



Get started

NetApp Cloud Tiering

NetApp

November 10, 2025

This PDF was generated from <https://docs.netapp.com/us-en/data-services-cloud-tiering/concept-cloud-tiering.html> on November 10, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Get started	1
Learn about NetApp Cloud Tiering	1
NetApp Console	1
Features	1
Supported object storage providers	2
Pricing and licenses	3
How Cloud Tiering works	4
Tier on-premises data to the cloud	6
Tier data from on-premises ONTAP clusters to Amazon S3 in NetApp Cloud Tiering	6
Tier data from on-premises ONTAP clusters to Azure Blob storage in NetApp Cloud Tiering	18
Tier data from on-premises ONTAP clusters to Google Cloud Storage in NetApp Cloud Tiering	24
Tiering data from on-premises ONTAP clusters to StorageGRID in NetApp Cloud Tiering	30
Tier data from on-premises ONTAP clusters to S3 object storage in NetApp Cloud Tiering	35
Set up licensing for NetApp Cloud Tiering	41
30-day free trial	41
Use a Cloud Tiering PAYGO subscription	41
Use an annual contract	42
Use a Cloud Tiering BYOL license	43
Apply Cloud Tiering licenses to clusters in special configurations	44
NetApp Cloud Tiering technical FAQ	45
Cloud Tiering service	45
Licenses and Costs	47
ONTAP	48
Object storage	49
Console agents	51
Tiering policies	52
Networking and security	53

Get started

Learn about NetApp Cloud Tiering

NetApp Cloud Tiering extends your data center to the cloud by automatically tiering inactive data from on-premises ONTAP clusters to object storage. This frees valuable space on the cluster for more workloads, without making changes to the application layer. Cloud Tiering can reduce costs in your data center and enables you to switch from a CAPEX model to an OPEX model.

Cloud Tiering leverages the capabilities of *FabricPool*. FabricPool is a NetApp Data Fabric technology that enables automated tiering of data to low-cost object storage. Active (hot) data remains on the local tier (on-premises ONTAP aggregates), while inactive (cold) data is moved to the cloud tier — all while preserving ONTAP data efficiencies.

Originally supported on AFF, FAS, and ONTAP Select systems with all-SSD aggregates, starting with ONTAP 9.8 you can tier data from aggregates consisting of HDDs in addition to high-performance SSDs. See [the considerations and requirements for using FabricPool](#) for details.

You can configure tiering for single-node clusters, HA-configured clusters, clusters in Tiering Mirror configurations, and MetroCluster configurations using FabricPool Mirror. Cloud Tiering licenses are shared among all of your clusters.

[Use the Cloud Tiering TCO calculator to see how much money you can save.](#)

NetApp Console

NetApp Cloud Tiering is accessible through the NetApp Console.

The NetApp Console provides centralized management of NetApp storage and data services across on-premises and cloud environments at enterprise scale. The Console is required to access and use NetApp data services. As a management interface, it enables you to manage many storage resources from one interface. Console administrators can control access to storage and services for all systems within the enterprise.

You don't need a license or subscription to start using the NetApp Console and you only incur charges when you need to deploy Console agents in your cloud to ensure connectivity to your storage systems or NetApp data services. However, some NetApp data services accessible from the Console are licensed or subscription-based.

Learn more about the [NetApp Console](#).

Features

Cloud Tiering offers automation, monitoring, reports, and a common management interface:

- Automation makes it easier to set up and manage data tiering from on-premises ONTAP clusters to the cloud.
- You can choose the default cloud provider storage class/access tier, or use lifecycle management to assign a more cost-effective tier to older tiered data.
- You can create connections to additional object stores that can be used for other aggregates in your cluster.

- Using the UI, you can drag object stores to an aggregate for tiering and for FabricPool mirroring.
- A single pane of glass removes the need to independently manage FabricPool across several clusters.
- Reports show the amount of active and inactive data on each cluster.
- A tiering health status helps you identify and correct issues as they occur.
- If you have Cloud Volumes ONTAP systems, you'll find them in the Clusters page so you get a full view of data tiering in your hybrid cloud infrastructure.

For more details about the value that Cloud Tiering provides, [check out the Cloud Tiering page on the NetApp Console website](#).



Cloud Volumes ONTAP systems are read-only from Cloud Tiering. [You set up tiering for Cloud Volumes ONTAP systems in the NetApp Console..](#)

Supported object storage providers

You can tier inactive data from an on-premises ONTAP system to the following object storage providers:

- Amazon S3
- Microsoft Azure Blob
- Google Cloud Storage
- NetApp StorageGRID
- S3-compatible object storage (for example, MinIO)

Cloud Tiering licenses can also be shared with your clusters that are tiering data to IBM Cloud Object Storage. The FabricPool configuration must be set up using System Manager or the ONTAP CLI, but [licensing for this type of configuration is completed using Cloud Tiering](#).



You can tier data from NAS volumes to the public cloud or to private clouds, like StorageGRID. When you tier data that is accessed by SAN protocols, NetApp recommends using private clouds due to connectivity considerations.

Object storage tiers

ONTAP clusters can tier inactive data to a single object store, or to multiple object stores. When you set up data tiering, you have the choice to add a new bucket/container or to select an existing bucket/container, along with a storage class or access tier.

- [Learn about supported AWS S3 storage classes](#)
- [Learn about supported Azure Blob access tiers](#)
- [Learn about supported Google Cloud storage classes](#)

Cloud Tiering uses the cloud provider default storage class/access tier for your inactive data. However, you can apply a lifecycle rule so that the data automatically transitions from the default storage class to another storage class after a certain number of days. This can help keep your costs down by moving very cold data to less expensive storage.



You can't select lifecycle rules for data tiered to StorageGRID or S3-compatible storage.

Pricing and licenses

Pay for Cloud Tiering through a pay-as-you-go subscription, an annual subscription, a bring-your-own NetApp tiering license, or a combination. A 30-day free trial is available for your first cluster if you don't have a license.

There are no charges when tiering data to StorageGRID. Neither a BYOL license or PAYGO registration is required.

[View pricing details.](#)

Because Cloud Tiering preserves the storage efficiencies of the source volume, you pay the cloud provider object storage costs for the tiered data after ONTAP efficiencies (for the smaller amount of data after deduplication and compression have been applied).

30-day free trial

If you don't have a Cloud Tiering license, a 30-day free trial of tiering starts when you set up tiering to your first cluster. After the 30-day free trial ends, you'll need to pay for tiering through a pay-as-you-go subscription, annual subscription, a BYOL license, or a combination.

If your free trial ends and you haven't subscribed or added a license, then ONTAP no longer tiers cold data to object storage. All previously tiered data remains accessible; meaning you can retrieve and use this data. When retrieved, this data is moved back to the performance tier from the cloud.

Pay-as-you-go subscription

Cloud Tiering offers consumption-based licensing in a pay-as-you-go model. After subscribing through your cloud provider's marketplace, you pay per GB for data that's tiered—there's no up-front payment. You are billed by your cloud provider through your monthly bill.

You should subscribe even if you have a free trial or if you bring your own license (BYOL):

- Subscribing ensures that there's no disruption of service after your free trial ends.

When the trial ends, you'll be charged hourly according to the amount of data that you tier.

- If you tier more data than allowed by your BYOL license, then data tiering continues through your pay-as-you-go subscription.

For example, if you have a 10 TB license, all capacity beyond the 10 TB is charged through the pay-as-you-go subscription.

You won't be charged from your pay-as-you-go subscription during your free trial or if you haven't exceeded your Cloud Tiering BYOL license.

[Learn how to set up a pay-as-you-go subscription.](#)

Annual contract

Cloud Tiering offers an annual contract when tiering inactive data to Amazon S3 or Azure. It's available in 1-, 2-, or 3-year terms.

Annual contracts are not currently supported when tiering to Google Cloud.

Bring your own license

Bring your own license by purchasing a **Cloud Tiering** license from NetApp (previously known as a "Cloud Tiering" license). You can purchase 1-, 2-, or 3-year term licenses and specify any amount of tiering capacity (starting at a minimum of 10 TiB). The BYOL Cloud Tiering license is a *floating* license that you can use across multiple on-premises ONTAP clusters. The total tiering capacity that you define in your Cloud Tiering license can be used by all of your on-premises clusters.

After you purchase a Cloud Tiering license, you'll need to add the license to the NetApp Console. [See how to use a Cloud Tiering BYOL license.](#)

As noted above, we recommend that you set up a pay-as-you-go subscription, even if you have purchased a BYOL license.




Starting August 2021 the old **FabricPool** license was replaced by the **Cloud Tiering** license. [Read more about how the Cloud Tiering license is different than the FabricPool license.](#)

How Cloud Tiering works

Cloud Tiering is a NetApp-managed service that uses FabricPool technology to automatically tier inactive (cold) data from your on-premises ONTAP clusters to object storage in your public cloud or private cloud. Connections to ONTAP take place from a Console agent.

The following image shows the relationship between each component:



At a high level, Cloud Tiering works like this:

1. You discover your on-premises cluster from the NetApp Console.
2. You set up tiering by providing details about your object storage, including the bucket/container, a storage

class or access tier, and lifecycle rules for the tiered data.

3. The Console configures ONTAP to use the object storage provider and discovers the amount of active and inactive data on the cluster.
4. You choose the volumes to tier and the tiering policy to apply to those volumes.
5. ONTAP starts tiering inactive data to the object store as soon as the data has reached the thresholds to be considered inactive (see [Volume tiering policies](#)).
6. If you have applied a lifecycle rule to the tiered data (only available for some providers), older tiered data is assigned to a more cost-effective tier after a certain number of days.

Volume tiering policies

When you select the volumes that you want to tier, you choose a *volume tiering policy* to apply to each volume. A tiering policy determines when or whether the user data blocks of a volume are moved to the cloud.

You can also adjust the **cooling period**. This is the number of days that user data in a volume must remain inactive before it is considered "cold" and moved to object storage. For tiering policies that allow you to adjust the cooling period, the valid values are:

- 2 to 183 days when using ONTAP 9.8 and later
- 2 to 63 days for earlier ONTAP versions

2 to 63 is the recommended best practice.

No Policy (None)

Keeps the data on a volume in the performance tier, preventing it from being moved to the cloud tier.

Cold snapshots (Snapshot only)

ONTAP tiers cold Snapshot blocks in the volume that are not shared with the active file system to object storage. If read, cold data blocks on the cloud tier become hot and are moved to the performance tier.

Data is tiered only after an aggregate has reached 50% capacity and when the data has reached the cooling period. The default number of cooling days is 2, but you can adjust this number.



Re-heated data is written back to the performance tier only if there is space. If the performance tier capacity is more than 70% full, blocks continue to be accessed from the cloud tier.

Cold user data & snapshots (Auto)

ONTAP tiers all cold blocks in the volume (not including metadata) to object storage. The cold data includes not just Snapshot copies, but also cold user data from the active file system.

- If read by random reads, cold data blocks on the cloud tier become hot and are moved to the performance tier.
- If read by sequential reads, such as those associated with index and antivirus scans, cold data blocks on the cloud tier stay cold and are not written to the performance tier.

This policy is available starting with ONTAP 9.4.

Data is tiered only after an aggregate has reached 50% capacity and when the data has reached the cooling period. The default number of cooling days is 31, but you can adjust this number.



Re-heated data is written back to the performance tier only if there is space. If the performance tier capacity is more than 70% full, blocks continue to be accessed from the cloud tier.

All user data (All)

All data (not including metadata) is immediately marked as cold and tiered to object storage as soon as possible. There is no need to wait 48 hours for new blocks in a volume to become cold. Blocks located in the volume prior to the All policy being set require 48 hours to become cold.

If read, cold data blocks on the cloud tier stay cold and are not written back to the performance tier. This policy is available starting with ONTAP 9.6.

Take the following into consideration before you choose this tiering policy:

- Tiering data immediately reduces storage efficiencies (inline only).
- You should use this policy only if you are confident that cold data on the volume will not change.
- Object storage is not transactional and will result in significant fragmentation if subjected to change.
- Consider the impact of SnapMirror transfers before assigning the All tiering policy to source volumes in data protection relationships.

Because data is tiered immediately, SnapMirror will read data from the cloud tier rather than the performance tier. This will result in slower SnapMirror operations—possibly slowing other SnapMirror operations later in queue—even if they are using different tiering policies.

- NetApp Backup and Recovery is similarly affected by volumes set with a tiering policy. [See tiering policy considerations with Backup and Recovery](#).

All DP user data (Backup)

All data on a data protection volume (not including metadata) is immediately moved to the cloud tier. If read, cold data blocks on the cloud tier stay cold and are not written back to the performance tier (starting with ONTAP 9.4).



This policy is available for ONTAP 9.5 or earlier. It was replaced with the **All** tiering policy starting with ONTAP 9.6.

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.

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.

vpce-0ff5c15df7e00fbab / RN-VPC-Endpoint

Details

Endpoint ID	vpce-0ff5c15df7e00fbab	Status	Available
VPC ID	vpc-673ae603 [VPC EMEA]	Status message	-
DNS names	<ul style="list-style-type: none"> *.vpce-0ff5c15df7e00fbab-yxs7it8v.us-west-2.amazonaws.com - (Z1YSA3EXCYUU9Z) *.vpce-0ff5c15df7e00fbab-yxs7it8v-us-west-2b.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	Enable	
	Tiering	■ Off	Enable	

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	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

↓

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%	0 Bytes	10 TiB
<small>Online</small>	<small>Volume size</small>	<small>Cold data</small>	<small>Snapshot size</small>	<small>Used size</small>

Select tiering policy

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

Adjust cooling days i 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.](#)



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:



i 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) 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 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	512 GiB 25% Cold data	0 Bytes	10 TiB Used size
----------	-------	----------------------------	---------	---------------------

Select tiering policy

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

Adjust cooling days 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.](#)

Set up licensing for NetApp Cloud Tiering

A 30-day free trial of NetApp Cloud Tiering starts when you set up tiering from your first cluster. After the free trial ends, you'll need to pay for Cloud Tiering through a pay-as-you-go or annual subscription from your cloud provider's marketplace, a BYOL license from NetApp, or a combination of both.

A few notes before you read any further:

- If you've already subscribed to Cloud Tiering (PAYGO) in your cloud provider's marketplace, then you're automatically subscribed to Cloud Tiering for on-premises ONTAP systems as well. You'll see an active subscription in the Cloud Tiering **On-Premises dashboard** tab. You won't need to subscribe again. You'll see an active subscription in the NetApp Console.
- The BYOL Cloud Tiering license (previously known as a "Cloud Tiering" license) is a *floating* license that you can use across multiple on-premises ONTAP clusters in your NetApp Console organization. This is different (and much easier) than in the past where you purchased a *FabricPool* license for each cluster.
- There are no charges when tiering data to StorageGRID, so neither a BYOL license nor a PAYGO registration is required. This tiered data doesn't count against the capacity purchased in your license.

[Learn more about how licensing works for Cloud Tiering.](#)

30-day free trial

If you don't have a Cloud Tiering license, a 30-day free trial of Cloud Tiering starts when you set up tiering to your first cluster. After the 30-day free trial ends, you'll need to pay for Cloud Tiering through a pay-as-you-go subscription, annual subscription, a BYOL license, or a combination.

If your free trial ends and you haven't subscribed or added a license, then ONTAP no longer tiers cold data to object storage. All previously tiered data remains accessible; meaning you can retrieve and use this data. When retrieved, this data is moved back to the performance tier from the cloud.

Use a Cloud Tiering PAYGO subscription

Pay-as-you-go subscriptions from your cloud provider's marketplace enable you to license the use of Cloud Volumes ONTAP systems and many Cloud Data Services, such as Cloud Tiering.

After you subscribe to Cloud Tiering, you can manage your subscriptions in the Console. [View and manage your subscriptions.](#)

Subscribing from the AWS Marketplace

Subscribe to Cloud Tiering from the AWS Marketplace to set up a pay-as-you-go subscription for data tiering from ONTAP clusters to AWS S3.

Steps

1. In the NetApp Console, select **Mobility > Cloud Tiering > On-Premises Dashboard**.
2. In the *Marketplace subscriptions* section, select **Subscribe** under Amazon Web Services and then select **Continue**.
3. Subscribe from the [AWS Marketplace](#), and then log back into the NetApp Console to complete the registration.

The following video shows the process:

[Subscribe from the AWS Marketplace](#)

Subscribing from the Azure Marketplace

Subscribe to Cloud Tiering from the Azure Marketplace to set up a pay-as-you-go subscription for data tiering from ONTAP clusters to Azure Blob storage.

Steps

1. In the NetApp Console, select **Mobility > Cloud Tiering > On-Premises Dashboard**.
2. In the *Marketplace subscriptions* section, select **Subscribe** under Microsoft Azure and then select **Continue**.
3. Subscribe from the [Azure Marketplace](#), and then log back into the NetApp Console to complete the registration.

The following video shows the process:

[Subscribe from the Azure Marketplace](#)

Subscribing from the Google Cloud Marketplace

Subscribe to Cloud Tiering from the Google Cloud Marketplace to set up a pay-as-you-go subscription for data tiering from ONTAP clusters to Google Cloud storage.

Steps

1. In the NetApp Console, select **Mobility > Cloud Tiering > On-Premises Dashboard**.
2. In the *Marketplace subscriptions* section, select **Subscribe** under Google Cloud and then select **Continue**.
3. Subscribe from the [Google Cloud Marketplace](#), and then log back into the NetApp Console to complete the registration.

The following video shows the process:

[Subscribe from the Google Cloud Marketplace](#)

Use an annual contract

Pay for Cloud Tiering annually by purchasing an annual contract. Annual contracts are available in 1-, 2-, or 3-year terms.

When tiering inactive data to AWS, you can subscribe to an annual contract from the [AWS Marketplace page](#). If you want to use this option, set up your subscription from the Marketplace page and then [associate the subscription with your AWS credentials](#).

When tiering inactive data to Azure, you can subscribe to an annual contract from the [Azure Marketplace page](#). If you want to use this option, set up your subscription from the Marketplace page and then [associate the subscription with your Azure credentials](#).

Annual contracts are not currently supported when tiering to Google Cloud.

Use a Cloud Tiering BYOL license

Bring-your-own licenses from NetApp provide 1-, 2-, or 3-year terms. The **BYOL Cloud Tiering** license (previously known as a "Cloud Tiering" license) is a *floating* license that you can use across multiple on-premises ONTAP clusters in your NetApp Console organization. The total tiering capacity defined in your Cloud Tiering license is shared among **all** of your on-premises clusters, making initial licensing and renewal easy. The minimum capacity for a tiering BYOL license starts at 10 TiB.

If you don't have a Cloud Tiering license, contact us to purchase one:

- Contact your NetApp sales representative
- Contact NetApp support.

Optionally, if you have an unassigned node-based license for Cloud Volumes ONTAP that you won't be using, you can convert it to a Cloud Tiering license with the same dollar-equivalence and the same expiration date. [Go here for details.](#)

You manage Cloud Tiering BYOL licenses in the Console. You can add new licenses and update existing licenses. [Learn how to manage licenses.](#)

Cloud Tiering BYOL licensing starting in 2021

The new **Cloud Tiering** license was introduced in August 2021 for tiering configurations that are supported within the NetApp Console using the Cloud Tiering service. The NetApp Console currently supports tiering to the following cloud storage: Amazon S3, Azure Blob storage, Google Cloud Storage, NetApp StorageGRID, and S3-compatible object storage.

The **FabricPool** license that you may have used in the past to tier on-premises ONTAP data to the cloud is being retained only for ONTAP deployments in sites that have no internet access (also known as "dark sites"), and for tiering configurations to IBM Cloud Object Storage. If you're using this type of configuration, you'll install a FabricPool license on each cluster using System Manager or the ONTAP CLI.



Note that tiering to StorageGRID does not require a FabricPool or Cloud Tiering license.

If you are currently using FabricPool licensing, you're not affected until your FabricPool license reaches its expiration date or maximum capacity. Contact NetApp when you need to update your license, or earlier to make sure there is no interruption in your ability to tier data to the cloud.

- If you're using a configuration that's supported in the Console, your FabricPool licenses will be converted to Cloud Tiering licenses and they'll appear in the Console. When those initial licenses expire, you'll need to update the Cloud Tiering licenses.
- If you're using a configuration that's not supported in the Console, then you'll continue using a FabricPool license. [See how to license tiering using System Manager.](#)

Here are some things you need to know about the two licenses:

Cloud Tiering license	FabricPool license
It is a <i>floating</i> license that you can use across multiple on-premises ONTAP clusters.	It is a per-cluster license that you purchase and license for every cluster.
It is registered in the NetApp Console.	It is applied to individual clusters using System Manager or the ONTAP CLI.

Cloud Tiering license	FabricPool license
Tiering configuration and management is done through the Cloud Tiering service in the NetApp Console.	Tiering configuration and management is done through System Manager or the ONTAP CLI.
Once tiering is configured, you can use the tiering service without a license for 30 days using the free trial.	Once configured, you can tier the first 10 TB of data for free.

Manage Cloud Tiering licenses

If your licensed term is nearing the expiration date, or if your licensed capacity is reaching the limit, you'll be notified in Cloud Tiering as well as in the Console.

You can update existing licenses, view license status, and add new licenses through the Console. [Learn about managing licenses](#).

Apply Cloud Tiering licenses to clusters in special configurations

ONTAP clusters in the following configurations can use Cloud Tiering licenses, but the license must be applied in a different manner than single-node clusters, HA-configured clusters, clusters in Tiering Mirror configurations, and MetroCluster configurations using FabricPool Mirror:

- Clusters that are tiered to IBM Cloud Object Storage
- Clusters that are installed in "dark sites"

Process for existing clusters that have a FabricPool license

When you [discover any of these special cluster types in Cloud Tiering](#), Cloud Tiering recognizes the FabricPool license and adds the license to the Console. Those clusters will continue tiering data as usual. When the FabricPool license expires, you'll need to purchase a Cloud Tiering license.

Process for newly created clusters

When you discover typical clusters in Cloud Tiering, you'll configure tiering using the Cloud Tiering interface. In these cases the following actions happen:

1. The "parent" Cloud Tiering license tracks the capacity being used for tiering by all clusters to make sure there is enough capacity in the license. The total licensed capacity and expiration date are shown in the Console.
2. A "child" tiering license is automatically installed on each cluster to communicate with the "parent" license.

 The licensed capacity and expiration date shown in System Manager or in the ONTAP CLI for the "child" license is not the real information, so don't be concerned if the information is not the same. These values are managed internally by the Cloud Tiering software. The real information is tracked in the Console.

For the two configurations listed above, you'll need to configure tiering using System Manager or the ONTAP CLI (not by using the Cloud Tiering interface). So in these cases you'll need to push the "child" license to these clusters manually from the Cloud Tiering interface.

Note that since data is tiered to two different object storage locations for Tiering Mirror configurations, you'll


need to purchase a license with enough capacity for tiering data to both locations.

Steps

1. Install and configure your ONTAP clusters using System Manager or the ONTAP CLI.

Do not configure tiering at this point.

2. [Purchase a Cloud Tiering license](#) for the capacity needed for the new cluster, or clusters.
3. In the Console [add the license to the digital wallet](#)[add the license].
4. In Cloud Tiering, [discover the new clusters](#).
5. From the Clusters page, select **...** for the cluster and select **Deploy License**.



6. In the *Deploy License* dialog, select **Deploy**.

The child license is deployed to the ONTAP cluster.

7. Return to System Manager or the ONTAP CLI and set up your tiering configuration.

[FabricPool Mirror configuration information](#)

[FabricPool MetroCluster configuration information](#)

[Tiering to IBM Cloud Object Storage information](#)

NetApp Cloud Tiering technical FAQ

This FAQ can help if you're just looking for a quick answer to a question regarding NetApp Cloud Tiering.

Cloud Tiering service

The following FAQs relate to how Cloud Tiering works.

What are the benefits of using the Cloud Tiering service?

Cloud Tiering addresses the challenges that come with rapid data growth, providing you with benefits such as:

- Effortless data center extension to the cloud, providing up to 50x more space
- Storage optimization, yielding an average storage savings of 70%
- Reduced total cost of ownership by 30%, on average
- No need to refactor applications

What kind of data is useful to tier to the cloud?

Essentially, any data that is considered inactive on both primary and secondary storage systems is a good target to move to the cloud. On primary systems, such data can include snapshots, historical records, and finished projects. On secondary systems, this includes all volumes that contain copies of primary data made for DR and backup purposes.

Can I tier data from both NAS volumes and SAN volumes?

Yes, you can tier data from NAS volumes to the public cloud or to private clouds, like StorageGRID. When you tier data that is accessed by SAN protocols, NetApp recommends using private clouds because SAN protocols are more sensitive to connectivity issues than NAS.

What is the definition of inactive data or infrequently used data, and how is that controlled?

The definition of what can also be referred to cold data is: "volume blocks (metadata excluded) that have not been accessed for some amount of time". The "amount of time" is determined by a tiering policy attribute named cooling-days.

Will Cloud Tiering retain my storage efficiency savings in the cloud tier?

Yes, the ONTAP volume-level storage efficiencies such as compression, deduplication, and compaction are preserved when moving data to the cloud tier.

What is the difference between FabricPool and Cloud Tiering?

FabricPool is the ONTAP tiering technology that can be self-managed through the ONTAP CLI and System Manager, or managed as-a-service through Cloud Tiering. Cloud Tiering turns FabricPool into a managed service with advanced automation processes, on both ONTAP and in the cloud, providing greater visibility and control over tiering across hybrid and multi-cloud deployments.

Can the data tiered to the cloud be used for disaster recovery or for backup/archive?

No. Since the volume's metadata is never tiered from the performance tier, the data stored in object storage cannot be accessed directly.

However, Cloud Tiering can be used to achieve cost-effective backup and DR by enabling it on secondary systems and SnapMirror destination volumes (DP volumes), to tier off all of the data (metadata excluded), thus reducing your data center footprint and TCO.

Is Cloud Tiering applied at the volume or aggregate level?

Cloud Tiering is enabled at the volume level by associating a tiering policy with each volume. Cold data identification is done at the block level.

How does Cloud Tiering determine which blocks to tier to the cloud?

The tiering policy associated with the volume is the mechanism that controls which blocks are tiered and when. The policy defines the type of data blocks (snapshots, user data, or both) and the cooling period. See [Volume Tiering Policies](#) for details.

How does Cloud Tiering affect the volume capacity?

Cloud Tiering has no effect on the volume's capacity but rather on the aggregate's performance tier usage.

Does Cloud Tiering enable Inactive Data Reporting?

Yes, Cloud Tiering enables Inactive Data Reporting (IDR) on each aggregate. This setting enables us to identify the amount of inactive data that can be tiered to low-cost object storage.

How long does it take IDR to show information from the moment I start running it?

IDR starts showing information after the configured cooling period has passed. Using ONTAP 9.7 and earlier, IDR had a non-adjustable cooling period of 31 days. Starting with ONTAP 9.8, the IDR cooling-period can be configured up to 183 days.

Licenses and Costs

The following FAQs relate to licensing and costs to use Cloud Tiering.

How much does using Cloud Tiering cost?

When tiering cold data to the public cloud:

- For the pay-as-you-go (PAYGO), usage-based subscription: \$0.05 per GB/Month.
- For the annual (BYOL), term-based subscription: starting from \$0.033 per GB/Month.

[See pricing details.](#)

When tiering cold data to a NetApp StorageGRID system (private cloud) there is no cost.

Can I have both a BYOL and PAYGO license for the same ONTAP cluster?

Yes. Cloud Tiering allows you to use a BYOL license, a PAYGO subscription, or a combination of both.

What happens if I have reached the BYOL capacity limit, or if my BYOL license expires?

If you reach the BYOL capacity limit or if your BYOL license expires, tiering of new cold data stops, tiering of new cold data stops. All previously tiered data remains accessible; meaning you can retrieve and use this data. When retrieved, this data is moved back to the performance tier from the cloud.

However, if you have a PAYGO marketplace subscription to the *BlueXP - Deploy & Manage Cloud Data Services*, new cold data will continue to be tiered to object storage and you'll pay for those charges on a per-use basis.

Does the Cloud Tiering license include the egress charges from the cloud provider?

No, it does not.

Is rehydration of on-premises system subject to the egress cost charged by the cloud providers?

Yes. All reads from the public cloud are subject to egress fees.

How can I estimate my cloud charges? Is there a “what if” mode for Cloud Tiering?

The best way to estimate how much a cloud provider will charge for hosting your data is to use their calculators: [AWS](#), [Azure](#) and [Google Cloud](#).

Are there any extra charges by the cloud providers for reading/retrieving data from the object storage to the on-premises storage?

Yes. Check [Amazon S3 Pricing](#), [Block Blob Pricing](#), and [Cloud Storage Pricing](#) for additional pricing incurred with data reading/retrieving.

How can I estimate my volumes' savings and get a cold data report before I enable Cloud Tiering?

To get an estimate, add your ONTAP cluster to the NetApp Console and inspect it through the Cloud Tiering Clusters page. Select **Calculate potential tiering savings** for the cluster to launch the [Cloud Tiering TCO calculator](#) to see how much money you can save.

How am I charged for tiering when I am using an ONTAP MetroCluster?

When used in MetroCluster environments, total tiering license is applied to both clusters' usage. For example, if you have a license for 100TiB of tiering, each cluster's used tiering capacity contributes to the total capacity of 100TiB.

ONTAP

The following questions relate to ONTAP.

Which ONTAP versions does Cloud Tiering support?

Cloud Tiering supports ONTAP version 9.2 and higher.

What types of ONTAP systems are supported?

Cloud Tiering is supported with single-node and high-availability AFF, FAS, and ONTAP Select clusters. Clusters in FabricPool Mirror configurations and MetroCluster configurations are also supported.

Can I tier data from FAS systems with HDDs only?

Yes, starting ONTAP 9.8 you can tier data from volumes hosted on HDD aggregates.

Can I tier data from an AFF joined to a cluster that has FAS nodes with HDDs?

Yes. Cloud Tiering can be configured to tier volumes hosted on any aggregate. The data tiering configuration is irrelevant to the type of controller used and whether the cluster is heterogeneous or not.

What about Cloud Volumes ONTAP?

If you have Cloud Volumes ONTAP systems, you'll find them in the Cloud Tiering Clusters page so you get a full view of data tiering in your hybrid cloud infrastructure. However, Cloud Volumes ONTAP systems are read-only from Cloud Tiering. You can't set up data tiering on Cloud Volumes ONTAP from Cloud Tiering. [You set up tiering for Cloud Volumes ONTAP systems from the ONTAP system in the NetApp Console.](#)

What other requirements are necessary for my ONTAP clusters?

It depends on where you tier the cold data. Refer to the following links for more details:

- [Tiering data to Amazon S3](#)
- [Tiering data to Azure Blob storage](#)

- [Tiering data to Google Cloud Storage](#)
- [Tiering data to StorageGRID](#)
- [Tiering data to S3 object storage](#)

Object storage

The following questions relate to object storage.

Which object storage providers are supported?

Cloud Tiering supports the following object storage providers:

- Amazon S3
- Microsoft Azure Blob
- Google Cloud Storage
- NetApp StorageGRID
- S3-compatible object storage (for example, MinIO)
- IBM Cloud Object Storage (the FabricPool configuration must be done using System Manager or the ONTAP CLI)

Can I use my own bucket/container?

Yes, you can. When you set up data tiering, you have the choice to add a new bucket/container or to select an existing bucket/container.

Which regions are supported?

- [Supported AWS regions](#)
- [Supported Azure regions](#)
- [Supported Google Cloud regions](#)

Which S3 storage classes are supported?

Cloud Tiering supports data tiering to the *Standard*, *Standard-Infrequent Access*, *One Zone-Infrequent Access*, *Intelligent Tiering*, and *Glacier Instant Retrieval* storage classes. See [Supported S3 storage classes](#) for more details.

Why are Amazon S3 Glacier Flexible and S3 Glacier Deep Archive not supported by Cloud Tiering?

The main reason Amazon S3 Glacier Flexible and S3 Glacier Deep Archive aren't supported is that Cloud Tiering is designed as a high-performance tiering solution, so data must be continuously available and quickly accessible for retrieval. With S3 Glacier Flexible and S3 Glacier Deep Archive, data retrieval can last anywhere between a few minutes to 48 hours.

Can I use other S3-compatible object storage services, such as MinIO, with Cloud Tiering?

Yes, configuring S3-compatible object storage through the Tiering UI is supported for clusters using ONTAP 9.8 and later. [See the details here](#).

Which Azure Blob access tiers are supported?

Cloud Tiering supports data tiering to the *Hot* or *Cool* access tiers for your inactive data. See [Supported Azure Blob access tiers](#) for more details.

Which storage classes are supported for Google Cloud Storage?

Cloud Tiering supports data tiering to the *Standard*, *Nearline*, *Coldline*, and *Archive* storage classes. See [Supported Google Cloud storage classes](#) for more details.

Does Cloud Tiering support the use of lifecycle management policies?

Yes. You can enable lifecycle management so that Cloud Tiering transitions data from the default storage class/access tier to a more cost-effective tier after a certain number of days. The lifecycle rule is applied to all objects in the selected bucket for Amazon S3 and Google Cloud storage, and to all containers in the selected storage account for Azure Blob.

Does Cloud Tiering use one object store for the entire cluster or one per aggregate?

In a typical configuration there is one object store for the entire cluster. Starting in August 2022, you can use the **Advanced Setup** page to add additional object stores for a cluster, and then attach different object stores to different aggregates, or attach 2 object stores to an aggregate for mirroring.

Can multiple buckets be attached to the same aggregate?

It is possible to attach up to two buckets per aggregate for the purpose of mirroring, where cold data is synchronously tiered to both buckets. The buckets can be from different providers and different locations. Starting in August 2022, you can use the **Advanced Setup** page to attach two object stores to a single aggregate.

Can different buckets be attached to different aggregates in the same cluster?

Yes. The general best practice is to attach a single bucket to multiple aggregates. However, when using the public cloud there is a maximum IOPS limitation for the object storage services, therefore multiple buckets must be considered.

What happens with the tiered data when you migrate a volume from one cluster to another?

When migrating a volume from one cluster to another, all the cold data is read from the cloud tier. The write location on the destination cluster depends on whether tiering was enabled and the type of tiering policy used on the source and destination volumes.

What happens with the tiered data when you move a volume from one node to another in the same cluster?

If the destination aggregate does not have an attached cloud tier, data is read from the cloud tier of the source aggregate and written entirely to the local tier of the destination aggregate. If the destination aggregate has an attached cloud tier, data is read from the cloud tier of the source aggregate and first written to the local tier of the destination aggregate, to facilitate quick cutover. Later, based on the tiering policy used, it is written to the cloud tier.

Starting with ONTAP 9.6, if the destination aggregate is using the same cloud tier as the source aggregate, the cold data does not move back to the local tier.

How can I bring my tiered data back on-premises to the performance tier?

Write back is generally performed on reads and depends on the tiering policy type. Prior to ONTAP 9.8, writing back of the entire volume can be done with a *volume move* operation. Starting with ONTAP 9.8, the Tiering UI has options to **Bring back all data** or **Bring back active file system**. [See how to move data back to the performance tier.](#)

When replacing an existing AFF/FAS controller with a new one, would the tiered data be migrated back on-prem?

No. During the “head swap” procedure, the only thing that changes is the aggregate’s ownership. In this case, it will be changed to the new controller without any data movement.

Can I use the cloud provider’s console or object storage explorers to look at the data tiered to a bucket? Can I use the data stored in the object storage directly without ONTAP?

No. The objects constructed and tiered to the cloud do not contain a single file but up to 1,024 4KB blocks from multiple files. A volume’s metadata always remains on the local tier.

Console agents

The following questions relate to the Console agent.

What is the Console agent?

The Console agent is software running on a compute instance either within your cloud account, or on-premises, that enables the NetApp Console to securely manage cloud resources. To use the Cloud Tiering service, you must deploy an agent.

Where does the Console agent need to be installed?

- When you tier data to S3, the agent can reside in an AWS VPC or on your premises.
- When you tier data to Blob storage, the agent can reside in an Azure VNet or on your premises.
- When you tier data to Google Cloud Storage, the agent must reside in a Google Cloud Platform VPC.
- When you tier data to StorageGRID or other S3-Compatible storage providers, the agent must reside on your premises.

Can I deploy the Console agent on-premises?

Yes. The agent software can be downloaded and manually installed on a Linux host in your network. [See how to install the agent in your premises.](#)

Is an account with a cloud service provider required before using Cloud Tiering?

Yes. You must have an account before you can define the object storage that you want to use. An account with a cloud storage provider is also required when setting up the agent in the cloud on a VPC or VNet.

What are the implications if the Console agent fails?

In the case of an agent failure, only the visibility into the tiered environments is impacted. All the data is accessible and newly identified cold data is automatically tiered to object storage.

Tiering policies

What are the available tiering policies?

There are four tiering policies:

- None: Classifies all data as always hot; preventing any data from the volume being moved to object storage.
- Cold Snapshots (Snapshot-only): Only cold snapshot blocks are moved to object storage.
- Cold User Data and Snapshots (Auto): Both cold snapshot blocks and cold user data blocks are moved to object storage.
- All User Data (All): Classifies all data as cold; immediately moving the entire volume to object storage.

[Learn more about Tiering Policies.](#)

At which point is my data is considered cold?

Since data tiering is done at the block level, a data block is considered cold after it hasn't been accessed for a certain period of time, which is defined by the tiering policy's minimum-cooling-days attribute. The applicable range is 2-63 days with ONTAP 9.7 and earlier, or 2-183 days starting with ONTAP 9.8.

What is the default cooling period for data before it is tiered to the cloud tier?

The default cooling period for the Cold Snapshot policy is 2 days, while the default cooling period for Cold User Data and Snapshots is 31 days. The cooling-days parameter is not applicable to the All tiering policy.

Is all the tiered data retrieved from object storage when I do a full backup?

During full backup all the cold data is read. The retrieval of the data depends on the tiering policy used. When using the All and Cold User Data and Snapshots policies, the cold data is not written back to the performance tier. When using the Cold Snapshots policy, only in case of an old snapshot being used for the backup will its cold blocks be retrieved.

Can you choose a tiering size per volume?

No. However, you can choose which volumes are eligible for tiering, the type of data to be tiered, and its cooling period. This is done by associating a tiering policy with that volume.

Is the All User Data policy the only option for data protection volumes?

No. Data protection (DP) volumes can be associated with any of the three policies available. The type of policy used on the source and destination (DP) volumes determines the write location of the data.

Does resetting the tiering policy of a volume to None rehydrate the cold data or just prevent future cold blocks from being moved to the cloud?

No rehydration takes place when a tiering policy is reset, but it will prevent new cold blocks from being moved to the cloud tier.

After tiering data to the cloud, can I change the tiering policy?

Yes. The behavior after the change depends on the new associated policy.

What should I do if I want to ensure certain data is not moved to the cloud?

Do not associate a tiering policy with the volume containing that data.

Where is the metadata of the files stored?

The metadata of a volumes is always stored locally, on the performance tier — it is never tiered to the cloud.

Networking and security

The following questions relate to networking and security.

What are the networking requirements?

- The ONTAP cluster initiates an HTTPS connection over port 443 to your object storage provider.
ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.
- For StorageGRID, the ONTAP cluster initiates an HTTPS connection over a user-specified port to StorageGRID (the port is configurable during tiering setup).
- An agent needs an outbound HTTPS connection over port 443 to your ONTAP clusters, to the object store, and to the Cloud Tiering service.

For more details, see:

- [Tiering data to Amazon S3](#)
- [Tiering data to Azure Blob storage](#)
- [Tiering data to Google Cloud Storage](#)
- [Tiering data to StorageGRID](#)
- [Tiering data to S3 object storage](#)

What tools can I use for monitoring and reporting in order to manage cold data stored in the cloud?

Other than Cloud Tiering, [Active IQ Unified Manager](#) and [digital advisor](#) can be used for monitoring and reporting.

What are the implications if the network link to the cloud provider fails?

In case of a network failure, the local performance tier remains online and hot data remains accessible. However, blocks that were already moved to the cloud tier will be inaccessible and applications will receive an error message when trying to access that data. Once connectivity is restored, all data will be seamlessly accessible.

Is there a network bandwidth recommendation?

The underlying FabricPool tiering technology read latency depends on connectivity to the cloud tier. Although tiering works on any bandwidth, it is recommended to place intercluster LIFs on 10 Gbps ports to provide adequate performance. There are no recommendations or bandwidth limitations for the agent.

Additionally, you can throttle the amount of network bandwidth that is used during the transfer of inactive data from the volume to object storage. The *Maximum transfer rate* setting is available when configuring your

cluster for tiering, and afterwards from the **Clusters** page.

Is there any latency when a user attempts to access tiered data?

Yes. Cloud tiers cannot provide the same latency as the local tier since latency depends on the connectivity. To estimate the latency and throughput of an object store, Cloud Tiering provides a Cloud Performance Test (based on the ONTAP object store profiler) that can be used after the object store is attached and before tiering is set up.

How is my data secured?

AES-256-GCM encryption is maintained on both the performance and cloud tiers. TLS 1.2 encryption is used to encrypt data over the wire as it moves between tiers, and to encrypt communication between the agent and both the ONTAP cluster and the object store.

Do I need an Ethernet port installed and configured on my AFF?

Yes. An intercluster LIF must be configured on an ethernet port, on each node within an HA pair that hosts volumes with data you plan to tier to the cloud. For more information, see the Requirements section for the cloud provider where you plan to tier data.

What permissions are required?

- [For Amazon, permissions are required to manage the S3 bucket.](#)
- For Azure, no extra permissions are needed outside of the permissions that you need to provide to the NetApp Console.
- For Google Cloud, Storage Admin permissions are needed for a service account that has storage access keys.
- For StorageGRID, S3 permissions are needed.
- For S3-compatible object storage, S3 permissions are needed.

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—with 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.