



Technical Report

FabricPool best practices

ONTAP 9.14.1

John Lantz, NetApp

January 2024 | TR-4598

Abstract

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

TABLE OF CONTENTS

| | |
|---|-----------|
| Overview | 5 |
| Primary use cases | 6 |
| Reclaim capacity on primary storage (Auto, Snapshot-Only, or All) | 6 |
| Shrink the secondary storage footprint (All)..... | 9 |
| Requirements | 10 |
| Platforms | 10 |
| Intercluster LIFs..... | 10 |
| Internet Protocol version..... | 11 |
| Transmission Control Protocol (TCP) connections | 11 |
| Volumes | 11 |
| Cloud Tiering license | 12 |
| Special configurations | 12 |
| Certificate authority certification..... | 13 |
| Architecture | 14 |
| Block temperature | 14 |
| Object creation | 15 |
| Data movement | 15 |
| Object storage | 20 |
| Configuration | 23 |
| Create a bucket/container | 23 |
| Add a cloud tier to ONTAP | 24 |
| Attach a cloud tier to a local tier..... | 28 |
| Volume tiering policies..... | 31 |
| Cloud retrieval | 33 |
| Volume tiering minimum cooling days | 34 |
| FabricPool Mirror | 35 |
| MetroCluster | 39 |
| Security | 41 |
| Interoperability | 42 |
| StorageGRID..... | 43 |
| Performance | 45 |
| Network connections | 45 |
| Object store profiler | 45 |

| | |
|---|-----------|
| Sequential read performance | 46 |
| Aggressive read-ahead | 46 |
| PUT throttling..... | 47 |
| SnapMirror concurrency | 47 |
| Low performance environments | 47 |
| Virtualized object storage | 48 |
| Sizing..... | 49 |
| Sizing the local tier | 49 |
| Sizing the cloud tier | 51 |
| Local tier space utilization | 51 |
| Volume space utilization..... | 52 |
| Available license capacity..... | 53 |
| Data migration..... | 54 |
| Cloud Write..... | 54 |
| Migration options | 56 |
| Data tiering within Cloud Volumes ONTAP..... | 56 |
| NetApp Private Storage for AWS | 57 |
| Where to find additional information | 58 |
| Version history..... | 59 |
| Contact us | 60 |

LIST OF TABLES

| | |
|--|----|
| Table 1) SnapMirror behavior..... | 17 |
| Table 2) Default unreclaimed space thresholds..... | 21 |
| Table 3) NetApp interoperability..... | 42 |
| Table 4) Third-party interoperability..... | 43 |
| Table 5) FabricPool byte-ranged GET sizes..... | 46 |
| Table 6) IDR behavior..... | 50 |

LIST OF FIGURES

| | |
|---|----|
| Figure 1) Before and after FabricPool..... | 5 |
| Figure 2) Reclaiming space with the auto-volume tiering policy..... | 6 |
| Figure 3) Reclaiming space with the Snapshot-Only volume tiering policy..... | 7 |
| Figure 4) Reclaiming space with the All volume tiering policy..... | 8 |
| Figure 5) Using the All volume tiering policy with secondary storage..... | 9 |
| Figure 6) Changing the volume tiering policy during a volume move..... | 19 |
| Figure 7) Possible cloud tier-to-local tier relationships in ONTAP 9.7..... | 23 |
| Figure 8) FabricPool containing one local tier and two cloud tiers..... | 36 |
| Figure 9) An example of the licensed capacity limit with a 100TB FabricPool license..... | 36 |
| Figure 10) An example of the licensed capacity limit with a 200TB FabricPool license..... | 37 |
| Figure 11) MetroCluster plus FabricPool..... | 40 |
| Figure 12) IDR in ONTAP System Manager..... | 50 |
| Figure 13) FabricPool space utilization information..... | 52 |
| Figure 14) License capacity..... | 54 |
| Figure 15) Migrating data without Cloud Write..... | 55 |
| Figure 16) Migrating data with Cloud Write..... | 55 |

Overview


FabricPool, first available in ONTAP 9.2, is a data fabric powered by NetApp 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. It delivers the benefits of cloud economics by tiering to public clouds such as Alibaba Cloud Object Storage Service, Amazon S3, Google Cloud Storage, IBM Cloud Object Storage, and Microsoft Azure Blob Storage, as well as to private clouds such as NetApp StorageGRID®.

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.

- ONTAP supports FabricPool on SSD and HDD local tiers (also known as storage aggregates in the ONTAP CLI). Flash Pool aggregates are not supported.
- NetApp ONTAP Select supports FabricPool. NetApp recommends using all-SSD FabricPool local tiers.
- NetApp Cloud Volumes ONTAP supports data tiering with Amazon S3, Google Cloud Storage, and Microsoft Azure Blob Storage.

Figure 1) Before and after FabricPool.



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 NetApp 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 [NetApp SnapMirror[®]](#) and [NetApp SnapVault[®]](#).

Reclaim capacity on primary storage (Auto, Snapshot-Only, or All)

Auto

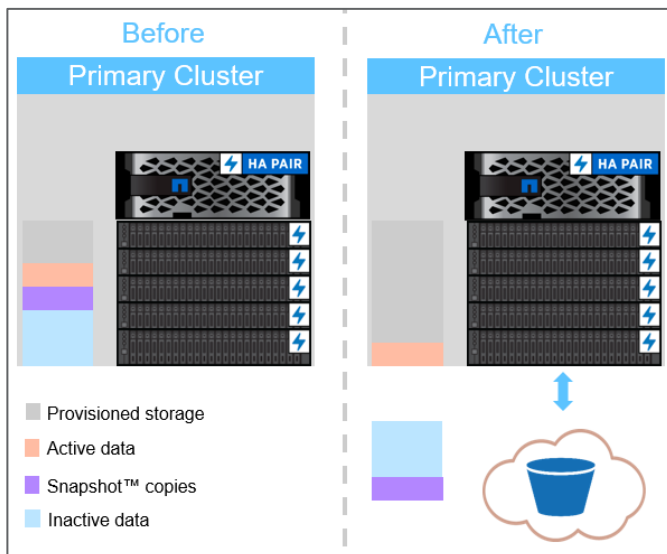
The majority of inactive (cold) data in storage environments is associated with unstructured data, accounting for more than 50% of total storage capacity in many storage environments.

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.

First available in ONTAP 9.4, the auto volume tiering policy shown in Figure 2 moves all cold blocks in the volume, not just blocks associated with NetApp 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.



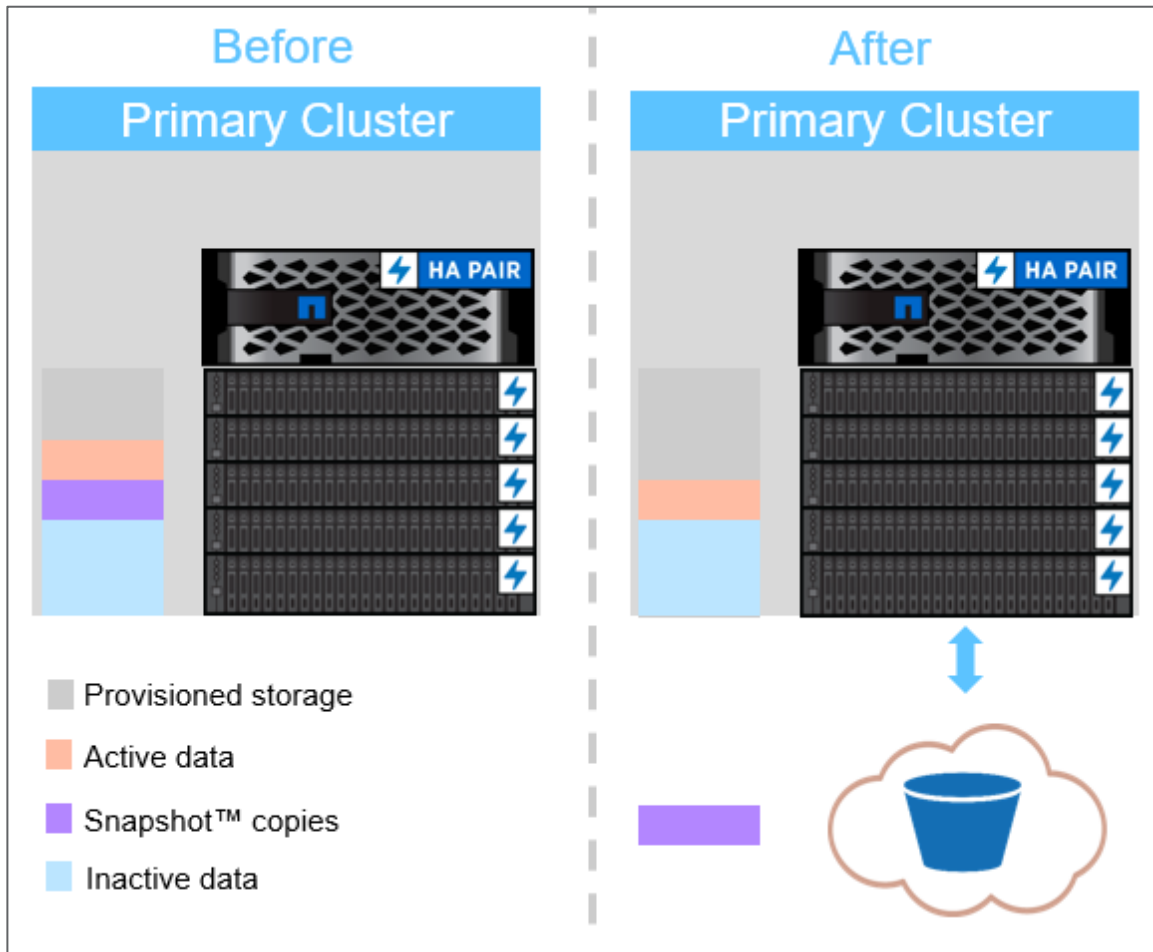
Snapshot-Only

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.

Note: 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.

Note: NetApp does not recommend using 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 NetApp SnapMirror (disaster recovery) or NetApp 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 disk 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, or in a software defined environment using ONTAP Select. You can tier data using FabricPool anywhere ONTAP can be deployed.

Requirements

FabricPool requires ONTAP 9.2 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.

Although installation and use of certificate authority (CA) certificates are recommended best practices, beginning in ONTAP 9.4, installation of CA certificates is not required for StorageGRID.

Platforms

- **AFF**
- **FAS**

FabricPool is supported on all platforms capable of running ONTAP 9.2, except for the following:

- FAS8020
- FAS2554, FAS2552, FAS2520

- **ONTAP Select**

Note: NetApp recommends using all-SSD FabricPool local tiers.

- **Cloud tiers**

- Alibaba Cloud Object Storage Service (Standard, Infrequent Access)
- 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)
- IBM Cloud Object Storage (Standard, Vault, Cold Vault, Flex)
- Microsoft Azure Blob Storage (Hot and Cool)
- NetApp StorageGRID 10.3 and later

Note: Glacier Flexible Retrieval and Glacier Deep Archive are not supported.

- **Data tiering**

- Amazon FSx for NetApp ONTAP
- Azure NetApp Files
- NetApp Cloud Volumes for Google Cloud
- NetApp Cloud Volumes ONTAP for AWS
- NetApp 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. NetApp recommends creating an intercluster LIF on additional HA pairs to seamlessly attach cloud tiers to local tiers on those nodes as well.

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

Note: 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.

Note: FabricPool Mirror will double this number as it connects to two different endpoints.

Volumes

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

```
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 provides 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 (such as Amazon S3) as cloud tiers for AFF and FAS systems. 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 pay-as-you-go subscriptions from cloud provider marketplaces, or 2-, 12-, 24-, 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.

For more information, see: [Applying Cloud Tiering licenses to clusters in special configurations](#).

Note: Legacy 12-month FabricPool licenses are required for NetApp MetroCluster, FabricPool Mirror, and dark site or other air-gapped environments which are not yet supported by Cloud Tiering.

Certificate authority certification

When FabricPool uses StorageGRID or other private clouds such as some IBM Cloud Object Storage 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.

Note: Beginning in ONTAP 9.4, [CA certificates are no longer required](#). However, using signed certificates from a third-party certificate authority remains 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
```

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 discs: 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.

Note: 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
```

Note: 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.

Note: 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>
```

Note: 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](#) 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](#) 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 * | All * | Cloud* |
| All | None | Local |
| None | Auto | Local |
| None | Snapshot-Only | Local |
| None | All | Cloud |
| None | None | Local |

*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.

Beginning in ONTAP 9.6, 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.

Note: 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)
- Unsynchronized FabricPool Mirror buckets
- When tiering to Amazon S3; source and destination aggregates use a different Amazon naming format:
 - FabricPool aggregates created in releases earlier than ONTAP 9.5 use the old format.
 - FabricPool aggregates created in ONTAP 9.5 and later use the current format.

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.

Note: 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 tier 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.

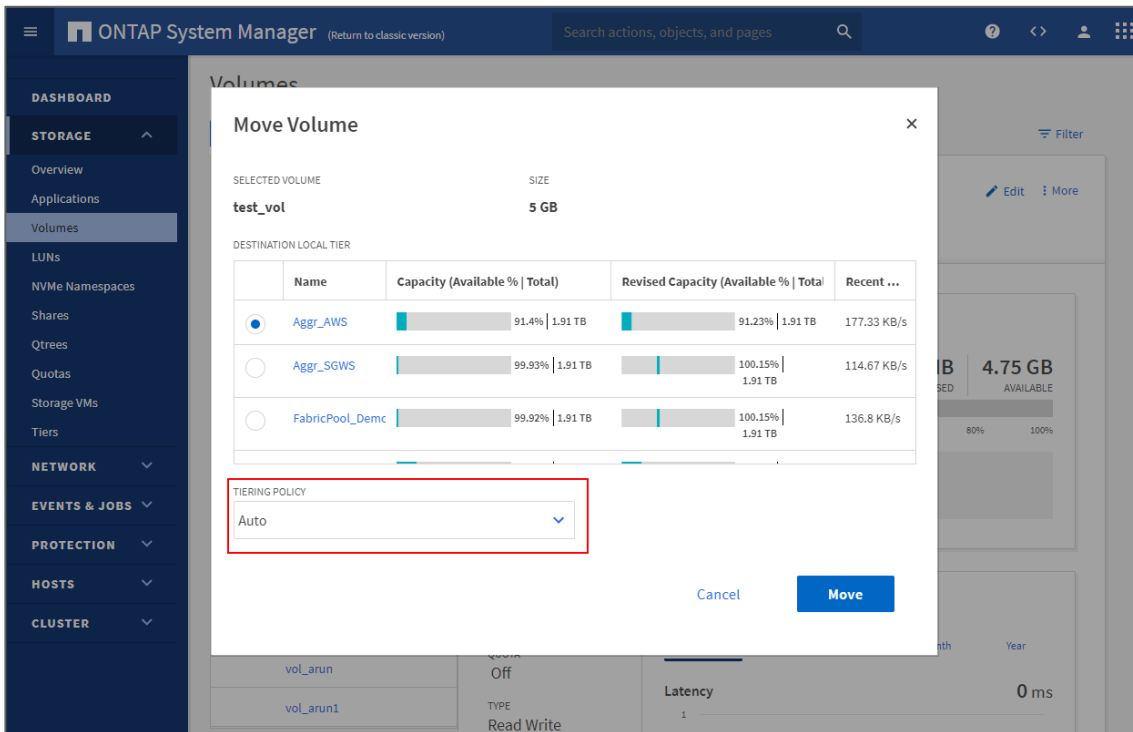
```
vol move start -vserver <name> -volume <name> -destination-aggregate <name> -tiering-policy none
```

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. NetApp recommends using tiering settings on the parent FlexVol that are either equal to or less aggressive than 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.

FabricPool cloud tiers

FabricPool supports object stores from multiple providers (Alibaba, Amazon, Google, IBM, Microsoft, NetApp StorageGRID, and so on) as cloud tiers.

More than one type of cloud tier can be used in a cluster. Usually, one cloud tier is attached to each local tier, but, beginning in ONTAP 9.7, FabricPool Mirror enables the attachment of two cloud tiers to a single local tier.

ONTAP S3

Beginning in ONTAP 9.8, ONTAP supports tiering to buckets created 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).

NetApp recommends using StorageGRID, NetApp's premier object store solution, when tiering more than 300TB of inactive data.

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 9.3 and earlier | ONTAP 9.4–9.7 | ONTAP 9.8 and later | Cloud Volumes ONTAP |
|------------------------------|-----------------------|---------------|---------------------|---------------------|
| Alibaba Cloud Object Storage | n/a | 15% | 20% | n/a |
| Amazon S3 | 0% | 20% | 20% | 30% |
| Google Cloud Storage | n/a | 12% | 20% | 35% |
| IBM Cloud Object Storage | n/a | 14% | 20% | n/a |
| Microsoft Azure Blob Storage | n/a | 15% | 25% | 35% |
| NetApp ONTAP S3 | n/a | n/a | 40% | n/a |
| StorageGRID | 0% | 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. Examples include Amazon's Standard-IA and Azure Blob Storage's Cool.

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%)
```

Note: Advanced privilege level is required.

Note: 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.

Note: Third-party deduplication has not been qualified by NetApp.

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, AFF volumes are created using adaptive compression by default. (-storage-efficiency-mode default)
TSSE must be manually enabled on volumes. (-storage-efficiency-mode efficient)

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.


Although FabricPool supports the attachment of two buckets (cloud tiers) per local tier using [FabricPool Mirror](#), attaching a single cloud tier is more common.

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. NetApp does not recommend attaching a single cloud tier to local tiers in multiple clusters.

Note: 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.



The screenshot shows a 'Create Bucket' dialog box. The title bar reads 'Create Bucket'. Below it is a section titled 'Bucket Details' with a question mark icon. There are two input fields: 'Name' containing the text 'fabricpool789' and 'Region' with a dropdown menu currently showing 'us-east-1'. At the bottom right of the dialog are two buttons: 'Cancel' and 'Save'.

Note: In releases earlier than StorageGRID 11.1, creating a bucket required using a third-party S3 client such as an S3 browser.

Note: 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.

ONTAP S3

You can find instructions for creating buckets in ONTAP S3 in [Provisioning Object Storage with System Manager](#).

Other object store providers

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

- [Alibaba Cloud Object Storage Service](#)
- [Amazon S3](#)
- [Google Cloud Storage](#)
- [IBM Cloud Object Storage](#)
- [Microsoft Azure Blob Storage](#)

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.

Note: 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
- FabricPool Mirror
- MetroCluster
- Dark site or otherwise air gapped environments

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`)

Note: 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.

Note: 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 or as a FabricPool Mirror.

Note: 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. (Using [FabricPool Mirror](#), you can attach a different cloud tier.)

8. Click Save.

Add Cloud Tier ✕

NAME

SERVER NAME (FQDN)

SSL

PORT

ACCESS KEY ID

SECRET KEY

CONTAINER NAME ?

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
-ip-space 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.

S3 compatible providers

Customers who want to use object stores that are not officially supported as a cloud tier can do so using `-provider-type S3-Compatible`. Customers must test and confirm that the object store meets their requirements.

NetApp does not support, nor is liable for any issues arising from any third-party Object Store Service, specifically where it does not have agreed support arrangements with the third party with whom the product originated. It is acknowledged and agreed that NetApp shall not be liable for any associated damage or otherwise be required to provide support on that 3rd party product.

Certificate authority certificate validation

CA certificates associated with private cloud object stores, such as StorageGRID and some IBM Cloud Object Storage 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 not recommended, but it is possible beginning in ONTAP 9.4.

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.

Note: Turning off certificate validation is not recommended.

8. Click Save.

SSL

Object store certificate [?](#)

CERTIFICATE

Copy the contents of the signed certificate, including the "BEGIN" and "END" tags, and then paste the contents in this box.

COMMON NAME (OPTIONAL)

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
-ipospace 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.

Note: 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. (Using [FabricPool Mirror](#), a different cloud tier can be 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 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.

Note: 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.

Note: 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. (Using [FabricPool Mirror](#), you can attach a different cloud tier.)

Attach Cloud Tiers

LOCAL TIER
aff_01_aggr1

ATTACH AS PRIMARY
AWS_GovCloud

Update Tiering Policy [Considerations](#)

i Displays the volumes of the selected local tier.

| <input type="checkbox"/> Volumes | Storage VM | Inactive Data Capacity | Tiering |
|---|---------------------|------------------------|---------|
| <input type="checkbox"/> OraDev_Vol | AFF_SAN_DEFAULT_SVM | 32.46 GB | None |
| <input type="checkbox"/> vol_sanluns01dev_02 | svm_sjb_sanluns01 | - | None |
| <input type="checkbox"/> vol_thin_1000G | svm_sjb_sqldb01prod | - | None |
| <input type="checkbox"/> vol_sanluns01prod_01 | svm_sjb_sanluns01 | - | None |

Mirror cloud tier

Save Cancel

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
```

Note: 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. (Using [FabricPool Mirror](#), you can attach a different cloud tier.)

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 (FAS 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.

Note: 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 might occur if cluster LIFs resources become saturated. To avoid this, NetApp recommends using 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.

Note: 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 [tiering-minimum-cooling-days](#). (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 [tiering-minimum-cooling-days](#). (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.

- In releases earlier than ONTAP 9.6, the Backup volume tiering policy functioned the same as the All policy with the exception that the Backup policy can only be set on data protection volumes (destination targets).

Note: NetApp does not recommend using 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.

Note: 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 from 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) |

Note: 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.

Note: 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>
```

Note: Advanced privilege level is required.

Note: 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.

FabricPool Mirror

You can use more than one type of cloud tier in a cluster, but usually a single cloud tier is attached to a single local tier. Beginning in ONTAP 9.7, FabricPool Mirror allows the attachment of two cloud tiers to a single local tier, creating additional options for data availability and movement.

When using FabricPool Mirror, data is mirrored across two buckets. During bucket synchronization, data must be read from the pre-existing primary bucket and written to the secondary bucket. This synchronization is necessary to achieve a mirrored state between the two buckets.

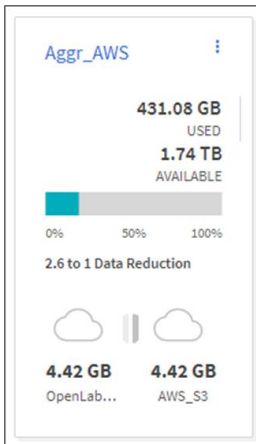
When both buckets are in a mirrored state, newly tiered data is synchronously tiered to both buckets. Because data is being tiered to two buckets synchronously, the effective throughput is half of standard single-bucket tiering. For example, if PUT operations to a single bucket take place at 600MBps, PUT operations to mirrored buckets in a FabricPool Mirror deployment take place at 300MBps.

Under normal circumstances, all GET operations take place from the primary bucket. Only if connectivity is interrupted to the primary bucket will GET operations take place from the secondary bucket.

If connectivity is lost to either bucket, tiering is temporarily suspended until connectivity is established.

Note: Although essential for FabricPool with NetApp MetroCluster, FabricPool Mirror is a stand-alone feature that does not require MetroCluster to use.

Figure 8) FabricPool containing one local tier and two cloud tiers.



When adding FabricPool Mirror to an existing FabricPool, data previously tiered to the original cloud tier read from the primary bucket is written to the newly attached secondary bucket.

Note: Public cloud charges for reads and writes apply to both buckets as normal.

Licensed capacity


During synchronization, the FabricPool capacity license applies equally to both cloud tiers.

For example, a cluster with a 100TB license and 75TB of data in bucket A is able to mirror 75TB of data to bucket B. A total of 150TB has been tiered even though the FabricPool license is only for 100TB of capacity.

After both tiers are mirrored, licensed capacity applies as normal and tiering to the cloud tier stops if the amount of data (used capacity) tiered to both cloud tiers is greater than the licensed capacity.

For example, a cluster with a 100TB FabricPool license and 75TB of tiered data in bucket A and 75TB of tiered data mirrored to bucket B uses 150TB of capacity (Figure 9). The cluster must increase the FabricPool license to >150TB to tier additional data.


Figure 9) An example of the licensed capacity limit with a 100TB FabricPool license.



Newly tiered data is synchronously tiered to both cloud tiers provided it is supported by the licensed capacity.

For example, a cluster with a 200TB license and 75TB of tiered data in bucket A and 75TB of tiered data mirrored to bucket B uses 150TB of capacity (Figure 10). FabricPool continues to tier data to both cloud tiers until both buckets contain 100TB of tiered data.

Figure 10) An example of the licensed capacity limit with a 200TB FabricPool license.



Multicloud solutions

Object stores are generally designed to provide 99.999999999% (11 9s) durability and 99.99% (4 9s) availability. Given 525,600 minutes in a year, 99.99% availability allows for 52.56 minutes of unavailability per year.

For customers who need better availability, FabricPool Mirror can be used to tier data to multiple cloud vendors for an additional level of resiliency as the likelihood that multiple cloud providers experiences outages at the same time is extremely rare. Using a single cloud vendor, but different regions or availability zones (for example, Google's europe-west3-a and europe-west6-a) as an alternative to using multiple providers.

During bucket synchronization, data must be read from the pre-existing primary bucket and written to the secondary bucket. Reads from public clouds are subject to egress fees.

Note: When applicable, after a synchronous mirror has been established between two cloud tiers, the total amount of data tiered to both cloud tiers is applied to the FabricPool license.

Enhanced data mobility

FabricPool has always allowed for data mobility, but in releases earlier than ONTAP 9.7, a volume move was required to nondisruptively move from one FabricPool to another. Although a simple process, volume moves require that the destination local tier has enough capacity to hold the volume being moved.

FabricPool Mirror can be used to change the cloud tier attached to a local tier without having to perform a volume move by:

1. Adding a secondary cloud tier to the local tier.
2. Waiting for a synchronous mirror to be established between the primary and secondary cloud tiers.
3. Swapping cloud tiers so the primary and secondary cloud tier relationships change. The secondary cloud tier becomes the new primary cloud tier and the original primary cloud tier becomes the new secondary cloud tier.
4. Deleting the FabricPool Mirror.

ONTAP System Manager

Attach

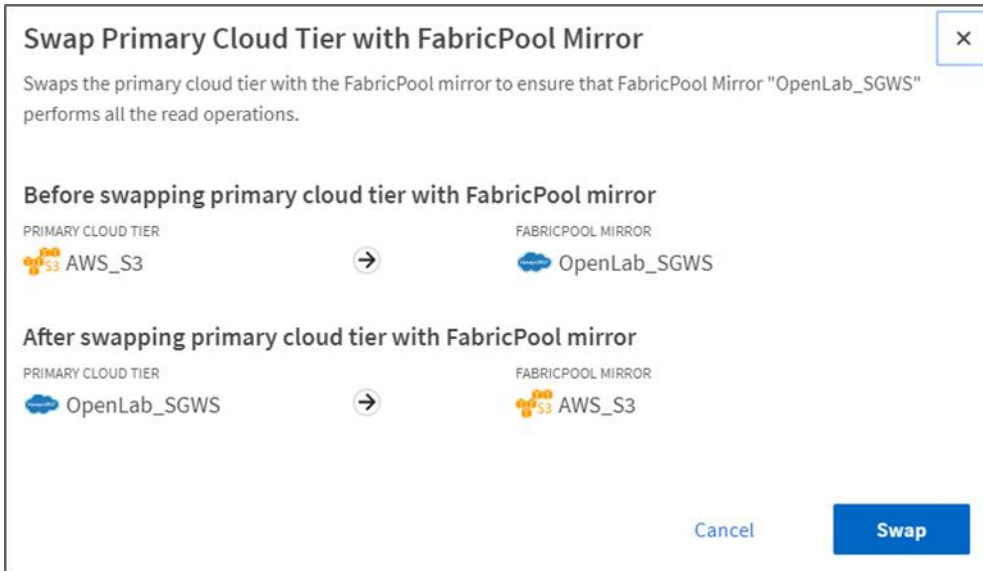
To attach an additional cloud tier, a FabricPool Mirror, to a local tier by using ONTAP System Manager, complete the following steps:

1. Launch ONTAP System Manager.
2. Click Attach FabricPool Mirror.
3. Select a cloud tier to use as the secondary cloud tier.
4. Click Save.

Swap

To swap cloud tiers so the primary and secondary cloud tier relationships change in a FabricPool Mirror using ONTAP System Manager, complete the following steps:

1. Launch ONTAP System Manager.
2. Click Storage.
3. Click Tiers.
4. Click the name of the local tier you wish to remove the FabricPool Mirror from.
5. Click More.
6. Click Swap Cloud Tiers.



7. Click Swap.

Delete

To remove the FabricPool Mirror using ONTAP System Manager, complete the following steps:

1. Launch ONTAP System Manager.
2. Click Storage.
3. Click Tiers.
4. Click the name of the local tier you wish to remove the FabricPool Mirror from.

5. Click More.
6. Click Delete FabricPool Mirror.
7. Click Delete.

ONTAP CLI

Attach

To attach an additional cloud tier, a FabricPool Mirror, to a local tier (storage aggregate) using the ONTAP CLI, run the following commands:

```
storage aggregate object-store mirror -aggregate <aggregate name> -name <object-store-name-2>
```

Swap

To swap cloud tiers so the primary and secondary cloud tier relationships change in a FabricPool Mirror using the ONTAP CLI, run the following commands:

```
storage aggregate object-store modify -aggregate <aggregate name> -name <object-store-name-2> -mirror-type primary
```

Delete

To remove the FabricPool Mirror by using the ONTAP CLI, run the following commands:

```
storage aggregate object-store unmirror -aggregate <aggregate name> -name <object-store-name-1>
```


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. Both FC and Ethernet (IP) MetroCluster configurations are used by thousands of enterprises worldwide for high availability, zero data loss, and nondisruptive operations both within and beyond the data center.

Beginning in ONTAP 9.7, MetroCluster's continuous availability extends to FabricPool using [FabricPool Mirror](#).

When using FabricPool Mirror, data is mirrored across two buckets so Cluster As local tiers can be connected to cloud tier A1 (primary) and cloud tier A2 (mirror) and Cluster Bs local tiers can be connected to cloud tier B1 (primary) and cloud tier B2 (mirror), as shown in Figure 16.

Figure 11) MetroCluster plus FabricPool.



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.

Unmirrored aggregates/local tiers

MetroCluster expects all local tiers, both traditional aggregates, and those using FabricPool, to be mirrored. Unmirrored aggregates in MetroCluster environments do not need to use FabricPool Mirror, but they generate messages warning that the aggregate is not mirrored, and when using FabricPool, that they are missing a FabricPool Mirror.

To view the MetroCluster error messages, run the following command:

```
metrocluster check show
```


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 NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE). Neither NSE, NVE, nor NAE 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 NVE/NAE remain encrypted when moved to the cloud tier. Client-side encryption keys are owned by ONTAP.

All objects not encrypted using NVE/NAE 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.

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

Disabling cloud tier encryption

Beginning in ONTAP 9.7, encrypting cold data at rest is no longer required. 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. NetApp highly recommends using client-side NVE or NAE 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-encryption false
```

Note: Advanced privilege level is required.

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) NetApp interoperability.

| Focus | Supported | Not supported |
|--------------------------|--|--|
| Cloud tier | <ul style="list-style-type: none"> ONTAP S3 9.8+ StorageGRID 10.3+ | ONTAP S3 in multiprotocol NAS volumes |
| Data protection | <ul style="list-style-type: none"> Cloud Backup Service MetroCluster MetroCluster SDS SnapMirror (XDP and DP) SnapMirror Synchronous SnapVault (XDP and DP) SVM-DR SVM Migrate StorageGRID replication and erasure coding <p>Note: For best results, use replication with StorageGRID 11.2 or later and erasure coding with StorageGRID 11.3 or later.</p> | <ul style="list-style-type: none"> 7-Mode Data Transition Using SnapMirror 7-Mode Transition Tool (7MTT) DP_Optimized license (DPO) Object versioning Secure Purge SMTape NetApp SnapLock® technology Cascading SnapMirror relationships using the All (or Backup) tiering policy. StorageGRID ILM policies other than replication and erasure coding StorageGRID Compliance buckets NetApp SyncMirror® technology Tamperproof Snapshot copies WORM |
| Encryption | <ul style="list-style-type: none"> NetApp Volume Encryption NetApp Storage Encryption Server-side encryption (AES-256) TLS 1.2 | – |
| Storage efficiency | <ul style="list-style-type: none"> Inline deduplication Inline compression Compaction Aggregate inline deduplication (local tier only) | – |
| Storage virtualization | – | NetApp FlexArray® technology |
| Quality of service (QoS) | QoS maximums (ceiling) | QoS minimums (floors) |
| Additional features | BlueXP Storage Class Lifecycle Management | <ul style="list-style-type: none"> Auto Balance Aggregate Flash Pools |

Table 4) Third-party interoperability.

| Focus | Supported | Not supported |
|-----------------|---|---|
| Cloud tier | <ul style="list-style-type: none"> 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, Archive) IBM Cloud Object Storage (including Cleversafe and SoftLayer) S3 in ONTAP 9.8 and later Microsoft Azure Blob Storage (Hot and Cool) StorageGRID 10.3+ | <ul style="list-style-type: none"> Alibaba Archive Amazon S3 Glacier Flexible Retrieval Amazon S3 Glacier Deep Archive Azure Archive IBM Archive |
| Data protection | Amazon's 99.999999999% multi-region durability | ILM policies |
| Encryption | Server-side encryption (AES-256) TLS 1.2 | – |

StorageGRID

Performance

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 the system.

NetApp recommends provisioning enough StorageGRID nodes to meet or exceed capacity and performance requirements.

Load balancing

StorageGRID includes an optional load balancer called the API Gateway Node. Although the API Gateway Node is low cost and requires no configuration, it is not as robust as other load-balancer options.

Grids that act as cloud tiers for applications that need immediate access to data should consider using appliance-based load balancers, such as the SG1000.

For additional information, see [TR-4626: StorageGRID Load Balancer Options](#).

Note: With StorageGRID 11.3, you can use high-availability (HA) groups to provide highly available data connections for S3 client. You can also use HA groups to provide highly available connections to the Grid Manager and the Tenant Manager. HA groups use virtual IP addresses (VIPs) to provide active-backup access to Gateway Node or Admin Node services.

Storage efficiency

The ONTAP volume-level storage efficiencies such as compression, deduplication, and compaction are preserved when moving data to the cloud tier. NetApp recommends disabling stored object compression in StorageGRID.

Security

All data encrypted by ONTAP NVE/NAE remains encrypted when moved to the cloud tier. Client-side encryption keys are owned by ONTAP.

All objects not encrypted by using ONTAP NVE/NAE are automatically encrypted by StorageGRID using AES-256-GCM encryption. No additional encryption is necessary. NetApp recommends disabling stored object encryption in StorageGRID.

Data protection

StorageGRID uses two-copy replication as the default ILM rule for data protection. Beginning with StorageGRID 11.2 and later, intrasite erasure coding using a 2+1 scheme is the recommended best practice for cost efficient data protection.

Erasure coding uses more CPU, but significantly less storage capacity, than replication. 4+1 and 6+1 schemes use even less capacity than 2+1, but at the cost of lower throughput and less flexibility when adding storage nodes during grid expansion.

Note: Single copy replication is not recommended due to lowered system availability and data durability. Geographically dispersed erasure coding such as 4+2 or 6+3 over multiple physical sites is not recommended due to additional latencies.

Information lifecycle management

FabricPool supports StorageGRID's information lifecycle management (ILM) policies for data replication and erasure coding to protect cloud tier data from failure.

Object tagging

Beginning in ONTAP 9.8, FabricPool also supports advanced ILM rules such as filtering based on object tags when StorageGRID is used as the cloud tier. ILM rules can be used in conjunction with tags to direct objects to specific nodes, change data protection policies (from replication to erasure coding), etc.

ONTAP supports up to four tag key=value pairs per volume.

To add tags to a volume using the ONTAP CLI, run the following commands:

```
volume modify <name> -tiering-object-tags <key1=value1>,<key2=value2>, <key3=value3>, etc.
```

Note: ILM can include 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. For example, FabricPool cloud tier data must not be expired/deleted or moved out of the bucket to other locations (Archive, Glacier, and so on).

Consistency controls

StorageGRID's [consistency controls](#) affects how the metadata that StorageGRID uses to track objects is distributed between nodes and the availability of objects for client requests. NetApp recommends using the default, read-after-new-write, consistency control for buckets used as FabricPool targets.

Note: Do not use the available consistency control for buckets used as FabricPool targets.

Virtualized nodes

In addition to performance-optimized hardware appliances, you can deploy StorageGRID nodes can be deployed as virtual machines (VMs) or Docker containers. Do not host virtualized nodes in ONTAP 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 VMs can be marked as cold and tiered into themselves, causing significant spikes in latency and reductions in throughput when read.

Performance

Network connections

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

Although direct connections provide better performance and lower data transfer charges, they are not required by FabricPool. Because performance can be significantly better when using direct connections, doing so using at least 10Gbps is the recommended best practice for FabricPool.

- [Alibaba Cloud Object Storage Service \(Express Connect\)](#)
- [Amazon S3 \(Direct Connect\)](#)
- [Google Cloud Storage \(Cloud Interconnect\)](#)
- [IBM Cloud Object Storage \(Direct Link\)](#)
- [Microsoft Azure Blob Storage \(ExpressRoute\)](#)

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, NetApp recommends using enough nodes to meet or exceed 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 the 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>
```

Note: Advanced privilege level is required.

2. View the results.

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

Note: 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 4MB 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.

Note: 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.

Note: 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:

```
volume modify -vserver <name> -volume <name> -aggressive-readahead-mode <file_prefetch>
```

Note: 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.

Note: 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]
```

Note: 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.

Note: If you are using more than one intercluster LIF on a node with different routing, NetApp recommends placing 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

SAN 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.

NetApp recommends 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.

Warning: Even a short disruption can be disastrous to production applications using SAN protocols. NetApp recommends using private networks and private clouds, like ONTAP S3 or StorageGRID object stores, when tiering data that is accessed by 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.

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.

Note: 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

First available in ONTAP 9.4, 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 9.6 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.
- ONTAP 9.4 and later
 - IDR is enabled by default on FabricPool local tiers in ONTAP 9.4.
 - IDR cannot be enabled in environments that do not support FabricPool: for example, root local tiers, local tiers with space provisioning, Flash Pools, and so on.


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 Note: 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 12) 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.

Note: Although IDR is enabled by default on all SSD local tiers (ONTAP 9.6 and later), 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

NetApp's recommended 1:10 local tier: 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 reaches 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 13) FabricPool space utilization information.



ONTAP 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
```

| Aggregate | Object Store Name | Provider | Type | Used Space | License 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.

ONTAP 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 14) License capacity.

| OWNER | STATE | SERIAL NUMBER | CAPACITY (AVAILABLE % TOTAL) | EXPIRATION DATE |
|-------|-----------|---------------|---|-----------------|
| aff | Compliant | 360000104 |  99% 1 TB | n/a |

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.

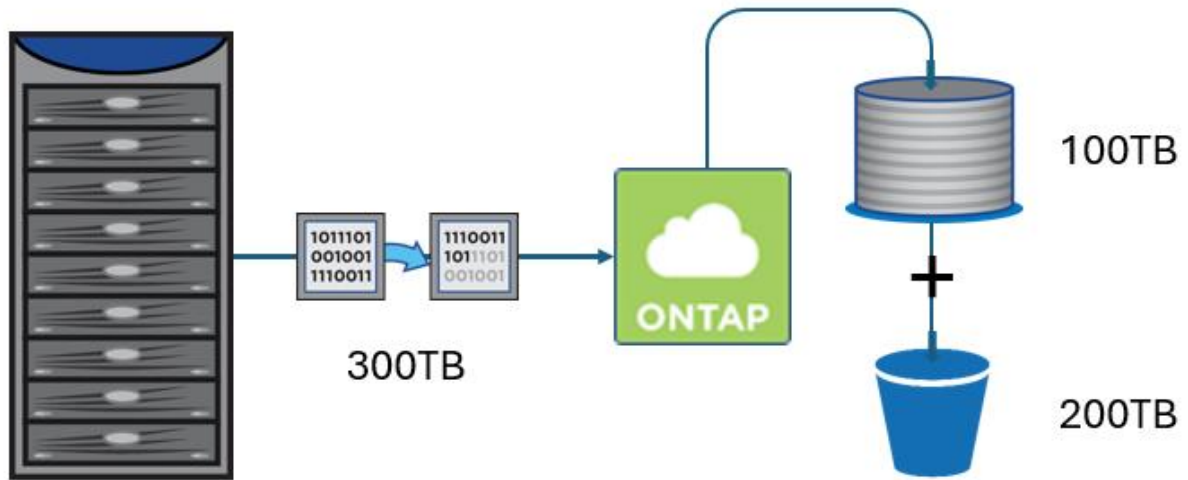
This type of scenario can happen in any environment that has been sized specifically for FabricPool aggregates, where a large data set is being migrated to a smaller local tier with the knowledge that the majority of the migrated data will be provisioned by the cloud tier.

Cloud Write

Beginning in ONTAP 9.13 (Cloud Volumes ONTAP and Amazon FSx for NetApp ONTAP) and ONTAP 9.14.1 (AFF, FAS, ONTAP Select), FabricPool supports Cloud Write, a feature specifically designed to avoid filling up the local tier before data can tier to the cloud tier.

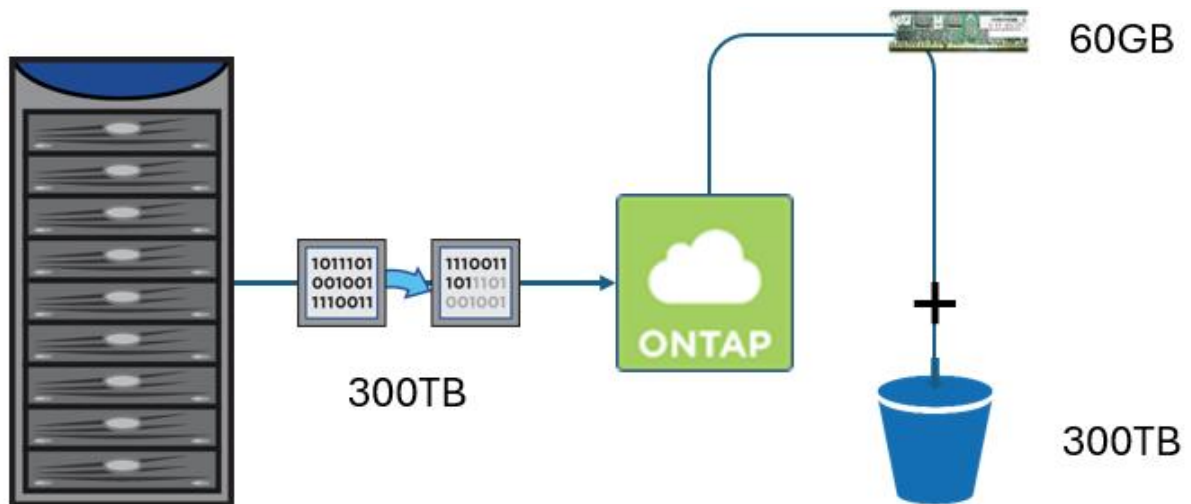
For example, in Figure 15, 300TB of data is being migrated to an ONTAP environment with the intention that 100TB will be active on the local tier and 200TB will be inactive and be provisioned on the cloud tier; but the migration will fail if the local tier reaches 98% capacity before data can tier out to the cloud tier.

Figure 15) 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 to enhance performance.

Figure 16) Migrating data with Cloud Write.



Note: 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
```

Note: 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 minimum 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.

Data tiering within Cloud Volumes ONTAP

Data Tiering within Cloud Volumes ONTAP, Amazon FSx for NetApp ONTAP, and Azure NetApp Files is based on FabricPool technology; however, it has different advantages and limitations.

Data Tiering documentation is located on the NetApp Cloud Docs site:

- [Data Tiering with Cloud Volumes ONTAP: Data Tiering Overview](#)
- [Data Tiering with Cloud Volumes ONTAP: Tiering Inactive Data to Low-Cost Object Storage](#)
- [Volume data tiering with FSx for ONTAP](#)

NetApp Private Storage for AWS

NetApp Private Storage (NPS) for AWS meets or exceeds all FabricPool best practices. The NPS for AWS solution is a high-performance cloud-connected storage architecture that allows enterprises to build an agile cloud infrastructure that combines the scalability and flexibility of the AWS cloud with the control and performance of NetApp storage.

NPS for AWS is typically deployed at one of the many AWS-approved Direct Connect partner colocation data centers (for example, Equinix). It uses AWS Direct Connect to provide a low-latency, highly available, dedicated connection between NetApp storage and the AWS cloud.

Where to find additional information

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

- ONTAP 9 Documentation
<http://docs.netapp.com/us-en/ontap/index.html>
- ONTAP Reference
<https://docs.netapp.com/us-en/ontap/concepts/manual-pages.html>
- StorageGRID 11.7 Documentation
<https://docs.netapp.com/us-en/storagegrid-117/>
- ONTAP S3 configuration overview with System Manager
https://docs.netapp.com/us-en/ontap/concept_object_provision_overview.html
- Setup licensing for BlueXP Tiering
https://docs.netapp.com/us-en/occm/task_licensing_cloud_tiering.html
- ONTAP FabricPool Licensing Overview
[https://kb.netapp.com/Advice_and_Troubleshooting/Data_Storage_Software/ONTAP_OS/ONTAP_FabricPool_\(FP\)_Licensing_Overview](https://kb.netapp.com/Advice_and_Troubleshooting/Data_Storage_Software/ONTAP_OS/ONTAP_FabricPool_(FP)_Licensing_Overview)
- TR-4015: SnapMirror Configuration and Best Practices Guide
<http://www.netapp.com/us/media/tr-4015.pdf>
- TR-4375: NetApp MetroCluster FC
<https://www.netapp.com/us/media/tr-4375.pdf>
- TR-4571: FlexGroup Volume Best Practices
<https://www.netapp.com/us/media/tr-4571.pdf>
- TR-4626: StorageGRID load balancer
<https://www.netapp.com/us/media/tr-4626.pdf>
- TR-4689: NetApp MetroCluster IP
<https://www.netapp.com/us/media/tr-4689.pdf>
- TR-4695: Database Storage Tiering with FabricPool
<https://www.netapp.com/us/media/tr-4695.pdf>
- TR-4814: ONTAP S3 Best Practices
<https://www.netapp.com/us/media/tr-4814.pdf>
- TR-4826: FabricPool with StorageGRID
<https://www.netapp.com/us/media/tr-4826.pdf>

Version history

| Version | Date | Document version history |
|---------|----------------|---|
| 3.5 | January 2024 | John Lantz: Updated for 9.14.1 Support for Cloud Write and Aggressive Readahead. Added details regarding large data migrations into FabricPool environments. |
| 3.4 | June 2023 | John Lantz: Updated for 9.13.1 Added details regarding PUT throttling and improved sequential read performance. |
| 3.3 | March 2023 | John Lantz: Added ONTAP S3 to the Default unreclaimed space thresholds table. Updated the StorageGRID consistency controls recommendation. |
| 3.2 | January 2023 | John Lantz: Added recommendations for StorageGRID consistency controls. |
| 3.1 | December 2022 | John Lantz: Updated for 9.12.1. Support for SVM migrate. Support for FabricPool, FlexGroup, and SVM-DR working in conjunction. (Prior to 9.12.1 any two of these features worked together, but not all three in conjunction.) |
| 3.0 | September 2022 | John Lantz: Added additional guidance regarding recommended best practices associated with the All volume tiering policy. |
| 2.9 | August 2022 | John Lantz: Additional minor updates for ONTAP 9.11.1. |
| 2.8 | July 2022 | John Lantz: Updated for 9.11.1. Added support for the Amazon S3 Glacier Instant Retrieval storage class. Data tiering now supported on Amazon FSx for NetApp ONTAP and Azure NetApp Files. |
| 2.7 | February 2022 | John Lantz: Updated for 9.10.1. Support for PUT throttling and temperature-sensitive storage efficiency (TSSE). |
| 2.6 | August 2021 | John Lantz: License installation changed from FabricPool licenses in System Manager to Cloud Tiering licenses in Cloud Manager. |
| 2.5 | June 2021 | John Lantz: Updated for 9.9.1. |
| 2.4 | December 2020 | John Lantz: Updated for 9.8. Support for HDD aggregates, ONTAP S3, cloud retrieval, and object tagging. Inactive data reporting can now be customized using the volume's minimum cooling days setting. |
| 2.3 | July 2020 | John Lantz: Added details regarding MetroCluster, StorageGRID, and FabricPool Mirror best practices. |
| 2.2 | January 2020 | John Lantz: Updated for ONTAP 9.7. Support for FabricPool Mirror, MetroCluster, S3-compatible providers, and NDMP. Writes from the cloud tier to the local tier are now disabled if local tier capacity is greater than 90%. Major updates to System Manager interfaces. NetApp terminology now refers to aggregates as local tiers. |
| 2.1 | September 2019 | John Lantz: Added details regarding the All tiering policy. |
| 2.0 | July 2019 | John Lantz: Updated for ONTAP 9.6. Support for Google Cloud Storage, Alibaba Cloud Object Storage Service, SVM-DR, and term-based licenses. All volume tiering policy replaces the now deprecated Backup volume |

| Version | Date | Document version history |
|---------|----------------|--|
| | | tiering policy. Added additional information regarding volume move enhancements, inactive data reporting, and maximum tiering capacities. |
| 1.9 | March 2019 | John Lantz: Added details regarding AES-256-GCM encryption and the need to avoid clock skew when attaching to private clouds. |
| 1.8 | January 2019 | John Lantz: Updated for ONTAP 9.5. Added support for FlexGroup volumes, client-side encryption, Amazon Commercial Cloud Services (C2S), IBM Cloud Object Storage, and the ability to change the aggregate fullness threshold. Aggregated Storage Tiering with Cloud Volumes ONTAP information. |
| 1.7 | August 2018 | John Lantz: Added additional information regarding Cloud Volumes ONTAP capacity and performance. |
| 1.6 | July 2018 | John Lantz: Cloud ONTAP renamed to Cloud Volumes ONTAP. Added additional information regarding metadata. |
| 1.5 | June 2018 | John Lantz: Support for tiering to Microsoft Azure Blob Storage, the Auto volume tiering policy, and io1 EBS volumes added to ONTAP Cloud. Writes from the cloud tier to the performance tier are disabled if performance tier capacity is greater than 70%. |
| 1.4 | May 2018 | John Lantz: Updated for ONTAP 9.4. Added Auto tiering policy, Microsoft Azure Blob support, inactive data reporting, and support for ONTAP Select Premium. |
| 1.3 | January 2018 | John Lantz: Updated for ONTAP 9.3. Added ONTAP Cloud functionality, AWS GovCloud S3, and additional interoperability details (QoS, StorageGRID, etc.). |
| 1.2 | September 2017 | John Lantz: Added details regarding connectivity requirements. |
| 1.1 | August 2017 | John Lantz: Added details regarding intercluster LIF requirements. |
| 1.0 | June 2017 | John Lantz: Initial commit. |

Contact us

Let us know how we can improve this technical report.

Contact us at doccomments@netapp.com.

Include TR-4598: FabricPool best practices in the subject line.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2023 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, non-sublicensable, 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).

TR-4598-0124