



# Use BlueXP tiering

## BlueXP tiering

NetApp  
February 11, 2024

# Table of Contents

- Use BlueXP tiering ..... 1
  - Managing data tiering for your clusters ..... 1
  - Managing object storage used for data tiering ..... 9
  - Measure network latency and throughput performance ..... 14
  - Get an overview of data tiering from your clusters ..... 15
  - Monitor the status of tiering alerts ..... 17

# Use BlueXP tiering

## Managing data tiering for your clusters

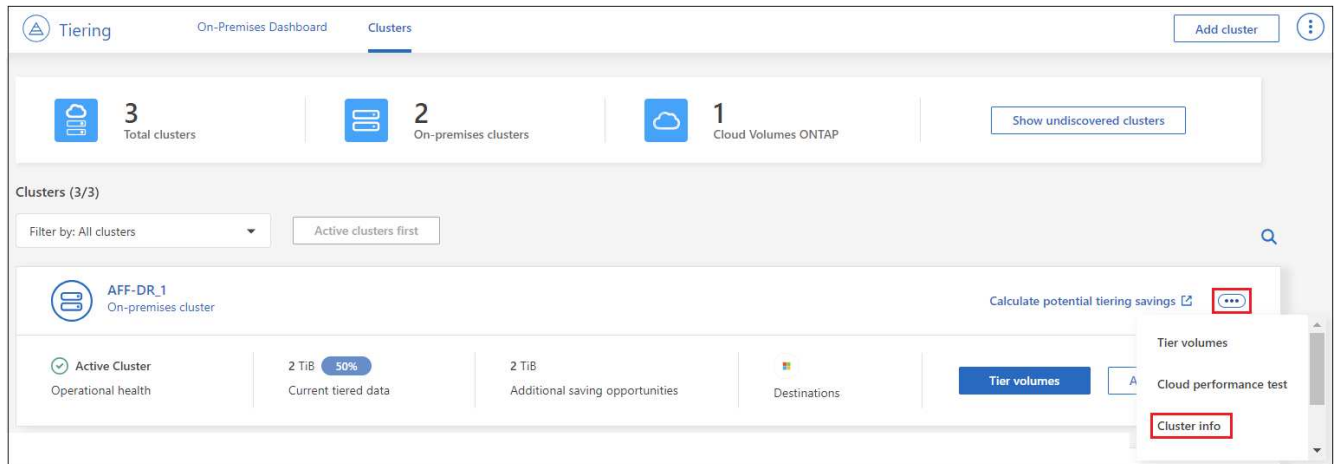
Now that you've set up data tiering from your on-prem ONTAP clusters, you can tier data from additional volumes, change a volume's tiering policy, discover additional clusters, and more.

### 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. BlueXP tiering provides this information for each cluster.

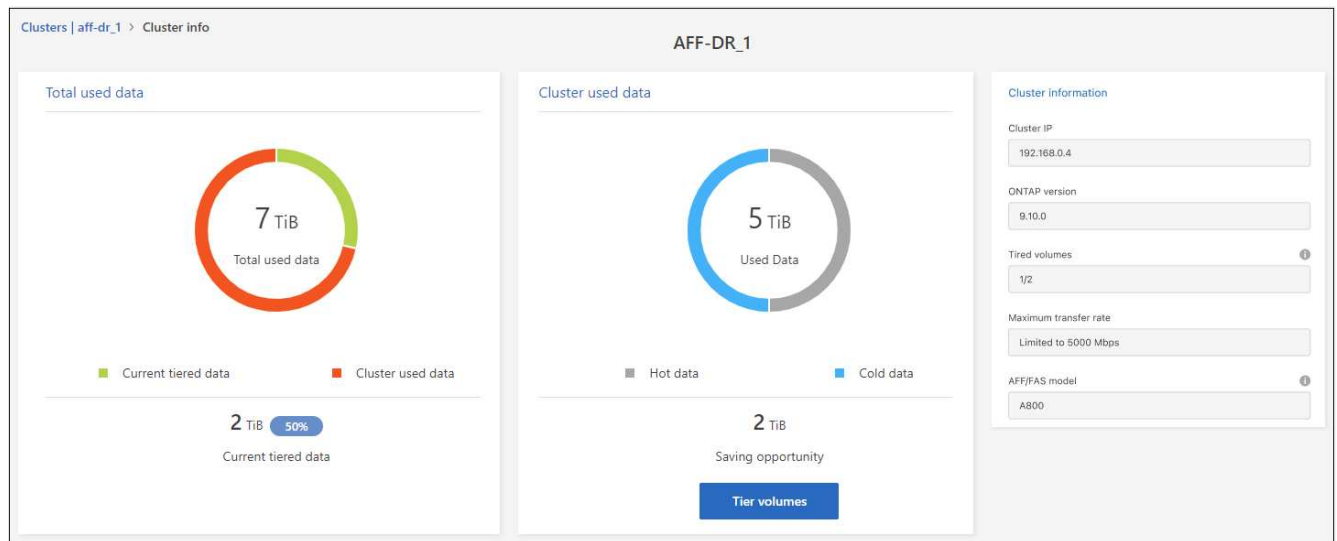
#### Steps

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



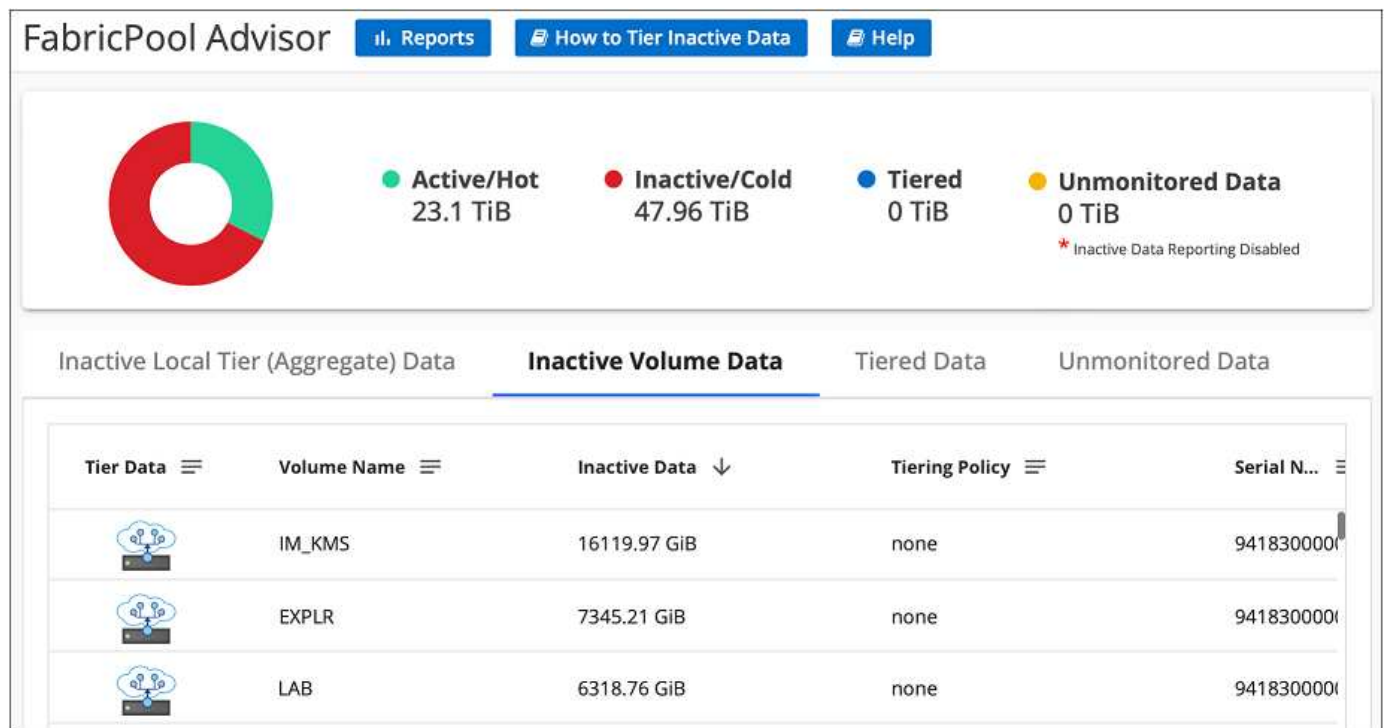
3. Review details about the cluster.

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 BlueXP 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 Digital Advisor](#) if you're familiar with this NetApp product. Just select **Cloud Recommendations** from the left navigation pane.



## Tiering data from additional volumes

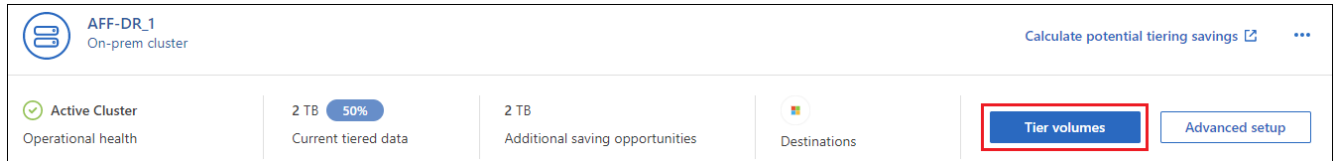
Set up data tiering for additional volumes at any time—for example, after creating a new volume.




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

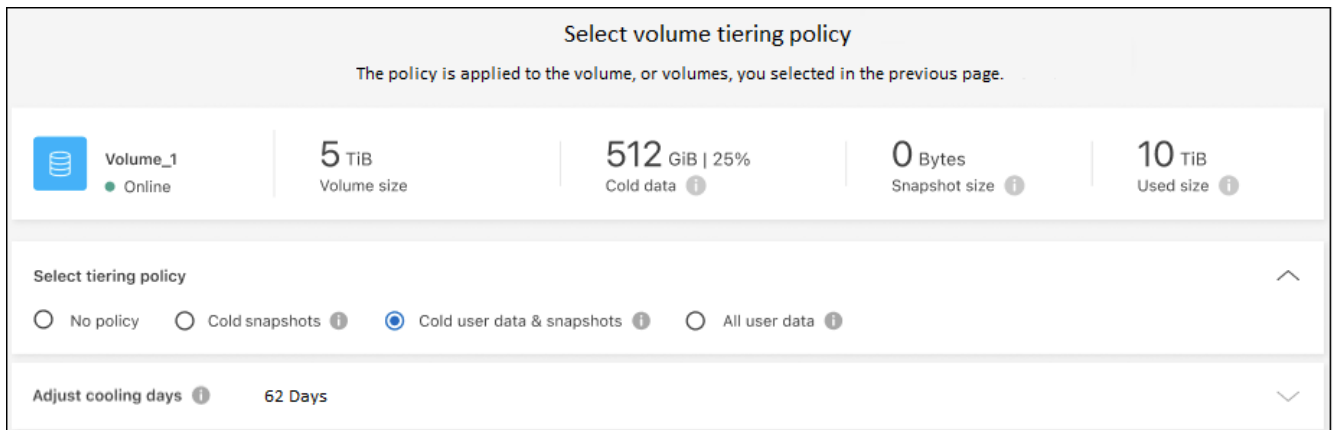
## Steps

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



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 click **Configure volumes**.
  - To select multiple volumes, check the box for each volume (☒ Volume\_1) and click **Configure volumes**.
  - To select a single volume, click 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 click **Apply**.

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



## Result

The selected volumes start to have their data tiered to the cloud.

## Changing 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 > Tiering**.
2. From the **Clusters** page, click **Tier volumes** for the cluster.

- Click the row for a volume, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

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

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

## Result

The tiering policy is changed and data begins to be tiered based on the new policy.

## Changing the network bandwidth available to upload inactive data to object storage

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

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

- 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 click **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

You can 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 click the [↓](#) button. BlueXP 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.

Volumes (16)

Tier Volumes

🔍

↓

<input type="checkbox"/>	Volume Name	Aggregate/s Name	SVM Name	Node/s Name	Volume Size	Cold Data (Estimated)	Tier Status	Tiering Policy	Used Size
<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

## Migrating 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 might 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

Make sure you have enough space in the performance tier for all the data that is being moved back from the cloud.

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

### Edit volume tiering policy

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



**Volume\_1**  
Online

**5 TiB**  
Volume size

**512 GiB | 25%**  
Cold data ⓘ

**0 Bytes**  
Snapshot size ⓘ

**10 TiB**  
Used size ⓘ

#### Select tiering policy

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

#### Adjust cooling days ⓘ

3 Days

#### Retrieve tiered data

☐ Don't bring back any data
☐ Bring back all data ⓘ
☒ Bring back active file system ⓘ

## 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.

## Managing tiering settings on aggregates

Each aggregate in your on-prem 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

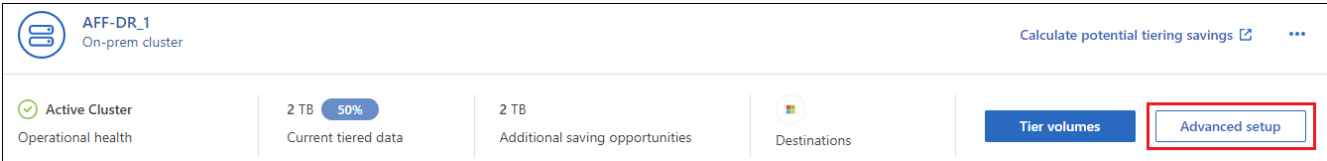
Inactive data reporting (IDR) uses a 31-day cooling period to determine which data is considered inactive. The amount of cold data that is tiered is dependent on the tiering policies set on volumes. This amount might be different than the amount of cold data detected by IDR using a 31-day cooling period.



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 was enabled on an aggregate.

Steps

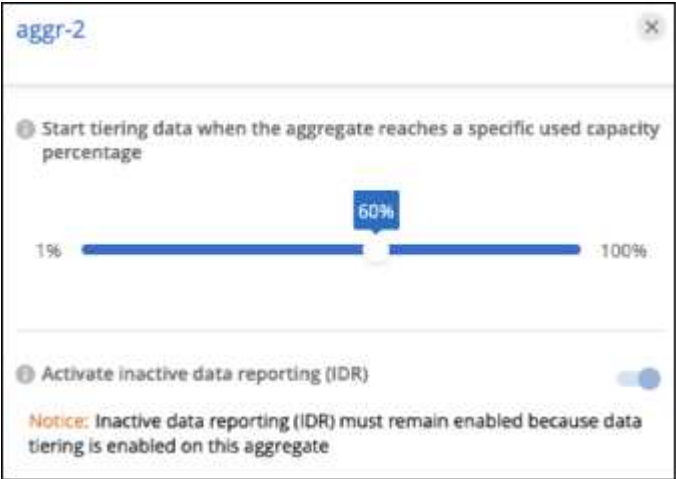
- 1. From the **Clusters** page, click **Advanced setup** for the selected cluster.



- 2. From the Advanced Setup page, click 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**.

## Fixing operational health

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

### 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 BlueXP has outbound connections to the BlueXP tiering service, to the object store, and to the ONTAP clusters that it discovers.

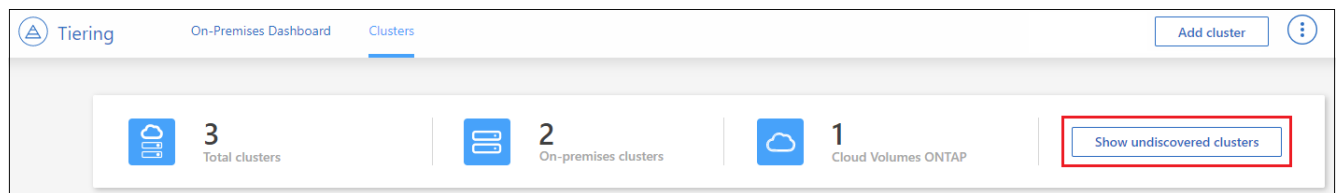
## Discovering additional clusters from BlueXP tiering

You can add your undiscovered on-prem ONTAP clusters to BlueXP 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.

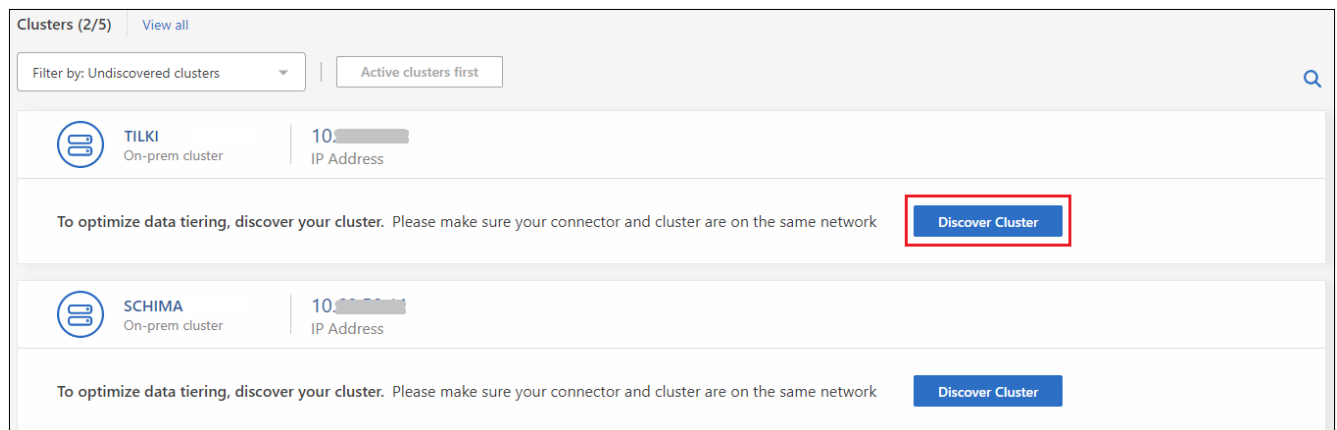
### Steps

1. From BlueXP tiering, click the **Clusters** tab.
2. To see any undiscovered clusters, click **Show undiscovered clusters**.



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

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



3. Click **Discover Cluster** for the cluster that you want to manage through BlueXP and implement data tiering.
4. In the *Cluster Details* page, enter the password for the admin user account and click **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 Working Environment Name, so just click **Go**.

## Result

BlueXP discovers the cluster and adds it to a working environment in the Canvas using the cluster name as the working environment name.

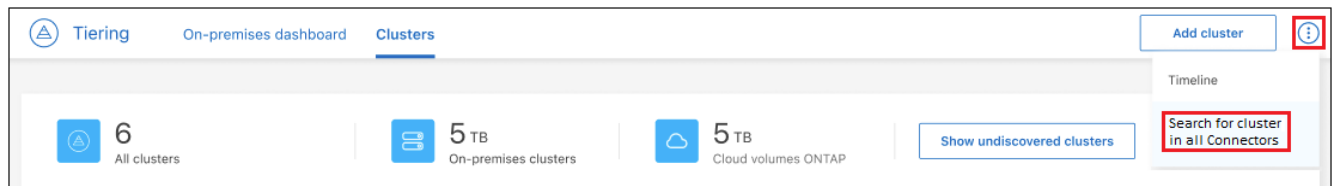
You can enable the Tiering service or other services for this cluster in the right panel.

## Search for a cluster across all BlueXP Connectors

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

### Steps

1. In the BlueXP tiering menu bar, click the action menu and select **Search for cluster in all Connectors**.



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

BlueXP tiering displays the name of the Connector if it is able to find the cluster.

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

## Managing object storage used for data tiering

After you've configured your on-prem ONTAP clusters to tier data to a particular object storage, you can perform additional object storage tasks. 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.

### Viewing object stores configured for a cluster

You might want to see all the object stores that have been configured for your cluster and to which aggregates they are attached. BlueXP tiering provides this information for each cluster.

### Steps

1. From the **Clusters** page, click 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.

### Object Store Information

Here you can see all the information on your object stores

Create New Object Store

#### ObjectStore#1

**GENERAL INFO**

IPspace	default
Server	bucket1.S3...
Access Key	AAVBNEQU...
Attached Aggregates:	aggr1
Used capacity	98TB

**BUCKET INFO**

Bucket Name	bucket1
Bucket Region	us-east-1
AWS Account ID	Subs20
Storage Class/Rule	S3 Glacier Ins ...

#### ObjectStore#2

**GENERAL INFO**

IPspace	default
Server	container.AZ...
Access Key	AAVBNEQU...
Attached Aggregates:	3 ⓘ
Used capacity	180TB

**CONTAINER INFO**

Container Name	Container1
Storage Account	SA2
Container Region	us-east-1
Storage Class/Rule	Hot(30d)-> C ...

## Adding a new object store

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

### Steps

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

### Object Store Information

Here you can see all the information on your object stores

Create New Object Store

#### ObjectStore#1

**GENERAL INFO**

IPspace	default
Server	bucket1.S3...
Access Key	AAVBNEQU...
Attached Aggregates:	aggr1
Used capacity	98TB

**BUCKET INFO**

Bucket Name	bucket1
Bucket Region	us-east-1
AWS Account ID	Subs20
Storage Class/Rule	S3 Glacier Ins ...

#### ObjectStore#2

**GENERAL INFO**

IPspace	default
Server	container.AZ...
Access Key	AAVBNEQU...
Attached Aggregates:	3 ⓘ
Used capacity	180TB

**CONTAINER INFO**

Container Name	Container1
Storage Account	SA2
Container Region	us-east-1
Storage Class/Rule	Hot(30d)-> C ...

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 click **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 click **Continue**.

The *fabric-pool* prefix is required because the IAM policy for the 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 Lifecycle:** BlueXP 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 click **Continue**. For example, the screenshot below shows that tiered data is assigned to the *Standard-IA* class from the *Standard* class after 45 days in object storage.

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

Storage Class Life Cycle Management Connectivity ⓘ

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

STORAGE CLASS SETUP ⓘ

**Standard**

☒ Move data from Standard to Standard-IA after 30 days in object store

☐ Keep data in this storage class

↓

Standard-IA No Time Limit

- Standard-IA
- Intelligent-Tiering
- One Zone-IA
- Glacier Instant Retrieval

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

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

Selecting the correct IPspace ensures that BlueXP 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.

## Attaching a second object store to an aggregate for mirroring

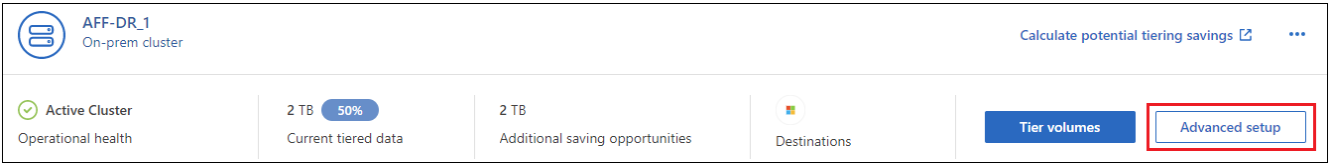
You can attach a second object store to an aggregate to create a FabricPool mirror to synchronously tier data to two object stores. You must have one object store already attached to the aggregate. [Learn more about FabricPool mirrors.](#)

When using 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.](#)

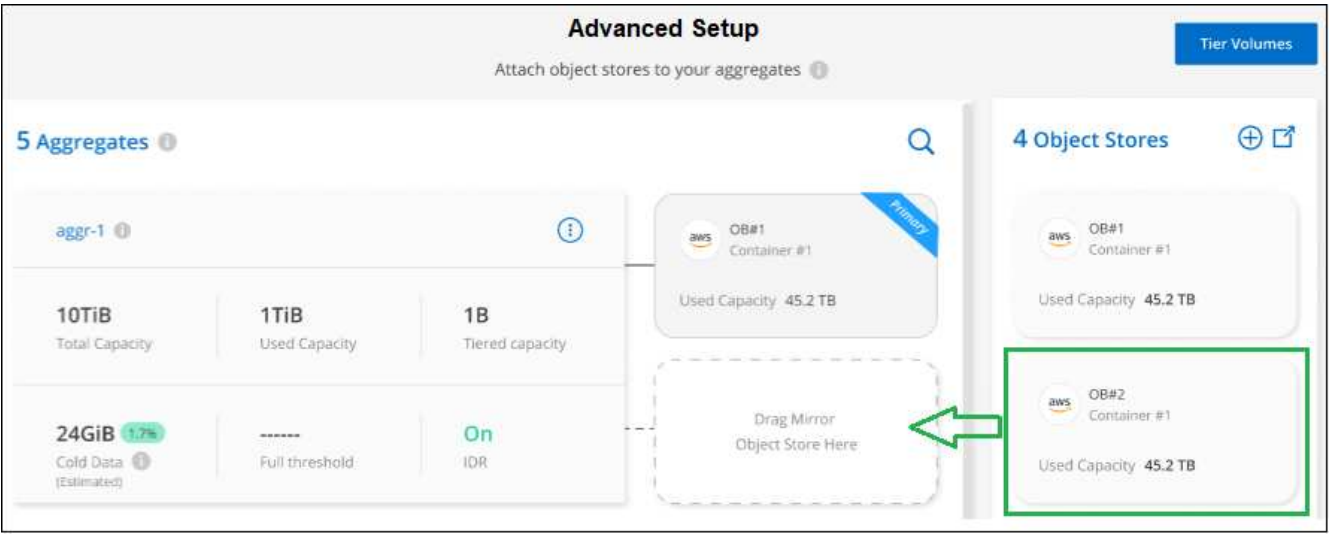
Note that when using StorageGRID as your object store in a MetroCluster configuration, both ONTAP systems can perform FabricPool tiering to a single StorageGRID system. Each ONTAP system must tier data to different buckets.

### Steps

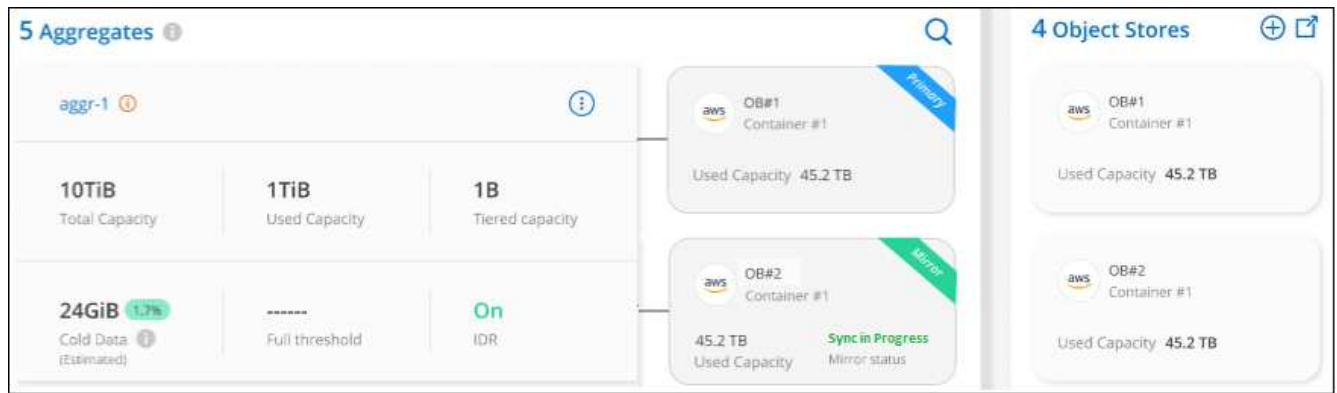
1. From the **Clusters** page, click **Advanced setup** for the selected cluster.



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



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



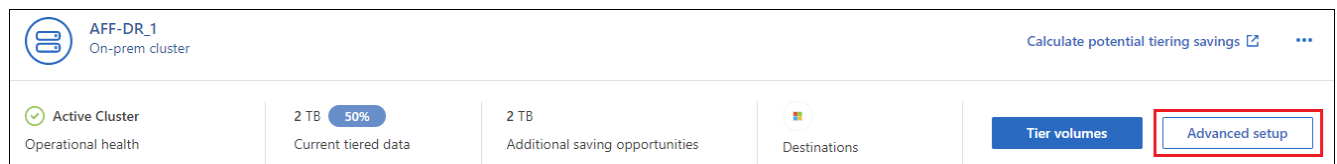
The Mirror status will appear as "Sync in progress" while the 2 object stores are synchronizing. The status will change to "Synchronized" when synchronization is complete.

## Swapping the primary and mirror object store

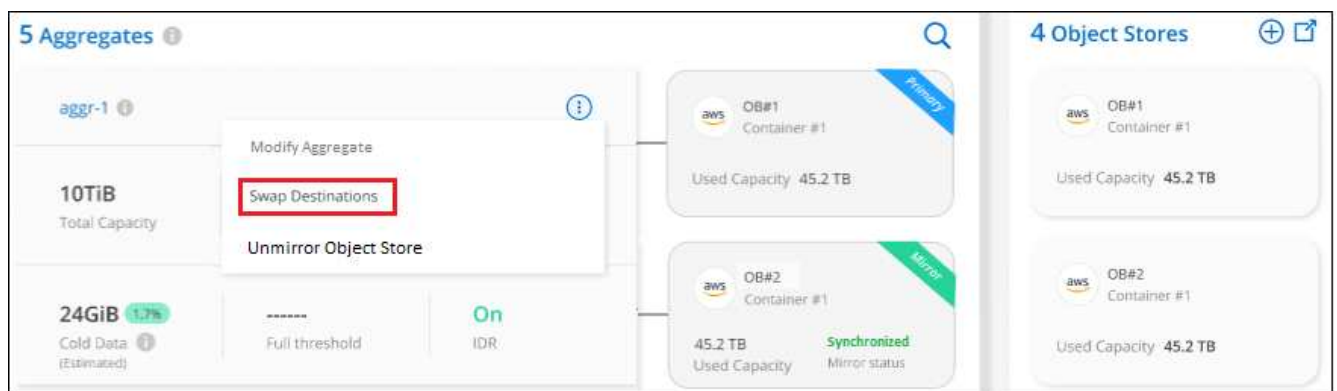
You can swap the primary and mirror object store for an aggregate. The object store mirror becomes the primary, and the original primary becomes the mirror.

### Steps

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



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



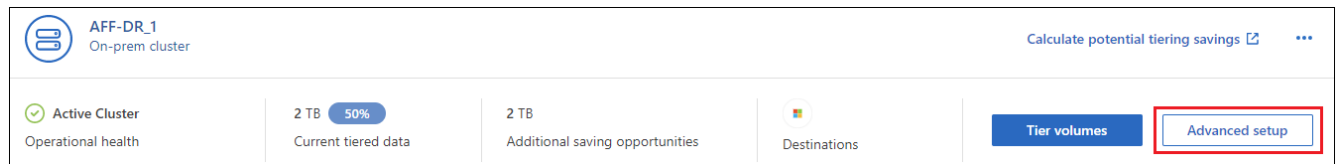
3. Approve the action in the dialog box and the primary and mirror objects stores are swapped.

## Removing a mirror object store from an aggregate

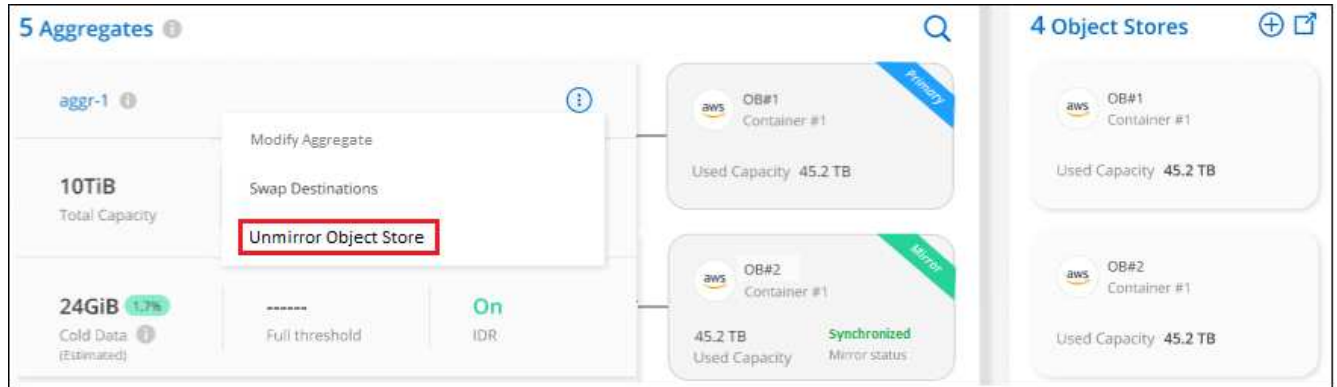
You can remove a FabricPool mirror if you no longer need to replicate to an additional object store.

### Steps

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



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



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.

## Migrating your tiered data to a different cloud provider

BlueXP 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:

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

## Measure network latency and throughput performance

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. The test also identifies any failures that occurred.

Here are sample performance results:

## Your cluster performance results

Node: aff-02

object-store-1

Last check: 03/28/2023 01:30 pm

[Recheck performance](#)

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 > Tiering**.
2. From the **Clusters** page, click the menu icon for a cluster and select **Cloud Performance Test**.
3. Review the details and click **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

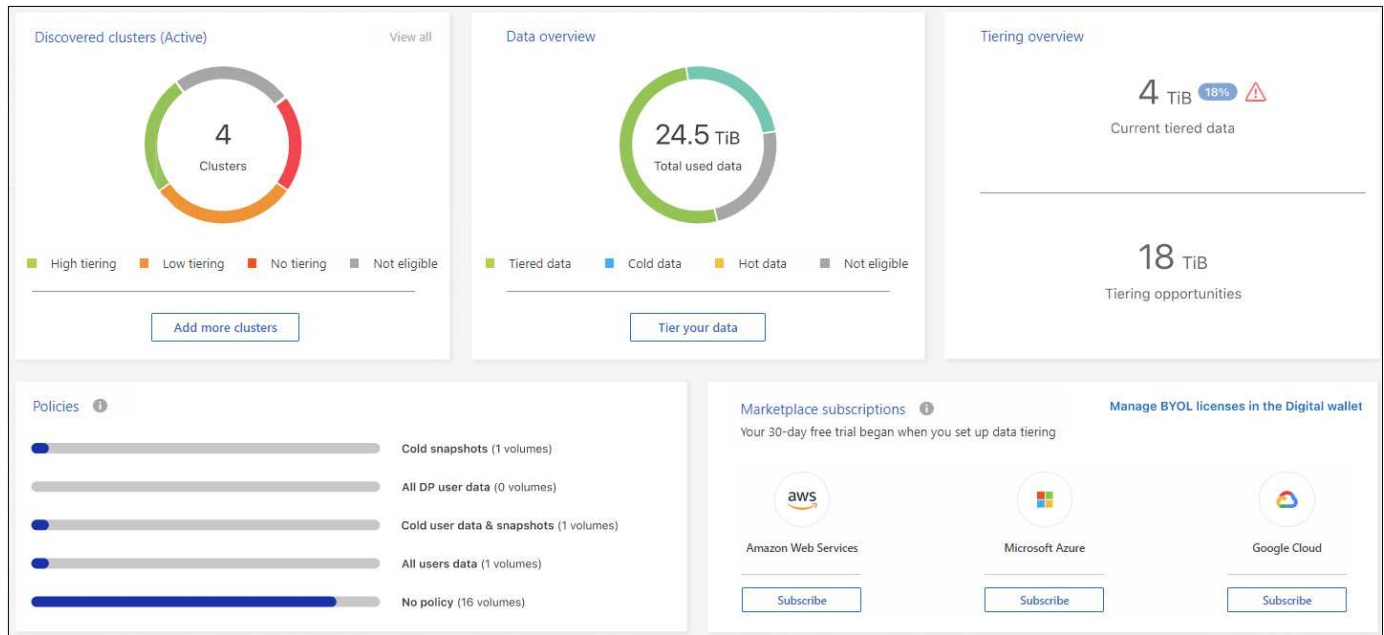
1. From the left navigation menu, select **Mobility > Tiering**.
2. From the **Clusters** page, click the menu icon for a cluster and select **Cloud Performance Test**.
3. Select a node from the drop-down list.
4. View the results or recheck the performance.

## Get an overview of data tiering from your clusters

BlueXP 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.

Just click **Tiering > On-Premises Dashboard**. BlueXP tiering provides the following details about your

environment.



## Discovered clusters

The number of on-premises clusters that BlueXP 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 indicates the amount of data that is being tiered, and more, 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

You can view the status of tiering alerts in the BlueXP 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 clicking the  in the BlueXP 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 increase your storage efficiency

This notification is a "Recommendation" to help make your systems more efficient and to save on 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 [BlueXP tiering total cost of ownership and savings calculator](#) to help you calculate your cost savings.

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

## 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.