



FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3

FlexPod

NetApp
January 21, 2025

Table of Contents

- FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3 1
 - TR-4801: FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3 1
 - FlexPod overview and architecture 1
 - FabricPool 3
 - FabricPool requirements 7
 - Configuration 11
 - Performance considerations 21
 - Cost of ownership 22
 - Conclusion 22
 - Where to find additional information. 22

FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3

TR-4801: FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3

Scott Kovacs, NetApp

Flash storage prices continue to fall, making it available to workloads and applications that were not previously considered candidates for flash storage. However, making the most efficient use of the storage investment is still critically important for IT managers. IT departments continue to be pressed to deliver higher-performing services with little or no budget increase. To help address these needs, NetApp FabricPool allows you to leverage cloud economics by moving infrequently used data off of expensive on-premises flash storage to a more cost-effective storage tier in the public cloud. Moving infrequently accessed data to the cloud frees up valuable flash storage space on AFF or FAS systems to deliver more capacity for business-critical workloads to the high-performance flash tier.

This technical report reviews the FabricPool data- tiering feature of NetApp ONTAP in the context of a FlexPod converged infrastructure architecture from NetApp and Cisco. You should be familiar with the FlexPod Datacenter converged infrastructure architecture and the ONTAP storage software to fully benefit from the concepts discussed in this technical report. Building on familiarity with FlexPod and ONTAP, we discuss FabricPool, how it works, and how it can be used to achieve more efficient use of on-premises flash storage. Much of the content in this report is covered in greater detail in [TR-4598 FabricPool Best Practices](#) and other ONTAP product documentation. The content has been condensed for a FlexPod infrastructure and does not completely cover all use cases for FabricPool. All features and concepts examined are available in ONTAP 9.6.

Additional information about FlexPod is available in [TR-4036 FlexPod Datacenter Technical Specifications](#).


FlexPod overview and architecture

FlexPod overview

FlexPod is a defined set of hardware and software that forms an integrated foundation for both virtualized and nonvirtualized solutions. FlexPod includes NetApp AFF storage, Cisco Nexus networking, Cisco MDS storage networking, the Cisco Unified Computing System (Cisco UCS), and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit into one data center rack, or it can be deployed according to a customer's data center design. Port density allows the networking components to accommodate multiple configurations.

One benefit of the FlexPod architecture is the ability to customize, or flex, the environment to suit a customer's requirements. A FlexPod unit can easily be scaled as requirements and demand change. A unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The FlexPod reference architecture highlights the resiliency, cost benefit, and ease of deployment of a Fibre Channel and IP-based storage solution. A storage system that is capable of serving multiple protocols across a single interface gives customers a choice and protects their investment because it is truly a wire-once architecture. The following figure shows many of the hardware components of FlexPod.

FlexPod Datacenter solution



FlexPod architecture

The following figure shows the components of a VMware vSphere and FlexPod solution and the network connections needed for Cisco UCS 6454 fabric interconnects. This design has the following components:

- Port-channelled 40Gb Ethernet connections between the Cisco UCS 5108 blade chassis and the Cisco UCS fabric interconnects
- 40Gb Ethernet connections between the Cisco UCS fabric interconnect and the Cisco Nexus 9000
- 40Gb Ethernet connections between the Cisco Nexus 9000 and the NetApp AFF A300 storage array


These infrastructure options expanded with the introduction of Cisco MDS switches sitting between the Cisco UCS fabric interconnect and the NetApp AFF A300. This configuration provides FC-booted hosts with 16Gb FC block-level access to shared storage. The reference architecture reinforces the wire-once strategy, because, as additional storage is added to the architecture, no recabling is required from the hosts to the Cisco UCS fabric interconnect.

Cisco Unified Computing System
 Cisco UCS 6332-16UP
 Fabric Interconnects,
 UCS B-Series Blade Servers
 with UCS VIC 1340 and UCS
 2304 Fabric Extender

Cisco Nexus 93180YC-EX

NetApp storage controllers AFF-A300

Cisco MDS 9148S



FabricPool

FabricPool overview

FabricPool is a hybrid storage solution in ONTAP that uses an all-flash (SSD) aggregate as a performance tier and an object store in a public cloud service as a cloud tier. This configuration enables policy-based data movement, depending on whether or not data is frequently accessed. FabricPool is supported in ONTAP for both AFF and all-SSD aggregates on FAS platforms. Data processing is performed at the block level, with frequently accessed data blocks in the all-flash performance tier tagged as hot and infrequently accessed blocks tagged as cold.

Using FabricPool helps to reduce storage costs without compromising performance, efficiency, security, or protection. FabricPool is transparent to enterprise applications and capitalizes on cloud efficiencies by lowering storage TCO without having to rearchitect the application infrastructure.

FlexPod can benefit from the storage tiering capabilities of FabricPool to make more efficient use of ONTAP flash storage. Inactive virtual machines (VMs), infrequently used VM templates, and VM backups from NetApp SnapCenter for vSphere can consume valuable space in the datastore volume. Moving cold data to the cloud tier frees space and resources for high-performance, mission-critical applications hosted on the FlexPod infrastructure.



Fibre Channel and iSCSI protocols generally take longer before experiencing a timeout (60 to 120 seconds), but they do not retry to establish a connection in the same way that NAS protocols do. If a SAN protocol times out, the application must be restarted. Even a short disruption could be disastrous to production applications using SAN protocols because there is no way to guarantee connectivity to public clouds. To avoid this issue, NetApp recommends using private clouds when tiering data that is accessed by SAN protocols.

In ONTAP 9.6, FabricPool integrates with all the major public cloud providers: Alibaba Cloud Object Storage Service, Amazon AWS S3, Google Cloud Storage, IBM Cloud Object Storage, and Microsoft Azure Blob Storage. This report focuses on Amazon AWS S3 storage as the cloud object tier of choice.

The composite aggregate

A FabricPool instance is created by associating an ONTAP flash aggregate with a cloud object store, such as an AWS S3 bucket, to create a composite aggregate. When volumes are created inside the composite aggregate, they can take advantage of the tiering capabilities of FabricPool. When data is written to the volume, ONTAP assigns a temperature to each of the data blocks. When the block is first written, it is assigned a temperature of hot. As time passes, if the data is not accessed, it undergoes a cooling process until it is finally assigned a cold status. These infrequently accessed data blocks are then tiered off the performance SSD aggregate and into the cloud object store.

The period of time between when a block is designated as cold and when it is moved to cloud object storage is modified by the volume tiering policy in ONTAP. Further granularity is achieved by modifying ONTAP settings that control the number of days required for a block to become cold. Candidates for data tiering are traditional volume snapshots, SnapCenter for vSphere VM backups and other NetApp Snapshot- based backups, and any infrequently used blocks in a vSphere datastore, such as VM templates and infrequently accessed VM data.

Inactive data reporting

Inactive data reporting (IDR) is available in ONTAP to help evaluate the amount of cold data that can be tiered from an aggregate. IDR is enabled by default in ONTAP 9.6 and uses a default 31-day cooling policy to determine which data in the volume is inactive.



The amount of cold data that is tiered depends on the tiering policies set on the volume. This amount may be different than the amount of cold data detected by IDR using the default 31-day cooling period.

Object creation and data movement

FabricPool works at the NetApp WAFL block level, cooling blocks, concatenating them into storage objects, and migrating those objects to a cloud tier. Each FabricPool object is 4MB and is composed of 1,024 4KB blocks. The object size is fixed at 4MB based on performance recommendations from leading cloud providers and cannot be changed. If cold blocks are read and made hot, only the requested blocks in the 4MB object are fetched and moved back to the performance tier. Neither the entire object nor the entire file is migrated back. Only the necessary blocks are migrated.



If ONTAP detects an opportunity for sequential readaheads, it requests blocks from the cloud tier before they are read to improve performance.

By default, data is moved to the cloud tier only when the performance aggregate is greater than 50% utilized. This threshold can be set to a lower percentage to allow a smaller amount of data storage on the performance

flash tier to be moved to the cloud. This might be useful if the tiering strategy is to move cold data only when the aggregate is nearing capacity.

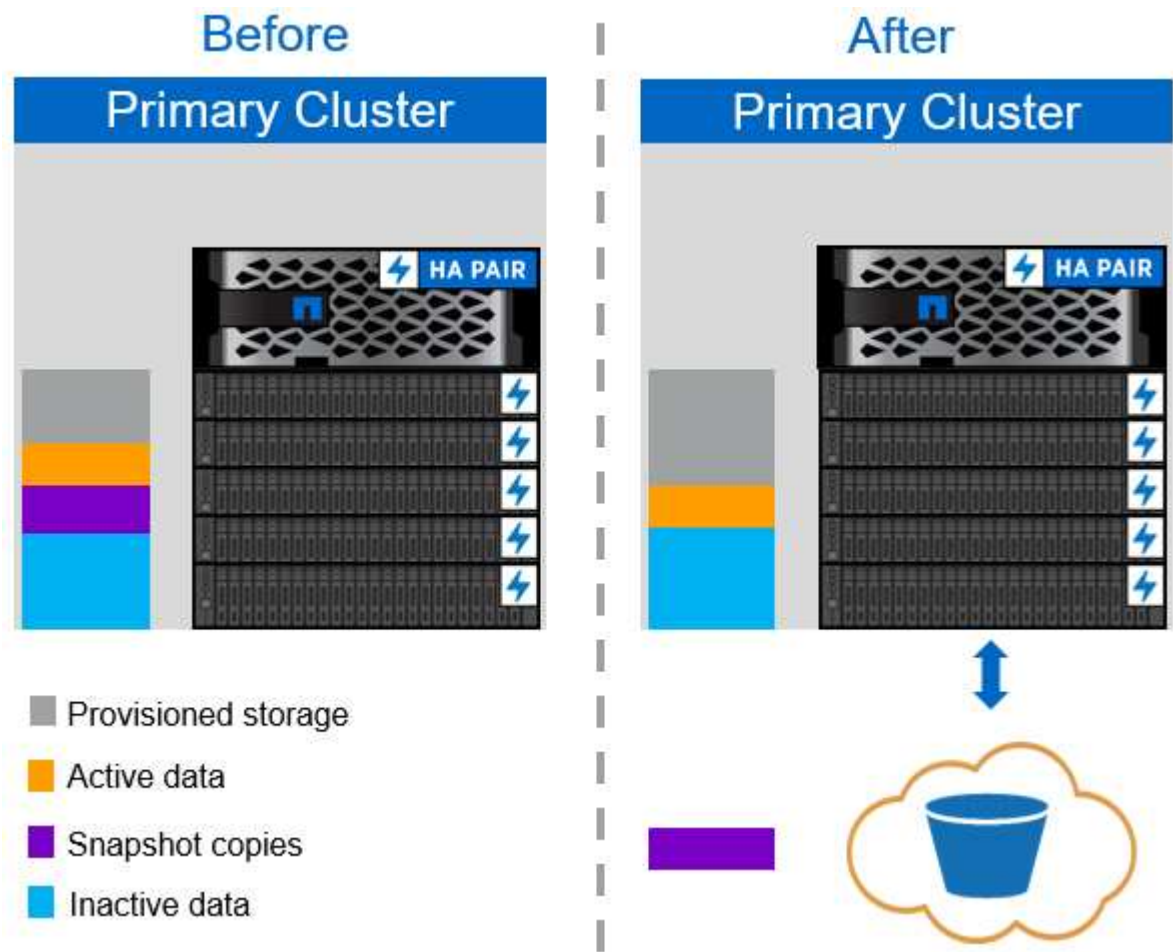
If performance tier utilization is at greater than 70% capacity, cold data is read directly from the cloud tier without being written back to the performance tier. By preventing cold data write-backs on heavily used aggregates, FabricPool preserves the aggregate for active data.

Reclaim performance tier space

As previously discussed, the primary use case for FabricPool is to facilitate the most efficient use of high-performance on-premises flash storage. Cold data in the form of volume snapshots and VM backups of the FlexPod virtual infrastructure can occupy a significant amount of expensive flash storage. Valuable performance- tier storage can be freed by implementing one of two tiering policies: Snapshot-Only or Auto.

Snapshot-Only tiering policy

The Snapshot-Only tiering policy, illustrated in the following figure, moves cold volume snapshot data and SnapCenter for vSphere backups of VMs that are occupying space but are not sharing blocks with the active file system into a cloud object store. The Snapshot-Only tiering policy moves cold data blocks to the cloud tier. If a restore is required, cold blocks in the cloud are made hot and moved back to the performance flash tier on the premises.



Auto tiering policy


The FabricPool Auto tiering policy, illustrated in the following figure, not only moves cold snapshot data blocks to the cloud, it also moves any cold blocks in the active file system. This can include VM templates and any

unused VM data in the datastore volume. Which cold blocks are moved is controlled by the `tiering-minimum-cooling-days` setting for the volume. If cold blocks in the cloud tier are randomly read by an application, those blocks are made hot and brought back to the performance tier. However, if cold blocks are read by a sequential process such as an antivirus scanner, the blocks remain cold and persist in the cloud object store; they are not moved back to the performance tier.

When using the Auto tiering policy, infrequently accessed blocks that are made hot are pulled back from the cloud tier at the speed of cloud connectivity. This may affect VM performance if the application is latency sensitive, which should be considered before using the Auto tiering policy on the datastore. NetApp recommends placing Intercluster LIFs on ports with a speed of 10GbE for adequate performance.



The object store profiler should be used to test latency and throughput to the object store before attaching it to a FabricPool aggregate.



All tiering policy

Unlike the Auto and Snapshot-only policies, the All tiering policy moves entire volumes of data immediately into the cloud tier. This policy is best suited to secondary data protection or archival volumes for which data must be kept for historical or regulatory purposes but is rarely accessed. The All policy is not recommended for VMware datastore volumes because any data written to the datastore is immediately moved to the cloud tier. Subsequent read operations are performed from the cloud and could potentially introduce performance issues for VMs and applications residing in the datastore volume.

Security

Security is a central concern for the cloud and for FabricPool. All the native security features of ONTAP are supported in the performance tier, and the movement of data is secured as it is transferred to the cloud tier. FabricPool uses the [AES-256-GCM](#) encryption algorithm on the performance tier and maintains this encryption end to end into the cloud tier. Data blocks that are moved to the cloud object store are secured with transport layer security (TLS) v1.2 to maintain data confidentiality and integrity between storage tiers.



Communicating with the cloud object store over an unencrypted connection is supported but not recommended by NetApp.

Data encryption

Data encryption is vital to the protection of intellectual property, trade information, and personally identifiable customer information. FabricPool fully supports both NetApp Volume Encryption (NVE) and NetApp Storage Encryption (NSE) to maintain existing data protection strategies. All encrypted data on the performance tier remains encrypted when moved to the cloud tier. Client-side encryption keys are owned by ONTAP and the server-side object store encryption keys are owned by the respective cloud object store. Any data not encrypted with NVE is encrypted with the AES-256-GCM algorithm. No other AES-256 ciphers are supported.



The use of NSE or NVE is optional and not required to use FabricPool.

FabricPool requirements

FabricPool requires ONTAP 9.2 or later and the use of SSD aggregates on any of the platforms listed in this section. Additional FabricPool requirements depend on the cloud tier being attached. For entry-level AFF platforms that have a fixed, relatively small capacity such as the NetApp AFF C190, FabricPool can be highly effective for moving inactive data to the cloud tier.

Platforms

FabricPool is supported on the following platforms:

- NetApp AFF
 - A800
 - A700S, A700
 - A320, A300
 - A220, A200
 - C190
 - AFF8080, AFF8060, and AFF8040
- NetApp FAS
 - FAS9000
 - FAS8200
 - FAS8080, FAS8060, and FAS8040
 - FAS2750, FAS2720

- FAS2650, FAS2620



Only SSD aggregates on FAS platforms can use FabricPool.

- Cloud tiers
 - Alibaba Cloud Object Storage Service (Standard, Infrequent Access)
 - Amazon S3 (Standard, Standard-IA, One Zone-IA, Intelligent-Tiering)
 - Amazon Commercial Cloud Services (C2S)
 - Google Cloud Storage (Multi-Regional, Regional, Nearline, Coldline)
 - IBM Cloud Object Storage (Standard, Vault, Cold Vault, Flex)
 - Microsoft Azure Blob Storage (Hot and Cool)

Intercluster LIFs

Cluster high-availability (HA) pairs that use FabricPool require two intercluster logical interfaces (LIFs) to communicate with the cloud tier. NetApp recommends creating an intercluster LIF on additional HA pairs to seamlessly attach cloud tiers to aggregates on those nodes as well.

The LIF that ONTAP uses to connect with the AWS S3 object store must be on a 10Gbps port.

If more than one Intercluster LIF is used on a node with different routing, NetApp recommends placing them in different IPspaces. During configuration, FabricPool can select from multiple IPspaces, but it is not able to select specific intercluster LIFs within an IPspace.




Disabling or deleting an intercluster LIF interrupts communication to the cloud tier.

Connectivity

FabricPool read latency is a function of connectivity to the cloud tier. Intercluster LIFs using 10Gbps ports, illustrated in the following figure, provide adequate performance. NetApp recommends validating the latency and throughput of the specific network environment to determine the effect it has on FabricPool performance.



When using FabricPool in low-performance environments, minimum performance requirements for client applications must continue to be met, and recovery time objectives should be adjusted accordingly.



Object store profiler

The object store profiler, an example of which is shown below and is available through the ONTAP CLI, tests the latency and throughput performance of object stores before they are attached to a FabricPool aggregate.



The cloud tier must be added to ONTAP before it can be used with the object store profiler.

Start the object store profiler from the advanced privilege mode in ONTAP with the following command:

```
storage aggregate object-store profiler start -object-store-name <name>
-node <name>
```

To view the results, run the following command:

```
storage aggregate object-store profiler show
```

Cloud tiers do not provide performance similar to that found on the performance tier (typically GB per second). Although FabricPool aggregates can easily provide SATA-like performance, they can also tolerate latencies as high as 10 seconds and low throughput for tiering solutions that do not require SATA-like performance.

```
bb09-a300-2::*> storage aggregate object-store profiler show
Object store config name: aws_infra_fp_bk_1
Node name: bb09-a300-2-1
Status: Active. Issuing GETs
Start time: 10/3/2019 12:37:24
```

Op	Size	Total	Failed	min	Latency (ms) max	avg	Throughput
PUT	4MB	1084	0	336	5951	2817	69.55MB
GET	4KB	158636	0	27	1132	41	32.22MB
GET	8KB	0	0	0	0	0	0B
GET	32KB	0	0	0	0	0	0B
GET	256KB	0	0	0	0	0	0B

5 entries were displayed.

Volumes

Storage thin provisioning is a standard practice for the FlexPod virtual infrastructure administrator. NetApp Virtual Storage Console (VSC) provisions storage volumes for VMware datastores without any space guarantee (thin provisioning) and with optimized storage efficiency settings per NetApp best practices. If VSC is used to create VMware datastores, no additional action is required, because no space guarantee should be assigned to the datastore volume.



FabricPool cannot attach a cloud tier to an aggregate that contains volumes using a space guarantee other than None (for example, Volume).

```
volume modify -space-guarantee none
```

Setting the `space-guarantee none` parameter provides thin provisioning for the volume. The amount of space consumed by volumes with this guarantee type grows as data is added instead of being determined by the initial volume size. This approach is essential for FabricPool because the volume must support cloud tier data that becomes hot and is brought back to the performance tier.

Licensing

FabricPool requires a capacity-based license when attaching third-party object storage providers (such as Amazon S3) as cloud tiers for AFF and FAS hybrid flash systems.

FabricPool licenses are available in perpetual or term-based (1-year or 3-year) format.

Tiering to the cloud tier stops when the amount of data (used capacity) stored on the cloud tier reaches the licensed capacity. Additional data, including SnapMirror copies to volumes using the All tiering policy, cannot be tiered until the license capacity is increased. Although tiering stops, data is still accessible from the cloud tier. Additional cold data remains on SSDs until the licensed capacity is increased.

A free 10TB capacity, term-based FabricPool license comes with the purchase of any new ONTAP 9.5 or later cluster, although additional support costs might apply. FabricPool licenses (including additional capacity for existing licenses) can be purchased in 1TB increments.

A FabricPool license can only be deleted from a cluster that contains no FabricPool aggregates.



FabricPool licenses are cluster-wide. You should have the UUID available when purchasing a license (`cluster identify show`). For additional licensing information, refer to the [NetApp Knowledgebase](#).

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.

The screenshot displays the OnCommand System Manager web interface. The left-hand navigation pane shows the 'Configuration' menu item selected, with a sub-menu 'Licenses' also highlighted. The main content area is titled 'Licenses' and contains a table of installed licenses. A modal dialog titled 'Add License Packages' is open in the foreground, featuring a text input field for 'Enter comma separated license keys', a 'Choose Files' button, and an 'Add' button at the bottom right. The background table lists various license packages such as '(DEPRECATED)-Cluster Base License', 'Trusted Platform Module License', 'FabricPool License', and 'NFS License'.

Package	Entitlement Risk	Description
(DEPRECATED)-Cluster Base License	-NA-	Installed on a cluster
Trusted Platform Module License	-NA-	No License Available
FabricPool License	-NA-	Installed on a cluster
NFS License	Medium risk	Medium risk
CIFS License		
ISCSI License		
FCP License		
SnapRestore License		
SnapMirror License		
FlexClone License		
SnapVault License		
SnapLock License		


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.



The screenshot shows the ONTAP System Manager interface. In the left sidebar, the 'Configuration' menu item is highlighted with a red box. The main panel displays the 'Licenses' section, with the 'Details' tab selected. A table lists various licenses, and the 'FabricPool License' row is highlighted with a red box. The table columns are: Package, Cluster/Node, Serial Number, Type, State, Legacy, Maximum Capacity, and Current Capacity.

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

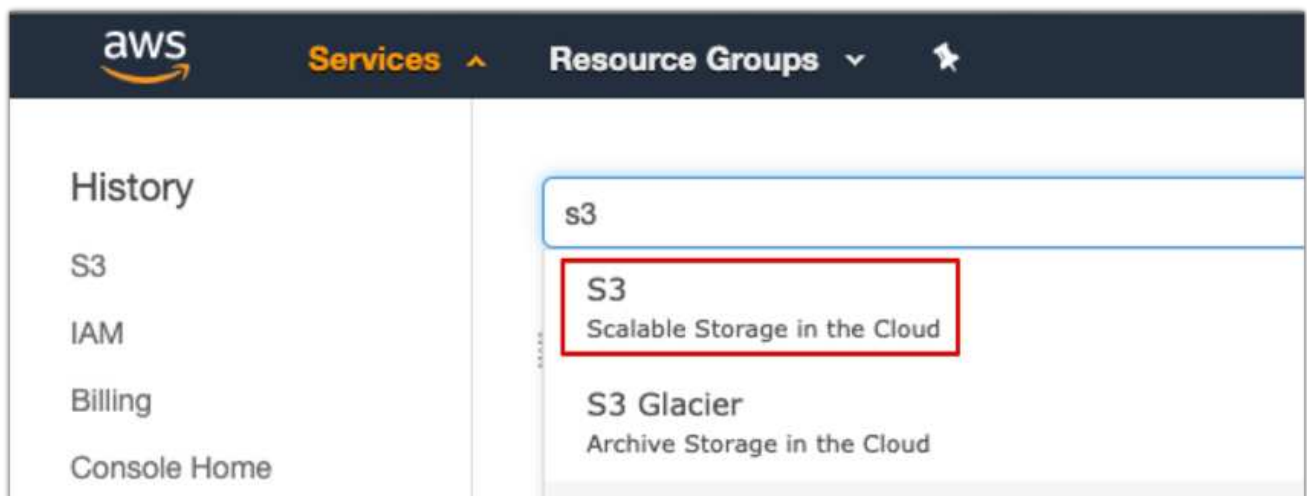
FabricPool supports the attachment of one bucket per aggregate, as illustrated in the following figure. A single bucket can be attached to a single aggregate, and a single bucket can be attached to multiple aggregates. However, a single aggregate cannot be attached to multiple buckets. Although a single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



The screenshot shows the 'Create bucket' dialog in the AWS S3 console. The dialog has a blue header with the title 'Create bucket' and a close button (X). Below the header is a progress bar with four steps: 1. Name and region, 2. Configure options, 3. Set permissions, and 4. Review. The first step, 'Name and region', is currently active. It contains three input fields: 'Bucket name' with the value 'flexpod-fp-bk-1', 'Region' with the value 'US East (Ohio)', and 'Copy settings from an existing bucket' with a dropdown menu showing 'Select bucket (optional)' and '4 Buckets'. At the bottom of the dialog, there are three buttons: 'Create' (highlighted with a red rectangle), 'Cancel', and '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

[Learn more about 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.

Performance considerations

Size the performance tier

When considering sizing, keep in mind that the performance tier should be capable of the following tasks:

- Supporting hot data
- Supporting cold data until the tiering scan moves the data to the cloud tier
- Supporting cloud tier data that becomes hot and is written back to the performance tier
- Supporting WAFL metadata associated with the attached cloud tier

For most environments, a 1:10 performance-to-capacity ratio on FabricPool aggregates is extremely conservative, while providing significant storage savings. For example, if the intent is to tier 200TB to the cloud tier, then the performance tier aggregate should be 20TB at a minimum.



Writes from the cloud tier to the performance tier are disabled if performance tier capacity is greater than 70%. If this occurs, blocks are read directly from the cloud tier.

Size the cloud tier

When considering sizing, the object store acting as the cloud tier should be capable of the following tasks:

- Supporting reads of existing cold data
- Supporting writes of new cold data
- Supporting object deletion and defragmentation

Cost of ownership

The [FabricPool Economic Calculator](#) is available through the independent IT analyst firm Evaluator Group to help project the cost savings between on premises and the cloud for cold data storage. The calculator provides a simple interface to determine the cost of storing infrequently accessed data on a performance tier versus sending it to a cloud tier for the remainder of the data lifecycle. Based on a 5-year calculation, the four key factors—source capacity, data growth, snapshot capacity, and the percentage of cold data—are used to determine storage costs over the time period.

Conclusion

The journey to the cloud varies between organizations, between business units, and even between business units within organizations. Some choose a fast adoption, while others take a more conservative approach. FabricPool fits into the cloud strategy of organizations no matter their size and regardless of their cloud adoption speed, further demonstrating the efficiency and scalability benefits of a FlexPod infrastructure.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- FabricPool Best Practices

www.netapp.com/us/media/tr-4598.pdf

- NetApp Product Documentation

<https://docs.netapp.com>

- TR-4036: FlexPod Datacenter Technical Specification

<https://www.netapp.com/us/media/tr-4036.pdf>

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—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.