



# **Manage volumes for FabricPool**

## **ONTAP 9**

NetApp  
February 12, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap/fabricpool/create-volume-task.html> on February 12, 2026. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

Manage volumes for FabricPool . . . . .	1
Create a volume on a FabricPool-enabled ONTAP local tier . . . . .	1
Move a volume to a FabricPool-enabled ONTAP local tier . . . . .	2
Destination local tier . . . . .	3
Optimized volume moves . . . . .	3
Enable ONTAP volumes in FabricPool to write directly to the cloud . . . . .	4
Enable writing directly to the cloud during volume creation . . . . .	5
Enable writing directly to the cloud on an existing volume . . . . .	5
Disable writing directly to the cloud on a volume . . . . .	6
Enable ONTAP volumes in FabricPool to perform aggressive read-aheads . . . . .	6
Enable aggressive read-ahead mode during volume creation . . . . .	7
Disable aggressive read-ahead mode . . . . .	7
View aggressive read-ahead mode on a volume . . . . .	7

# 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

- 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_FabricPool`, 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
```

## 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.