



Manage volumes for FabricPool

ONTAP 9

NetApp
September 19, 2024

Table of Contents

- Manage volumes for FabricPool 1
 - Create a volume for FabricPool 1
 - Move a volume to FabricPool 2
 - Enable and disable volumes to write directly to the cloud 3
 - Enable and disable aggressive read-ahead mode 5

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
```


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.