Technical Report

FabricPool Best Practices

ONTAP 9.6

John Lantz, NetApp
August 2019 | 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.

NetApp®

TABLE OF CONTENTS

Overview

FabricPool, first available in ONTAP 9.2, is a NetApp Data Fabric technology that enables automated tiering of data to low-cost object storage tiers either on or off premises.

Unlike manual tiering solutions, FabricPool reduces 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 AFF systems and all-SSD aggregates on FAS systems.
- ONTAP Select supports FabricPool. NetApp recommends using all-SSD FabricPool aggregates.

Figure 1) Before and after FabricPool.

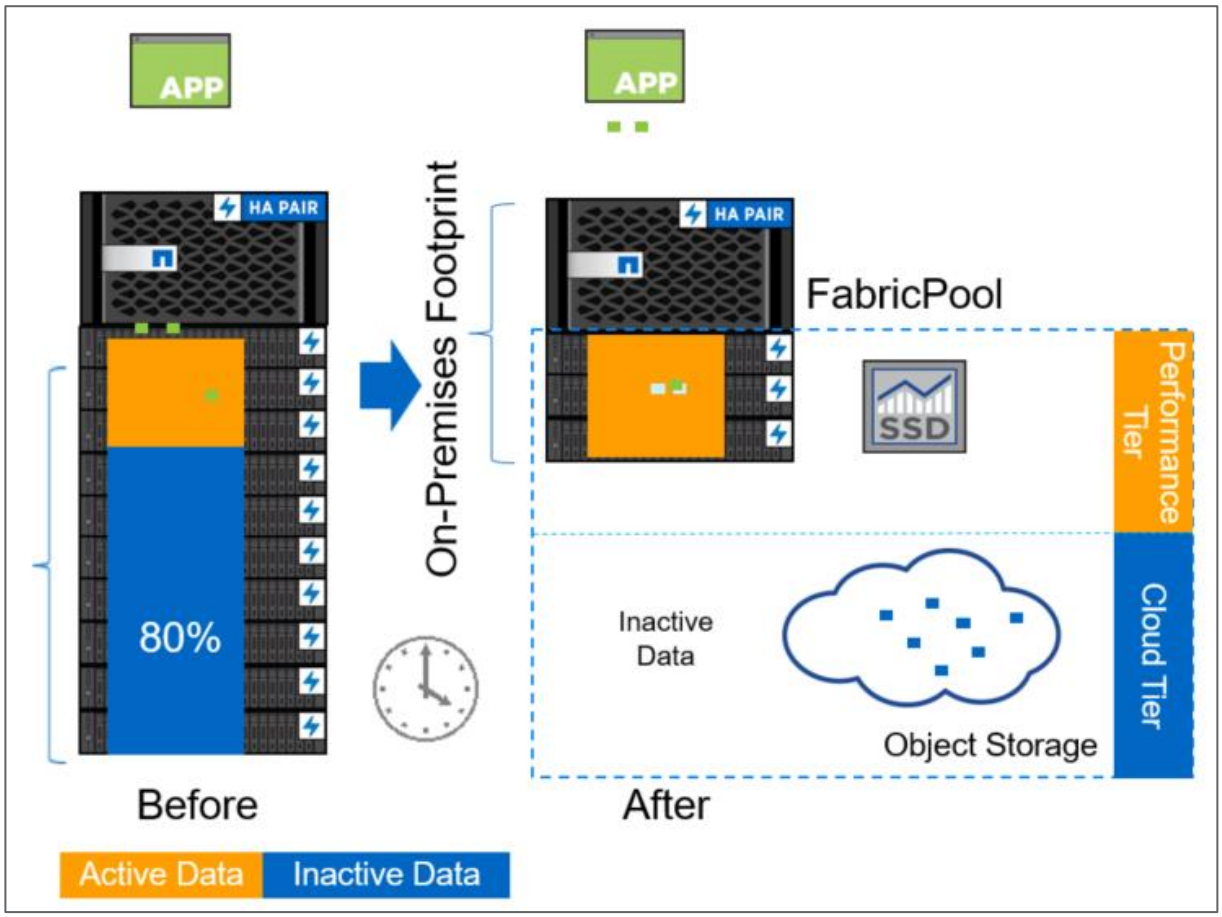

1 Primary Use Cases

The primary purpose of FabricPool is to reduce storage footprints and associated costs. Active data remains on high-performance SSDs 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 datacenters, it is not a backup solution. Access control lists (ACLs), directory structures, and WAFL® (Write Anywhere File Layout) metadata always stays on the performance tier. If a catastrophic disaster destroys the performance tier, a new environment cannot be created using the data on the cloud tier because it contains no WAFL metadata.

For complete data protection, consider using existing ONTAP technologies such as SnapMirror® and SnapVault®.

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

Auto Tiering Policy

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 SSDs, and tiering this data to a low-cost object store is an easy way to reclaim existing SSD capacity and reduce the amount of required SSD capacity moving forward.

First available in ONTAP 9.4, the Auto volume tiering policy moves all cold blocks in the volume, not just blocks associated with Snapshot copies, to the cloud tier.

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

Figure 2) Reclaiming space with the Auto volume tiering policy.

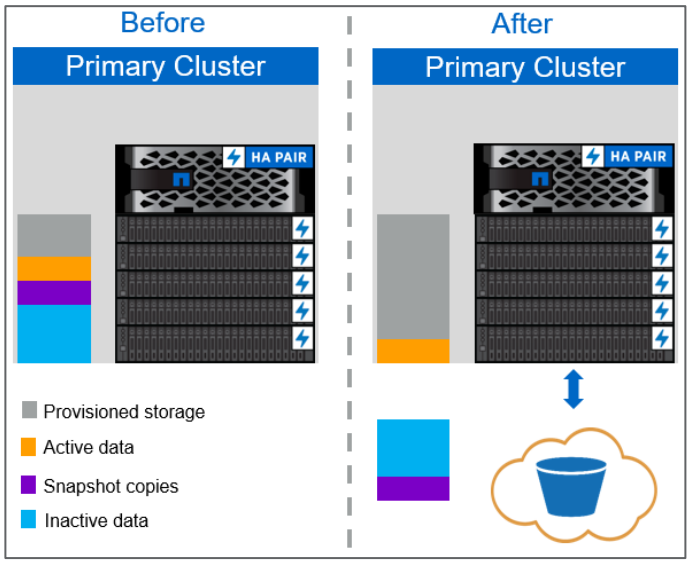

Snapshot-Only Tiering Policy

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 SSDs.

Snapshot-Only, a volume tiering policy for FabricPool, is an easy way to reclaim storage space on SSDs. When configured to use this policy, cold Snapshot 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 performance tier.

Note: The FabricPool Snapshot-Only volume tiering policy reduces the amount of storage used by Snapshot copies on SSDs. 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.

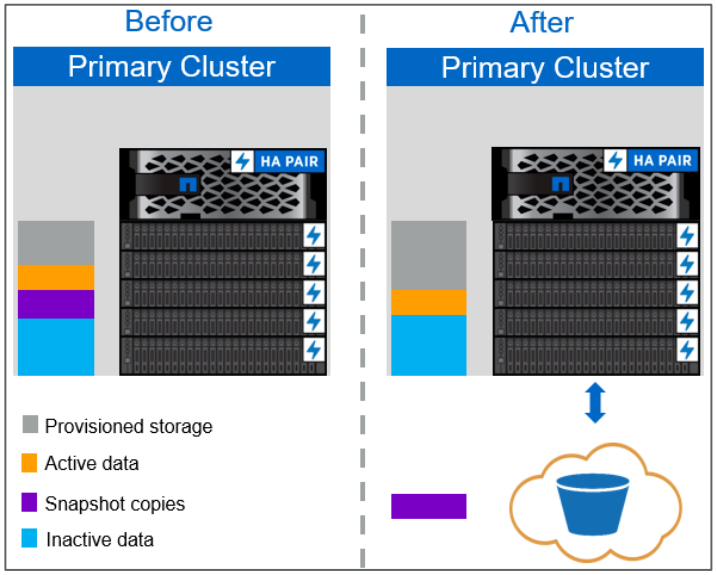

All Tiering Policy

In addition to cold primary data in active volumes (Auto) and snapshots (Snapshot-Only), one of the most common uses of FabricPool is to move entire volumes of data to low-cost clouds. Completed projects, legacy reports, or historical records—any dataset that must be retained but is unlikely to be frequently 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 policy is primarily used with secondary data and data protection volumes, but it can be also be used to tier all data in read/write volumes, provided the volume is not subject to frequent transactional operations.

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 performance tier.

Figure 4) Reclaiming space with the All volume tiering policy.

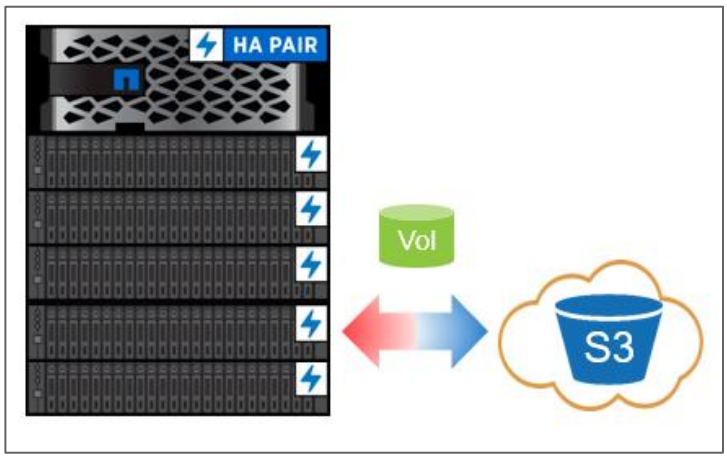

1.2 Shrink the Secondary Storage Footprint (All)

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, are an inefficient use of high-performance SSDs, and are expensive for large datasets even when using HDDs. FabricPool's 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 performance tier.

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

Figure 5) Using the All volume tiering policy with secondary storage.

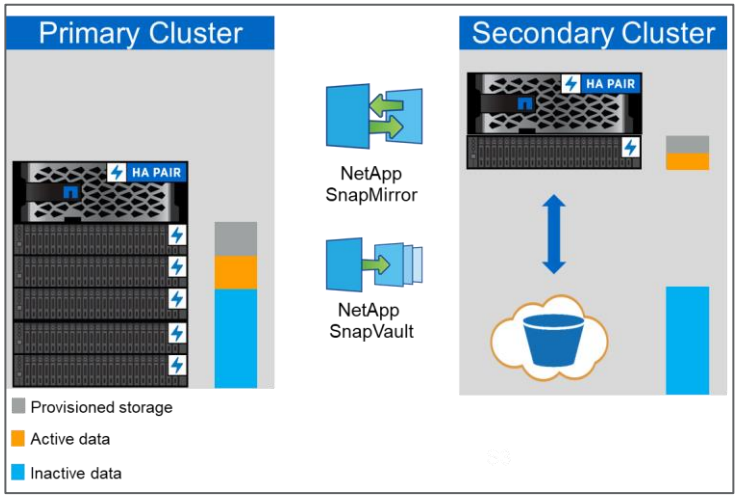

Figure 5 depicts the secondary as a traditional cluster running ONTAP. The secondary can also be in the cloud using Cloud ONTAP Volumes, or in a software defined environment using ONTAP Select. Data can be tiered using FabricPool anywhere ONTAP can be deployed.

2 Requirements

FabricPool requires ONTAP 9.2 or later and the use of SSD aggregates on any of the platforms listed in the next section. Additional FabricPool requirements depend on the cloud tier being attached.

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

2.1 Platforms

FabricPool is supported on a variety of platforms:

- **AFF**
  - A800
  - A700S, A700
  - A320, A300
  - A220, A200
  - C190
  - AFF8080, AFF8060, AFF8040
- **FAS**
  - FAS9000
  - FAS8200
  - FAS8080, FAS8060, FAS8040
  - FAS2750, FAS2720
  - FAS2650, FAS2620

  Note: Only SSD aggregates on FAS platforms can use FabricPool.
- **ONTAP Select**

  Note: NetApp recommends using all-SSD FabricPool aggregates.
- **Cloud tiers**
  - Alibaba Cloud Object Storage Service (Standard, Infrequent Access)
  - Amazon S3 (Standard, Standard-IA, One Zone-IA, Intelligent-Tiering)
  - Amazon Commercial Cloud Services (C2S)
  - Google Cloud Storage (Multi-Regional, Regional, Nearline, Coldline)
  - IBM Cloud Object Storage (Standard, Vault, Cold Vault, Flex)
  - Microsoft Azure Blob Storage (Hot and Cool)
  - StorageGRID 10.3+
- **Data Tiering with Cloud Volumes ONTAP**
  - Amazon S3
  - Google Cloud Storage
  - Microsoft Azure Blob Storage

2.2 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 aggregates on those nodes as well.

If you are using more than one IC 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 IC LIFs within an IPspace.

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

2.3 Volumes

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

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

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

FlexGroup Volumes

All aggregates used by a NetApp FlexGroup volume must be FabricPool aggregates.

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

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—decreasing performance. QoS Min must be turned off on volumes in FabricPool aggregates. Alternatively, tiering must be turned off (`-tiering-policy none`) on volumes that need QoS Min.

2.4 FabricPool 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 hybrid flash systems. A FabricPool license is not required when using StorageGRID as the cloud tier or when using Amazon S3, google Cloud Storage, or Microsoft Azure Blob Storage as the cloud tier for Cloud Volumes for ONTAP.

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

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 SSDs until the licensed capacity is increased.

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

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

Note: FabricPool licenses are cluster wide. Have your UUID available when purchasing a license (`cluster identity show`). For additional licensing information, refer to the NetApp Knowledgebase.

Installation

After you acquire a NetApp license file (NLF), you can install it by using OnCommand® System Manager. To do so, complete the following steps:

1. Click Configurations.
2. Click Cluster.
3. Click Licenses.
4. Click Add.
5. Click Choose Files to browse and select a file.
6. Click Add.

Figure 6) FabricPool license.

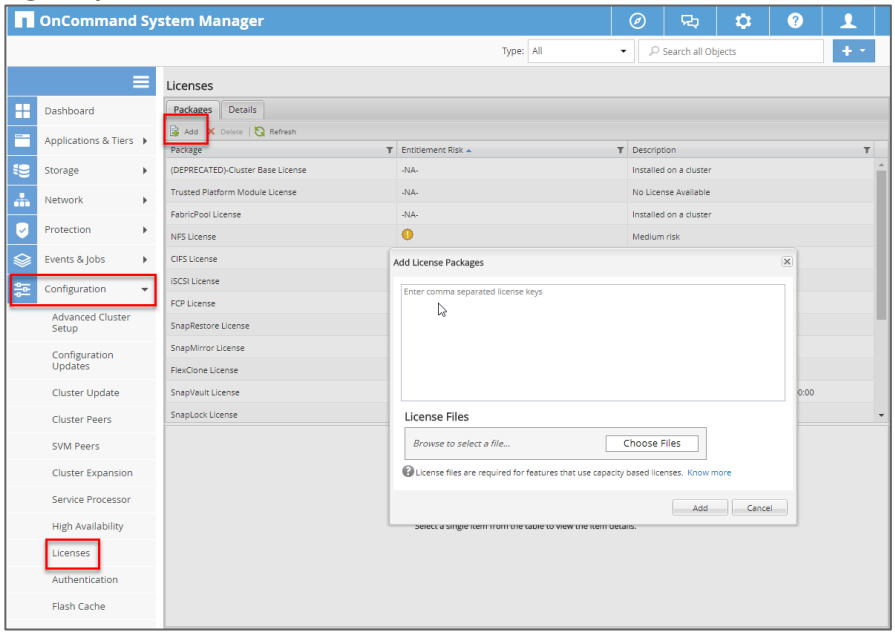

License Capacity

Licensed capacity can be viewed using the ONTAP CLI as well as OnCommand System Manager.

To see the licensed capacity, run the following command in the ONTAP CLI:

```
system license show-status
```

In OnCommand System Manager, complete the following steps:

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

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

Figure 7) License capacity.

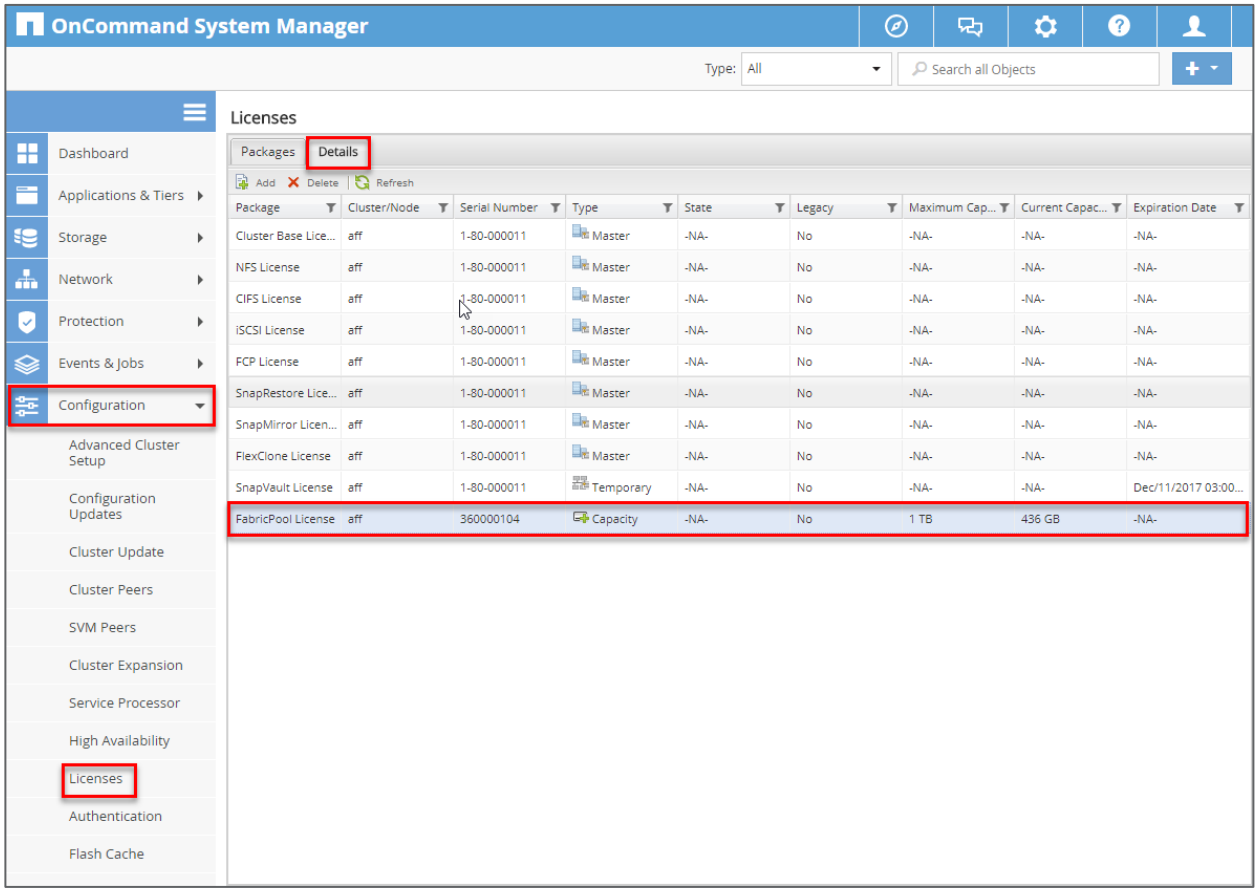

2.5 Certificate Authority Certification

When FabricPool uses StorageGRID or other private clouds like some IBM Cloud Object Storage environments as a cloud tier, it must use a Transport Layer Security (TLS) connection. Using FabricPool without TLS configuration is supported but not recommended.

Note: FabricPool requires that CA certificates use the same fully qualified domain name (FQDN) as the cloud tier server with which they are associated, but the default StorageGRID CA certificates use a common name (CN) that isn't based on the server's FQDN. This approach causes certificate-based errors that prohibit StorageGRID from being attached to aggregates.

Errors might include the following examples:

- Unable to add 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 as a cloud tier, you must replace the certificates in the grid with certificates that use the correct FQDN.

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

Note: Starting 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.

Installation

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

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

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

Install Certificates to ONTAP

In 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
```

3 Architecture

FabricPool works by associating an external object store with an aggregate in ONTAP, creating a composite aggregate: a FabricPool. Volumes inside the composite aggregate can then take advantage of the FabricPool by keeping active (hot) data on performative SSDs (the performance 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 migrates data is extremely useful when architecting storage solutions.

3.1 Block Temperature

When a block is written to an SSD, 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.

3.2 Object Creation

FabricPool works at the WAFL block level, cooling blocks, concatenating them into objects, and migrating those objects to a cloud tier. 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. If cold blocks are read and made hot, only the requested blocks in the 4MB object are fetched. Neither the entire object nor the entire file is migrated back. Only the necessary blocks are migrated.

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

3.3 Data Movement

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.

Tiering Fullness Threshold

By default, tiering to the cloud tier only happens if the performance tier aggregate is >50% full. There is little reason to tier cold data to a cloud tier if the performance 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 performance tier before tiering takes place. This may be useful for large aggregates that contain little hot/active data.

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

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

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

Note: Advanced privilege level is required.

Write-Back Prevention

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

SnapMirror Behavior

Movement of data from the cloud tier to the performance 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	Performance > Performance Cloud > Cloud
Auto	Snapshot-Only	Performance
Auto	All	Cloud
Auto	None	Performance
Snapshot-Only	Auto	Performance > Performance Cloud > Cloud
Snapshot-Only	Snapshot-Only	Performance > Performance Cloud > Cloud
Snapshot-Only	All	Cloud
Snapshot-Only	None	Performance
All	Auto	Performance
All	Snapshot-Only	Performance
All *	All *	Cloud*
All	None	Performance
None	Auto	Performance
None	Snapshot-Only	Performance
None	All	Cloud
None	None	Performance

*Cascading SnapMirror relationships are not supported when using the All volume tiering policy.

Volume Move

Volume move (`vol move`) is the way that ONTAP moves a volume nondisruptively from one aggregate (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 aggregate (the attached cloud tier) and the volume (volume tiering policies) levels can have a major impact on functionality.

Destination Aggregates

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

Starting in ONTAP 9.6, if a volume move's destination aggregate uses the same bucket as the source aggregate, data on the source volume that is stored in the bucket does not move back to the performance tier. This results in significant network efficiencies. (Setting the tiering policy to None will result in cold data being moved to the performance tier.)

If a volume move's destination aggregate has an attached cloud tier, data on the source volume that is stored on the cloud tier is first migrated to the performance tier on the destination aggregate. It is then migrated to the cloud tier on the destination aggregate if this approach is appropriate for the volume's tiering policy. Migrating data to the performance 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 aggregate resets the inactivity period of blocks on the performance tier. For example, a volume using the Auto volume tiering policy with data on the performance tier that has been inactive for 20 days 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 aggregate's performance 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 performance 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 aggregate's performance tier first.

If both source and destination volumes use the Snapshot-Only policy, and the Snapshot block is being read from the source aggregate's cloud tier, then FabricPool knows the Snapshot blocks are cold and moves the cold blocks to the destination aggregate's 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 migrated to the destination aggregate's cloud tier. There is no need to wait 48 hours for blocks in the volume to become cold. Metadata is always stored on the performance 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 migrated to the destination aggregate's performance tier.

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

OnCommand System Manager

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

1. Navigate to the Volumes page.
2. Select the volume you want to move.
3. Click Actions.
4. Click Move.
5. Select a destination aggregate.
6. Select a tiering policy.
7. Click Move.

Figure 8) Changing the volume tiering policy during a volume move.

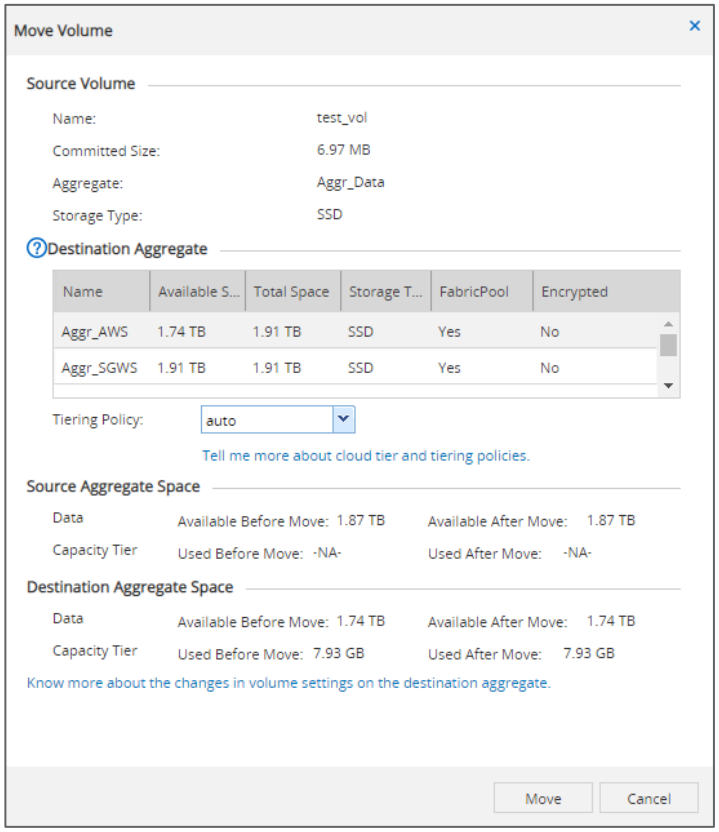

ONTAP CLI

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 of the parent FlexVol volume. After a FlexVol volume has been created, the volume tiering policy can be changed.

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

3.4 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. An ONTAP performance tier currently has a maximum volume size of 100TB and a maximum aggregate size of 800TB. Object stores have no such limits, and buckets with petabytes of data in them are not uncommon.

FabricPool Object Stores

FabricPool currently supports object stores from multiple providers (Alibaba, Amazon, Google, IBM, Microsoft, NetApp, etc.) as cloud tiers.

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

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.

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

- 12% Google Cloud Storage
- 14% IBM Cloud Object Storage
- 15% Alibaba Cloud Object Storage Service
- 15% Microsoft Azure Blob Storage
- 20% Amazon S3
- 40% StorageGRID

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 here.

Increasing the Threshold

Alternatively, consider increasing unreclaimed space thresholds if object fragmentation is resulting in significantly more object store capacity being used then 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, it is possible for 80% of total capacity in the cloud tier to be unreferenced by ONTAP.

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

In situations such as these, it might be advantageous to increase the unreclaimed space threshold—or increasing volume minimum cooling days—to reduce capacity being 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%)
```

(Advanced privilege level required.)

Note: Prior to ONTAP 9.4, object deletion effectively takes place at 0% free.

ONTAP Storage Efficiencies

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

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

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

4 Configuration

After FabricPool's basic requirements have been met, attaching an object store to an aggregate 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 an aggregate.
4. Set volume tiering policies.

4.1 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 an aggregate as a cloud tier.

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

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

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

Figure 9) Bucket-to-aggregate relationships.

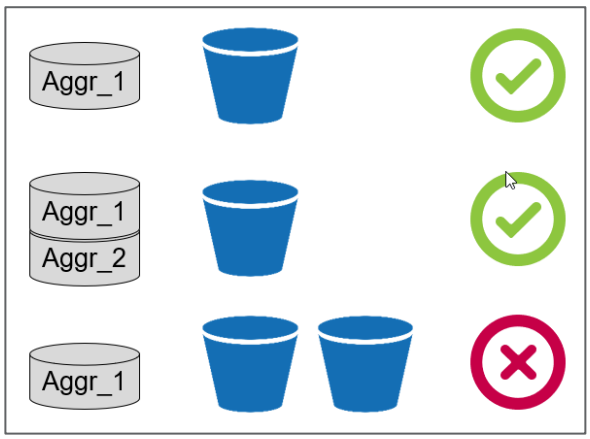

StorageGRID

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. Login with your tenant account ID, username and password
3. Select S3.
4. Select Buckets.
5. Click Create Bucket.
6. Provide a DNS compliant name.
7. Click Save.

Note: Prior to StorageGRID 11.1, creating a bucket required using a third-party S3 client such as an S3 browser.

Additional StorageGRID Settings

FabricPool supports StorageGRID's Information Lifecycle Management (ILM) policies for data replication and erasure coding to protect cloud tier data from failure. However, FabricPool does not support advanced ILM rules such as filtering based on user metadata or tags.

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, etc.).

StorageGRID uses two-copy replication as a default data protection solution. As of StorageGRID 11.2+, intra-site erasure coding using an 4+1 or 6+1 scheme is the recommended best practice for cost efficient data protection. Erasure coding uses more CPU, but significantly less storage capacity, than replication. Single copy replication is not recommended due to lowered system availability and data durability. Geographically dispersed erasure coding such as 4+2 over three physical sites is also not recommended due to WAN latencies.

ONTAP storage efficiencies such as compression, deduplication, and compaction are preserved when moving data to the cloud tier. We recommend disabling StorageGRID compression and encryption.

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 aggregate.

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
  https://www.alibabacloud.com/help/doc-detail/31885.htm
- Amazon S3
  https://docs.aws.amazon.com/AmazonS3/latest/gsg/CreatingABucket.html
- Google Cloud Storage
  https://cloud.google.com/storage/docs/creating-buckets
- IBM Cloud Object Storage
  https://cloud.ibm.com/docs/services/cloud-object-storage?topic=cloud-object-storage-getting-started
- Microsoft Azure Blob Storage
  https://docs.microsoft.com/en-us/azure/storage/blobs/storage-quickstart-blobs-portal

Other Object Store Provider Settings

FabricPool does not support information lifecycle management (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 will result prevent the Cleversafe bucket from being attached to the aggregate.

4.2 Add a Cloud Tier to ONTAP

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

FabricPool supports Amazon S3, IBM Object Cloud Storage, Microsoft Azure Blob Storage, and StorageGRID 10.3+ object stores as cloud tiers.

You need the following information:

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

OnCommand System Manager

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

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

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

7. Click Save and Attach Aggregates.

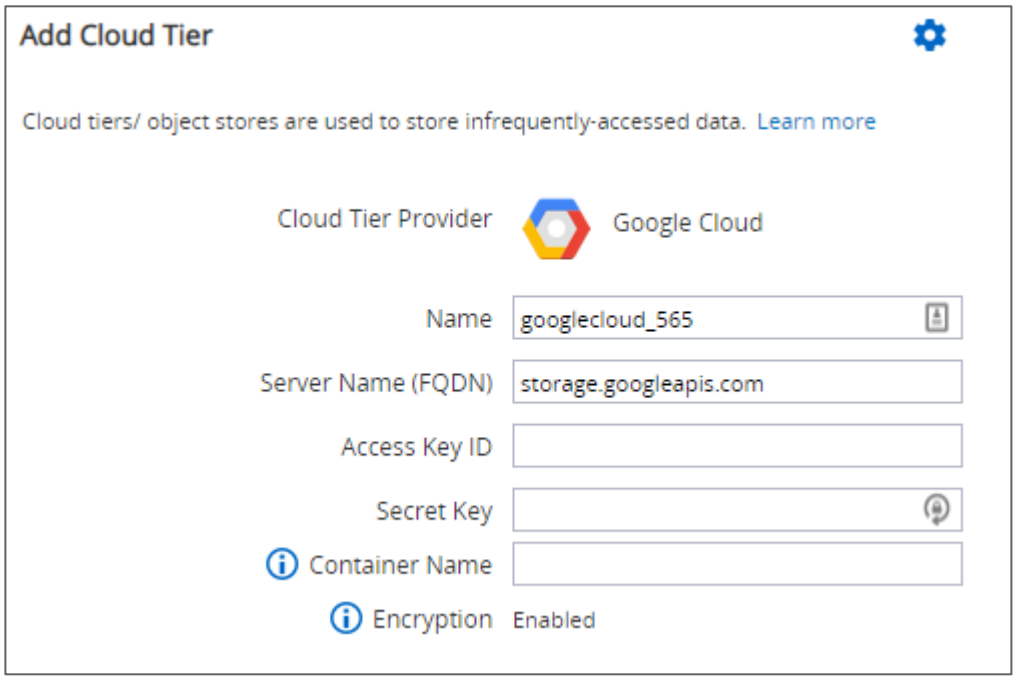

ONTAP CLI

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

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

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 aggregates as a cloud tier. 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 starting in ONTAP 9.4.

OnCommand System Manager

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

1. Launch OnCommand System Manager.
2. Click Storage.
3. Click Aggregates & Disks.
4. Click Cloud Tiers.
5. Select an object store provider.
6. Complete the text fields as required for your object store provider.
7. Click the Object Store Certificate button to turn it off.
8. Click Save and Attach Aggregates.

Figure 10) Object store certificate.

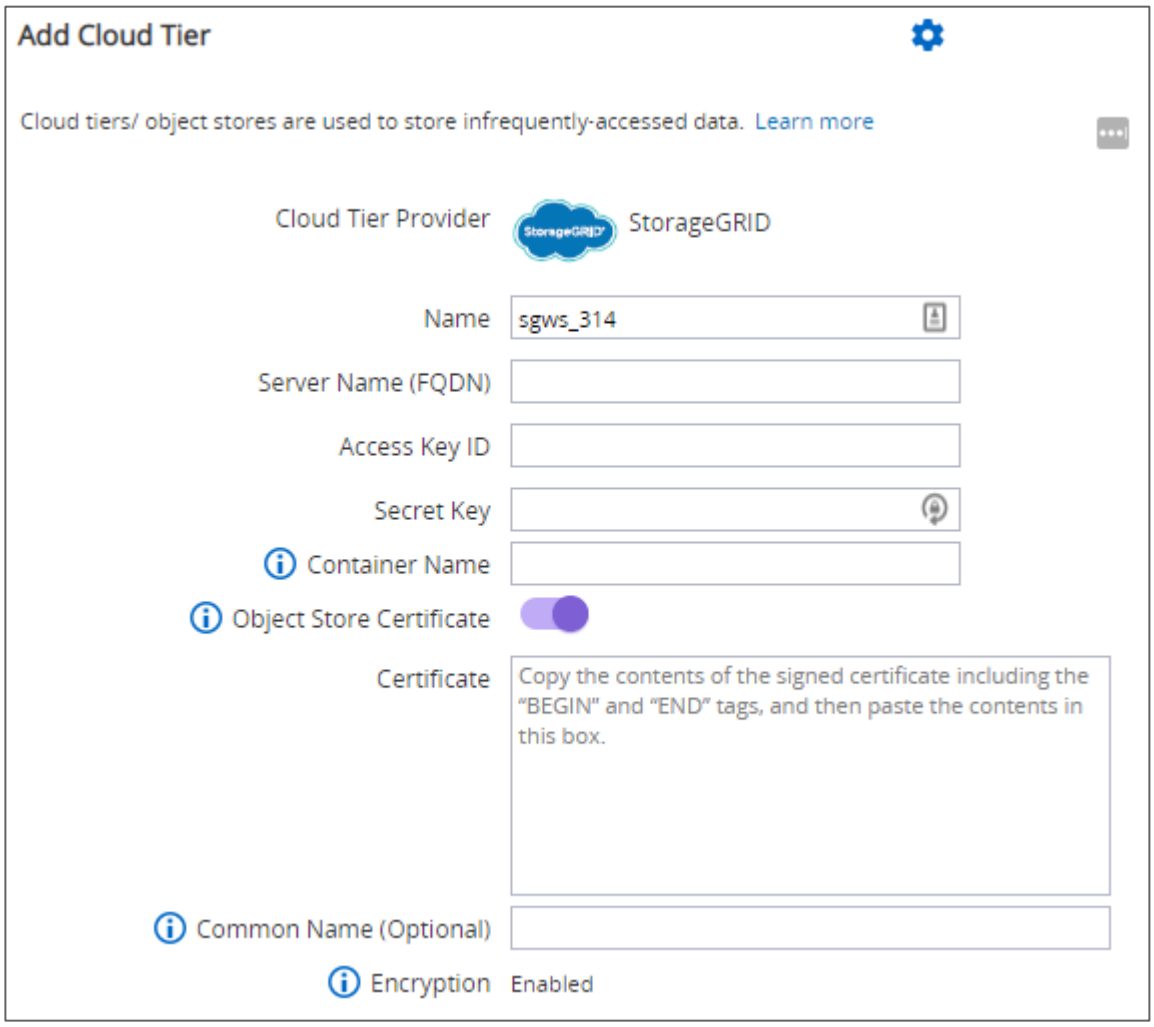

ONTAP CLI

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

```
object-store config create
-object-store-name <name>
-provider-type <AWS/Azure_Cloud/IBM_COS/SGWS>
-port <443/8082> (AWS&Azure_Cloud&IBM_COS/SGWS)
-server <name>
-container-name <bucket-name>
-access-key <string>
-secret-password <string>
-ssl-enabled true
-ipspace default
-is-certificate-validation-enabled false
```

4.3 Attach a Cloud Tier to an Aggregate

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

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

Note: Attaching a cloud tier to an aggregate is a permanent action. A cloud tier cannot be unattached from an aggregate after being attached to an aggregate.

Thin Provisioning

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

FlexGroup Volumes

All aggregates used by a FlexGroup volume must be FabricPool aggregates.

It is important to consider how bucket-to-aggregate relationships might affect performance when planning storage architectures. Many object store providers set a maximum number of supported IOPS at the bucket/container level. Environments that require maximum performance should use multiple buckets to reduce the possibility that object store IOPS limitations affect performance across multiple FabricPool aggregates. Attaching a single bucket/container to all FabricPool aggregates in a cluster might be more beneficial to environments that value manageability over performance.

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

OnCommand System Manager

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

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

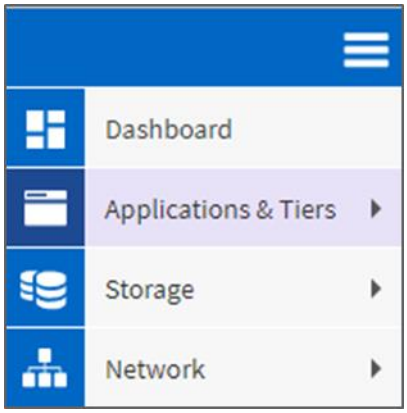

3. Click Storage Tiers.
4. Click an aggregate.
5. Click Actions and select Attach Cloud Tier.

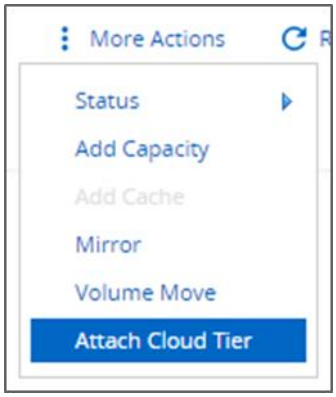

6. Select a cloud tier.
7. View and update the tiering policies for the volumes on the aggregate (optional). By default, volume tiering policies are set as Snapshot-Only.
8. Click Save.

ONTAP CLI

To attach a cloud tier to an aggregate 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
```

To list the aggregates used by a FlexGroup, and attach a cloud tier to those aggregates 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
```

4.4 Set Volume Tiering Policies

By default, volumes use the None volume tiering policy. After volume creation, the volume tiering policy can be changed using OnCommand 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 aggregate is >50% utilized, it takes approximately 31 days for inactive blocks to become cold. The Auto cooling period is adjustable between 2 days and 63 days using `tiering-minimum-cooling-days`.
  - When cold blocks in a volume with a tiering policy set to Auto are read randomly, they are made hot and written to the performance tier.
  - When cold blocks in a volume with a tiering policy set to Auto are read sequentially, they stay cold and remain on the cloud tier. They are not written to the performance tier.
- **Snapshot-Only**
  - Cold Snapshot blocks in the volume that are not shared with the active file system are moved to the cloud tier. Assuming the aggregate 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 63 days using `tiering-minimum-cooling-days`.
  - When cold blocks in a volume with a tiering policy set to Snapshot-Only are read, they are made hot and written to the performance tier.
- **All**
  - All data blocks (not including metadata) placed in the volume are immediately moved to the cloud tier. There is no need to wait 48 hours for new blocks in a volume using the All tiering policy to become cold.
  - Blocks located in the volume prior to the All policy being set require 48 hours 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 performance tier.

- Prior to 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: 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, and the addition of storage inefficiencies.

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.

- **None (default)**
  - Volumes set to use none as their tiering policy do not tier cold data to the cloud tier.
  - Setting the tiering policy to none prevents new tiering. Volume data that has previously been moved to the cloud tier remains in the cloud tier until it becomes hot and is automatically moved back to the performance tier.
  - When cold blocks in a volume with a tiering policy set to none are read, they are made hot and written to the performance tier.

OnCommand System Manager

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

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

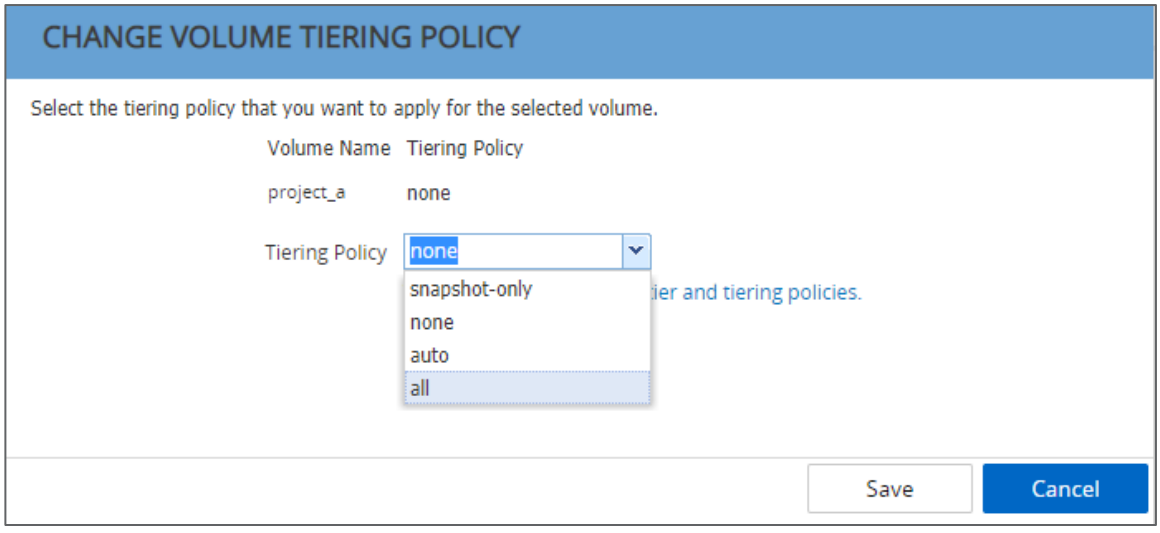

ONTAP CLI

To change a volume's tiering policy 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.

4.5 Set Volume Tiering Minimum Cooling Days

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

Auto

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

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

If you would like to reduce this value from the default 31-days, be aware that data should no longer be active before being marked as cold. For example, if a 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, and the addition of storage inefficiencies.

Snapshot-Only

The default tiering-minimum-cooling-days setting for the Snapshot-Only tiering policy is 2 days. A 2-day minimum 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 <2-63>
```

(Advanced privilege level required.)

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

4.6 Security

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

Performance Tier

FabricPool supports NetApp Storage Encryption (NSE) and NetApp Volume Encryption (NVE). Neither NSE nor NVE are required to use FabricPool aggregates.

Over the Wire

Objects moving between performance and cloud tiers are encrypted by using TLS 1.2. 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 remain encrypted when moved to the cloud tier. Client-side encryption keys are owned by ONTAP.

All objects not encrypted using NVE are encrypted server-side using AES-256-GCM encryption. Server-side encryption keys are owned by the respective object store.

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

5 Interoperability

In general, ONTAP functionality is unchanged on FabricPool aggregates. Although ONTAP must create and transfer objects and blocks between performance and cloud tiers, data protection, efficiency, and security are nearly identical to standard aggregates 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 2 and Table 3 are unique to FabricPool aggregates.

Table 2) NetApp interoperability.

Focus	Supported	Not Supported
Cloud tier	StorageGRID 10.3+	HDD aggregates
Data protection	SnapMirror (XDP and DP) SnapMirror Synchronous SnapVault (XDP and DP) SVM-DR StorageGRID replication and erasure coding **Note:** For best performance, use StorageGRID 11.2+ when enabling replication and 11.3+ when enabling erasure coding	7-Mode Data Transition Using SnapMirror 7-Mode Transition Tool (7MTT) DP_Optimized license (DPO) Dump MetroCluster NetApp SyncMirror® technology Cascading SnapMirror relationships using the All (or Backup) tiering policy. NDMP on any volume in a FabricPool aggregate (including NDMPcopy: check with your backup vendor to determine if NDMP is required) SMTape NetApp SnapLock® technology StorageGRID ILM policies other than replication and erasure coding StorageGRID object versioning
Encryption	NetApp Volume Encryption NetApp Storage Encryption Server-side encryption (AES-256) TLS 1.2	
Storage efficiency	Inline deduplication Inline compression Compaction Aggregate inline deduplication (performance tier only)	
Storage virtualization		NetApp FlexArray® technology
Quality of service (QoS)	QoS maximums (ceiling)	QoS minimums (floors)
Additional features		Auto Balance Aggregate

Table 3) Third-party interoperability.

Focus	Supported	Not Supported
Cloud tier	Alibaba Cloud Object Storage Service (Standard, Infrequent Access) Amazon S3 (Standard, Standard-IA, One Zone-IA, Intelligent-Tiering) Amazon Commercial Cloud Services (C2S) Google Cloud Storage (Multi-Regional, Regional, Nearline, Coldline) IBM Cloud Object Storage (including Cleversafe and SoftLayer) Microsoft Azure Blob Storage (Hot and Cool) StorageGRID 10.3+	Alibaba Archive Amazon Glacier Azure Archive IBM Archive
Data protection	Amazon's 99.999999999% multi-region durability	ILM policies
Encryption	Server-side encryption (AES-256) TLS 1.2	–

6 Performance

6.1 Sizing the Performance Tier

When considering sizing, the performance tier should be capable of the following tasks:

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

For most environments, a 1:10 performance:capacity ratio on FabricPool aggregates is extremely conservative while providing significant storage savings.

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

6.2 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

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 an aggregate.

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 the aggregate volumes. This amount might be different than the amount of cold data detected by IDR using a 31-day cooling period.

- IDR is enabled by default on all SSD aggregates in ONTAP 9.6.
- IDR is enabled by default on FabricPool aggregates in ONTAP 9.4 and ONTAP 9.5.
- IDR can be enabled on non-FabricPool aggregates using the ONTAP CLI. This includes HDD aggregates starting in ONTAP 9.6.
- IDR cannot be enabled for situations in which FabricPool cannot be enabled: for example, root, aggregates, MetroCluster, and so on.

OnCommand System Manager

IDR is displayed on the Storage Tiers page in OnCommand System Manager.

Figure 11) Inactive data reporting.

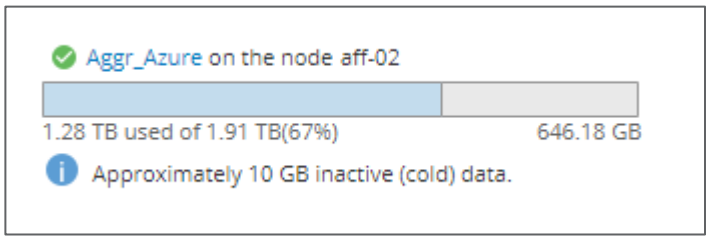

ONTAP CLI

To enable IDR on a non-FabricPool aggregate, 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
```

6.3 Connectivity

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.

Object Store Profiler

Starting 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 aggregates.

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

Start the object store profiler.

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

(Advanced privilege level required.)

View the results.

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

Note: Cloud tiers do not provide performance similar to that found on the performance tier (typically GB per second).

Although FabricPool aggregates can easily provide SATA-like performance, they can also tolerate latencies as high as 10 seconds and low throughputs for tiering solutions that do not need 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.

Network Connections

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 10Gbps is the recommended best practice for FabricPool.

- Alibaba Cloud Object Storage Service (Express Connect)
  https://www.alibabacloud.com/product/express-connect/pricing
- Amazon S3 (Direct Connect)
  https://aws.amazon.com/directconnect/faqs/
- Google Cloud Storage (Cloud Interconnect)
  https://cloud.google.com/interconnect/docs/
- IBM Cloud Object Storage (Direct Link)
  https://www.ibm.com/cloud/direct-link
- Microsoft Azure Blob Storage (ExpressRoute)
  https://azure.microsoft.com/en-us/pricing/details/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.

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 IC 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 IC LIFs within an IPspace.

Loss of Connectivity

If for any reason connectivity to the cloud is lost, the FabricPool performance 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

Fibre Channel and iSCSI protocols generally take longer before experiencing a timeout (60-120 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.

Even a short disruption could be disastrous to production applications using SAN protocols because there is no way to guarantee connectivity to public clouds. To avoid this, NetApp recommends using private clouds, like StorageGRID, when tiering data that is accessed by SAN protocols.

- **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.

6.4 Capacity

Aggregates

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

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

For example, if a volume move takes place at 2GBps but tiering takes place at 500MBps, 50TB completes the volume move to the performance tier in ~7 hours. However, ~28 hours are required for tiering to an object store. The performance tier must have enough capacity to store the data before it is tiered. Aggregate space utilization can be determined by using OnCommand System Manager or the ONTAP CLI.

OnCommand System Manager

In OnCommand System Manager, FabricPool space utilization is displayed on the Storage Tiers tab. Details include performance tier maximum capacity, used capacity, and external tier used capacity.

Figure 10) Storage tiers aggregate-level information.

Additionally, used capacity as a percentage of the licensed capacity is displayed at the bottom of the Storage Tiers tab.

Figure 11) Licensed capacity.

ONTAP CLI

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

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

Example:

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

Volumes

FlexVol volumes in a FabricPool aggregate 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 performance tier and 99TB on the cloud tier has reached the 100TB maximum FlexVol size, even though only 1TB is stored on the performance 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 performance tier aggregate reaches 98% capacity, FabricPool stops tiering cold data to the cloud tier. If the performance tier reaches 70% capacity, cold data is read directly from the cloud tier without migrating to the performance tier.

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

ONTAP CLI

View FabricPool volume space utilization details using the ONTAP CLI.

```
volume show-footprint
```

Total, performance, 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 SSDs 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 OnCommand System Manager, complete the following steps:

1. Click Configurations.
2. Click Licenses.
3. Click Details.
4. Click FabricPool License.
5. Current capacity is listed in the Current Capacity column.

Figure 12) License capacity.

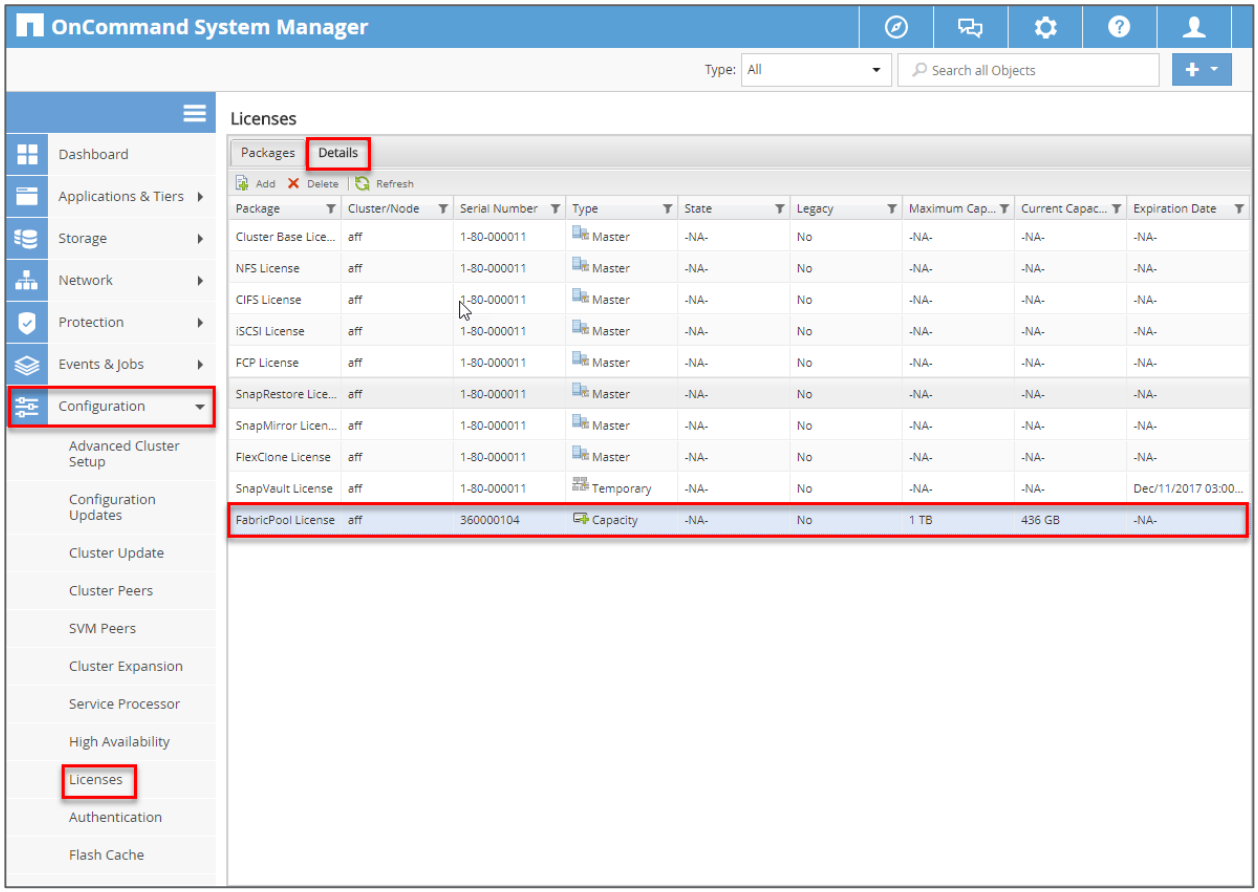

7 Data Tiering within Cloud Volumes ONTAP

Data Tiering within Cloud Volumes ONTAP is based on FabricPool technology; however, it has different advantages and limitations.

Data Tiering within the Cloud Volumes ONTAP documentation is now located on the NetApp Cloud Docs site.

- (Data Tiering with Cloud Volumes ONTAP) Data Tiering Overview
  https://docs.netapp.com/us-en/occm/concept_data_tiering.html
- (Data Tiering with Cloud Volumes ONTAP) Tiering Inactive Data to Low-Cost Object Storage
  https://docs.netapp.com/us-en/occm/task_tiering.html
- (Data Tiering with Cloud Volumes ONTAP) Cloud Manager and Cloud Volumes ONTAP documentation
  https://docs.netapp.com/us-en/occm/

8 Cloud Tiering Service

The Cloud Tiering Service is based on FabricPool technology; however, it has different advantages and limitations.

Cloud Tiering Service documentation is located on the NetApp Cloud Central site.

- Cloud Tiering with Cloud Volumes ONTAP
  https://cloud.netapp.com/cloud-tiering

9 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:

- Cluster Management Using OnCommand System Manager
  https://library.netapp.com/ecm/ecm_download_file/ECMLP2854936
- Commands: Manual Page Reference
  https://library.netapp.com/ecm/ecm_download_file/ECMLP2838138
- Configuring StorageGRID certificates for ONTAP clients using FabricPool
  http://docs.netapp.com/sgws-112/topic/com.netapp.doc.sg-admin/GUID-E1AF31C7-BDA2-495C-ABFE-C3A45A12B026.html
- ONTAP 9 Disks and Aggregates Power Guide
  https://library.netapp.com/ecm/ecm_download_file/ECMLP2496263
- ONTAP FabricPool Licensing Overview
  https://kb.netapp.com/support/s/article/ka21A0000008qb3QAA/ONTAP-FabricPool-FP-Licensing-Overview
- StorageGRID Administrator Guide
  https://library.netapp.com/ecm/ecm_download_file/ECMLP2848253
- TR-4015: SnapMirror Configuration and Best Practices Guide
  http://www.netapp.com/us/media/tr-4015.pdf
- TR-4075: DataMotion for Volumes
  http://www.netapp.com/us/media/tr-4075.pdf
- TR-4133: NetApp Private Storage for Amazon Web Services (AWS)
  http://www.netapp.com/us/media/tr-4133.pdf
- TR-4183: SnapVault Best Practices Guide
  https://www.netapp.com/us/media/tr-4183.pdf
- TR-4571: FlexGroup Volume Best Practices
  https://www.netapp.com/us/media/tr-4571.pdf
- TR-4695: Database Storage Tiering with FabricPool
  https://www.netapp.com/us/media/tr-4695.pdf
- NVA-0009: NetApp Private Storage for Cloud
  http://www.netapp.com/us/media/nva-0009.pdf

Version History

Version	Date	Document Version History
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 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

Trademark Information

TR-4598-0819