



Manage FabricPool

ONTAP 9

NetApp
February 20, 2026

Table of Contents

- Manage FabricPool 1
 - Analyze inactive ONTAP data with inactive data reporting 1
 - Procedure 1
 - Manage volumes for FabricPool 3
 - Create a volume on a FabricPool-enabled ONTAP local tier 3
 - Move a volume to a FabricPool-enabled ONTAP local tier 5
 - Enable ONTAP volumes in FabricPool to write directly to the cloud 7
 - Enable ONTAP volumes in FabricPool to perform aggressive read-aheads 8
- Manage ONTAP FabricPool volumes with user-created custom tags 10
 - Assign a new tag during volume creation 10
 - Modify an existing tag 11
 - Delete a tag 12
 - View existing tags on a volume 13
 - Check object tagging status on FabricPool volumes 13
- Monitor space utilization of a FabricPool-enabled ONTAP local tier 14
- Modify an ONTAP volume's tiering policy and minimum cooling period 17
- Archive volumes with FabricPool (video) 18
- Modify an ONTAP volume's default FabricPool tiering policy 18
- Set thresholds on ONTAP FabricPool per-node put rate 18
- Customize ONTAP FabricPool object deletion and defragmentation 19
 - Unreclaimed space threshold 20
 - Change the unreclaimed space threshold 20
- Promote ONTAP data to the performance tier 21
 - Promote all data from a FabricPool volume to the performance tier 21
 - Promote file system data to the performance tier 21
 - Check the status of a performance tier promotion 22
 - Trigger scheduled migration and tiering 22

Manage FabricPool

Analyze inactive ONTAP data with 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.

Before you begin

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 snapshots.
- 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.

Related information

- [storage aggregate modify](#)

Manage volumes for FabricPool

Create a volume on a FabricPool-enabled ONTAP local tier

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



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see [Disks and local tiers](#).

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.

Before you begin

- 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 local tier. 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 a FabricPool-enabled ONTAP local tier

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

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



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see [Disks and local tiers](#).

Destination local tier

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

Beginning with ONTAP 9.8, when a volume has [inactive data reporting](#) enabled, FabricPool will use the volume’s heat map to immediately queue cold data to begin tiering as soon as it is written to the destination local tier.

Prior to ONTAP 9.8, moving a volume to another local tier resets the inactivity period of blocks on the local tier. For example, a volume using the Auto volume tiering policy with data on the local tier that has been inactive for 20 days, but had not yet tiered, will have the temperature of the data reset to 0 days after a volume move.

Optimized volume moves

Beginning with ONTAP 9.6, if a volume move’s destination local tier uses the same bucket as the source local tier, data on the source volume that is stored in the bucket does not move back to the local tier. Tiered data stays at rest and only hot data needs to be moved from one local tier to another. This optimized volume move results in significant network efficiencies.

For example, a 300TB optimized volume move means that even though 300TB of cold data moves from one

local tier to another, it will not trigger 300TB of reads and 300TB of writes to the object store.

Unoptimized volume moves generate additional network and compute traffic (reads/GETs and writes/PUTs), increasing demands on the ONTAP cluster and object store, potentially raising costs when tiering to public object stores.

Some configurations are incompatible with optimized volume moves:



- Changing tiering policy during volume move
- Source and destination local tiers using different encryption keys
- FlexClone volumes
- FlexClone parent volumes
- MetroCluster (supports optimized volume moves in ONTAP 9.8 and later)
- Unsynchronized FabricPool Mirror buckets

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

Writing data to the local tier first improves the performance of the volume move and reduces cutover time. If a volume tiering policy is not specified when performing a volume move, the destination volume uses the tiering policy of the source volume.

If a different tiering policy is specified when performing the volume move, the destination volume is created with the specified tiering policy and the volume move is not optimized.

Volume metadata

Regardless of whether a volume move is optimized, ONTAP stores a significant amount of metadata about the location, storage efficiency, permissions, usage patterns, etc., of all data, both local and tiered. Metadata always stays on the local tier and is not tiered. When a volume is moved from one local tier to another, this information needs to be moved to the destination local tier as well.

Duration

Volume moves still take time to complete and the expectation should be that an optimized volume move will take approximately the same amount of time as moving an equal amount of non-tiered data.

It is important to understand that "throughput" reported by the `volume move show` command does not represent throughput in terms of data being moved from the cloud tier, but volume data being updated locally.



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

Steps

1. Use the `volume move start` command to move a volume from a source local tier to a destination local tier.

Example of moving a volume

The following example moves a volume named `myvol2` of `vs1` SVM to `dest_FabricPool1`, a FabricPool-enabled local tier.

```
cluster1::> volume move start -vserver vs1 -volume myvol2
-destination-aggregate dest_FabricPool
```

Enable ONTAP volumes in FabricPool 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 -vserver <svm name> -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 -vserver vs1 -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 -vserver <svm name> -volume <volume name> -is-cloud-write-enabled true
```

The following example modifies the volume named vol1 to enable cloud write:

```
volume modify -vserver vs1 -volume vol1 -is-cloud-write-enabled true
```

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 on a volume:

```
volume modify -vserver <svm name> -volume <volume name> -is-cloud-write-enabled false
```

The following example disables cloud write mode on the volume named vol1:

```
volume modify -vserver vs1 -volume vol1 -is-cloud-write-enabled false
```

Enable ONTAP volumes in FabricPool to perform aggressive read-aheads

Beginning with ONTAP 9.14.1, you can enable and disable aggressive read-ahead mode on volumes in FabricPools. In ONTAP 9.13.1, aggressive read-ahead mode was introduced only on cloud platforms. Beginning with ONTAP 9.14.1, aggressive read-ahead mode is available on all platforms that FabricPool supports, including on-premises platforms. The feature is disabled by default.

When aggressive read-ahead is *disabled*, FabricPool only reads the file blocks that a client application needs; it does not need to read the entire file. This can result in reduced network traffic, especially for large GB-sized and TB-sized files. *Enabling* aggressive read-ahead on a volume turns this functionality off, and FabricPool preemptively reads the entire file sequentially from the object store, increasing GET throughput and reducing

the latency of client reads on the file. By default, when tiered data is read sequentially it stays cold and is not written to the local tier.

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

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
```

Manage ONTAP FabricPool volumes with user-created custom tags

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 four 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.

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

Example 1. Steps

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.

Example 2. Steps

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 modify -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.

Example 3. Steps

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.

Steps

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.

Steps

1. Use the `vol show` command with the `-fields needs-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 space utilization of a FabricPool-enabled ONTAP local tier

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.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see [Disks and local tiers](#).

About this task

Beginning with ONTAP 9.18.1, the `storage aggregate show-space` command changes how Logical Referenced Capacity and Logical Unreferenced Capacity is reported. Logical Referenced Capacity reports referenced blocks in all objects and unreferenced blocks in fragmented objects. Logical Unreferenced Capacity reports only unreferenced blocks in objects that have crossed the fullness threshold and are eligible for object deletion and defragmentation.

For example, when you use the default aggregate fullness threshold of 40% for ONTAP S3 and StorageGRID, 60% of the blocks in an object must be unreferenced before the blocks are reported as unreferenced capacity.

In releases earlier than ONTAP 9.18.1, Logical Referenced Capacity reports referenced blocks in all objects (both full and fragmented objects). Logical Unreferenced Capacity reports unreferenced blocks in all objects.

Steps

1. Monitor the space utilization for FabricPool-enabled local tiers 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 a local tier	<code>storage aggregate show</code> with the <code>-instance</code> parameter
Details of space utilization within an local tiers, including the object store's referenced capacity	<code>storage aggregate show-space</code> with the <code>-instance</code> parameter
Space utilization of the object stores that are attached to the local tiers, including how much license space is being used	<code>storage aggregate object-store show-space</code>
A list of volumes in a local tier 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.

Related information

- [storage aggregate object](#)
- [storage aggregate show](#)

- [storage aggregate show-space](#)

Modify an ONTAP volume's tiering policy and 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.

Before you begin

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



When in an SVM DR relationship, source and destination volumes do not need to use FabricPool aggregates, but they must use the same tiering policy.

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.

[NetApp video: Archiving volumes with FabricPool \(backup + volume move\)](#)

Related information

[NetApp TechComm TV: FabricPool playlist](#)

Modify an ONTAP volume's default FabricPool 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.

Before you begin

- 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
```

Set thresholds on ONTAP FabricPool per-node put rate

As a storage admin, you can use PUT throttling to set an upper threshold on the maximum per-node put rate.

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

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



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

Before you begin

Advanced privilege level is required.

Steps

1. Throttle FabricPool PUT operations using the ONTAP CLI:

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

Related information

- [storage aggregate object-store put-rate-limit modify](#)

Customize ONTAP FabricPool object deletion and defragmentation

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

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

You can customize the unreclaimed space threshold percentage and set it to different default levels for different object stores. The default settings are:

Object Store	ONTAP 9.8 and later	ONTAP 9.7 to 9.4	ONTAP 9.3 and earlier	Cloud Volumes ONTAP
Amazon S3	20%	20%	0%	30%
Google Cloud Storage	20%	12%	n/a	35%

Microsoft Azure Blob Storage	25%	15%	n/a	35%
NetApp ONTAP S3	40%	n/a	n/a	n/a
NetApp StorageGRID	40%	40%	0%	n/a

Unreclaimed space threshold

Changing the default unreclaimed space threshold settings will increase or decrease the accepted amount of object fragmentation. Reducing fragmentation will reduce the amount of physical capacity used by the cloud tier at the expense of additional object store resources (reads and writes).

Threshold reduction

To avoid additional expenses, consider reducing the unreclaimed space thresholds when using object store pricing schemes that reduce the cost of storage but increase the cost of reads. Examples include Amazon's Standard-IA and Azure Blob Storage's Cool.

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

Threshold increases

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

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

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



As the system defragments objects and increases their storage efficiency, it might fragment the underlying files by writing referenced blocks to new, more efficient objects. If you significantly increase the unreclaimed space threshold, you can create objects that are more storage efficient but have reduced sequential read performance.

This additional activity results in increased costs from third party S3 providers, such as AWS, Azure, and Google.

NetApp recommends avoiding increasing the unreclaimed space threshold above 60%.

Change the unreclaimed space threshold

You can customize the unreclaimed space threshold percentage for different object stores.

Before you begin

Advanced privilege level is required.

Steps

1. To change the default unreclaimed space threshold, customize and run the following command:

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

Related information

- [storage aggregate object-store modify](#)

Promote ONTAP data to the performance tier

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 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 tier and promote it to the performance tier.

Steps

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 in the cloud tier and promote it to the performance tier.

Steps

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.

Steps

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.

Steps

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
```

Copyright information

Copyright © 2026 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.