



# Concepts

## Cloud Tiering

NetApp  
June 10, 2024

# Table of Contents

- Concepts ..... 1
  - Cloud Tiering overview ..... 1
  - How Cloud Tiering works ..... 1
  - How licensing works ..... 4
  - Savings opportunities ..... 5
  - Accounts ..... 6

# 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](#).



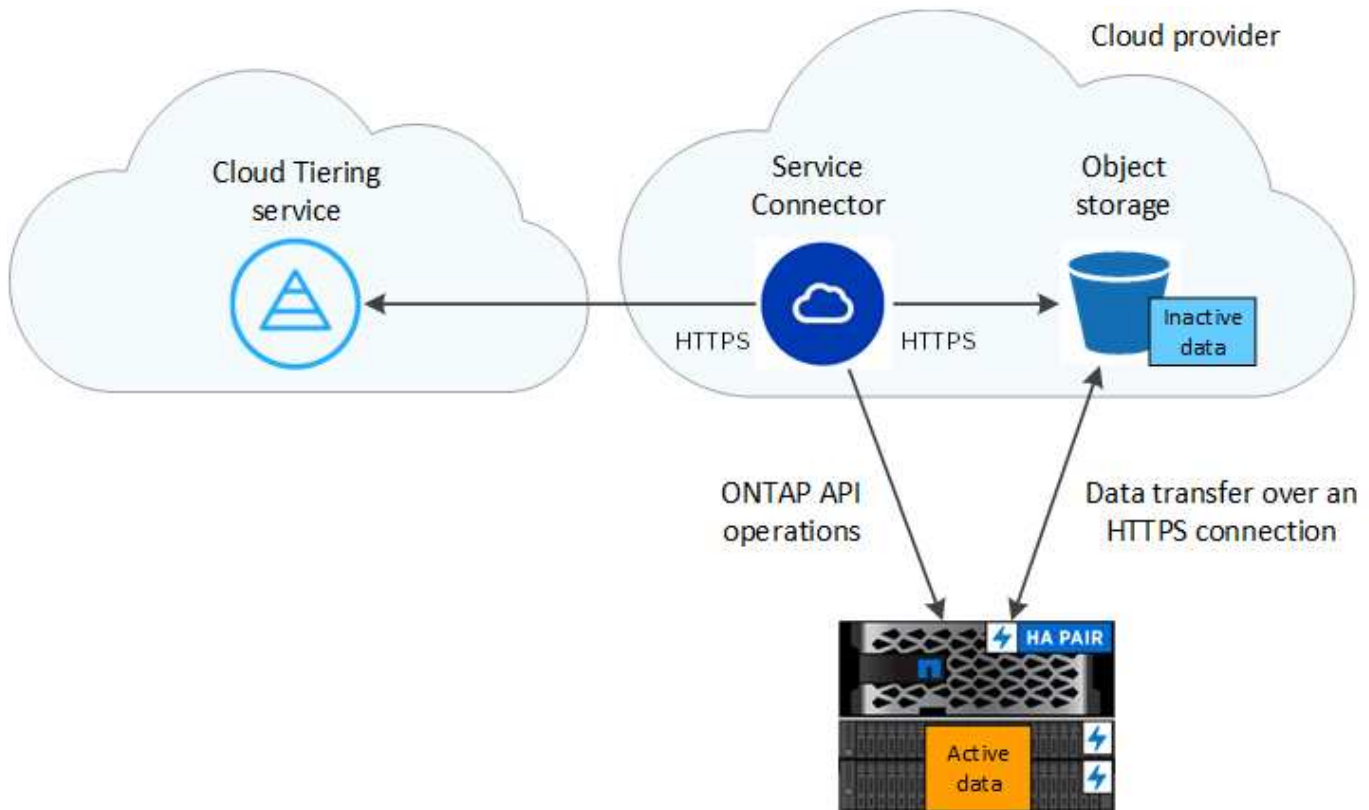
While Cloud Tiering can significantly reduce storage footprints, it is not a backup solution.

## 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, a status indicator for 'Operational Health' (Active Cluster), a 'Current Savings' section indicating that cold data has not been tiered yet, and a 'Saving Opportunities' section showing an estimated amount of 9.55 TB with a 'Set up Tiering' button. A 'Last Tiering Setup Status' message states that tiering has not been setup through Cloud Tiering.

AFF-GCP		Last Tiering Setup Status
<a href="#">More Info &gt;</a>		Tiering hasn't been setup through Cloud Tiering
<b>Operational Health</b> ✓ Active Cluster	<b>Current Savings</b> Cold data has not been tiered yet	<b>Saving Opportunities</b> Estimated Amount 9.55 TB <a href="#">Set up Tiering</a>

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

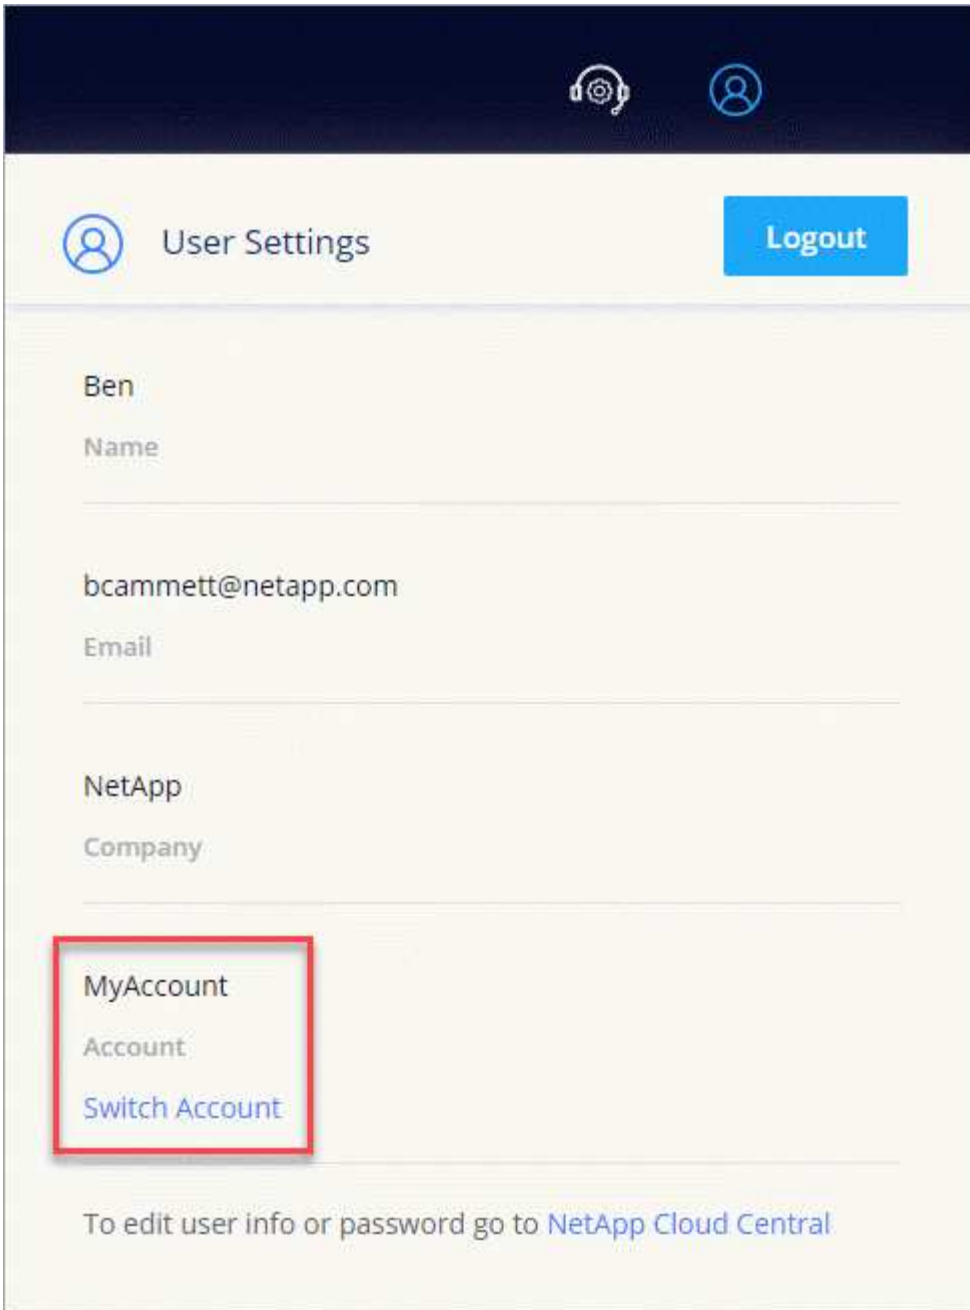
Each Cloud Central user is associated with one or more Cloud Central accounts. An account enables multi-tenancy: multiple users can manage data tiering in a single account.

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](#)

## 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.