



Configuration

FlexPod

NetApp
March 25, 2024

Table of Contents

- Configuration 1
 - Software revisions 1
 - Install FabricPool license 1
 - Create AWS S3 bucket 3
 - Add a cloud tier to ONTAP 5
 - Attach a cloud tier to an ONTAP aggregate 7
 - Set volume tiering policy 9
 - Set volume tiering minimum cooling days 10

Configuration

Software revisions

The following table illustrates validated hardware and software versions.

Layer	Device	Image	Comments
Storage	NetApp AFF A300	ONTAP 9.6P2	
Compute	Cisco UCS B200 M5 blade servers with Cisco UCS VIC 1340	Release 4.0(4b)	
Network	Cisco Nexus 6332-16UP fabric interconnect	Release 4.0(4b)	
	Cisco Nexus 93180YC-EX switch in NX-OS standalone mode	Release 7.0(3)I7(6)	
Storage network	Cisco MDS 9148S	Release 8.3(2)	
Hypervisor		VMware vSphere ESXi 6.7U2	ESXi 6.7.0,13006603
		VMware vCenter Server	vCenter server 6.7.0.30000 Build 13639309
Cloud provider		Amazon AWS S3	Standard S3 bucket with default options

The basic requirements for FabricPool are outlined in [FabricPool Requirements](#). After all the basic requirements have been met, complete the following steps to configure FabricPool:

1. Install a FabricPool license.
2. Create an AWS S3 object store bucket.
3. Add a cloud tier to ONTAP.
4. Attach the cloud tier to an aggregate.
5. Set the volume tiering policy.


Next: [Install FabricPool license](#).

Install FabricPool license

After you acquire a NetApp license file, you can install it with OnCommand System Manager. To install the license file, complete the following steps:

1. Click Configurations.
2. Click Cluster.
3. Click Licenses.

4. Click Add.
5. Click Choose Files to browse and select a file.
6. Click Add.



License capacity

You can view the license capacity by using either the ONTAP CLI or OnCommand System Manager. To see the licensed capacity, run the following command in the ONTAP CLI:

```
system license show-status
```

In OnCommand System Manager, complete the following steps:

1. Click Configurations.
2. Click Licenses.
3. Click the Details tab.

ONTAP System Manager

Preview the new experience

Type: All Search all Objects

Events & Jobs

Configuration

Advanced Cluster Setup

Cluster

Authentication

Configuration Updates

Expansion

Service Processor

High Availability

Licenses

Update

Licenses

Packages Details

+ Add - Delete Refresh

Package	Cluster/Node	Serial Number	Type	State	Legacy	Maximum Capacity	Current Capacity
Cluster Base License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
NFS License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
CIFS License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
iSCSI License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
FCP License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
SnapRestore License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
FlexClone License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
SnapManagerSuite L...	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
FabricPool License	cie-na300-g1325		Capacity	-NA-	No	10 TB	0 Byte

Maximum capacity and current capacity are listed on the FabricPool License row.

Next: [Create AWS S3 bucket.](#)

Create AWS S3 bucket


Buckets are object store containers that hold data. You must provide the name and location of the bucket in which data is stored before it can be added to an aggregate as a cloud tier.



Buckets cannot be created using OnCommand System Manager, OnCommand Unified Manager, or ONTAP.

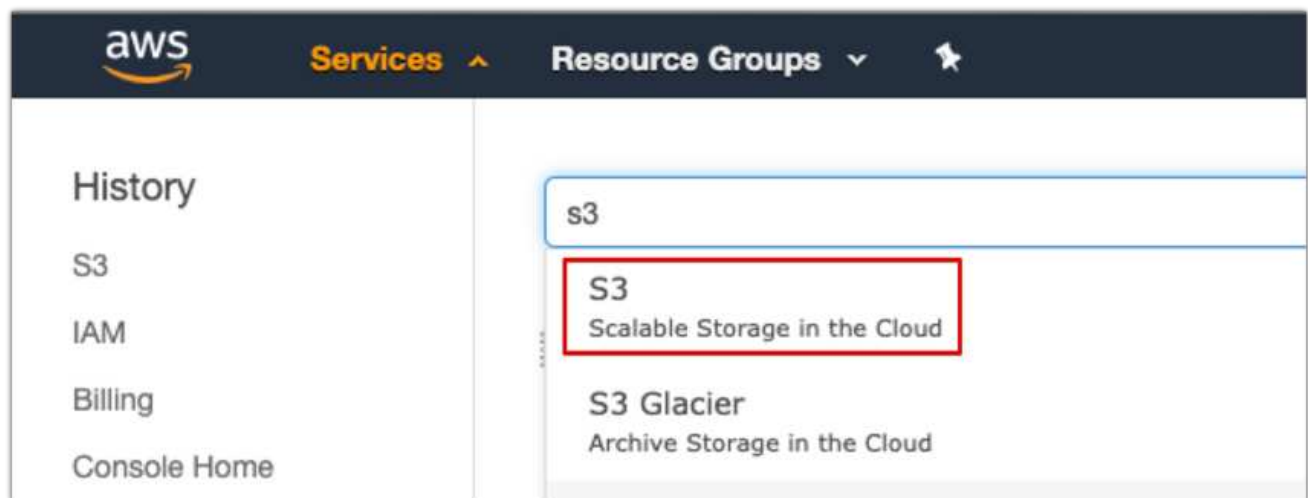
FabricPool supports the attachment of one bucket per aggregate, as illustrated in the following figure. A single bucket can be attached to a single aggregate, and a single bucket can be attached to multiple aggregates. However, a single aggregate cannot be attached to multiple buckets. Although a single bucket can be attached to multiple aggregates in a cluster, NetApp does not recommend attaching a single bucket to aggregates in multiple clusters.

When planning a storage architecture, consider how the bucket-to-aggregate relationship might affect performance. Many object store providers set a maximum number of supported IOPS at the bucket or container level. Environments that require maximum performance should use multiple buckets to reduce the possibility that object-store IOPS limitations might affect performance across multiple FabricPool aggregates. Attaching a single bucket or container to all FabricPool aggregates in a cluster might be more beneficial to environments that value manageability over cloud-tier performance.



Create an S3 bucket

1. In the AWS management console from the home page, enter S3 in the search bar.
2. Select S3 Scalable Storage in the Cloud.



3. On the S3 home page, select Create Bucket.
4. Enter a DNS-compliant name and choose the region to create the bucket.

Create bucket

1 Name and region 2 Configure options 3 Set permissions 4 Review

Name and region

Bucket name ⓘ

flexpod-fp-bk-1

Region

US East (Ohio)

Copy settings from an existing bucket

Select bucket (optional) 4 Buckets

Create Cancel Next

5. Click Create to create the object store bucket.

Next: [Add a cloud tier to ONTAP](#)

Add a cloud tier to ONTAP

Before an object store can be attached to an aggregate, it must be added to and identified by ONTAP. This task can be completed with either OnCommand System Manager or the ONTAP CLI.

FabricPool supports Amazon S3, IBM Object Cloud Storage, and Microsoft Azure Blob Storage object stores as cloud tiers.

You need the following information:

- Server name (FQDN); for example, `s3.amazonaws.com`
- Access key ID
- Secret key
- Container name (bucket name)

OnCommand System Manager

To add a cloud tier with OnCommand System Manager, complete the following steps:

1. Launch OnCommand System Manager.
2. Click Storage.
3. Click Aggregates & Disks.
4. Click Cloud Tiers.
5. Select an object store provider.
6. Complete the text fields as required for the object store provider.

In the Container Name field, enter the object store's bucket or container name.

7. Click Save and Attach Aggregates.

Add Cloud Tier



Cloud tiers/ object stores are used to store infrequently-accessed data. [Learn more](#)

Cloud Tier Provider  Amazon S3

Type


Name

Server Name (FQDN)

Access Key ID

Secret Key

 Container Name

 Encryption ☒ Enabled

ONTAP CLI

To add a cloud tier with the ONTAP CLI, enter the following commands:


```
object-store config create
-object-store-name <name>
-provider-type <AWS>
-port <443/8082> (AWS)
-server <name>
-container-name <bucket-name>
-access-key <string>
-secret-password <string>
-ssl-enabled true
-ipSPACE default
```

Next: [Attach a cloud tier to an ONTAP aggregate.](#)

Attach a cloud tier to an ONTAP aggregate

After an object store has been added to and identified by ONTAP, it must be attached to an aggregate to create a FabricPool. This task can be completed by using either OnCommand System Manager or the ONTAP CLI.

More than one type of object store can be connected to a cluster, but only one type of object store can be attached to each aggregate. For example, one aggregate can use Google Cloud, and another aggregate can use Amazon S3, but one aggregate cannot be attached to both.




Attaching a cloud tier to an aggregate is a permanent action. A cloud tier cannot be unattached from an aggregate that it has been attached to.


OnCommand System Manager

To attach a cloud tier to an aggregate by using OnCommand System Manager, complete the following steps:

1. Launch OnCommand System Manager.
2. Click Applications & Tiers.



3. Click Storage Tiers.
4. Click an aggregate.
5. Click Actions and select Attach Cloud Tier.



6. Select a cloud tier.
7. View and update the tiering policies for the volumes on the aggregate (optional). By default, the volume tiering policy is set as Snapshot-Only.
8. Click Save.

ONTAP CLI

To attach a cloud tier to an aggregate by using the ONTAP CLI, run the following commands:

```
storage aggregate object-store attach
-aggregate <name>
-object-store-name <name>
```

Example:

```
storage aggregate object-store attach -aggregate aggr1 -object-store-name
- aws_infra_fp_bk_1
```

[Next: Set volume tiering policy.](#)

Set volume tiering policy

By default, volumes use the None volume tiering policy. After volume creation, the volume tiering policy can be changed by using OnCommand System Manager or the ONTAP CLI.

When used with FlexPod, FabricPool provides three volume tiering policies, Auto, Snapshot-Only, and None.

- **Auto**

- All cold blocks in the volume are moved to the cloud tier. Assuming that the aggregate is more than 50% utilized, it takes approximately 31 days for inactive blocks to become cold. The Auto cooling period is adjustable between 2 days and 63 days by using the `tiering-minimum-cooling-days` setting.
- When cold blocks in a volume with a tiering policy set to Auto are read randomly, they are made hot and written to the performance tier.
- When cold blocks in a volume with a tiering policy set to Auto are read sequentially, they stay cold and remain on the cloud tier. They are not written to the performance tier.

- **Snapshot-Only**

- Cold snapshot blocks in the volume that are not shared with the active file system are moved to the cloud tier. Assuming that the aggregate is more than 50% utilized, it takes approximately 2 days for inactive snapshot blocks to become cold. The Snapshot-Only cooling period is adjustable from 2 to 63 days by using the `tiering-minimum-cooling-days` setting.
- When cold blocks in a volume with a tiering policy set to Snapshot-Only are read, they are made hot and written to the performance tier.

- **None (Default)**

- Volumes set to use None as their tiering policy do not tier cold data to the cloud tier.
- Setting the tiering policy to None prevents new tiering.
- Volume data that has previously been moved to the cloud tier remains in the cloud tier until it becomes hot and is automatically moved back to the performance tier.

OnCommand System Manager

To change a volume's tiering policy by using OnCommand System Manager, complete the following steps:

1. Launch OnCommand System Manager.
2. Select a volume.
3. Click More Actions and select Change Tiering Policy.
4. Select the tiering policy to apply to the volume.
5. Click Save.

CHANGE VOLUME TIERING POLICY

Select the tiering policy that you want to apply for the selected volume.

Volume Name	Tiering Policy
affa3..._fp_1	auto

Tiering Policy auto

- snapshot-only
- none
- auto
- all

[er and tiering policies.](#)

Save Cancel

ONTAP CLI

To change a volume's tiering policy by using the ONTAP CLI, run the following command:

```
volume modify -vserver <svm_name> -volume <volume_name>  
-tiering-policy <auto|snapshot-only|all|none>
```

Next: [Set volume tiering minimum cooling days.](#)

Set volume tiering minimum cooling days

The `tiering-minimum-cooling-days` setting determines how many days must pass before inactive data in a volume using the Auto or Snapshot-Only policy is considered cold and eligible for tiering.

Auto

The default `tiering-minimum-cooling-days` setting for the Auto tiering policy is 31 days.

Because reads keep block temperatures hot, increasing this value might reduce the amount of data that is eligible to be tiered and increase the amount of data kept on the performance tier.

If you would like to reduce this value from the default 31 days, be aware that data should no longer be active

before being marked as cold. For example, if a multiday workload is expected to perform a significant number of writes on day 7, the volume's `tiering-minimum-cooling-days` setting should be set no lower than 8 days.



Object storage is not transactional like file or block storage. Making changes to files that are stored as objects in volumes with overly aggressive minimum cooling days can result in the creation of new objects, the fragmentation of existing objects, and the addition of storage inefficiencies.

Snapshot-Only

The default `tiering-minimum-cooling-days` setting for the Snapshot-Only tiering policy is 2 days. A 2-day minimum gives additional time for background processes to provide maximum storage efficiency and prevents daily data-protection processes from having to read data from the cloud tier.

ONTAP CLI

To change a volume's `tiering-minimum-cooling-days` setting by using the ONTAP CLI, run the following command:

```
volume modify -vserver <svm_name> -volume <volume_name> -tiering-minimum  
-cooling-days <2-63>
```

The advanced privilege level is required.



Changing the tiering policy between Auto and Snapshot-Only (or vice versa) resets the inactivity period of blocks on the performance tier. For example, a volume using the Auto volume tiering policy with data on the performance tier that has been inactive for 20 days will have the performance tier data inactivity reset to 0 days if the tiering policy is set to Snapshot-Only.

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.