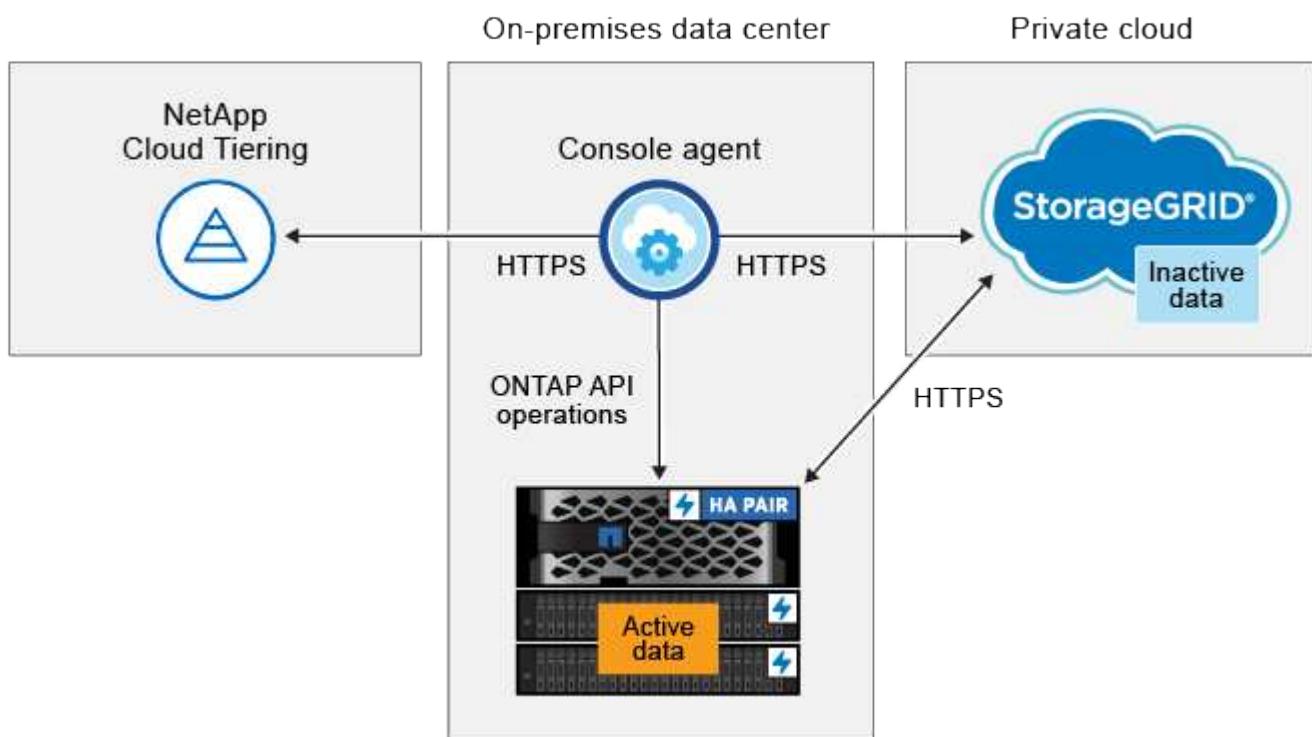
## Set up tiering

In the NetApp Console, select an on-premises system, select **Enable** for Cloud Tiering, and follow the prompts to tier data to StorageGRID.

## Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between the agent and StorageGRID is for object storage setup only.

## Prepare your ONTAP clusters

Your ONTAP clusters must meet the following requirements when tiering data to StorageGRID.

### Supported ONTAP platforms

- When using ONTAP 9.8 and later: You can tier data from AFF systems, or FAS systems with all-SSD aggregates or all-HDD aggregates.
- When using ONTAP 9.7 and earlier: You can tier data from AFF systems, or FAS systems with all-SSD aggregates.

## Supported ONTAP version

ONTAP 9.4 or later

## Licensing

A Cloud Tiering license isn't required in your NetApp Console organization, nor is a FabricPool license required on the ONTAP cluster, when tiering data to StorageGRID.

## Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over a user-specified port to the StorageGRID Gateway Node (the port is configurable during tiering setup).

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- An inbound connection is required from the agent, which must reside on your premises.

A connection between the cluster and the Cloud Tiering service is not required.

- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to tier. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

When you set up data tiering, Cloud Tiering prompts you for the *IPspace* to use. You should choose the *IPspace* that each LIF is associated with. That might be the "Default" *IPspace* or a custom *IPspace* that you created. Learn more about [LIFs](#) and [IPspaces](#).

## Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. Refer to the ONTAP documentation for [functionality or features not supported by FabricPool](#).



Cloud Tiering supports FlexGroup volumes, starting with ONTAP 9.5. Setup works the same as any other volume.

## Discover an ONTAP cluster

You need to add an on-premises ONTAP system to the NetApp Console before you can start tiering cold data.

[Learn how to discover a cluster.](#)

## Prepare StorageGRID

StorageGRID must meet the following requirements.

## Supported StorageGRID versions

StorageGRID 10.3 and later is supported.

## S3 credentials

When you set up tiering to StorageGRID, you need to provide Cloud Tiering with an S3 access key and secret key. Cloud Tiering uses the keys to access your buckets.

These access keys must be associated with a user who has the following permissions:

```
"s3>ListAllMyBuckets",  
"s3>ListBucket",  
"s3GetObject",  
"s3PutObject",  
"s3DeleteObject",  
"s3CreateBucket"
```

## Object versioning

You must not enable StorageGRID object versioning on the object store bucket.

## Create or switch Console agents

The Console agent is required to tier data to the cloud. When tiering data to StorageGRID, an agent must be available on your premises.

You must have the Organization admin role to create an agent.

- [Learn about agents](#)
- [Install and set up an agent on-premises](#)
- [Switch between agents](#)

## Prepare networking for the Console agent

Ensure that the agent has the required networking connections.

### Steps

1. Ensure that the network where the agent is installed enables the following connections:

- An HTTPS connection over port 443 to the Cloud Tiering service ([see the list of endpoints](#))
- An HTTPS connection over port 443 to your StorageGRID system
- An HTTPS connection over port 443 to your ONTAP cluster management LIF

## Tier inactive data from your first cluster to StorageGRID

After you prepare your environment, start tiering inactive data from your first cluster.

### What you'll need

- [An on-premises system added to the NetApp Console](#).
- The FQDN of the StorageGRID Gateway Node, and the port that will be used for HTTPS communications.
- An AWS access key that has the required S3 permissions.

### Steps

1. Select the on-premises ONTAP system.
2. Click **Enable** for Cloud Tiering from the right panel.

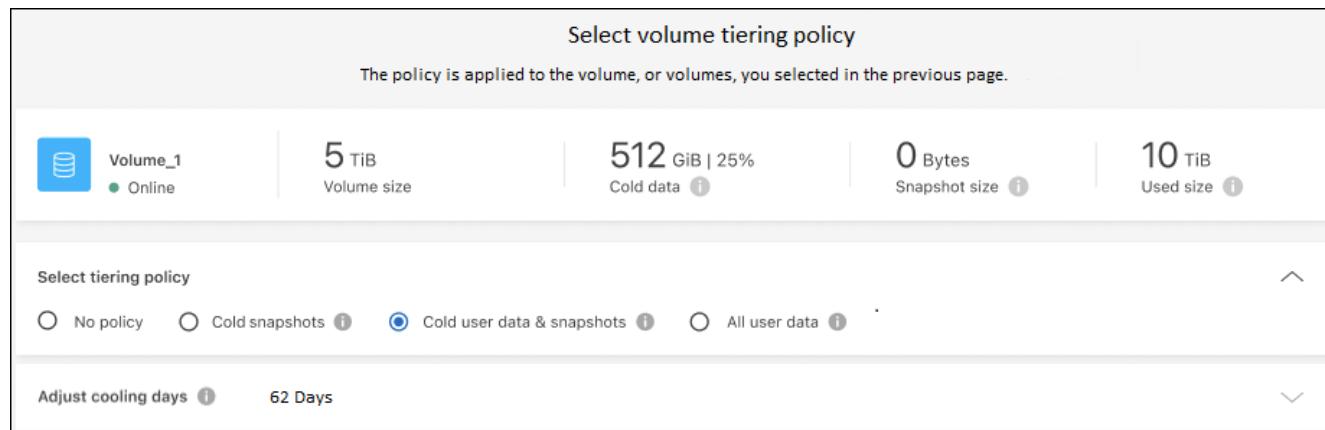
If the StorageGRID tiering destination exists as a system in the NetApp Console you can drag the cluster onto the StorageGRID system to initiate the setup wizard.



3. **Define Object Storage Name:** Enter a name for this object storage. It must be unique from any other object storage you may be using with aggregates on this cluster.
4. **Select Provider:** Select **StorageGRID** and select **Continue**.
5. Complete the steps on the **Create Object Storage** pages:
  - a. **Server:** Enter the FQDN of the StorageGRID Gateway Node, the port that ONTAP should use for HTTPS communication with StorageGRID, and the access key and secret key for an account that has the required S3 permissions.
  - b. **Bucket:** Add a new bucket or select an existing bucket that starts with the prefix *fabric-pool* and select **Continue**.  
 The *fabric-pool* prefix is required because the IAM policy for the agent enables the instance to perform S3 actions on buckets named with that exact prefix. For example, you could name the S3 bucket *fabric-pool-AFF1*, where AFF1 is the name of the cluster.
  - c. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and select **Continue**.  
 Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to StorageGRID object storage.  
 You can also set the network bandwidth available to upload inactive data to object storage by defining the "Maximum transfer rate". Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.
6. On the *Tier Volumes* page, select the volumes that you want to configure tiering for and launch the Tiering Policy page:
  - To select all volumes, check the box in the title row ( **Volume Name**) and select **Configure volumes**.
  - To select multiple volumes, check the box for each volume ( **Volume\_1**) and select **Configure volumes**.
  - To select a single volume, select the row (or icon) for the volume.

7. In the *Tiering Policy* dialog, select a tiering policy, optionally adjust the cooling days for the selected volumes, and select **Apply**.

[Learn more about volume tiering policies and cooling days.](#)



#### What's next?

You can review information about the active and inactive data on the cluster. [Learn more about managing your tiering settings.](#)

You can also create additional object storage in cases where you may want to tier data from certain aggregates on a cluster to different object stores. Or if you plan to use FabricPool Mirroring where your tiered data is replicated to an additional object store. [Learn more about managing object stores.](#)

## Tier data from on-premises ONTAP clusters to S3 object storage in NetApp Cloud Tiering

Free space on your on-premises ONTAP clusters by tiering inactive data in NetApp Cloud Tiering to any object storage service which uses the Simple Storage Service (S3) protocol.

At this time, MinIO object storage has been qualified.

Customers who want to use object stores that are not officially supported as a cloud tier can do so using these instructions. Customers must test and confirm that the object store meets their requirements.



NetApp does not support, nor is liable, for any issues arising from any third-party Object Store Service, specifically where it does not have agreed support arrangements with the third party with whom the product originated. It is acknowledged and agreed that NetApp shall not be liable for any associated damage or otherwise be required to provide support on that third-party product.

## Quick start

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

**1**

## Prepare to tier data to S3-compatible object storage

You need the following:

- A source on-premises ONTAP cluster that's running ONTAP 9.8 or later that you have added to the NetApp Console, and a connection over a user-specified port to the destination S3-compatible object storage. [Learn how to discover a cluster.](#)
- The FQDN, Access Key, and Secret Key for the object storage server so that the ONTAP cluster can access the bucket.
- A Console agent installed on your premises.
- Networking for the agent that enables an outbound HTTPS connection to the source ONTAP cluster, to the S3-compatible object storage, and to the Cloud Tiering service.

**2**

## Set up tiering

In the Console, select an on-premises system, select **Enable** for the Tiering service, and follow the prompts to tier data to S3-compatible object storage.

**3**

## Set up licensing

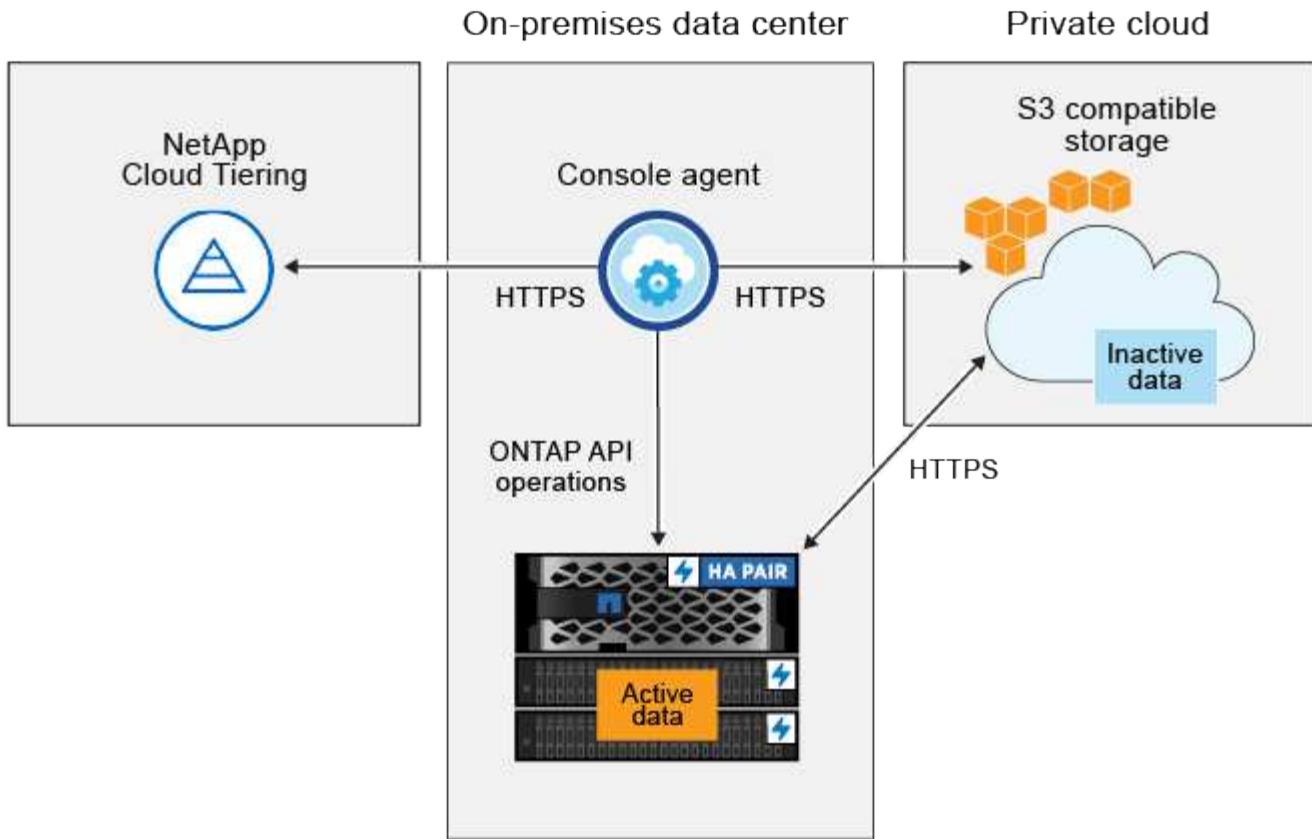
Pay for Cloud Tiering through a pay-as-you-go subscription from your cloud provider, a Cloud Tiering bring-your-own-license, or a combination of both:

- To subscribe to the PAYGO offering from the [AWS Marketplace](#), [Azure Marketplace](#), or [GCP Marketplace](#), select **Subscribe** and follow the prompts.
- To pay using a Cloud Tiering BYOL license, [contact us if you need to purchase one](#), and then [add it to the NetApp Console](#).

## Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between the agent and the S3-compatible object storage server is for object storage setup only.

## Prepare your ONTAP clusters

Your source ONTAP clusters must meet the following requirements when tiering data to S3-compatible object storage.

### Supported ONTAP platforms

You can tier data from AFF systems, or FAS systems with all-SSD aggregates or all-HDD aggregates.

### Supported ONTAP version

ONTAP 9.8 or later

### Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over a user-specified port to S3-compatible object storage (the port is configurable during tiering setup).

The source ONTAP system reads and writes data to and from object storage. The object storage never initiates, it just responds.

- An inbound connection is required from the agent, which must reside on your premises.

A connection between the cluster and the Cloud Tiering service is not required.

- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to tier. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

When you set up data tiering, Cloud Tiering prompts you for the IPspace to use. You should choose the IPspace that each LIF is associated with. That might be the "Default" IPspace or a custom IPspace that you created. Learn more about [LIFs](#) and [IPspaces](#).

## Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. Refer to the ONTAP documentation for [functionality or features not supported by FabricPool](#).



Cloud Tiering supports both FlexVol and FlexGroup volumes.

## Discover an ONTAP cluster

You need to add your on-premises ONTAP system to the Console before you can start tiering cold data.

[Learn how to discover a cluster.](#)

## Prepare S3-compatible object storage

S3-compatible object storage must meet the following requirements.

### S3 credentials

When you set up tiering to S3-compatible object storage, you're prompted to create an S3 bucket or to select an existing S3 bucket. You need to provide Cloud Tiering with an S3 access key and secret key. Cloud Tiering uses the keys to access your bucket.

These access keys must be associated with a user who has the following permissions:

```
"s3>ListAllMyBuckets",  
"s3>ListBucket",  
"s3.GetObject",  
"s3:PutObject",  
"s3>DeleteObject",  
"s3:CreateBucket"
```

## Create or switch agents

A Console agent is required to tier data to the cloud. When tiering data to S3-compatible object storage, an agent must be available on your premises. You'll either need to install a new agent or make sure that the currently selected agent resides on-premises.

- [Learn about agents](#)
- [Install and set up an agent on-premises](#)
- [Switch between agents](#)

## Prepare networking for the Console agent

Ensure that the agent has the required networking connections.

### Steps

1. Ensure that the network where the agent is installed enables the following connections:
  - An HTTPS connection over port 443 to the Cloud Tiering service ([see the list of endpoints](#))
  - An HTTPS connection over port 443 to S3-compatible object storage
  - An HTTPS connection over port 443 to your ONTAP cluster management LIF

## Tiering inactive data from your first cluster to S3-compatible object storage

After you prepare your environment, start tiering inactive data from your first cluster.

### What you'll need

- [An on-premises system added to NetApp Console](#).
- The FQDN of the S3-compatible object storage server and the port that will be used for HTTPS communications.
- An access key and secret key that has the required S3 permissions.

### Steps

1. Select the on-premises ONTAP system.
2. Click **Enable** for the Cloud Tiering service from the right panel.



3. **Define Object Storage Name:** Enter a name for this object storage. It must be unique from any other object storage you may be using with aggregates on this cluster.
4. **Select Provider:** Select **S3 Compatible** and select **Continue**.
5. Complete the steps on the **Create Object Storage** pages:
  - a. **Server:** Enter the FQDN of the S3-compatible object storage server, the port that ONTAP should use for HTTPS communication with the server, and the access key and secret key for an account that has the required S3 permissions.
  - b. **Bucket:** Add a new bucket or select an existing bucket and select **Continue**.
  - c. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and select **Continue**.

Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your S3-compatible object storage.

You can also set the network bandwidth available to upload inactive data to object storage by defining the "Maximum transfer rate". Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.

6. On the *Success* page select **Continue** to set up your volumes now.
7. On the *Tier Volumes* page, select the volumes that you want to configure tiering for and select **Continue**:

- To select all volumes, check the box in the title row ( **Volume Name**) and select **Configure volumes**.
- To select multiple volumes, check the box for each volume ( **Volume\_1**) and select **Configure volumes**.
- To select a single volume, select the row (or icon) for the volume.

Tier Volumes									
Volumes (16) <small>(1)</small>   2 selected	Volume Name	Aggregate/s Name	SVM Name	Node/s Name	Volume Size	Snapshot Size	Cold Data...	Tier Status	Tiering Policy
<input type="checkbox"/>	volume_1	aggr-1	svm_1	volume_1_node	20 TB	0 B	10 TB   10 %	Not Tiered Volume	No Policy
<input checked="" type="checkbox"/>	volume_2	aggr-1	svm_1	volume_2_node	15 TB	205 GB	2.8 TB   70 %	Not Tiered Volume	No Policy
<input checked="" type="checkbox"/>	volume_3	aggr-1	svm_1	volume_3_node	8 TB	0 B	716.8 GB   70 %	Not Tiered Volume	No Policy
<input type="checkbox"/>	volume_4	aggr-1	svm_1	volume_4_node	3 TB	3 GB	35 GB   70 %	Not Tiered Volume	No Policy

8. In the *Tiering Policy* dialog, select a tiering policy, optionally adjust the cooling days for the selected volumes, and select **Apply**.

[Learn more about volume tiering policies and cooling days.](#)

Select volume tiering policy  
The policy is applied to the volume, or volumes, you selected in the previous page.

Volume_1	5 TiB Volume size	512 GiB   25% Cold data <small>i</small>	0 Bytes Snapshot size <small>i</small>	10 TiB Used size <small>i</small>
----------	----------------------	---	---	--------------------------------------

Select tiering policy

No policy    Cold snapshots i    Cold user data & snapshots i    All user data i

Adjust cooling days i   62 Days

## What's next?

[Be sure to subscribe to the Cloud Tiering service.](#)

You can review information about the active and inactive data on the cluster. [Learn more about managing your tiering settings.](#)

You can also create additional object storage in cases where you may want to tier data from certain aggregates on a cluster to different object stores. Or if you plan to use FabricPool Mirroring where your tiered data is

replicated to an additional object store. [Learn more about managing object stores.](#)

## Copyright information

Copyright © 2025 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.