



FabricPool tier management

ONTAP 9

NetApp
March 08, 2024

This PDF was generated from <https://docs.netapp.com/us-en/ontap/fabricpool/index.html> on March 08, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- FabricPool tier management 1
 - FabricPool tier management overview 1
 - Benefits of storage tiers by using FabricPool 2
 - Considerations and requirements for using FabricPool 2
 - About FabricPool tiering policies 6
 - FabricPool management workflow 9
 - Configure FabricPool 10
 - Manage FabricPool 26
 - Manage FabricPool mirrors 45
 - Commands for managing aggregates with FabricPool 52

FabricPool tier management

FabricPool tier management overview

You can use FabricPool to automatically tier data depending on how frequently the data is accessed.

FabricPool is a hybrid storage solution that uses an all flash (all SSD) aggregate as the performance tier and an object store as the cloud tier. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

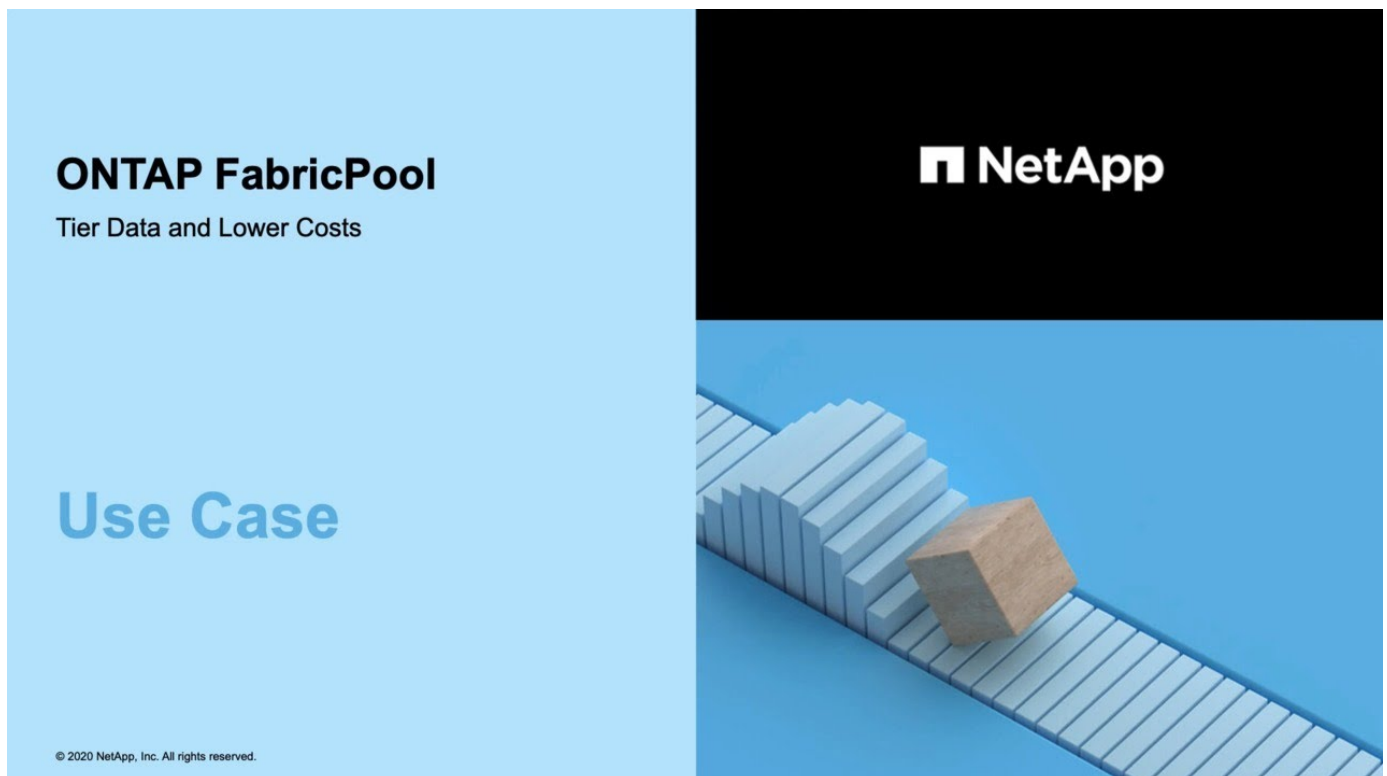
The cloud tier can be located on NetApp StorageGRID or ONTAP S3 (beginning with ONTAP 9.8), or one of the following service providers:

- Alibaba cloud
- Amazon S3
- Amazon Commercial Cloud Services
- Google Cloud
- IBM cloud
- Microsoft Azure Blob Storage



Beginning with ONTAP 9.7, additional object store providers that support generic S3 APIs can be used by selecting the S3_Compatible object store provider.

Tier Data and Lower Costs Use Case video



Related information

See also the [NetApp Cloud Tiering](#) documentation.

Benefits of storage tiers by using FabricPool

Configuring an aggregate to use FabricPool enables you to use storage tiers. You can efficiently balance the performance and cost of your storage system, monitor and optimize the space utilization, and perform policy-based data movement between storage tiers.

- You can optimize storage performance and reduce storage cost by storing data in a tier based on whether the data is frequently accessed.

- Frequently accessed (“hot”) data is stored in the *performance tier*.

The performance tier uses high-performance primary storage, such as an all flash (all SSD) aggregate of the storage system.

- Infrequently accessed (“cold”) data is stored in the *cloud tier*, also known as the *capacity tier*.

The cloud tier uses an object store that is less costly and does not require high performance.

- You have the flexibility in specifying the tier in which data should be stored.

You can specify one of the supported tiering policy options at the volume level. The options enable you to efficiently move data across tiers as data becomes hot or cold.

Types of FabricPool tiering policies

- You can choose one of the supported object stores to use as the cloud tier for FabricPool.
- You can monitor the space utilization in a FabricPool-enabled aggregate.
- You can see how much data in a volume is inactive by using inactive data reporting.
- You can reduce the on-premise footprint of the storage system.

You save physical space when you use a cloud-based object store for the cloud tier.

Considerations and requirements for using FabricPool

You should familiarize yourself with a few considerations and requirements about using FabricPool.

General considerations and requirements

- You must be running ONTAP 9.2 at the minimum to use FabricPool.
- You must be running ONTAP 9.4 or later releases for the following FabricPool functionality:
 - The `auto tiering policy`
 - Specifying the tiering minimum cooling period
 - Inactive data reporting (IDR)

- Using Microsoft Azure Blob Storage for the cloud as the cloud tier for FabricPool
- Using FabricPool with ONTAP Select
- You must be running ONTAP 9.5 or later releases for the following FabricPool functionality:
 - Specifying the tiering fullness threshold
 - Using IBM Cloud Object Storage as the cloud tier for FabricPool
 - NetApp Volume Encryption (NVE) of the cloud tier, enabled by default.
- You must be running ONTAP 9.6 or later releases for the following FabricPool functionality:
 - The `all` tiering policy
 - Inactive data reporting enabled manually on HDD aggregates
 - Inactive data reporting enabled automatically for SSD aggregates when you upgrade to ONTAP 9.6 and at time aggregate is created, except on low end systems with less than 4 CPU, less than 6 GB of RAM, or when WAFL-buffer-cache size is less than 3 GB.

ONTAP monitors system load, and if the load remains high for 4 continuous minutes, IDR is disabled, and is not automatically enabled. You can reenable IDR manually, however, manually enabled IDR is not automatically disabled.

- Using Alibaba Cloud Object Storage as the cloud tier for FabricPool
- Using Google Cloud Platform as the cloud tier for FabricPool
- Volume move without cloud tier data copy
- You must be running ONTAP 9.7 or later releases for the following FabricPool functionality:
 - Non transparent HTTP and HTTPS proxy to provide access to only whitelisted access points, and to provide auditing and reporting capabilities.
 - FabricPool mirroring to tier cold data to two object stores simultaneously
 - FabricPool mirrors on MetroCluster configurations
 - NDMP dump and restore enabled by default on FabricPool attached aggregates.



If the backup application uses a protocol other than NDMP, such as NFS or SMB, all data being backed up in the performance tier becomes hot and can affect tiering of that data to the cloud tier. Non-NDMP reads can cause data migration from the cloud tier back to the performance tier.

NDMP Backup and Restore Support for FabricPool

- You must be running ONTAP 9.8 or later for the following FabricPool functionality:
 - Cloud migration control to enable you to override the default tiering policy
 - Promoting data to the performance tier
 - FabricPool with SnapLock Enterprise. FabricPool with SnapLock Enterprise requires a Feature Product Variance Request (FPVR). To create an FPVR, please contact your sales team.
 - Minimum cooling period maximum of 183 days
 - Object tagging using user-created custom tags
 - FabricPools on HDD platforms and aggregates

HDD FabricPools are supported with SAS, FSAS, BSAS and MSATA disks only on systems with 6 or

more CPU cores, including the following models:

- FAS9000
- FAS8700
- FAS8300
- FAS8200
- FAS8080
- FAS8060
- FAS8040
- FAS2750
- FAS2720
- FAS2650
- FAS2620

Check [Hardware Universe](#) for the latest supported models.

- FabricPool is supported on all platforms capable of running ONTAP 9.2 except for the following:
 - FAS8020
 - FAS2554
 - FAS2552
 - FAS2520

- FabricPool supports the following aggregate types:

- On AFF systems, you can use only all flash (all SSD) aggregates for FabricPool.

You cannot use Flash Pool aggregates, which contain both SSDs and HDDs.

- On FAS systems, you can use either all flash (all SSD) or HDD aggregates for FabricPool.
- On Cloud Volumes ONTAP and ONTAP Select, you can use either SSD or HDD aggregates for FabricPool.

However, using SSD aggregates is recommended.

- FabricPool supports using the following object stores as the cloud tier:

- NetApp StorageGRID 10.3 or later
- NetApp ONTAP S3 (ONTAP 9.8 and later)
- Alibaba Cloud Object Storage
- Amazon Web Services Simple Storage Service (AWS S3)
- Google Cloud Storage
- IBM Cloud Object Storage
- Microsoft Azure Blob Storage for the cloud

- The object store “bucket” (container) you plan to use must have already been set up, must have at least 10 GB of storage space, and must not be renamed.
- HA pairs that use FabricPool require intercluster LIFs to communicate with the object store.

- You cannot detach a cloud tier from a local tier after it is attached; however, you can use [FabricPool mirror](#) to attach a local tier to a different cloud tier.
- If you use throughput floors (QoS Min), the tiering policy on the volumes must be set to `none` before the aggregate can be attached to FabricPool.

Other tiering policies prevent the aggregate from being attached to FabricPool. A QoS policy will not enforce throughput floors when FabricPool is enabled.

- You should follow the best practice guidelines for using FabricPool in specific scenarios.

[NetApp Technical Report 4598: FabricPool Best Practices in ONTAP 9](#)

Additional considerations when using Cloud Volumes ONTAP

Cloud Volumes ONTAP does not require a FabricPool license, regardless of the object store provider you are using.

Additional considerations for tiering data accessed by SAN protocols

When tiering data that is accessed by SAN protocols, NetApp recommends using private clouds, like StorageGRID, due to connectivity considerations.

Important

You should be aware that when using FabricPool in a SAN environment with a Windows host, if the object storage becomes unavailable for an extended period of time when tiering data to the cloud, files on the NetApp LUN on the Windows host might become inaccessible or disappear. See the Knowledge Base article [During FabricPool S3 object store unavailable Windows SAN host reported filesystem corruption](#).

Functionality or features not supported by FabricPool

- Object stores with WORM enabled and object versioning enabled.
- Information lifecycle management (ILM) policies that are applied to object store buckets

FabricPool supports StorageGRID's Information Lifecycle Management policies only 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 typically includes various movement and deletion policies. These policies can be disruptive to the data in the cloud tier of FabricPool. Using FabricPool with ILM policies that are configured on object stores can result in data loss.

- 7-Mode data transition using the ONTAP CLI commands or the 7-Mode Transition Tool
- FlexArray Virtualization
- RAID SyncMirror, except in a MetroCluster configuration
- SnapLock volumes when using ONTAP 9.7 and earlier releases
- Tape backup using SMTape for FabricPool-enabled aggregates
- The Auto Balance functionality
- Volumes using a space guarantee other than `none`

With the exception of root SVM volumes and CIFS audit staging volumes, FabricPool does not support attaching a cloud tier to an aggregate that contains volumes using a space guarantee other than `none`. For

example, a volume using a space guarantee of `volume (-space-guarantee volume)` is not supported.

- Clusters with DP_Optimized license
- Flash Pool aggregates

About FabricPool tiering policies

FabricPool tiering policies enable you to move data efficiently across tiers as data becomes hot or cold. Understanding the tiering policies helps you select the right policy that suits your storage management needs.

Types of FabricPool tiering policies

FabricPool tiering policies determine when or whether the user data blocks of a volume in FabricPool are moved to the cloud tier, based on the volume “temperature” of hot (active) or cold (inactive). The volume “temperature” increases when it is accessed frequently and decreases when it is not. Some tiering policies have an associated tiering minimum cooling period, which sets the time that user data in a volume of FabricPool must remain inactive for the data to be considered “cold” and moved to the cloud tier.

After a block has been identified as cold, it is marked as eligible to be tiered. A daily 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.



Data in volumes using the `all` tiering policy is immediately marked as cold and begins tiering to the cloud tier as soon as possible. It does not need to wait for the daily tiering scan to run.

You can use the `volume object-store tiering show` command to view the tiering status of a FabricPool volume. For more information, see the [Command Reference](#).

The FabricPool tiering policy is specified at the volume level. Four options are available:

- The `snapshot-only` tiering policy (the default) moves user data blocks of the volume Snapshot copies that are not associated with the active file system to the cloud tier.

The tiering minimum cooling period is 2 days. You can modify the default setting for the tiering minimum cooling period with the `-tiering-minimum-cooling-days` parameter in the advanced privilege level of the `volume create` and `volume modify` commands. Valid values are 2 to 183 days using ONTAP 9.8 and later. If you are using a version of ONTAP earlier than 9.8, valid values are 2 to 63 days.

- The `auto` tiering policy, supported only on ONTAP 9.4 and later releases, moves cold user data blocks in both the Snapshot copies and the active file system to the cloud tier.

The default tiering minimum cooling period is 31 days and applies to the entire volume, for both the active file system and the Snapshot copies.

You can modify the default setting for the tiering minimum cooling period with the `-tiering-minimum-cooling-days` parameter in the advanced privilege level of the `volume create` and `volume modify` commands. Valid values are 2 to 183 days.

- The `all` tiering policy, supported only with ONTAP 9.6 and later, moves all user data blocks in both the active file system and Snapshot copies to the cloud tier. It replaces the `backup` tiering policy.

The `all` volume tiering policy should not be used on read/write volumes that have normal client traffic.

The tiering minimum cooling period does not apply because the data moves to the cloud tier as soon as the tiering scan runs, and you cannot modify the setting.

- The `none` tiering policy keeps a volume's data in the performance tier and does not move cold 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.

The tiering minimum cooling period does not apply because the data never moves to the cloud tier, and you cannot modify the setting.

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.

The `volume show` command output shows the tiering policy of a volume. A volume that has never been used with FabricPool shows the `none` tiering policy in the output.

What happens when you modify the tiering policy of a volume in FabricPool

You can modify the tiering policy of a volume by performing a `volume modify` operation. You must understand how changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

- Changing the tiering policy from `snapshot-only` or `none` to `auto` causes ONTAP to send user data blocks in the active file system that are already cold to the cloud tier, even if those user data blocks were not previously eligible for the cloud tier.
- Changing the tiering policy to `all` from another policy causes ONTAP to move all user blocks in the active file system and in the Snapshot copies to the cloud as soon as possible. Prior to ONTAP 9.8, blocks needed to wait until the next tiering scan ran.

Moving blocks back to the performance tier is not allowed.

- Changing the tiering policy from `auto` to `snapshot-only` or `none` does not cause active file system blocks that are already moved to the cloud tier to be moved back to the performance tier.

Volume reads are needed for the data to be moved back to the performance tier.

- Any time you change the tiering policy on a volume, the tiering minimum cooling period is reset to the default value for the policy.

What happens to the tiering policy when you move a volume

- Unless you explicitly specify a different tiering policy, a volume retains its original tiering policy when it is moved in and out of a FabricPool-enabled aggregate.

However, the tiering policy takes effect only when the volume is in a FabricPool-enabled aggregate.

- The existing value of the `-tiering-minimum-cooling-days` parameter for a volume moves with the volume unless you specify a different tiering policy for the destination.

If you specify a different tiering policy, then the volume uses the default tiering minimum cooling period for that policy. This is the case whether the destination is FabricPool or not.

- You can move a volume across aggregates and at the same time modify the tiering policy.
- You should pay special attention when a `volume move` operation involves the `auto` tiering policy.

Assuming that both the source and the destination are FabricPool-enabled aggregates, the following table summarizes the outcome of a `volume move` operation that involves policy changes related to `auto`:

When you move a volume that has a tiering policy of...	And you change the tiering policy with the move to...	Then after the volume move...
<code>all</code>	<code>auto</code>	All data is moved to the performance tier.
<code>snapshot-only</code> , <code>none</code> , or <code>auto</code>	<code>auto</code>	Data blocks are moved to the same tier of the destination as they previously were on the source.
<code>auto</code> or <code>all</code>	<code>snapshot-only</code>	All data is moved to the performance tier.
<code>auto</code>	<code>all</code>	All user data is moved to the cloud tier.
<code>snapshot-only</code> , <code>auto</code> or <code>all</code>	<code>none</code>	All data is kept at the performance tier.

What happens to the tiering policy when you clone a volume

- Beginning with ONTAP 9.8, a clone volume always inherits both the tiering policy and the cloud retrieval policy from the parent volume.

In releases earlier than ONTAP 9.8, a clone inherits the tiering policy from the parent except when the parent has the `all` tiering policy.

- If the parent volume has the `never` cloud retrieval policy, its clone volume must have either the `never` cloud retrieval policy or the `all` tiering policy, and a corresponding cloud retrieval policy default.
- The parent volume cloud retrieval policy cannot be changed to `never` unless all its clone volumes have a cloud retrieval policy `never`.

When you clone volumes, keep the following best practices in mind:

- The `-tiering-policy` option and `tiering-minimum-cooling-days` option of the clone only controls the tiering behavior of blocks unique to the clone. Therefore, we recommend using tiering settings on the parent FlexVol that are either move the same amount of data or move less data than any of the clones
- The cloud retrieval policy on the parent FlexVol should either move the same amount of data or should move more data than the retrieval policy of any of the clones

How tiering policies work with cloud migration

FabricPool cloud data retrieval is controlled by tiering policies that determine data retrieval from the cloud tier to performance tier based on the read pattern. Read patterns can be either sequential or random.

The following table lists the tiering policies and the cloud data retrieval rules for each policy.

Tiering policy	Retrieval behavior
none	Sequential and random reads
snapshot-only	Sequential and random reads
auto	Random reads
all	No data retrieval

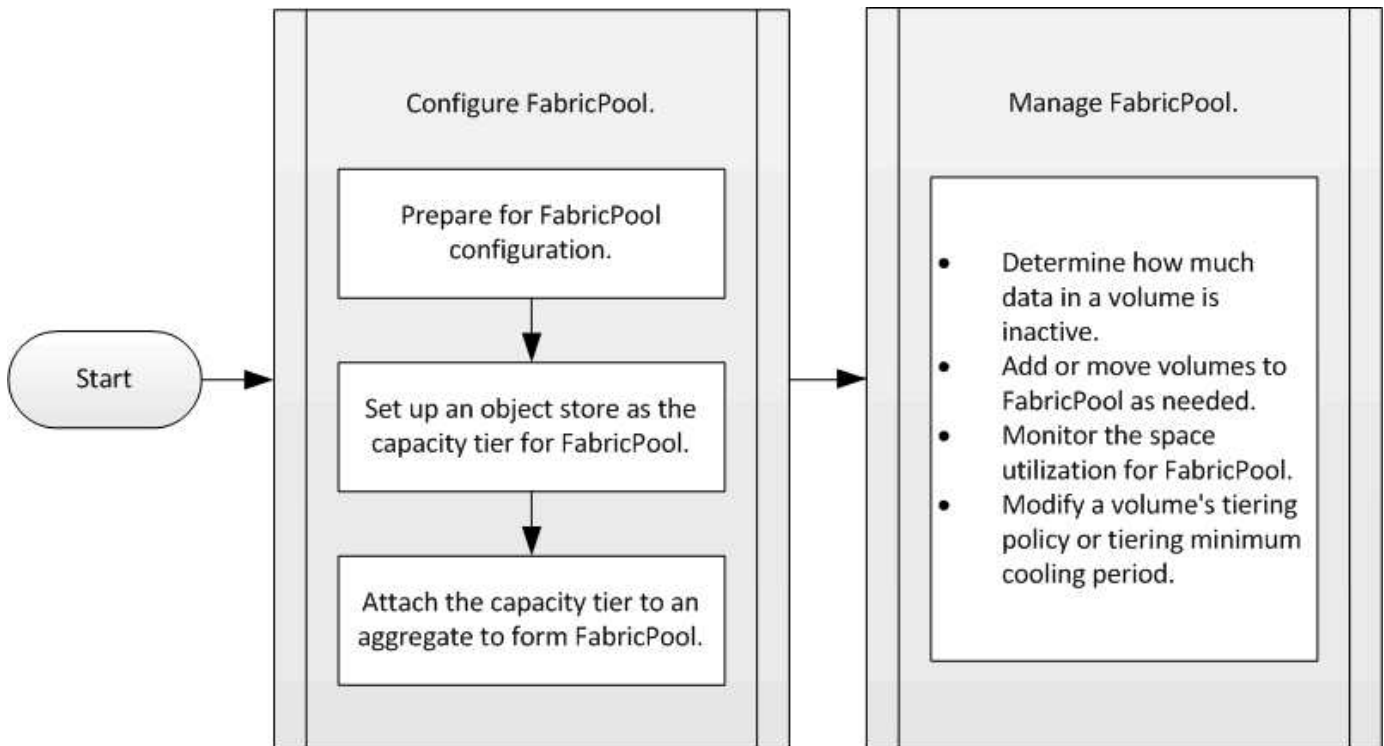
Beginning with ONTAP 9.8, the cloud migration control `cloud-retrieval-policy` option overrides the default cloud migration or retrieval behavior controlled by the tiering policy.

The following table lists the supported cloud retrieval policies and their retrieval behavior.

Cloud retrieval policy	Retrieval behavior
default	Tiering policy decides what data should be pulled back, so there is no change to cloud data retrieval with “default,” <code>cloud-retrieval-policy</code> . This policy is the default value for any volume regardless of the hosted aggregate type.
on-read	All client-driven data read is pulled from cloud tier to performance tier.
never	No client-driven data is pulled from cloud tier to performance tier
promote	<ul style="list-style-type: none">• For tiering policy “none,” all cloud data is pulled from the cloud tier to the performance tier• For tiering policy “snapshot-only,” AFS data is pulled.

FabricPool management workflow

You can use the FabricPool workflow diagram to help you plan the configuration and management tasks.



Configure FabricPool

Prepare for FabricPool configuration

Prepare for FabricPool configuration overview

Configuring FabricPool helps you manage which storage tier (the local performance tier or the cloud tier) data should be stored based on whether the data is frequently accessed.

The preparation required for FabricPool configuration depends on the object store you use as the cloud tier.

Add a connection to the cloud

Beginning with ONTAP 9.9.0, you can use System Manager to add a connection to the cloud.

You start by using NetApp Cloud Insights to configure a collector. During the configuration process, you copy a pairing code that is generated by Cloud Insights, and then you log on to a cluster using System Manager. There, you add a cloud connection using that pairing code. The rest of the process is completed in Cloud Insights.



If you choose the option to use a proxy server when adding a connection from Cloud Volumes ONTAP to Cloud Insights Service, you must ensure that the URL <https://example.com> is accessible from the proxy server. The message "The HTTP Proxy configuration is not valid" is displayed when <https://example.com> is not accessible.

Steps

1. In Cloud Insights, during the process to configure a collector, copy the generated pairing code.

2. Using System Manager with ONTAP 9.9.0 or later, log on to the cluster.
3. Go to **Cluster > Settings**.
4. In the Cloud Connections section, select **Add** to add a connection.
5. Enter a name for the connection, and paste the pairing code in the space provided.
6. Select **Add**.
7. Return to Cloud Insights to complete the configuration of the collector.

For additional information about Cloud Insights, refer to [Cloud Insights documentation](#).

Install a FabricPool license

The FabricPool license you might have used in the past is changing and is being retained only for configurations that aren't supported within BlueXP. Starting August 21, 2021, Cloud Tiering BYOL licensing was introduced for tiering configurations that are supported within BlueXP using the Cloud Tiering service.

[Learn more about the new Cloud Tiering BYOL licensing.](#)

Configurations that are supported by BlueXP must use the Digital Wallet page in BlueXP to license tiering for ONTAP clusters. This requires you to set up a BlueXP account and set up tiering for the particular object storage provider you plan to use. BlueXP currently supports tiering to the following object storage: Amazon S3, Azure Blob storage, Google Cloud Storage, S3-compatible object storage, and StorageGRID.

[Learn more about the Cloud tiering service.](#)

You can download and activate a FabricPool license using System Manager if you have one of the configurations that is not supported within BlueXP:

- ONTAP installations in Dark Sites
- ONTAP clusters that are tiering data to IBM Cloud Object Storage or Alibaba Cloud Object Storage

The FabricPool license is a cluster-wide license. It includes an entitled usage limit that you purchase for object storage that is associated with FabricPool in the cluster. The usage across the cluster must not exceed the capacity of the entitled usage limit. If you need to increase the usage limit of the license, you should contact your sales representative.

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

A term-based FabricPool license with 10 TB of free capacity is available for first time FabricPool orders for existing clusters configurations not supported within BlueXP. Free capacity is not available with perpetual licenses. A license is not required if you use NetApp StorageGRID or ONTAP S3 for the cloud tier. Cloud Volumes ONTAP does not require a FabricPool license, regardless of the provider you are using.

This task is supported only by uploading the license file to the cluster using System Manager.

Steps

1. Download the NetApp License File (NLF) for the FabricPool license from the [NetApp Support Site](#).
2. Perform the following actions using System Manager to upload the FabricPool license to the cluster:
 - a. In the **Cluster > Settings** pane, on the **Licenses** card, click ➔.

- b. On the **License** page, click [+ Add](#).
- c. In the **Add License** dialog box, click **Browse** to select the NLF you downloaded, and then click **Add** to upload the file to the cluster.

Related information

[ONTAP FabricPool \(FP\) Licensing Overview](#)

[NetApp Software License Search](#)

[NetApp TechComm TV: FabricPool playlist](#)

Install a CA certificate if you use StorageGRID

Unless you plan to disable certificate checking for StorageGRID, you must install a StorageGRID CA certificate on the cluster so that ONTAP can authenticate with StorageGRID as the object store for FabricPool.

About this task

ONTAP 9.4 and later releases enable you to disable certificate checking for StorageGRID.

Steps

1. Contact your StorageGRID administrator to obtain the StorageGRID system's CA certificate.
2. Use the `security certificate install` command with the `-type server-ca` parameter to install the StorageGRID CA certificate on the cluster.

The fully qualified domain name (FQDN) you enter must match the custom common name on the StorageGRID CA certificate.

Update an expired certificate

To update an expired certificate, the best practice is to use a trusted CA to generate the new server certificate. In addition, you should ensure that the certificate is updated on the StorageGRID server and on the ONTAP cluster at the same time to keep any downtime to a minimum.

Related information

[StorageGRID Resources](#)

Install a CA certificate if you use ONTAP S3

Unless you plan to disable certificate checking for ONTAP S3, you must install a ONTAP S3 CA certificate on the cluster so that ONTAP can authenticate with ONTAP S3 as the object store for FabricPool.

Steps

1. Obtain the ONTAP S3 system's CA certificate.
2. Use the `security certificate install` command with the `-type server-ca` parameter to install the ONTAP S3 CA certificate on the cluster.

The fully qualified domain name (FQDN) you enter must match the custom common name on the ONTAP S3 CA certificate.

Update an expired certificate

To update an expired certificate, the best practice is to use a trusted CA to generate the new server certificate. In addition, you should ensure that the certificate is updated on the ONTAP S3 server and on the ONTAP cluster at the same time to keep any downtime to a minimum.

Related information

[S3 configuration](#)

Set up an object store as the cloud tier for FabricPool

Set up an object store as the cloud tier for FabricPool overview

Setting up FabricPool involves specifying the configuration information of the object store (StorageGRID, ONTAP S3, Alibaba Cloud Object Storage, Amazon S3, Google Cloud Storage, IBM Cloud Object Storage, or Microsoft Azure Blob Storage for the cloud) that you plan to use as the cloud tier for FabricPool.

Set up StorageGRID as the cloud tier

If you are running ONTAP 9.2 or later, you can set up StorageGRID as the cloud tier for FabricPool. When tiering data that is accessed by SAN protocols, NetApp recommends using private clouds, like StorageGRID, due to connectivity considerations.

Considerations for using StorageGRID with FabricPool

- You need to install a CA certificate for StorageGRID, unless you explicitly disable certificate checking.
- You must not enable StorageGRID object versioning on the object store bucket.
- A FabricPool license is not required.
- If a StorageGRID node is deployed in a virtual machine with storage assigned from a NetApp AFF system, confirm that the volume does not have a FabricPool tiering policy enabled.

Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

About this task

Load balancing is enabled for StorageGRID in ONTAP 9.8 and later. When the server's hostname resolves to more than one IP address, ONTAP establishes client connections with all the IP addresses returned (up to a maximum of 16 IP addresses). The IP addresses are picked up in a round-robin method when connections are established.

Procedures

You can set up StorageGRID as the cloud tier for FabricPool with ONTAP System Manager or the ONTAP CLI.

System Manager

1. Click **Storage > Tiers > Add Cloud Tier** and select StorageGRID as the object store provider.
2. Complete the requested information.
3. If you want to create a cloud mirror, click **Add as FabricPool Mirror**.

A FabricPool mirror provides a method for you to seamlessly replace a data store, and it helps to ensure that your data is available in the event of disaster.

CLI

1. Specify the StorageGRID configuration information by using the `storage aggregate object-store config create` command with the `-provider-type SGWS` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access StorageGRID with the provided information.
 - You use the `-access-key` parameter to specify the access key for authorizing requests to the StorageGRID object store.
 - You use the `-secret-password` parameter to specify the password (secret access key) for authenticating requests to the StorageGRID object store.
 - If the StorageGRID password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in StorageGRID without interruption.

- Setting the `-is-certificate-validation-enabled` parameter to `false` disables certificate checking for StorageGRID.

```
cluster1::> storage aggregate object-store config create
-object-store-name mySGWS -provider-type SGWS -server mySGWSserver
-container-name mySGWScontainer -access-key mySGWSkey
-secret-password mySGWSpass
```

2. Display and verify the StorageGRID configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the StorageGRID configuration information for FabricPool.

Set up ONTAP S3 as the cloud tier

If you are running ONTAP 9.8 or later, you can set up ONTAP S3 as the cloud tier for FabricPool.

What you'll need

You must have the ONTAP S3 server name and the IP address of its associated LIFs on the remote cluster.

There must be intercluster LIFs on the local cluster.

About this task

Load balancing is enabled for ONTAP S3 servers in ONTAP 9.8 and later. When the server's hostname resolves to more than one IP address, ONTAP establishes client connections with all the IP addresses returned (up to a maximum of 16 IP addresses). The IP addresses are picked up in a round-robin method when connections are established.

Procedures

You can set up ONTAP S3 as the cloud tier for FabricPool with ONTAP System Manager or the ONTAP CLI.

System Manager

1. Click **Storage > Tiers > Add Cloud Tier** and select ONTAP S3 as the object store provider.
2. Complete the requested information.
3. If you want to create a cloud mirror, click **Add as FabricPool Mirror**.

A FabricPool mirror provides a method for you to seamlessly replace a data store, and it helps to ensure that your data is available in the event of disaster.

CLI

1. Add entries for the S3 server and LIFs to your DNS server.

Option	Description
If you use an external DNS server	Give the S3 server name and IP addresses to the DNS server administrator.
If you use your local system's DNS hosts table	Enter the following command: <pre>dns host create -vserver svm_name -address ip_address -hostname s3_server_name</pre>

2. Specify the ONTAP S3 configuration information by using the `storage aggregate object-store config create` command with the `-provider-type ONTAP_S3` parameter.
 - The `storage aggregate object-store config create` command fails if the local ONTAP system cannot access the ONTAP S3 server with the information provided.
 - You use the `-access-key` parameter to specify the access key for authorizing requests to the ONTAP S3 server.
 - You use the `-secret-password` parameter to specify the password (secret access key) for authenticating requests to the ONTAP S3 server.
 - If the ONTAP S3 server password is changed, you should immediately update the corresponding password stored in the local ONTAP system.

Doing so enables access to the data in the ONTAP S3 object store without interruption.

- Setting the `-is-certificate-validation-enabled` parameter to `false` disables certificate checking for ONTAP S3.

```
cluster1::> storage aggregate object-store config create  
-object-store-name myS3 -provider-type ONTAP_S3 -server myS3server  
-container-name myS3container -access-key myS3key  
-secret-password myS3pass
```

3. Display and verify the ONTAP_S3 configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the `ONTAP_S3` configuration information for FabricPool.

Set up Alibaba Cloud Object Storage as the cloud tier

If you are running ONTAP 9.6 or later, you can set up Alibaba Cloud Object Storage as the cloud tier for FabricPool.

Considerations for using Alibaba Cloud Object Storage with FabricPool

- You might need a FabricPool license.

Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool. If you need additional capacity on an AFF system, if you use Alibaba Cloud Object Storage on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.

- On AFF and FAS systems and ONTAP Select, FabricPool supports the following Alibaba Object Storage Service classes:
 - Alibaba Object Storage Service Standard
 - Alibaba Object Storage Service Infrequent Access

[Alibaba Cloud: Introduction to storage classes](#)

Contact your NetApp sales representative for information about storage classes not listed.

Steps

1. Specify the Alibaba Cloud Object Storage configuration information by using the `storage aggregate object-store config create` command with the `-provider-type AliCloud` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access Alibaba Cloud Object Storage with the provided information.
 - You use the `-access-key` parameter to specify the access key for authorizing requests to the Alibaba Cloud Object Storage object store.
 - If the Alibaba Cloud Object Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Alibaba Cloud Object Storage without interruption.

```
storage aggregate object-store config create my_ali_oss_store_1
-provider-type AliCloud -server oss-us-east-1.aliyuncs.com
-container-name my-ali-oss-bucket -access-key DXJRXHPXHYXA9X31X3JX
```

2. Display and verify the Alibaba Cloud Object Storage configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the Alibaba Cloud Object Storage configuration information for FabricPool.

Set up Amazon S3 as the cloud tier

If you are running ONTAP 9.2 or later, you can set up Amazon S3 as the cloud tier for FabricPool. If you are running ONTAP 9.5 or later, you can set up Amazon Commercial Cloud Services (C2S) for FabricPool.

Considerations for using Amazon S3 with FabricPool

- You might need a FabricPool license.
 - Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool.

If you need additional capacity on an AFF system, if you use Amazon S3 on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.

If you order FabricPool for the first time for an existing cluster, a FabricPool license with 10 TB of free capacity is available.
- It is recommended that the LIF that ONTAP uses to connect with the Amazon S3 object server be on a 10 Gbps port.
- On AFF and FAS systems and ONTAP Select, FabricPool supports the following Amazon S3 storage classes:
 - Amazon S3 Standard
 - Amazon S3 Standard - Infrequent Access (Standard - IA)
 - Amazon S3 One Zone - Infrequent Access (One Zone - IA)
 - Amazon S3 Intelligent-Tiering
 - Amazon Commercial Cloud Services
 - Beginning with ONTAP 9.11.1, Amazon S3 Glacier Instant Retrieval (FabricPool does not support Glacier Flexible Retrieval or Glacier Deep Archive)

[Amazon Web Services Documentation: Amazon S3 Storage Classes](#)

Contact your sales representative for information about storage classes not listed.

- On Cloud Volumes ONTAP, FabricPool supports tiering from General Purpose SSD (gp2) and Throughput Optimized HDD (st1) volumes of Amazon Elastic Block Store (EBS).

Steps

1. Specify the Amazon S3 configuration information by using the `storage aggregate object-store config create` command with the `-provider-type AWS_S3` parameter.

- You use the `-auth-type CAP` parameter to obtain credentials for C2S access.

When you use the `-auth-type CAP` parameter, you must use the `-cap-url` parameter to specify the full URL to request temporary credentials for C2S access.

- The `storage aggregate object-store config create` command fails if ONTAP cannot access Amazon S3 with the provided information.
- You use the `-access-key` parameter to specify the access key for authorizing requests to the Amazon S3 object store.
- You use the `-secret-password` parameter to specify the password (secret access key) for

authenticating requests to the Amazon S3 object store.

- If the Amazon S3 password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Amazon S3 without interruption.

```
cluster1::> storage aggregate object-store config create
-object-store-name my_aws_store -provider-type AWS_S3
-server s3.amazonaws.com -container-name my-aws-bucket
-access-key DXJRXHPXHYXA9X31X3JX
```

```
cluster1::> storage aggregate object-store config create -object
-store-name my_c2s_store -provider-type AWS_S3 -auth-type CAP -cap
-url
https://123.45.67.89/api/v1/credentials?agency=XYZ&mission=TESTACCT&r
ole=S3FULLACCESS -server my-c2s-s3server-fqdn -container my-c2s-s3-
bucket
```

2. Display and verify the Amazon S3 configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the Amazon S3 configuration information for FabricPool.

Set up Google Cloud Storage as the cloud tier

If you are running ONTAP 9.6 or later, you can set up Google Cloud Storage as the cloud tier for FabricPool.

Additional considerations for using Google Cloud Storage with FabricPool

- You might need a FabricPool license.

Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool. If you need additional capacity on an AFF system, if you use Google Cloud Storage on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.

- It is recommended that the LIF that ONTAP uses to connect with the Google Cloud Storage object server be on a 10 Gbps port.
- On AFF and FAS systems and ONTAP Select, FabricPool supports the following Google Cloud Object storage classes:
 - Google Cloud Multi-Regional
 - Google Cloud Regional
 - Google Cloud Nearline
 - Google Cloud Coldline

Steps

1. Specify the Google Cloud Storage configuration information by using the `storage aggregate object-store config create` command with the `-provider-type GoogleCloud` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access Google Cloud Storage with the provided information.
 - You use the `-access-key` parameter to specify the access key for authorizing requests to the Google Cloud Storage object store.
 - If the Google Cloud Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Google Cloud Storage without interruption.

```
storage aggregate object-store config create my_gcp_store_1 -provider
-type GoogleCloud -container-name my-gcp-bucket1 -access-key
GOOGAUZZUV2USCFGHGQ511I8
```

2. Display and verify the Google Cloud Storage configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the Google Cloud Storage configuration information for FabricPool.

Set up IBM Cloud Object Storage as the cloud tier

If you are running ONTAP 9.5 or later, you can set up IBM Cloud Object Storage as the cloud tier for FabricPool.

Considerations for using IBM Cloud Object Storage with FabricPool

- You might need a FabricPool license.

Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool. If you need additional capacity on an AFF system, if you use IBM Cloud Object Storage on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.

If you order FabricPool for the first time for an existing cluster, a FabricPool license with 10 TB of free capacity is available.

- It is recommended that the LIF that ONTAP uses to connect with the IBM Cloud object server be on a 10 Gbps port.

Steps

1. Specify the IBM Cloud Object Storage configuration information by using the `storage aggregate object-store config create` command with the `-provider-type IBM_COS` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access IBM Cloud Object Storage with the provided information.

- You use the `-access-key` parameter to specify the access key for authorizing requests to the IBM Cloud Object Storage object store.
- You use the `-secret-password` parameter to specify the password (secret access key) for authenticating requests to the IBM Cloud Object Storage object store.
- If the IBM Cloud Object Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in IBM Cloud Object Storage without interruption.

```
storage aggregate object-store config create
-object-store-name MyIBM -provider-type IBM_COS
-server s3.us-east.objectstorage.softlayer.net
-container-name my-ibm-cos-bucket -access-key DXJRXHPXHYXA9X31X3JX
```

2. Display and verify the IBM Cloud Object Storage configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the IBM Cloud Object Storage configuration information for FabricPool.

Set up Azure Blob Storage for the cloud as the cloud tier

If you are running ONTAP 9.4 or later, you can set up Azure Blob Storage for the cloud as the cloud tier for FabricPool.

Considerations for using Microsoft Azure Blob Storage with FabricPool

- You might need a FabricPool license.

Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool. If you need additional capacity on an AFF system, if you use Azure Blob Storage on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.

If you order FabricPool for the first time for an existing cluster, a FabricPool license with 10 TB of free capacity is available.

- A FabricPool license is not required if you are using Azure Blob Storage with Cloud Volumes ONTAP.
- It is recommended that the LIF that ONTAP uses to connect with the Azure Blob Storage object server be on a 10 Gbps port.
- FabricPool currently does not support Azure Stack, which is on-premises Azure services.
- At the account level in Microsoft Azure Blob Storage, FabricPool supports only hot and cool storage tiers.

FabricPool does not support blob-level tiering. It also does not support tiering to Azure's archive storage tier.

About this task

FabricPool currently does not support Azure Stack, which is on-premises Azure services.

Steps

1. Specify the Azure Blob Storage configuration information by using the `storage aggregate object-store config create` command with the `-provider-type Azure_Cloud` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access Azure Blob Storage with the provided information.
 - You use the `-azure-account` parameter to specify the Azure Blob Storage account.
 - You use the `-azure-private-key` parameter to specify the access key for authenticating requests to Azure Blob Storage.
 - If the Azure Blob Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Azure Blob Storage without interruption.

```
cluster1::> storage aggregate object-store config create
-object-store-name MyAzure -provider-type Azure_Cloud
-server blob.core.windows.net -container-name myAzureContainer
-azure-account myAzureAcct -azure-private-key myAzureKey
```

2. Display and verify the Azure Blob Storage configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the Azure Blob Storage configuration information for FabricPool.

Set up object stores for FabricPool in a MetroCluster configuration

If you are running ONTAP 9.7 or later, you can set up a mirrored FabricPool on a MetroCluster configuration to tier cold data to object stores in two different fault zones.

About this task

- FabricPool in MetroCluster requires that the underlying mirrored aggregate and the associated object store configuration must be owned by the same MetroCluster configuration.
- You cannot attach an aggregate to an object store that is created in the remote MetroCluster site.
- You must create object store configurations on the MetroCluster configuration that owns the aggregate.

Before you begin

- The MetroCluster configuration is set up and properly configured.
- Two objects stores are set up on the appropriate MetroCluster sites.
- Containers are configured on each of the object stores.
- IP spaces are created or identified on the two MetroCluster configurations and their names match.

Step

1. Specify the object store configuration information on each MetroCluster site by using the `storage object-store config create` command.

In this example, FabricPool is required on only one cluster in the MetroCluster configuration. Two object

store configurations are created for that cluster, one for each object store bucket.

```
storage aggregate
  object-store config create -object-store-name mcc1-ostore-config-s1
-provider-type SGWS -server
  <SGWS-server-1> -container-name <SGWS-bucket-1> -access-key <key>
-secret-password <password> -encrypt
  <true|false> -provider <provider-type> -is-ssl-enabled <true|false>
ipSPACE
  <IPSpace>
```

```
storage aggregate object-store config create -object-store-name mcc1-
ostore-config-s2
  -provider-type SGWS -server <SGWS-server-2> -container-name <SGWS-
bucket-2> -access-key <key> -secret-password <password> -encrypt
<true|false> -provider <provider-type>
  -is-ssl-enabled <true|false> ipSPACE <IPSpace>
```

This example sets up FabricPool on the second cluster in the MetroCluster configuration.

```
storage aggregate
  object-store config create -object-store-name mcc2-ostore-config-s1
-provider-type SGWS -server
  <SGWS-server-1> -container-name <SGWS-bucket-3> -access-key <key>
-secret-password <password> -encrypt
  <true|false> -provider <provider-type> -is-ssl-enabled <true|false>
ipSPACE
  <IPSpace>
```

```
storage aggregate
  object-store config create -object-store-name mcc2-ostore-config-s2
-provider-type SGWS -server
  <SGWS-server-2> -container-name <SGWS-bucket-4> -access-key <key>
-secret-password <password> -encrypt
  <true|false> -provider <provider-type> -is-ssl-enabled <true|false>
ipSPACE
  <IPSpace>
```

Test object store throughput performance before attaching to a local tier

Before you attach an object store to a local tier, you can test the object store's latency and throughput performance by using object store profiler.

Before you being

- You must add the cloud tier to ONTAP before you can use it with the object store profiler.
- You must be at the ONTAP CLI advanced privilege mode.

Steps

1. Start the object store profiler:

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

2. View the results:

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

Attach the cloud tier to a local tier (aggregate)

After setting up an object store as the cloud tier, you specify the local tier (aggregate) to use by attaching it to FabricPool. In ONTAP 9.5 and later, you can also attach local tiers (aggregates) that contain qualified FlexGroup volume constituents.

About this task

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. However, you can use [FabricPool mirror](#) to attach a local tier to a different cloud tier.

Before you begin

When you use the ONTAP CLI to set up an aggregate for FabricPool, the aggregate must already exist.




When you use System Manager to set up a local tier for FabricPool, you can create the local tier and set it up to use for FabricPool at the same time.

Steps

You can attach a local tier (aggregate) to a FabricPool object store with ONTAP System Manager or the ONTAP CLI.

System Manager

1. Navigate to **Storage > Tiers**, select a cloud tier, then click .
2. Select **Attach local tiers**.
3. Under **Add as Primary** verify that the volumes are eligible to attach.
4. If necessary, select **Convert volumes to thin provisioned**.
5. Click **Save**.

CLI

To attach an object store to an aggregate with the CLI:

1. **Optional:** To see how much data in a volume is inactive, follow the steps in [Determining how much data in a volume is inactive by using inactive data reporting](#).

Seeing how much data in a volume is inactive can help you decide which aggregate to use for FabricPool.

2. Attach the object store to an aggregate by using the `storage aggregate object-store attach` command.

If the aggregate has never been used with FabricPool and it contains existing volumes, then the volumes are assigned the default `snapshot-only` tiering policy.

```
cluster1::> storage aggregate object-store attach -aggregate myaggr
-object-store-name Amazon01B1
```

You can use the `allow-flexgroup true` option to attach aggregates that contain FlexGroup volume constituents.

3. Display the object store information and verify that the attached object store is available by using the `storage aggregate object-store show` command.

```
cluster1::> storage aggregate object-store show
```

Aggregate	Object Store Name	Availability State
-----	-----	-----
myaggr	Amazon01B1	available

Tier data to local bucket


Beginning with ONTAP 9.8, you can tier data to local object storage using ONTAP S3.

Tiering data to a local bucket provides a simple alternative to moving data to a different local tier. This procedure uses an existing bucket on the local cluster, or you can let ONTAP automatically create a new storage VM and a new bucket.

Keep in mind that once you attach to a local tier (aggregate) the cloud tier cannot be unattached.

An S3 license is required for this workflow, which creates a new S3 server and new bucket, or uses existing ones. A FabricPool license is not required for this workflow.

Step

1. Tier data to a local bucket: click **Tiers**, select a tier, then click .
2. If necessary, enable thin provisioning.
3. Choose an existing tier or create a new one.
4. If necessary, edit the existing tiering policy.

Manage FabricPool

Manage FabricPool overview

To help you with your storage tiering needs, ONTAP enables you to display how much data in a volume is inactive, add or move volumes to FabricPool, monitor the space utilization for FabricPool, or modify a volume's tiering policy or tiering minimum cooling period.

Determine how much data in a volume is inactive by using inactive data reporting

Seeing how much data in a volume is inactive enables you to make good use of storage tiers. Information in inactive data reporting helps you decide which aggregate to use for FabricPool, whether to move a volume in to or out of FabricPool, or whether to modify the tiering policy of a volume.

What you'll need

You must be running ONTAP 9.4 or later to use the inactive data reporting functionality.

About this task

- Inactive data reporting is not supported on some aggregates.

You cannot enable inactive data reporting when FabricPool cannot be enabled, including the following instances:


- Root aggregates
- MetroCluster aggregates running ONTAP versions earlier than 9.7
- Flash Pool (hybrid aggregates, or SnapLock aggregates)
- Inactive data reporting is enabled by default on aggregates where any volumes have adaptive compression enabled.
- Inactive data reporting is enabled by default on all SSD aggregates in ONTAP 9.6.
- Inactive data reporting is enabled by default on FabricPool aggregate in ONTAP 9.4 and ONTAP 9.5.
- You can enable inactive data reporting on non-FabricPool aggregates using the ONTAP CLI, including HDD aggregates, beginning with ONTAP 9.6.

Procedure

You can determine how much data is inactive with ONTAP System Manager or the ONTAP CLI.

System Manager

1. Choose one of the following options:

- When you have existing HDD aggregates, navigate to **Storage > Tiers** and click  for the aggregate on which you want to enable inactive data reporting.
- When no cloud tiers are configured, navigate to **Dashboard** and click the **Enable inactive data reporting** link under **Capacity**.

CLI

To enable inactive data reporting with the CLI:

1. If the aggregate for which you want to see inactive data reporting is not used in FabricPool, enable inactive data reporting for the aggregate by using the `storage aggregate modify` command with the `-is-inactive-data-reporting-enabled true` parameter.

```
cluster1::> storage aggregate modify -aggregate aggr1 -is-inactive
-data-reporting-enabled true
```

You need to explicitly enable the inactive data reporting functionality on an aggregate that is not used for FabricPool.

You cannot and do not need to enable inactive data reporting on a FabricPool-enabled aggregate because the aggregate already comes with inactive data reporting. The `-is-inactive-data-reporting-enabled` parameter does not work on FabricPool-enabled aggregates.

The `-fields is-inactive-data-reporting-enabled` parameter of the `storage aggregate show` command shows whether inactive data reporting is enabled on an aggregate.

2. To display how much data is inactive on a volume, use the `volume show` command with the `-fields performance-tier-inactive-user-data,performance-tier-inactive-user-data-percent` parameter.

```
cluster1::> volume show -fields performance-tier-inactive-user-
data,performance-tier-inactive-user-data-percent

vserver volume performance-tier-inactive-user-data performance-tier-
inactive-user-data-percent
-----
-----
vsim1    vol0    0B                                0%
vs1      vs1rv1 0B                                0%
vs1      vv1     10.34MB                             0%
vs1      vv2     10.38MB                             0%
4 entries were displayed.
```

- The `performance-tier-inactive-user-data` field displays how much user data stored in the aggregate is inactive.

- The `performance-tier-inactive-user-data-percent` field displays what percent of the data is inactive across the active file system and Snapshot copies.
- For an aggregate that is not used for FabricPool, inactive data reporting uses the tiering policy to decide how much data to report as cold.
 - For the `none` tiering policy, 31 days is used.

- For the `snapshot-only` and `auto`, inactive data reporting uses `tiering-minimum-cooling-days`.

- For the `ALL` policy, inactive data reporting assumes the data will tier within a day.

Until the period is reached, the output shows “-” for the amount of inactive data instead of a value.

- On a volume that is part of FabricPool, what ONTAP reports as inactive depends on the tiering policy that is set on a volume.
 - For the `none` tiering policy, ONTAP reports the amount of the entire volume that is inactive for at least 31 days. You cannot use the `-tiering-minimum-cooling-days` parameter with the `none` tiering policy.
 - For the `ALL`, `snapshot-only`, and `auto` tiering policies, inactive data reporting is not supported.

Manage volumes for FabricPool

Create a volume for FabricPool

You can add volumes to FabricPool by creating new volumes directly in the FabricPool-enabled aggregate or by moving existing volumes from another aggregate to the FabricPool-enabled aggregate.

When you create a volume for FabricPool, you have the option to specify a tiering policy. If no tiering policy is specified, the created volume uses the default `snapshot-only` tiering policy. For a volume with the `snapshot-only` or `auto` tiering policy, you can also specify the tiering minimum cooling period.

What you'll need

- Setting a volume to use the `auto` tiering policy or specifying the tiering minimum cooling period requires ONTAP 9.4 or later.
- Using FlexGroup volumes requires ONTAP 9.5 or later.
- Setting a volume to use the `all` tiering policy requires ONTAP 9.6 or later.
- Setting a volume to use the `-cloud-retrieval-policy` parameter requires ONTAP 9.8 or later.

Steps

1. Create a new volume for FabricPool by using the `volume create` command.

- The `-tiering-policy` optional parameter enables you to specify the tiering policy for the volume.

You can specify one of the following tiering policies:

- `snapshot-only` (default)

- `auto`
- `all`
- `backup` (deprecated)
- `none`

Types of FabricPool tiering policies

- The `-cloud-retrieval-policy` optional parameter enables cluster administrators with the advanced privilege level to override the default cloud migration or retrieval behavior controlled by the tiering policy.

You can specify one of the following cloud retrieval policies:

- `default`

The tiering policy determines what data is pulled back, so there is no change to cloud data retrieval with `default` `cloud-retrieval-policy`. This means the behavior is the same as in pre-ONTAP 9.8 releases:

- If the tiering policy is `none` or `snapshot-only`, then “default” means that any client-driven data read is pulled from the cloud tier to performance tier.
- If the tiering policy is `auto`, then any client-driven random read is pulled but not sequential reads.
- If the tiering policy is `all` then no client-driven data is pulled from the cloud tier.

- `on-read`

All client-driven data reads are pulled from the cloud tier to performance tier.

- `never`

No client-driven data is pulled from the cloud tier to performance tier

- `promote`

- For tiering policy `none`, all cloud data is pulled from the cloud tier to the performance tier
- For tiering policy `snapshot-only`, all active filesystem data is pulled from the cloud tier to the performance tier.

- The `-tiering-minimum-cooling-days` optional parameter in the advanced privilege level enables you to specify the tiering minimum cooling period for a volume that uses the `snapshot-only` or `auto` tiering policy.

Beginning with ONTAP 9.8, you can specify a value between 2 and 183 for the tiering minimum cooling days. If you are using a version of ONTAP earlier than 9.8, you can specify a value between 2 and 63 for the tiering minimum cooling days.

Example of creating a volume for FabricPool

The following example creates a volume called “myvol1” in the “myFabricPool” FabricPool-enabled aggregate. The tiering policy is set to `auto` and the tiering minimum cooling period is set to 45 days:


```
cluster1::*> volume create -vserver myVS -aggregate myFabricPool
-volume myvol1 -tiering-policy auto -tiering-minimum-cooling-days 45
```

Related information

[FlexGroup volumes management](#)

Move a volume to FabricPool

When you move a volume to FabricPool, you have the option to specify or change the tiering policy for the volume with the move. Beginning with ONTAP 9.8, when you move a non-FabricPool volume with inactive data reporting enabled, FabricPool uses a heat map to read tierable blocks, and moves cold data to the capacity tier on the FabricPool destination.

What you'll need

You must understand how changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

[What happens to the tiering policy when you move a volume](#)

About this task

If a non-FabricPool volume has inactive data reporting enabled, when you move a volume with tiering-policy `auto` or `snapshot-only` to a FabricPool, FabricPool reads the temperature tierable blocks from a heat map file and uses that temperature to move the cold data directly to the capacity tier on the FabricPool destination.

You should not use the `-tiering-policy` option on volume move if you are using ONTAP 9.8 and you want FabricPools to use inactive data reporting information to move data directly to the capacity tier. Using this option causes FabricPools to ignore the temperature data and instead follow the move behavior of releases prior to ONTAP 9.8.

Step

1. Use the `volume move start` command to move a volume to FabricPool.

The `-tiering-policy` optional parameter enables you to specify the tiering policy for the volume.

You can specify one of the following tiering policies:

- `snapshot-only` (default)
- `auto`
- `all`
- `none`

[Types of FabricPool tiering policies](#)

Example of moving a volume to FabricPool

The following example moves a volume named "myvol2" of the "vs1" SVM to the "dest_FabricPool" FabricPool-enabled aggregate. The volume is explicitly set to use the `none` tiering policy:

```
cluster1::> volume move start -vserver vs1 -volume myvol2  
-destination-aggregate dest_FabricPool -tiering-policy none
```

Enable and disable volumes to write directly to the cloud

Beginning with ONTAP 9.14.1, you can enable and disable writing directly to the cloud on a new or existing volume in a FabricPool to allow NFS clients to write data directly to the cloud without waiting for tiering scans. SMB clients still write to the performance tier in a cloud write enabled volume. Cloud-write mode is disabled by default.

Having the ability to write directly to the cloud is helpful for cases like migrations, for example, where large amounts of data are transferred to a cluster than the cluster can support on the local tier. Without cloud-write mode, during a migration, smaller amounts of data are transferred, then tiered, then transferred and tiered again, until the migration is complete. Using cloud-write mode, this type of management is no longer required because the data is never transferred to the local tier.

Before you begin

- You should be a cluster or SVM administrator.
- You must be at the advanced privilege level.
- The volume must be a read-write type volume.
- The volume must have the ALL tiering policy.

Enable writing directly to the cloud during volume creation

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Create a volume and enable cloud-write mode:

```
volume create -volume <volume name> -is-cloud-write-enabled <true|false>  
-aggregate <local tier name>
```

The following example creates a volume named vol1 with cloud write enabled on the FabricPool local tier (aggr1):

```
volume create -volume vol1 -is-cloud-write-enabled true -aggregate aggr1
```

Enable writing directly to the cloud on an existing volume

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Modify a volume to enable cloud-write mode:

```
volume modify -volume <volume name> -is-cloud-write-enabled <true|false>  
-aggregate <local tier name>
```

The following example modifies a volume named vol1 with cloud write enabled on the FabricPool local tier (aggr1):

```
volume modify -volume vol1 -is-cloud-write-enabled true -aggregate aggr1
```

Disable writing directly to the cloud on a volume

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Disable cloud-write mode:

```
volume modify -volume <volume name> -is-cloud-write-enabled <true|false>  
-aggregate <aggregate name>
```

The following example creates a volume named vol1 with cloud write enabled:

```
volume modify -volume vol1 -is-cloud-write-enabled false -aggregate  
aggr1
```

Enable and disable aggressive read-ahead mode

Beginning with ONTAP 9.14.1, you can enable and disable aggressive read-ahead mode on volumes in FabricPools that provide support for media and entertainment, such as movie streaming workloads. Aggressive read-ahead mode is available in ONTAP 9.14.1 on all on-premises platforms that support FabricPool. The feature is disabled by default.

About this task

The `aggressive-readahead-mode` command has two options:

- `none`: read-ahead is disabled.

- `file_prefetch`: the system reads the entire file into memory ahead of the client application.

Before you begin

- You should be a cluster or SVM administrator.
- You must be at the advanced privilege level.

Enable aggressive read-ahead mode during volume creation

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Create a volume and enable aggressive read-ahead mode:

```
volume create -volume <volume name> -aggressive-readahead-mode  
<none|file_prefetch>
```

The following example creates a volume named `vol1` with aggressive read-ahead enabled with the `file_prefetch` option:

```
volume create -volume vol1 -aggressive-readahead-mode file_prefetch
```

Disable aggressive read-ahead mode

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Disable aggressive read-ahead mode:

```
volume modify -volume <volume name> -aggressive-readahead-mode none
```

The following example modifies a volume named `vol1` to disable aggressive read-ahead mode:

```
volume modify -volume vol1 -aggressive-readahead-mode none
```

View aggressive read-ahead mode on a volume

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. View the aggressive read-ahead mode:

```
volume show -fields aggressive-readahead-mode
```

Object tagging using user-created custom tags

Object tagging using user-created custom tags overview

Beginning with ONTAP 9.8, FabricPool supports object tagging using user-created custom tags to enable you to classify and sort objects for easier management. If you are a user with the admin privilege level, you can create new object tags, and modify, delete, and view existing tags.

Assign a new tag during volume creation

You can create a new object tag when you want to assign one or more tags to new objects that are tiered from a new volume you create. You can use tags to help you classify and sort tiering objects for easier data management. Beginning with ONTAP 9.8, you can use System Manager to create object tags.

About this task

You can set tags only on FabricPool volumes attached to StorageGRID. These tags are retained during a volume move.

- A maximum of 4 tags per volume is allowed
- In the CLI, each object tag must be a key-value pair separated by an equal sign ("")
- In the CLI, multiple tags must be separated by a comma (",")
- Each tag value can contain a maximum of 127 characters
- Each tag key must start with either an alphabetic character or an underscore.

Keys must contain only alphanumeric characters and underscores, and the maximum number of characters allowed is 127.

Procedure

You can assign object tags with ONTAP System Manager or the ONTAP CLI.

System Manager

1. Navigate to **Storage > Tiers**.
2. Locate a storage tier with volumes you want to tag.
3. Click the **Volumes** tab.
4. Locate the volume you want to tag and in the **Object Tags** column select **Click to enter tags**.
5. Enter a key and value.
6. Click **Apply**.

CLI

1. Use the `volume create` command with the `-tiering-object-tags` option to create a new volume with the specified tags. You can specify multiple tags in comma-separated pairs:

```
volume create [ -vserver <vserver name> ] -volume <volume_name>
-tiering-object-tags <key1=value1> [
    ,<key2=value2>,<key3=value3>,<key4=value4> ]
```

The following example creates a volume named `fp_volume1` with three object tags.

```
vol create -volume fp_volume1 -vserver vs0 -tiering-object-tags
project=fabricpool,type=abc,content=data
```

Modify an existing tag

You can change the name of a tag, replace tags on existing objects in the object store, or add a different tag to new objects that you plan to add later.

About this task

Using the `volume modify` command with the `-tiering-object-tags` option replaces existing tags with the new value you provide.

Procedure

System Manager

1. Navigate to **Storage > Tiers**.
2. Locate a storage tier with volumes containing tags you want to modify.
3. Click the **Volumes** tab.
4. Locate the volume with tags you want to modify, and in the **Object Tags** column click the tag name.
5. Modify the tag.
6. Click **Apply**.

CLI

1. Use the `volume modify` command with the `-tiering-object-tags` option to modify an existing tag.

```
volume modify [ -vserver <vserver name> ] -volume <volume_name>  
-tiering-object-tags <key1=value1> [ ,<key2=value2>,  
<key3=value3>,<key4=value4> ]
```

The following example changes the name of the existing tag `type=abc` to `type=xyz`.

```
vol create -volume fp_volume1 -vserver vs0 -tiering-object-tags  
project=fabricpool,type=xyz,content=data
```

Delete a tag

You can delete object tags when you no longer want them set on a volume or on objects in the object store.

Procedure

You can delete object tags with ONTAP System Manager or the ONTAP CLI.

System Manager

1. Navigate to **Storage > Tiers**.
2. Locate a storage tier with volumes containing tags you want to delete.
3. Click the **Volumes** tab.
4. Locate the volume with tags you want to delete, and in the **Object Tags** column click the tag name.
5. To delete the tag, click the trash can icon.
6. Click **Apply**.

CLI

1. Use the `volume modify` command with the `-tiering-object-tags` option followed by an empty value (`""`) to delete an existing tag.

The following example deletes the existing tags on `fp_volume1`.

```
vol modify -volume fp_volume1 -vserver vs0 -tiering-object-tags ""
```

View existing tags on a volume

You can view the existing tags on a volume to see what tags are available before appending new tags to the list.

Step

1. Use the `volume show` command with the `-tiering-object-tags` option to view existing tags on a volume.

```
volume show [ -vserver <vserver name> ] -volume <volume_name> -fields  
-tiering-object-tags
```

Check object tagging status on FabricPool volumes

You can check if tagging is complete on one or more FabricPool volumes.

Step

1. Use the `vol show` command with the `-fieldsneeds-object-retagging` option to see if tagging is in progress, if it has completed, or if tagging is not set.

```
vol show -fields needs-object-retagging [ -instance | -volume <volume  
name>]
```

One of the following values is displayed:

- `true` — the object tagging scanner has not yet to run or needs to run again for this volume
- `false` — the object tagging scanner has completed tagging for this volume
- `<->` — the object tagging scanner is not applicable for this volume. This happens for volumes that are not residing on FabricPools.

Monitor the space utilization for FabricPool

You need to know how much data is stored in the performance and cloud tiers for FabricPool. That information helps you determine whether you need to change the tiering policy of a volume, increase the FabricPool licensed usage limit, or increase the storage space of the cloud tier.

Steps

1. Monitor the space utilization for FabricPool-enabled aggregates by using one of the following commands to display the information:

If you want to display...	Then use this command:
The used size of the cloud tier in an aggregate	<code>storage aggregate show with the -instance parameter</code>
Details of space utilization within an aggregate, including the object store's referenced capacity	<code>storage aggregate show-space with the -instance parameter</code>
Space utilization of the object stores that are attached to the aggregates, including how much license space is being used	<code>storage aggregate object-store show-space</code>
A list of volumes in an aggregate and the footprints of their data and metadata	<code>volume show-footprint</code>

In addition to using CLI commands, you can use Active IQ Unified Manager (formerly OnCommand Unified Manager), along with FabricPool Advisor, which is supported on ONTAP 9.4 and later clusters, or System Manager to monitor the space utilization.

The following example shows ways of displaying space utilization and related information for FabricPool:

```
cluster1::> storage aggregate show-space -instance
```

```
Aggregate: MyFabricPool
...
Aggregate Display Name:
MyFabricPool
...
Total Object Store Logical Referenced
Capacity: -
Object Store Logical Referenced Capacity
Percentage: -
...
Object Store
Size: -
Object Store Space Saved by Storage
Efficiency: -
Object Store Space Saved by Storage Efficiency
Percentage: -
Total Logical Used
Size: -
Logical Used
Percentage: -
Logical Unreferenced
Capacity: -
Logical Unreferenced
Percentage: -
```

```
cluster1::> storage aggregate show -instance
```

```
Aggregate: MyFabricPool
...
Composite: true
Capacity Tier Used Size:
...
```

```
cluster1::> volume show-footprint
```

```
Vserver : vs1
```

```
Volume : rootvol
```

Feature	Used	Used%
Volume Footprint	KB	%
Volume Guarantee	MB	%
Flexible Volume Metadata	KB	%
Delayed Frees	KB	%
Total Footprint	MB	%

```
Vserver : vs1
```

```
Volume : vol
```

Feature	Used	Used%
Volume Footprint	KB	%
Footprint in Performance Tier	KB	%
Footprint in Amazon01	KB	%
Flexible Volume Metadata	MB	%
Delayed Frees	KB	%
Total Footprint	MB	%
...		

2. Take one of the following actions as needed:

If you want to...	Then...
Change the tiering policy of a volume	Follow the procedure in Managing storage tiering by modifying a volume's tiering policy or tiering minimum cooling period .
Increase the FabricPool licensed usage limit	Contact your NetApp or partner sales representative. NetApp Support
Increase the storage space of the cloud tier	Contact the provider of the object store that you use for the cloud tier.

Manage storage tiering by modifying a volume's tiering policy or tiering minimum cooling period

You can change the tiering policy of a volume to control whether data is moved to the

cloud tier when it becomes inactive (*cold*). For a volume with the `snapshot-only` or `auto` tiering policy, you can also specify the tiering minimum cooling period that user data must remain inactive before it is moved to the cloud tier.

What you'll need

Changing a volume to the `auto` tiering policy or modifying the tiering minimum cooling period requires ONTAP 9.4 or later.

About this task

Changing the tiering policy of a volume changes only the subsequent tiering behavior for the volume. It does not retroactively move data to the cloud tier.

Changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

What happens when you modify the tiering policy of a volume in FabricPool

Steps

1. Modify the tiering policy for an existing volume by using the `volume modify` command with the `-tiering-policy` parameter:

You can specify one of the following tiering policies:

- `snapshot-only` (default)
- `auto`
- `all`
- `none`

Types of FabricPool tiering policies

2. If the volume uses the `snapshot-only` or `auto` tiering policy and you want to modify the tiering minimum cooling period, use the `volume modify` command with the `-tiering-minimum-cooling-days` optional parameter in the advanced privilege level.

You can specify a value between 2 and 183 for the tiering minimum cooling days. If you are using a version of ONTAP earlier than 9.8, you can specify a value between 2 and 63 for the tiering minimum cooling days.

Example of modifying the tiering policy and the tiering minimum cooling period of a volume

The following example changes the tiering policy of the volume “myvol” in the SVM “vs1” to `auto` and the tiering minimum cooling period to 45 days:

```
cluster1::> volume modify -vserver vs1 -volume myvol  
-tiering-policy auto -tiering-minimum-cooling-days 45
```

Archive volumes with FabricPool (video)

This video shows a quick overview of using System Manager to archive a volume to a cloud tier with FabricPool.

Related information

[NetApp TechComm TV: FabricPool playlist](#)

Use cloud migration controls to override a volume's default tiering policy

You can change a volume's default tiering policy for controlling user data retrieval from the cloud tier to performance tier by using the `-cloud-retrieval-policy` option introduced in ONTAP 9.8.

What you'll need

- Modifying a volume using the `-cloud-retrieval-policy` option requires ONTAP 9.8 or later.
- You must have the advanced privilege level to perform this operation.
- You should understand the behavior of tiering policies with `-cloud-retrieval-policy`.

[How tiering policies work with cloud migration](#)

Step

1. Modify the tiering policy behavior for an existing volume by using the `volume modify` command with the `-cloud-retrieval-policy` option:

```
volume create -volume <volume_name> -vserver <vserver_name> - tiering-  
policy <policy_name> -cloud-retrieval-policy
```

```
vol modify -volume fp_volume4 -vserver vs0 -cloud-retrieval-policy  
promote
```

Promote data to the performance tier

Promote data to the performance tier overview

Beginning with ONTAP 9.8, if you are a cluster administrator at the advanced privilege level, you can proactively promote data to the performance tier from the cloud tier using a combination of the `tiering-policy` and the `cloud-retrieval-policy` setting.

About this task

You might do this if you want to stop using FabricPool on a volume, or if you have a `snapshot-only` tiering policy and you want to bring restored Snapshot copy data back to the performance tier.

Promote all data from a FabricPool volume to the performance tier

You can proactively retrieve all data on a FabricPool volume in the Cloud and promote it to the performance tier.

Step

1. Use the `volume modify` command to set `tiering-policy` to `none` and `cloud-retrieval-policy` to `promote`.

```
volume modify -vserver <vserver-name> -volume <volume-name> -tiering  
-policy none -cloud-retrieval-policy promote
```

Promote file system data to the performance tier

You can proactively retrieve active file system data from a restored Snapshot copy in the cloud tier and promote it to the performance tier.

Step

1. Use the `volume modify` command to set `tiering-policy` to `snapshot-only` and `cloud-retrieval-policy` to `promote`.

```
volume modify -vserver <vserver-name> -volume <volume-name> -tiering  
-policy snapshot-only cloud-retrieval-policy promote
```

Check the status of a performance tier promotion

You can check the status of performance tier promotion to determine when the operation is complete.

Step

1. Use the `volume object-store` command with the `tiering` option to check the status of the performance tier promotion.

```
volume object-store tiering show [ -instance | -fields <fieldname>, ...  
] [ -vserver <vserver name> ] *Vserver  
[[-volume] <volume name>] *Volume [ -node <nodename> ] *Node Name [ -vol  
-dsid <integer> ] *Volume DSID  
[ -aggregate <aggregate name> ] *Aggregate Name
```

```

volume object-store tiering show v1 -instance

Vserver: vs1
Volume: v1
Node Name: node1
Volume DSID: 1023
Aggregate Name: a1
State: ready
Previous Run Status: completed
Aborted Exception Status: -
Time Scanner Last Finished: Mon Jan 13 20:27:30 2020
Scanner Percent Complete: -
Scanner Current VBN: -
Scanner Max VBNs: -
Time Waiting Scan will be scheduled: -
Tiering Policy: snapshot-only
Estimated Space Needed for Promotion: -
Time Scan Started: -
Estimated Time Remaining for scan to complete: -
Cloud Retrieve Policy: promote

```

Trigger scheduled migration and tiering

Beginning with ONTAP 9.8, you can trigger a tiering scan request at any time when you prefer not to wait for the default tiering scan.

Step

1. Use the `volume object-store` command with the `trigger` option to request migration and tiering.

```

volume object-store tiering trigger [ -vserver <vserver name> ] *VServer
Name [-volume] <volume name> *Volume Name

```

Manage FabricPool mirrors

Manage FabricPool mirrors overview

To ensure data is accessible in data stores in the event of a disaster, and to enable you to replace a data store, you can configure a FabricPool mirror by adding a second data store to synchronously tier data to two data stores . You can add a second data store to new or existing FabricPool configurations, monitor the mirror status, display FabricPool mirror details, promote a mirror, and remove a mirror. You must be running ONTAP 9.7 or later.

Create a FabricPool mirror

To create a FabricPool mirror, you attach two object stores to a single FabricPool. You can create a FabricPool mirror either by attaching a second object store to an existing, single object store FabricPool configuration, or you can create a new, single object store FabricPool configuration and then attach a second object store to it. You can also create FabricPool mirrors on MetroCluster configurations.

What you'll need

- You must have already created the two object stores using the `storage aggregate object-store config` command.
- If you are creating FabricPool mirrors on MetroCluster configurations:
 - You must have already set up and configured the MetroCluster
 - You must have created the object store configurations on the selected cluster.

If you are creating FabricPool mirrors on both clusters in a MetroCluster configuration, you must have created object store configurations on both of the clusters.

- If you are not using on premises object stores for MetroCluster configurations, you should ensure that one of the following scenarios exists:
 - Object stores are in different availability zones
 - Object stores are configured to keep copies of objects in multiple availability zones

[Setting up object stores for FabricPool in a MetroCluster configuration](#)

About this task

The object store you use for the FabricPool mirror must be different from the primary object store.

The procedure for creating a FabricPool mirror is the same for both MetroCluster and non-MetroCluster configurations.

Steps

1. If you are not using an existing FabricPool configuration, create a new one by attaching an object store to an aggregate using the `storage aggregate object-store attach` command.

This example creates a new FabricPool by attaching an object store to an aggregate.

```
cluster1::> storage aggregate object-store attach -aggregate aggr1 -name  
my-store-1
```

2. Attach a second object store to the aggregate using the `storage aggregate object-store mirror` command.

This example attaches a second object store to an aggregate to create a FabricPool mirror.


```
cluster1::> storage aggregate object-store mirror -aggregate aggr1 -name
my-store-2
```

Monitor FabricPool mirror resync status

When you replace a primary object store with a mirror, you might have to wait for the mirror to resync with the primary data store.

About this task

If the FabricPool mirror is in sync, no entries are displayed.

Step

1. Monitor mirror resync status using the `storage aggregate object-store show-resync-status` command.

```
aggr1::> storage aggregate object-store show-resync-status
-aggregate aggr1
```

Aggregate	Primary	Mirror	Complete Percentage
-----	-----	-----	-----
aggr1	my-store-1	my-store-2	40%

Display FabricPool mirror details

You can display details about a FabricPool mirror to see what object stores are in the configuration and whether the object store mirror is in sync with the primary object store.

Step

1. Display information about a FabricPool mirror using the `storage aggregate object-store show` command.

This example displays the details about the primary and mirror object stores in a FabricPool mirror.

```
cluster1::> storage aggregate object-store show
```

Aggregate	Object Store Name	Availability	Mirror Type
-----	-----	-----	-----
aggr1	my-store-1	available	primary
	my-store-2	available	mirror

This example displays details about the FabricPool mirror, including whether the mirror is degraded due to a resync operation.

```
cluster1::> storage aggregate object-store show -fields mirror-type,is-mirror-degraded
```

aggregate	object-store-name	mirror-type	is-mirror-degraded
aggr1	my-store-1	primary	-
	my-store-2	mirror	false

Promote a FabricPool mirror

You can reassign the object store mirror as the primary object store by promoting it. When the object store mirror becomes the primary, the original primary automatically becomes the mirror.

What you'll need

- The FabricPool mirror must be in sync
- The object store must be operational

About this task

You can replace the original object store with an object store from a different cloud provider. For instance, your original mirror might be an AWS object store, but you can replace it with an Azure object store.

Step

1. Promote an object store mirror by using the `storage aggregate object-store modify -aggregate` command.

```
cluster1::> storage aggregate object-store modify -aggregate aggr1 -name my-store-2 -mirror-type primary
```

Remove a FabricPool mirror

You can remove a FabricPool mirror if you no longer need to replicate an object store.

What you'll need

The primary object store must be operational, otherwise, the command fails.

Step

1. Remove an object store mirror in a FabricPool by using the `storage aggregate object-store unmirror -aggregate` command.

```
cluster1::> storage aggregate object-store unmirror -aggregate aggr1
```

Replace an existing object store using a FabricPool mirror

You can use FabricPool mirror technology to replace one object store with another one. The new object store does not have to use the same cloud provider as the original object store.

About this task

You can replace the original object store with an object store that uses a different cloud provider. For instance, your original object store might use AWS as the cloud provider, but you can replace it with an object store that uses Azure as the cloud provider, and vice versa. However, the new object store must retain the same object size as the original.

Steps

1. Create a FabricPool mirror by adding a new object store to an existing FabricPool using the `storage aggregate object-store mirror` command.

```
cluster1::> storage aggregate object-store mirror -aggregate aggr1 -name my-AZURE-store
```

2. Monitor the mirror resync status using the `storage aggregate object-store show-resync-status` command.

```
cluster1::> storage aggregate object-store show-resync-status -aggregate aggr1
```

Aggregate	Primary	Mirror	Complete Percentage
-----	-----	-----	-----
aggr1	my-AWS-store	my-AZURE-store	40%

3. Verify the mirror is in sync using the `storage aggregate object-store> show -fields mirror-type,is-mirror-degraded` command.

```
cluster1::> storage aggregate object-store show -fields mirror-type,is-mirror-degraded
```

aggregate	object-store-name	mirror-type	is-mirror-degraded
-----	-----	-----	-----
aggr1	my-AWS-store	primary	-
	my-AZURE-store	mirror	false

4. Swap the primary object store with the mirror object store using the `storage aggregate object-store modify` command.

```
cluster1::> storage aggregate object-store modify -aggregate aggr1 -name
my-AZURE-store -mirror-type primary
```

5. Display details about the FabricPool mirror using the `storage aggregate object-store show -fields mirror-type,is-mirror-degraded` command.

This example displays the information about the FabricPool mirror, including whether the mirror is degraded (not in sync).

```
cluster1::> storage aggregate object-store show -fields mirror-type, is-
mirror-degraded
```

aggregate	object-store-name	mirror-type	is-mirror-degraded
-----	-----	-----	-----
aggr1	my-AZURE-store	primary	-
	my-AWS-store	mirror	false

6. Remove the FabricPool mirror using the `storage aggregate object-store unmirror` command.

```
cluster1::> storage aggregate object-store unmirror -aggregate aggr1
```

7. Verify that the FabricPool is back in a single object store configuration using the `storage aggregate object-store show -fields mirror-type,is-mirror-degraded` command.

```
cluster1::> storage aggregate object-store show -fields mirror-type,is-
mirror-degraded
```

aggregate	object-store-name	mirror-type	is-mirror-degraded
-----	-----	-----	-----
aggr1	my-AZURE-store	primary	-

Replace a FabricPool mirror on a MetroCluster configuration

If one of the object stores in a FabricPool mirror is destroyed or becomes permanently unavailable on a MetroCluster configuration, you can make the object store the mirror if it is not the mirror already, remove the damaged object store from FabricPool mirror, and then add a new object store mirror to the FabricPool.

Steps

1. If the damaged object store is not already the mirror, make the object store the mirror with the `storage aggregate object-store modify` command.

```
storage aggregate object-store modify -aggregate -aggregate fp_aggr1_A01  
-name mccl_ostore1 -mirror-type mirror
```

2. Remove the object store mirror from the FabricPool by using the `storage aggregate object-store unmirror` command.

```
storage aggregate object-store unmirror -aggregate <aggregate name>  
-name mccl_ostore1
```

3. You can force tiering to resume on the primary data store after you remove the mirror data store by using the `storage aggregate object-store modify` with the `-force-tiering-on-metrocluster true` option.

The absence of a mirror interferes with the replication requirements of a MetroCluster configuration.

```
storage aggregate object-store modify -aggregate <aggregate name> -name  
mccl_ostore1 -force-tiering-on-metrocluster true
```

4. Create a replacement object store by using the `storage aggregate object-store config create` command.

```
storage aggregate object-store config create -object-store-name  
mccl_ostore3 -cluster clusterA -provider-type SGWS -server <SGWS-server-  
1> -container-name <SGWS-bucket-1> -access-key <key> -secret-password  
<password> -encrypt <true|false> -provider <provider-type> -is-ssl  
-enabled <true|false> ipspace <IPSpace>
```

5. Add the object store mirror to the FabricPool mirror using the `storage aggregate object-store mirror` command.

```
storage aggregate object-store mirror -aggregate aggr1 -name
mcc1_ostore3-mc
```

6. Display the object store information using the `storage aggregate object-store show` command.

```
storage aggregate object-store show -fields mirror-type,is-mirror-
degraded
```

aggregate	object-store-name	mirror-type	is-mirror-degraded
aggr1	mcc1_ostore1-mc	primary	-
	mcc1_ostore3-mc	mirror	true

7. Monitor the mirror resync status using the `storage aggregate object-store show-resync-status` command.

```
storage aggregate object-store show-resync-status -aggregate aggr1
```

Aggregate	Primary	Mirror	Complete Percentage
aggr1	mcc1_ostore1-mc	mcc1_ostore3-mc	40%

Commands for managing aggregates with FabricPool

You use the `storage aggregate object-store` commands to manage object stores for FabricPool. You use the `storage aggregate` commands to manage aggregates for FabricPool. You use the `volume` commands to manage volumes for FabricPool.

If you want to...	Use this command:
Define the configuration for an object store so that ONTAP can access it	<code>storage aggregate object-store config create</code>
Modify object store configuration attributes	<code>storage aggregate object-store config modify</code>
Rename an existing object store configuration	<code>storage aggregate object-store config rename</code>

Delete the configuration of an object store	<code>storage aggregate object-store config delete</code>
Display a list of object store configurations	<code>storage aggregate object-store config show</code>
Attach a second object store to a new or existing FabricPool as a mirror	<code>storage aggregate object-store mirror</code> with the <code>-aggregate</code> and <code>-name</code> parameter in the admin privilege level
Remove an object store mirror from an existing FabricPool mirror	<code>storage aggregate object-store unmirror</code> with the <code>-aggregate</code> and <code>-name</code> parameter in the admin privilege level
Monitor FabricPool mirror resync status	<code>storage aggregate object-store show-resync-status</code>
Display FabricPool mirror details	<code>storage aggregate object-store show</code>
Promote an object store mirror to replace a primary object store in a FabricPool mirror configuration	<code>storage aggregate object-store modify</code> with the <code>-aggregate</code> parameter in the admin privilege level
Test the latency and performance of an object store without attaching the object store to an aggregate	<code>storage aggregate object-store profiler start</code> with the <code>-object-store-name</code> and <code>-node</code> parameter in the advanced privilege level
Monitor the object store profiler status	<code>storage aggregate object-store profiler show</code> with the <code>-object-store-name</code> and <code>-node</code> parameter in the advanced privilege level
Abort the object store profiler when it is running	<code>storage aggregate object-store profiler abort</code> with the <code>-object-store-name</code> and <code>-node</code> parameter in the advanced privilege level
Attach an object store to an aggregate for using FabricPool	<code>storage aggregate object-store attach</code>
Attach an object store to an aggregate that contains a FlexGroup volume for using FabricPool	<code>storage aggregate object-store attach</code> with the <code>allow-flexgroup true</code>
Display details of the object stores that are attached to FabricPool-enabled aggregates	<code>storage aggregate object-store show</code>

Display the aggregate fullness threshold used by the tiering scan	<code>storage aggregate object-store show</code> with the <code>-fields tiering-fullness-threshold</code> parameter in the advanced privilege level
Display space utilization of the object stores that are attached to FabricPool-enabled aggregates	<code>storage aggregate object-store show-space</code>
Enable inactive data reporting on an aggregate that is not used for FabricPool	<code>storage aggregate modify</code> with the <code>-is -inactive-data-reporting-enabled true</code> parameter
Display whether inactive data reporting is enabled on an aggregate	<code>storage aggregate show</code> with the <code>-fields is-inactive-data-reporting-enabled</code> parameter
Display information about how much user data is cold within an aggregate	<code>storage aggregate show-space</code> with the <code>-fields performance-tier-inactive-user-data,performance-tier-inactive-user-data-percent</code> parameter
<p>Create a volume for FabricPool, including specifying the following:</p> <ul style="list-style-type: none"> • The tiering policy • The tiering minimum cooling period (for the <code>snapshot-only</code> or <code>auto</code> tiering policy) 	<p><code>volume create</code></p> <ul style="list-style-type: none"> • You use the <code>-tiering-policy</code> parameter to specify the tiering policy. • You use the <code>-tiering-minimum-cooling -days</code> parameter in the advanced privilege level to specify the tiering minimum cooling period.
<p>Modify a volume for FabricPool, including modifying the following:</p> <ul style="list-style-type: none"> • The tiering policy • The tiering minimum cooling period (for the <code>snapshot-only</code> or <code>auto</code> tiering policy) 	<p><code>volume modify</code></p> <ul style="list-style-type: none"> • You use the <code>-tiering-policy</code> parameter to specify the tiering policy. • You use the <code>-tiering-minimum-cooling -days</code> parameter in the advanced privilege level to specify the tiering minimum cooling period.
<p>Display FabricPool information related to a volume, including the following:</p> <ul style="list-style-type: none"> • The tiering minimum cooling period • How much user data is cold 	<p><code>volume show</code></p> <ul style="list-style-type: none"> • You use the <code>-fields tiering-minimum-cooling-days</code> parameter in the advanced privilege level to display the tiering minimum cooling period. • You use the <code>-fields performance-tier-inactive-user-data,performance-tier-inactive-user-data-percent</code> parameter to display how much user data is cold.

Move a volume in to or out of FabricPool	volume move start You use the <code>-tiering-policy</code> optional parameter to specify the tiering policy for the volume.
Modify the threshold for reclaiming unreferenced space (the defragmentation threshold) for FabricPool	storage aggregate object-store modify with the <code>-unreclaimed-space-threshold</code> parameter in the advanced privilege level
<p>Modify the threshold for the percent full the aggregate becomes before the tiering scan begins tiering data for FabricPool</p> <p>FabricPool continues to tier cold data to a cloud tier until the local tier reaches 98% capacity.</p>	storage aggregate object-store modify with the <code>-tiering-fullness-threshold</code> parameter in the advanced privilege level
Display the threshold for reclaiming unreferenced space for FabricPool	storage aggregate object-store show or storage aggregate object-store show-space command with the <code>-unreclaimed-space-threshold</code> parameter in the advanced privilege level

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.