Fsas Technologies Inc.

P3AG-5662-04ENZ0

ETERNUS AX series All-Flash Arrays, ETERNUS AC series All-Flash Arrays, ETERNUS HX series Hybrid Arrays

FabricPool Best Practices ONTAP 9.14.1

Table of Contents

List of Figures

List of Tables

Preface

This document describes best practices for the ONTAP software component, FabricPool. The capabilities, requirements, implementation, and best practices for this software are covered in this document.

Fourth Edition
March 2025

Trademarks

Third-party trademark information related to this product is available at:
https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html
Trademark symbols such as ™ and ® are omitted in this document.

About This Manual

Intended Audience

This manual is intended for system administrators who configure and manage operations of the ETERNUS AX/AC/HX, or field engineers who perform maintenance. Refer to this manual as required.

Related Information and Documents

The latest information for the ETERNUS AX/AC/HX is available at:
https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html

Document Conventions

Notice Symbols

The following notice symbols are used in this manual:

Caution	Indicates information that you need to observe when using the ETERNUS AX/AC/HX. Make sure to read the information.
Note	Indicates information and suggestions that supplement the descriptions included in this manual.

1. Overview

FabricPool is a Data Fabric technology that enables automated tiering of data to low-cost object storage tiers either on or off premises.

Unlike manual tiering solutions, FabricPool reduces the total cost of ownership by automating the tiering of data to lower the cost of storage. Cloud Volumes ONTAP delivers the economic benefits of the cloud by supporting data tiring with Amazon S3, Google Cloud Storage, and Microsoft Azure Blob Storage.

FabricPool is transparent to applications and allows enterprises to take advantage of cloud economics without sacrificing performance or having to rearchitect solutions to leverage storage efficiency.

- ETERNUS AX/AC/HX supports FabricPool on SSD and HDD local tiers (also known as storage aggregates in the ONTAP CLI). FlashPool aggregates are not supported.
- Cloud Volumes ONTAP supports data tiering with Amazon S3, Google Cloud Storage, and Microsoft Azure Blob Storage.

Figure 1 Before and after FabricPool

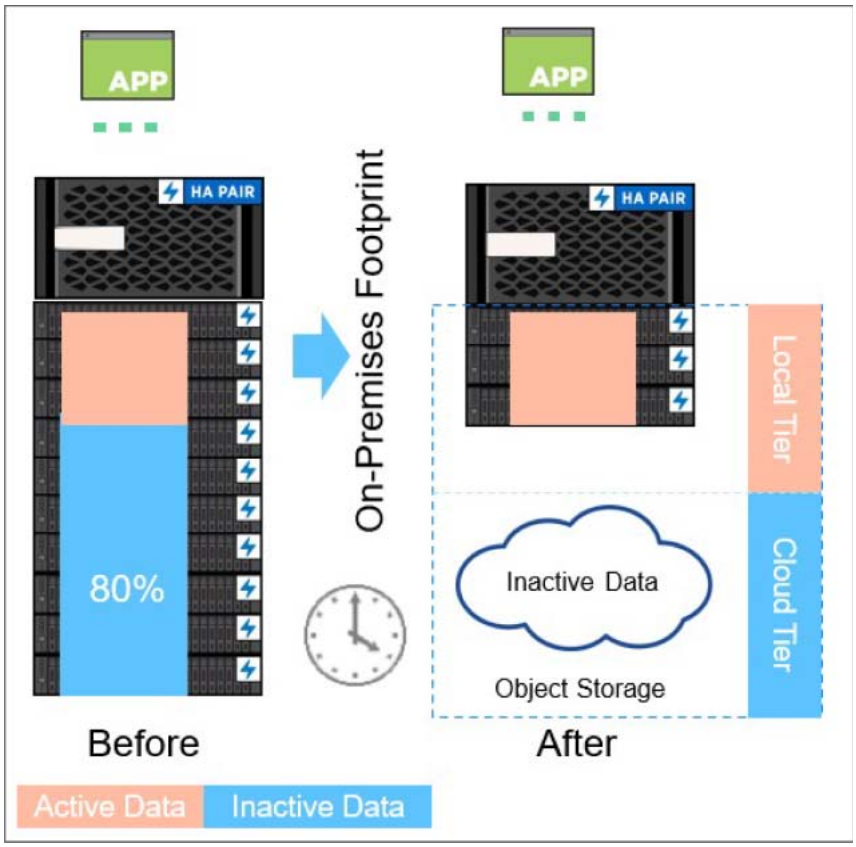

2. Primary Use Cases

The primary purpose of FabricPool is to reduce storage footprints and associated costs. Active data remains on high-performance local tiers, and inactive data is tiered to low-cost object storage while preserving ONTAP functionality and data efficiencies.

FabricPool has two primary use cases:

- Reclaim capacity on primary storage
- Shrink the secondary storage footprint

Although FabricPool can significantly reduce storage footprints in primary and secondary data centers, it is not a backup solution. Access control lists (ACLs), directory structures, and WAFL metadata always stay on the local tier. If a catastrophic disaster destroys the local tier, a new environment cannot be created using the data on the cloud tier because it contains no WAFL metadata.

For complete data protection, consider using existing ONTAP technologies such as SnapMirror. For details, refer to ETERNUS AX/AC/HX series SnapMirror Configuration and Best Practices Guide in the manual site.

Reclaim Capacity on Primary Storage (Auto, Snapshot-Only, or All)

Auto

The majority of inactive (cold) data in storage environments is associated with unstructured data, accounting for more than 50% of total storage capacity in many storage environments.

Infrequently accessed data associated with productivity software, completed projects, and old datasets is an inefficient use of high-performance storage capacity, and tiering this data to a low-cost object store is an easy way to reclaim existing local capacity and reduce the amount of required local capacity moving forward.

The auto volume tiering policy shown in Figure 2 moves all cold blocks in the volume, not just blocks associated with Snapshot copies, to the cloud tier.

If read by random reads, cold data blocks on the cloud tier become hot and are moved to the local 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 local tier.

Figure 2 Reclaiming space with the Auto volume tiering policy

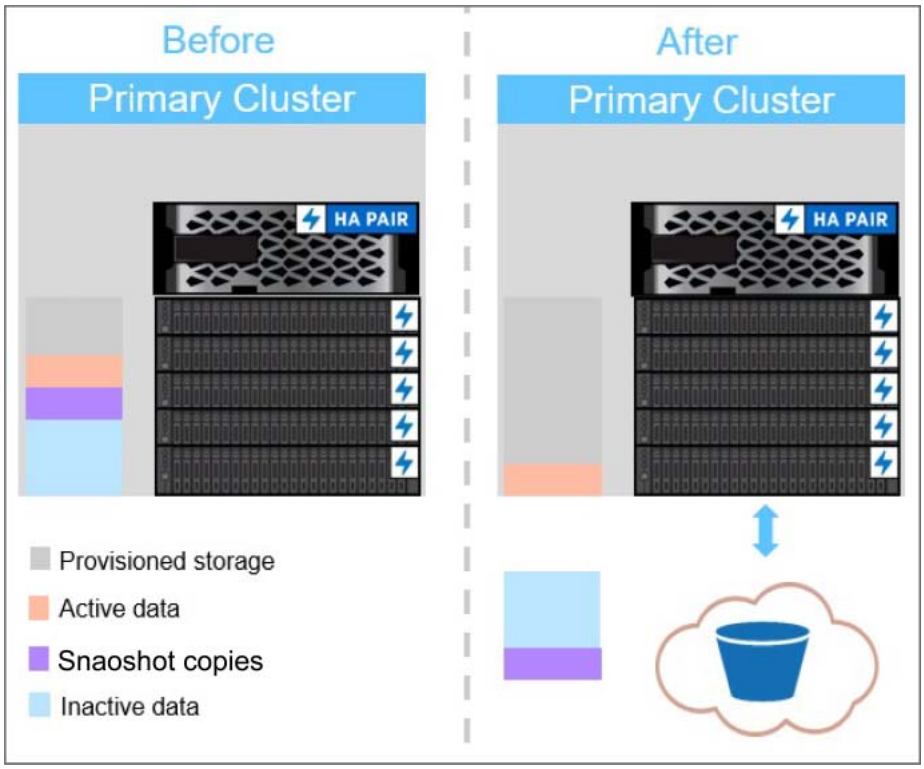

Snapshot-Only

Snapshot copies can frequently consume more than 10% of a typical storage environment. Although essential for data protection and disaster recovery, these point-in-time copies are rarely used and are an inefficient use of high-performance storage.

Snapshot-Only, a volume tiering policy for FabricPool, is an easy way to reclaim storage space on high-performance storage. When configured to use this policy, cold Snapshot copy blocks in the volume that are not shared with the active file system are moved to the cloud tier. If read, cold data blocks on the cloud tier become hot and are moved to the local tier.

Caution

The FabricPool Snapshot-Only volume tiering policy, as shown in Figure 3, reduces the amount of storage used by Snapshot copies on the local tier. It does not increase the maximum number of Snapshot copies allowed by ONTAP, which remains 1,023.

Figure 3 Reclaiming space with the Snapshot-Only volume tiering policy

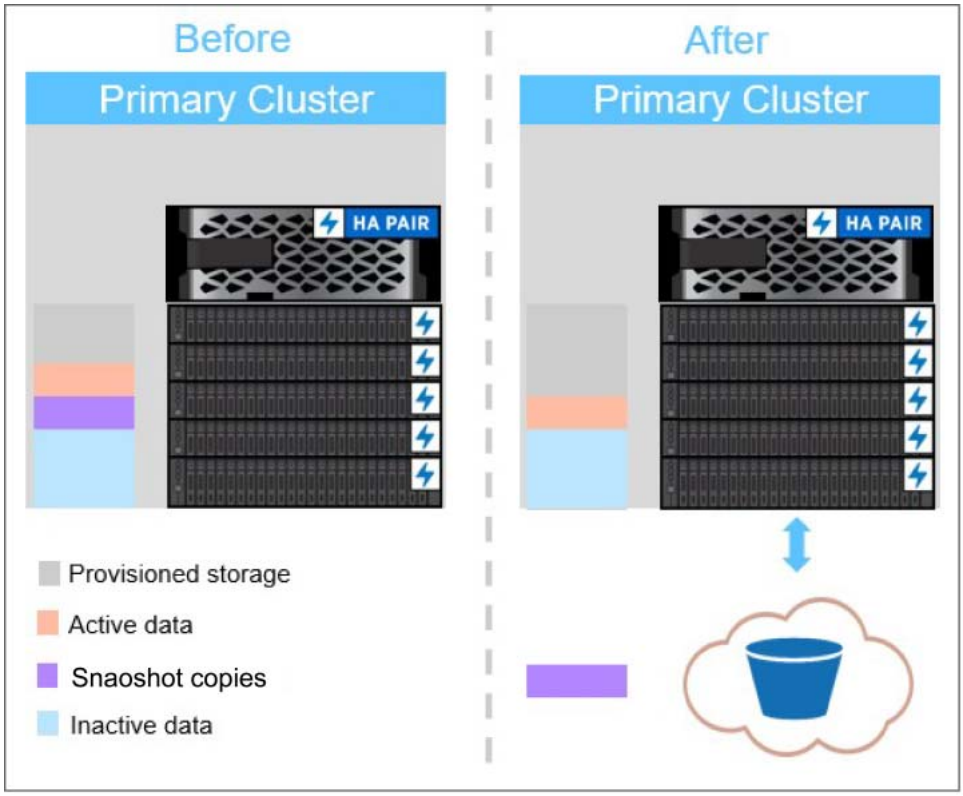

All

In addition to cold primary data in active volumes (Auto) and snapshots (Snapshot-Only), another use of FabricPool is to move entire volumes of secondary (backup and recovery) data to low-cost clouds. Completed projects, legacy reports, or historical records—any dataset that must be retained but is unlikely to be read—are ideal candidates to be tiered to low-cost object storage.

Moving entire volumes is accomplished by setting the All volume tiering policy on a volume. The All volume tiering policy, as shown in Figure 4, is primarily used with secondary data and data protection volumes.

Caution

It is not recommended to use the All volume tiering policy with primary data (read/write volumes). SAN LUNs, in particular, should not be hosted from volumes using the All volume tiering policy.

Data in volumes using the All tiering policy, (excluding data illegible for tiering) is immediately marked as cold and tiered to the cloud as soon as possible. There is no waiting for a minimum number of days to pass before the data is made cold and tiered. If read, cold data blocks on the cloud tier stay cold and are not written back to the local tier.

Figure 4 Reclaiming space with the All volume tiering policy

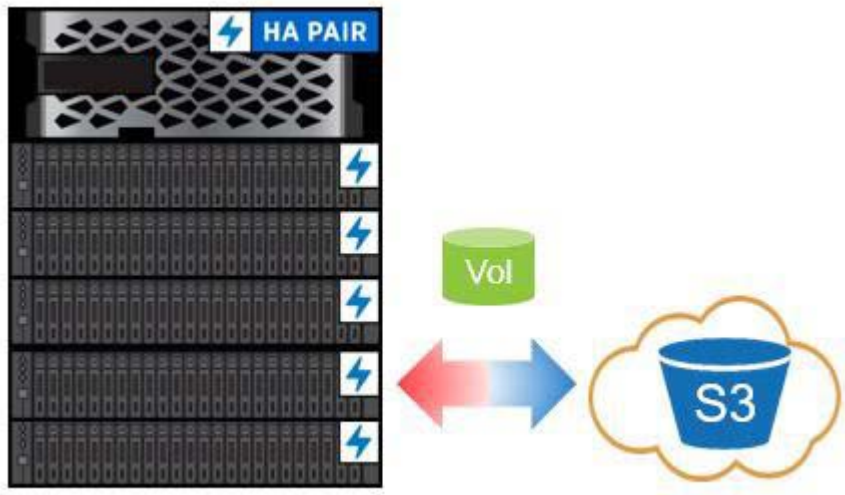

Shrink the Secondary Storage Footprint (All)

Secondary data includes data protection volumes that are SnapMirror (disaster recovery) or SnapVault (backup) destination targets. This data is frequently stored on secondary clusters that share a 1:1 or greater ratio with the primary data that they are protecting (one baseline copy and multiple Snapshot copies). For large datasets, this approach can be prohibitively expensive, forcing users to make expensive decisions about the data they need to protect.

Like Snapshot copies, data protection volumes are infrequently used and are an inefficient use of high-performance storage. The FabricPool All volume tiering policy changes this paradigm.

Instead of 1:1 primary-to-backup ratios, the FabricPool All policy allows users to significantly reduce the number of drive shelves on their secondary clusters, tiering most of the backup data to low-cost object stores. ACLs, directory structures, and WAFL metadata remains on the secondary cluster's local tier.

If read, cold data blocks in volumes using the All policy are not written back to the local tier. This reduces the need for high-capacity secondary storage local tiers.

Figure 5 Using the All volume tiering policy with secondary storage

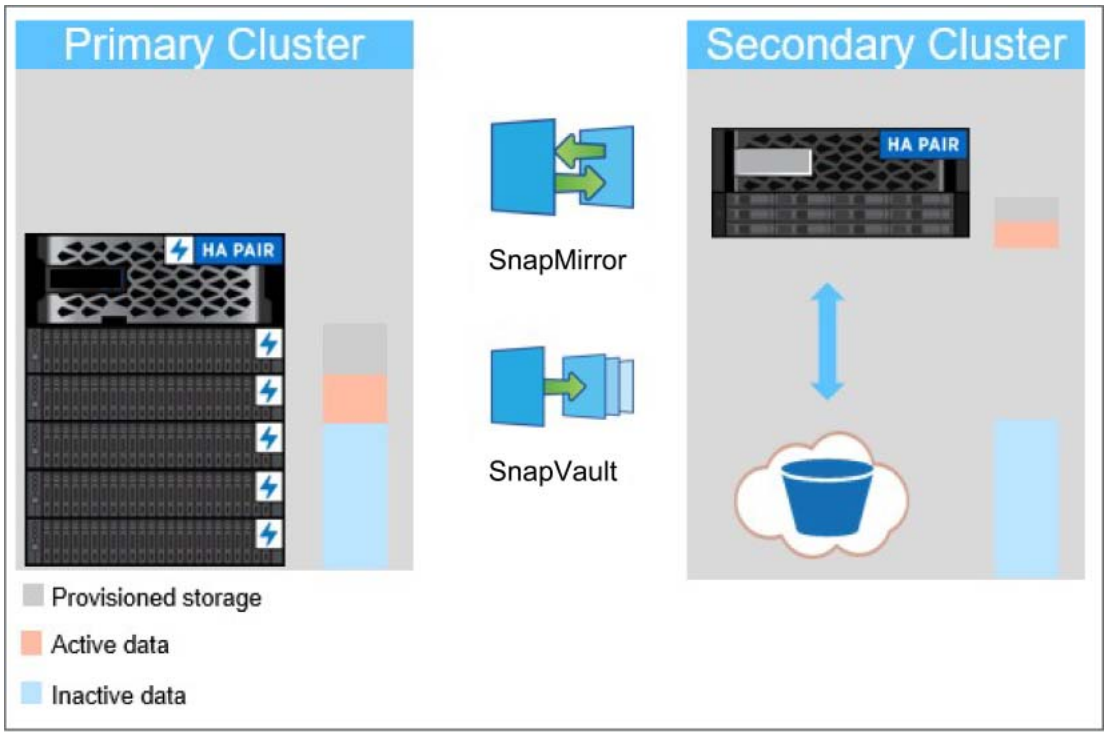

Figure 5 illustrates the secondary cluster as a traditional cluster running ONTAP. The secondary cluster can also be in the cloud using Cloud ONTAP Volumes. You can tier data using FabricPool anywhere ONTAP can be deployed.

3. Requirements

FabricPool requires ONTAP 9.7 or later. Additional FabricPool requirements depend on the version of ONTAP being used and the cloud tier being attached.

In releases earlier than ONTAP 9.8, FabricPool is only supported on SSD local tiers.

Platforms

FabricPool is supported by all the ETERNUS AX/AC/HX series platforms:

- Cloud tiers
  - Object Storage of FUJITSU Hybrid IT Service FJcloud-O
  - Amazon S3 (Standard, Standard-IA, One Zone-IA, Intelligent-Tiering, Glacier Instant Retrieval)
  - Amazon Commercial Cloud Services (C2S)
  - Google Cloud Storage (Multi-Regional/Regional, Nearline/Coldline/Archive)
  - Microsoft Azure Blob Storage (Hot and Cool)
  - StorageGRID 10.3 and later
- Data tiering
  - Cloud Volumes for Google Cloud
  - Cloud Volumes ONTAP for AWS
  - Cloud Volumes ONTAP for Azure

Intercluster LIFs

Cluster high-availability (HA) pairs that use FabricPool require two intercluster LIFs to communicate with the cloud tier. It is also recommended to create an intercluster LIF on additional HA pairs to seamlessly attach cloud tiers to local tiers on those nodes.

If you are using more than one intercluster LIF on a node with different routing, it is recommended to place them in different IPspaces. During configuration, FabricPool can select from multiple IPspaces, but it is unable to select specific intercluster LIFs within an IPspace.

Caution

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

Internet Protocol Version

Beginning in ONTAP 9.9.1, FabricPool supports IPv6. Prior to ONTAP 9.9.1, FabricPool only supports IPv4.

Transmission Control Protocol (TCP) Connections

Object store infrastructure must be capable of supporting at least 700 TCP connections. FabricPool can use 1600-3800 TCP connections per node, per object store endpoint. Server-side load balancers, firewalls, and proxies must be sized to appropriately handle FabricPool traffic.

Volumes

FabricPool cannot attach a cloud tier to a local tier that contains volumes by using a space guarantee other than None.

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

Setting the `space-guarantee none` parameter assures thin provisioning of the volume. The amount of space consumed by volumes with this guaranteed 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 local tier.

FlexGroup Volumes

When provisioning FlexGroup volumes on FabricPool local tiers (storage aggregates), automatic processes in ONTAP System Manager require that the FlexGroup volume uses FabricPool local tiers on every cluster node. This is a recommended best practice but is not a requirement when manually provisioning FlexGroup volumes.

Provisioning FlexGroup constituent volumes on heterogeneous local tiers (some using FabricPool, some not using FabricPool) is not recommended and will result in unpredictable tiering and performance.

Quality of Service Minimums

FabricPool and quality of service minimums (QoS Min) goals are mutually exclusive; QoS Min guarantees performance minimums, whereas FabricPool sends blocks to an object store and decreasing performance. QoS Min must be turned off on volumes in FabricPool local tiers. Alternatively, tiering must be turned off (`-tiering-policy none`) on volumes that require QoS Min.

Cloud Tiering License

FabricPool requires a capacity-based license when attaching third-party object storage providers as cloud tiers for the ETERNUS AX/AC/HX series. A Cloud Tiering license is not required when using StorageGRID or ONTAP S3 as the cloud tier or when using Amazon S3, Google Cloud Storage, or Microsoft Azure Blob Storage as the cloud tier for Cloud Volumes ONTAP.

New Cloud Tiering licenses (including add-on or extensions to preexisting FabricPool licenses) are activated in the Cloud Manager Digital Wallet. You can set up and configure tiering by using the Cloud Tiering service.

Cloud Tiering licenses are available in 12- and 36-month term-based licenses.

You can purchase Cloud Tiering licenses (including additional capacity for existing licenses) in 1TB increments.

Note

Cloud Tiering licenses are attached to a customer's account and the total tiering capacity can be used across multiple clusters.

Licensed Capacity

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 remains accessible from the cloud tier. Additional cold data remains on the local tier until the licensed capacity is increased.

Special Configurations

ONTAP clusters tiering to endpoints other than Amazon S3, Google Cloud Storage, and Microsoft Azure Blob Storage can use Cloud Tiering licenses but the license must be applied in a different manner than typical single-node and HA-configured ONTAP clusters.

Certificate Authority Certification

When FabricPool uses StorageGRID or other private clouds such as some FUJITSU Hybrid IT Service FJcloud-O environments as a cloud tier, it must use a Transport Layer Security (TLS) connection. CA certificates associated with private cloud object stores should be installed on ONTAP before attaching them to local tiers. Using CA certificates creates a trusted relationship between ONTAP and the object store and helps to secure access to management interfaces, gateway nodes, and storage.

Caution

Using signed certificates from a third-party certificate authority is the recommended best practice. Failure to install a CA certificate results in an error unless certificate validation is turned off.

FQDN

FabricPool requires that CA certificates use the same fully qualified domain name (FQDN) as the cloud tier server with which they are associated.

In releases earlier than StorageGRID 11.3, the default CA certificates use a common name (CN) that is not based on the server's FQDN. Using the common name causes certificate-based errors that prohibit StorageGRID from being attached to ONTAP local tiers.

Errors might include the following examples:

- Unable to add a cloud tier. Cannot verify the certificate provided by the object store server. The certificates might not be installed on the cluster. Do you want to add the certificate now?
- Cannot verify the certificate provided by the object store server.

To avoid these errors and successfully attach StorageGRID 11.2 or earlier releases as a cloud tier, you must replace the certificates in the grid with certificates that use the correct FQDN.

Although you can use self-signed certificates, using signed certificates from a third-party certificate authority is the recommended best practice.

Installation

To install CA certificates in ONTAP, complete the following steps:

1. Retrieve the CA certificates.
2. Install the certificates into ONTAP.

Retrieve CA Certificates

Retrieve the Root CA certificate and, if they exist, any intermediate CA certificates in Base-64 encoded format (sometimes also called PEM format) from the Certification Authority who created the certificate.

If you followed the procedure for StorageGRID SSL Certificate Configuration these are the certificates in the `chain.pem` file.

To retrieve the certificate for a StorageGRID endpoint, complete the following steps:

1. Open the StorageGRID Administration console.
2. Select Configuration > Load Balancer Endpoints.
3. Select your endpoint and click Edit Endpoint.
4. Copy the certificate PEM, including:

```
-----BEGIN CERTIFICATE----- and -----END CERTIFICATE-----
```

To retrieve the certificate when using a third-party load balancer, complete the following steps:

1. Run the following command:

```
openssl s_client -connect <FQDN> -showcerts
```

2. Copy the certificate, including:

```
-----BEGIN CERTIFICATE----- and -----END CERTIFICATE-----
```

Install Certificates to ONTAP

In ONTAP System Manager, when adding a new Cloud Tier of type StorageGRID, you can paste the CA certificate. If there is an intermediate CA which issued the StorageGRID certificate, then this must be the intermediate CA certificate. If the StorageGRID certificate was issued directly by the Root CA, then you must use the Root CA certificate.

To install the Root certificates (and any intermediate certificates) to ONTAP, run the following command:

```
security certificate install -vserver <name> -type server-ca
```

4. Architecture

FabricPool works by associating a cloud tier (an external object store) with a local tier (storage aggregate) in ONTAP, creating a composite collection of drives: a FabricPool. Volumes inside the FabricPool can then take advantage of the tiering by keeping active (hot) data on high-performance storage (the local tier) and tiering inactivate (cold) data to the external object store (the cloud tier).

Although only a basic level of understanding is necessary to configure and use FabricPool, understanding how FabricPool determines block temperature, creates objects, and writes data is extremely useful when architecting storage solutions.

Block Temperature

When a block is written to the local tier, it is assigned a temperature value indicating that it is hot. Over time, a background cooling scan cools blocks, making hot blocks warm and eventually turning blocks cold if they have not been read. Assuming no activity, a block becomes cold based on the time set by the `tiering-minimum-cooling-days` setting.

Caution

The All volume tiering policy is an exception to this rule. Blocks in volumes using the All tiering policy are immediately identified as cold and marked for tiering.

Object Creation

FabricPool works at the WAFL block level, cooling blocks, concatenating them into objects, and writing those objects to a cloud tier. Prior to storage efficiencies being applied, each FabricPool object is 4MB and 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. Given ONTAP storage efficiencies, the actual size of a FabricPool object may be less than 4MB.

Data Movement

Tiering Data to an Object Store

After a block has been identified as cold, it is marked for tiering. During this time, a background tiering scan looks for cold blocks. When enough 4KB blocks from the same volume have been collected, they are concatenated into a 4MB object and moved to the cloud tier based on the volume tiering policy.

To view the status of the tiering scan, run the following command:

```
volume object-store tiering show
```

Caution

Advanced privilege level is required.
The volume object-store tiering show command includes a number of optional field parameters that are not displayed by default but may be useful for troubleshooting. For more information, see: https://docs.netapp.com/us-en/ontap-cli-9141/volume-object-store-tiering-show.html

Tiering Fullness Threshold

By default, tiering to the cloud tier only happens if the local tier is >50% full. There is little reason to tier cold data to a cloud tier if the local tier is being underutilized.

In ONTAP 9.5, the 50% tiering fullness threshold is adjustable. Setting the threshold to a lower number reduces the amount of data required to be stored on the local tier before tiering takes place. This may be useful for large local tiers that contain little hot/active data.

Setting the threshold to a higher number increases the amount of data required to be stored on the local tier before tiering takes place. This may be useful for solutions designed to tier only when local tiers are near maximum capacity.

Caution

The All volume tiering policy ignores the tiering fullness threshold. Blocks in volumes using the All tiering policy are tiered irrespective of the tiering fullness threshold.

To change the tiering fullness threshold, run the following command:

```
storage aggregate object-store modify -aggregate <name> -tiering-fullness-threshold <#>
(0%-99%) -object-store-name <name>
```

Caution

Advanced privilege level is required.

Reading Data from an Object Store

When a client application reads data that has been tiered, FabricPool initiates hundreds of concurrent byte-ranged GET operations ranging from 4KB to 288KB. These operations are extremely network efficient as neither the entire object, nor the entire file, needs to be read-only the necessary WAFL blocks.

After being read from the cloud tier, data is immediately passed to the client application.

Random Reads

In order to improve performance, when cold blocks are read from the cloud tier randomly, they are made hot and written back to the local tier. The next read of the same block will come directly from the local tier.

Note

This is the default behavior for volumes using the Auto tiering policy. Write back behavior is dependent on the "Volume Tiering Policies" and "Cloud Retrieval" policies.

Sequential Reads

In order to improve performance, if ONTAP detects an opportunity for sequential readaheads, it requests WAFL blocks from the cloud tier before they are read by the client application. When blocks are read from the cloud tier sequentially, they stay cold and remain on the cloud tier.

Note

This is the default behavior for volumes using the Auto tiering policy. Write back behavior is dependent on the "Volume Tiering Policies" and "Cloud Retrieval" policies.

Write-Back Prevention

If the local tier is at >90% capacity, cold data is read directly from the cloud tier without being written back to the local tier. By preventing cold data write-backs on heavily utilized local tiers, FabricPool preserves the local tier for active data.

Prior to ONTAP 9.7, write-back prevention took place when the local tier was at 70% capacity.

SnapMirror Behavior

Movement of data from the cloud tier to the local tier can take place any time a block is read.

Table 1 SnapMirror behavior

Source Volume Tiering Policy	Destination Volume Tiering Policy	Write Location
Auto	Auto	Local > Local Cloud > Cloud
Auto	Snapshot-Only	Local
Auto	All	Cloud
Auto	None	Local
Snapshot-Only	Auto	Local > Local Cloud > Cloud
Snapshot-Only	Snapshot-Only	Local > Local Cloud > Cloud
Snapshot-Only	All	Cloud
Snapshot-Only	None	Local
All	Auto	Local
All	Snapshot-Only	Local
All*1	All*1	Cloud*1
All	None	Local
None	Auto	Local
None	Snapshot-Only	Local
None	All	Cloud
None	None	Local

*1: Cascading SnapMirror relationships are not supported when using the All volume tiering policy. Only the final destination volume should use the All volume tiering policy.

Volume Move

Volume move (`vol move`) is the way that ONTAP moves a volume nondisruptively from one local tier (source) to another (destination). Volume moves can be performed for a variety of reasons, although the most common reasons are hardware lifecycle management, cluster expansion, and load balancing.

It is important to understand how volume move works with FabricPool because the changes that take place at both the local tier, the attached cloud tier, and the volume (volume tiering policies) can have a major impact on functionality.

Destination Local Tier

If a volume move's destination local tier does not have an attached cloud tier, data on the source volume that is stored on the cloud tier is written to the local tier on the destination local tier.

For the ETERNUS AX/AC/HX, if a volume move's destination local tier uses the same bucket as the source local tier, data on the source volume that is stored in the bucket does not move back to the local tier. This optimized volume move results in significant network efficiencies.

Caution

Some configurations are incompatible with optimized volume moves:

- Changing tiering policy during volume move
- Source and destination aggregates use different encryption keys
- FlexClone volumes
- FlexClone parent volumes
- MetroCluster (supports optimized volume moves in ONTAP 9.8 and later)

If a volume move's destination local tier has an attached cloud tier, data on the source volume that is stored on the cloud tier is first written to the local tier on the destination local tier. It is then written to the cloud tier on the destination local tier if this approach is appropriate for the volume's tiering policy. Moving data to the local tier first improves the performance of the volume move and reduces cutover time.

If a volume tiering policy is not specified when performing a volume move, the destination volume uses the tiering policy of the source volume. If a different tiering policy is specified when performing the volume move, the destination volume is created with the specified tiering policy.

Caution

When in an SVM DR relationship, source and destination volumes must use the same tiering policy.

Minimum Cooling Days

Moving a volume to another local tier resets the inactivity period of blocks on the local tier. For example, a volume using the Auto volume tiering policy with data on the local tier that has been inactive for 20 days has data inactivity reset to 0 days after a volume move.

Auto

If `-tiering-policy auto` is specified during the volume move, data movement is variable, but all data moves to the destination local tier first.

If the source volume uses the Auto, None, or Snapshot-Only policy, blocks are moved to the same tier that they existed on prior to the move. If the source volume uses the All policy, all data is moved to the local tier.

```
vol move start -vserver <name> -volume <name> -destination-aggregate <name> -tiering-policy
auto
```

Snapshot-Only

If `-tiering-policy snapshot-only` is specified during the volume move, data movement is variable, but data moves to the destination local first.

If both source and destination volumes use the Snapshot-Only policy, and the Snapshot block is being read from the source cloud tier, then FabricPool knows the Snapshot blocks are cold and moves the cold blocks to the destination cloud tier.

```
vol move start -vserver <name> -volume <name> -destination-aggregate <name> -tiering-policy
snapshot-only
```

All

If `-tiering-policy all` is specified during the volume move, data is immediately identified as cold and written to the destination cloud tier. There is no need to wait 48 hours for blocks in the volume to become cold. Metadata is always stored on the local tier.

```
vol move start -vserver <name> -volume <name> -destination-aggregate <name> -tiering-policy
all
```

None

If `-tiering-policy none` is specified during the volume move, data is written to the destination local tier.

```
vol move start -vserver <name> -volume <name> -destination-aggregate <name> -tiering-policy
none
```

ONTAP System Manager

To perform a volume, move with ONTAP System Manager, complete the following steps:

1. Click STORAGE.
2. Click Volumes.
3. Select the volume you want to move.
4. Click More.
5. Click Move.
6. Select a destination local tier.
7. Select a tiering policy.
8. Click Move.

Figure 6 Changing the volume tiering policy during a volume move

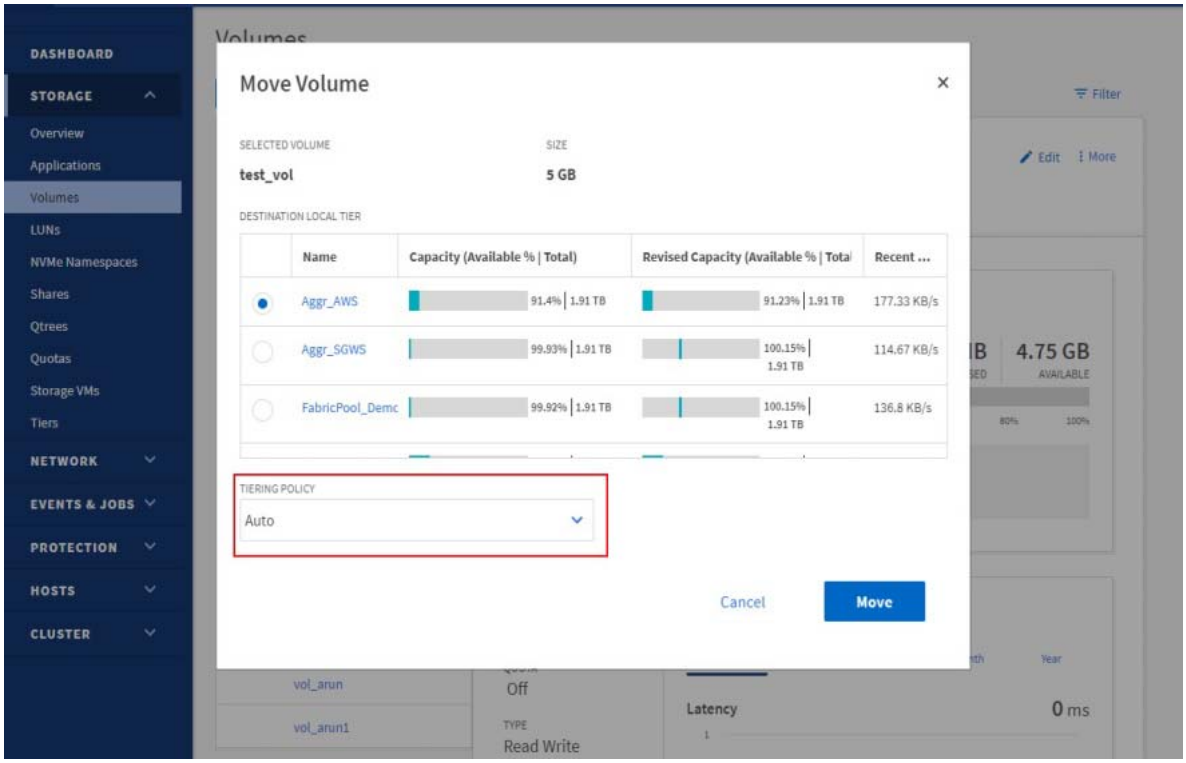

- ONTAP CLI

To perform a volume move using the ONTAP CLI, run the following command:

```
vol move start -vserver <name> -volume <name> -destination-aggregate <name> -tiering-policy
<policy>
```

FlexClone Volumes

FlexClone volumes are copies of a parent FlexVol volume. Newly created FlexClone volumes inherit the volume tiering policy and the tiering-minimum-cooling days setting of the parent FlexVol volume. After a FlexVol volume has been created, you can change the volume tiering policy (see "Volume Tiering Policies").

The tiering policy and tiering-minimum-cooling-days of the clone volume only controls the tiering behavior of blocks unique to the clone. It is recommended to use tiering settings on the parent FlexVol that are either less aggressive than or equally aggressive to any of the clones. As a best practice, this keeps more data owned by the parent volume on the local tier, increasing the performance of the clone volumes.

FlexClone volumes that copy data protection destination volumes using the All tiering policy do not inherit the volume tiering policy of their parent. Instead, they are created using the Snapshot-Only policy.

If a FlexClone volume is split (`volume clone split`) from its parent volume, the copy operation writes the FlexClone volume's blocks to the local tier.

FlexGroup Volumes

A FlexGroup volume is a single namespace that is made up of multiple constituent member volumes but is managed as a single volume. Individual files in a FlexGroup volume are allocated to individual member volumes and are not striped across volumes or nodes.

FlexGroup volumes are not constrained by the 100TB and two-billion file limitations of FlexVol volumes. Instead, FlexGroup volumes are only limited by the physical maximums of the underlying hardware and have been tested to 20PB and 400 billion files. Architectural maximums could be higher.

Volume tiering policies are set at the FlexGroup volume level—they cannot be set on the various constituent/member volumes that compose the FlexGroup volume.

When provisioning FlexGroup volumes on FabricPool local tiers, automatic processes require that the FlexGroup volume uses FabricPool local tier on every cluster node. This is a recommended best practice but not a requirement when manually provisioning FlexGroup volumes.

Object Storage

Object storage is a storage architecture that manages data as objects, as opposed to other storage architectures such as file or block storage. Objects are kept inside a single container (such as a bucket) and are not nested as files inside a directory inside other directories.

Although object storage is generally less performative than file or block storage, it is significantly more scalable. ONTAP currently has a maximum volume size of 100TB and a maximum local tier size of 800TB. Object stores have no such limits, and buckets with petabytes of data in them are not uncommon.

ONTAP S3

Beginning in ONTAP 9.8, ONTAP supports tiering to buckets creating using ONTAP S3, allowing for ONTAP to ONTAP tiering as well. FabricPool can tier to buckets located on the local cluster (a local bucket using cluster LIFs) or buckets located on a remote cluster (a traditional FabricPool cloud tier).

Object Deletion and Defragmentation

FabricPool does not delete blocks from attached object stores. Instead, FabricPool deletes entire objects after a certain percentage of the blocks in the object are no longer referenced by ONTAP.

For example, there are 1,024 4KB blocks in a 4MB object tiered to Amazon S3. Defragmentation and deletion do not occur until less than 205 4KB blocks (20% of 1,024) are being referenced by ONTAP. When enough (1,024) blocks have zero references, their original 4MB objects are deleted, and a new object is created.

You can customize this percentage, the unreclaimed space threshold, but is set to different default levels for different object stores. The default settings are as follows:

Table 2 Default unreclaimed space thresholds.

Object store	ONTAP 9.7	ONTAP 9.8 and later	Cloud Volumes ONTAP
Microsoft Azure Blob Storage	15%	25%	35%
Amazon S3	20%	20%	30%
Google Cloud Storage	12%	20%	35%
Object Storage of FUJITSU Hybrid IT Service FJcloud-O	40%	40%	N/A
ONTAP S3	N/A	40%	N/A
StorageGRID	40%	40%	N/A

Unreclaimed Space Threshold

Object defragmentation reduces the amount of physical capacity used by the cloud tier at the expense of additional object store resources (reads and writes).

Reducing the Threshold

To avoid additional expenses, consider reducing the unreclaimed space thresholds when using object store pricing schemes that reduce the cost of storage but increase the cost of reads.

For example, tiering a volume of 10 year old projects that has been saved for legal reasons might be less expensive when using a pricing scheme such as Standard-IA or Cool than it would be when using standard pricing schemes. Although reads are more expensive for such a volume, including reads required by object defragmentation, they are unlikely to occur frequently.

Increasing the Threshold

Alternatively, consider increasing unreclaimed space thresholds if object fragmentation causes significantly more object store capacity to be used than necessary for the data being referenced by ONTAP. For example, using an unreclaimed space threshold of 20% in a worst-case scenario where all objects are equally fragmented to the maximum allowable extent means that it is possible for 80% of total capacity in the cloud tier to be unreferenced by ONTAP. For example:

- 2TB referenced by ONTAP + 8TB unreferenced by ONTAP = 10TB total capacity used by the cloud tier.

In this situation, it might be advantageous to increase the unreclaimed space threshold—or increase volume minimum cooling days—to reduce the capacity used by unreferenced blocks.

To change the default unreclaimed space threshold, run the following command:

```
storage aggregate object-store modify -aggregate <name> -object-store-name <name>
-unreclaimed-space-threshold <%> (0%-99%)
```

Caution

- Advanced privilege level is required.
- As objects are defragged and made more storage efficient, underlying files might become more fragmented as referenced blocks are written to new, more efficient objects. For this reason, significantly increasing the unreclaimed space threshold results in objects with increased storage efficiency but possibly reduced sequential read performance.

ONTAP Storage Efficiencies

Storage efficiencies such as compression, deduplication, and compaction are preserved when moving data to the cloud tier, reducing required object storage capacity and transport costs.

Aggregate inline deduplication is supported on the local tier, but associated storage efficiencies are not carried over to objects stored on the cloud tier.

When using the All volume tiering policy, storage efficiencies associated with background deduplication processes may be reduced as data is likely to be tiered before the additional storage efficiencies can be applied.

Caution

Third-party deduplication has not been qualified.

Temperature-Sensitive Storage Efficiency

Beginning in ONTAP 9.8, temperature-sensitive storage efficiency (TSSE) uses temperature scans to determine how hot or cold data is and compresses larger or smaller blocks of data accordingly — making storage efficiency more efficient.

Beginning in ONTAP 9.10.1, TSSE is supported on volumes located on FabricPool-enabled local tiers (storage aggregates). TSSE compression-based storage efficiencies are preserved when tiering to cloud tiers. Although more efficient, smaller blocks will require smaller GETs, reducing GET performance from the cloud tier.

Note

Beginning with ONTAP 9.10.1, ETERNUS AX series volumes are created using adaptive compression by default.
(`-storage-efficiency-mode default`)
TSSE must be manually enabled on volumes.
(`-storage-efficiency-mode efficient`)

5. Configuration

After the FabricPool basic requirements have been met, attaching a cloud tier to a local tier in ONTAP requires the following four steps:

1. Create a bucket/container on the object store.
2. Add a cloud tier using the bucket to ONTAP.
3. Attach the cloud tier to a local tier.
4. Set volume tiering policies.

Create a Bucket/Container

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 a local tier as a cloud tier.

Buckets cannot be created using ONTAP System Manager, Active IQ, or ONTAP.

A single cloud tier can be attached to a single local tier, and a single cloud tier can be attached to multiple local tiers. Attaching a single cloud tier to multiple local tiers in a cluster is the general best practice.
It is not recommend to attach a single cloud tier to local tiers in multiple clusters.

Caution

Consider how cloud tier-to-local tier relationships might affect performance when planning storage architectures. Many public object store providers set a maximum number of supported IOPS at the bucket/container level. Environments that require maximum performance from public object stores should use multiple buckets to reduce the possibility that object-store IOPS limitations affect performance across multiple local tiers tiering to the same cloud tier.

Attaching a cloud tier to all FabricPool local tiers is the general best practice and provides significant benefits to environments that value manageability over public object store cloud tier performance.

Figure 7 Possible cloud tier-to-local tier relationships in ONTAP 9.7

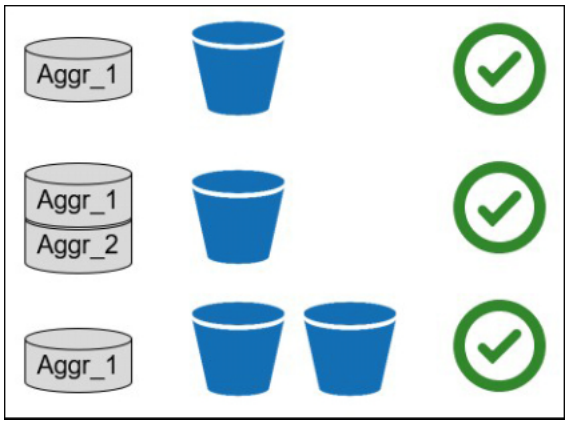

StorageGRID

To create a bucket in StorageGRID, complete the following steps using the StorageGRID Tenant Manager:

1. Open the Admin Node in a web browser (for example, https://admin.company.com/?accountId=###).
2. Log in with your tenant account ID, user name, and password.
3. Select S3.
4. Select Buckets.
5. Click Create Bucket.
6. Provide a DNS compliant name.
7. Click Save.

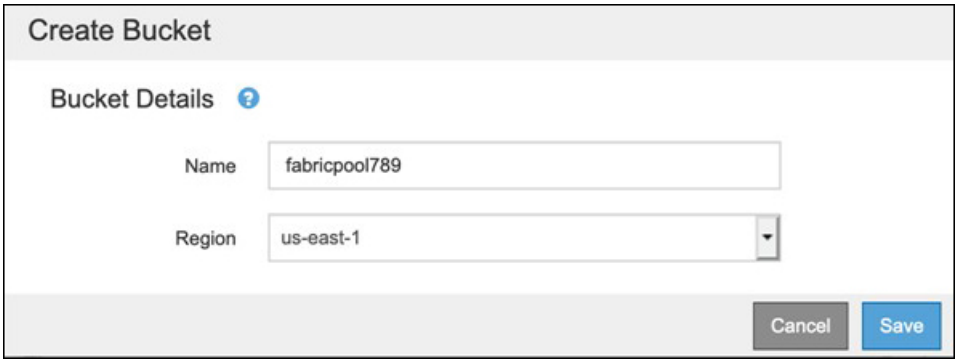

Caution

- In releases earlier than StorageGRID 11.1, creating a bucket required using a third-party S3 client such as an S3 browser.
- ONTAP and StorageGRID system clocks must not be out of sync by more than a few minutes. Significant clock skew prevents the StorageGRID bucket from being attached to the local tier.

Other Object Store Providers

Instructions for creating buckets on other object store providers can be found on their respective sites:

- Amazon S3
- Google Cloud Storage
- Microsoft Azure Blob Storage

To create buckets on FUJITSU Hybrid IT Service FJcloud-O, contact our support division.

Other Object Store Provider Settings

Outside of BlueXP and StorageGRID, FabricPool does not support ILM policies applied to object store buckets.

ILM typically includes various movement and deletion policies based on geography, storage class, retention, and other categories that would be disruptive to FabricPool cloud tier data. FabricPool has no knowledge of ILM policies or configurations set on external object stores, and misconfiguration of ILM policies can result in data loss.

Caution

ONTAP and private cloud system clocks must not be out of sync by more than a few minutes. Significant clock skew prevents the Cleversafe bucket from being attached to the local tier.

Add a Cloud Tier to ONTAP

Before a cloud tier can be attached to a local tier, it must be added to and identified by ONTAP. You can complete this task by using Cloud Manager's Cloud Tiering Service.

ONTAP System Manager

FabricPool licenses continue to be supported for ONTAP environments and third-party object storage providers not supported by Cloud Manager. These environments include:

- Alibaba Cloud Object Storage Service
- Amazon Commercial Cloud Services
- IBM Cloud Object Storage
- MetroCluster
- Dark site or otherwise air-gapped environments which are not yet supported by Cloud Tiering
  When adding a cloud tier by using System Manager or the CLI, you need the following information:
- Server name (FQDN) (for example, `s3.amazonaws.com`)

  Caution

  Azure might require the account prefix (for example, `accountprefix.blob.core.windows.net`)
- Access key ID
- Secret key
- Container name (bucket name)

To add a cloud tier using ONTAP System Manager, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click Tiers.
4. Click Add Cloud Tier.
5. Select an object store provider.

6 Complete the text fields as required for your object store provider.

Caution

Enter the object store's bucket/container name in the Container Name field.

7 (Optional; cloud tiers can be attached to local tiers later if desired.) Add the cloud tier to local tiers as a primary cloud.

Caution

Attaching a cloud tier to a local tier is a permanent action. A cloud tier cannot be unattached from a local tier after being attached.

8 Click Save.

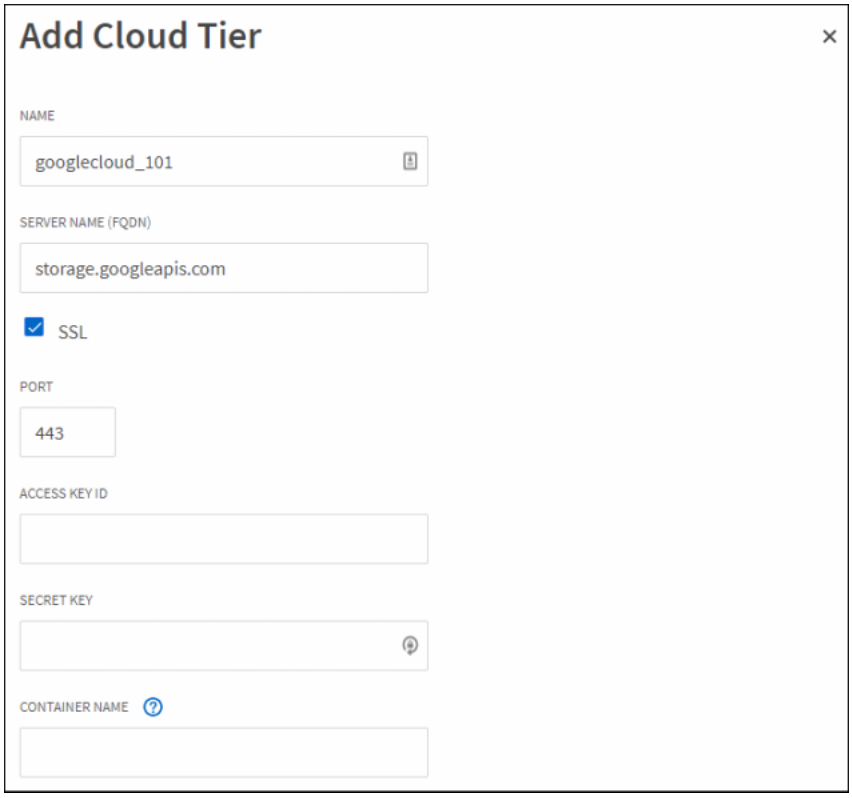

ONTAP CLI

To add a cloud tier by using the ONTAP CLI, run the following commands:

```
object-store config create
-object-store-name <name>
-provider-type <AliCloud/AWS/Azure_Cloud/CAP/GoogleCloud/IBM_COS/ONTAP_S3/S3_Compatible/SGWS>
-port <443/8082> (public clouds/SGWS)
-server <name>
-container-name <bucket-name>
-access-key <string>
-secret-password <string>
-ssl-enabled true
-ipspace default
-is-certificate-validation-enabled true
-use-http-proxy false
-url-style <path-style/virtual-hosted-style>
```

ONTAP S3 Local Buckets

Beginning in ONTAP 9.8, ONTAP supports tiering to buckets created using ONTAP S3, allowing for ONTAP to ONTAP tiering. Buckets located on the local cluster are known to ONTAP automatically and are available as an option when attaching a cloud tier to a local tier.

Certificate Authority Certificate Validation

CA certificates associated with private cloud object stores, such as StorageGRID and FUJITSU Hybrid IT Service FJcloud-O environments, should be installed on ONTAP before attaching them to local tiers. Using CA certificates creates a trusted relationship between ONTAP and the object store and helps to secure access to management interfaces, gateway nodes, and storage.

Failure to install a CA certificate results in an error unless certificate validation is turned off. Turning off certificate validation is possible, but not recommended.

ONTAP System Manager

CA certificate validation can be turned off when adding a StorageGRID cloud tier using ONTAP System Manager. To do so, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click Tiers.
4. Click Add Cloud Tier.
5. Select an object store provider.
6. Complete the text fields as required for your object store provider.
7. Click the Object Store Certificate button to turn it off.

Caution

Turning off certificate validation is not recommended.

8. Click Save.

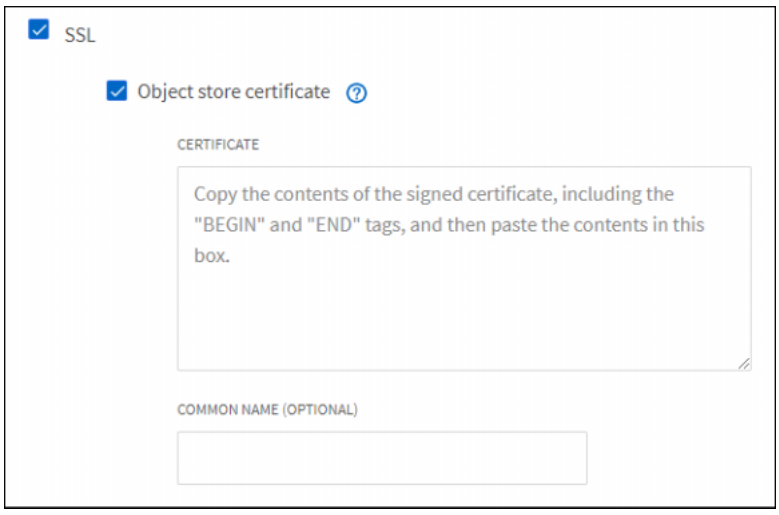

- ONTAP CLI

  You can turn off CA certificate validation when adding a private cloud tier by using the ONTAP CLI. To do so, run the following commands:

```
object-store config create
-object-store-name <name>
-provider-type <IBM_COS/ONTAP_S3/S3_Compatible/SGWS>
-port <443/8082> (other providers/SGWS)
-server <name>
-container-name <bucket-name>
-access-key <string>
-secret-password <string>
-ssl-enabled true
-ipspace default
-is-certificate-validation-enabled false
-use-http-proxy false
-url-stle <path-style/virtual-hosted-stle>
```

Attach a Cloud Tier to a Local Tier

After an object store has been added to and identified by ONTAP as a cloud tier, it can be attached to a local tier to create a FabricPool. You can complete this task by using either ONTAP System Manager or the ONTAP CLI.

Caution

Attaching a cloud tier to a local tier is a permanent action. A cloud tier cannot be unattached from a local tier after being attached.

Thin Provisioning

FabricPool cannot attach a cloud tier to a local tier that contains volumes using a space guarantee other than none (for example, `volume`). For additional information, see FabricPool's requirements.

FlexGroup Volumes

When provisioning FlexGroup volumes on FabricPool local tiers (storage aggregates), automatic processes in ONTAP System Manager require that the FlexGroup volume uses FabricPool local tiers on every cluster node. This is a recommended best practice but is not a requirement when manually provisioning Flex-Group volumes.

Provisioning FlexGroup constituent volumes on heterogeneous local tiers (some using FabricPool, some not using FabricPool) is not recommended and will result in unpredictable tiering and performance.

Caution

Consider how cloud tier-to-local tier relationships might affect performance when planning storage architectures. Many public object store providers set a maximum number of supported IOPS at the bucket/container level. Environments that require maximum performance from public object stores should use multiple buckets to reduce the possibility that object-store IOPS limitations affect performance across multiple local tiers tiering to the same cloud tier.

Attaching a single cloud tier endpoint to all FabricPool local tiers is the general best practice and provides significant benefits to environments that value manageability over minor gains in public object store cloud tier performance.

ONTAP System Manager

To attach a cloud tier to a local tier using ONTAP System Manager, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click the name of a local tier.
4. Click More.
5. Click Attach Cloud Tiers.
6. Select the primary cloud tier to attach.
7. Select volumes to set tiering policies.
8. Click Save.

Caution

Attaching a cloud tier to a local tier is a permanent action. A cloud tier cannot be unattached from a local tier after being attached.

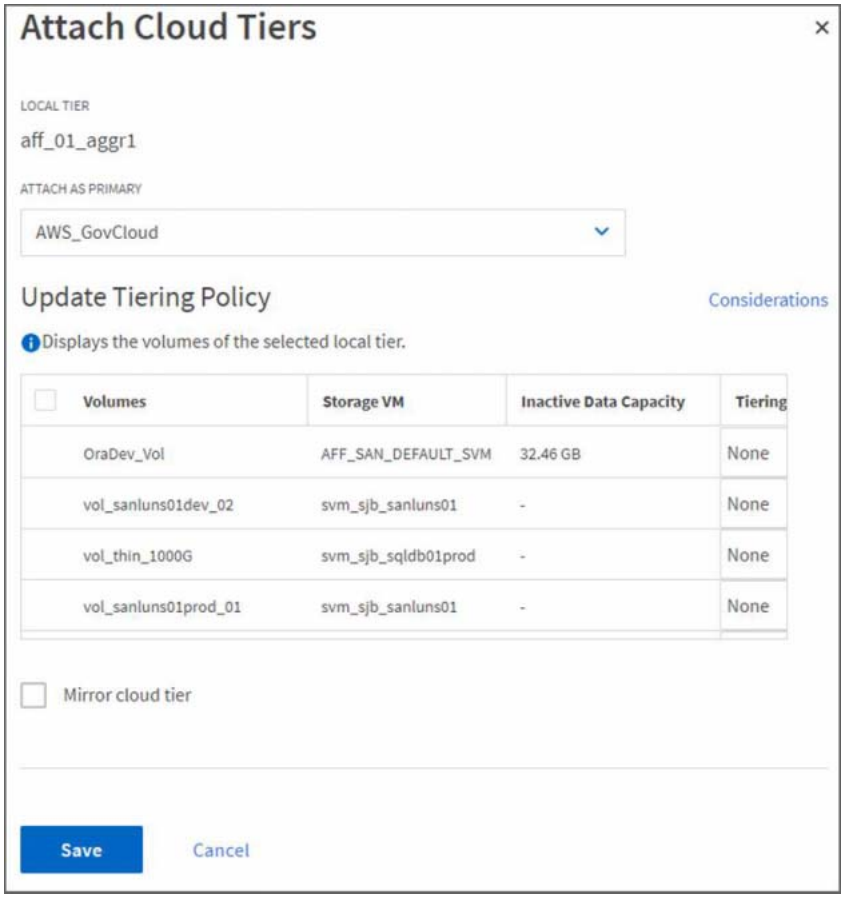

ONTAP CLI

To attach a cloud tier to a local tier (storage 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_fabricpool_bucket
```

Caution

Attaching a cloud tier to a local tier is a permanent action. A cloud tier cannot be unattached from a local tier after being attached.

FlexGroup Volumes

To list the local tiers used by a FlexGroup volume, and attach a cloud tier to those local tiers by using the ONTAP CLI, run the following commands:

```
volume show -volume <name> -fields aggr-list
```

Then:

```
storage aggregate object-store attach
-aggregate <name>
-object-store-name <name>
-allow-flexgroup true
```

ONTAP S3 Local Buckets

To attach a local bucket to a local tier by using ONTAP System Manager, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click the name of a local tier.
4. Click More.
5. Click Tier to Local Bucket.

6. Select Existing or New.

   If selecting New, a new SVM and bucket is created. If available, System Manager selects low-cost media (ETERNUS HX series HDD) for the bucket.
7. Select bucket capacity.
8. Click Save.

   When a new bucket is created, its secret key is displayed. Save/download this key for future use because it is not displayed again.

Caution

- Unlike local tiers attached to cloud tiers where FabricPool uses intercluster LIFs to communicate with the cloud tier, when a local tier is attached to a local bucket, FabricPool uses cluster LIFs for intracluster traffic. When ONTAP S3 is used as a cloud tier, verify that the Lifs associated with these IP Addresses have ‘data-s3-server’ policy associated with them.
- Performance degradation may occur if cluster LIF resources become saturated. To avoid this, it is recommended to use four-node or greater clusters when tiering to a local bucket—the recommended best practice being an HA pair for the local tier and an HA pair for the local bucket. Tiering to local buckets on a single HA pair is not recommended.

Volume Tiering Policies

By default, volumes use the None volume tiering policy. The exception to this are newly created FlexVol volumes on FabricPool aggregates which use the Snapshot-Only volume tiering policy.

After volume creation, the volume tiering policy can be changed using ONTAP System Manager or the ONTAP CLI.

FabricPool provides four volume tiering policies, as described in the following sections.

Caution

When used by FlexGroup volumes, the volume tiering policy is set at the FlexGroup volume level. Volume tiering policies cannot be set on the various constituent/member volumes that compose the FlexGroup volume.

- **Auto:**
  - All cold blocks in the volume are moved to the cloud tier. Assuming the local tier is >50% utilized, it takes approximately 31 days for inactive blocks to become cold. The Auto cooling period is adjustable between 2 days and 183 days by using the minimum-cooling-days setting. (63-day maximum in releases earlier than ONTAP 9.8.)
  - When cold blocks in a volume with a tiering policy set to Auto are read randomly, they are made hot and written to the local 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 local tier.
  - Object storage is not transactional like file or block storage. Making changes to files being stored as objects in volumes with overly aggressive minimum cooling days can result in the creation of new objects, fragmentation of existing objects, decreased read performance, and the addition of storage inefficiencies.
- **Snapshot-Only:**
  - Cold Snapshot blocks in the volume that are not shared with the active file system are moved to the cloud tier. Assuming the local tier is >50% utilized, it takes approximately two days for inactive Snapshot blocks to become cold. The Snapshot-Only cooling period is adjustable from 2 to 183 days using the minimum-cooling-days setting. (63-day maximum prior to ONTAP 9.8.)
  - When read, cold blocks associated with Snapshot copies stay cold and are not written back to the local tier.
- **All:**
  - All data blocks (not including metadata) placed in the volume are immediately marked as cold and moved to the cloud tier as soon as possible. There is no need to wait 48 hours for new blocks in a volume using the All tiering policy to become cold.
  - When cold blocks in a volume with a tiering policy set to All are read, they remain cold and stay on the cloud tier. They are not written to the local tier.
  - Object storage is not transactional like file or block storage. Making changes to files being stored as objects in volumes using the All tiering policy can result in the creation of new objects, fragmentation of existing objects, decreased read performance, and the addition of storage inefficiencies.

  Caution

  It is not recommend to use the All volume tiering policy with primary data (read/write volumes). SAN LUNs, in particular, should not be hosted from volumes using the All volume tiering policy.

  Because the All tiering policy tiers data as soon as possible, storage efficiencies that rely on background processes, like deduplication, might not have enough time to be applied. Inline storage efficiencies like compression and compaction are still applied.

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 reads data from the cloud tier rather than the local tier. This results in slower SnapMirror operations—possibly slowing other SnapMirror operations later in queue—even if they are using different tiering policies.

- **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 local tier.
  - When cold blocks in a volume with a tiering policy set to None are read, they are made hot and written to the local tier.

ONTAP System Manager

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

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click Volumes.
4. Select a volume.
5. Click Edit.
6. Select the tiering policy you want to apply to the volume.

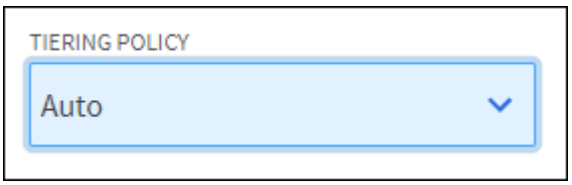

7. Click Save.

Caution

Beginning in ONTAP 9.8, changing the tiering policy to All, Auto, or Snapshot-Only triggers the background tiering scan to immediately run.

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>
```

Note

The default volume tiering policy is None. The exception to this is are newly created FlexVol volumes on FabricPool aggregates which use the Snapshot-Only volume tiering policy.

Cloud Retrieval

When using the Auto volume tiering policy, if cold blocks are read sequentially, they stay cold and remain on the cloud tier. For most client applications this is desirable behavior and prevents deep file scans common to antivirus and analytics applications from writing cold data back to the local tier.

Beginning in ONTAP 9.8, volumes can set cloud retrieval policies to override this default behavior.

FabricPool provides four cloud retrieval policies, as described in the following sections.

- **Default (default):**
  - When cold blocks in a volume are read, they use the default behavior of their volume tiering policy.
- **Never:**
  - When cold blocks in a volume with a cloud retrieval policy set to Never are read, they remain cold and stay on the cloud tier. They are not written to the local tier.
  - Setting the cloud retrieval policy to Never is similar to the All tiering policy in that data is not allowed to return to the local tier but differs from the All tiering policy in that it continues to use the volume's tiering-minimum-cooling-days setting rather than being tiered as soon as possible.
  - For example, a volume using the Auto tiering policy's default setting would not mark data as cold until after 31-days of inactivity. After 31-days, the inactive data would be tiered to object storage and would not come back when read because the volume's cloud retrieval policy had been set to Never.
- **On-Read:**
  - When cold blocks in a volume with a cloud retrieval policy set to On-Read are read, randomly or sequentially, they are made hot and written to the local tier.
  - Applications that use sequential reads triggers write-backs to the local tier by setting the volume cloud retrieval policy to On-Read. This can be beneficial for applications that need local-tier performance for previously cold data that is now being read by active workloads.
- **Promote:**
  - Setting the cloud retrieval policy to Promote immediately queues tiered data to return to the local tier—provided the tiering policy allows it. For example:
    - Bring all data back to the local tier:

	Tiering policy	Cloud retrieval policy
Before	Auto	Default
After	None (Cold blocks are not tiered)	Promote (Previously tiered blocks return to the local tier)

    - Bring the active file system back to the local tier, but keep snapshot copies on the cloud tier:

	Tiering policy	Cloud retrieval policy
Before	Auto	Default
After	Snapshot-Only (Only cold Snapshot blocks is tiered)	Promote (Previously tiered, non-Snapshot blocks, return to the local tier)

Caution

Promote GETs operations are automatically placed at a lower priority (bullied) by all other workloads. They will not compete with other client applications for compute or network resources, but data retrieval may be slow.

Consider using a "Volume Move" instead of Promote if bringing data back rapidly is a priority.

ONTAP CLI

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

```
volume modify -vserver <svm_name> -volume <volume_name> -cloud-retrieval-policy
<default|never|on-read|promote>
```

Note

Advanced privilege level is required.

Volume Tiering Minimum Cooling Days

FabricPool is not an ILM policy that permanently archives data after a set period of time. FabricPool is a high-performance tiering solution that makes data immediately accessible and dynamically moves data to and from the cloud tier-based client application activity.

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.

Caution

- Increasing -tiering-minimum-cooling-days increases the footprint of inactive data on the local tier: data takes longer before it is marked inactive and eligible for tiering to the cloud tier. Additionally, if data is read from the cloud tier, made hot, and written back to the local tier, it takes longer to become inactive again and tiered back to the cloud.
- Although 60-day, 90-day, or 180-day minimum cooling policies may be needed to conform to SLAs that require data to stay on a specific tier of storage (SLAs that are time-based rather than activity based), they are not recommended as a best practice.

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 local 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 multi-day workload is expected to perform a significant number of writes on day seven, the volume's tiering-minimum-cooling-days setting should be set no lower than eight days.

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

Snapshot-Only

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

ONTAP CLI

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

```
volume modify -vserver <svm_name> -volume <volume_name> -tiering-minimum-cooling-days <2-183>
```

Caution

- Advanced privilege level is required.
- Changing the tiering policy between Auto and Snapshot-Only (or vice versa) resets the `tiering-minimum-cooling-days` parameter to its default setting for the target policy. For example, a volume using the Auto volume tiering policy with data on the local tier that has been inactive for 20 days has the `tiering-minimum-cooling-days` parameter reset to 2 days if the tiering policy is set to Snapshot-Only.

MetroCluster

MetroCluster provides continuous data availability across geographically separated data centers for mission-critical applications. MetroCluster continuous availability and disaster recovery software runs on ONTAP data management software. Ethernet (IP) MetroCluster configuration is used by thousands of enterprises worldwide for high availability, zero data loss, and nondisruptive operations both within and beyond the data center.

Note

In order to successfully create a FabricPool local tier in MetroCluster, primary and mirror buckets must be accessible from both clusters.

Licensed capacity

FabricPool licensed capacity applies equally to both buckets in a MetroCluster configuration.

Security

FabricPool maintains AES-256-GCM encryption on the local tier, on the cloud tier, and over the wire when moving data between the tiers.

Local Tier

FabricPool supports Storage Encryption (SE), Volume Encryption (VE), and Aggregate Encryption (AE). Neither SE, VE, nor AE are required to use FabricPool.

Over the Wire

Objects moving between local and cloud tiers are encrypted by using TLS 1.2 using AES-256-GCM. Other encryption modes, such as CCM, are not supported. To some extent, encryption affects connectivity (latency) because object stores must use CPU cycles to decrypt the data. Communicating with object stores without TLS encryption is supported but is not recommended.

Cloud Tier

All objects encrypted by VE/AE remain encrypted when moved to the cloud tier. Client-side encryption keys are owned by ONTAP.

All objects not encrypted using VE/AE are automatically encrypted server-side using AES-256-GCM encryption. No additional encryption is necessary. Server-side encryption keys are owned by the respective object store.

Caution

FabricPool requires the use of the AES-256-GCM authenticated encryption. Other encryption modes, such as CCM, are not supported.

Disabling Cloud Tier Encryption

Using FabricPool without encrypting data at rest is not recommended but may be required by low performance S3 compatible object storage providers who cannot provide server-side encryption and low latency at the same time. It is highly recommended using client-side VE or AE encryption in these circumstances as encrypting data at rest remains the recommended best practice.

To disable cloud tier encryption, run the following command:

```
storage aggregate object-store config modify -serverside-encyption false
```

Caution

Advanced privilege level is required.

6. Interoperability

In general, ONTAP functionality is unchanged on FabricPool local tiers. Although ONTAP must create and transfer objects and blocks between local and cloud tiers, data protection, efficiency, and security are nearly identical to standard local tiers in ONTAP. The primary differentiators are performance and cost, with object stores being slower and less expensive.

The exceptions to normal interoperability listed in Table 3 and Table 4 are unique to FabricPool local tiers.

Table 3 Interoperability

Focus	Supported	Not Supported
Cloud tier	• ONTAP S3 9.8+ • StorageGRID 10.3+	ONTAP S3 in multiprotocol NAS volumes
Data protection	• MetroCluster • MetroCluster SDS • SnapMirror (XDP and DP) • SnapMirror Synchronous • SnapVault (XDP and DP) • SVM-DR • SVM Migrate • StorageGRID replication and erasure coding **Caution** For best results, use replication with StorageGRID 11.2 or later and erasure coding with StorageGRID 11.3 or later.	• 7-Mode Data Transition Using SnapMirror • 7-Mode Transition Tool (7MTT) • DP_Optimized license (DPO) • Object versioning • Secure Purge • SMTape • SnapLock technology • Cascading SnapMirror relationships using the All (or Backup) tiering policy. • StorageGRID ILM policies other than replication and erasure coding • StorageGRID Compliance buckets • SyncMirror technology • Tamperproof Snapshot copies • WORM
Encryption	• Volume Encryption • Storage Encryption • Server-side encryption (AES-256) • TLS 1.2	-
Storage efficiency	• Inline deduplication • Inline compression • Compaction • Aggregate inline deduplication • (local tier only)	-
Storage virtualization	-	FlexArray technology
Quality of service (QoS)	QoS maximums (ceiling)	QoS minimums (floors)
Additional features	BlueXP Storage Class Lifecycle Management	• Auto Balance Aggregate • Flash Pools

Table 4 Third-party interoperability

Focus	Supported	Not Supported
Cloud tier	• Amazon S3 (Standard, Standard-IA, One Zone-IA, Intelligent-Tiering) • Amazon Commercial Cloud Services (C2S) • Google Cloud Storage (Multi-Regional/Regional, Nearline/Coldline/Archive) • S3 in ONTAP 9.8 and later • Microsoft Azure Blob Storage (Hot and Cool) • StorageGRID 10.3+	-
Data protection	Amazon's 99.999999999% multi-region durability	ILM policies
Encryption	Server-side encryption (AES-256) TLS 1.2	-

7. Performance

Network Connections

FabricPool read latency is a function of connectivity to the cloud tier. LIFs using 10 Gbps ports provide adequate performance. It is recommended to validate the latency and throughput of your specific network environment to determine the impact it has on the FabricPool performance.

Because performance can be significantly better, using 10Gbps or 25Gbps is the recommended best practice for FabricPool.

- StorageGRID

  Unlike public clouds that might set a maximum number of supported IOPS at the bucket/container level, StorageGRID performance scales with the number of nodes in a system. For acceptable performance targets, it is recommended to use enough nodes to meet or exceed the FabricPool connectivity requirements.

Object Store Profiler

Beginning in ONTAP 9.4, an object store profiler is available through the CLI that lets you test latency and throughput performance of object stores before you attach them to FabricPool local tiers.
You must "Add a Cloud Tier to ONTAP" before you can use it with the object store profiler.

1. Start the object store profiler.

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

Caution

Advanced privilege level is required.

2. View the results.

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

Caution

Object store profiler results are a measurement of connectivity between ONTAP and the cloud tier object store by using 4MB PUT operations and random-read byte-ranged GET operations ranging from 4KB to 256KB. (Only internal ONTAP features, such as SnapMirror, can make use of 256KB GET operations, third-party clients cannot.)

Object store profiler results are not an indicator of client application performance and do not consider competing workloads or unique client application behavior.

Table 5 FabricPool byte-ranged GET sizes

	ONTAP 9.8 and earlier	ONTAP 9.9.1+
Random reads	4KB 8KB 32KB 256KB	4KB 8KB 32KB 256KB
Sequential reads	4KB 8KB 32KB 256KB	36KB 40KB 64KB 288KB

Sequential Read Performance

ONTAP's adaptive readahead algorithms are designed to anticipate what data will be requested next and read it into memory before the read request arrives.

Beginning in ONTAP 9.13.1, FabricPool performance was improved by increasing the concurrency and parallelism of byte-ranged GETs during sequential reads. Customers can expect significant improvements in both multi-file and single-file sequential read performance.

Caution

Customers using private object stores should consider performance headroom on the object store and if throttling FabricPool PUTs may be necessary. Although most object stores used as cloud tiers are dedicated to FabricPool traffic, not all are, and FabricPool may bully other object store clients.

Aggressive Read-Ahead

One of the advantages of FabricPool's block-based tiering is its network efficiency. FabricPool only reads the WAFL blocks that the client application needs–it does not need to read the entire file. This can result in a substantial reduction in network traffic–especially for large GB-sized and TB-sized files.

Enabling aggressive read-ahead on a volume turns this functionality off and preemptively reads the entire file sequentially from the object store?increasing GET throughput?and reducing the latency of client reads on the file. By default, when tiered data is read sequentially it stays cold and is not written to the local tier.

Aggressive read-ahead trades network efficiency for increased performance of tiered data.

Caution

Additional network traffic may result in additional costs when tiering to public clouds, particularly when using storage classes that reduce the cost of storage but increase the cost of reads such as Amazon's Standard-IA and Azure Blob Storage's Cool.

To enable aggressive read-ahead, run the following command:

```
To enable aggressive read-ahead, run the following command: volume modify -vserver <name>
-volume <name> -aggressive-readahead-mode <file_prefetch>
```

Caution

Advanced privilege level is required.

PUT Throttling

PUT throttling enables storage administrators to set an upper threshold on the maximum per node put rate.

PUT throttling is useful when network resources or the object store endpoint are resource constrained. Although rare, resource constraints can occur with underpowered object stores or during the first days of FabricPool usage when TB or PB of cold data begins to tier out.

PUT throttling is per node. The minimum PUT throttling put-rate-limit is 8MB/s. Setting the put-rate-limit to a value less than 8MB/s will result in 8MB/s throughput on that node. Multiple nodes, tiering concurrently, may consume more bandwidth and potentially saturate a network link with extremely limited capacity.

Caution

FabricPool PUT operations do not compete for resources with other applications. FabricPool PUT operations are automatically placed at a lower priority (bullied) by client applications and other ONTAP workloads, such as SnapMirror. PUT throttling using `put-rate-limit` may be useful for reducing network traffic associated with FabricPool tiering but it is unrelated to concurrent ONTAP traffic.

To throttle FabricPool PUT operations by using the ONTAP CLI, run the following command:

```
storage aggregate object-store put-rate-limit modify -node <name> -default <true|false> -put-
rate-bytes-limit <integer>[KB|MB|GB|TB|PB]
```

Caution

Advanced privilege level is required.

SnapMirror Concurrency

Because concurrent SnapMirror and SnapVault replication operations share the network link to the cloud tier, initialization and RTO are dependent on the available bandwidth and latency to the cloud tier. Performance degradation might occur if connectivity resources become saturated.

Proactive configuration of multiple LIFs can significantly decrease this type of network saturation.

Caution

If you are using more than one intercluster LIF on a node with different routing, it is recommended to place them in different IPspaces. During configuration, FabricPool can select from multiple IPspaces, but it is unable to select specific intercluster LIFs within an IPspace.

Low Performance Environments

Although cloud tiers can provide SATA-like throughput and sub 100ms latencies, performance varies between providers. FabricPool can tolerate latencies as high as 10 seconds and low throughputs for tiering solutions that do not need, or cloud tiers that cannot provide, SATA-like performance.

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

Loss of Connectivity

If for any reason connectivity to the cloud is lost, the FabricPool local tier remains online, but applications receive an error message when attempting to get data from the cloud tier. Cold blocks that exist exclusively on the cloud tier remain unavailable until connectivity is reestablished.

NAS Protocols

NFS and SMB protocols generally retry every five seconds until a connection is reestablished.

Error messages include the following:

- **SMB**
  `STATUS_INTERNAL_ERROR`
  Client applications might or might not retry upon receiving this error (this is client dependent). The client does not have to remount.
- **NFS**
  v3: `EJUKEBOX`
  v4: `EDELAY`
  NFS client applications retry after five seconds. The NFS client hangs until connectivity is reestablished if it gets the same error after a retry.

SAN Protocols

FC and iSCSI protocols generally take longer before experiencing a timeout (120 seconds), but they do not retry to establish a connection in the same way NAS protocols do. If a SAN protocol times out, the application must be restarted. This behavior is consistent for all SAN transport protocols supported by ONTAP.

- **SAN**
  `UNRECOVERED READ ERROR/RECOMMEND REWRITE THE DATA`
  If the host is connected to the ONTAP LUN and the LUN is configured in a RAID set on the host (for example, Volume Manager), the host RAID subsystem might be able to recover the data from parity, and the data is rewritten to a new location. If the host is unable to recover this data, then the application on the host might need to be restarted so that the read can be retried.

It is recommended using the following guidance when tiering data in volumes hosting LUNs:

- **Snapshot-Only**
  Snapshot-Only is an acceptable tiering policy for most SAN use cases.
- **Auto**
  Auto should only be used for non-critical applications.
- **All**
  All should not be used on volumes hosting LUNs.

Caution

Even a short disruption can be disastrous to production applications using SAN protocols. It is recommended to use private networks and private clouds, such as ONTAP S3 or StorageGRID object stores when tiering data that is accessed via SAN protocols.

Virtualized Object Storage

Do not host virtualized object stores, (sometimes referred to as bare metal object storage) in volumes that tier inactive data. Set the tiering policy on those volumes to None.

Failure to set the tiering policy to None can place the virtualized object store at risk as blocks associated with the virtual machines may be marked as cold and tiered into themselves, causing significant spikes in latency and reductions in throughput when read.

8. Sizing

Sizing the Local Tier

When considering sizing, the local 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 local tier
- Supporting WAFL metadata associated with the attached cloud tier

For most environments, a 1 : 10 :: local tier : cloud tier ratio is extremely conservative while providing significant storage savings.

Caution

Writes from the cloud tier to the local tier are disabled if local tier capacity is greater than 90%. If this occurs, blocks are read directly from the cloud tier.

Inactive Data Reporting

Inactive data reporting (IDR) is an excellent tool for determining the amount of inactive (cold) data that can be tiered from a local tier.

By default, IDR uses a 31-day cooling period to determine what data is considered inactive. The amount of cold data that is tiered is dependent on the tiering policies set on volumes. In releases earlier than ONTAP 9.8, IDR used a fixed 31-day cooling period.

- ONTAP 9.8 and later
  - IDR cooling period can be adjusted using the volume -tiering-minimum-cooling-days setting.
- ONTAP 9.7 and later
  - IDR is enabled by default on all non-FabricPool SSD local tiers.
  - IDR can be enabled on HDD local tiers using the ONTAP CLI.

Table 6 IDR behavior

FabricPool aggregate	Tiering policy	Behavior	Window
Yes	None	Reports all cold data	31 days
	Snapshot-Only	Reports all cold data	ONTAP 9.8 and later: does not report ONTAP 9.7 and earlier: 31 days **Caution** Snapshot blocks would have already been tiered by using the default two-day setting.
	Auto	Does not report IDR	Inactive data has already tiered
	All/Backup	Does not report IDR	Inactive data has already tiered
No	None	Reports all cold data	31 days
	Snapshot-Only	Reports all cold data	ONTAP 9.8 and later: -tiering-minimum-cooling-days setting ONTAP 9.7 and earlier: 31 days
	Auto	Reports all cold data	ONTAP 9.8 and later: -tiering-minimum-cooling-days setting ONTAP 9.7 and earlier: 31 days
	All/Backup	Reports all cold data	ONTAP 9.8 and later: reports all user data as cold after the first scan finishes ONTAP 9.7 and earlier: 31 days

ONTAP System Manager

IDR is displayed on the local tiers overview in ONTAP System Manager.

Figure 8 IDR in ONTAP System Manager

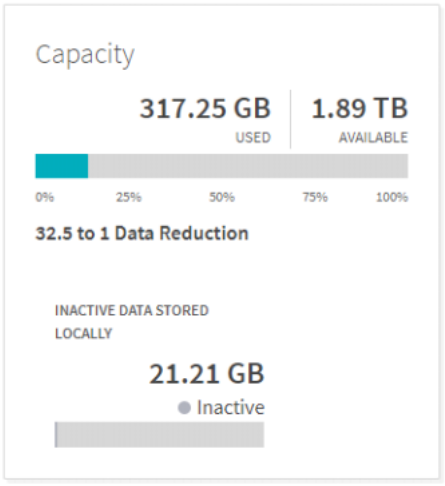

ONTAP CLI

To enable IDR on a non-FabricPool local tier, run the following command:

```
storage aggregate modify -aggregate <name> -is-inactive-data-reporting-enabled true
```

To display IDR by using the ONTAP CLI, run the following command:

```
storage aggregate show-space -fields performance-tier-inactive-user-data, performance-tier-
inactive-user-data-percent
```

To display IDR on a single volume by using the ONTAP CLI, run the following command:

```
Volume show -fields performance-tier-inactive-user-data, performance-tier-inactive-user-data-percent
```

The `performance-tier-inactive-user-data-percent` field displays what percent of the volume's total capacity is inactive, not the percent of the volume's used capacity.

Caution

Although IDR is enabled by default on all SSD local tiers, if a client workload needs 100% of system resources, it automatically turns off, resetting cooling days to zero. If this happens, IDR is not automatically turned back on.

To avoid automated process shutting off IDR in order to free up resources for other workloads, manually enable `-is-inactive-data-reporting-enabled` to `true`.

Maximum Tiering Capacity

1:10 local tier is recommended: cloud tier ratio is conservative. FabricPool continues to tier cold data to a cloud tier until the local tier reaches 98% capacity. For example, an 800TB local tier reachs 98% capacity at 784TB. Given a dataset using 5% metadata, 15.6PB could have been tiered to the cloud before reaching 784TB on the local tier.

Sizing 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
- Supporting at least 700 TCP connections

Local Tier Space Utilization

ONTAP System Manager

In ONTAP System Manager, FabricPool space utilization is displayed on the local tiers overview. Details include local tier maximum capacity, used capacity, and external tier used capacity.

Figure 9 FabricPool space utilization information

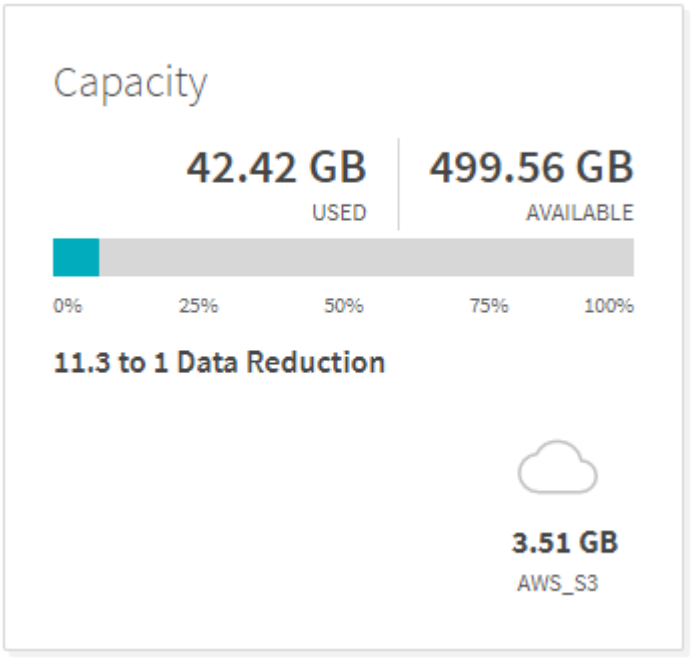

OTNAP CLI

To view FabricPool space utilization details using the ONTAP CLI, run the following command:

```
storage aggregate object-store show-space
```

Example:

```
storage aggregate object-store show-space
                                                                  License
Aggregate      Object Store Name Provider Type Used Space    Space Used%
-------------- ----------------- ------------- ------------- ------------
aggr1          aws_bucket         AWS_S3        423.3GB      41%
1 entries were displayed.
```

Volume Space Utilization

FlexVol volumes in a FabricPool local tier cannot exceed the 100TB maximum volume size for FlexVols regardless of what tier the data is located on. For example, a FlexVol with 1TB on the local tier and 99TB on the cloud tier has reached the 100TB maximum FlexVol size, even though only 1TB is stored on the local tier.

Unlike FlexVol volumes, FlexGroup volumes have virtually no capacity or file count constraints outside of the physical limits of hardware or the total volume limits of ONTAP.

If the local tier reaches 98% capacity, FabricPool stops tiering cold data to the cloud tier. If the local tier reaches 90% capacity, cold data is read directly from the cloud tier without being written back to the local tier.

FabricPool volume space utilization can be determined by using ONTAP System Manager or the ONTAP CLI.

OTNAP CLI

View FabricPool volume space utilization details using the ONTAP CLI.

```
volume show-footprint
```

Total, local tier (performance tier), and cloud tier (using the bucket name) footprints are displayed.

```
Vserver : svm_fabricpool
Volume  : project_b

Feature                                  Used Used%
-------------------------------  ---------- -----
Volume Data Footprint                16.84GB 1%
       Footprint in Performance Tier 131.7MB 1%
       Footprint in my-bucket        16.74GB 99%
Volume Guarantee                         0B 0%
Flexible Volume Metadata             429.1MB 0%
Delayed Frees                        27.60MB 0%
Total Footprint                      17.29GB 1%
```

Available License Capacity

A capacity warning is triggered when the cloud tier reaches 85% of the maximum capacity set by the capacity-based license. Tiering to the cloud tier stops when the amount of data (used capacity) stored on the third-party 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 remains accessible from the cloud tier. Cold data remains on the local tier until the licensed capacity is increased.

To view the capacity status of the FabricPool license using the ONTAP CLI, run the following command:

```
system license show-status
```

Example:

```
system license show-status
Status    License              Scope     Detailed Status
--------- -------------------  --------- ----------------------
valid
          NFS                  site      -
          CIFS                 site      -
          iSCSI                site      -
          FCP                  site      -
          SnapRestore          site      -
          SnapMirror           site      -
          FlexClone            site      -
          FabricPool           cluster   The system is using 423.3GB, and can use up to 10TB.
not-installed
          SnapVault            -         -
          SnapLock             -         -
          SnapManagerSuite     -         -
          SnapProtectApps      -         -
          V_StorageAttach      -         -
          Insight_Balance      -         -
          OCShift              -         -
          TPM                  -         -
          VE                   -         -
          DP_Optimized         -         -
not-applicable
          Cloud                -         -
          Select               -         -
20 entries were displayed.
```

To view the capacity status of the FabricPool license using ONTAP System Manager, complete the following steps:

1. Click CLUSTER.
2. Click Settings.
3. Click FabricPool License.
4. Current capacity is listed in the Current Capacity column.

Figure 10 License capacity

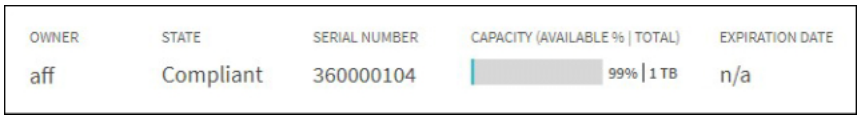

9. Data Migration

Because of the difference in ingress and egress rates, it is possible run out of space on a small local tier when attempting to migrate more data to it than it has capacity to hold. Data is usually coming into the local tier at a faster rate than it can be converted into objects and tiered out.

For example, if 50TB of data is migrated to the local tier at 2GBps, the migration will not be complete for ~7 hours. If all 50TB of data is inactive and tiered to the cloud tier at 600MBps, ~24 hours are required before all the data will be tiered to the object store.

The local tier must have enough capacity to store the data before it is tiered. Local space utilization can be determined by using ONTAP System Manager or the ONTAP CLI.

This type of scenario can happen in any environment that has been sized specifically for FabricPool aggregates, where a large data set is being migrated to a smaller local tier with the knowledge that the majority of the migrated data will be provisioned by the cloud tier.

Cloud Write

Beginning in ONTAP 9.13 (Cloud Volumes ONTAP) and ONTAP 9.14.1 (ETERNUS AX/AC/HX series), FabricPool supports Cloud Write, a feature specifically designed to avoid filling up the local tier before data can tier to the cloud tier.

For example, in Figure 11, 300TB of data is being migrated to an ONTAP environment with the intention that 100TB will be active on the local tier and 200TB will be inactive and be provisioned on the cloud tier; but the migration will fail if the local tier reaches 98% capacity before data can tier out to the cloud tier.

Figure 11 Migrating data without Cloud Write

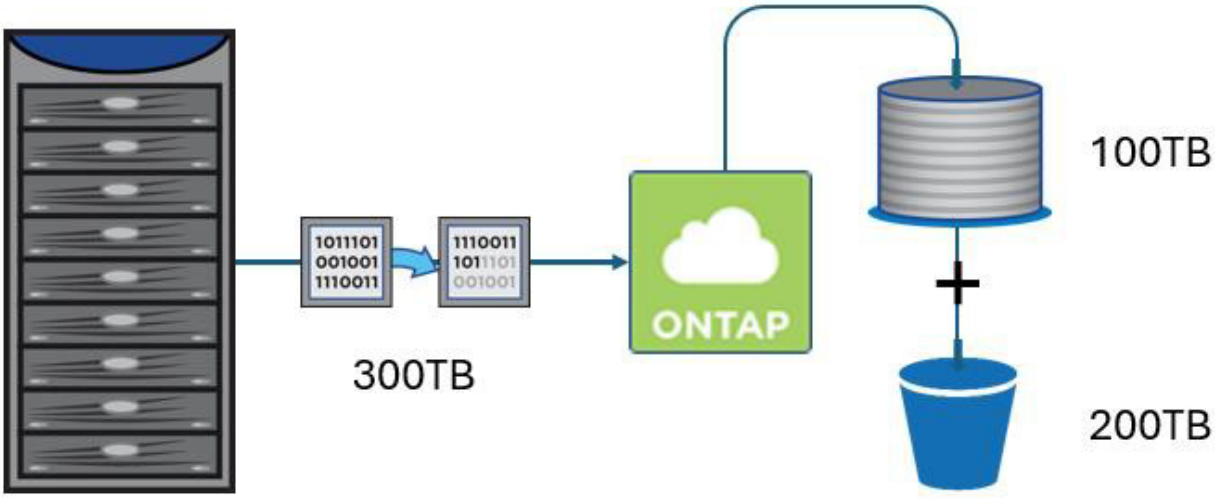

By enabling Cloud Write on the destination volume, the above scenario will be avoided. Incoming data bypasses the local tier, is written to a 60GB transfer log, and immediately begins tiering to the cloud tier as rapidly as possible. Ingress of client writes automatically scales to match the egress of data to the cloud tier, and unlike normal FabricPool PUT operations, Cloud Write PUTs are prioritized in order enhance performance.

Figure 12 Migrating data with Cloud Write

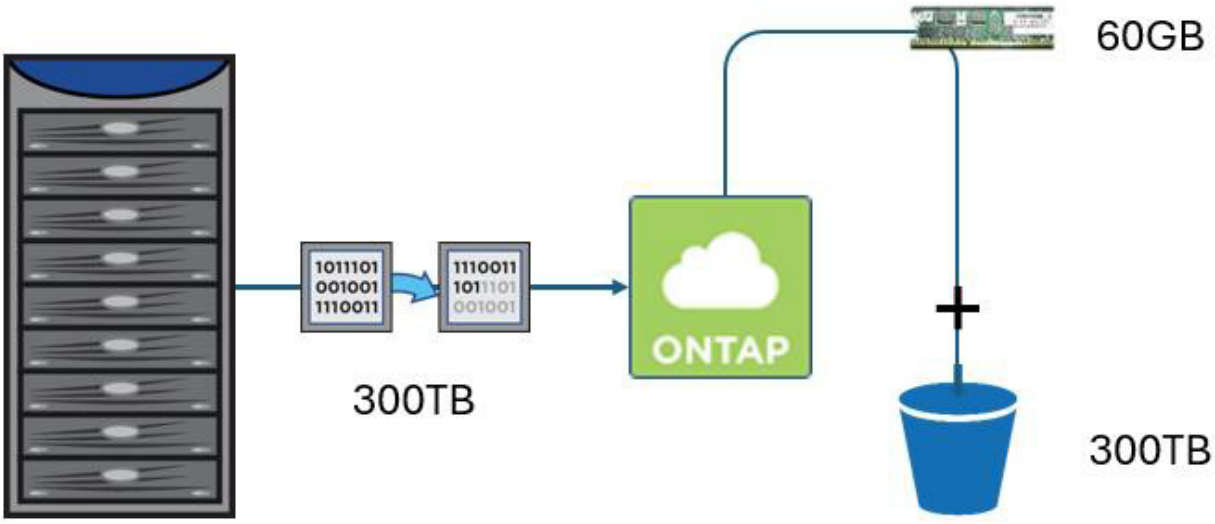

Caution

Cloud write only supports NFS-based migrations. Data migrations using other protocols will be written to the local tier as normal.

To enable Cloud Write on a volume, run the following command:

```
volume modify -vserver <name> -volume <name> -tiering-policy all -is-cloud-write-enabled
true
```

Caution

Advanced privilege level is required.

Cloud Write is intended for large data migrations and should be turned off after the migration is complete. If not disabled, Cloud Write, and the All tiering policy, will have a negative impact on workloads running on the volume.

Turning off Cloud Write and changing the tiering policy to Auto or Snapshot-Only will allow random client reads to the cloud tier to write data back to the high-performance local tier where active data belongs.

Migration Options

Using the All Volume Tiering Policy

Whether it is used in conjunction with Cloud Write or not, when the All volume tiering policy is used on a volume it will tier data as quickly as possible. Incoming data will take advantage of inline storage efficiencies, but the data will be tiered out of the local tier before additional background storage efficiencies such as TSSE can be applied.

Because all data in the volume is marked as inactive and tiered, when the migration is over and the volume tiering policy is changed to Auto or Snapshot-Only with the intention of serving active data, previously tiered data will need to be read from the object store before it is written back to the local tier.

Using the Auto Volume Tiering Policy

It is preferable to use the Auto volume tiering policy when the migration destination volume has enough capacity to hold all the data being migrated. Tiered data will be more storage efficient and active data will remain on the local tier, significantly improving performance and reducing network traffic.

By temporarily adjusting the volume's tiering minimum cooling days value to 2-days, storage and network efficiencies can be retained without needing to provision the migrated data on the local tier for a month before it would tier normally.

For example:

1. Set the volume to use the Auto tiering policy and a tiering minimum cooling days of 2.
2. Migrate data to the volume.
3. Wait for the inactive data to tier out.
4. Change the tiering minim cooling days to 31.

The migrated data will need to remain on the local tier for at least two days, but in the long term the data will be more efficient and require less network traffic than it would have using the All tiering policy.

ETERNUS AX series All-Flash Arrays, ETERNUS AC series All-Flash Arrays,
ETERNUS HX series Hybrid Arrays FabricPool Best Practices ONTAP 9.14.1

P3AG-5662-04ENZ0

Date of issuance: March 2025
Issuance responsibility: Fsas Technologies Inc.

Fsas Technologies Inc.