



Use NetApp Cloud Tiering

NetApp Cloud Tiering

NetApp

November 10, 2025

Table of Contents

Use NetApp Cloud Tiering	1
Manage data tiering for your clusters in NetApp Cloud Tiering	1
Review tiering info for a cluster	1
Tier data from additional volumes	2
Change a volume's tiering policy	3
Change the network bandwidth available to upload inactive data to object storage	4
Download a tiering report for your volumes	5
Migrate data from the cloud tier back to the performance tier	5
Manage tiering settings on aggregates	6
Fix operational health	8
Discover additional clusters from Cloud Tiering	8
Search for a cluster across all Console agents	9
Manage object storage used for data tiering in NetApp Cloud Tiering	9
View object stores configured for a cluster	9
Add a new object store	10
Attach a second object store to an aggregate for mirroring	12
Swap the primary and mirror object store	13
Remove a mirror object store from an aggregate	13
Migrate your tiered data to a different cloud provider	14
Measure network latency and throughput performance in NetApp Cloud Tiering	14
Get an overview of data tiering from your clusters in NetApp Cloud Tiering	15
Monitor the status of tiering alerts from NetApp Cloud Tiering	17

Use NetApp Cloud Tiering

Manage data tiering for your clusters in NetApp Cloud Tiering

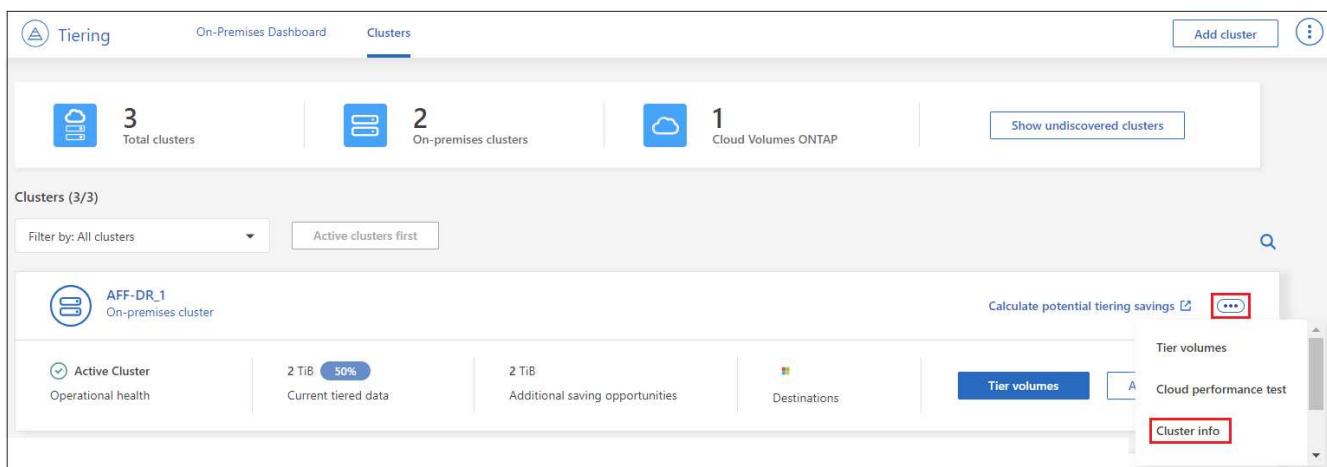
Now that you've set up data tiering from your on-premises ONTAP clusters, you can tier data from additional volumes, change a volume's tiering policy, discover additional clusters, and more by using NetApp Cloud Tiering.

Review tiering info for a cluster

Check the data in the cloud tier, on disks, or the amount of hot and cold data on the cluster's 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

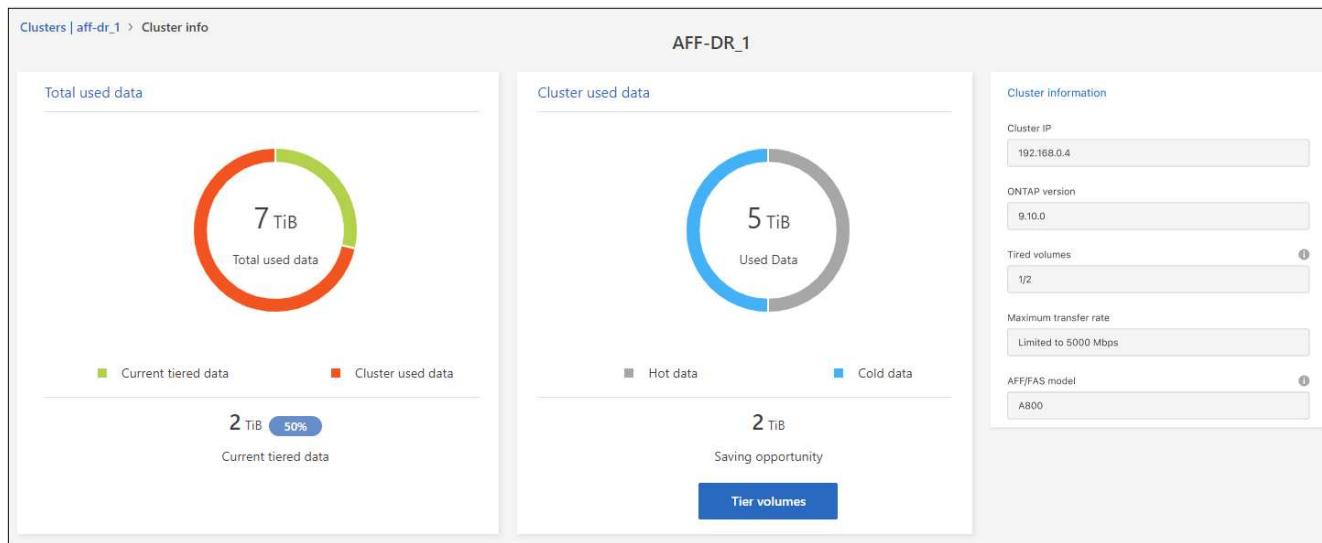
1. From the left navigation menu, select **Mobility > Cloud Tiering**.
2. From the **Clusters** page, select the menu icon **...** for a cluster and select **Cluster info**.



The screenshot shows the 'Clusters' page of the NetApp Cloud Tiering interface. At the top, there are summary counts: 3 Total clusters, 2 On-premises clusters, and 1 Cloud Volumes ONTAP. A 'Show undiscovered clusters' button is also present. Below this, a table lists clusters: 'AFF-DR_1' (On-premises cluster, Active Cluster, Operational health, 2 TiB Current tiered data, 2 TiB Additional saving opportunities, Destinations, Tier volumes, Cloud performance test, Cluster info). A 'Calculate potential tiering savings' button is located to the right of the cluster table. The 'Clusters (3/3)' section has a 'Filter by: All clusters' dropdown and an 'Active clusters first' button. A search icon is also present.

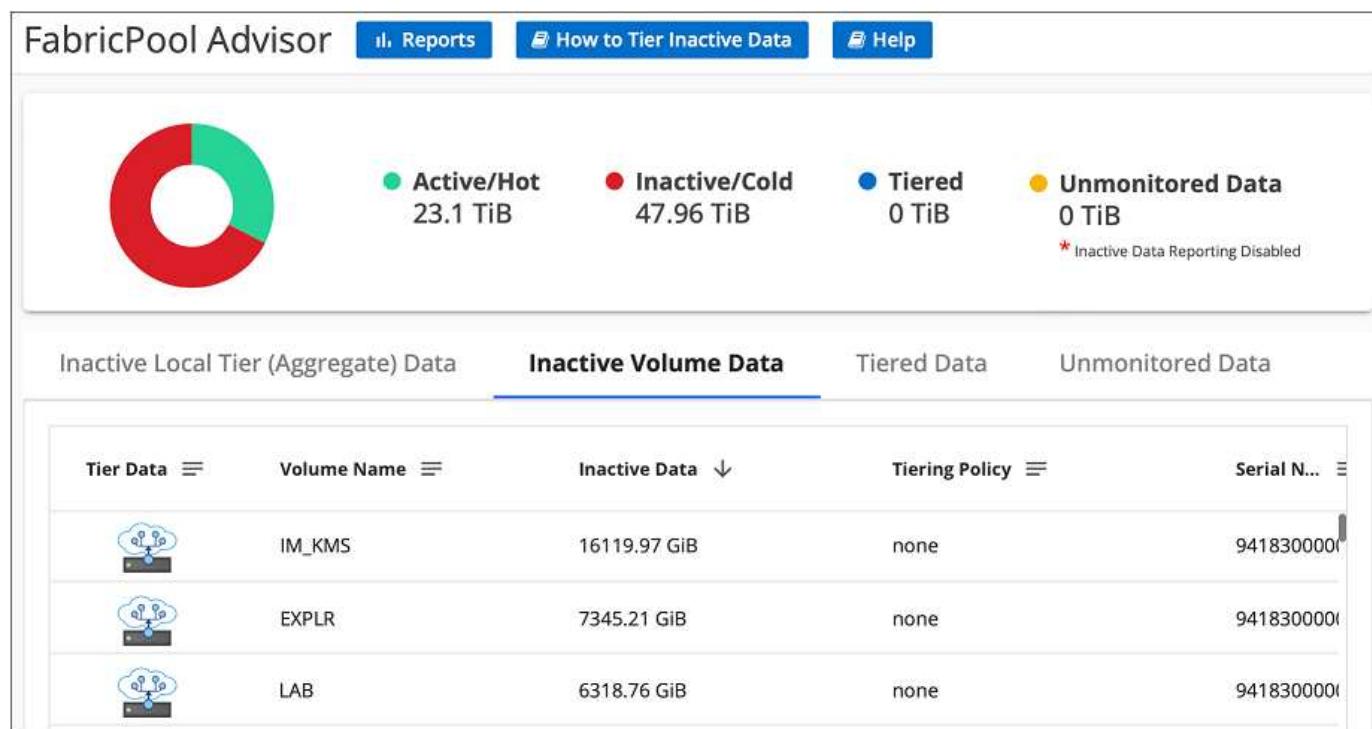
3. Review details about the cluster.

Here's an example:



Note that the display is different for Cloud Volumes ONTAP systems. While Cloud Volumes ONTAP volumes can have data tiered to the cloud, they do not use the Cloud Tiering service. [Learn how to tier inactive data from Cloud Volumes ONTAP systems to low-cost object storage](#).

You can also [view tiering information for a cluster from Active IQ Digital Advisor \(also known as Digital Advisor\)](#) if you're familiar with this NetApp product. Select **Cloud Recommendations** from the left navigation pane.



Tier data from additional volumes

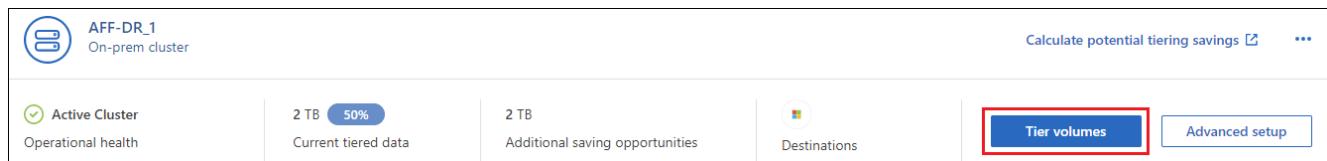
Set up data tiering for additional volumes at any time—for example, after creating a new volume.



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

Steps

1. From the left navigation menu, select **Mobility > Cloud Tiering**.
2. From the **Clusters** page, select **Tier volumes** for the cluster.



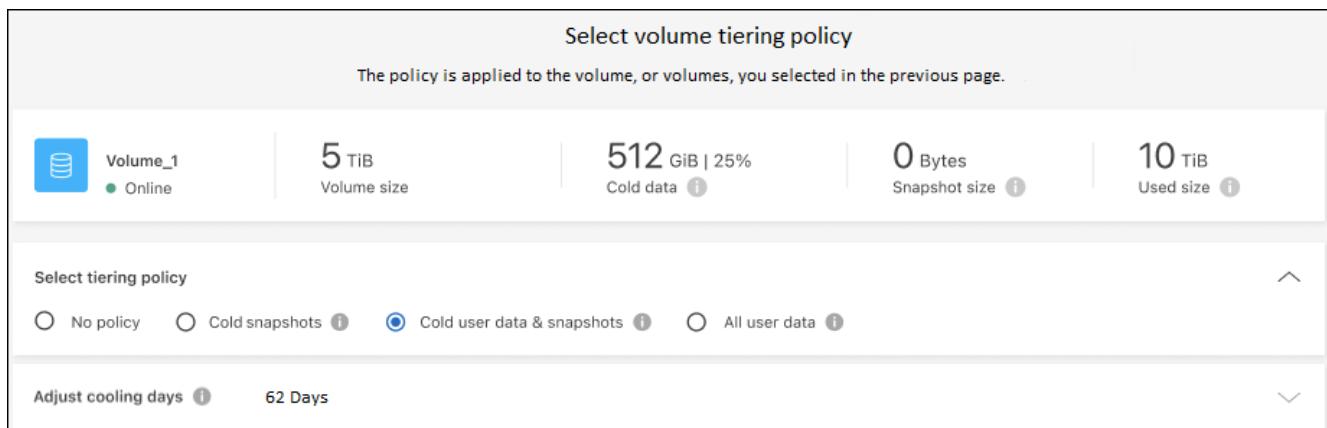
The screenshot shows the ONTAP Mobility > Cloud Tiering interface. At the top, it displays 'AFF-DR_1 On-prem cluster'. Below this, there are several status indicators: 'Active Cluster' (green checkmark), 'Operational health' (green), '2 TB 50%' (Current tiered data), '2 TB' (Additional saving opportunities), 'Destinations' (cloud icon), and 'Tier volumes' (button highlighted with a red box). There is also an 'Advanced setup' button.

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

- To select all volumes, check the box in the title row (**Volume Name**) and select **Configure volumes**.
- To select multiple volumes, check the box for each volume (**Volume_1**) and select **Configure volumes**.
- To select a single volume, select the row (or  icon) for the volume.

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

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



The screenshot shows the 'Select volume tiering policy' dialog. At the top, it says 'Select volume tiering policy' and 'The policy is applied to the volume, or volumes, you selected in the previous page.' Below this, it shows details for 'Volume_1': '5 TiB Volume size', '512 GiB | 25% Cold data', '0 Bytes Snapshot size', and '10 TiB Used size'. In the 'Select tiering policy' section, 'Cold user data & snapshots' is selected. In the 'Adjust cooling days' section, the value is set to '62 Days'.

Result

ONTAP starts tiering the selected volumes' data to the cloud.

Change a volume's tiering policy

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—it does not retroactively move data to the cloud tier.

Steps

1. From the left navigation menu, select **Mobility > CloudTiering**.

2. From the **Clusters** page, select **Tier volumes** for the cluster.
3. Click the row for a volume, select a tiering policy, optionally adjust the cooling days, and select **Apply**.

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



If you see options to "Retrieve Tiered Data", see [Migrate data from the cloud tier back to the performance tier](#) for details.

Result

ONTAP changes the tiering policy and begins tiering data based on the new policy.

Change the network bandwidth available to upload inactive data to object storage

When you activate Cloud Tiering for a cluster, by default, ONTAP can use an unlimited amount of bandwidth to transfer the inactive data from volumes in system to object storage. If tiering traffic affects user workloads, throttle the network bandwidth used during the transfer. You can choose a value between 1 and 10,000 Mbps as the maximum transfer rate.

1. From the left navigation menu, select **Mobility > Tiering**.
2. From the **Clusters** page, select the menu icon **...** for a cluster and select **Maximum transfer rate**.

3. In the *Maximum transfer rate* page, select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit. Then select **Apply**.

Maximum transfer rate

Specify the amount of network bandwidth that can be used to upload tiered data to object storage

- Unlimited
- Limited

Limited to: 10000 Mbps

1000



Apply

Cancel

This setting does not affect the bandwidth allocated to any other clusters that are tiering data.

Download a tiering report for your volumes

Download a report of the Tier Volumes page so you can review the tiering status of all the volumes on the clusters you are managing. Just select the button. Cloud Tiering generates a .CSV file that you can review and send to other groups as needed. The .CSV file includes up to 10,000 rows of data.

Tier Volumes								
	Volume Name	Aggregate/s Name	SVM Name	Node/s Name	Volume Size	Cold Data (Estimated)	Tier Status	Tiering Policy
<input type="checkbox"/>	volume_1	aggr-1	svm_1	volume_1_node	20 TB	10 TB 10 %	Tiered Volume	Cold snapshots 10 TB
<input type="checkbox"/>	volume_10	soft_restricted_aggr	svm_4	volume_10_node	10 TB	358.4 GB 70 %	Unavailable for Tiering	No Policy 512 GB
<input type="checkbox"/>	volume_11	aggr-1	svm_5	volume_11_node	10 TB	358.4 GB 70 %	Tiered Volume	Cold snapshots 512 GB
<input type="checkbox"/>	volume_12	aggr-1	svm_6	volume_12_node	10 TB	358.4 GB 70 %	Not Tiered Volume	No Policy 512 GB
<input type="checkbox"/>	volume_13	aggr-1	svm_7	volume_13_node	10 TB	5 MB 0 %	Tiered Volume	Cold snapshots 512 GB

Migrate data from the cloud tier back to the performance tier

Tiered data that is accessed from the cloud may be "re-heated" and moved back to the performance tier. However, if you want to proactively promote data to the performance tier from the cloud tier, you can do this in the *Tiering Policy* dialog. This capability is available when using ONTAP 9.8 and greater.

You can do this if you want to stop using tiering on a volume, or if you decide to keep all user data on the performance tier, but keep Snapshot copies on the cloud tier.

There are two options:

Option	Description	Affect on Tiering Policy
Bring back all data	Retrieves all volume data and Snapshot copies tiered in the cloud and promotes them to the performance tier.	Tiering policy is changed to "No policy".

Option	Description	Affect on Tiering Policy
Bring back active file system	Retrieves only active file system data tiered in the cloud and promotes it to the performance tier (Snapshot copies remain in the cloud).	Tiering policy is changed to "Cold snapshots".

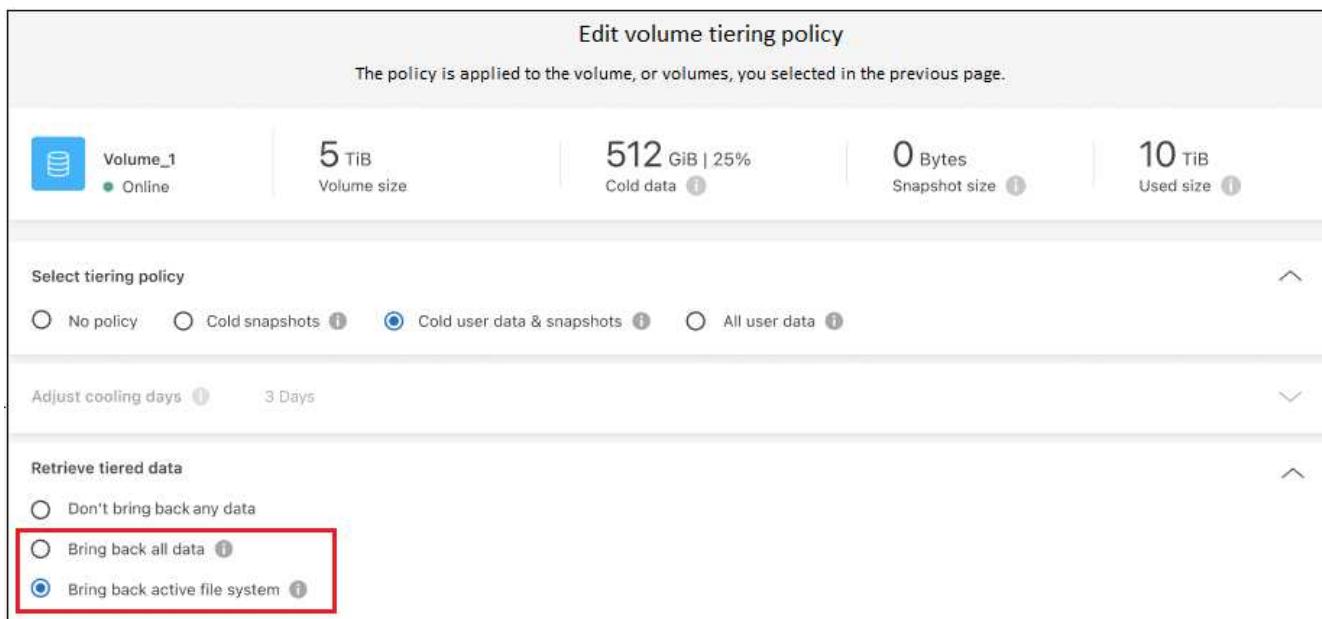


You may be charged by your cloud provider based on that amount of data transferred off the cloud.

Steps

Ensure the performance tier has enough space for data moved back from the cloud.

1. From the left navigation menu, select **Mobility > Cloud Tiering**.
2. From the **Clusters** page, select **Tier volumes** for the cluster.
3. Click the icon for the volume, choose the retrieval option you want to use, and select **Apply**.



Result

The tiering policy is changed and the tiered data starts to be migrated back to the performance tier. Depending on the amount of data in the cloud, the transfer process could take some time.

Manage tiering settings on aggregates

Each aggregate in your on-premises ONTAP systems 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

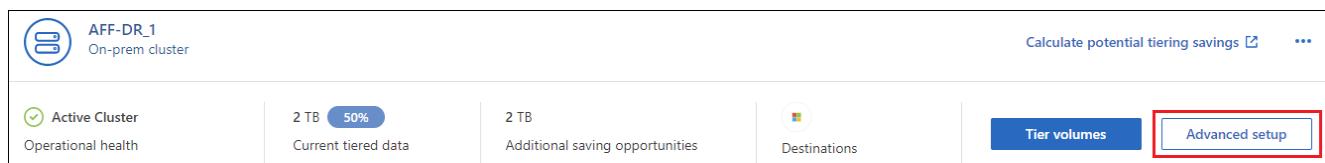
Inactive data reporting (IDR) uses a 31-day cooling period to determine which data is considered inactive. The amount of cold data that is tiered is dependent on the tiering policies set on volumes. This amount might be different than the amount of cold data detected by IDR using a 31-day cooling period.



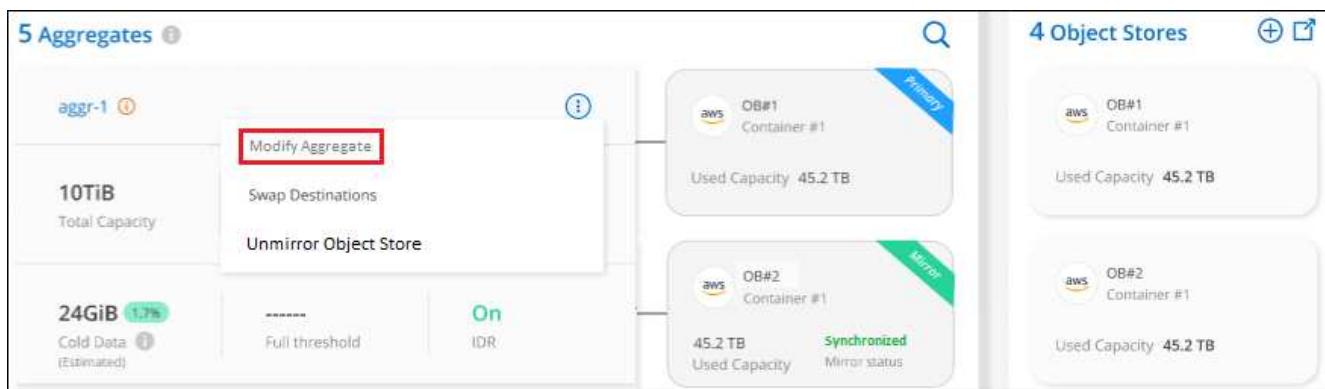
It's best to keep IDR enabled because it helps to identify your inactive data and savings opportunities. IDR must remain enabled if data tiering is enabled on an aggregate.

Steps

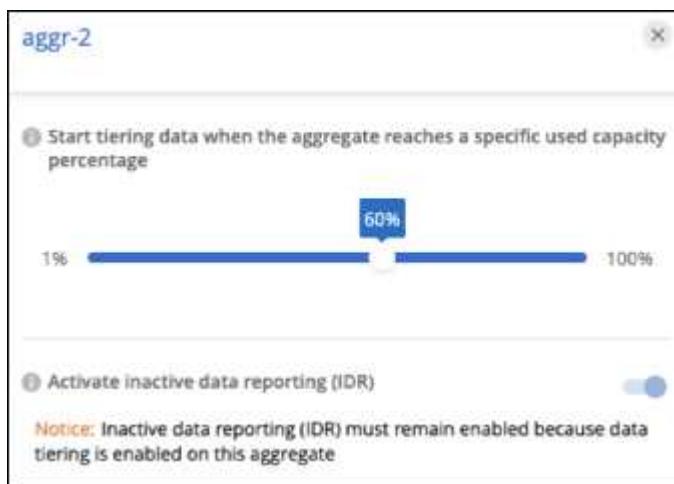
1. From the **Clusters** page, select **Advanced setup** for the selected cluster.



2. From the Advanced Setup page, select the menu icon for the aggregate and select **Modify Aggregate**.



3. In the dialog that is displayed, modify the fullness threshold and choose whether to enable or disable inactive data reporting.



4. Click **Apply**.

Fix operational health

If failures occur, Cloud Tiering displays a "Failed" operational health status on the Cluster Dashboard. The health reflects the status of the ONTAP system and the NetApp Console.

Steps

1. Identify any clusters that have an operational health of "Failed."
2. Hover over the informational "i" icon 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 Console has outbound connections to the Cloud Tiering service, to the object store, and to the ONTAP clusters that it discovers.

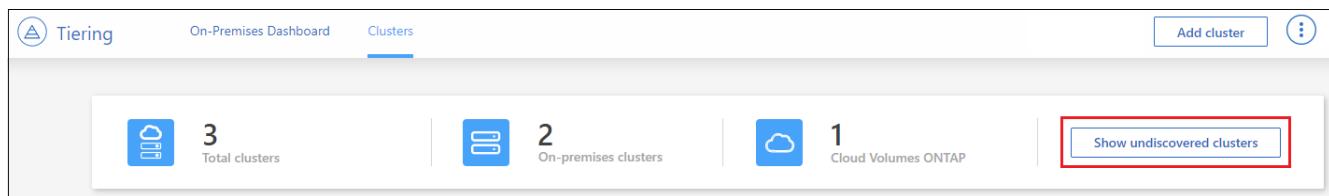
Discover additional clusters from Cloud Tiering

You can add your undiscovered on-premises ONTAP clusters to the Console from the Tiering *Cluster* page so that you can enable tiering for the cluster.

Note that buttons also appear on the Tiering *On-Prem dashboard* page for you to discover additional clusters.

Steps

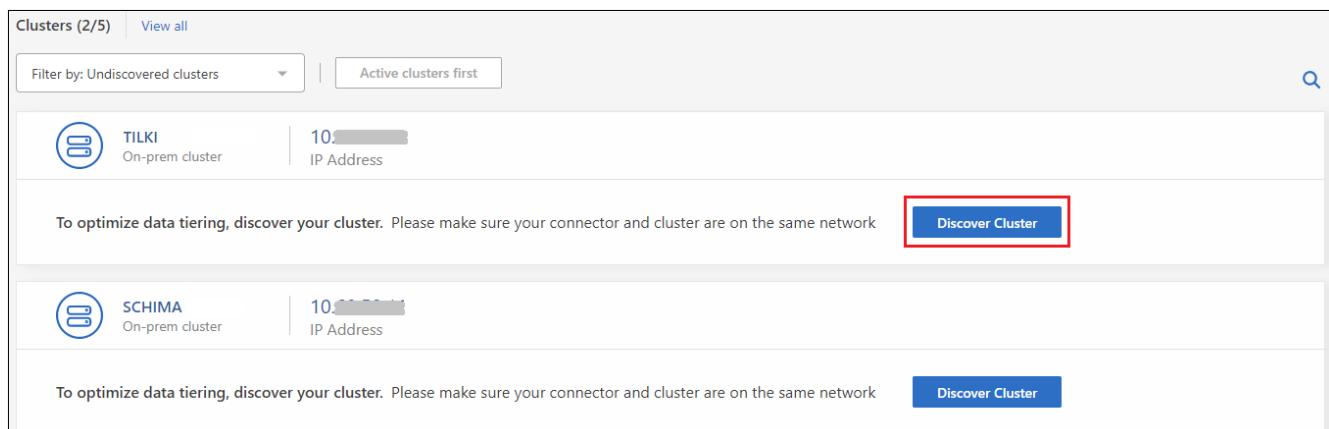
1. From Cloud Tiering, select the **Clusters** tab.
2. To see any undiscovered clusters, select **Show undiscovered clusters**.



The screenshot shows the Cloud Tiering On-Prem dashboard with the 'Clusters' tab selected. At the top, there are three summary counts: '3 Total clusters', '2 On-premises clusters', and '1 Cloud Volumes ONTAP'. Below these counts is a red box highlighting the 'Show undiscovered clusters' button.

If your NSS credentials are saved in the Console, the clusters in your account are displayed in the list.

If your NSS credentials are not saved, you are first prompted to add your credentials before you can see the undiscovered clusters.



The screenshot shows the 'Clusters (2/5)' page. It lists two undiscovered clusters: 'TILKI' (On-prem cluster) and 'SCHIMA' (On-prem cluster). Each cluster entry includes an IP address and a 'Discover Cluster' button. The 'Discover Cluster' button for the TILKI cluster is highlighted with a red box.

3. Click **Discover Cluster** for the cluster that you want to manage through the Console and implement data

tiering.

4. In the *Cluster Details* page, enter the password for the admin user account and select **Discover**.

Note that the cluster management IP address is populated based on information from your NSS account.

5. In the *Details & Credentials* page the cluster name is added as the System Name, so select **Go**.

Result

The Console discovers the cluster and adds it to the Systems page using the cluster name as the system name.

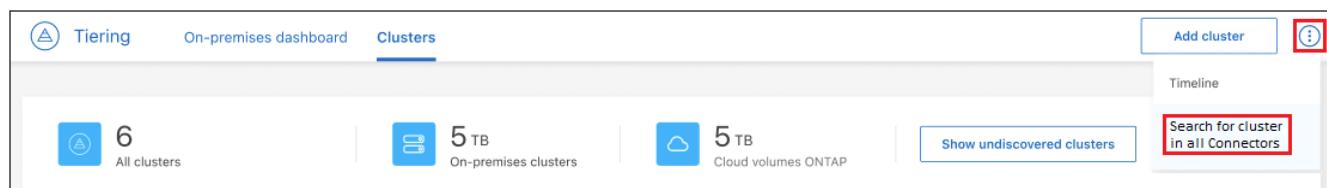
You can enable the Tiering service or other services for this cluster in the right panel.

Search for a cluster across all Console agents

If you are using multiple agents to manage all the storage in your environment, some clusters on which you want to implement tiering may be in another agent. If you are not sure which agent is managing a certain cluster, you can search across all agents using Cloud Tiering.

Steps

1. In the Cloud Tiering menu bar, select the action menu and select **Search for cluster in all agents**.



2. In the displayed Search dialog, enter the name of the cluster and select **Search**.

Cloud Tiering displays the name of the agent if it is able to find the cluster.

3. [Switch to the agent and configure tiering for the cluster](#).

Manage object storage used for data tiering in NetApp Cloud Tiering

After you've configured your on-premises ONTAP clusters to tier data to a particular object storage, you can perform additional object storage tasks using NetApp Cloud Tiering. You can add new object storage, mirror your tiered data to a secondary object storage, swap the primary and mirror object storage, remove a mirrored object store from an aggregate, and more.

View object stores configured for a cluster

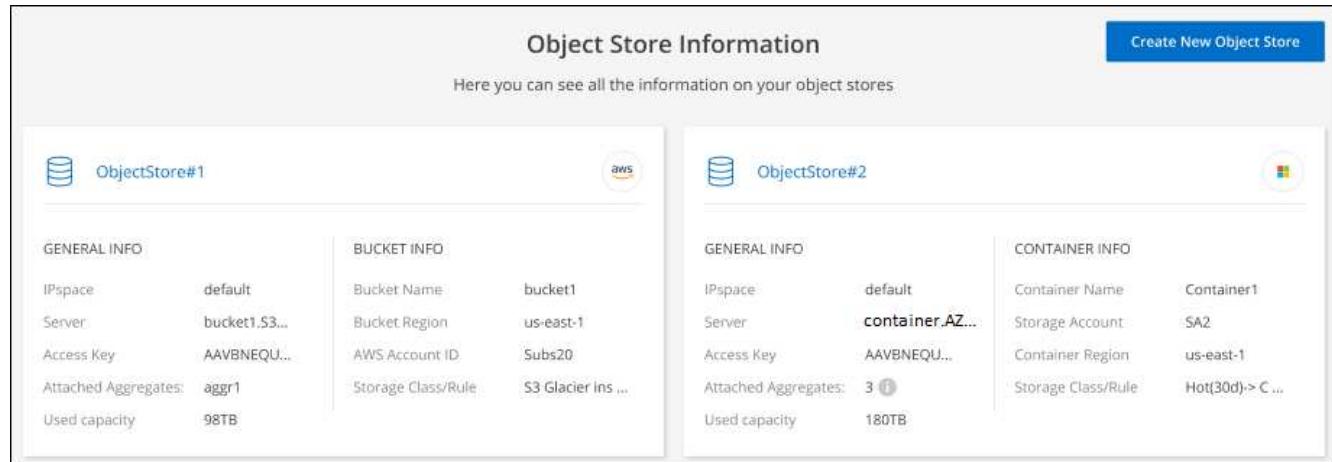
You can view all the object stores that have been configured for each cluster and to which aggregates they are attached.

Steps

1. From the **Clusters** page, select the menu icon for a cluster and select **Object Store Info**.

2. Review details about the object stores.

This example shows both an Amazon S3 and Azure Blob object store attached to different aggregates on a cluster.



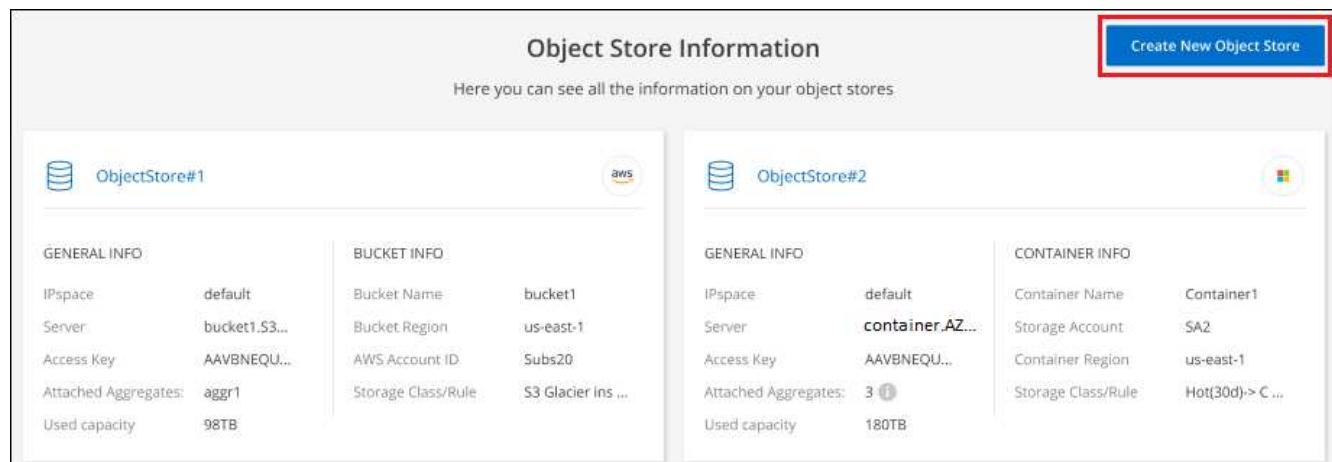
ObjectStore#1		ObjectStore#2					
GENERAL INFO		BUCKET INFO		GENERAL INFO		CONTAINER INFO	
IPspace	default	Bucket Name	bucket1	IPspace	default	Container Name	Container1
Server	bucket1.S3...	Bucket Region	us-east-1	Server	container.AZ...	Storage Account	SA2
Access Key	AAVBNEQU...	AWS Account ID	Subs20	Access Key	AAVBNEQU...	Container Region	us-east-1
Attached Aggregates:	aggr1	Storage Class/Rule	S3 Glacier Ins ...	Attached Aggregates:	3	Storage Class/Rule	Hot(30d)-> C ...
Used capacity	98TB			Used capacity	180TB		

Add a new object store

You can add a new object store for aggregates in your cluster. After you create it, you can attach it to an aggregate.

Steps

1. From the **Clusters** page, select the menu icon for a cluster and select **Object Store Info**.
2. From the Object Store Information page, select **Create New Object Store**.



ObjectStore#1		ObjectStore#2					
GENERAL INFO		BUCKET INFO		GENERAL INFO		CONTAINER INFO	
IPspace	default	Bucket Name	bucket1	IPspace	default	Container Name	Container1
Server	bucket1.S3...	Bucket Region	us-east-1	Server	container.AZ...	Storage Account	SA2
Access Key	AAVBNEQU...	AWS Account ID	Subs20	Access Key	AAVBNEQU...	Container Region	us-east-1
Attached Aggregates:	aggr1	Storage Class/Rule	S3 Glacier Ins ...	Attached Aggregates:	3	Storage Class/Rule	Hot(30d)-> C ...
Used capacity	98TB			Used capacity	180TB		

The object store wizard starts. The example below shows how to create an object store in Amazon S3.

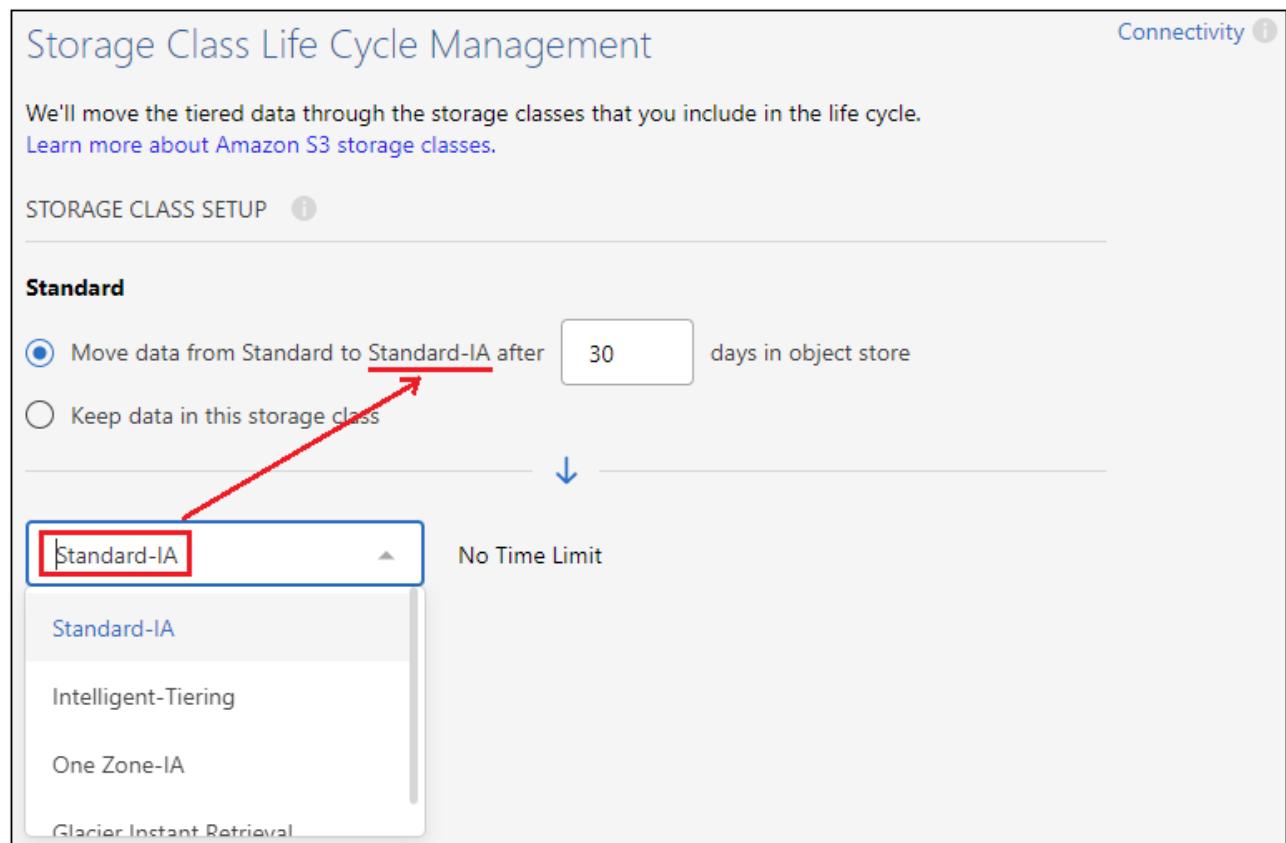
3. **Define Object Storage Name:** Enter a name for this object storage. It must be unique from any other object storage you may be using with aggregates on this cluster.
4. **Select Provider:** Select the provider, for example **Amazon Web Services**, and select **Continue**.
5. Complete the steps on the **Create Object Storage** pages:
 - a. **S3 Bucket:** Add a new S3 bucket or select an existing S3 bucket that starts with the prefix *fabric-pool*. Then enter the AWS Account ID that provides access to the bucket, select the bucket region, and select **Continue**.

The *fabric-pool* prefix is required because the IAM policy for the Console agent enables the instance to perform S3 actions on buckets named with that exact prefix. For example, you could name the S3 bucket *fabric-pool-AFF1*, where AFF1 is the name of the cluster.

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

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

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



The screenshot shows the 'Storage Class Life Cycle Management' interface. In the 'STORAGE CLASS SETUP' section, under the 'Standard' heading, there are two options: 'Move data from Standard to Standard-IA after' and 'Keep data in this storage class'. The first option is selected, with '30' days in object store. A red arrow points from the 'Standard-IA' label in the dropdown menu to the selected radio button. The dropdown menu also shows 'Standard-IA' highlighted with a red box. Other options in the dropdown include 'Standard-IA', 'Intelligent-Tiering', 'One Zone-IA', and 'Glacier Instant Retrieval'. The 'Connectivity' tab is visible in the top right corner.

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

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

The IAM user must be in the same AWS account as the bucket that you selected or created on the **S3 Bucket** page. See the required permissions in the section about activating tiering.

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

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

The object store is created.

Now you can attach the object store to an aggregate in your cluster.

Attach a second object store to an aggregate for mirroring

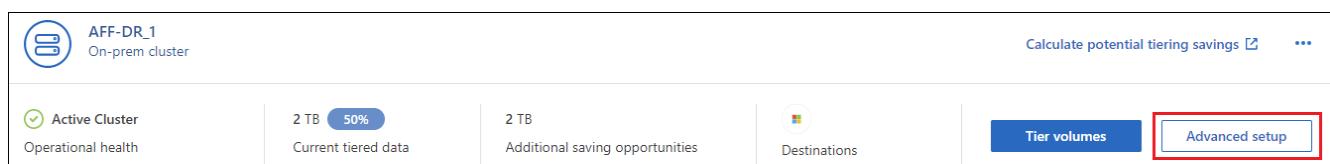
You can attach a second object store to an aggregate to create a FabricPool mirror to synchronously tier data to two object stores. You must have one object store already attached to the aggregate. [Learn more about FabricPool mirrors.](#)

When you use a MetroCluster configuration, it's a best practice to use object stores in the public cloud that are in different availability zones. [Learn more about MetroCluster requirements in the ONTAP documentation.](#) Within a MetroCluster, it is not recommended to use unmirrored aggregates, doing so will give you an error message.

When you use StorageGRID as your object store in a MetroCluster configuration, both ONTAP systems can perform FabricPool tiering to a single StorageGRID system. Each ONTAP system must tier data to different buckets.

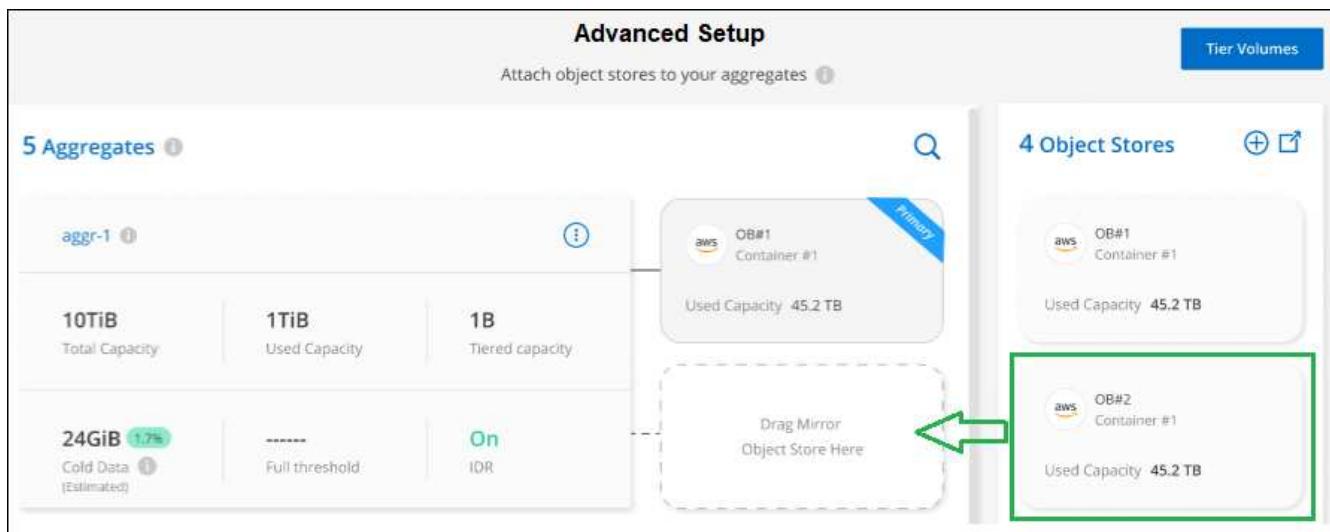
Steps

1. From the **Clusters** page, select **Advanced setup** for the selected cluster.



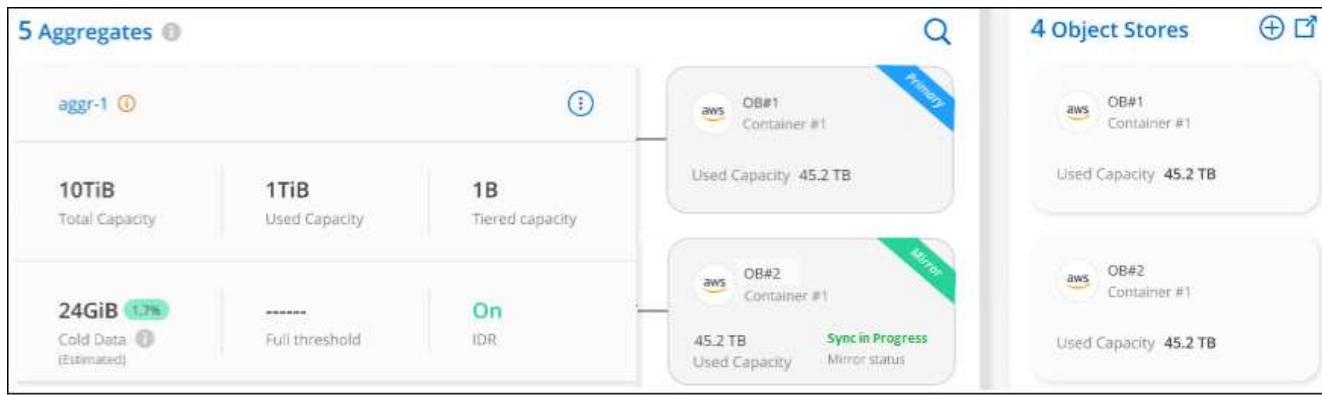
The screenshot shows the ONTAP Clusters page for an 'AFF-DR_1 On-prem cluster'. The 'Advanced setup' button is highlighted with a red box. Other visible buttons include 'Tier volumes' and 'Calculate potential tiering savings'.

2. From the Advanced Setup page, drag the object store you want to use to the location for the mirror object store.



The screenshot shows the 'Advanced Setup' page. On the left, there are details for 'aggr-1': Total Capacity 10TiB, Used Capacity 1TiB, Tiered capacity 1B, and a cold data threshold of 24GiB (1.7%). On the right, there are two object stores: 'OB#1 Container #1' (aws, Used Capacity 45.2 TB) and 'OB#2 Container #1' (aws, Used Capacity 45.2 TB). A green arrow points from the 'OB#2' box to a dashed box labeled 'Drag Mirror Object Store Here', indicating where to drop the second object store to create a mirror.

3. In the Attach Object Store dialog, select **Attach** and the second object store is attached to the aggregate.



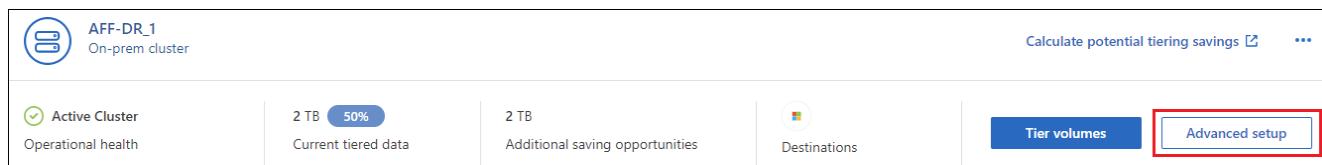
The Mirror status will appear as "Sync in progress" while the 2 object stores are synchronizing. The status will change to "Synchronized" when synchronization is complete.

Swap the primary and mirror object store

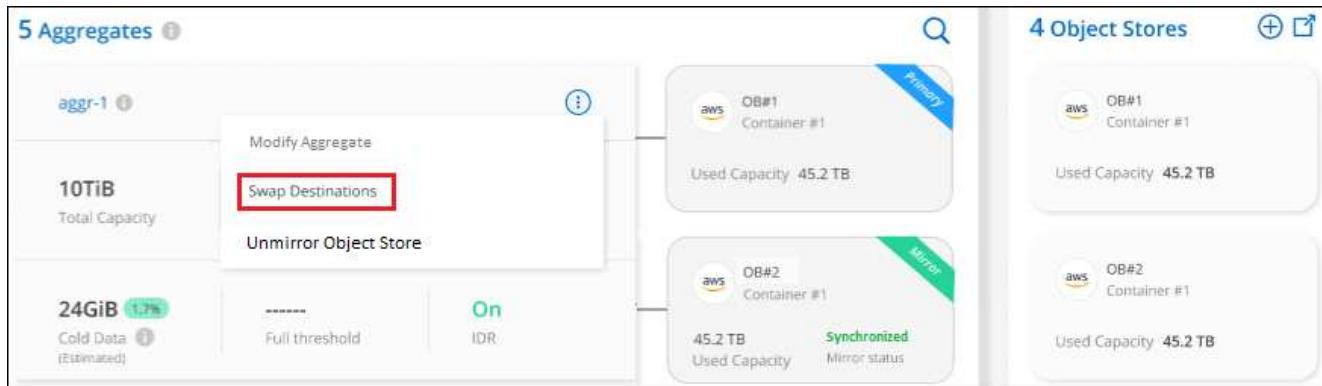
You can swap the primary and mirror object store for an aggregate. The object store mirror becomes the primary, and the original primary becomes the mirror.

Steps

1. From the **Clusters** page, select **Advanced setup** for the selected cluster.



2. From the Advanced Setup page, select the menu icon for the aggregate and select **Swap Destinations**.



3. Approve the action in the dialog box and the primary and mirror objects stores are swapped.

Remove a mirror object store from an aggregate

You can remove a FabricPool mirror if you no longer need to replicate to an additional object store.

Steps

1. From the **Clusters** page, select **Advanced setup** for the selected cluster.

AFF-DR_1
On-prem cluster

Active Cluster
Operational health

2 TB 50%
Current tiered data

2 TB
Additional saving opportunities

Destinations

Tier volumes

Advanced setup

- From the Advanced Setup page, select the menu icon for the aggregate and select **Unmirror Object Store**.

5 Aggregates

aggr-1

10TiB Total Capacity

24GiB 1.7% Cold Data (Estimated)

Modify Aggregate Swap Destinations Unmirror Object Store

----- Full threshold On IDR:

aws OB#1 Container #1 Primary
Used Capacity 45.2 TB

aws OB#2 Container #1 Mirror
45.2 TB Used Capacity Synchronized Mirror status

4 Object Stores

aws OB#1 Container #1
Used Capacity 45.2 TB

aws OB#2 Container #1
Used Capacity 45.2 TB

The mirror object store is removed from the aggregate and the tiered data is no longer replicated.



When removing the mirror object store from a MetroCluster configuration, you'll be prompted whether you want to remove the primary object store as well. You can choose to keep the primary object store attached to the aggregate, or to remove it.

Migrate your tiered data to a different cloud provider

Cloud Tiering enables you to easily migrate your tiered data to a different cloud provider. For example, if you want to move from Amazon S3 to Azure Blob, you can follow the steps listed above in this order:

1. Add an Azure Blob object store.
2. Attach this new object store as the mirror to the existing aggregate.
3. Swap the primary and mirror object stores.
4. Unmirror the Amazon S3 object store.

Measure network latency and throughput performance in NetApp Cloud Tiering

Run a Cloud Performance Test to measure network latency and throughput performance from an ONTAP cluster to an object store before and after setting up data tiering in NetApp Cloud Tiering. The test also identifies any failures that occurred.

Here are sample performance results:

Your cluster performance results

Operation	Size	Avg.Latency (ms)	Throughput
PUT	4 MB	502	408.06 MB
GET	4 KB	79	15.05 MB
GET	8 KB	197	28.35 MB
GET	32 KB	291	109.71 MB
GET	256 KB	361	714.39 MB

Notice: We recommend that you run this check when the cluster is under 50% CPU utilization.

Before you get started

It's best to run this check when the cluster is under 50% CPU utilization.

Steps for a cluster that hasn't been set up for tiering

1. From the left navigation menu, select **Mobility > Cloud Tiering**.
2. From the **Clusters** page, select the menu icon for a cluster and select **Cloud Performance Test**.
3. Review the details and select **Continue**.
4. Follow the prompts to provide the required information.

The information that you need to provide is the same as if you were setting up tiering on the cluster.

5. Optionally continue to the Tier Volumes wizard to complete the setup.

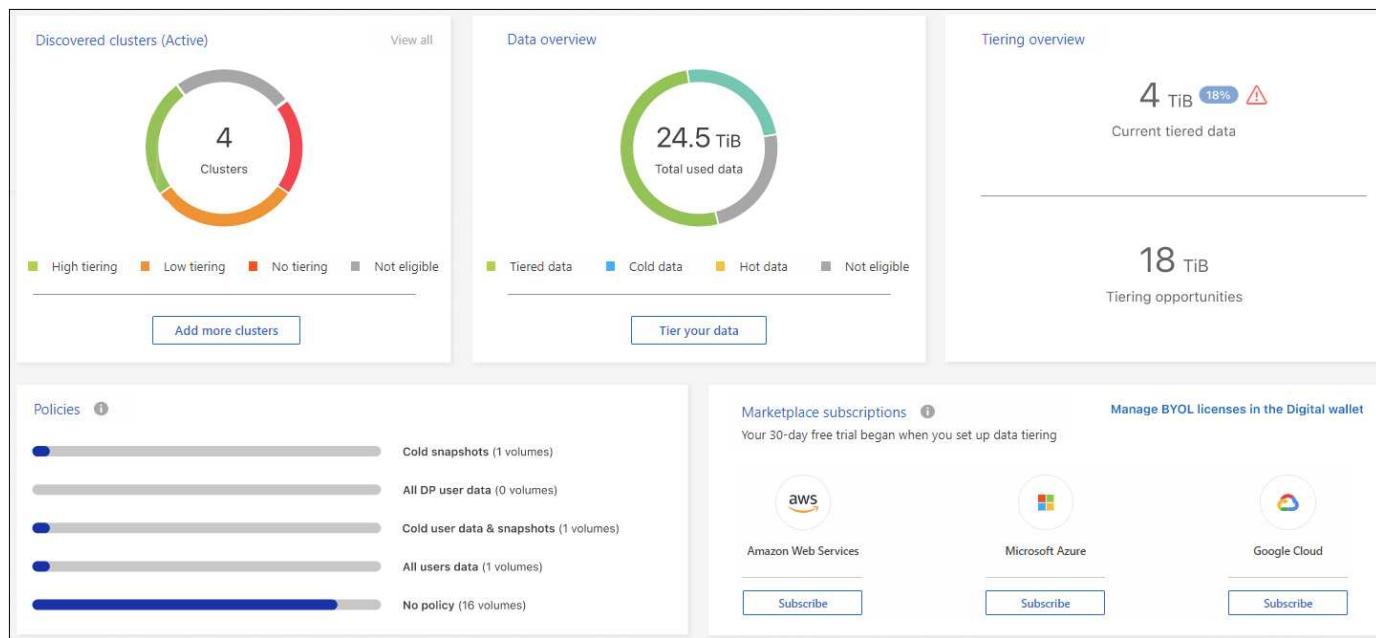
Steps for a cluster that has been set up for tiering

1. From the left navigation menu, select **Mobility > Cloud Tiering**.
2. From the **Clusters** page, select the menu icon for a cluster and select **Cloud Performance Test**.
3. Select a node from the drop-down list.
4. View the results or recheck the performance.

Get an overview of data tiering from your clusters in NetApp Cloud Tiering

NetApp Cloud Tiering provides an aggregated view of data tiering from each of your on-premises clusters. This overview provides a clear picture of your environment and enables you to take proper actions.

Select **Cloud Tiering > On-Premises Dashboard** to view the following details about your environment.



Discovered clusters

The number of on-premises clusters that Cloud Tiering has discovered. The chart provides an overview of the tiering status for these clusters.

- High tiering - Clusters that are tiering over 20% of their cold data
- Low tiering - Clusters that are tiering less than 20% of their cold data
- No tiering - Clusters that aren't tiering any data
- Not eligible - Clusters that don't support data tiering

Data overview

The amount of data that is being used by all discovered clusters. The chart shows how much data is being tiered for these clusters.

- Tiered data - Total cold data being tiered to the cloud
- Cold data - Total cold data that is not being tiered
- Hot data - Total hot data that is active being used
- Not eligible - Total data that is not being tiered because the cluster or volume doesn't support data tiering

Tiering overview

The amount of data that is currently being tiered, and the amount of cold data that could potentially be tiered.

Policies

The number of times that each tiering policy has been applied to a volume.

Marketplace subscriptions

The number of clusters associated with each type of Marketplace Subscription and an indication about your subscription status.

Monitor the status of tiering alerts from NetApp Cloud Tiering

You can view the status of tiering alerts from NetApp Cloud Tiering in the NetApp Console Notification Center.

The Notification Center tracks the progress of tiering incidents so you can verify whether they have been resolved or not. You can display the notifications by selecting the  in the Console menu bar.

At this time, there is one tiering event that will appear as a notification:

Tier additional data from cluster <name> to object storage to save storage space

This notification is a "Recommendation" to improve system efficiency and reduce storage costs. It indicates that a cluster is tiering less than 20% of its cold data - including clusters that are tiering no data. It provides a link to the [Cloud Tiering total cost of ownership and savings calculator](#) to help you calculate your cost savings.

The NetApp Console does not send an email for this notification.

[Learn more about the Notification Center.](#)

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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