



Tier data to StorageGRID

Cloud Tiering

NetApp
March 06, 2021

Table of Contents

- Tier data to StorageGRID 1
 - Quick start for tiering inactive data to StorageGRID 1
 - Preparing to tier inactive data to StorageGRID 1
 - Tiering inactive data from your first cluster to StorageGRID 6

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

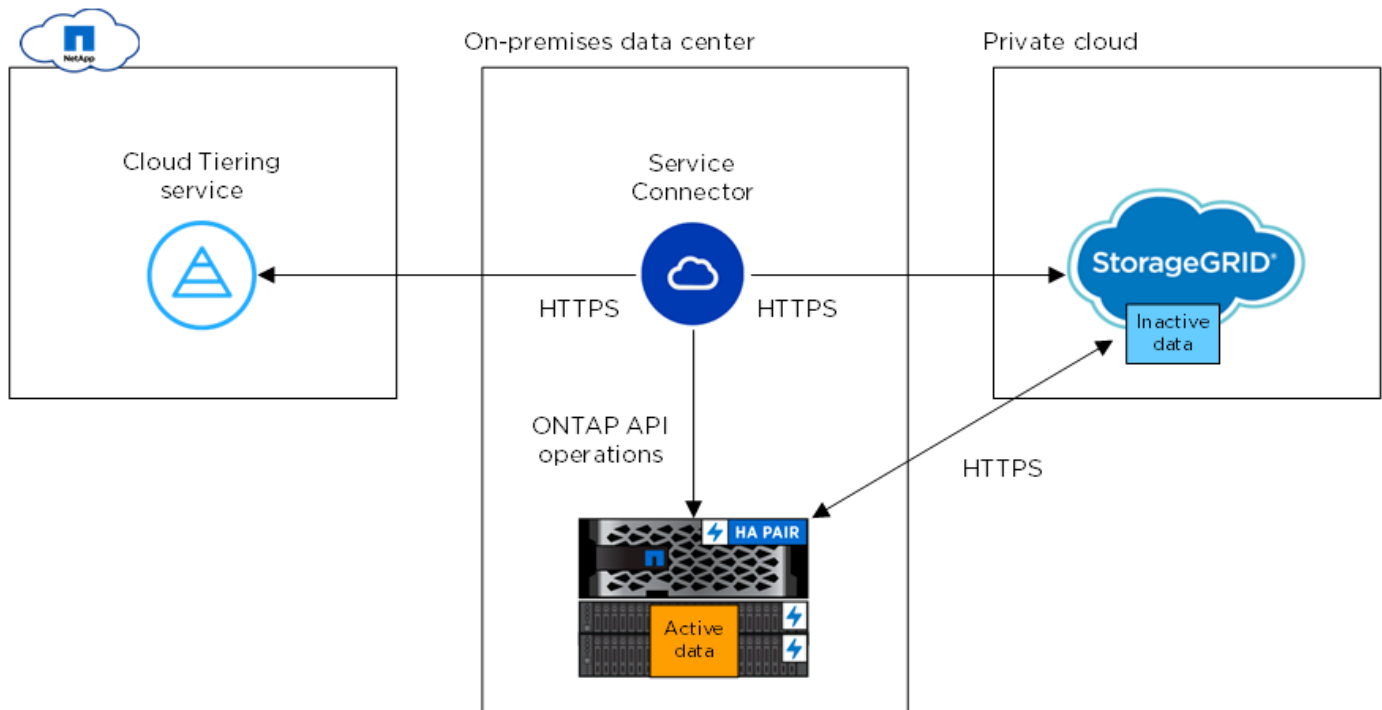
Log in to [NetApp Cloud Central](#), start a free trial of the Cloud Tiering service, and click **Let's Start, Discover Your First Cluster**.

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.



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.

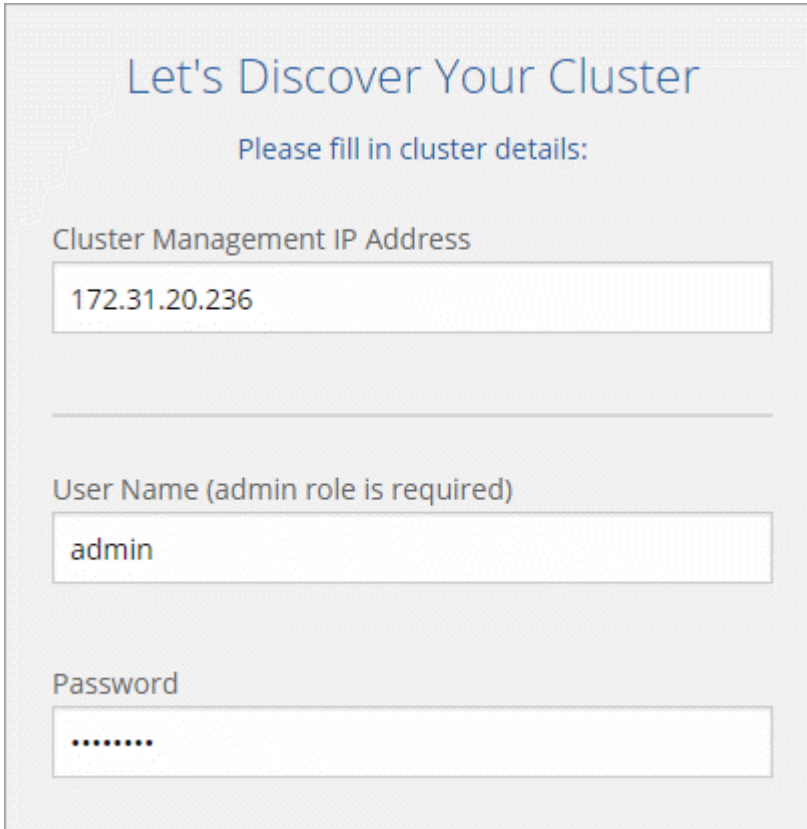
The Service Connector uses this account to send API operations to the ONTAP cluster.

- [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

172.31.20.236

User Name (admin role is required)

admin

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.](#)

Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.