



Cloud Tiering documentation

Cloud Tiering

NetApp
June 10, 2024

Table of Contents

- Cloud Tiering documentation 1
 - Discover what's new 1
 - Learn about Cloud Tiering 1
 - Get started 1
 - Automate with APIs 1
 - Get help and connect with peers 1
- What's new in Cloud Tiering 2
 - 7 Sept 2020 2
 - 5 Aug 2020 2
 - 8 July 2020 3
 - 8 June 2020 5
 - 18 May 2020 6
 - 1 Apr 2020 6
 - 25 Dec 2019 7
 - 3 Nov 2019 7
 - 8 Sept 2019 8
 - 7 Aug 2019 9
 - 4 July 2019 9
 - 10 June 2019 9
 - 5 May 2019 10
 - 7 Apr 2019 11
- Concepts 13
 - Cloud Tiering overview 13
 - How Cloud Tiering works 13
 - How licensing works 16
 - Savings opportunities 17
 - Accounts 18
- Get started 20
 - Supported object storage providers 20
 - Tier data to AWS S3 20
 - Tier data to Azure Blob storage 32
 - Tier data to Google Cloud Storage 41
 - Tier data to StorageGRID 49
 - Set up licensing for Cloud Tiering 56
- Managing data tiering from your clusters 59
 - Discovering additional clusters 59
 - Tiering data from additional volumes 60
 - Changing a volume's tiering policy and cooling period 60
 - Managing tiering settings on aggregates 61
 - Reviewing tiering info for a cluster 62
 - Fixing operational health 64
 - Removing a failed cluster 64
- How to get help and find more information 65

- Cloud Tiering APIs 66
 - Getting started 66
 - API reference 66
- Cloud Tiering technical FAQ 67
 - ONTAP 67
 - Object storage 67
 - NetApp Service Connector 68
 - Networking 69
 - Permissions 69
- Legal notices 71
 - Copyright 71
 - Trademarks 71
 - Patents 71
 - Privacy policy 71
 - Open source 71

Cloud Tiering documentation

Based on NetApp FabricPool technology, Cloud Tiering identifies infrequently-used data in your ONTAP clusters and automatically and seamlessly moves that data to low-cost object storage in the cloud.

Discover what's new

[What's new in Cloud Tiering](#)

Learn about Cloud Tiering

- [What it is](#)
- [How it works](#)
- [Licensing](#)

Get started

- [Quick start for AWS](#)
- [Quick start for Azure](#)
- [Quick start for GCP](#)
- [Quick start for StorageGRID](#)

Automate with APIs

- [Get started with APIs](#)
- [API reference](#)

Get help and connect with peers

[NetApp Community: Cloud Data Services](#)

What's new in Cloud Tiering

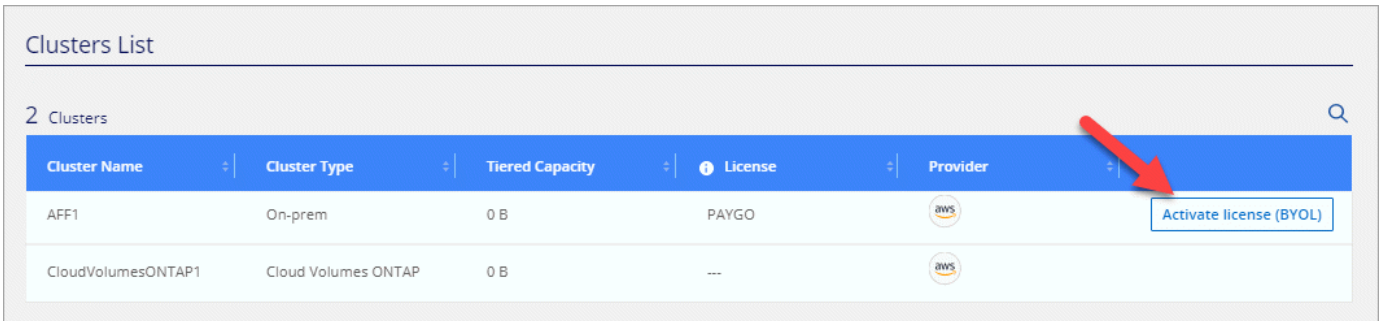
NetApp periodically updates Cloud Tiering to bring you new features, enhancements, and bug fixes.

7 Sept 2020

This Cloud Tiering update includes the following enhancements.



Activate licenses (BYOL)

You can now activate a FabricPool license on an ONTAP cluster directly from Cloud Tiering. [Learn more.](#)



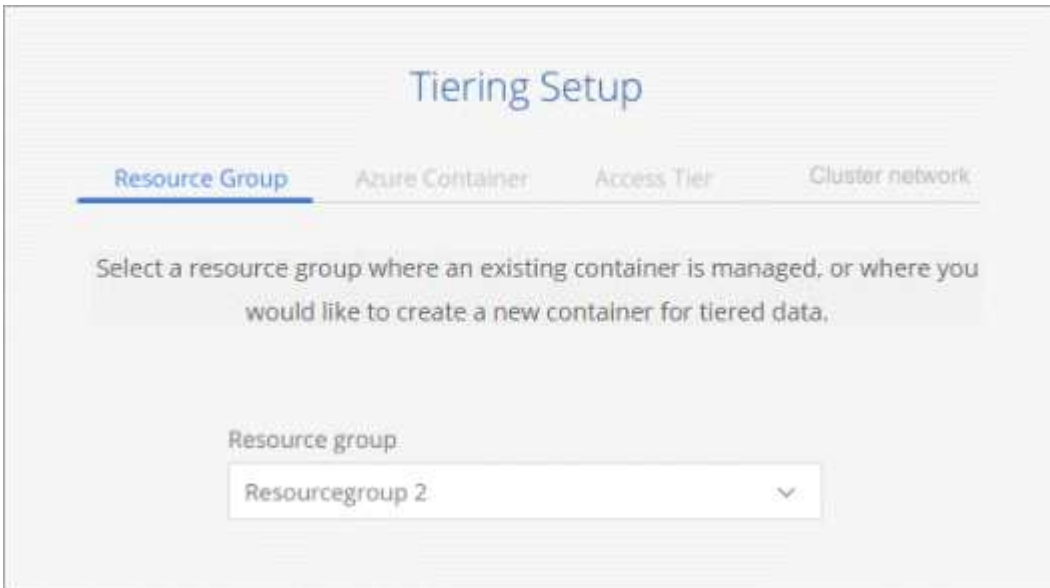
Clusters List

2 Clusters

Cluster Name	Cluster Type	Tiered Capacity	License	Provider	
AFF1	On-prem	0 B	PAYGO		Activate license (BYOL)
CloudVolumesONTAP1	Cloud Volumes ONTAP	0 B	---		

Azure resource group selection

A new step is available when you set up data tiering to Azure Blob storage. You now select a resource group first and then choose a Blob container.



Tiering Setup

Resource Group Azure Container Access Tier Cluster network

Select a resource group where an existing container is managed, or where you would like to create a new container for tiered data.

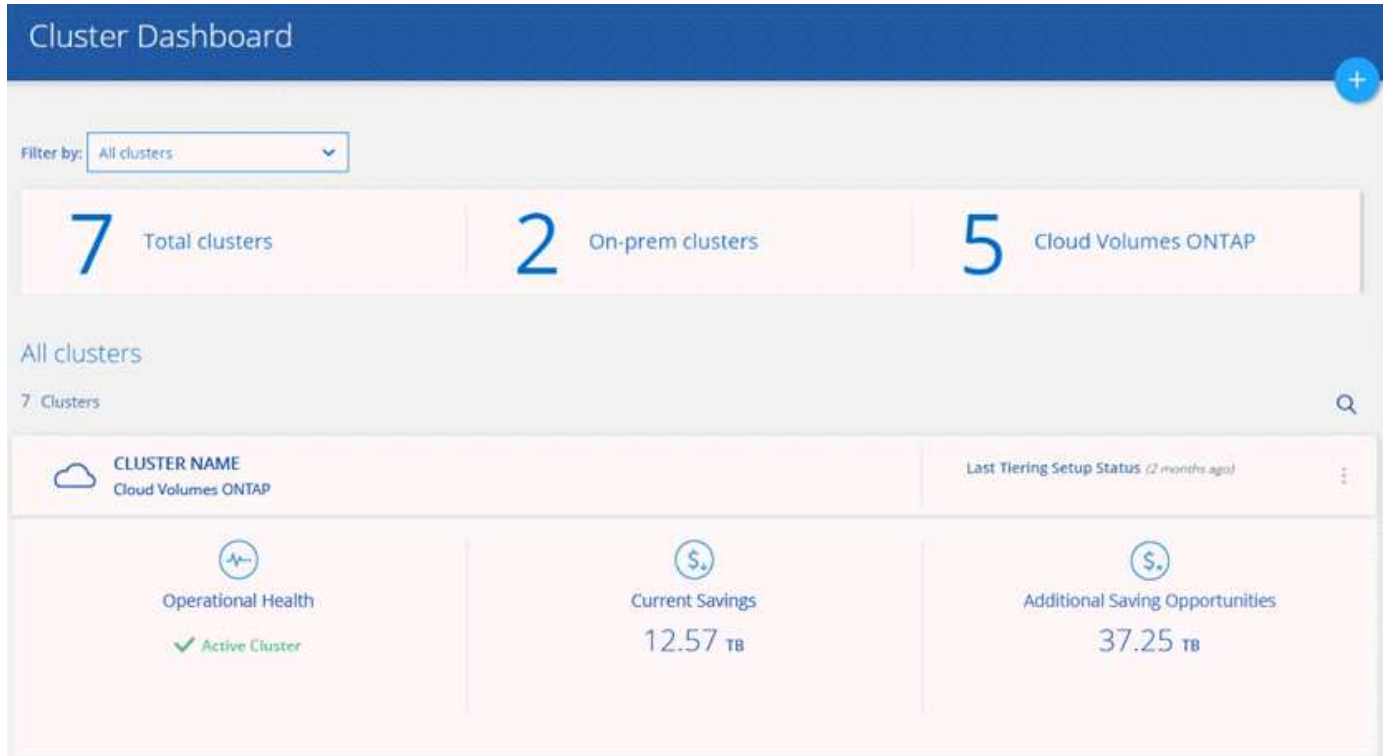
Resource group

Resourcegroup 2

5 Aug 2020

If you have Cloud Volumes ONTAP systems, you'll find them in the Cluster Dashboard so you see a full view of data tiering in your hybrid cloud infrastructure.

From the Cluster Dashboard, you can view tiering information similar to an on-prem ONTAP cluster: operational health, current savings, savings opportunities, details about volumes and aggregates, and more.



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'll still set up tiering the same way: from the working environment in Cloud Manager.

8 July 2020

This Cloud Tiering update includes the following new features and enhancements.

Tiering settings for aggregates

Each aggregate has two tiering settings that you can now adjust from Cloud Tiering: the tiering fullness threshold and whether inactive data reporting is enabled. [Learn more.](#)



Inactive data reporting for discovered clusters

Inactive data reporting is now automatically enabled on ONTAP clusters that you discover through Cloud Manager. This enhancement makes it easier for Cloud Tiering to show you the potential savings from tiering cold data.

Remove clusters

If the health of a cluster is failed, you can now remove it from the dashboard.

2 Clusters

Unknown Cluster

Last Tiering Setup Status
Tiering hasn't been setup through Cloud Tiering

Operational Health
× Failed ⓘ

Current Savings
No Available Data

Saving Opportunities
No Available Data

Remove Cluster

Discovery of unsupported clusters

Cloud Tiering now provides details about discovered clusters that aren't eligible for tiering.

ONPREM_1

⚠ Not Eligible for Tiering
You need to upgrade ONTAP to version 9.4 or later.

Current Savings
No Available Data

Additional Saving Opportunities
37.25 TB

8 June 2020

This Cloud Tiering update includes the following new features and enhancements.

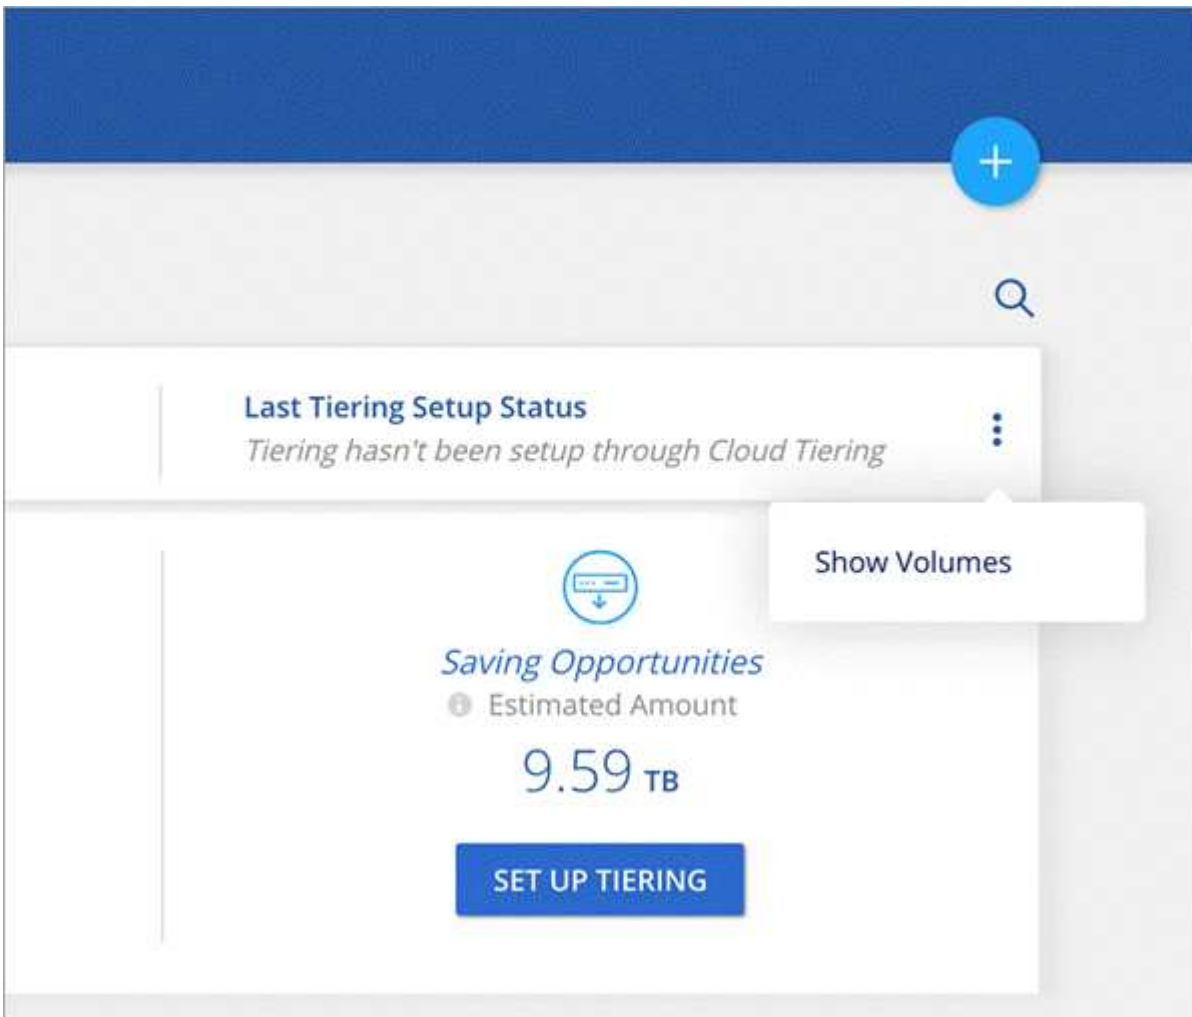
Support for StorageGRID

You can now tier inactive data from your ONTAP clusters to StorageGRID.

[Learn how to get started.](#)

View volumes before setup

You can now view details about a cluster's volumes before you set up tiering on the volumes.



Multi-tenancy using Cloud Central accounts

Cloud Tiering now supports multi-tenancy through Cloud Central accounts. Accounts enable multiple users to manage the same clusters in an account.

[Learn more about accounts.](#)

18 May 2020

This update includes the following new features and enhancements.

Cooling days

Cloud Tiering now enables you to adjust the cooling period that's associated with a volume tiering policy. The cooling period is the number of days that user data in a volume must remain inactive before it's considered "cold" and moved to object storage. Adjusting the cooling days enables you to optimize tiering for your environment.

[Learn how to adjust the cooling period.](#)

Rediscovery of clusters

We've made improvements to the Cloud Tiering service that requires you to deploy new Service Connectors and rediscover your clusters.

The benefits include improved performance, the ability to release new features faster, improved error handling, an improved discovery process, and more.

Cloud Tiering will prompt you to rediscover your clusters when you log in. If you need help, contact us by using the in-product chat.

1 Apr 2020

We're pleased to announce the following new features and enhancements.

Simplified cluster setup

We enhanced how you set up data tiering on a new cluster. The simplified wizard walks you through three steps: discover the cluster, set up the cluster for tiering, and then choose the volumes that you want to tier. The following pages provide more details:

- [Tiering data to AWS S3](#)
- [Tiering data to Azure Blob storage](#)
- [Tiering data to Google Cloud storage](#)

Ability to add or select an existing bucket/container

When you set up data tiering on a cluster, you now have the choice to add a new bucket/container or to select an existing bucket/container. Prior to this change, Cloud Tiering automatically created it for you.

Support for additional S3 storage classes

Cloud Tiering now supports two additional S3 storage classes: *Zone-IA* and *Intelligent*.

When you set up data tiering, Cloud Tiering can apply a lifecycle rule so the tiered data transitions from the *Standard* storage class to another storage class after 30 days. [Learn more.](#)

25 Dec 2019

This update includes the following new features and enhancements.

Changing a volume's tiering policy

You can now change a volume's tiering policy after you set the initial tiering policy. For example, you can easily change from *Cold snapshots* to *Cold user data*.

[Learn how to change a volume's tiering policy.](#)

Licensing enhancements

- Support for ONTAP FabricPool licenses

Cloud Tiering now supports the FabricPool licenses that you add to an ONTAP cluster. If a license is available on the cluster, then Cloud Tiering displays the license in the **Licensing** page and enables you to tier data to the cloud based on that license.

- Support for pay-as-you-go from the Azure Marketplace

You can now pay-as-you-go when tiering cold data to Azure Blob storage.

- Support for pay-as-you-go from the GCP Marketplace

You can now pay-as-you-go when tiering cold data to Google Cloud storage.

- Support for a combination of PAYGO and BYOL licensing

You can combine a FabricPool license with a pay-as-you-go subscription to pay for data that you tier to the cloud. If your FabricPool license expires or if you tier more data than allowed by the license, then data tiering is never interrupted—it continues through your PAYGO subscription.

For more details:

- [Learn how licensing works](#)
- [Learn how to set up pay-as-you-go subscriptions and BYOL](#)

3 Nov 2019

This update includes the following new features and enhancements.

Support for Google Cloud

You can now tier inactive data from your ONTAP clusters to Google Cloud Storage.

[Learn how to get started](#)

Support for additional tiering policies

You can now select the following tiering policies when setting up tiering:

- All user data (All)
- All DP user data (Backup)

[Learn about these tiering policies.](#)

Tiering policy per volume

You can now choose a different tiering policy for each volume when you set up tiering.

Tier Volumes Learn how much you can save with each Tiering Policy

1 - 3 of 3 Volumes 🔍

Volume Name ↑	SVM Name ↑	Volume Size ↑	Used Size ↑	Cold Data ↑	Tier Status [3] ▾	Tiering Policy ↑
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ... 70 %	✓ Tiered Volume	All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ... 70 %	✓ Tiered Volume	Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B 0 %	✓ Tiered Volume	Cold snapshots

8 Sept 2019

This Cloud Tiering update includes the following new features and enhancements.

Tiering setup status

Cloud Tiering now shows the status of tiering setup for each cluster. For example, the status might indicate that Cloud Tiering is setting up the object store, or that it successfully set up 145 volumes for data tiering. The status also identifies if any failures occurred during setup.

AFF-2 Tiering Setup Status (Last Operation A few seconds ago)
✓ 2 Volumes Tiered Successfully

[More Info >](#)

Operational Health

✓ Active Cluster

Current Savings

Cold data has not been tiered yet

Saving Opportunities

1 TB

Set up Tiering

Integration with Cloud Manager AWS subscriptions

If you use NetApp Cloud Manager and you've already [subscribed through its new AWS Marketplace offering](#), then you're automatically subscribed to Cloud Tiering, as well. You'll see an active subscription in Cloud Tiering in the **Licensing** tab. You won't need to subscribe.

If you've already subscribed through Cloud Tiering, then this change has no impact to you. You're all set.

7 Aug 2019

This update includes the following new features and enhancements.

On-premises Service Connector for data tiering to S3


You can now install the Service Connector on an on-premises Linux host when you tier cold data to AWS S3. [Learn more.](#)

Auto discovery of clusters

If you discovered ONTAP clusters through NetApp Cloud Manager, the clusters are automatically added to Cloud Tiering if they support data tiering.

Object storage connectivity check

If Cloud Tiering finds a connectivity problem with the object storage bucket, the tiering health indicator in the dashboard provides details about the problem.



4 July 2019

Cloud Tiering was updated to fix a few bugs.

10 June 2019

This update includes the following new features and enhancements.

Cloud Tiering is now generally available

The Controlled Availability release of Cloud Tiering has completed—Cloud Tiering is now available for customer use from [NetApp Cloud Central](#). A 30-day free trial is available for both AWS and Azure. It starts when you set up tiering to your first cluster.

Pay-as-you-go from the AWS Marketplace

After your free trial starts, subscribe to the Cloud Tiering service to ensure that there's no disruption of service after the trial ends. When it ends, you'll be charged hourly according to the amount of data that you tier.

[Learn how to subscribe from the AWS Marketplace.](#)



We're planning to add Cloud Tiering to the Azure Marketplace as soon as Azure supports SaaS pricing.

Support for FlexGroup volumes

You can now tier inactive data from FlexGroup volumes to object storage, starting with ONTAP 9.5. Setup works the same as any other volume.

5 May 2019

This update includes the following new features and enhancements.


Support for Microsoft Azure

You can now tier inactive data from your ONTAP clusters to Azure Blob storage.

- [Learn how to tier inactive data to Azure](#)
- [Review support for Azure Blob access tiers and Azure regions](#)

Ability to choose an IPspace for connections to object storage

When you set up tiering for an ONTAP cluster, you now need to 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.



To understand the requirements for the IPspace and the associated intercluster LIFs, refer to ONTAP cluster requirements:

- [Preparing to tier inactive data to AWS S3](#)
- [Preparing to tier inactive data to Azure Blob storage](#)

7 Apr 2019

This update includes the following new features and enhancements.

- [Support for FAS systems with all-SSD aggregates](#)
- [Support for additional versions of ONTAP](#)
- [Ability to choose the type of cold data that you want to tier](#)
- [Ability to choose an S3 storage class](#)

Support for FAS systems with all-SSD aggregates

In addition to AFF systems, Cloud Tiering now supports FAS systems that have one or more all-SSD aggregates.

Support for additional versions of ONTAP

Cloud Tiering now supports ONTAP 9.2 and 9.3. This is in addition to supporting ONTAP 9.4 and later.

Ability to choose the type of cold data that you want to tier

For ONTAP 9.4 and later, you can now choose the type of cold data that you want to tier. You can tier *all cold data* or just *Snapshot copies*.

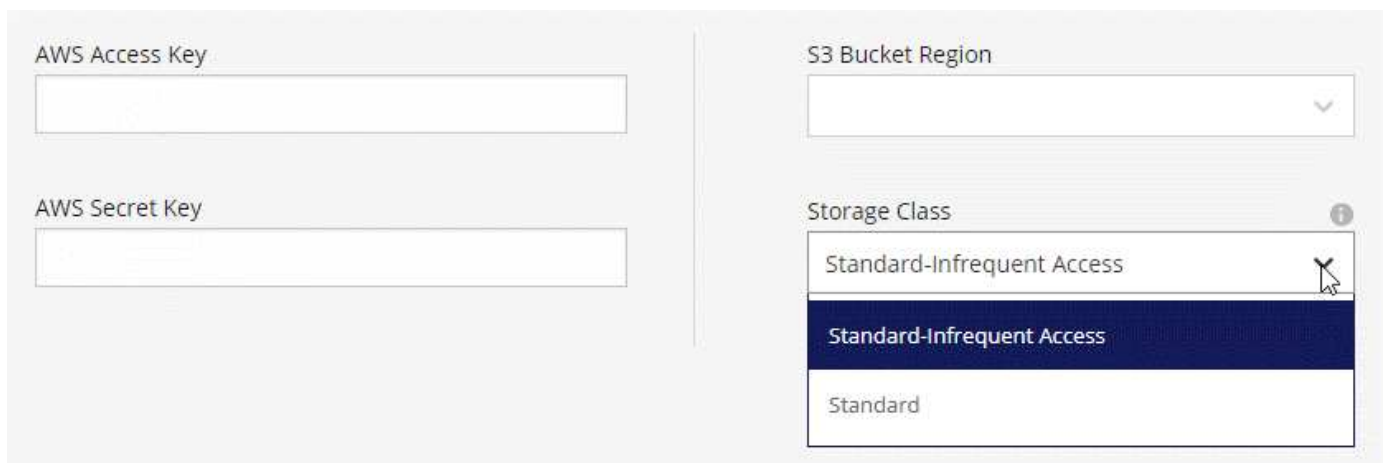
If you have an ONTAP 9.2 or 9.3 system, you can tier Snapshot copies only.

The option is available after you click **Set up Tiering**.

- [Learn more about the cold data that you can tier](#)
- [Learn how to tier data from your first cluster](#)
- [Learn how to tier data from additional volumes](#)

Ability to choose an S3 storage class

When you set up data tiering on a cluster for the first time, you can now choose to tier cold data to the S3 *Standard* storage class or to the *Standard-Infrequent Access* storage class. [Learn about these S3 storage classes](#).



The screenshot shows a configuration interface for S3 storage classes. On the left, there are two input fields: "AWS Access Key" and "AWS Secret Key". On the right, there are two dropdown menus: "S3 Bucket Region" and "Storage Class". The "Storage Class" dropdown is open, showing three options: "Standard-Infrequent Access" (highlighted in dark blue), "Standard-Infrequent Access" (in white), and "Standard" (in white). An information icon (i) is visible next to the "Storage Class" label.

Learn how to set up data tiering on a cluster.

Concepts

Cloud Tiering overview

NetApp's Cloud Tiering service extends your data center to the cloud by automatically tiering inactive data from ONTAP clusters to object storage. Without compromising on manageability and performance, Cloud Tiering efficiently manages your storage pool by seamlessly placing your data at the right tier at the right time, based on its usage.

The Cloud Tiering service leverages the capabilities of *FabricPool*. FabricPool is a NetApp Data Fabric technology that enables automated tiering of data to low-cost object storage. Active data remains on high-performance SSDs, while inactive data is tiered to low-cost object storage while preserving ONTAP data efficiencies.

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

- Automation makes it easier to set up and manage data tiering from ONTAP clusters to the cloud
- 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

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



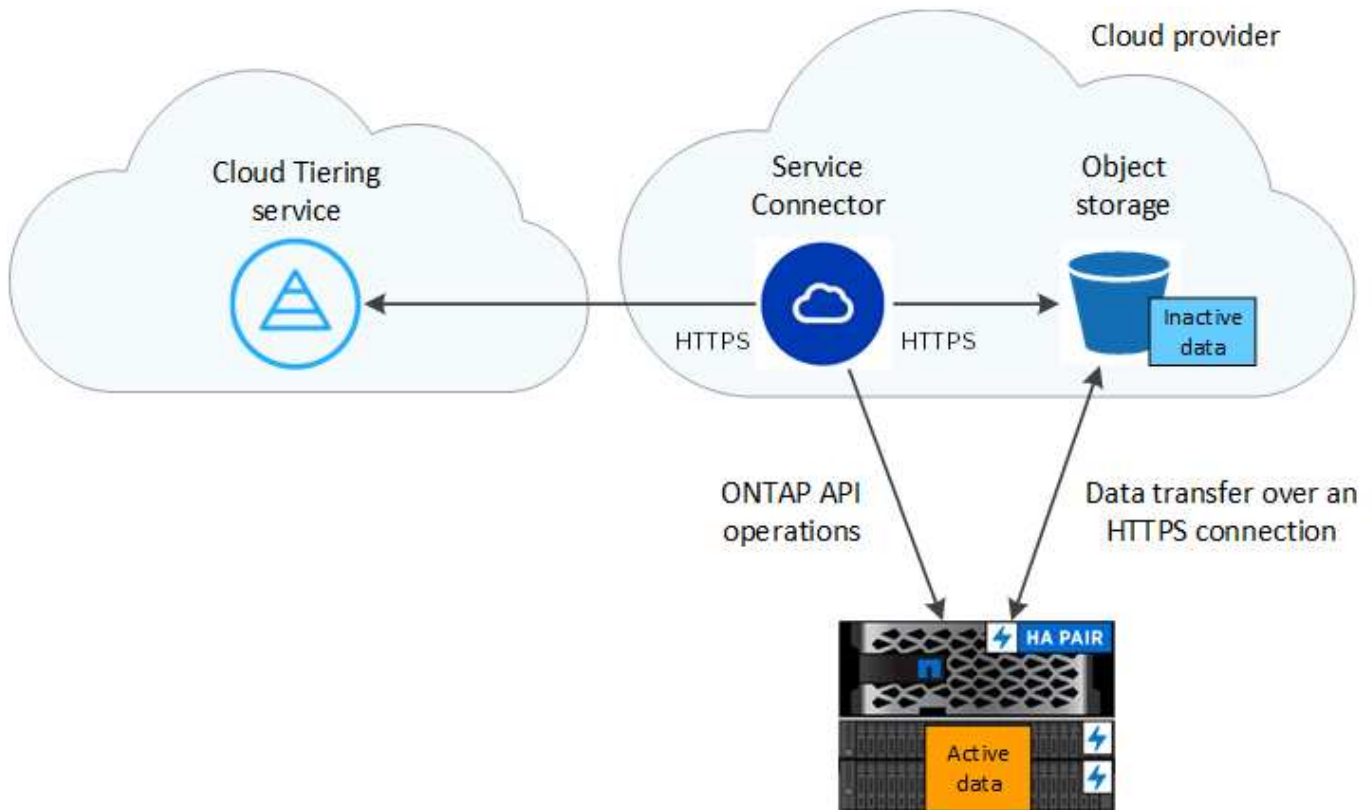
While Cloud Tiering can significantly reduce storage footprints, it is not a backup solution.

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 the NetApp Service Connector.

Overview

The following image shows the relationship between each component:



At a high level, Cloud Tiers works like this:

1. You provide Cloud Tiers with details about your cluster and deploy a NetApp Service Connector.
 - When tiering to S3, the Service Connector can be in an AWS VPC or on your premises.
 - When tiering to Blob storage, the Service Connector must be in an Azure VNet.
 - When tiering to Google Cloud Storage, the Service Connector must reside in a Google Cloud Platform VPC.
 - When tiering to StorageGRID, the Service Connector must reside on your premises.
2. You provide details about your object storage, including the bucket/container and a storage class or access tier.
3. The Service Connector 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](#)).

NetApp Service Connector

The Service Connector is software that communicates with ONTAP clusters to discover the amount of active and inactive data on the cluster and to set up data tiering. Cloud Tiers prompts you to deploy the Service Connector when you discover your first ONTAP cluster. Connections to ONTAP take place from the Service Connector. A single Service Connector can discover multiple ONTAP clusters.

Object storage

Each ONTAP cluster tiers inactive data to a single object store. 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 when tiering to the public cloud.

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

Volume tiering policies

When you select the volumes that you want to tier, you also 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.

No tiering policy

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

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 the number of days.



Writes from the cloud tier to the performance tier are disabled if performance tier capacity is greater than 70%. If this occurs, blocks are accessed directly from the cloud tier.

Cold user data (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.

Data is tiered only after an aggregate has reached 50% capacity and when the data has reached the cooling period. The cooling period is the time that user data in a volume must remain inactive for the data to be considered "cold" and moved to the object store. The default number of cooling days is 31, but you can adjust the number of days.



Writes from the cloud tier to the performance tier are disabled if performance tier capacity is greater than 70%. If this occurs, blocks are accessed directly 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. Note that 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.

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.

How licensing works

Pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license called *FabricPool*, or a combination of both. A 30-day free trial is available for your first cluster if you don't have a license.



A license isn't required when tiering data to StorageGRID.

30-day free trial

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

If your free trial ends and you haven't subscribed or added a license, then ONTAP no longer tiers cold data to object storage, but existing data is still available for access.

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 FabricPool 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 FabricPool license.

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

Bring your own license

Bring your own license by purchasing an ONTAP FabricPool license from NetApp. You can purchase term-based or perpetual licenses.

After you purchase a FabricPool license, you'll need to add it to the cluster, [which you can do directly from Cloud Tiering](#).

After you activate the license through Cloud Tiering, if you purchase additional add-on capacity at a later time, the license on the cluster is automatically updated with the new capacity. There's no need to apply a new NetApp License File (NLF) to the cluster.

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

[Contact us to purchase a license.](#)

Savings opportunities

Cloud Tiering first shows you the *estimated* space that you can save on each cluster. 31 days after you discover the cluster, it shows you a precise value, if you're running ONTAP 9.4 or later.

The initial savings opportunity that Cloud Tiering displays is an estimate based on industry standards. We estimate that 70% of the data on the cluster is cold data that you can move to lower-cost object storage.

Cloud Tiering displays an estimate because a 31-day cooling period is needed to determine which data is considered inactive. For ONTAP 9.4 and later, Cloud Tiering replaces the estimate with a precise value after this 31-day cooling period.

Here's an example of an estimate showing 9.55 TB of capacity savings:

The screenshot displays a dashboard for cluster AFF-GCP. It includes a 'More Info' link, an 'Operational Health' section with a green checkmark and 'Active Cluster' status, a 'Current Savings' section with a message 'Cold data has not been tiered yet', and a 'Saving Opportunities' section showing an estimated amount of 9.55 TB and a 'Set up Tiering' button. A 'Last Tiering Setup Status' message indicates that tiering has not been setup through Cloud Tiering.

Operational Health	Current Savings	Saving Opportunities
Active Cluster	Cold data has not been tiered yet	Estimated Amount: 9.55 TB

Here's an example of a precise value for a different cluster that has not been set up for tiering:

AFF-2
[More Info >](#)

Last Tiering Setup Status
Tiering hasn't been setup through Cloud Tiering

Operational Health
✓ Active Cluster

Current Savings
Cold data has not been tiered yet

Saving Opportunities
1 TB
[Set up Tiering](#)

And here's an example of a precise value on a cluster that is tiering some of it's volumes, but not all:

AFF-3
[More Info >](#)

Last Tiering Setup Status *(A few seconds ago)*
✓ 1 Volume was successfully set up

Operational Health
✓ Active Cluster

Current Savings
2 TB | \$ 9,677

Additional Saving Opportunities
2 TB
[Set up Tiering](#)


Accounts

Each Cloud Central user is associated with one or more Cloud Central accounts. An account enables multi-tenancy: multiple users can manage data tiering in a single account.

For example, two users might be associated with the same Cloud Central account. Both of those users can see the same clusters that were discovered in that account.

But if those two users are associated with *separate* Cloud Central accounts, then the users would only see the clusters in the account that they are associated with.

If a user is associated with multiple Cloud Central accounts, you can change to a different account at any time from the User Settings menu in Cloud Tiering.



If you want to associate a user to a specific account, you can use Cloud Central APIs, Cloud Manager's user interface, or contact us for help using the in-product chat.

[NetApp Cloud Central Services API](#)

Get started

Supported object storage providers

Cloud Tiering can tier inactive data from an ONTAP cluster to a public cloud (AWS S3, Microsoft Azure Blob storage, or Google Cloud Storage) or to a private cloud (StorageGRID).

AWS S3

- [Quick start for AWS](#)
- [Supported S3 storage classes and regions](#)

Azure Blob storage

- [Quick start for Azure](#)
- [Supported Azure Blob access tiers and regions](#)

Google Cloud Storage

- [Quick start for Google Cloud Storage](#)
- [Supported Google Cloud storage classes and regions](#)

StorageGRID using the S3 protocol

[Quick start for StorageGRID](#)

Tier data to AWS S3

Quick start for tiering inactive data to AWS

Getting started with Cloud Tiering in AWS includes a few steps.



Prepare to tier data to AWS

You need the following:

- An AFF or FAS system with all-SSD aggregates running ONTAP 9.2 or later, and an HTTPS connection to AWS S3.
- An AWS account that has an access key and [the required permissions](#) so the ONTAP cluster can tier inactive data in and out of AWS S3.
- A location for the Service Connector: either [an AWS VPC](#) or [an on-premises Linux host](#).

With either option, the Service Connector needs an outbound HTTPS connection to the ONTAP cluster, to S3 storage, and to the Cloud Tiering service.

2

Tier inactive data from your first cluster

Log in to [NetApp Cloud Central](#), start a free trial of the Cloud Tiering service, and click **Let's Start, Discover Your First Cluster**.

3

Set up licensing

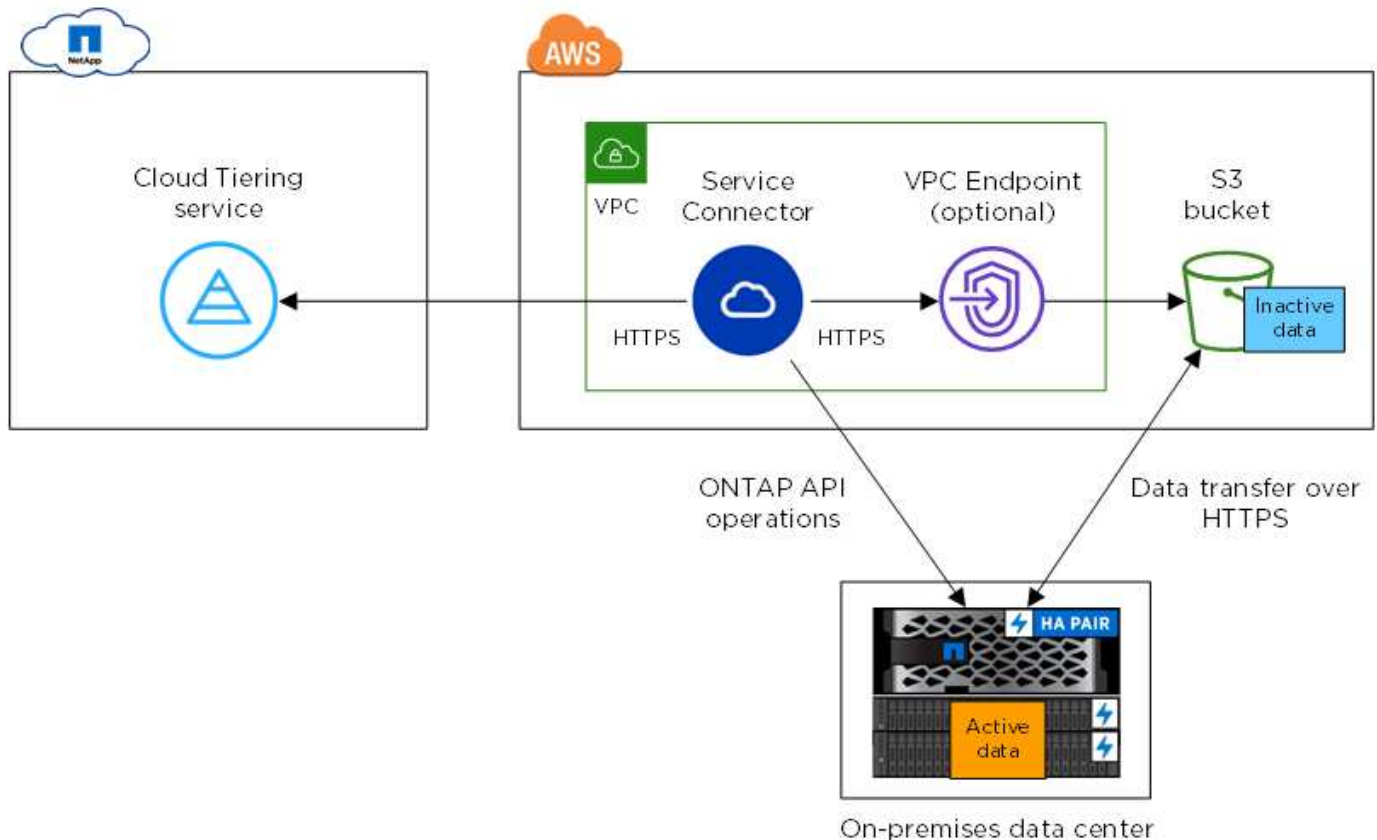
After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license, or a combination of both:

- To subscribe from the AWS Marketplace, click **Licensing**, click **Subscribe**, and then follow the prompts.
- To add a tiering license, [contact us if you need to purchase one](#), and then add it to your cluster from ONTAP System Manager.

Preparing to tier inactive data to AWS S3

Before you use Cloud Tiering, verify support for your ONTAP cluster, prepare your object storage, and set up a location for the Service Connector.

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



Communication between the Service Connector and S3 is for object storage setup only. The Service Connector can reside on your premises, instead of in the cloud.

Preparing your ONTAP clusters

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

Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

Supported ONTAP version

ONTAP 9.2 or later

Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 to AWS S3.

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

Although AWS Direct Connect provides better performance and lower data transfer charges, it is not required between the ONTAP cluster and AWS S3. Because performance is significantly better when using AWS Direct Connect, doing so is the recommended best practice.

- An inbound connection is required from the NetApp Service Connector, which can reside in an AWS VPC 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 tiered volumes. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

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.

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. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to 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.

Choosing a location for the Service Connector

The Service Connector is NetApp software that communicates with your ONTAP clusters. You can deploy the Service Connector on your premises or in an AWS VPC.

Be sure to set up the Service Connector in the same AWS account to which you want to tier data.

- [Installing the Service Connector on prem](#)
- [Preparing to deploy the Service Connector in an AWS VPC](#)

Preparing to deploy the Service Connector in an AWS VPC

Cloud Tiering guides you through the process of deploying the Service Connector on an EC2 instance. Make sure that your AWS account and networking are set up.

Setting up an AWS account for the Service Connector

The AWS account where you deploy the EC2 instance must have permissions and an access key. Cloud Tiering tiers data to an S3 bucket that resides in the same AWS account as the Service Connector.

Steps

1. Provide [the permissions in this policy](#) to the IAM user.

[AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#)

2. Create or locate an access key that you can provide to Cloud Tiering.

These credentials are used by the Cloud Tiering service to launch the EC2 instance in AWS. Providing your keys is secure and private. NetApp does not save them.

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

Setting up AWS networking for the Service Connector

The Service Connector needs a connection to your ONTAP clusters, to AWS S3, and to the Cloud Tiering service.

Steps

1. Identify a VPC for the Service Connector that enables the following connections:
 - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to S3
 - An HTTPS connection over port 443 to your ONTAP clusters

Cloud Tiering enables you to deploy the EC2 instance with a public IP address and you can configure it to use your own proxy server.

You don't need to create your own security group because Cloud Tiering can do that for you. The security group that Cloud Tiering creates has no inbound connectivity and open outbound connectivity.

2. If needed, enable a VPC Endpoint to S3.

A VPC Endpoint to S3 is recommended if you have a Direct Connect or VPN connection from your ONTAP cluster to the VPC and you want communication between the Service Connector and S3 to stay in your AWS internal network.

Preparing AWS S3 for data tiering

When you set up data tiering to a new cluster, Cloud Tiering prompts you to create an S3 bucket or select an existing S3 bucket in the AWS account where you set up the Service Connector.

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.

Steps

1. Provide the following permissions to the IAM user:

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

[AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#)

2. Create or locate an 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](#)

Installing the Service Connector on your premises for tiering to AWS S3

If you want to tier cold data to AWS S3, you can deploy the Service Connector either on your premises or in an AWS VPC. This page describes how to install the on-premises Service Connector.

To deploy the Service Connector in AWS, [follow the prompts in Cloud Tiering when discovering your first cluster](#).

Understanding the relationship between the Service Connector and Cloud Manager

To install the Service Connector, you need to download and install [NetApp Cloud Manager software](#). You need to do this because the Service Connector is part of Cloud Manager.

Verifying host requirements

[Refer to Connector host requirements in the Cloud Manager documentation](#).

Preparing your networking

The Service Connector needs a connection to your ONTAP clusters, to AWS S3, and to the Cloud Tiering service.

Steps

1. Set up an on-premises location for the Service Connector that enables the following connections:
 - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to S3
 - An HTTPS connection over port 443 to your ONTAP clusters
2. Ensure that outbound internet access is allowed to those endpoints:

- <http://dev.mysql.com/get/mysql-community-release-el7-5.noarch.rpm>
- <https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm>
- <https://s3.amazonaws.com/aws-cli/awscli-bundle.zip>

The installer accesses these URLs during the installation process.

Providing permissions to an AWS account

After you install the Service Connector, you need to provide access keys for an AWS account. That account needs specific permissions so the Service Connector can set up data tiering to AWS S3 on your behalf.

Cloud Tiering tiers data to an S3 bucket that resides in this AWS account.

Steps

1. From the AWS IAM console, create an IAM policy by copying and pasting the permissions below.

For step-by-step instructions, refer to [AWS Documentation: Creating IAM Policies](#).

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "sts:DecodeAuthorizationMessage",
        "s3:ListBucket",
        "ec2:DescribeRegions",
        "cloudformation:CreateStack",
        "cloudformation>DeleteStack",
        "cloudformation:DescribeStacks",
        "cloudformation:DescribeStackEvents",
        "cloudformation:ValidateTemplate"
      ],
      "Resource": "*"
    },
    {
      "Sid": "fabricPoolPolicy",
      "Effect": "Allow",
      "Action": [
        "s3:DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:CreateBucket",
        "s3:GetBucketTagging",
        "s3:PutBucketTagging"
      ],
      "Resource": "arn:aws:s3:::fabric-pool*"
    }
  ]
}

```

2. Attach the policy to an IAM role or an IAM user.

For step-by-step instructions, refer to the following:

- [AWS Documentation: Creating IAM Roles](#)
- [AWS Documentation: Adding and Removing IAM Policies](#)

Result

The account now has the required permissions. You need to provide access keys for the AWS account after you install the Service Connector.

Installing the Service Connector on an on-premises Linux host

After you verify system and network requirements, download and install the software on a supported Linux host.

About this task

- Root privileges are not required for installation.
- The Service Connector installs the AWS command line tools (awscli) to enable recovery procedures from NetApp support.

If you receive a message that installing the awscli failed, you can safely ignore the message. The Service Connector can operate successfully without the tools.

- The installer that is available on the NetApp Support Site might be an earlier version. After installation, the software automatically updates itself if a new version is available.

Steps

1. Download the installation script for Cloud Manager 3.8.4 or later from the [NetApp Support Site](#), and then copy it to the Linux host.

[Why do I need to install Cloud Manager?](#)

2. Assign permissions to execute the script.

Example

```
chmod +x OnCommandCloudManager-V3.8.4.sh
```

3. Run the installation script:

```
./OnCommandCloudManager-V3.8.4.sh [silent] [proxy=ipaddress]  
[proxyport=port] [proxyuser=user_name] [proxypwd=password]
```

silent runs the installation without prompting you for information.

proxy is required if the host is behind a proxy server.

proxyport is the port for the proxy server.

proxyuser is the user name for the proxy server, if basic authentication is required.

proxypwd is the password for the user name that you specified.

4. Unless you specified the silent parameter, type **Y** to continue the script, and then enter the HTTP and HTTPS ports when prompted.

If you change the HTTP and HTTPS ports, you must ensure that users can access the Cloud Manager web console from a remote host:

- Modify the security group to allow inbound connections through the ports.
- Specify the port when you enter the URL to the web console.

The Service Connector is now installed. At the end of the installation, the Cloud Manager service (occm) restarts twice if you specified a proxy server.

5. Open a web browser and enter the following URL:

`https://ipaddress:port`

ipaddress can be localhost, a private IP address, or a public IP address, depending on the configuration of the host.

port is required if you changed the default HTTP (80) or HTTPS (443) ports. For example, if the HTTPS port was changed to 8443, you would enter `https://ipaddress:8443`

6. Sign up at NetApp Cloud Central or log in.
7. After you log in, set up Cloud Manager:
 - a. Specify the Cloud Central account to associate with this Cloud Manager system. This should be the same account that you specified when you ran the pre-installation script.

[Learn about Cloud Central accounts.](#)

- b. Enter a name for the system.

[A screenshot that shows the set up Cloud Manager screen that enables you to select a Cloud Central account and name the Cloud Manager system.]

After you finish

Add an AWS account to Cloud Manager that has the required permissions.

Adding the AWS account to Cloud Manager

After you provide an AWS account with the required permissions, you need to add AWS access keys to Cloud Manager. This enables the Service Connector to set up data tiering to AWS S3 on your behalf.

Cloud Tiering tiers data to an S3 bucket that resides in this AWS account.

Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.

[A screenshot that shows the Settings icon in the upper right of the Cloud Manager console.]

2. Click **Add Credentials** and select **AWS**.
3. Select **AWS keys**.
4. Confirm that the policy requirements have been met and then click **Create Account**.

Result

The Service Connector is now installed with the permissions that it needs to tier cold data from your ONTAP systems to AWS S3. You should now see the Service Connector when you [set up tiering to a new cluster](#).

Tiering inactive data from your first cluster to AWS S3

After you prepare your AWS environment, just log in to Cloud Tiering and start tiering inactive data from your first cluster.

What you'll need

- To discover the cluster, you'll need the following:
 - The cluster management IP address.
 - The user name and password of an ONTAP account that has administrator-level privileges.

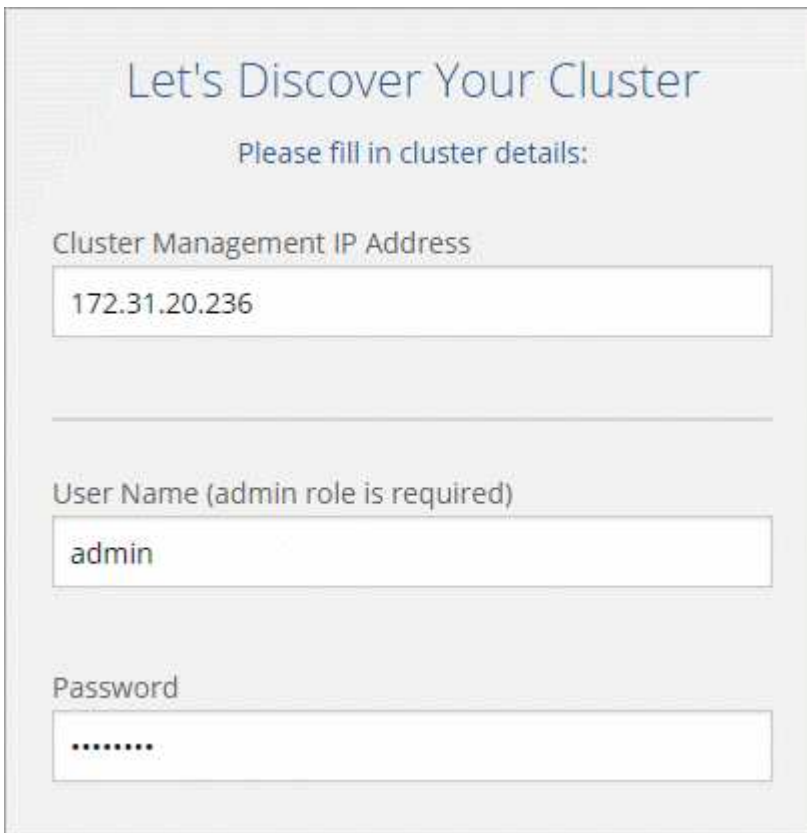
The Service Connector uses this account to send API operations to the ONTAP cluster.

- To deploy the Service Connector in AWS, you'll need the following:
 - The AWS region, VPC, and subnet in which the Service Connector will reside.
 - An AWS access key for an IAM user who has the required permissions.
- To set up tiering to S3, you'll need an AWS access key for an IAM user who has the required S3 permissions.

If you haven't met these requirements, see [Preparing your environment](#).

Steps

1. Log in to [NetApp Cloud Central](#).
2. Select the Cloud Tiering service.
3. Click **Let's Start, Discover Your First Cluster**.
4. Complete the steps on the **Discover Cluster** page:
 - a. Enter the cluster management IP address and the user name and password of an account that has administrator-level privileges.



Let's Discover Your Cluster

Please fill in cluster details:

Cluster Management IP Address

User Name (admin role is required)

Password

- b. Click **Discover Cluster**.

If you already have an existing Service Connector, then Cloud Tiering automatically attempts to use that Service Connector. Cloud Tiering moves on to the next step, if any existing Service Connector has connectivity to the cluster.

- c. If Cloud Tiering prompts you to create a Service Connector, click **Create your first Service Connector**.

If you have an existing Service Connector and Cloud Tiering can't use it, then you'll need to do one of two things:

- Click **Add Service Connector** to create a new Service Connector.
- Check the status and connectivity of any existing Service Connectors and then try to discover the cluster again.

- d. If you clicked the button to create a Service Connector, follow the prompts to deploy it in AWS:

- **Select Provider:** Select **Amazon Web Services** as the target location for the Service Connector.
- **AWS Credentials:** Enter the AWS access key ID and secret key for an IAM user that has [the required permissions](#) to deploy the Service Connector.
- **Location:** Select the AWS region, VPC, and subnet for the Service Connector EC2 instance.

Remember, the Service Connector must have a constant connection to the ONTAP cluster and a constant internet connection to the Cloud Tiering service.

- **Network:** Select a key pair to use for the EC2 instance, choose whether to assign a public IP, and specify an HTTP proxy, if one is required for outbound connectivity.
- **Security Group:** Select **Create a new security group** so Cloud Tiering can create the security group, or select your own. Then click **Go**.

The security group that Cloud Tiering creates has no inbound connectivity and open outbound connectivity.

Leave the page open until the deployment is complete.

- e. Back on the Discover Cluster page, select the Service Connector that you just created.

5. Complete the steps on the **Tiering Setup** page:

- a. **S3 Bucket:** Add a new S3 bucket or select an existing S3 bucket that starts with the prefix *fabric-pool* and click **Continue**.

The *fabric-pool* prefix is required because the IAM policy for the Service Connector 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.

- b. **Storage Class:** Select the S3 storage class that you want to transition the data to after 30 days and click **Continue**.

If you choose Standard, then the data remains in that storage class.


- c. **Credentials:** Enter the access key ID and secret key for an IAM user who has [the required S3 permissions](#).

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

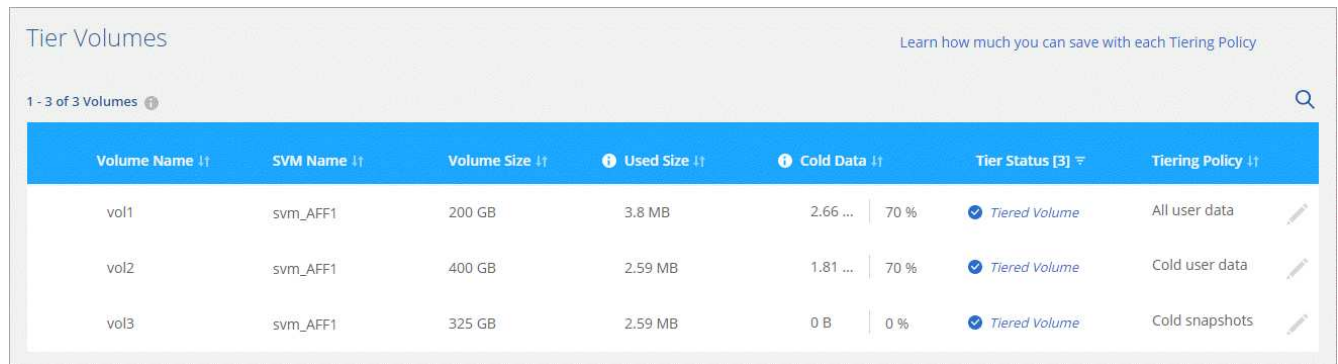
- d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

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

If you haven't reviewed requirements for the IPspace and the associated intercluster LIFs, see [ONTAP cluster requirements](#).

6. Click **Continue** to select the volumes that you want to tier.
7. For each volume, click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

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



The screenshot shows the 'Tier Volumes' interface. At the top, it says 'Tier Volumes' and 'Learn how much you can save with each Tiering Policy'. Below that, it indicates '1 - 3 of 3 Volumes'. The main part of the interface is a table with the following columns: Volume Name, SVM Name, Volume Size, Used Size, Cold Data, Tier Status [3], and Tiering Policy. The table contains three rows of data:

Volume Name	SVM Name	Volume Size	Used Size	Cold Data	Tier Status [3]	Tiering Policy
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ...	70 %	Tiered Volume All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ...	70 %	Tiered Volume Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B	0 %	Tiered Volume Cold snapshots

8. When you're done, click **Close**.

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 from the AWS Marketplace.](#)

You can also add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters](#).

Supported S3 storage classes and regions

Cloud Tiering supports several S3 storage classes and most regions.

Supported S3 storage classes

Cloud Tiering can apply a lifecycle rule so the data transitions from the *Standard* storage class to another storage class after 30 days. You can choose from the following storage classes:

- Standard-Infrequent Access
- One Zone-IA
- Intelligent

If you choose Standard, then the data remains in that storage class.

[Learn about S3 storage classes.](#)

Supported AWS regions

Cloud Tiering supports the following AWS regions.

Asia Pacific

- Mumbai
- Seoul
- Singapore
- Sydney
- Tokyo

Europe

- Frankfurt
- Ireland
- London
- Paris
- Stockholm

North America

- Canada Central
- GovCloud (US-West) – starting with ONTAP 9.3
- US East (N. Virginia)
- US East (Ohio)
- US West (N. California)
- US West (Oregon)

South America

- São Paulo

Tier data to Azure Blob storage

Quick start for tiering inactive data to Azure

Getting started with Cloud Tiering in Microsoft Azure includes a few steps.



Prepare to tier data to Azure Blob storage

You need the following:

- An AFF or FAS system with all-SSD aggregates running ONTAP 9.4 or later, and an HTTPS connection to Azure Blob storage.
- An Azure account that has the [required permissions](#) to deploy the Service Connector in a VNet.

The Service Connector needs an outbound HTTPS connection to the ONTAP cluster in your data center, to Azure Blob storage, and to the Cloud Tiering service.



Tier inactive data from your first cluster

Log in to [NetApp Cloud Central](#), start a free trial of the Cloud Tiering service, and click **Let's Start, Discover Your First Cluster**.



Set up licensing


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

- To subscribe from the Azure Marketplace, click **Licensing**, click **Subscribe**, and then follow the prompts.
- To add a tiering license, [contact us if you need to purchase one](#), and then add it to your cluster from ONTAP System Manager.

Preparing to tier inactive data to Azure Blob storage

Before you use Cloud Tiering, verify support for your ONTAP cluster, provide the required permissions, and set up your networking.

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



Communication between the Service Connector and Blob storage is for object storage setup only.

Preparing your ONTAP clusters

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

Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

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 is not required between the ONTAP cluster and Azure Blob storage. Because performance is significantly better when using ExpressRoute, doing so is the recommended best practice.

- An inbound connection is required from the NetApp Service Connector, which resides in an Azure VNet.

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

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

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

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.

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. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to 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.

Preparing to deploy the Service Connector in Azure

The Service Connector is NetApp software that communicates with your ONTAP clusters. Cloud Tiering guides you through the process of deploying the Service Connector on an Azure virtual machine.

A few steps are required before you can deploy the Service Connector in Azure. You'll need to provide the required permissions and set up your networking.

It's important to note that Cloud Tiering tiers data to a Blob container that resides in the same Azure subscription as the Service Connector. So be sure to complete these steps in the Azure subscription where both the Service Connector and Blob container should reside.

Steps

1. [Grant Azure permissions.](#)
2. [Set up networking.](#)

Granting Azure permissions

Ensure that your Azure account has the required permissions to deploy the NetApp Service Connector in an Azure VNet.



During deployment, Cloud Tiering creates and assigns a role to the Service Connector that provides the required permissions so ONTAP can tier inactive data to Azure Blob storage.

Steps

1. Create a custom role using the NetApp Cloud Central policy:
 - a. Download the [Cloud Central policy for Azure.](#)



Right-click the link and click **Save link as...** to download the file.

- b. Modify the JSON file by adding your Azure subscription ID to the assignable scope.

Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
]
```

- c. Use the JSON file to create a custom role in Azure.

The following example shows how to create a custom role using the Azure CLI 2.0:

```
az role definition create --role-definition C:\Policy_for_Setup_As_Service_Azure.json
```

You should now have a custom role called *Azure SetupAsService*.

2. Assign the role to the user who will deploy the Service Connector from Cloud Tiering:

- a. Open the **Subscriptions** service and select the user's subscription.
- b. Click **Access control (IAM)**.
- c. Click **Add > Add role assignment** and then add the permissions:
 - Select the **Azure SetupAsService** role.



Azure SetupAsService is the default name provided in the [Cloud Central policy](#). If you chose a different name for the role, then select that name instead.

- Assign access to an **Azure AD user, group, or application**.
- Select the user account.
- Click **Save**.

Result

The Azure user now has the permissions required to deploy the Service Connector.

Setting up Azure networking

Cloud Tiering prompts you for the Azure VNet where the Service Connector should be deployed. Make sure that the VNet provides the required networking connections.

Steps

1. Identify a VNet for the Service Connector that enables the following connections:
 - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to Azure Blob storage
 - An HTTPS connection over port 443 to your ONTAP clusters

Cloud Tiering enables you to deploy the virtual machine with a public IP address and you can configure it to use your own proxy server.

You don't need to create your own network security group because Cloud Tiering can do that for you. The security group that Cloud Tiering creates has no inbound connectivity and open outbound connectivity.

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 Service Connector and Blob storage to stay in your virtual private network.

Tiering inactive data from your first cluster to Azure Blob storage

After you prepare your Azure environment, just log in to Cloud Tiering and start tiering inactive data from your first cluster.

What you'll need

- To discover the cluster, you'll need the following:
 - The cluster management IP address.
 - The user name and password of an ONTAP account that has administrator-level privileges.

The Service Connector uses this account to send API operations to the ONTAP cluster.

- To deploy the Service Connector in Azure, you'll need the following:
 - An Azure account that has the required permissions to deploy the Service Connector virtual machine.
 - The Azure subscription, region, VNet, and subnet in which the Service Connector will reside.

If you haven't met these requirements, see [Preparing to tier data to Azure](#).

Steps

1. Log in to [NetApp Cloud Central](#).
2. Select the Cloud Tiering service.
3. Click **Let's Start, Discover Your First Cluster**.
4. Complete the steps on the **Discover Cluster** page:
 - a. Enter the cluster management IP address and the user name and password of an account that has administrator-level privileges.

Let's Discover Your Cluster

Please fill in cluster details:

Cluster Management IP Address

172.31.20.236

User Name (admin role is required)

admin

Password

.....

- b. Click **Discover Cluster**.

If you already have an existing Service Connector, then Cloud Tiering automatically attempts to use that Service Connector. Cloud Tiering moves on to the next step, if any existing Service Connector has connectivity to the cluster.

- c. If Cloud Tiering prompts you to create a Service Connector, click **Create your first Service Connector**.

If you have an existing Service Connector and Cloud Tiering can't use it, then you'll need to do one of two things:

- Click **Add Service Connector** to create a new Service Connector.
- Check the status and connectivity of any existing Service Connectors and then try to discover the cluster again.

- d. If you clicked the button to create a Service Connector, follow the prompts to deploy it in Azure:

- **Select Provider:** Select **Microsoft Azure** as the target location for the Service Connector.

When prompted, sign in and accept the permissions request from Microsoft.

The form is owned and hosted by Microsoft. Your credentials are not provided to NetApp.

- **Virtual Machine Authentication:** Enter a name for the virtual machine and choose an authentication method.
- **Basic Settings:** Select an Azure subscription, choose a region, and specify a new or existing resource group for the virtual machine.
- **Network:** Select a VNet and subnet, choose whether to assign a public IP address, and specify an

HTTP proxy, if one is required for outbound connectivity.

Remember, the Service Connector must have a constant connection to the ONTAP cluster and a constant internet connection to the Cloud Tiering service.

- **Security Group:** Select **Create a new security group** so Cloud Tiering can create the security group, or select your own. Then click **Go**.

The security group that Cloud Tiering creates has no inbound connectivity and open outbound connectivity.

Leave the page open until the deployment is complete.

- e. Back on the Discover Cluster page, select the Service Connector that you just created.

5. Complete the steps on the **Tiering Setup** page:

- a. **Resource Group:** Select a resource group where an existing container is managed, or where you would like to create a new container for tiered data.
- b. **Azure Container:** Add a new Blob container to a storage account or select an existing container and click **Continue**.


The storage account and containers that appear in this step belong to the resource group that you selected in the previous step.

- c. **Access Tier:** Select the access tier that you want to use for the tiered data and click **Continue**.
- d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

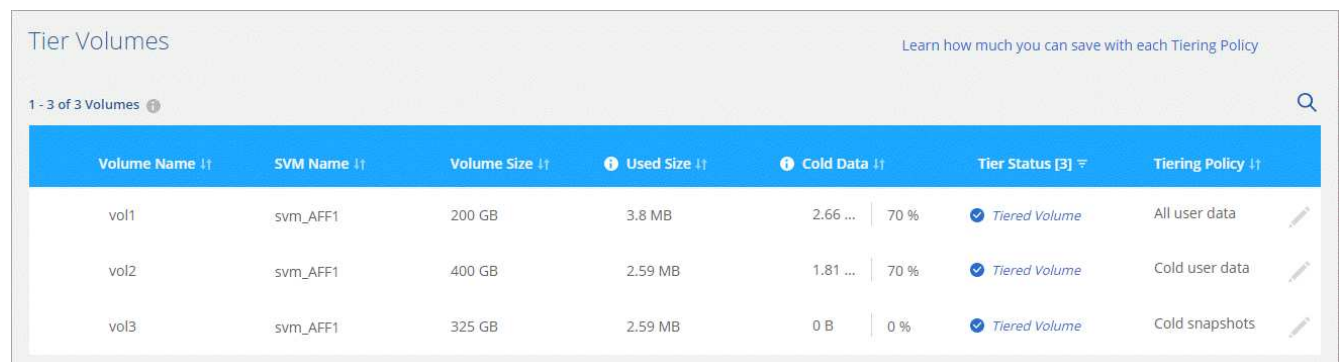
Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.

If you haven't reviewed requirements for the IPspace and the associated intercluster LIFs, see [ONTAP cluster requirements](#).

6. Click **Continue** to select the volumes that you want to tier.

7. For each volume, click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

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



Volume Name	SVM Name	Volume Size	Used Size	Cold Data	Tier Status [3]	Tiering Policy
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ... 70 %	✓ Tiered Volume	All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ... 70 %	✓ Tiered Volume	Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B 0 %	✓ Tiered Volume	Cold snapshots

8. When you're done, click **Close**.

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 from the Azure Marketplace.](#)

You can also add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters](#).

Supported Azure Blob access tiers and regions

Cloud Tiering supports the *Hot* access tier and most regions.

Supported Azure Blob access tiers

When you set up data tiering to Azure, Cloud Tiering automatically uses the *Hot* access tier for your inactive data.

Supported Azure regions

Cloud Tiering supports the following Azure regions.

Africa

- South Africa North

Asia Pacific

- Australia East
- Australia Southeast
- East Asia
- Japan East
- Japan West
- Korea Central
- Korea South
- Southeast Asia

Europe

- France Central
- Germany Central
- Germany Northeast
- North Europe
- UK South
- UK West
- West Europe

North America

- Canada Central
- Canada East
- Central US
- East US
- East US 2
- North Central US
- South Central US
- West US
- West US 2
- West Central US

South America

- Brazil South

Tier data to Google Cloud Storage

Quick start for tiering inactive data to Google Cloud Storage

Start tiering inactive data to Google Cloud Storage by completing a few steps.



Prepare to tier data to Google Cloud Storage

You need the following:

- An AFF or FAS system with all-SSD aggregates that's running ONTAP 9.6 or later and has an HTTPS connection to Google Cloud Storage.
- A service account that has the predefined Storage Admin role and storage access keys.
- A service account that includes the permissions defined in the [Service Connector policy for GCP](#).
- Google Cloud APIs enabled in your project: Cloud Deployment Manager V2 API, Cloud Resource Manager API, and Compute Engine API.
- A GCP user that has the [required permissions](#) to deploy the Service Connector in a Google Cloud Platform VPC.

The VPC where you deploy the Service Connector must provide an outbound HTTPS connection to the ONTAP cluster in your data center, to Google Cloud Storage, and to the Cloud Tiering service.

For step-by-step instructions, see [Preparing to tier inactive data to Google Cloud Storage](#).



Tier inactive data from your first cluster

Log in to [NetApp Cloud Central](#), start a free trial of the Cloud Tiering service, and click **Let's Start, Discover**

Your First Cluster.

For step-by-step instructions, see [Tiering inactive data from your first cluster to Google Cloud Storage](#).

3 Set up licensing

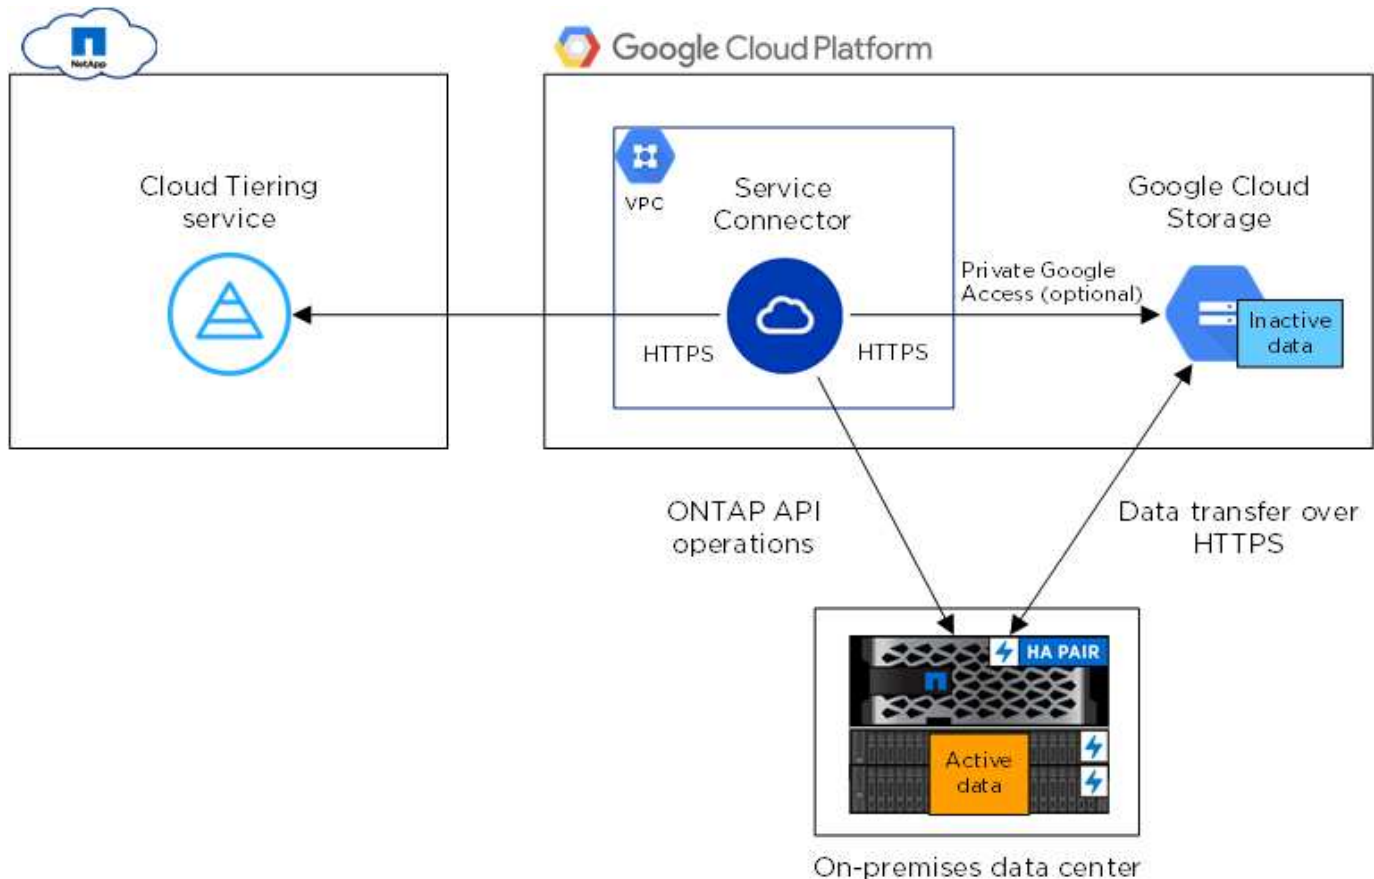
After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license, or a combination of both:

- To subscribe from the GCP Marketplace, click **Licensing**, click **Subscribe**, and then follow the prompts.
- To add a tiering license, [contact us if you need to purchase one](#), and then add it to your cluster from ONTAP System Manager.

Preparing to tier inactive data to Google Cloud Storage

Before you start tiering data, verify support for your ONTAP cluster, provide the required permissions, and set up your networking.

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



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

Preparing your ONTAP clusters

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

Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

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 is not required between the ONTAP cluster and Google Cloud Storage. Because performance is significantly better when using Google Cloud Interconnect, doing so is the recommended best practice.

- An inbound connection is required from the NetApp Service Connector, which resides in an 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 tiered volumes. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

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.

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. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to ONTAP documentation for [functionality or features not supported by FabricPool.](#)



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

Preparing to deploy the Service Connector in GCP

The Service Connector is NetApp software that communicates with your ONTAP clusters. Cloud Tiering guides you through the process of deploying the Service Connector on a GCP virtual machine instance.

A few steps are required before you can deploy the Service Connector in GCP. You'll need to provide the required permissions, set up a service account, and set up your networking.

It's important to note that Cloud Tiering tiers data to a Google Cloud bucket that resides in the same project as the Service Connector. So be sure to complete these steps in the project where both the Service Connector and bucket should reside.

Steps

1. [Set up GCP permissions.](#)
2. [Set up a service account.](#)
3. [Set up networking.](#)

Setting up GCP permissions

Ensure that your GCP user has the required permissions to deploy the NetApp Service Connector in a Google Cloud Platform VPC. You also need to enable a few APIs in the project.

Steps

1. Ensure that the GCP user who deploys the Service Connector has the permissions in the [Cloud Central policy for GCP](#).

[You can create a custom role using the YAML file](#) and then attach it to the user. You'll need to use the gcloud command line to create the role.

2. [Enable the following Google Cloud APIs in your project:](#)
 - Cloud Deployment Manager V2 API
 - Cloud Resource Manager API
 - Compute Engine API

Result

The GCP user now has the permissions required to deploy the Service Connector in GCP from Cloud Tiering.

Setting up a service account

When you deploy the Service Connector from Cloud Tiering, you need to select a service account to associate with the VM instance. This service account needs specific permissions to enable management of tiering.

Steps

1. [Create a role in GCP](#) that includes the permissions defined in the [Service Connector policy for GCP](#).

You'll need to use the gcloud command line to create the role.

2. [Create a GCP service account and apply the custom role that you just created.](#)

Setting up GCP networking

Cloud Tiering prompts you for the VPC where the Service Connector should be deployed. Make sure that the VPC provides the required networking connections.

Steps

1. Identify a VPC for the Service Connector that enables the following connections:
 - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to Google Cloud Storage
 - An HTTPS connection over port 443 to your ONTAP clusters

Cloud Tiering enables you to deploy the virtual machine with a public IP address and you can configure it to use your own proxy server.

You don't need to create your own firewall rules for the instance because Cloud Tiering can do that for you. The firewall rules that Cloud Tiering creates allows inbound connectivity over HTTP, HTTPS, and SSH. Outbound connectivity is open.

2. Optional: Enable Private Google Access on the subnet where you plan to deploy the Service Connector.

[Private Google Access](#) is recommended if you have a direct connection from your ONTAP cluster to the VPC and you want communication between the Service Connector 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).

Preparing Google Cloud Storage for data tiering

When you set up tiering, you need to provide Cloud Tiering with 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.

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 click **Interoperability**. If you haven't already done so, click **Enable interoperability access**.
 - b. Under **Access keys for service accounts**, click **Create a key for a service account**, select the service account that you just created, and click **Create Key**.

You'll need to [enter the keys in Cloud Tiering](#) later when you set up tiering.

Tiering inactive data from your first cluster to Google Cloud Storage

After you prepare your Google Cloud environment, just log in to Cloud Tiering and start tiering inactive data from your first cluster.

What you'll need

- To discover the cluster, you'll need the following:
 - The cluster management IP address.
 - The user name and password of an ONTAP account that has administrator-level privileges.

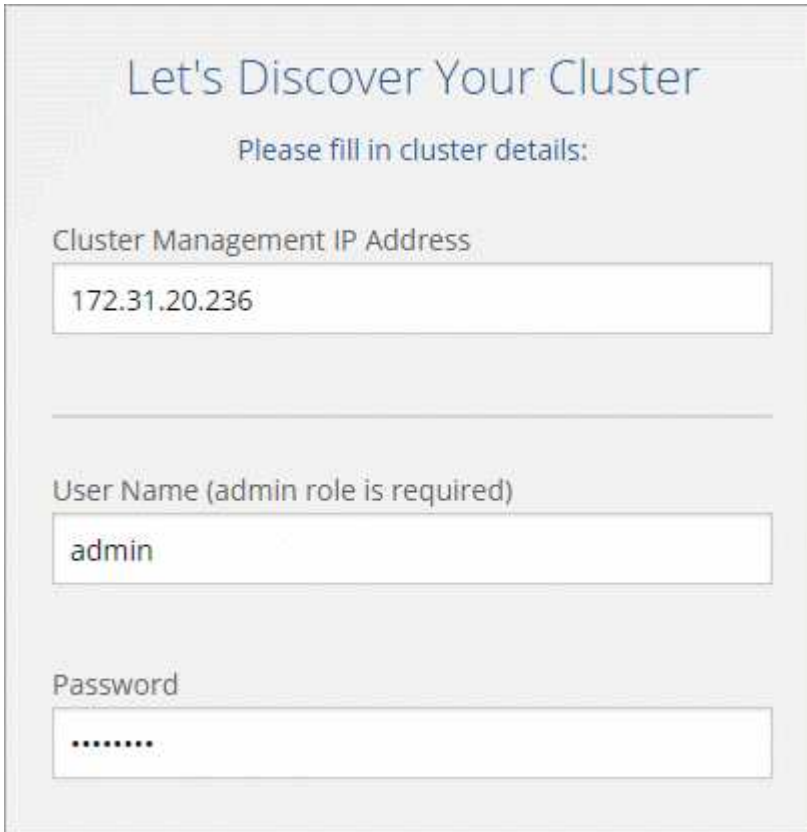
The Service Connector uses this account to send API operations to the ONTAP cluster.

- To deploy the Service Connector in GCP, you'll need the following:
 - A Google account that has the required permissions to deploy the Service Connector virtual machine.
 - The project, region, VPC, and subnet in which the Service Connector will reside.
 - A service account that includes the permissions defined in the [Service Connector policy for GCP](#).
- To set up tiering to Google Cloud Storage, you'll need storage access keys for a service account that has the Storage Admin role.

If you haven't met these requirements, see [Preparing to tier data to Google Cloud Storage](#).

Steps

1. Log in to [NetApp Cloud Central](#).
2. Select the Cloud Tiering service.
3. Click **Let's Start, Discover Your First Cluster**.
4. Complete the steps on the **Discover Cluster** page:
 - a. Enter the cluster management IP address and the user name and password of an account that has administrator-level privileges.



The screenshot shows a web form titled "Let's Discover Your Cluster" with the instruction "Please fill in cluster details:". It contains three input fields: "Cluster Management IP Address" with the value "172.31.20.236", "User Name (admin role is required)" with the value "admin", and "Password" with masked characters "*****".

- b. Click **Discover Cluster**.

If you already have an existing Service Connector, then Cloud Tiering automatically attempts to use that Service Connector. Cloud Tiering moves on to the next step, if any existing Service Connector has connectivity to the cluster.

- c. If Cloud Tiering prompts you to create a Service Connector, click **Create your first Service Connector**.

If you have an existing Service Connector and Cloud Tiering can't use it, then you'll need to do one of two things:

- Click **Add Service Connector** to create a new Service Connector.
 - Check the status and connectivity of any existing Service Connectors and then try to discover the cluster again.
- d. If you clicked the button to create a Service Connector, follow the prompts to deploy it in GCP:
 - **Select Provider:** Select **Google Cloud Platform** as the target location for the Service Connector.

When prompted, sign in and accept the permissions request from Google. The form is owned and hosted by Google. Your credentials are not provided to NetApp.

- **Basic Settings:** Enter a name for the virtual machine, select a project, and then select a service account that includes the permissions defined in the [Service Connector policy for GCP](#).
- **Location:** Specify networking for the virtual machine—a region, zone, VPC, and subnet, and then choose whether you want to assign a public IP address and specify an HTTP proxy for outbound connectivity.

Remember, the Service Connector must have a constant connection to the ONTAP cluster and a constant internet connection to the Cloud Tiering service.

- **Firewall Policy:** Select **Create a new firewall policy** so Cloud Tiering can create the security group, or select an existing firewall policy. Then click **Go**.

The firewall policy that Cloud Tiering creates allows inbound HTTP, HTTPS, and SSH connectivity. It has open outbound connectivity.

Leave the page open until the deployment is complete.

e. Back on the Discover Cluster page, select the Service Connector that you just created.


5. Complete the steps on the **Tiering Setup** page:

- Bucket:** Add a new Google Cloud Storage bucket or select an existing bucket and click **Continue**.
- Storage Class:** Select the storage class that you want to use for the tiered data and click **Continue**.
- Credentials:** Enter the storage access key and secret key for a service account that has the Storage Admin role.
- Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

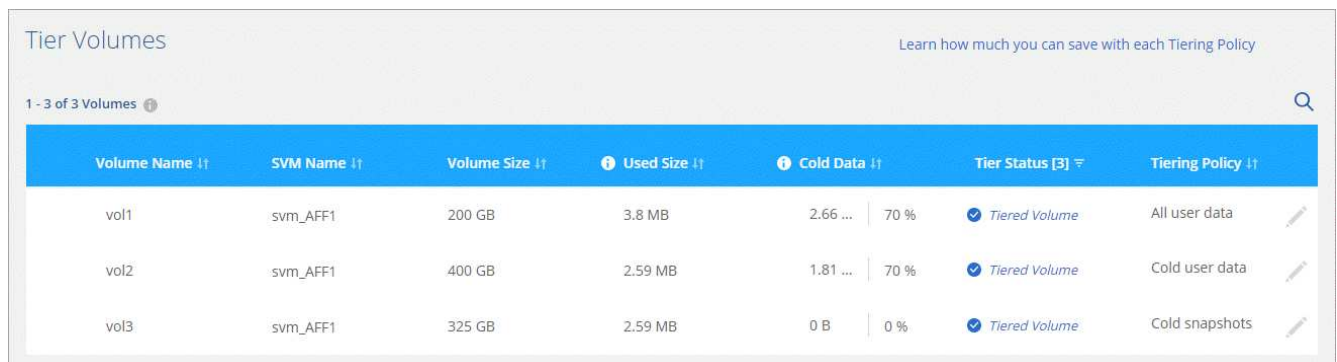
Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.




If you haven't reviewed requirements for the IPspace and the associated intercluster LIFs, see [ONTAP cluster requirements](#).

6. Click **Continue** to select the volumes that you want to tier.

7. For each volume, click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

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



Volume Name ↑↓	SVM Name ↑↓	Volume Size ↑↓	Used Size ↑↓	Cold Data ↑↓	Tier Status [3] ⇅	Tiering Policy ↑↓
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ... 70 %	✓ Tiered Volume	All user data 
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ... 70 %	✓ Tiered Volume	Cold user data 
vol3	svm_AFF1	325 GB	2.59 MB	0 B 0 %	✓ Tiered Volume	Cold snapshots 

8. When you're done, click **Close**.

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 from the GCP Marketplace](#).

You can also add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters](#).

Supported Google Cloud storage classes and regions

Cloud Tiering supports the Standard storage class and most Google Cloud regions.

Supported access tiers

Cloud Tiering uses the *Standard* access tier for your inactive data.

Supported Google Cloud regions

Cloud Tiering supports the following regions.

Americas

- Iowa
- Los Angeles
- Montreal
- N. Virginia
- Oregon
- Sao-Paulo
- South Carolina

Asia Pacific

- Hong Kong
- Mumbai
- Osaka
- Singapore
- Sydney
- Taiwan
- Tokyo

Europe

- Belgium
- Finland

- Frankfurt
- London
- Netherlands
- Zurich

Tier data to StorageGRID

Quick start for tiering inactive data to StorageGRID

Getting started with Cloud Tiering by tiering data from an ONTAP cluster to StorageGRID includes a few steps.



Prepare to tier data to StorageGRID

You need the following:

- An AFF or FAS system with all-SSD aggregates running ONTAP 9.4 or later, and a connection over a user-specified port to StorageGRID.
- StorageGRID 10.3 or later with AWS access keys that have S3 permissions.
- A Service Connector installed on an on-premises Linux host.

The Service Connector needs an outbound HTTPS connection to the ONTAP cluster, to StorageGRID, and to the Cloud Tiering service.



Tier inactive data from your first cluster


Log in to [NetApp Cloud Central](#), start a free trial of the Cloud Tiering service, and click **Let's Start, Discover Your First Cluster**.

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

Preparing to tier inactive data to StorageGRID

Before you use Cloud Tiering, verify support for your ONTAP cluster, prepare StorageGRID, and install a Service Connector on an on-premises Linux host.

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



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

Preparing your ONTAP clusters

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

Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

Supported ONTAP version

ONTAP 9.4 or later

Licensing

A FabricPool license isn't 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 StorageGRID (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 NetApp Service Connector, which resides 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 tiered volumes. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

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.

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. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to 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.

Preparing StorageGRID

StorageGRID must meet the following requirements.

Supported StorageGRID versions

StorageGRID 10.3 and later are 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",  
"s3:GetObject",  
"s3:PutObject",  
"s3:DeleteObject",  
"s3:CreateBucket"
```

Object versioning

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

Installing the Service Connector on-prem for StorageGRID

To tier data to StorageGRID, you need to install a Service Connector on an on-prem Linux host.

Understanding the relationship between the Service Connector and Cloud Manager

To install the Service Connector, you need to download and install [NetApp Cloud Manager software](#). You need to do this because the Service Connector is part of Cloud Manager.

Verifying host requirements

The Service Connector is supported on a Linux host that meets the following requirements.

[Refer to Connector host requirements in the Cloud Manager documentation.](#)

Preparing your networking

The Service Connector needs a connection to your ONTAP clusters, to StorageGRID, and to the Cloud Tiering service.

Steps

1. Set up an on-premises location for the Service Connector that enables the following connections:
 - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to StorageGRID
 - An HTTPS connection over port 443 to your ONTAP clusters
2. Ensure that outbound internet access is allowed to those endpoints:
 - <http://dev.mysql.com/get/mysql-community-release-el7-5.noarch.rpm>
 - <https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm>
 - <https://s3.amazonaws.com/aws-cli/awscli-bundle.zip>

The installer accesses these URLs during the installation process.

Installing the Service Connector on an on-premises Linux host

After you verify system and network requirements, download and install the software on a supported Linux host.

About this task

- Root privileges are not required for installation.
- The Service Connector installs the AWS command line tools (awscli) to enable recovery procedures from NetApp support.

If you receive a message that installing the awscli failed, you can safely ignore the message. The Service Connector can operate successfully without the tools.

- The installer that is available on the NetApp Support Site might be an earlier version. After installation, the software automatically updates itself if a new version is available.

Steps

1. Download the installation script for Cloud Manager 3.8.4 or later from the [NetApp Support Site](#), and then copy it to the Linux host.

[Why do I need to install Cloud Manager?](#)

2. Assign permissions to execute the script.

Example

```
chmod +x OnCommandCloudManager-V3.8.4.sh
```

3. Run the installation script:

```
./OnCommandCloudManager-V3.8.4.sh [silent] [proxy=ipaddress]
[proxyport=port] [proxyuser=user_name] [proxypwd=password]
```

silent runs the installation without prompting you for information.

proxy is required if the host is behind a proxy server.

proxyport is the port for the proxy server.

proxyuser is the user name for the proxy server, if basic authentication is required.

proxypwd is the password for the user name that you specified.

4. Unless you specified the silent parameter, type **Y** to continue the script, and then enter the HTTP and HTTPS ports when prompted.

If you change the HTTP and HTTPS ports, you must ensure that users can access the Cloud Manager web console from a remote host:

- Modify the security group to allow inbound connections through the ports.
- Specify the port when you enter the URL to the web console.

The Service Connector is now installed. At the end of the installation, the Cloud Manager service (occm) restarts twice if you specified a proxy server.

5. Open a web browser and enter the following URL:

`https://ipaddress:port`

ipaddress can be localhost, a private IP address, or a public IP address, depending on the configuration of the host.

port is required if you changed the default HTTP (80) or HTTPS (443) ports. For example, if the HTTPS port was changed to 8443, you would enter `https://ipaddress:8443`

6. Sign up at NetApp Cloud Central or log in.

7. After you log in, set up Cloud Manager:

- a. Specify the Cloud Central account to associate with this Cloud Manager system. This should be the same account that you specified when you ran the pre-installation script.

[Learn about Cloud Central accounts.](#)

- b. Enter a name for the system.

[A screenshot that shows the set up Cloud Manager screen that enables you to select a Cloud Central account and name the Cloud Manager system.]

Result

The Service Connector is now installed and setup. You can use it to discover a cluster in Cloud Tiering.

Tiering inactive data from your first cluster to StorageGRID

After you prepare your environment, just log in to Cloud Tiering and start tiering inactive data from your first cluster.

What you'll need

- To discover the cluster, you'll need the following:
 - The cluster management IP address.
 - The user name and password of an ONTAP account that has administrator-level privileges.

The Service Connector uses this account to send API operations to the ONTAP cluster.

- [A Service Connector installed in your on-premises network.](#)
- To set up tiering to StorageGRID, you'll need the following:
 - The FQDN of the server.
 - An AWS access key and secret key for an account that has the required S3 permissions.

If you haven't met these requirements, see [Preparing to tier data to StorageGRID](#).

Steps

1. Log in to [NetApp Cloud Central](#).
2. Select the Cloud Tiering service.
3. Click **Let's Start, Discover Your First Cluster**.
4. Complete the steps on the **Discover Cluster** page:
 - a. Enter the cluster management IP address and the user name and password of an account that has administrator-level privileges.

Let's Discover Your Cluster

Please fill in cluster details:

Cluster Management IP Address

User Name (admin role is required)

Password

- b. Click **Discover Cluster**.

Cloud Tiering automatically uses the on-premises Service Connector to discover the cluster.


5. Complete the steps on the **Tiering Setup** page:

- a. **Choose your provider:** Select StorageGRID.
- b. **Server:** Enter the FQDN of the StorageGRID server, enter the port that ONTAP should use for HTTPS communication with StorageGRID, and enter the access key and secret key for an AWS account that has the required S3 permissions.
- c. **Bucket:** Add a new bucket or select an existing bucket for the tiered data.
- d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

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

If you haven't reviewed requirements for the IPspace and the associated intercluster LIFs, see [ONTAP cluster requirements](#).

6. Click **Continue** to select the volumes that you want to tier.

7. For each volume, click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

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

Tier Volumes Learn how much you can save with each Tiering Policy

1 - 3 of 3 Volumes 🔍

Volume Name ↑	SVM Name ↑	Volume Size ↑	Used Size ↑	Cold Data ↑	Tier Status [3] ⇅	Tiering Policy ↑
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ... 70 %	✓ Tiered Volume	All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ... 70 %	✓ Tiered Volume	Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B 0 %	✓ Tiered Volume	Cold snapshots

8. When you're done, click **Close**.

Result

You've successfully set up data tiering from volumes on the cluster to StorageGRID.

What's next?

You can add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters](#).

Set up licensing for Cloud Tiering

Pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license, or a combination of both. If you want to pay as you go, then you need to subscribe from the marketplace for the cloud provider to which you want to tier cold data. There's no need to subscribe from every marketplace.

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



A license isn't required when tiering data to StorageGRID.

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.



If you use NetApp Cloud Manager and you've already [subscribed through its new AWS Marketplace offering](#), then you're automatically subscribed to Cloud Tiering, as well. You'll see an active subscription in Cloud Tiering in the **Licensing** tab. You won't need to subscribe.

Steps

1. In Cloud Tiering, click **Licensing**.
2. Click **Subscribe** under AWS Marketplace and then click **Continue**.
3. Subscribe from the AWS Marketplace, and then log back in to Cloud Central to complete the registration.

The following video shows the process:

▶ https://docs.netapp.com/us-en/cloud-tiering//media/video_subscribing_aws.mp4 (video)

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 Cloud Tiering, click **Licensing**.
2. Click **Subscribe** under Azure Marketplace and then click **Continue**.
3. Subscribe from the Azure Marketplace, and then log back in to Cloud Central to complete the registration.

The following video shows the process:

► https://docs.netapp.com/us-en/cloud-tiering//media/video_subscribing_azure.mp4 (video)

Subscribing from the GCP Marketplace

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

Steps

1. In Cloud Tiering, click **Licensing**.
2. Click **Subscribe** under GCP Marketplace and then click **Continue**.
3. Subscribe from the GCP Marketplace, and then log back in to Cloud Central to complete the registration.

The following video shows the process:

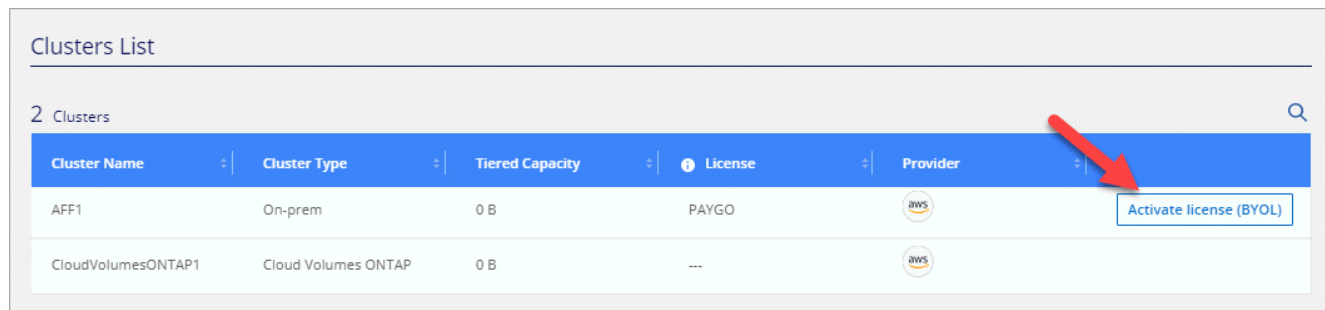
► https://docs.netapp.com/us-en/cloud-tiering//media/video_subscribing_gcp.mp4 (video)

Adding a tiering license to ONTAP

Bring your own license by purchasing an ONTAP FabricPool license from NetApp.

Steps

1. If you don't have a FabricPool license, [contact us to purchase one](#).
2. In Cloud Tiering, go to the **Licensing** page.
3. In the Clusters List table, click **Activate license (BYOL)** for an on-prem ONTAP cluster.



Cluster Name	Cluster Type	Tiered Capacity	License	Provider	
AFF1	On-prem	0 B	PAYGO	aws	Activate license (BYOL)
CloudVolumesONTAP1	Cloud Volumes ONTAP	0 B	---	aws	

4. Enter the serial number of the license and then enter the NetApp Support Site account that's associated with the serial number.

5. Click **Activate license**.

Result

Cloud Tiering registers the license and installs it on the cluster.

After you finish

If you purchase additional add-on capacity at a later time, the license on the cluster is automatically updated with the new capacity. There's no need to apply a new NetApp License File (NLF) to the cluster.

Managing data tiering from your clusters

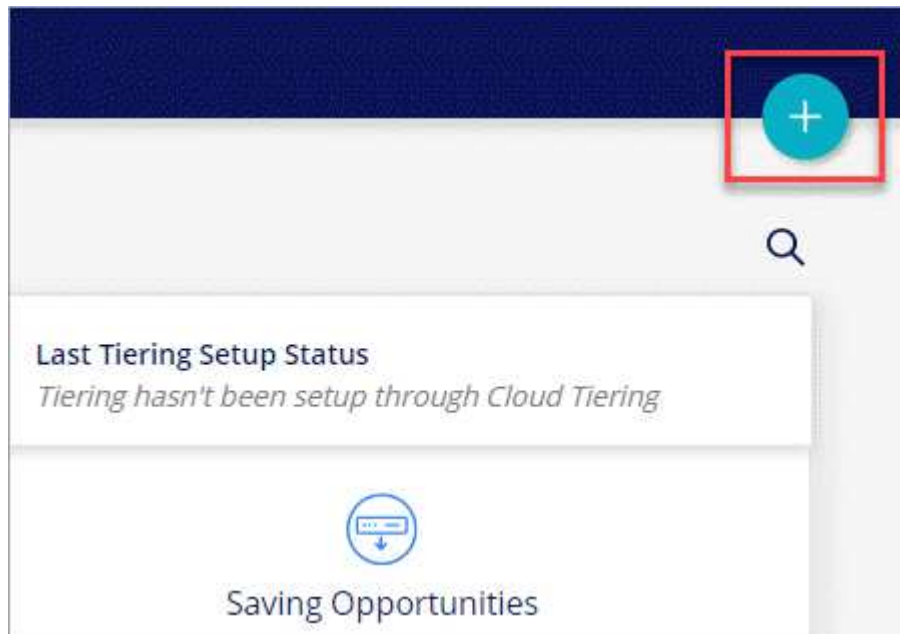
Now that you've discovered and tiered data from your first ONTAP cluster, you can tier data from additional volumes, discover additional clusters, and more.

Discovering additional clusters

Add additional ONTAP clusters at any time to start tiering inactive data from those clusters.

Steps

1. From the Cluster Dashboard, click the following icon:



2. Enter the cluster management IP address and the user name and password of an account that has administrator-level privileges.
3. Click **Discover Cluster**.

If you already have an existing Service Connector, then Cloud Tiering automatically attempts to use that Service Connector.

4. If you have an existing Service Connector and Cloud Tiering can't use it, then you'll need to do one of two things:
 - Click **Add Service Connector** to create a new Service Connector.
 - Check the status and connectivity of any existing Service Connectors and then try to discover the cluster again.

What's next?


Click **Tier Volumes** to start tiering inactive data from the cluster.



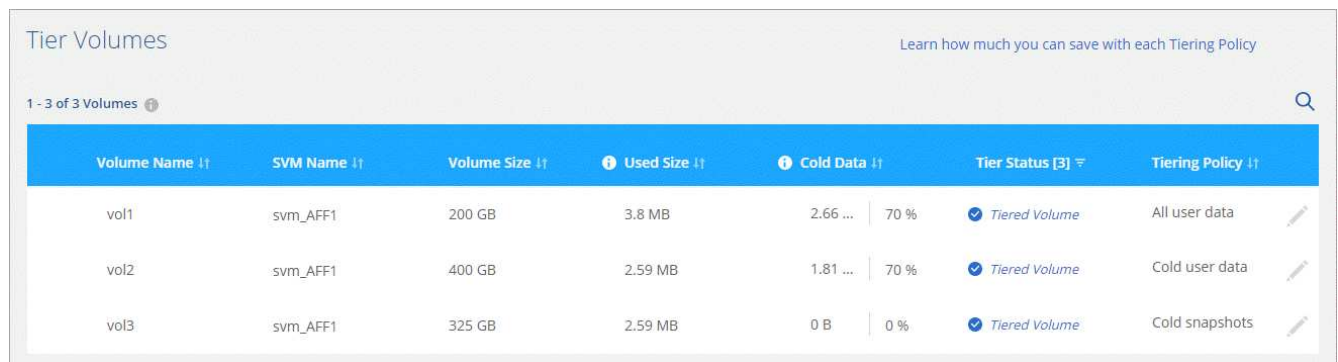
Tiering data from additional volumes




Set up data tiering for additional volumes at any time—for example, after creating a new volume.

Steps

1. From the **Cluster Dashboard**, click **Tier Volumes**.
2. For each volume, click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

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



Volume Name ↑	SVM Name ↑	Volume Size ↑	Used Size ↑	Cold Data ↑	Tier Status [3] ▾	Tiering Policy ↑
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ... 70 %	✓ Tiered Volume	All user data 
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ... 70 %	✓ Tiered Volume	Cold user data 
vol3	svm_AFF1	325 GB	2.59 MB	0 B 0 %	✓ Tiered Volume	Cold snapshots 



You don't need to configure the object storage because it was already configured when you initially set up tiering for the cluster. ONTAP will tier inactive data from these volumes to the same object store.


3. When you're done, click **Close**.

Changing a volume's tiering policy and cooling period

Changing the tiering policy for a volume changes how ONTAP tiers cold data to object storage. The change starts from the moment that you change the policy—it changes only the subsequent tiering behavior for the volume.

Steps

1. From the **Cluster Dashboard**, click **Tier Volumes** for the cluster.

2. Click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

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

Managing tiering settings on aggregates

Each aggregate has two settings that you can adjust: the tiering fullness threshold and whether inactive data reporting is enabled.

Tiering fullness threshold

Setting the threshold to a lower number reduces the amount of data required to be stored on the performance tier before tiering takes place. This might be useful for large aggregates that contain little active data.

Setting the threshold to a higher number increases the amount of data required to be stored on the performance tier before tiering takes place. This might be useful for solutions designed to tier only when aggregates are near maximum capacity.

Inactive data reporting

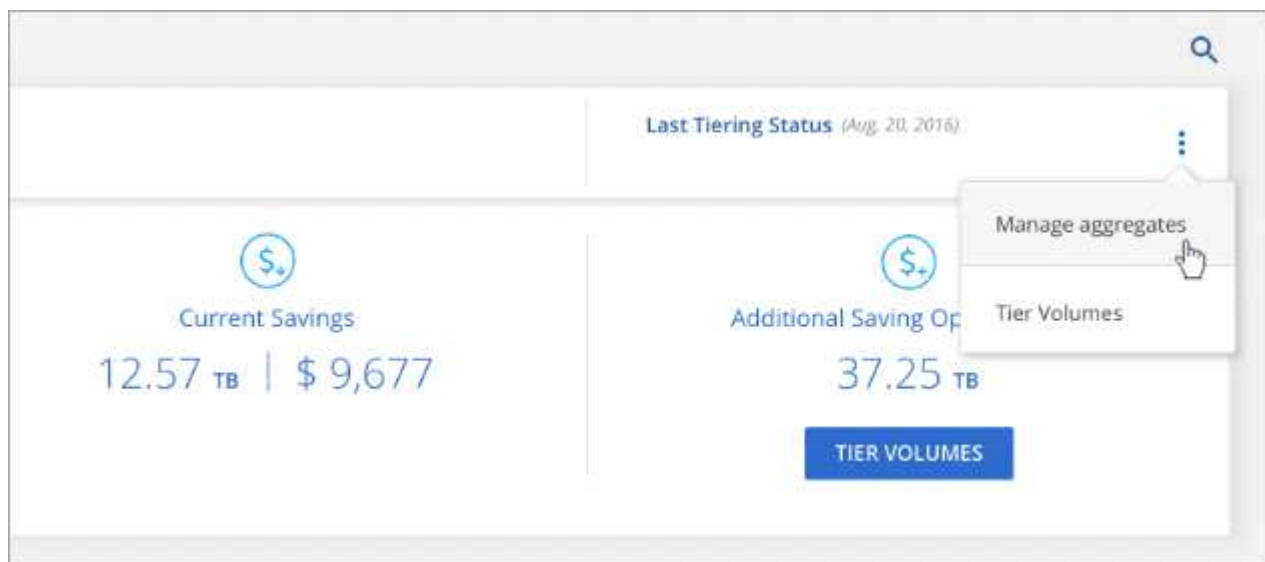
Inactive data reporting (IDR) uses a 31-day cooling period to determine which data is considered inactive. The amount of cold data that is tiered is dependent on the tiering policies set on volumes. This amount might be different than the amount of cold data detected by IDR using a 31-day cooling period.




It's best to keep IDR enabled because it helps to identify your inactive data and savings opportunities. IDR must remain enabled if data tiering was enabled on an aggregate.

Steps

1. From the **Cluster Dashboard**, click menu icon for a cluster and select **Manage Aggregates**.



2. On the **Manage Aggregates** page, click the  icon for an aggregate in the table.
3. Modify the fullness threshold and choose whether to enable or disable inactive data reporting.




4. Click **Apply**.

Reviewing tiering info for a cluster

You might want to see how much data is in the cloud tier and how much data is on disks. Or, you might want to see the amount of hot and cold data on the cluster's disks. Cloud Tiering provides this information for each cluster.


Steps

1. From the **Cluster Dashboard**, click **More info** for a cluster.



2. Review details about the cluster.

Here's an example:




Fixing operational health

Failures can happen. When they do, Cloud Tiering displays a "Failed" operational health status on the Cluster Dashboard. The health reflects the status of the ONTAP system and the Service Connector.

Steps

1. Identify any clusters that have an operational health of "Failed."



2. Hover over the **i** icon to see the failure reason.
3. Correct the issue:
 - a. Verify that the ONTAP cluster is operational and that it has an inbound and outbound connection to your object storage provider.
 - b. Verify that the Service Connector is running and that it has outbound connections to the Cloud Tiering service, to the object store, and to the ONTAP clusters that it discovers.




The name of the Service Connector instance/virtual machine is prefixed with "Service-connector."

Removing a failed cluster

If the health of a cluster is failed, you can remove it from the dashboard to focus on the operational clusters.

Steps

1. From the **Cluster Dashboard**, identify any clusters that have an operational health of "Failed."
2. Click menu icon for a cluster and select **Remove Cluster**.




How to get help and find more information

Support for Cloud Tiering is covered only through the resources listed below. We do not offer phone or web ticketing at this time through mysupport.netapp.com.

- Chat services

Chat with NetApp cloud experts to get product assistance. The chat icon is available in the lower right of the Cloud Tiering interface:



- [NetApp Community: Cloud Data Services](#)

In this forum, you can use labels and filters to look at Cloud Tiering topics. If you'd like to ask a question, click **Register** in the upper-right corner to sign up.

- [NetApp Cloud Central](#)

Find more information about Cloud Tiering, as well as additional NetApp products and solutions for the cloud.

- [NetApp Product Documentation](#)

Search NetApp product documentation for instructions, resources, and answers.

Cloud Tiering APIs

The Cloud Tiering capabilities that are available through the web UI are also available through GraphQL APIs.

Getting started

To get started, obtain a user token and then add the token to the Authorization header when making API calls.

Steps

1. Go to [Cloud Central's API Documentation page](#).
2. Click **Learn how to authenticate**.
3. Follow the instructions to acquire and use an access token.

API reference

Documentation for each API is available from [NetApp Cloud Central](#).

Cloud Tiering technical FAQ

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

ONTAP

The following questions relate to ONTAP.

What are the requirements for my ONTAP cluster?

It depends on where you tier the cold data. Refer to the following:

- [Preparing to tier inactive data to AWS S3](#)
- [Preparing to tier inactive data to Azure Blob storage](#)
- [Preparing to tier inactive data to Google Cloud Storage](#)
- [Preparing to tier inactive data to StorageGRID](#)

Does Cloud Tiering enable inactive data reporting?

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

Can I tier data from NAS volumes and SAN volumes?

You can use Cloud Tiering to tier data from NAS volumes to the public cloud and from SAN volumes to a private cloud using StorageGRID.

What about Cloud Volumes ONTAP?

If you have Cloud Volumes ONTAP systems, you'll find them in the Cluster Dashboard so you see a full view of data tiering in your hybrid cloud infrastructure.

From the Cluster Dashboard, you can view tiering information similar to an on-prem ONTAP cluster: operational health, current savings, savings opportunities, details about volumes and aggregates, and more.

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'll still set up tiering the same way: from the working environment in Cloud Manager.

Object storage

The following questions relate to object storage.

Which object storage providers are supported?

Amazon S3, Azure Blob storage, Google Cloud Storage, and StorageGRID using the S3 protocol are supported.

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 public cloud 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-IA*, or *Intelligent* storage class. See [Supported S3 storage classes](#) for more details.

Which Azure Blob access tiers are supported?

Cloud Tiering automatically uses the *Hot* access tier for your inactive data.

Which storage classes are supported for Google Cloud Storage?

Cloud Tiering uses the *Standard* storage class for inactive data.

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

One object store for the entire cluster.

Can I apply policies to my object store to move data around independent of tiering?

No, Cloud Tiering does not support object lifecycle management rules that move or delete data from object stores.

NetApp Service Connector

The following questions relate to the NetApp Service Connector.

What is the Service Connector?

The Service Connector is NetApp software that communicates with ONTAP clusters to discover information about active and inactive data, and to set up data tiering. For more details, see [How Cloud Tiering works](#).

Where can I run the Service Connector?

- When tiering cold data to S3, the Service Connector can reside in an AWS VPC or on your premises.
- When tiering cold data to Blob storage, the Service Connector must reside in an Azure VNet.
- When tiering cold data to Google Cloud Storage, the Service Connector must reside in a Google Cloud Platform VPC.

- When tiering cold data to StorageGRID, the Service Connector must reside on an on premises Linux host.

How do you name the instance/virtual machine for the Service Connector?

The name of the Service Connector is prefixed with "Service-connector."

Networking

The following questions relate to networking.

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).
- The NetApp Service Connector 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:

- [Preparing to tier inactive data to AWS S3](#)
- [Preparing to tier inactive data to Azure Blob storage](#)
- [Preparing to tier inactive data to Google Cloud Storage](#)
- [Preparing to tier inactive data to StorageGRID](#)

Permissions

The following questions relate to permissions.

What permissions are required in AWS?

Permissions are needed to install the Service Connector:

- [These permissions are required to deploy the Service Connector in an AWS VPC](#)
- [These permissions are required when you deploy the Service Connector on an on-premises Linux host](#)

A different set of permissions are required [to manage the S3 bucket](#).

What permissions are required in Azure?

Permissions are needed [to deploy the Service Connector in an Azure VNet](#).

During deployment, Cloud Tiering creates and assigns a role to the Service Connector that provides the required permissions so ONTAP can tier inactive data to Azure Blob storage.

What permissions are required in Google Cloud Platform?

- Permissions are needed for the GCP user who will deploy the Service Connector in GCP from Cloud Tiering.
- Permissions are needed for a service account that has storage access keys.
- Permissions are needed for a service account that you'll associate with the Service Connector VM instance.

For details, see [Preparing to tier inactive data to Google Cloud Storage](#).

What permissions are required for StorageGRID?

[S3 permissions are needed](#).

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<http://www.netapp.com/us/legal/copyright.aspx>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/us/media/patents-page.pdf>

Privacy policy

<https://www.netapp.com/us/legal/privacypolicy/index.aspx>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

[Notice for Cloud Tiering](#)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

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

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

Trademark information

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