

FlexGroup volumes management

ONTAP 9

NetApp October 30, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap/flexgroup/index.html on October 30, 2025. Always check docs.netapp.com for the latest.

Table of Contents

FlexGroup volumes management	1
Learn about ONTAP FlexGroup volumes management with the CLI	1
Learn about ONTAP FlexGroup volumes.	1
Supported and unsupported configurations for ONTAP FlexGroup volumes	2
Features supported beginning with ONTAP 9.16.1	2
Features supported beginning with ONTAP 9.15.1	2
Features supported beginning with ONTAP 9.14.1	2
Features supported beginning with ONTAP 9.13.1	
Features supported beginning with ONTAP 9.12.1	3
Features supported beginning with ONTAP 9.11.1	
Features supported beginning with ONTAP 9.10.1	3
Features supported beginning with ONTAP 9.9.1	3
Features supported beginning with ONTAP 9.8	
Features supported beginning with ONTAP 9.7	4
Features supported beginning with ONTAP 9.6.	
Features supported beginning with ONTAP 9.5.	
Features supported beginning with ONTAP 9.4.	5
Features supported in ONTAP 9.3 and earlier	
Unsupported FlexGroup volume configurations in ONTAP 9	6
FlexGroup volume setup	7
ONTAP FlexGroup volume setup workflow	7
Enable 64-bit NFSv3 identifiers on ONTAP SVMs with FlexGroups	7
Provision an ONTAP FlexGroup volume automatically	
Create ONTAP FlexGroup volumes	. 11
Manage FlexGroup volumes	. 15
Monitor the space usage of ONTAP FlexGroup volumes	. 15
Increase the size of ONTAP FlexGroup volumes	. 17
Reduce the size of ONTAP FlexGroup volumes	
Configure ONTAP FlexGroup volumes to automatically grow and shrink their size	. 20
Delete directories asynchronously from ONTAP FlexGroup volumes	. 21
Manage client rights to delete ONTAP directories asynchronously with FlexGroups	
Create qtrees with ONTAP FlexGroup volumes	
Use quotas for ONTAP FlexGroup volumes	
Enable storage efficiency on ONTAP FlexGroup volumes	
Protect ONTAP FlexGroup volumes using snapshots	. 36
Move constituents from ONTAP FlexGroup volumes.	. 38
Use aggregates in FabricPool for existing ONTAP FlexGroup volumes	
Balance ONTAP FlexGroup volumes by redistributing file data	
Rebalance ONTAP FlexGroup volumes by moving files	. 45
Data protection for FlexGroup volumes	
Data protection for ONTAP FlexGroup volumes workflow summary	. 52
Create SnapMirror relationships for ONTAP FlexGroup volumes	
Create SnapVault relationships for ONTAP FlexGroup volumes	. 57

Create unified data protection relationships for ONTAP FlexGroup volumes	59
Create SVM disaster recovery relationships for ONTAP FlexGroup volumes	62
Transition ONTAP FlexGroup SnapMirror relationships to SVM DR	64
Convert ONTAP FlexVol volumes to FlexGroup volumes within an SVM-DR relationship	66
Considerations for creating SnapMirror cascade and fanout relationships for ONTAP FlexGroup	
volumes	68
Considerations for creating SnapVault backup relationships and unified data protection relationships	
for ONTAP FlexGroup volumes	69
Monitor SnapMirror data transfers for ONTAP FlexGroup volumes	70
Manage data protection operations for FlexGroup volumes	71
Disaster recovery for FlexGroup volumes	71
Expand FlexGroup volumes in a SnapMirror relationship	77
Perform a SnapMirror single file restore from an ONTAP FlexGroup volume	81
Restore ONTAP FlexGroup volumes from SnapVault backups	84
Disable SVM protection on ONTAP FlexGroup volumes	85
Enable SVM protection on ONTAP FlexGroup volumes	86
Convert FlexVol volumes to FlexGroup volumes	88
Learn about converting ONTAP FlexVol volumes to FlexGroup volumes	88
Convert ONTAP FlexVol volumes to ONTAP FlexGroup volumes	89
Convert ONTAP FlexVol volume SnapMirror relationships to ONTAP FlexGroup volume SnapMirror	
relationships	92

FlexGroup volumes management

Learn about ONTAP FlexGroup volumes management with the CLI

You can set up, manage, and protect FlexGroup volumes for scalability and performance. A FlexGroup volume is a scale-out volume that provides high performance along with automatic load distribution.

You can configure FlexGroup volumes if the following are true:

- You want to use best practices, not explore every available option.
- You have cluster administrator privileges, not SVM administrator privileges.



Beginning with ONTAP 9.5, FlexGroup volumes replace Infinite Volumes, which are not supported in ONTAP 9.5 or later releases.

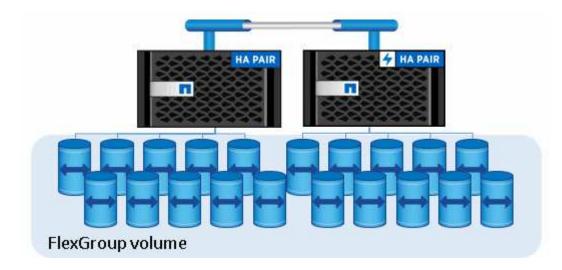
See the supported and unsupported configurations for FlexGroup volumes for more information.

Related information

Conceptual information about FlexVol volumes is applicable to FlexGroup volumes. Information about FlexVol volumes and ONTAP technology is available in the ONTAP Reference Library and in Technical Reports (TRs).

Learn about ONTAP FlexGroup volumes

A FlexGroup volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. A FlexGroup volume contains several member volumes (constituents) that automatically and transparently share the traffic. *Member volumes* are the underlying FlexVol volumes that make up a FlexGroup volume.



FlexGroup volumes provide the following benefits:

High scalability

Multiple FlexGroup volumes can be provisioned on a cluster as long as the number of member volumes does not exceed the node or cluster limits.

Beginning with ONTAP 9.12.1P2, the maximum capacity for a single FlexGroup volume is 60PB, with 400 billion files on a 10-node cluster when large volume support is enabled. Without large volume support, the maximum capacity for a single FlexGroup volume is 20PB.



Although the maximum capacity of a single FlexGroup volume is 60PB (200 member volumes x 300TB = 60PB), best performance is achieved when the used capacity of member volumes remains below 80% (200 member volumes x 240TB = 48PB).

· High performance

FlexGroup volumes can use the resources of the cluster to serve workloads that have high throughput and low latency.

Simplified management

A FlexGroup volume is a single namespace container that can be managed in a similar way as FlexVol volumes.

Supported and unsupported configurations for ONTAP FlexGroup volumes

You should be aware of the ONTAP features that are supported and not supported with FlexGroup volumes in ONTAP 9.

Features supported beginning with ONTAP 9.16.1

· Advanced capacity balancing

Features supported beginning with ONTAP 9.15.1

Automatic provisioning enhancements

Features supported beginning with ONTAP 9.14.1

• Snapshot tagging: Support for creating, modifying and deleting snapshot tags (SnapMirror labels and comments) for snapshots on FlexGroup volumes using the volume snapshot command.

Features supported beginning with ONTAP 9.13.1

- Autonomous Ransomware Protection (ARP) for FlexGroup volumes, including the following supported functionality:
 - FlexGroup expand operations: A new member volume inherits Autonomous Ransomware Protection attributes.
 - FlexVol to FlexGroup conversions: Conversions of FlexVols with active Autonomous Ransomware Protection is possible.
 - · FlexGroup rebalancing: Autonomous Ransomware Protection is supported during disruptive and non-

disruptive rebalancing operations.

- Schedule a single FlexGroup rebalancing operation.
- SnapMirror fanout relationships with SVM DR on FlexGroup volumes. Supports fanout to eight sites.

Features supported beginning with ONTAP 9.12.1

- FlexGroup rebalancing
- · SnapLock for SnapVault
- SnapMirror Cloud
- FabricPool, FlexGroup, and SVM DR working in conjunction. (In releases earlier than ONTAP 9.12.1, any two of these features worked together, but not all three in conjunction.)
- Large volume support increases FlexGroup volume member size from a maximum of 100TB to a maximum of 300TB.

Features supported beginning with ONTAP 9.11.1

SnapLock volumes

SnapLock does not support the following features with FlexGroup volumes:

- · Legal-hold
- Event-based retention
- SnapLock for SnapVault

You configure SnapLock at the FlexGroup level. You cannot configure SnapLock at the member volume level.

· Client asynchronous directory delete

Features supported beginning with ONTAP 9.10.1

- Convert a FlexVol volume to a FlexGroup volume within an SVM DR relationship
- SVM DR FlexClone support for FlexGroup volumes

Features supported beginning with ONTAP 9.9.1

SVM disaster recovery

Cloning a FlexGroup volume that is part of an SVM DR relationship is not supported.

• SnapMirror fanout relationships of 2 or more (A to B, A to C), with a maximum of 8 fanout legs.

Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes

• SnapMirror cascading relationships up to two levels (A to B to C)

Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes

Features supported beginning with ONTAP 9.8

- Restoring a single file from a FlexGroup SnapMirror vault or from a UDP destination
 - Restore can be from a FlexGroup volume of any geometry to FlexGroup volume of any geometry
 - Only one file per restore operation is supported
- Converting volumes transitioned from 7-Mode systems to FlexGroup volumes

For more information, see the NetApp Knowledge Base: How To Convert a Transitioned FlexVol to FlexGroup.

- NFSv4.2
- Asynchronous delete of files and directories
- Files System Analytics (FSA)
- FlexGroup as a VMware vSphere datastore
- Additional support for tape backup and restore using NDMP, including the following features:
 - NDMP restartable backup extension (RBE) and Snapshot Management Extension (SSME)
 - Environment variables EXCLUDE and MULTI_SUBTREE_NAMES support FlexGroup backups
 - Introduction of IGNORE CTIME MTIME environment variable for FlexGroup backups
 - Individual file recovery in a FlexGroup using the NDMP_SNAP_RECOVER message, which is part of extension 0x2050 Dump and restore sessions are aborted during an upgrade or revert.

Features supported beginning with ONTAP 9.7

- FlexClone volume
- NFSv4 and NFSv4.1
- pNFS
- Tape backup and restore by using NDMP

You must be aware of the following points for NDMP support on FlexGroup volumes:

• The NDMP_SNAP_RECOVER message in the extension class 0x2050 can be used only for recovering an entire FlexGroup volume.

Individual files in a FlexGroup volume cannot be recovered.

- NDMP restartable backup extension (RBE) is not supported for FlexGroup volumes.
- Environment variables EXCLUDE and MULTI_SUBTREE_NAMES are not supported for FlexGroup volumes.
- The ndmpcopy command is supported for data transfer between FlexVol and FlexGroup volumes.

If you revert from Data ONTAP 9.7 to an earlier version, the incremental transfer information of the previous transfers is not retained and therefore, you must perform a baseline copy after reverting.

- VMware vStorage APIs for Array Integration (VAAI)
- Conversion of a FlexVol volume to a FlexGroup volume
- FlexGroup volumes as FlexCache origin volumes

Features supported beginning with ONTAP 9.6

- · Continuously available SMB shares
- MetroCluster configurations
- Renaming a FlexGroup volume (volume rename command)
- Shrinking or reducing the size of a FlexGroup volume (volume size command)
- · Elastic sizing
- NetApp aggregate encryption (NAE)
- Cloud Volumes ONTAP

Features supported beginning with ONTAP 9.5

- · ODX copy offload
- · Storage-Level Access Guard
- Enhancements to change notifications for SMB shares

Change notifications are sent for changes to the parent directory on which the changenotify property is set and for changes to all of the subdirectories in that parent directory.

- FabricPool
- · Quota enforcement
- · Qtree statistics
- Adaptive QoS for files in FlexGroup volumes
- FlexCache (cache only; FlexGroup as origin supported in ONTAP 9.7)

Features supported beginning with ONTAP 9.4

- FPolicy
- · File auditing
- Throughput floor (QoS Min) and adaptive QoS for FlexGroup volumes
- Throughput ceiling (QoS Max) and throughput floor (QoS Min) for files in FlexGroup volumes

You use the volume file modify command to manage the QoS policy group that is associated with a file.

- Relaxed SnapMirror limits
- SMB 3.x multichannel

Features supported in ONTAP 9.3 and earlier

- · Antivirus configuration
- · Change notifications for SMB shares

Notifications are sent only for changes to the parent directory on which the changenotify property is set. Change notifications are not sent for changes to subdirectories in the parent directory.

- · Qtrees
- Throughput ceiling (QoS Max)
- Expand the source FlexGroup volume and destination FlexGroup volume in a SnapMirror relationship
- SnapVault backup and restore
- Unified data protection relationships
- Autogrow option and autoshrink option
- · Inode count factored to ingest
- · Volume encryption
- Aggregate inline deduplication (cross-volume deduplication)
- NetApp volume encryption (NVE)
- SnapMirror technology
- Snapshots
- · Digital Advisor
- · Inline adaptive compression
- Inline deduplication
- · Inline data compaction
- AFF
- · Quota reporting
- NetApp Snapshot technology
- SnapRestore software (FlexGroup level)
- Hybrid aggregates
- · Constituent or member volume move
- Postprocess deduplication
- NetApp RAID-TEC technology
- · Per-aggregate consistency point
- Sharing FlexGroup with FlexVol volume in the same SVM

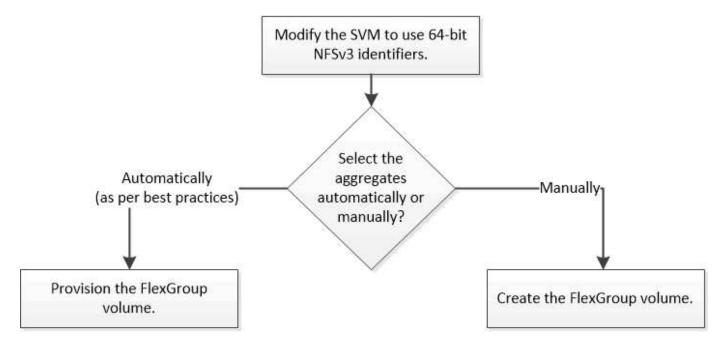
Unsupported FlexGroup volume configurations in ONTAP 9

Unsupported protocols	Unsupported data protection features	Other unsupported ONTAP features
 pNFS (ONTAP 9.6 and earlier) SMB 1.0 SMB transparent failover (ONTAP 9.5 and earlier) SAN 	 SnapLock volumes (ONTAP 9.10.1 and earlier) SMTape SnapMirror synchronous SVM DR with FlexGroup volumes containing FabricPools (ONTAP 9.11.1 and earlier) 	 Remote Volume Shadow Copy Service (VSS) SVM data mobility

FlexGroup volume setup

ONTAP FlexGroup volume setup workflow

You can either provision a FlexGroup volume where ONTAP automatically selects the aggregates based on the best practices for optimum performance, or create a FlexGroup volume by manually selecting the aggregates and configuring it for data access.



Before you begin

You must have created the SVM with NFS and SMB added to the list of allowed protocols for the SVM.

About this task

You can automatically provision a FlexGroup volume only on clusters with four nodes or less. On clusters with more than four nodes, you must create a FlexGroup volume manually.

Enable 64-bit NFSv3 identifiers on ONTAP SVMs with FlexGroups

To support the high file count of FlexGroup volumes and to avoid file ID collisions, you should enable 64-bit file identifiers on the SVM on which the FlexGroup volume must be created.

Steps

- 1. Log in to the advanced privilege level: set -privilege advanced
- 2. Modify the SVM to use 64-bit NFSv3 FSIDs and file IDs: vserver nfs modify -vserver svm_name -v3-64bit-identifiers enabled

After you finish

All of the clients must be remounted. This is required because the file system IDs change, and the clients might receive stale file handle messages when attempting NFS operations.

Provision an ONTAP FlexGroup volume automatically

When you create a FlexGroup volume, you can choose to have ONTAP automatically provision the FlexGroup volume by selecting the underlying local tiers (aggregates). Local tiers are selected based on the best practices for optimum performance and capacity.

Before you begin

Each node in the cluster must have at least one local tier.



When creating a FlexGroup volume that will tier inactive data, each node must have at least one local tier with FabricPool enabled.

About this task

ONTAP selects two local tiers with the largest amount of usable space on each node to create the FlexGroup volume. If two local tiers are not available, ONTAP selects one local tier per node to create the FlexGroup volume.

Beginning with ONTAP 9.15.1, when you automatically provision a FlexGroup volume, ONTAP uses balanced placement (BP) to choose the local tiers and FlexGroup member (constituent) volumes layout. One aspect of BP is how it limits over-provisioning local tiers when creating 'none' guaranteed (thin-provisioned) FlexGroup volumes. The size of the overall FlexGroup volume is limited by the amount of free space on the local tier, although the limit is higher than it is for 'volume' guaranteed (thick-provisioned) FlexGroup volumes. When you create a FlexGroup volume using REST APIs or auto-provision-as with the ONTAP CLI, provisioning might fail because of insufficient space due to this limit. You can avoid this by creating smaller FlexGroup volumes, or by creating a FlexGroup volume and selecting the local tiers manually using the aggr-list

parameter.

Steps

1. Provision the FlexGroup volume:

```
volume create -vserver svm_name -volume fg_vol_name -auto-provision-as
flexgroup -size fg_size [-encrypt true] [-qos-policy-group
qos_policy_group_name] [-support-tiering true] [-granular-data advanced]
```

Beginning with ONTAP 9.16.1, you can enable advanced capacity balancing (-granular-data advanced in the CLI) to write data across multiple FlexGroup member volumes when files are larger than 10GB.

Beginning with ONTAP 9.5, you can create FlexGroup volumes on local tiers with FabricPool enabled. To automatically provision a FlexGroup volume on local tiers with FabricPool enabled, you must set the -support-tiering parameter to true. The volume guarantee must be always set to none for FabricPool. You can also specify the tiering policy and tiering minimum cooling period for the FlexGroup volume.

Disk and aggregate management

You can specify a throughput ceiling (QoS Max) for FlexGroup volumes. This limits the performance resources that the FlexGroup volume can consume. Beginning with ONTAP 9.4, you can specify throughput floors (QoS Min) and adaptive QoS for FlexGroup volumes.

Performance management

You can set the <code>-encrypt</code> parameter to <code>true</code> if you want to enable encryption on the FlexGroup volume. For creating an encrypted volume, you must have installed the volume encryption license and the key manager.



You must enable encryption on FlexGroup volumes at the time of creation. You cannot enable encryption on existing FlexGroup volumes.

Encryption of data at rest

The size parameter specifies the size of the FlexGroup volume in KB, MB, GB, TB, or PB.

The following example shows how to provision a FlexGroup volume of size 400 TB:

```
cluster-1::> volume create -vserver vs0 -volume fg -auto-provision-as flexgroup -size 400TB Warning: The FlexGroup "fg" will be created with the following number of constituents of size 25TB: 16. The constituents will be created on the following aggregates: aggr1,aggr2 Do you want to continue? \{y|n\}: y [Job 34] Job succeeded: Successful
```

The following example shows how to create a QoS policy group for throughput ceiling and how to apply it to a FlexGroup volume:

cluster1::> qos policy-group create -policy group pg-vs1 -vserver vs1 -max-throughput 5000iops

```
cluster-1::> volume create -vserver vs0 -volume fg -auto-provision-as flexgroup -size 400TB -qos-policy-group pg-vs1
Warning: The FlexGroup "fg" will be created with the following number of constituents of size 25TB: 16.
The constituents will be created on the following aggregates:
aggr1,aggr2
Do you want to continue? {y|n}: y
[Job 34] Job succeeded: Successful
```

The following example shows how to provision a FlexGroup volume of size 400 TB on local tiers with FabricPool enabled:

```
cluster-1::> volume create -vserver vs0 -volume fg -auto-provision-as flexgroup -size 400TB -support-tiering true -tiering-policy auto Warning: The FlexGroup "fg" will be created with the following number of constituents of size 25TB: 16. The constituents will be created on the following aggregates: aggr1, aggr2 Do you want to continue? \{y|n\}: y [Job 34] Job succeeded: Successful
```

The FlexGroup volume is created with eight member volumes on each node in the cluster. The member volumes are distributed equally between the two largest local tiers on each node.

By default, the FlexGroup volume is created with the volume space guarantee setting except on AFF systems. For AFF systems, by default the FlexGroup volume is created with the none space guarantee.

2. Mount the FlexGroup volume with a junction path:

```
volume mount -vserver vserver_name -volume vol_name -junction-path
junction_path
```

```
cluster1::> volume mount -vserver vs0 -volume fg2 -junction-path /fg2
```

After you finish

You should mount the FlexGroup volume from the client.

If you are running ONTAP 9.6 or earlier and if the storage virtual machine (SVM) has both NFSv3 and NFSv4 configured, mounting the FlexGroup volume from the client might fail. In such cases, you must explicitly specify the NFS version when mounting the FlexGroup volume from the client.

```
# mount -t nfs -o vers=3 192.53.19.64:/fg2 /mnt/fg2
# ls /mnt/fg2
file1 file2
```

Related information

· qos policy-group create

Create ONTAP FlexGroup volumes

You can create a FlexGroup volume by manually selecting the local tiers (aggregates) on which the FlexGroup volume must be created, and then specifying the number of member volumes (constituents) on each local tier.

Alternatively, you can choose to have ONTAP automatically provision the FlexGroup volume by selecting the local tiers and letting ONTAP set the number of member volumes based on the best practices for optimum performance and capacity.

About this task

You must be aware of the space required in the local tiers for creating a FlexGroup volume.

You must consider the following guidelines when creating a FlexGroup volume for obtaining the best performance results with a FlexGroup volume:

A FlexGroup volume should use local tiers that are on identical hardware systems.

The use of identical hardware systems helps in providing predictable performance across the FlexGroup volume. Note: C-Series r1 and C-Series r2 systems are not identical systems. For example, the AFF C80 r1 and AFF C80 r2 are not identical.

A FlexGroup volume should span local tiers using the same disk type and RAID group configurations.

For consistent performance, you must ensure that all of the local tiers are made of all SSDs, all HDDs, or all Flash Pool (hybrid) local tiers. Additionally, the local tiers should have the same number of drives and RAID groups across the FlexGroup volume.

A FlexGroup volume can span parts of a cluster.

A FlexGroup volume does not have to be configured to span the entire cluster, but doing so can take greater advantage of the hardware resources that are available.

- When creating a FlexGroup volume, it is best if the local tiers on which the FlexGroup volume is deployed have the following characteristics:
 - Approximately the same amount of free space should be available across multiple local tier, especially when using thin provisioning.
 - Approximately 3 percent of the free space should be reserved for local tier metadata after creation of the FlexGroup volume.
- For FAS systems, it is best to have two local tiers per node and for AFF systems, you must have one local tier per node for the FlexGroup volume.
- For each FlexGroup volume, you should create at least eight member volumes that are distributed over two or more local tiers on FAS systems, and over one or more local tiers on AFF systems.

- Beginning with ONTAP 9.9.1, SnapMirror fanout relationships of two or more FlexGroup volumes are supported, with a maximum of eight fanout legs. System Manager does not support SnapMirror cascading FlexGroup volume relationships.
- When you use System Manager to create a FlexGroup volume, ONTAP automatically selects the local tiers required for creating the FlexGroup volume.
- Beginning with ONTAP 9.8, when you provision storage, QoS is enabled by default. You can disable QoS, or choose a custom QoS policy during the provisioning process or at a later time.

Before you begin

• Beginning with ONTAP 9.13.1, you can create volumes with capacity analytics and Activity Tracking enabled. To enable capacity or Activity Tracking, issue the volume create command with -analytics -state or -activity-tracking-state set to on.

To learn more about capacity analytics and Activity Tracking, see Enable File System Analytics. Learn more about volume create in the ONTAP command reference.

System Manager

Using System Manager, you can create a FlexGroup volume.

Steps

- Navigate to Storage > Volumes and select + Add.
- 2. In the Add volume window, enter a volume name and size, then select More options.
- 3. In the Storage and optimization section, select Distribute volume data across the cluster (FlexGroup).



If you are running ONTAP 9.8 or later and you want to disable QoS or choose a custom QoS policy, click **More Options**, and then under **Storage and Optimization**, select **Performance Service Level**.

4. Complete the remaining information for the volume and select **Save**.

CLI

1. Create the FlexGroup volume:

```
volume create -vserver <svm_name> -volume <flexgroup_name> -aggr
-list aggr1,aggr2,.. -aggr-list-multiplier <constituents_per_aggr>
-size <fg_size> [-encrypt true] [-qos-policy-group
qos_policy_group_name] [-granular-data advanced]
```

 The -aggr-list parameter specifies the list of local tiers to be used for FlexGroup member volumes.

For consistent performance across the FlexGroup volume, all of the local tiers must use the same disk type and RAID group configurations.

• The -aggr-list-multiplier parameter specifies the number of member volumes that will be created on each local tier listed with the -aggr-list parameter.

The default value of the -aggr-list-multiplier parameter is 4.

- The size parameter specifies the size of the FlexGroup volume in KB, MB, GB, TB, or PB.
- Beginning with ONTAP 9.16.1, you can enable advanced capacity balancing (-granular-data advanced in the CLI) to write data across multiple FlexGroup member volumes when files are larger than 10GB.
- Beginning with ONTAP 9.5, you can create FlexGroup volumes using local tiers with FabricPool enabled.

To create a FlexGroup volume for FabricPool, all the local tiers specified with the <code>-aggr-list</code> parameter must have FabricPool enabled. The volume guarantee must be always set to <code>none</code> when using FabricPool. You can also specify the tiering policy and tiering minimum cooling period for the FlexGroup volume.

Disk and aggregate management

 Beginning with ONTAP 9.4, you can specify throughput floors (QoS Min) and adaptive QoS for FlexGroup volumes.

Performance management

- You can specify a throughput ceiling (QoS Max) for FlexGroup volumes, which limits the performance resources that the FlexGroup volume can consume.
- You can set the -encrypt parameter to true if you want to enable encryption on the FlexGroup volume.

For creating an encrypted volume, you must have installed the volume encryption license and the key manager.



You must enable encryption on FlexGroup volumes at the time of creation. You cannot enable encryption on existing FlexGroup volumes.

Encryption of data at rest

```
cluster-1::> volume create -vserver vs0 -volume fg2 -aggr-list aggr1,aggr2,aggr3,aggr1 -aggr-list-multiplier 2 -size 500TB

Warning: A FlexGroup "fg2" will be created with the following number of constituents of size 62.50TB: 8.

Do you want to continue? {y|n}: y

[Job 43] Job succeeded: Successful
```

In the previous example, if you want to create the FlexGroup volume for FabricPool, all local tiers (aggr1, aggr2, and aggr3) must have FabricPool enabled. Mount the FlexGroup volume with a junction path: volume mount -vserver vserver_name -volume vol_name -junction-path junction_path

```
\verb|cluster1::> volume mount - vserver vs0 - volume fg2 - junction-path / fg|
```

After you finish

You should mount the FlexGroup volume from the client.

If you are running ONTAP 9.6 or earlier and if the storage virtual machine (SVM) has both NFSv3 and NFSv4 configured, mounting the FlexGroup volume from the client might fail. In such cases, you must explicitly specify the NFS version when you are mounting the FlexGroup volume from the client.

```
# mount -t nfs -o vers=3 192.53.19.64:/fg /mnt/fg2
# ls /mnt/fg2
file1 file2
```

Manage FlexGroup volumes

Monitor the space usage of ONTAP FlexGroup volumes

You can view a FlexGroup volume and its constituents, and monitor the space used by the FlexGroup volume.

About this task

Beginning with ONTAP 9.6, elastic sizing is supported. ONTAP automatically grows a constituent of a FlexGroup volume if it is running out of space by shrinking any other constituent in the FlexGroup volume that has free space by an equivalent amount. Elastic sizing avoids any out-of-space errors that are generated because of one or more FlexGroup constituent volumes running out of space.



Beginning with ONTAP 9.9.1, logical space reporting and enforcement is also available for FlexGroup volumes. For more information, see Logical space reporting and enforcement for volumes.

Step

1. View the space used by the FlexGroup volume and its constituents: volume show -vserver vserver name -volume-style-extended [flexgroup | flexgroup-constituent]

Available	e Used%	Aggregate			Size
 c 1	fg1 0001	aggr3	online	ВМ	31 25GR
12.97GB	- 	aggij	OHITHE	1777	31 . 230b
		aggr1	online	RW	31.25GB
12.98GB	- 	99			
vs1		aggr1	online	RW	31.25GB
13.00GB		- 55			
vs1		aggr3	online	RW	31.25GB
12.88GB	- 	3 3			
vs1		aggr1	online	RW	31.25GB
13.00GB	- 	3.3			
vs1	fg1 0006	aggr3	online	RW	31.25GB
12.97GB	<u>—</u> 56%				
vs1	fg10007	aggr1	online	RW	31.25GB
13.01GB	56%				
vs1	fg10008	aggr1	online	RW	31.25GB
13.01GB	56%				
vs1	fg10009	aggr3	online	RW	31.25GB
12.88GB	56%				
vs1	fg10010	aggr1	online	RW	31.25GB
13.01GB	56%				
vs1	fg10011	aggr3	online	RW	31.25GB
12.97GB	56%				
vs1	fg10012	aggr1	online	RW	31.25GB
13.01GB					
vs1		aggr3	online	RW	31.25GB
12.95GB					
vs1	<u> </u>	aggr3	online	RW	31.25GB
12.97GB					
vs1	<u> </u>	aggr3	online	RW	31.25GB
12.88GB					04 05
vs1 13.01GB	- 	aggr1	online	RW	31.25GB

You can use the available space and percentage space used to monitor the space usage of the FlexGroup volume.

Increase the size of ONTAP FlexGroup volumes

You can increase the size of a FlexGroup volume either by adding more capacity to all existing member volumes (constituents) of the FlexGroup volume or by expanding the FlexGroup volume with new member volumes. A FlexGroup volume cannot have more than 200 member volumes.

You can also increase the size of an individual volume within a FlexGroup volume if needed.

Before you begin

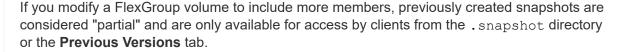
Sufficient space must be available in the aggregates.

About this task

If you want to add more space, you can increase the collective size of the FlexGroup volume. Increasing the size of a FlexGroup volume resizes the existing member volumes of the FlexGroup volume.

If you want to improve performance, you can expand the FlexGroup volume. You might want to expand a FlexGroup volume and add new member volumes in the following situations:

- New nodes have been added to the cluster.
- New local tiers (aggregates) have been created on the existing nodes.
- The existing member volumes of the FlexGroup volume have reached the maximum FlexVol size for the hardware (100TB or 300TB if large volume support has been enabled), and therefore the FlexGroup volume cannot be resized without adding additional member volumes.





If a snapshot is considered "partial", it cannot be used in SnapRestore operations. However, partial snapshots can be used to restore individual files from .snapshot directories or the **Previous Versions** tab.

In releases earlier than ONTAP 9.3, do not expand FlexGroup volumes after a SnapMirror relationship is established. If you expand the source FlexGroup volume after breaking the SnapMirror relationship in releases earlier than ONTAP 9.3, you must perform a baseline transfer to the destination FlexGroup volume once again. Beginning with ONTAP 9.3, you can expand FlexGroup volumes that are in a SnapMirror relationship.

Steps

1. Increase the size of the FlexGroup volume by increasing the capacity or performance of the FlexGroup volume, as required:

If you want to increase the	Then do this		
Capacity of the FlexGroup volume	Resize all the member volumes of the FlexGroup volume:		
	<pre>volume modify -vserver <svm_name> -volume <fg_name> -size <new_size></new_size></fg_name></svm_name></pre>		

Performance to the FlexGroup volume

Expand the FlexGroup volume by adding new member volumes (constituents):

volume expand -vserver vserver_name
-volume fg_name -aggr-list aggregate
name,... [-aggr-list-multiplier
constituents per aggr]

The default value of the -aggr-list -multiplier parameter is 1.

When expanding a FlexGroup volume using FabricPool, all local tiers (aggregates) must be attached to the same cloud tier.

Assuming existing aggregates (local tiers) or member volumes have not reached their maximum capacities (100/300TB or two billion files each), it is preferable to increase the overall size of the FlexGroup volume rather than adding additional member volumes.

Use volume expand only if increasing the existing volume size or file count is not an option or if the FlexGroup is being expanded across new hardware. The same number of member volumes should be added to all nodes in order to ensure consistent performance. For example, if an existing FlexGroup volume has 8 member volumes with four member volumes per node, adding two members per node will result in 12 member volumes, six member volumes per node.

When adding new members to new nodes, try to maintain a consistent number of member volumes per node as in the existing nodes. For example, if an existing FlexGroup volume has 8 member volumes with four member volumes per node, if the FlexGroup volumes is expanded to the new node, four member volumes should be added, resulting in a 12 member FlexGroup volume.

Adding new members to a FlexGroup volume changes the ingest heuristics to favor the new, empty, member volumes and can affect overall system performance for new data ingest until the new member volumes become balanced with pre-existing member volumes.

Examples

Example of increasing the capacity of the existing member volumes

The following example shows how to add 20 TB space to a FlexGroup volume volX:

```
cluster1::> volume modify -vserver svm1 -volume volX -size +20TB
```

If the FlexGroup volume has 16 member volumes, the space of each member volumes is increased by 1.25 TB.

Example of improving performance by adding new member volumes

The following example shows how to add four additional member volumes, two to each of the underlying local tiers (aggregates) to FlexGroup volume fg1:

```
cluster1::> volume expand -vserver svm1 -volume fg1 -aggr-list aggr1,aggr2
-aggr-list-multiplier 2
```

The size of the new member volumes is the same as that of the existing member volumes.

Increase the size of an individual volume

If you want to increase the size of an individual member volume in a FlexGroup volume, you can use the volume resize command.

Step

1. Increase the size of a single FlexGroup member volume:

```
volume size -volume <volume_name> -vserver <svm1> -new-size <new_size>
```

The following example increases the size of FlexGroup member volume FG_0003 to 3.7GB:

```
volume size -volume FG__0003 -vserver svm1 -new-size 3.7GB
vol size: Volume "svm1:FG__0003" size set to 3.70g.
```

Reduce the size of ONTAP FlexGroup volumes

Beginning with ONTAP 9.6, you can resize a FlexGroup volume to a value lower than its current size to free up the unused space from the volume. When you reduce the size of a FlexGroup volume, ONTAP automatically resizes all of the FlexGroup constituents.

Step

- 1. Check the current FlexGroup volume size: 'volume size -vserver vserver_name -volume fg_name'
- 2. Reduce the size of the FlexGroup volume: volume size -vserver vserver_name -volume fg name new size

When you specify the new size, you can specify either a lower value than the current size or a negative value using the minus sign (-) by which the current size of the FlexGroup volume is reduced.



If automatic shrinking is enabled for the volume (volume autosize command), the minimum autosize is set to the new size of the volume.

The following example displays the current volume size for the FlexGroup volume named volX and resizes the volume to 10TB:

```
cluster1::> volume size -vserver svm1 -volume volX
(volume size)
vol size: FlexGroup volume 'svm1:volX' has size 15TB.

cluster1::> volume size -vserver svm1 -volume volX 10TB
(volume size)
vol size: FlexGroup volume 'svm1:volX' size set to 10TB.
```

The following example displays the current volume size for the FlexGroup volume named volX and reduces the size of the volume by 5TB:

```
cluster1::> volume size -vserver svm1 -volume volX
  (volume size)
vol size: FlexGroup volume 'svm1:volX' has size 15TB.

cluster1::> volume size -vserver svm1 -volume volX -5TB
  (volume size)
vol size: FlexGroup volume 'svm1:volX' size set to 10TB.
```

Configure ONTAP FlexGroup volumes to automatically grow and shrink their size

Beginning with ONTAP 9.3, you can configure FlexGroup volumes to automatically grow and shrink according to how much space they currently require.

Before you begin

The FlexGroup volume must be online.

About this task

You can autosize FlexGroup volumes in two modes:

Increase the size of the volume automatically (grow mode)

Automatic growing helps prevent a FlexGroup volume from running out of space, if the aggregate can supply more space. You can configure the maximum size for the volume. The increase is automatically triggered based on the amount of data being written to the volume in relation to the current amount of used space and any thresholds set.

By default, the maximum size a volume can grow to is 120% of the size at which autogrow is enabled. If you need to ensure that the volume can grow to be larger than that, you must set the maximum size for the volume accordingly.

• Shrink the size of the volume automatically (grow shrink mode)

Automatic shrinking prevents a volume from being larger than needed, freeing space in the aggregate for use by other volumes.

Autoshrink can only be used in combination with autogrow to meet changing space demands and is not

available alone. When autoshrink is enabled, ONTAP automatically manages the shrinking behavior of a volume to prevent an endless loop of autogrow and autoshrink actions.

As a volume grows, the maximum number of files it can contain might be automatically increased. When a volume is shrunk, the maximum number of files it can contain is left unchanged, and a volume cannot be automatically shrunk below the size that corresponds to its current maximum number of files. For this reason, it might not be possible to automatically shrink a volume all the way to its original size.

Steps

1. Configure the volume to grow and shrink its size automatically: volume autosize -vserver vserver_name -volume vol_name -mode [grow | grow_shrink]

You can also specify the maximum size, minimum size, and thresholds for growing or shrinking the volume.

The following command enables automatic size changes for a volume called fg1. The volume is configured to grow to a maximum size of 5 TB when it is 70% full.

```
cluster1::> volume autosize -volume fg1 -mode grow -maximum-size 5TB
-grow-threshold-percent 70
vol autosize: volume "vs_src:fg1" autosize settings UPDATED.
```

Delete directories asynchronously from ONTAP FlexGroup volumes

Beginning with ONTAP 9.8, you can delete directories from Linux and Windows client shares asynchronously (that is, in the background). Cluster and SVM administrators can perform asynchronous delete operations on both FlexVol and FlexGroup volumes.

About this task

You must be a cluster administrator or an SVM administrator using the advanced privilege mode.

Beginning with ONTAP 9.8, you can use asynchronous delete functionality using the ONTAP CLI. Beginning with ONTAP 9.9.1, you can use this functionality with System Manager. For more information about this process, see Take corrective action based on ONTAP analytics in FSA.

Beginning with ONTAP 9.11.1, a storage administrator can grant rights on a volume to allow NFS and SMB clients to perform asynchronous delete operations. For more information, see Manage client rights to delete directories asynchronously.

You can use the volume file async-delete show command to check the status of in-progress asynchronous delete jobs, and, beginning with ONTAP 9.17.1, the status of asynchronous delete jobs issued from clients is also displayed.

Delete directories asynchronously

You can use System Manager or the ONTAP CLI to delete directories asynchronously.

System Manager

Beginning with ONTAP 9.10.1

- 1. Select **Storage > Volumes** and select the desired volume name.
- 2. In the individual volume page, select the **File system** tab, and then select the **Explorer** tab.
- 3. In the **Explorer** view, select the desired directory.
- 4. To delete, hover over a file or folder, and the delete poption appears.

You can only delete one object at a time.



When directories and files are deleted, the new storage capacity values are not displayed immediately.

In ONTAP 9.9.1

- 1. Select **Storage > Volumes**.
- Select the desired volume, then select Explorer.
- 3. In the **Explorer** view, select the desired directory.
- 4. To delete, hover over a file or folder, and the delete option appears.



Use the CLI to perform an asynchronous delete

1. Enter advanced privilege mode:

set -privilege advanced

2. Delete directories on a FlexVol or FlexGroup volume:

volume file async-delete start -vserver <SVM_name> -volume <volume_name>
-path <file path> -throttle <throttle>

The minimum throttle value is 10, the maximum is 100,000, and the default is 5000. Lower throttle values use less resources, which can result in a slower deletion rate, while higher throttle values use more resources, but can result in a faster deletion rate.

The following example deletes the directory named d2, which is located in the directory named d1.

cluster::*> volume file async-delete start -vserver vs1 -volume vol1 -path $\mathrm{d}1/\mathrm{d}2$

3. (Optional) Check the status of the in-progress async delete jobs:

volume file async-delete show

4. Verify that the directory was deleted:

event log show

The following example shows output for the event log when the directory is successfully deleted.

```
Time Node Severity Event

7/7/2025 09:04:04 cluster-vsim NOTICE
asyncDelete.message.success: Async delete job on path d1/d2 of
volume (MSID: 2162149232) was completed. Number of files deleted: 7,
Number of directories deleted: 5. Total number of bytes deleted:
135168.
```

Learn more about event log show in the ONTAP command reference.

Cancel a directory delete job

1. Enter advanced privilege mode:

```
set -privilege advanced
```

2. Verify that the directory delete is in progress:

```
volume file async-delete show
```

If the SVM, volume, JobID, and path of your directory is displayed, you can cancel the job.

3. Cancel the directory delete:

```
volume file async-delete cancel -vserver <SVM_name> -volume <volume_name>
-jobid <job_id>
```

Manage client rights to delete ONTAP directories asynchronously with FlexGroups

Beginning with ONTAP 9.11.1, storage administrators can grant rights on a volume to allow NFS and SMB clients to perform asynchronous delete operations. When asynchronous delete is enabled on the cluster, Linux client users can use the mv command and Windows client users can use the rename command to delete a directory on the specified volume by moving it to a hidden directory that by default is named .ontaptrashbin.

Rights are granted per volume. NFS client users should have root access on the NFS client and superuser access to the NFS export.

You can move directories only. You cannot move files to the .ontaptrashbin directory.

Learn about using ONTAP to delete directories asynchronously from FlexGroup volumes.

Enable client asynchronous directory delete

Steps

- 1. From the cluster CLI, enter advanced privilege mode: -privilege advance
- 2. Enable client asynchronous delete at the mountpoint of a volume and, if desired, provide an alternate name for the trashbin directory:

volume file async-delete client enable volume *volname* vserver *vserverName* trashbinname *name*

Example using the default trashbin name:

Example specifying an alternate trashbin name:

3. Verify client asynchronous delete is enabled:

```
volume file async-delete client show
```

Example:

Disable client asynchronous directory delete

Steps

1. From the cluster CLI, disable client asychronous directory delete:

volume file async-delete client disable volume *volname* vserver *vserverName*

Example:

```
cluster1::*> volume file async-delete client disable -volume vol1
-vserver vs1

Success: Asynchronous directory delete client disabled successfully on volume.
```

2. Verify client asynchronous delete is disabled:

```
volume file async-delete client show
```

Example:

Create qtrees with ONTAP FlexGroup volumes

Beginning with ONTAP 9.3, you can create qtrees with FlexGroup volumes. Qtrees enable you to partition your FlexGroup volumes into smaller segments that you can manage individually.

About this task

- If the source FlexGroup volume has qtrees in a SnapMirror relationship, the destination cluster must be running ONTAP 9.3 or later (a version of ONTAP software that supports qtrees).
- Beginning with ONTAP 9.5, qtree statistics are supported for FlexGroup volumes.

Steps

1. Create a qtree in the FlexGroup volume:

```
volume qtree create -vserver <vserver_name> -volume <volume_name> -qtree
<qtree_name>
```

You can optionally specify the security style, SMB oplocks, UNIX permissions, and export policy for the qtree.

```
\label{cluster1::} {\tt cluster1::} {\tt volume} \  \, {\tt qtree} \  \, {\tt create} \  \, {\tt -vserver} \  \, {\tt vs0} \  \, {\tt -volume} \  \, {\tt fg1} \  \, {\tt -qtree} \  \, {\tt qtree1} \\ {\tt -security-style} \  \, {\tt mixed}
```

Related information

Logical storage management

Use quotas for ONTAP FlexGroup volumes

In ONTAP 9.4 and earlier, you can apply quotas rules to FlexGroup volumes only for reporting purposes, but not for enforcing quota limits. Beginning with ONTAP 9.5, you can enforce limits on quota rules that are applied to FlexGroup volumes.

About this task

• Beginning with ONTAP 9.5, you can specify hard, soft, and threshold limit quotas for FlexGroup volumes.

You can specify these limits to constrain the amount of space, the number of files that a specific user, group, or qtree can create, or both. Quota limits generate warning messages in the following scenarios:

 When usage exceeds a configured soft limit, ONTAP issues a warning message, but further traffic is still allowed.

If usage later drops below the configured soft limit again, an all-clear message is issued.

When usage exceeds a configured threshold limit, ONTAP issues a second warning message.

No all-clear administrative message is issued when usage later drops below a configured threshold limit.

- If usage reaches a configured hard limit, ONTAP prevents further resource consumption by rejecting traffic.
- In ONTAP 9.5, quota rules cannot be created or activated on the destination FlexGroup volume of a SnapMirror relationship.
- During quota initialization, quotas are not enforced, and there are no notifications of breached quotas following quota initialization.

To check if quotas were breached during quota initialization, you can use the volume quota report command.

Quota targets and types

Quotas have a type: they can be either user, group, or tree. Quota targets specify the user, group, or qtree for which the quota limits are applied.

The following table lists the kinds of quota targets, what types of quotas each quota target is associated with, and how each quota target is represented:

Quota target	Quota type	How target is represented	Notes	
user	user quota	UNIX user name UNIX UID Windows user name in pre-Windows 2000 format Windows SID	User quotas can be applied for a specific volume or qtree.	
group	group quota	UNIX group name UNIX GID	Group quotas can be applied for a specific volume or qtree. ONTAP does not apply group quotas based on Windows IDs.	
qtree	tree quota	qtree name	Tree quotas are applied to a particular volume and do not affect qtrees in other volumes.	
" "	user quotagroup quota tree quota	Double quotation marks ("")	A quota target of "" denotes a <i>default quota</i> . For default quotas, the quota type is determined by the value of the type field.	

Behavior of FlexGroup volumes when quota limits are exceeded

Beginning with ONTAP 9.5, quota limits are supported on FlexGroup volumes. There are some differences in the way quota limits are enforced on a FlexGroup volume when compared to a FlexVol volume.

FlexGroup volumes might show the following behaviors when the quota limits are exceeded:

• The space and file usage in a FlexGroup volume might reach up to 5 percent higher than the configured hard limit before the quota limit is enforced by rejecting further traffic.

To provide the best performance, ONTAP might allow the space consumption to exceed the configured hard limit by a small margin before the quota enforcement begins. This additional space consumption does not exceed 5 percent of the configured hard limits, 1 GB, or 65536 files, whichever is lower.

- After the quota limit is reached, if a user or administrator deletes some files or directories such that the quota usage is now below the limit, the subsequent quota-consuming file operation might resume with a delay (might take up to 5 seconds to resume).
- When the total space and file usage of a FlexGroup volume exceed the configured quota limits, there might be a slight delay in logging an event log message.
- You might get "no space" errors if some constituents of the FlexGroup volume get full, but the quota limits are not reached.
- Operations, such as renaming a file or directory or moving files between qtrees, on quota targets, for which quota hard limits are configured, might take longer when compared to similar operations on FlexVol volumes.

Examples of quota enforcement for FlexGroup volumes

You can use the examples to understand how to configure quotas with limits in ONTAP 9.5 and later.

Example 1: Enforcing a quota rule with disk limits

1. You should create a quota policy rule of type user with both an achievable soft disk limit and hard disk limit.

```
cluster1::> volume quota policy rule create -vserver vs0 -policy-name
default -volume FG -type user -target "" -qtree "" -disk-limit 1T -soft
-disk-limit 800G
```

2. You can view the quota policy rule:

```
cluster1::> volume quota policy rule show -vserver vs0 -policy-name
default -volume FG
Vserver: vs0
                       Policy: default
                                               Volume: FG
                                          Soft
                                                         Soft
                                          Disk Files
                                                        Files
                      User
                                 Disk
               Qtree Mapping
                               Limit
                                         Limit Limit
                                                        Limit
Type
    Target
Threshold
user ""
              11 11
                      off
                                   1TB
                                         800GB
```

3. To activate the new quota rule, you initialize quotas on the volume:

```
cluster1::> volume quota on -vserver vs0 -volume FG -foreground true [Job 49] Job succeeded: Successful
```

4. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

```
cluster1::> volume quota report -vserver vs0 -volume FG
Vserver: vs0
                                ----Disk---- ----Files-----
                                                            Quota
Volume
                 Type ID
                                Used Limit
                                             Used Limit
        Tree
Specifier
-----
FG
                                 50GB
                 user
                       root
                                                 1
FG
                 user
                                 800GB
                                         1TB
                                                 0
2 entries were displayed.
```

After the hard disk limit is reached, the quota policy rule target (user, in this case) is blocked from writing more data to the files.

Example 2: Enforcing a quota rule for multiple users

1. You should create a quota policy rule of type user, where multiple users are specified in the quota target (UNIX users, SMB users, or a combination of both) and where the rule has both an achievable soft disk limit and hard disk limit.

```
cluster1::> quota policy rule create -vserver vs0 -policy-name default
-volume FG -type user -target "rdavis, ABCCORP\RobertDavis" -qtree ""
-disk-limit 1TB -soft-disk-limit 800GB
```

2. You can view the quota policy rule:

```
cluster1::> quota policy rule show -vserver vs0 -policy-name default
-volume FG
                        Policy: default
Vserver: vs0
                                                 Volume: FG
                                            Soft
                                                            Soft
                       User
                                  Disk
                                            Disk Files
                                                           Files
Type
      Target
               Qtree
                       Mapping
                                  Limit
                                           Limit Limit
                                                           Limit
Threshold
      "rdavis, ABCCORP\RobertDavis" "" off 1TB 800GB - -
user
```

3. To activate the new quota rule, you initialize quotas on the volume:

```
cluster1::> volume quota on -vserver vs0 -volume FG -foreground true [Job 49] Job succeeded: Successful
```

4. You can verify that the quota state is active:

5. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

The quota limit is shared among all users listed in the quota target.

After the hard disk limit is reached, users listed in the quota target are blocked from writing more data to the files.

Example 3: Enforcing quota with user mapping enabled

1. You should create a quota policy rule of type user, specify a UNIX user or a Windows user as the quota target with user-mapping set to on, and create the rule with both an achievable soft disk limit and hard disk limit.

The mapping between UNIX and Windows users must be configured earlier by using the vserver name-mapping create command.

```
cluster1::> quota policy rule create -vserver vs0 -policy-name default
-volume FG -type user -target rdavis -qtree "" -disk-limit 1TB -soft
-disk-limit 800GB -user-mapping on
```

2. You can view the quota policy rule:

```
cluster1::> quota policy rule show -vserver vs0 -policy-name default
-volume FG
             Policy: default Volume: FG
Vserver: vs0
                                  Soft
                                               Soft
                  User Disk
                                 Disk Files
                                             Files
                          Limit Limit Limit
Type Target Qtree Mapping
                                             Limit
Threshold
_____
user rdavis "" on
                           1TB
                                800GB
```

3. To activate the new quota rule, you initialize quotas on the volume:

```
cluster1::> volume quota on -vserver vs0 -volume FG -foreground true
[Job 49] Job succeeded: Successful
```

4. You can verify that the quota state is active:

5. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

The quota limit is shared between the user listed in the quota target and its corresponding Windows or UNIX user.

After the hard disk limit is reached, both the user listed in the quota target and its corresponding Windows or UNIX user is blocked from writing more data to the files.

Example 4: Verifying the qtree size when quota is enabled

1. You should create a quota policy rule of type tree and where the rule has both an achievable soft disk limit and hard disk limit.

```
cluster1::> quota policy rule create -vserver vs0 -policy-name default
-volume FG -type tree -target tree_4118314302 -qtree "" -disk-limit 48GB
-soft-disk-limit 30GB
```

2. You can view the quota policy rule:

```
Cluster1::> quota policy rule show -vserver vs0

Vserver: vs0

Policy: default

Soft

Soft

User

Disk

Disk

Files

Files

Type

Target

Qtree

Mapping

Limit

Limit

Limit

Limit

Threshold

Threshold

Three

tree tree_4118314302 "" - 48GB - 20 -
```

3. To activate the new quota rule, you initialize quotas on the volume:

```
cluster1::> volume quota on -vserver vs0 -volume FG -foreground true
[Job 49] Job succeeded: Successful
```

a. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

The quota limit is shared between the user listed in the quota target and its corresponding Windows or UNIX user.

4. From an NFS client, use the df command to view the total space usage, available space, and the used space.

```
scsps0472342001# df -m /t/10.53.2.189/FG-3/tree_4118314302
Filesystem 1M-blocks Used Available Use% Mounted on
10.53.2.189/FG-3 49152 31078 18074 63% /t/10.53.2.189/FG-3
```

With hard limit, the space usage is calculated from an NFS client as follows:

- Total space usage = hard limit for tree
- Free space = Hard limit minus qtree space usage Without hard limit, the space usage is calculated from an NFS client as follows:
- Space usage = quota usage
- Total space = Sum of quota usage and physical free space in the volume
- 5. From the SMB share, use Windows Explorer to view the total space usage, available space, and the used space.

From an SMB share, you should be aware of the following considerations for calculating the space usage:

- The user quota hard limit for the user and group is taken into consideration for calculating the total available space.
- The minimum value among the free space of the tree quota rule, the user quota rule, and the group quota rule is considered as the free space for the SMB share.
- The total space usage is variable for SMB and depends on the hard limit that corresponds to the minimum free space among the tree, user, and group.

Apply rules and limits on the FlexGroup volume

Steps

1. Create quota rules for targets: volume quota policy rule create -vserver vs0 -policy -name quota_policy_of_the_rule -volume flexgroup_vol -type {tree|user|group} -target target_for_rule -qtree qtree_name [-disk-limit hard_disk_limit_size] [-file-limit hard_limit_number_of_files] [-threshold

threshold_disk_limit_size] [-soft-disk-limit soft_disk_limit_size] [-softfile-limit soft limit number of files]

- The quota target type can be user, group, or tree for FlexGroup volumes.
- A path is not supported as the target when creating quota rules for FlexGroup volumes.
- Beginning with ONTAP 9.5, you can specify hard disk limit, hard file limit, soft disk limit, soft file limit, and threshold limit quotas for FlexGroup volumes.

In ONTAP 9.4 and earlier, you cannot specify the disk limit, file limit, threshold for disk limit, soft disk limit, or soft file limit when you create quota rules for FlexGroup volumes.

The following example shows a default quota rule being created for the user target type:

```
cluster1::> volume quota policy rule create -vserver vs0 -policy-name
quota_policy_vs0_1 -volume fg1 -type user -target "" -qtree ""
```

The following example shows a tree quota rule being created for the qtree named qtree1:

```
cluster1::> volume quota policy rule create -policy-name default -vserver
vs0 -volume fg1 -type tree -target "qtree1"
```

1. Activate the quotas for the specified FlexGroup volume: volume quota on -vserver svm_name -volume flexgroup vol -foreground true

```
cluster1::> volume quota on -vserver vs0 -volume fg1 -foreground true
```

1. Monitor the state of quota initialization: volume quota show -vserver svm name

FlexGroup volumes might show the mixed state, which indicates that all of the constituent volumes are not in the same state yet.

```
cluster1::> volume quota show -vserver vs0

Scan

Vserver Volume State Status

-----
vs0 fg1 initializing 95%
vs0 vol1 off -
2 entries were displayed.
```

1. View the quota report for the FlexGroup volume with active quotas: volume quota report -vserver svm name -volume flexgroup vol

You cannot specify a path with the volume quota report command for FlexGroup volumes.

The following example shows the user quota for the FlexGroup volume fg1:

cluster1:: Vserver:		quota rep	port -vse	erver vs0	-volume	e fg1		
				Di	sk	File	es	
Quota	_	_						
Volume	Tree	Type	ID	Used	Limit	Used	Limit	
Specifier								
fg1		user	*	0B	-	0	-	*
fg1		user	root	1GB	-	1	_	*
2 entrie	s were di	isplayed.						

The following example shows the tree quota for the FlexGroup volume fq1:

cluster1 Vserver:		ne quota	report -	vserver vs0	-volu	ıme fg1		
				Disk		File	es	Quota
Volume	Tree	Type	ID	Used L	imit	Used	Limit	
Specifie	er							
	-							
fg1	qtree1	tree	1	68KB	-	18	-	
qtree1								
fg1		tree	*	0B	-	0	-	*
2 entrie	es were di	isplayed.						

Results

The quota rules and limits are applied on the FlexGroup volume.

The usage might reach up to 5 percent higher than a configured hard limit before ONTAP enforces the quota by rejecting further traffic.

Related information

ONTAP command reference

Enable storage efficiency on ONTAP FlexGroup volumes

You can run deduplication and data compression together or independently on a FlexGroup volume to achieve optimal space savings.

Before you begin

The FlexGroup volume must be online.

Steps

1. Enable storage efficiency on the FlexGroup volume: volume efficiency on -vserver svm_name -volume volume name

Storage efficiency operations are enabled on all the constituents of the FlexGroup volume.

If a FlexGroup volume is expanded after storage efficiency is enabled on the volume, storage efficiency is automatically enabled on the new constituents.

2. Enable the required storage efficiency operation on the FlexGroup volume by using the volume efficiency modify command.

You can enable inline deduplication, postprocess deduplication, inline compression, and postprocess compression on FlexGroup volumes. You can also set the type of compression (secondary or adaptive) and specify a schedule or efficiency policy for the FlexGroup volume.

3. If you are not using schedules or efficiency policies for running the storage efficiency operations, start the efficiency operation: volume efficiency start -vserver svm name -volume volume name

If deduplication and data compression are enabled on a volume, data compression is run initially followed by deduplication. This command fails if any efficiency operation is already active on the FlexGroup volume.

4. Verify the efficiency operations that are enabled on the FlexGroup volume: volume efficiency show -vserver svm name -volume volume name

```
cluster1::> volume efficiency show -vserver vs1 -volume fq1
             Vserver Name: vs1
              Volume Name: fg1
              Volume Path: /vol/fg1
                    State: Enabled
                   Status: Idle
                 Progress: Idle for 17:07:25
                     Type: Regular
                 Schedule: sun-sat@0
             Compression: true
      Inline Compression: true
Incompressible Data Detection: false
      Constituent Volume: false
Compression Quick Check File Size: 524288000
           Inline Dedupe: true
         Data Compaction: false
```

Protect ONTAP FlexGroup volumes using snapshots

You can create snapshot policies that automatically manage the creation of snapshots or

you can manually create snapshots for FlexGroup volumes. A valid snapshot is created for a FlexGroup volume only after ONTAP can successfully create a snapshot for each constituent of the FlexGroup volume.

About this task

- If you have multiple FlexGroup volumes associated with a snapshot policy, you should ensure that the FlexGroup volumes schedules do not overlap.
- Beginning with ONTAP 9.8, the maximum number of snapshots supported on a FlexGroup volume is 1023.



Beginning with ONTAP 9.8, the volume snapshot show command for FlexGroup volumes reports snapshot size using logical blocks, rather than calculating the youngest owned blocks. This new size calculation method might make the snapshot size appear larger than calculations in earlier versions of ONTAP.

Steps

1. Create a snapshot policy or manually create a snapshot:

If you want to create a	Enter this	s command		
Snapshot policy	volume snapshot policy create			
	i	The schedules that are associated with the snapshot policy of a FlexGroup volume must have an interval greater than 30 minutes.		
	default	u create a FlexGroup volume, the snapshot policy is applied to the p volume.		
Snapshot manually	volume	snapshot create		
	i	After you create a snapshot for a FlexGroup volume, you cannot modify the attributes of the snapshot. If you want to modify the attributes, you must delete and then re-create the snapshot.		

Client access to the FlexGroup volume is briefly quiesced when a snapshot is created.

1. Verify that a valid snapshot is created for the FlexGroup volume: volume snapshot show -volume volume name -fields state

```
cluster1::> volume snapshot show -volume fg -fields state
vserver volume snapshot state
-----
fg_vs fg hourly.2016-08-23_0505 valid
```

2. View the snapshots for the constituents of the FlexGroup volume: volume snapshot show -is -constituent true

cluster1	::> volum	ne snapshot show -is-constituent true		
Block Vserver Used%		Snapshot	Size To	otal%
fg_vs 27%	fg0001	hourly.2016-08-23_0505	72MB	0%
218	fg0002	hourly.2016-08-23_0505	72MB	0%
27%	fg0003	hourly.2016-08-23 0505	72MB	0%
27%		110 all 1 1 20 10 00 20 _0000	, 2115	0 0
	fg0016	hourly.2016-08-23 0505	72MB	0%
27%		<u></u>		

Move constituents from ONTAP FlexGroup volumes

You can move the constituents of a FlexGroup volume from one aggregate to another to balance the load when certain constituents experience more traffic. Moving constituents also helps in freeing up space on an aggregate for resizing the existing constituents.

Before you begin

To move a FlexGroup volume constituent that is in a SnapMirror relationship, you must have initialized the SnapMirror relationship.

About this task

You cannot perform a volume move operation while the constituents of the FlexGroup volume are being expanded.

Steps

1. Identify the FlexGroup volume constituent that you want to move:

volume show -vserver svm name -is-constituent true

cluster1	::> volume sh	ow -vserver v	s2 -is-cons	stituent t	rue
Vserver	Volume	Aggregate	State	Type	Size
Availabl	e Used%				
vs2	fg1	-	online	RW	400TB
15.12TB	62%				
vs2	fg10001	aggr1	online	RW	25TB
8.12MB	59%				
vs2	fg10002	aggr2	online	RW	25TB
2.50TB	90%				

2. Identify an aggregate to which you can move the FlexGroup volume constituent:

volume move target-aggr show -vserver svm name -volume vol constituent name

The available space in the aggregate that you select must be greater than the size of the FlexGroup volume constituent that you are moving.

3. Verify that the FlexGroup volume constituent can be moved to the intended aggregate:

volume move start -vserver svm_name -volume vol_constituent_name -destination
-aggregate aggr_name -perform-validation-only true

cluster1::> volume move start -vserver vs2 -volume fg1_0002 -destination
-aggregate node12a_aggr3 -perform-validation-only true
Validation succeeded.

4. Move the FlexGroup volume constituent:

volume move start -vserver svm_name -volume vol_constituent_name -destination
-aggregate aggr name [-allow-mixed-aggr-types {true|false}]

The volume move operation runs as a background process.

Beginning with ONTAP 9.5, you can move FlexGroup volume constituents from a Fabric Pool to a non-Fabric Pool, or vice versa by setting the <code>-allow-mixed-aggr-types</code> parameter to <code>true</code>. By default, the <code>-allow-mixed-aggr-types</code> option is set to <code>false</code>.



You cannot use the volume move command for enabling encryption on FlexGroup volumes.

 $\label{eq:cluster1::} cluster1::> volume move start -vserver vs2 -volume fg1_002 -destination -aggregate node12a_aggr3$



If the volume move operation fails due to an active SnapMirror operation, you should abort the SnapMirror operation by using the <code>snapmirror</code> abort <code>-h</code> command. In some cases, the SnapMirror abort operation might also fail. In such situations, you should abort the volume move operation and retry later. Learn more about <code>snapmirror</code> abort in the ONTAP command reference.

5. Verify the state of the volume move operation:

```
volume move show -volume vol constituent name
```

The following example shows the state of a FlexGroup constituent volume that completed the replication phase and is in the cutover phase of the volume move operation:

Use aggregates in FabricPool for existing ONTAP FlexGroup volumes

Beginning with ONTAP 9.5, FabricPool is supported for FlexGroup volumes. If you want to use aggregates in FabricPool for your existing FlexGroup volumes, you can either convert the aggregates on which the FlexGroup volume resides to aggregates in FabricPool or migrate the FlexGroup volume constituents to aggregates in FabricPool.

Before you begin

- The FlexGroup volume must have space-guarantee set to none.
- If you want to convert the aggregates on which the FlexGroup volume resides to aggregates in FabricPool, the aggregates must be using all SSD disks.

About this task

If an existing FlexGroup volume resides on non-SSD aggregates, you must migrate the FlexGroup volume constituents to aggregates in FabricPool.

Choices

- To convert the aggregates on which the FlexGroup volume resides to aggregates in FabricPool, perform the following steps:
 - a. Set the tiering policy on the existing FlexGroup volume: volume modify -volume flexgroup name -tiering-policy [auto|snapshot|none|backup]

```
cluster-2::> volume modify -volume fg1 -tiering-policy auto
```

b. Identify the aggregates on which the FlexGroup volume resides: volume show -volume flexgroup_name -fields aggr-list

```
cluster-2::> volume show -volume fg1 -fields aggr-list
vserver volume aggr-list
-----
vs1 fg1 aggr1,aggr3
```

c. Attach an object store to each aggregate listed in the aggregate list: storage aggregate objectstore attach -aggregate aggregate name -name object-store-name -allow -flexgroup true

You must attach all of the aggregates to an object store.

```
cluster-2::> storage aggregate object-store attach -aggregate aggr1
-object-store-name Amazon01B1
```

- To migrate the FlexGroup volume constituents to aggregates in FabricPool, perform the following steps:
 - a. Set the tiering policy on the existing FlexGroup volume: volume modify -volume flexgroup_name -tiering-policy [auto|snapshot|none|backup]

```
cluster-2::> volume modify -volume fg1 -tiering-policy auto
```

b. Move each constituent of the FlexGroup volume to an aggregate in FabricPool in the same cluster: volume move start -volume constituent-volume -destination-aggregate FabricPool_aggregate -allow-mixed-aggr-types true

You must move all FlexGroup volume constituents to aggregates in FabricPool (in case the FlexGroup volume constituents are on mixed aggregate types) and ensure that all the constituents are balanced across the nodes in the cluster.

Related information

- · Disk and aggregate management
- storage aggregate object-store attach

Balance ONTAP FlexGroup volumes by redistributing file data

Beginning with ONTAP 9.16.1, you can enable advanced capacity balancing to enable data distribution between FlexGroup member volumes when very large files grow and consume space on one member volume.

Advanced capacity balancing expands the granular data functionality introduced in ONTAP 9.12.1, which enables ONTAP to rebalance FlexGroup volumes by moving files to other members. Beginning with ONTAP 9.16.1, when advanced capacity balancing is enabled with the <code>-granular-data</code> advanced option, both the "basic" file rebalancing capabilities as well as the advanced capacity capabilities are activated.



Both file rebalancing and advanced capacity balancing are disabled by default. After these features are enabled they cannot be disabled. If you need to disable capacity balancing, you must restore from a snapshot that was created before advanced capacity balancing was enabled.

Advanced capacity balancing is triggered by new writes reaching to 10GB or 1% of a volume's free space.

How files are distributed

If a file is created or grows large enough to trigger advanced capacity balancing, the file is distributed in stripes between 1GB and 10GB across the member FlexGroup volumes.

When advanced capacity balancing is enabled, ONTAP will not retroactively stripe existing large files. If an existing large file continues to grow after advanced capacity balancing is enabled, new content in existing large files might be striped across member FlexGroup volumes depending on the file's size and available space.

One way advanced capacity balancing determines stripe width is by using the amount of free space available on the member volume. Advanced capacity balancing creates a file stripe that is 1% of the available free space available. This means that stripes can start out larger if more space is available, and they become smaller as the FlexGroup fills up.

Supported protocols

Advanced capacity balancing is supported with the following protocols:

- NFSv3, NFSv4, NFSv4.1
- pNFS
- SMB

Enable advanced capacity balancing

Advanced capacity balancing is disabled by default. You must enable advanced capacity balancing to automatically balance FlexGroup capacity. Keep in mind that you cannot disable this feature once you enable it, but you can restore from a snapshot created before advanced capacity balancing was enabled.

Before you begin

- All nodes in the cluster must be running ONTAP 9.16.1 or later.
- You cannot revert to a release earlier than ONTAP 9.16.1 if advanced capacity balancing is enabled. If you
 need to revert, you must first restore from a snapshot created before advanced capacity balancing was
 enabled.
- If NFS copy offload has been enabled (vserver nfs -vstorage enabled) on an SVM, you cannot enable advanced capacity balancing on a FlexGroup volume. Similarly, if you have advanced capacity balancing enabled on any FlexGroup volume in an SVM, you cannot enable NFS copy offload.
- Advanced capacity balancing is not supported with FlexCache write-back.
- SnapMirror transfers are not supported with ONTAP versions earlier than ONTAP 9.16.1 when advanced capacity balancing is enabled on volumes in clusters running ONTAP 9.16.1 or later.

About this task

During creation of DP destination volumes using either of the granular data options (basic or advanced), the destination displays the setting as "disabled" until the SnapMirror transfer completes. After the transfer completes, the DP destination displays granular data as "enabled".

Enable advanced capacity balancing during FlexGroup creation

Steps

You can use System Manager or the ONTAP CLI to enable advanced capacity balancing when you create a new FlexGroup volume.

- 1. Navigate to **Storage > Volumes** and click + Add
- 2. In the Add volume window, enter the volume name and size. Then click More Options.
- 3. Under Storage and optimization, select Distribute volume data across the cluster (FlexGroup).
- 4. Select Advanced capacity balancing.
- 5. Finish configuring the volume and click Save.

CLI

1. Create a volume with advanced capacity balancing enabled:

```
volume create -vserver <svm name> -volume <volume name> -size <volume
size> -auto-provision-as flexgroup -junction-path /<path> -granular
-data advanced
```

Example:

```
volume create -vserver vs0 -volume newvol-size 1TB -auto-provision
-as flexgroup -junction-path /newvol -granular-data advanced
```

Enable advanced capacity balancing on existing FlexGroup volumes

Steps

You can use System Manager or the ONTAP CLI to enable advanced capacity balancing.

- 1. Navigate to **Storage > Volumes**, click **!**, and choose **Edit > Volume**.
- 2. In the Edit volume window, under Storage and optimization, select Advanced capacity balancing.
- 3. Click Save.

CLI

1. Modify an existing FlexGroup volume to enable advanced capacity balancing:

```
volume modify -vserver <svm name> -volume <volume name> -granular
-data advanced
```

Example:

volume modify -vserver vs0 -volume newvol -granular-data advanced

Rebalance ONTAP FlexGroup volumes by moving files

Beginning with ONTAP 9.12.1, you can rebalance FlexGroup volumes by non-disruptively moving files from one constituent in a FlexGroup to another constituent.

FlexGroup rebalancing helps redistribute capacity when imbalances develop over time due to the addition of new files and file growth. After you manually start the rebalance operation, ONTAP selects the files and moves them automatically and non-disruptively.



You should be aware that FlexGroup rebalancing degrades system performance when large numbers of files are moved as part of a single rebalancing event or over multiple rebalancing events because of the creation of multi-part inodes. Every file moved as part of a rebalancing event has 2 multi-part inodes associated with that file. The larger the number of files with multipart inodes as a percentage of the total number of files in a FlexGroup, the greater the performance impact. Certain use cases, such as a FlexVol to FlexGroup conversion, can result in a significant amount of multi-part inode creation.

Rebalancing is available only when all nodes in the cluster are running ONTAP 9.12.1 or later releases. You must enable granular data functionality on any FlexGroup volume that runs the rebalancing operation. Once that functionality is enabled, you cannot revert to ONTAP 9.11.1 and earlier versions unless you delete this volume or restore from a snapshot that was created before the setting was enabled.

Beginning with ONTAP 9.14.1, ONTAP introduces an algorithm to non-disruptively and proactively move files in volumes that have granular data enabled without user interaction. The algorithm operates in very specific, targeted scenarios to alleviate performance bottlenecks. Scenarios where this algorithm might act include very heavy write load on a particular set of files on one node in the cluster or a continually growing file in a very hot parent directory.

Beginning with ONTAP 9.16.1, you can also enable advanced capacity balancing to redistribute a large file's data between FlexGroup member volumes.

FlexGroup rebalancing considerations

You should be aware of how FlexGroup rebalancing works and how it interacts with other ONTAP features.

FlexVol to FlexGroup conversion

It is recommended that you *not* use automatic FlexGroup rebalancing after a FlexVol to FlexGroup conversion. Instead, you can redistribute existing files by using the volume rebalance file-move start command, available in ONTAP 9.10.1 and later. This operation is non-disruptive by default (-is-disruptive false). If some busy files cannot be moved, you can rerun the command in disruptive mode (-is-disruptive true) during a planned maintenance window. Learn more about volume rebalance file-move start in the ONTAP command reference.

Rebalancing with the automatic FlexGroup rebalancing feature can degrade performance when moving large numbers of files, like when you perform a FlexVol to FlexGroup conversion, and as much as 50 to 85% of the data on the FlexVol volume is moved to a new constituent.

· Minimum and maximum file size

File selection for automatic rebalancing is based on blocks saved. The minimum file size considered for rebalancing is 100 MB by default (can be configured as low as 20 MB using the min-file-size parameter shown below) and the maximum file size is 100 GB.

· Files in snapshots

You can configure FlexGroup rebalancing to only consider files to be moved which are not currently present in any snapshots. When rebalancing is started, a notification displays if a snapshot operation is scheduled anytime during a rebalancing operation.

Snapshots are restricted if a file is being moved and is undergoing framing at the destination. A snapshot restore operation is not allowed while file rebalancing is in progress.

Any snapshot created after the granular-data option is enabled cannot be replicated to a system running ONTAP 9.11.1 and earlier versions because ONTAP 9.11.1 and earlier versions do not support multi-part inodes.

SnapMirror operations

FlexGroup rebalancing should take place between scheduled SnapMirror operations. A SnapMirror operation might fail if a file is being relocated before a SnapMirror operation begins if that file move does not complete within the 24-minute SnapMirror retry period. Any new file relocation that begins after a SnapMirror transfer has started will not fail.

File-based compression storage efficiency

With file-based compression storage efficiency, the file is decompressed before it's moved to the destination, so the compression savings is lost. The compression savings is regained after a manually initiated background scanner runs on the FlexGroup volume after rebalancing. However, if any file is associated with a snapshot on any volume, the file will be ignored for compression.

Deduplication

Moving deduplicated files can cause increased overall usage for the FlexGroup volume. During file rebalancing, only unique blocks are moved to the destination, freeing that capacity on the source. Shared blocks remain on the source and are copied to the destination. While this achieves the goal of reducing the

used capacity on a nearly full source constituent, it can also lead to increased overall usage on the FlexGroup volume due to copies of shared blocks on the new destinations. This is also possible when files that are part of a snapshot are moved. The space savings is not fully recognized until the snapshot schedule recycles and there are no longer copies of the files in snapshots.

FlexClone volumes

If file rebalancing is in progress when a FlexClone volume is created, the rebalancing will not be performed on the FlexClone volume. Rebalancing on the FlexClone volume should be performed after it is created.

· File move

When a file is moved during a FlexGroup rebalancing operation, the file size is reported as part of quota accounting on both the source and destination constituents. Once the move is completed, quota accounting returns to normal, and the file size is only reported on the new destination.

· Autonomous Ransomware Protection

Beginning with ONTAP 9.13.1, Autonomous Ransomware Protection is supported during disruptive and non-disruptive rebalance operations.

Object store volumes

Volume capacity rebalancing is not supported on object store volumes, such as S3 buckets.

Enable FlexGroup rebalancing

Beginning with ONTAP 9.12.1, you can enable automatic non-disruptive FlexGroup volume rebalancing to redistribute files between FlexGroup constituents.

Beginning with ONTAP 9.13.1, you can schedule a single FlexGroup rebalancing operation to begin at a date and time in the future.

Before you begin

You must have enabled the granular-data option on the FlexGroup volume before enabling FlexGroup rebalancing. You can enable it by using one of these methods:

- When you create FlexGroup volume using the volume create command
- By modifying an existing FlexGroup volume to enable the setting using the volume modify command
- Setting it automatically when FlexGroup rebalancing is initiated using the volume rebalance command



If you are using ONTAP 9.16.1 or later and FlexGroup advanced capacity balancing is enabled using either the granular-data advanced option in the ONTAP CLI or using System Manager, FlexGroup rebalancing is also enabled.

Steps

You can manage FlexGroup rebalancing by using ONTAP System Manager or the ONTAP CLI.

- 1. Navigate to **Storage > Volumes** and locate the FlexGroup volume to rebalance.
- 2. Select view the volume details.
- 3. Under FlexGroup Balance Status select Rebalance.



The **Rebalance** option is only available when the FlexGroup status is out of balance.

- 4. In the **Rebalance Volume** window, change the default settings as needed.
- 5. To schedule the rebalancing operation, select **Rebalance Later** and enter the date and time.

CLI

1. Start automatic rebalancing:

```
volume rebalance start -vserver <SVM name> -volume <volume name>
```

Optionally, you can specify the following options:

[[-max-runtime] <time interval>] Maximum Runtime

[-max-threshold <percent>] Maximum Imbalance Threshold per Constituent

[-min-threshold <percent>] Minimum Imbalance Threshold per Constituent

[-max-file-moves <integer>] Maximum Concurrent File Moves per Constituent

[-min-file-size {<integer>[KB|MB|GB|TB|PB]}] Minimum file size

[-start-time <mm/dd/yyyy-00:00:00>] Schedule rebalance start date and time

[-exclude-snapshots {true|false}] Exclude files stuck in snapshots

Example:

```
volume rebalance start -vserver vs0 -volume fg1
```

Modify FlexGroup rebalance configurations

You can change a FlexGroup rebalancing configuration to update the imbalance threshold, number of concurrent files moves minimum file size, maximum runtime, and to include or exclude snapshots. Options to modify your FlexGroup rebalancing schedule are available beginning with ONTAP 9.13.1.

- 1. Navigate to **Storage > Volumes** and locate the FlexGroup volume to rebalance.
- 2. Select view the volume details.
- 3. Under FlexGroup Balance Status select Rebalance.



The **Rebalance** option is only available when the FlexGroup status is out of balance.

4. In the **Rebalance Volume** window, change the default settings as needed.

CLI

1. Modify automatic rebalancing:

volume rebalance modify -vserver <SVM name> -volume <volume name>

You can specify one or more of the following options:

[[-max-runtime] <time interval>] Maximum Runtime

[-max-threshold <percent>] Maximum Imbalance Threshold per Constituent

[-min-threshold <percent>] Minimum Imbalance Threshold per Constituent

[-max-file-moves <integer>] Maximum Concurrent File Moves per Constituent

[-min-file-size {<integer>[KB|MB|GB|TB|PB]}] Minimum file size

[-start-time <mm/dd/yyyy-00:00:00>] Schedule rebalance start date and time

[-exclude-snapshots {true|false}] Exclude files stuck in snapshots

Stop FlexGroup rebalance

After FlexGroup rebalancing is enabled or scheduled, you can stop it at any time.

System Manager

- 1. Navigate to **Storage > Volumes** and locate the FlexGroup volume.
- 2. Select view the volume details.
- 3. Select **Stop Rebalance**.

CLI

1. Stop FlexGroup rebalancing:

volume rebalance stop -vserver <SVM name> -volume <volume name>

View FlexGroup rebalance status	
You can display the status about a FlexGroup rebalance operation, the FlexGroup rebalance configuration, the rebalance operation time, and the rebalance instance details.	

- 1. Navigate to **Storage > Volumes** and locate the FlexGroup volume.
- 2. Select view the FlexGroup details.
- 3. FlexGroup Balance Status is displayed near the bottom of the details pane.
- 4. To view information about the last rebalance operation, select Last Volume Rebalance Status.

CLI

1. View the status of a FlexGroup rebalance operation:

```
volume rebalance show
```

Example of rebalance state:

Example of rebalance configuration details:

Example of rebalance time details:

Example of rebalance instance details:

```
> volume rebalance show -instance
Vserver Name: vs0
Volume Name: fq1
Is Constituent: false
Rebalance State: idle
Rebalance Notice Messages: -
Total Size: 4GB
AFS Used Size: 115.3MB
Constituent Target Used Size: -
Imbalance Size: 8KB
Imbalance Percentage: 0%
Moved Data Size: -
Maximum Constituent Imbalance Percentage: 1%
Rebalance Start Time: Wed Jul 20 16:06:11 2022
Rebalance Stop Time: -
Rebalance Runtime: 0h1m32s
Rebalance Maximum Runtime: 6h0m0s
Maximum Imbalance Threshold per Constituent: 20%
Minimum Imbalance Threshold per Constituent: 5%
Maximum Concurrent File Moves per Constituent: 25
Minimum File Size: 4KB
Exclude Files Stuck in snapshots: true
```

Data protection for FlexGroup volumes

Data protection for ONTAP FlexGroup volumes workflow summary

You can create SnapMirror disaster recovery (DR) relationships for FlexGroup volumes. You can also backup and restore FlexGroup volumes by using SnapVault technology, and you can create a unified data protection relationship that uses the same destination for backup and DR.

About this task

The SnapMirror relationship type is always XDP for FlexGroup volumes. The type of data protection that is provided by a SnapMirror relationship is determined by the replication policy that you use. You can use either the default policy or a custom policy of the required type for the replication relationship that you want to create.



Peer the clusters and SVMs

If the clusters and SVMs are not already peered, create the cluster peers and the SVM peers.



Create a job schedule

You must create a job schedule to determine when SnapMirror updates will take place.



Depending on the type of data protection, follow one of these paths:

• If SnapMirror DR:

Create a SnapMirror relationship. When you create the relationship, you can select the default policy MirrorAllSnapshots or a custom policy of type async-mirror.

• If SnapMirror vault:

Create a SnapMirror vault relationship. When you create the relationship, you can select the default policy XDPDefault or a custom policy of type vault.

If unified data protection:

Create a unified data protection relationship. When you create the relationship, you can select the default policy MirrorAndVault or a custom policy of type mirror-vault.

Create SnapMirror relationships for ONTAP FlexGroup volumes

You can create a SnapMirror relationship between the source FlexGroup volume and the destination FlexGroup volume on a peered SVM for replicating data for disaster recovery. You can use the mirror copies of the FlexGroup volume to recover data when a disaster occurs.

Before you begin

You must have created the cluster peering relationship and SVM peering relationship.

Cluster and SVM peering

About this task

- Beginning with ONTAP 9.9.1, you can use the ONTAP CLI to create SnapMirror cascade and fanout relationships for FlexGroup volumes. For details, see Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes.
- You can create both intercluster SnapMirror relationships and intracluster SnapMirror relationships for FlexGroup volumes.

• Beginning with ONTAP 9.3, you can expand FlexGroup volumes that are in a SnapMirror relationship.

If you are using a version of ONTAP earlier than ONTAP 9.3, do not expand FlexGroup volumes after a SnapMirror relationship is established; however, you can increase the capacity of FlexGroup volumes after establishing a SnapMirror relationship. If you expand the source FlexGroup volume after breaking the SnapMirror relationship in releases earlier than ONTAP 9.3, you must perform a baseline transfer to the destination FlexGroup volume.

Steps

- 1. Create a destination FlexGroup volume of type DP that has the same number of constituents as that of the source FlexGroup volume:
 - a. From the source cluster, determine the number of constituents in the source FlexGroup volume: volume show -volume volume name* -is-constituent true

server	Volume	Aggregate S	State	Type	Size
Available		33 3			
7SS	srcFG	_	online	RW	400TB
72.86GB	56%				
7SS	srcFG0001	Aggr_cmode	online	RW	25GB
0.86TB	56%				
7SS	srcFG0002	aggr1	online	RW	25TB
0.86TB	56%				
7SS	srcFG0003	Aggr_cmode	online	RW	25TB
0.72TB	57%				
7SS	srcFG0004	aggr1	online	RW	25TB
.0.73TB	57%				
7SS	srcFG0005	Aggr_cmode	online	RW	25TB
0.67TB	57%				
7SS	srcFG0006	aggr1	online	RW	25TB
.0.64TB	57%				
7SS	srcFG0007	Aggr_cmode	online	RW	25TB
.0.63TB	57%				

b. From the destination cluster, create a destination FlexGroup volume of type DP with the same number of constituents as that of the source FlexGroup volume.

cluster2::> volume create -vserver vsd -aggr-list aggr1,aggr2 -aggr -list-multiplier 8 -size 400TB -type DP dstFG Warning: The FlexGroup volume "dstFG" will be created with the following number of constituents of size 25TB: 16. Do you want to continue? $\{y|n\}$: y [Job 766] Job succeeded: Successful

c. From the destination cluster, verify the number of constituents in the destination FlexGroup volume: volume show -volume volume name* -is-constituent true

Vserver	Volume	Aggregate	State	Type	Size
Available	used%				
			2.1		400
vsd		_	online	DP	400TB
172.86GB					
	dstFG0001	Aggr_cmode	online	DP	25GB
10.86TB	56%				
vsd	dstFG0002	aggr1	online	DP	25TB
10.86TB	56%				
vsd	dstFG0003	Aggr_cmode	online	DP	25TB
10.72TB	57%				
vsd	dstFG 0004	aggr1	online	DP	25TB
10.73TB	<u>—</u> 57%				
vsd	dstFG 0005	Aggr cmode	online	DP	25TB
10.67TB		33 _			
	dstFG 0006	aggr1	online	DP	25TB
10.64TB		ر ر			_
	dstFG 0007	Agar cmode	online	DΡ	25TB
10.63TB		11991_0000	01111110	Di	2016
10.0315	J 1 70				

2. Create a job schedule: job schedule cron create -name job_name -month month -dayofweek day of week -day day of month -hour hour -minute minute

For the <code>-month</code>, <code>-dayofweek</code>, and <code>-hour</code> options, you can specify all to run the job every month, every day of the week, and every hour, respectively.

The following example creates a job schedule named my weekly that runs on Saturdays at 3:00 a.m.:

```
cluster1::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

3. Create a custom policy of type async-mirror for the SnapMirror relationship: snapmirror policy create -vserver SVM -policy snapmirror policy -type async-mirror

If you do not create a custom policy, you should specify the MirrorAllSnapshots policy for SnapMirror relationships.

4. From the destination cluster, create a SnapMirror relationship between the source FlexGroup volume and the destination FlexGroup volume: snapmirror create -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -type XDP -policy snapmirror_policy -schedule sched name

SnapMirror relationships for FlexGroup volumes must be of type XDP.

If you specify a throttle value for the SnapMirror relationship for the FlexGroup volume, each constituent uses the same throttle value. The throttle value is not divided among the constituents.



You cannot use SnapMirror labels of snapshots for FlexGroup volumes.

In ONTAP 9.4 and earlier, if the policy is not specified with the <code>snapmirror</code> <code>create</code> command, the <code>MirrorAllSnapshots</code> policy is used by default. In ONTAP 9.5, if the policy is not specified with the <code>snapmirror</code> <code>create</code> command, the <code>MirrorAndVault</code> policy is used by default.

cluster2::> snapmirror create -source-path vss:srcFG -destination-path
vsd:dstFG -type XDP -policy MirrorAllSnapshots -schedule hourly
Operation succeeded: snapmirror create for the relationship with
destination "vsd:dstFG".

5. From the destination cluster, initialize the SnapMirror relationship by performing a baseline transfer: snapmirror initialize -destination-path dest_svm:dest_flexgroup

After the baseline transfer is completed, the destination FlexGroup volume is updated periodically based on the schedule of the SnapMirror relationship.

cluster2::> snapmirror initialize -destination-path vsd:dstFG
Operation is queued: snapmirror initialize of destination "vsd:dstFG".



If you have created any SnapMirror relationship between FlexGroup volumes with the source cluster running ONTAP 9.3 and the destination cluster running ONTAP 9.2 or earlier, and if you create any qtrees in the source FlexGroup volume, the SnapMirror updates fail. To recover from this situation, you must delete all of the non-default qtrees in the FlexGroup volume, disable the qtree functionality on the FlexGroup volume, and then delete all of the snapshots that are enabled with the qtree functionality.

After you finish

You should set up the destination SVM for data access by setting up required configurations such as LIFs and export policies.

Related information

- · snapmirror create
- · snapmirror initialize
- · snapmirror policy create
- snapmirror update

Create SnapVault relationships for ONTAP FlexGroup volumes

You can configure a SnapVault relationship and assign a SnapVault policy to the relationship to create a SnapVault backup.

Before you begin

You must be aware of the considerations for creating a SnapVault relationship for FlexGroup volumes.

Steps

- 1. Create a destination FlexGroup volume of type DP that has the same number of constituents as that of the source FlexGroup volume:
 - a. From the source cluster, determine the number of constituents in the source FlexGroup volume: volume show -volume volume name* -is-constituent true

server	Volume	Aggregate	State	Type	Size
Available					
7SS	src	-	online	RW	400TB
72.86GB	56%				
7SS	src0001	Aggr_cmode	online	RW	25GB
0.86TB	56%				
7SS	src0002	aggr1	online	RW	25TB
0.86TB	56%				
7SS	src0003	Aggr_cmode	online	RW	25TB
0.72TB	57%				
7SS	src0004	aggr1	online	RW	25TB
.0.73TB	57%				
7SS	src0005	Aggr_cmode	online	RW	25TB
0.67TB	57%				
7SS	src0006	aggr1	online	RW	25TB
0.64TB	57%				
7SS	src0007	Aggr_cmode	online	RW	25TB
.0.63TB	57%				

b. From the destination cluster, create a destination FlexGroup volume of type DP with the same number of constituents as that of the source FlexGroup volume.

cluster2::> volume create -vserver vsd -aggr-list aggr1,aggr2 -aggr -list-multiplier 8 -size 400TB -type DP dst Warning: The FlexGroup volume "dst" will be created with the following number of constituents of size 25TB: 16. Do you want to continue? $\{y|n\}$: y [Job 766] Job succeeded: Successful

c. From the destination cluster, verify the number of constituents in the destination FlexGroup volume: volume show -volume volume_name* -is-constituent true

/server	Volume	Aggregate	State	Type	Size
Available	Used%				
7sd	dst	_	online	RW	400TB
172.86GB	56%				
7sd	dst0001	Aggr_cmode	online	RW	25GB
L0.86TB	56%				
7sd	dst0002	aggr1	online	RW	25TB
L0.86TB	56%				
7sd	dst0003	Aggr_cmode	online	RW	25TB
L0.72TB	57%				
7sd	dst0004	aggr1	online	RW	25TB
L0.73TB	57%				
/sd	dst0005	Aggr_cmode	online	RW	25TB
L0.67TB	57%				
/sd	dst0006	aggr1	online	RW	25TB
L0.64TB	57%				
/sd	dst0007	Aggr_cmode	online	RW	25TB
L0.63TB	57%				

2. Create a job schedule: job schedule cron create -name job_name -month month -dayofweek day of week -day day of month -hour hour -minute minute

For -month, -dayofweek, and -hour, you can specify all to run the job every month, day of the week, and hour, respectively.

The following example creates a job schedule named my weekly that runs on Saturdays at 3:00 a.m.:

```
cluster1::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

- 3. Create a SnapVault policy, and then define a rule for the SnapVault policy:
 - a. Create a custom policy of type vault for the SnapVault relationship: snapmirror policy create -vserver svm name -policy policy name -type vault
 - b. Define a rule for the SnapVault policy that determines which snapshots are transferred during initialization and update operations: snapmirror policy add-rule -vserver svm_name -policy policy_for_rule snapmirror-label snapmirror-label -keep retention count -schedule schedule

If you do not create a custom policy, you should specify the XDPDefault policy for SnapVault relationships.

4. Create a SnapVault relationship: snapmirror create -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -type XDP -schedule schedule_name -policy XDPDefault

In ONTAP 9.4 and earlier, if the policy is not specified with the snapmirror create command, the MirrorAllSnapshots policy is used by default. In ONTAP 9.5, if the policy is not specified with the snapmirror create command, the MirrorAndVault policy is used by default.

```
cluster2::> snapmirror create -source-path vss:srcFG -destination-path
vsd:dstFG -type XDP -schedule Daily -policy XDPDefault
```

Learn more about snapmirror create in the ONTAP command reference.

5. From the destination cluster, initialize the SnapVault relationship by performing a baseline transfer: snapmirror initialize -destination-path dest svm:dest flexgroup

```
cluster2::> snapmirror initialize -destination-path vsd:dst
Operation is queued: snapmirror initialize of destination "vsd:dst".
```

Related information

- · snapmirror create
- · snapmirror initialize
- snapmirror policy add-rule
- snapmirror policy create

Create unified data protection relationships for ONTAP FlexGroup volumes

Beginning with ONTAP 9.3, you can create and configure SnapMirror unified data protection relationships to configure disaster recovery and archiving on the same destination volume.

Before you begin

You must be aware of the considerations for creating unified data protection relationships for FlexGroup volumes.

Considerations for creating a SnapVault backup relationship and a unified data protection relationship for FlexGroup volumes

Steps

- 1. Create a destination FlexGroup volume of type DP that has the same number of constituents as that of the source FlexGroup volume:
 - a. From the source cluster, determine the number of constituents in the source FlexGroup volume: volume show -volume volume name* -is-constituent true

cluster1:	:> volume show	-volume srcF	G* -is-cons	stituent tru	ıe
Vserver	Volume	Aggregate	State	Туре	Size
Available	Used%				
VSS		-	online	RW	400TB
172.86GB					
	srcFG0001	Aggr_cmode	online	RW	25GB
10.86TB					
	srcFG0002	aggr1	online	RW	25TB
10.86TB					
	srcFG0003	Aggr_cmode	online	RW	25TB
10.72TB					
	srcFG0004	aggr1	online	RW	25TB
10.73TB					
	srcFG0005	Aggr_cmode	online	RW	25TB
10.67TB					
	srcFG0006	aggr1	online	RW	25TB
10.64TB					
	srcFG0007	Aggr_cmode	online	RW	25TB
10.63TB	57%				

b. From the destination cluster, create a destination FlexGroup volume of type DP with the same number of constituents as that of the source FlexGroup volume.

```
cluster2::> volume create -vserver vsd -aggr-list aggr1,aggr2 -aggr
-list-multiplier 8 -size 400TB -type DP dstFG

Warning: The FlexGroup volume "dstFG" will be created with the following number of constituents of size 25TB: 16.
Do you want to continue? {y|n}: y
[Job 766] Job succeeded: Successful
```

c. From the destination cluster, verify the number of constituents in the destination FlexGroup volume: volume show -volume volume name* -is-constituent true

Veckion	770]1170	7 caroasto	2+2+0	Time	Size
vserver Available	Volume	Aggregate	State	туре	Size
Avallable					
vsd	dstFG	-	online	RW	400TB
172.86GB	56%				
vsd	dstFG0001	Aggr_cmode	online	RW	25GB
10.86TB	56%				
vsd	dstFG0002	aggr1	online	RW	25TB
10.86TB	56%				
vsd	dstFG0003	Aggr_cmode	online	RW	25TB
10.72TB	57%				
vsd	dstFG0004	aggr1	online	RW	25TB
10.73TB	57%				
vsd	dstFG0005	Aggr_cmode	online	RW	25TB
10.67TB	57%				
vsd	dstFG0006	aggr1	online	RW	25TB
10.64TB	57%				
vsd	dstFG0007	Aggr_cmode	online	RW	25TB
10.63TB	57%				

2. Create a job schedule: job schedule cron create -name job_name -month month -dayofweek day of week -day day of month -hour hour -minute minute

For the -month, -dayofweek, and -hour options, you can specify all to run the job every month, every day of the week, and every hour, respectively.

The following example creates a job schedule named my weekly that runs on Saturdays at 3:00 a.m.:

```
cluster1::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

- 3. Create a custom policy of type mirror-vault, and then define a rule for the mirror and vault policy:
 - a. Create a custom policy of type mirror-vault for the unified data protection relationship: snapmirror policy create -vserver svm_name -policy policy_name -type mirror-vault
 - b. Define a rule for the mirror and vault policy that determines which snapshots are transferred during initialization and update operations: snapmirror policy add-rule -vserver svm_name -policy policy_for_rule snapmirror-label snapmirror-label -keep retention_count -schedule schedule

If you do not specify a custom policy, the MirrorAndVault policy is used for unified data protection relationships.

4. Create a unified data protection relationship: snapmirror create -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -type XDP -schedule schedule name -policy MirrorAndVault

In ONTAP 9.4 and earlier, if the policy is not specified with the <code>snapmirror</code> <code>create</code> command, the <code>MirrorAllSnapshots</code> policy is used by default. In ONTAP 9.5, if the policy is not specified with the <code>snapmirror</code> <code>create</code> command, the <code>MirrorAndVault</code> policy is used by default.

```
cluster2::> snapmirror create -source-path vss:srcFG -destination-path
vsd:dstFG -type XDP -schedule Daily -policy MirrorAndVault
```

Learn more about snapmirror create in the ONTAP command reference.

5. From the destination cluster, initialize the unified data protection relationship by performing a baseline transfer: snapmirror initialize -destination-path dest svm:dest flexgroup

```
cluster2::> snapmirror initialize -destination-path vsd:dstFG
Operation is queued: snapmirror initialize of destination "vsd:dstFG".
```

Related information

- · snapmirror create
- · snapmirror initialize
- · snapmirror policy add-rule
- snapmirror policy create

Create SVM disaster recovery relationships for ONTAP FlexGroup volumes

Beginning with ONTAP 9.9.1, you can create SVM disaster recovery (SVM DR) relationships using FlexGroup volumes. An SVM DR relationship provides redundancy and the ability to recover FlexGroup volumes in the event of a disaster by synchronizing and replicating the SVM configuration and its data. A SnapMirror license is required for SVM DR.

Before you begin

You cannot create a FlexGroup SVM DR relationship with the following applies.

- A FlexClone FlexGroup configuration exists
- The FlexGroup volume is part of a cascading relationship
- The FlexGroup volume is part of a fanout relationship, and your cluster is running an ONTAP version earlier than ONTAP 9.12.1. (Beginning with ONTAP 9.13.1, fanout relationships are supported.)

About this task

• All nodes in both clusters must be running the same ONTAP version as the node on which SVM DR support was added (ONTAP 9.9.1 or later).

- The SVM DR relationship between the primary and secondary sites should be healthy and should have enough space on both the primary and secondary SVMs to support the FlexGroup volumes.
- Beginning with ONTAP 9.12.1, FabricPool, FlexGroup, and SVM DR can work in conjunction. In releases earlier than ONTAP 9.12.1, any two of these features worked together, but not all three in conjunction.
- When you create a FlexGroup SVM DR relationship in which the FlexGroup volume is part of a fanout relationship, you should be aware of the following requirements:
 - The source and destination cluster must be running ONTAP 9.13.1 or later.
 - SVM DR with FlexGroup volumes supports SnapMirror fanout relationships to eight sites.

For information about creating an SVM DR relationship, see Manage SnapMirror SVM replication.

Steps

1. Create an SVM DR relationship, or use an existing relationship.

Replicate an entire SVM configuration

2. Create a FlexGroup volume on the primary site with the required number of constituents.

Creating a FlexGroup volume.

Wait until FlexGroup and all of its constituents are created before proceeding.

3. To replicate the FlexGroup volume, update the SVM at the secondary site: snapmirror update -destination-path destination svm name: -source-path source svm name:

You can also check if a scheduled SnapMirror update already exists by entering snapmirror show -fields schedule

4. From the secondary site, verify that the SnapMirror relationship is healthy: snapmirror show

5. From the secondary site, verify that the new FlexGroup volume and its constituents exist: snapmirror show -expand

cluster2::> snapmirror show -expand							
Progress Source Last		Destination	Mirror	Relationship	Total		
Path	Туре	Path	State	Status	Progress	Healthy	
Updated							
vs1:	XDP	vs1_dst:	Snapmir				
wal·fa ara	VDD	wel detifa	ıra	Idle	_	true	-
vsi.ig_sic	ADP	vs1_dst:fg_s	Snapmir	rored			
			-	Idle	_	true	-
vs1:fg_src_	_						
	XDP	vs1_dst:fg_s	rc0001 Snapmir				
			BIIGPIIII	Idle	_	true	-
vs1:fg_src_	_						
	XDP	vs1_dst:fg_s					
			Snapmir	rorea Idle	_	true	_
vs1:fg_src_	_0003			10:10		0200	
	XDP	vs1_dst:fg_s					
			Snapmir	rored Idle		+ 2110	
vsl:fg src	0004			idie	_	true	-
3	_	vs1_dst:fg_s	rc0004				
			Snapmir				
6 entries w	ere d	isplayed		Idle	_	true	-
6 entries were displayed.							

Related information

- snapmirror show
- snapmirror update

Transition ONTAP FlexGroup SnapMirror relationships to SVM DR

You can create a FlexGroup SVM DR relationship by transitioning an existing FlexGroup volume SnapMirror relationship.

Before you begin

- The FlexGroup volume SnapMirror relationship is in a healthy state.
- The source and destination FlexGroup volumes have the same name.

Steps

- 1. From the SnapMirror destination, resynchronize the FlexGroup level SnapMirror relationship: snapmirror resync
- 2. Create the FlexGroup SVM DR SnapMirror relationship. Use the same SnapMirror policy which is configured on the FlexGroup volume SnapMirror relationships: snapmirror create -destination -path dest_svm: -source-path src_svm: -identity-preserve true -policy MirrorAllSnapshots



You must use the -identity-preserve true option of the snapmirror create command when you create your replication relationship.

Learn more about snapmirror create in the ONTAP command reference.

3. Verify the relationship is broken off: snapmirror show -destination-path dest_svm: -source -path src svm:

4. Stop the destination SVM: vserver stop -vserver vs name

```
vserver stop -vserver fg_vs_renamed
[Job 245] Job is queued: Vserver Stop fg_vs_renamed.
[Job 245] Done
```

5. Resynchronize the SVM SnapMirror relationship: snapmirror resync -destination-path dest_svm: -source-path src_svm:

```
snapmirror resync -destination-path fg_vs_renamed: -source-path fg_vs: Warning: This Vserver has volumes which are the destination of FlexVol or FlexGroup SnapMirror relationships. A resync on the Vserver SnapMirror relationship will cause disruptions in data access
```

6. Verify that the SVM DR level SnapMirror relationship reaches a healthy idle state: snapmirror show

-expand

7. Verify that the FlexGroup SnapMirror relationship is in a healthy state: snapmirror show

Related information

- · snapmirror create
- snapmirror resync
- snapmirror show

Convert ONTAP FlexVol volumes to FlexGroup volumes within an SVM-DR relationship

Beginning with ONTAP 9.10.1, you can convert a FlexVol volume to a FlexGroup volume on an SVM-DR source.

Before you begin

- The FlexVol volume that is being converted must be online.
- The operations and configurations on the FlexVol volume must be compatible with the conversion process.

An error message is generated if the FlexVol volume has any incompatibility, and the volume conversion is cancelled. You can take corrective actions and retry the conversion. For more details, see Considerations for converting FlexVol volumes to FlexGroup volumes

Steps

- 1. Login using advance privilege mode: set -privilege advanced
- 2. From the destination, update the SVM-DR relationship:

```
snapmirror update -destination-path <destination_svm_name>: -source-path
<source_svm_name>:
```



You must enter a colon (:) after the SVM name in the -destination-path option.

3. Ensure that the SVM-DR relationship is in a SnapMirrored state and is not broken-off:

```
snapmirror show
```

4. From the destination SVM, verify that the FlexVol volume is ready for conversion:

```
volume conversion start -vserver <svm_name> -volume <vol_name> -check
-only true
```

If this command generates any errors other than "This is a destination SVMDR volume," you can take the appropriate corrective action, run the command again, and continue the conversion.

5. From the destination, disable transfers on the SVM-DR relationship:

```
snapmirror quiesce -destination-path <dest_svm>:
```



You must enter a colon (:) after the SVM name in the -destination-path option.

6. From the source cluster, start the conversion:

```
volume conversion start -vserver <svm_name> -volume <vol_name>
```

7. Verify that the conversion is successful:

```
volume show <vol_name> -fields volume-style-extended,state
```

```
cluster-1::*> volume show my_volume -fields volume-style-extended, state

vserver volume state volume-style-extended
------
vs0 my_volume online flexgroup
```

8. From the destination cluster, resume transfers for the relationship:

```
snapmirror resume -destination-path <dest_svm>:
```



You must enter a colon (:) after the SVM name in the -destination-path option.

9. From the destination cluster, perform an update to propagate the conversion to the destination:

```
snapmirror update -destination-path <dest_svm>:
```



You must enter a colon (:) after the SVM name in the <code>-destination-path</code> option.

10. Ensure that the SVM-DR relationship is in a SnapMirrored state and is not broken off:

```
snapmirror show
```

11. Ensure the conversion occurred on the destination:

```
volume show <vol_name> -fields volume-style-extended,state
```

```
cluster-2::*> volume show my_volume -fields volume-style-extended, state

vserver volume state volume-style-extended
-----
vs0_dst my_volume online flexgroup
```

Related information

- snapmirror resume
- · snapmirror quiesce
- snapmirror show
- snapmirror update

Considerations for creating SnapMirror cascade and fanout relationships for ONTAP FlexGroup volumes

There are support considerations and limitations you should keep in mind when creating SnapMirror cascade and fanout relationships for FlexGroup volumes.

Considerations for creating cascading relationships

- Each relationship can be either an inter cluster or intra cluster relationship.
- All asynchronous policy types, including async-mirror, mirror-vault, and vault, are supported for both relationships.
- · Only "MirrorAllSnapshots," not "MirrorLatest" async-mirror policies are supported.
- · Long-term retention snapshots are not supported.

Learn more about long-term retention snapshots.

- Concurrent updates of cascaded XDP relationships is supported.
- Supports removing A to B and B to C and resync A to C or resync C to A.
- A and B FlexGroup volumes also support fanout when all nodes are running ONTAP 9.9.1 or later.
- Restore operations from B or C FlexGroup volumes are supported.
- Transfers on FlexGroup relationships are not support while the destination is the source of a restore relationship.
- The destination of a FlexGroup restore cannot be the destination of any other FlexGroup relationship.
- FlexGroup file restore operations have the same restrictions as regular FlexGroup restore operations.
- All nodes in the cluster where the B and C FlexGroup volumes reside must be running ONTAP 9.9.1 or later.
- · All expand and auto expand functionality is supported.

- In a cascade configuration such as A to B to C, if A to B and B to C have different numbers of constituent SnapMirror relationships, then an abort operation from the source is not supported for the B to C SnapMirror relationship.
- System Manager does not support cascading relationships regardless of the ONTAP version.
- When converting an A to B to C set of FlexVol relationship to a FlexGroup relationship, you must convert the B to C hop first.
- All FlexGroup cascade configurations for relationships with policy types supported by REST are also supported by REST APIs in cascading FlexGroup configurations.
- As with FlexVol relationships, FlexGroup cascading is not supported by the snapmirror protect command.

Considerations for creating fanout relationships

- Two or more FlexGroup fanout relationships are supported; for example, A to B, A to C, with a maximum of 8 fanout legs.
- Each relationship can be either intercluster or intracluster.
- · Concurrent updates are supported for the two relationships.
- All expand and auto expand functionality is supported.
- If the fanout legs of the relationship have different numbers of constituent SnapMirror relationships, then an abort operation from the source is not supported for the A to B and A to C relationships.
- All nodes in the cluster where the source and destination FlexGroup volumes reside must be running ONTAP 9.9.1 or later.
- All asynchronous policy types currently supported for FlexGroup SnapMirror are supported in fanout relationships.
- You can perform restore operations from B to C FlexGroup volumes.
- All fanout configurations with policy types supported by rest are also supported for REST APIs in FlexGroup fanout configurations.

Related information

snapmirror protect

Considerations for creating SnapVault backup relationships and unified data protection relationships for ONTAP FlexGroup volumes

You must be aware of the considerations for creating a SnapVault backup relationship and unified data protection relationship for FlexGroup volumes.

- You can resynchronize a SnapVault backup relationship and a unified data protection relationship by using the -preserve option that enables you to preserve snapshots on the destination volume that are newer than the latest common snapshot.
- Long-term retention is not supported with FlexGroup volumes.

Long-term retention enables creating snapshots directly on the destination volume without requiring to store the snapshots on the source volume.

The snapshot command expiry-time option is not supported for FlexGroup volumes.

- Storage efficiency cannot be configured on the destination FlexGroup volume of a SnapVault backup relationship and unified data protection relationship.
- You cannot rename snapshots of a SnapVault backup relationship and unified data protection relationship for FlexGroup volumes.
- A FlexGroup volume can be the source volume of only one backup relationship or restore relationship.

A FlexGroup volume cannot be the source of two SnapVault relationships, two restore relationships, or a SnapVault backup relationship and a restore relationship.

• If you delete a snapshot on the source FlexGroup volume and re-create a snapshot with the same name, the next update transfer to the destination FlexGroup volume fails if the destination volume has a snapshot of the same name.

This is because snapshots cannot be renamed for FlexGroup volumes.

Monitor SnapMirror data transfers for ONTAP FlexGroup volumes

You should periodically monitor the status of the FlexGroup volume SnapMirror relationships to verify that the destination FlexGroup volume is updated periodically as per the specified schedule.

About this task

You must perform this task from the destination cluster.

Steps

1. View the SnapMirror relationship status of all FlexGroup volume relationships: snapmirror show -relationship-group-type flexgroup

cluster2::>	· snapı	mirror show -	relation	ship-group-type	flexgroup		
Progress							
Source		Destination	Mirror	Relationship	Total		
Last							
Path	Type	Path	State	Status	Progress	Healthy	
Updated							
vss:s	XDP	vsd:d	Snapmir	rored			
				Idle	-	true	-
vss:s2	XDP	vsd:d2	Uniniti	alized			
				Idle	-	true	-
2 entries w	ere d	isplayed.					

Related information

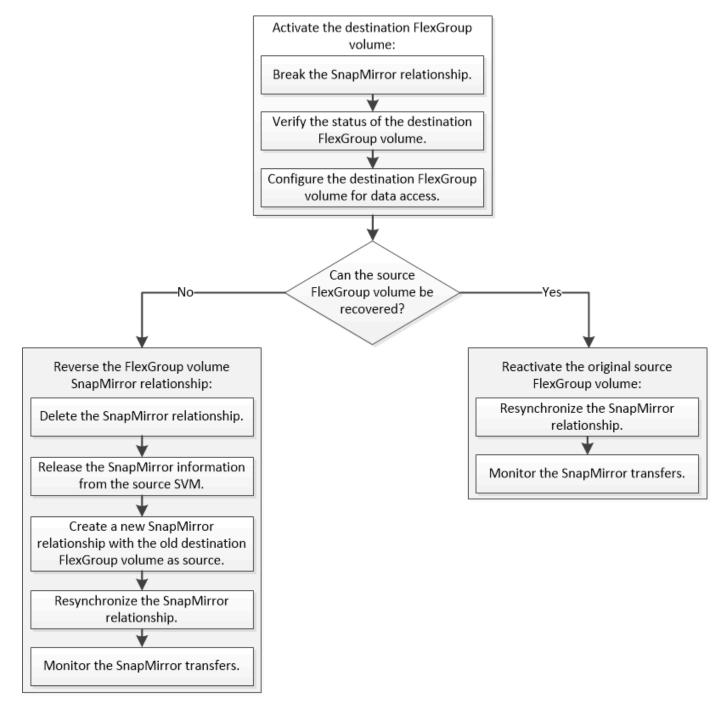
snapmirror show

Manage data protection operations for FlexGroup volumes

Disaster recovery for FlexGroup volumes

Disaster recovery workflow for ONTAP FlexGroup volumes

When a disaster strikes on the source FlexGroup volume, you should activate the destination FlexGroup volume and redirect client access. Depending on whether the source FlexGroup volume can be recovered, you should either reactivate the source FlexGroup volume or reverse the SnapMirror relationship.



About this task

Client access to the destination FlexGroup volume is blocked for a brief period when some SnapMirror operations, such as SnapMirror break and resynchronization, are running. If the SnapMirror operation fails, it is possible that some of the constituents remain in this state and access to the FlexGroup volume is denied. In such cases, you must retry the SnapMirror operation.

Activate the destination ONTAP FlexGroup volume

When the source FlexGroup volume is unable to serve data due to events such as data corruption, accidental deletion or an offline state, you must activate the destination FlexGroup volume to provide data access until you recover the data on the source FlexGroup volume. Activation involves stopping future SnapMirror data transfers and breaking the SnapMirror relationship.

About this task

You must perform this task from the destination cluster.

Steps

1. Disable future transfers for the FlexGroup volume SnapMirror relationship: snapmirror quiesce dest svm:dest flexgroup

```
cluster2::> snapmirror quiesce -destination-path vsd:dst
```

Break the FlexGroup volume SnapMirror relationship: snapmirror break dest svm:dest flexgroup

```
cluster2::> snapmirror break -destination-path vsd:dst
```

3. View the status of the SnapMirror relationship: snapmirror show -expand

cluster2::> snapmirror show -expand									
Progress Source Last		Destination N	Mirror	Relationship	Total				
		Path S		Status	Progress	Healthy			
vss:s	XDP	vsd:dst	Broke	n-off					
				Idle	-	true	-		
vss:s0001	XDP	vsd:dst0001	Broke						
vss.s 0002	XDD	vsd:dst 0002	Brokei	Idle	_	true	_		
V33:30002	ZDI	v5a.a5c0002	DIORCI	Idle	_	true	_		
vss:s0003	XDP	vsd:dst0003	Broke	n-off					
				Idle	-	true	-		
vss:s0004	XDP	vsd:dst0004	Broke						
weete 0005	AUD	vsd:dst 0005	Broke	Idle	_	true	-		
v33 . 30003	ADI	vsa:asc0005	DIORCI	Idle	_	true	_		
vss:s0006	XDP	vsd:dst0006	Broke	n-off					
				Idle	-	true	-		
vss:s0007	XDP	vsd:dst0007	Broke						
0000	VDD		Desplay	Idle	-	true	-		
vss:s0008	XDP	vsd:dst0008	Broke	n-oii Idle	_	true	_		
				1410		CLUC			

The SnapMirror relationship status of each constituent is Broken-off.

^{4.} Verify that the destination FlexGroup volume is read/write: volume show -vserver svm_name

	::> volume sho Volume e Used%	ow -vserver vs Aggregate		Type	Size
vsd	dst	-	online	**RW**	2GB
1.54GB	22%				
vsd	d2	_	online	DP	2GB
1.55GB	22%				
vsd	root vs0	aggr1	online	RW	100MB
94.02MB	_ 5%				
3 entries	s were display	red.			

5. Redirect clients to the destination FlexGroup volume.

Related information

- snapmirror break
- snapmirror quiesce
- · snapmirror show

Reactivate the original source ONTAP FlexGroup volume after disaster

When the source FlexGroup volume becomes available, you can resynchronize the original source and original destination FlexGroup volumes. Any new data on the destination FlexGroup volume is lost.

About this task

Any active quota rules on the destination volume are deactivated and the quota rules are deleted before resynchronization is performed.

You can use the volume quota policy rule create and volume quota modify commands to create and reactivate quota rules after the resynchronization operation is complete.

Steps

- 1. From the destination cluster, resynchronize the FlexGroup volume SnapMirror relationship: snapmirror resync -destination-path dst_svm:dest_flexgroup
- 2. View the status of the SnapMirror relationship: snapmirror show -expand

cluster2::> snapmirror show -expand								
Progress Source Last		Destination N	lirror	Relationship	Total			
Path Updated	Туре	Path S	State	Status	Progress	Healthy		
vss:s	XDP	vsd:dst	Snapm	irrored Idle	_	true	_	
vss:s0001	XDP	vsd:dst0001	Snapm	irrored				
vss:s 0002	XDP	vsd:dst 0002	Snapm	Idle irrored	_	true	_	
	1121	.54.4560002	onapiii	Idle	-	true	_	
vss:s0003	XDP	vsd:dst0003	Snapm	irrored				
vss:s 0004	XDP	vsd:dst 0004	Snapm	Idle irrored	_	true	_	
_			-	Idle	-	true	-	
vss:s0005	XDP	vsd:dst0005	Snapm	irrored Idle				
vss:s 0006	XDP	vsd:dst 0006	Snapm	irrored	_	true	_	
			-	Idle	-	true	-	
vss:s0007	XDP	vsd:dst0007	Snapm	irrored Idle		+ 2010		
vss:s 0008	XDP	vsd:dst 0008	Snapm	irrored		true	_	
_		<u>—</u>	-	Idle	-	true	-	

The SnapMirror relationship status of each constituent is Snapmirrored.

Related information

- snapmirror resync
- · snapmirror show

Reverse SnapMirror relationships between ONTAP FlexGroup volumes during disaster recovery

When a disaster disables the source FlexGroup volume of a SnapMirror relationship, you can use the destination FlexGroup volume to serve data while you repair or replace the source FlexGroup volume. After the source FlexGroup volume is online, you can make the original source FlexGroup volume a read-only destination and reverse the SnapMirror relationship.

About this task

Any active quota rules on the destination volume are deactivated and the quota rules are deleted before

resynchronization is performed.

You can use the volume quota policy rule create and volume quota modify commands to create and reactivate quota rules after the resynchronization operation is complete.

Steps

1. On the original destination FlexGroup volume, remove the data protection mirror relationship between the source FlexGroup volume and the destination FlexGroup volume: snapmirror delete -destination -path svm name:volume name

```
cluster2::> snapmirror delete -destination-path vsd:dst
```

2. On the original source FlexGroup volume, remove the relationship information from the source FlexGroup volume: snapmirror release -destination-path svm_name:volume_name -relationship -info-only

After deleting a SnapMirror relationship, you must remove the relationship information from the source FlexGroup volume before attempting a resynchronization operation.

```
cluster1::> snapmirror release -destination-path vsd:dst -relationship
-info-only true
```

3. On the new destination FlexGroup volume, create the mirror relationship: snapmirror create -source-path src_svm_name:volume_name -destination-path dst svm name:volume name -type XDP -policy MirrorAllSnapshots

```
cluster1::> snapmirror create -source-path vsd:dst -destination-path
vss:src -type XDP -policy MirrorAllSnapshots
```

Learn more about snapmirror create in the ONTAP command reference.

4. On the new destination FlexGroup volume, resynchronize the source FlexGroup: snapmirror resync -source-path svm name:volume name

```
cluster1::> snapmirror resync -source-path vsd:dst
```

5. Monitor the SnapMirror transfers: snapmirror show -expand

Progress							
Source		Destination N	Mirror	Relationshi	p Total		
Last							
Path	Type	Path	State	Status	Progress	Healthy	
Updated							
vsd:dst	VDD	vss:src	Cn	apmirrored			
vsa.ast	ADE	V55.51C	511	Idle	_	true	
vss·dst 00	001 XDP	vss:src 000	01 Sn	apmirrored		crue	
	701 1121	**************************************	01 011	Idle	_	true	
vsd:dst 00	002 XDP	vss:src 000	02 Sn	apmirrored			
				Idle	_	true	
vsd:dst0	003 XDP	vss:src000	03 Sn	apmirrored			
				Idle	_	true	
vsd:dst0	004 XDP	vss:src000	04 Sn	apmirrored			
				Idle	_	true	
vsd:dst0	005 XDP	vss:src000	05 Sn	apmirrored			
				Idle	_	true	
vsd:dst0	006 XDP	vss:src000	06 Sn	apmirrored			
1 1	007 1100	0.04	07 0	Idle	_	true	
vsa:ast0	JU/ XDP	vss:src000	J/ Sn	apmirrored Idle		+ 2017.0	
77ed·de+ 01	1U8 AUD	vss:src 000	18 Sn	apmirrored	_	true	
v.s.u.s.t0(JUU ADP	vss.sic000	00 511	Idle		true	

The SnapMirror relationship status of each constituent shows as Snapmirrored that indicates that the resynchronization was successful.

Related information

- snapmirror create
- snapmirror delete
- snapmirror release
- snapmirror resync
- snapmirror show

Expand FlexGroup volumes in a SnapMirror relationship

Expand ONTAP FlexGroup volumes in a SnapMirror relationship

Beginning with ONTAP 9.3, you can expand the source FlexGroup volume and

destination FlexGroup volume that are in a SnapMirror relationship by adding new constituents to the volumes. You can expand the destination volumes either manually or automatically.

About this task

• After expansion, the number of constituents in the source FlexGroup volume and destination FlexGroup volume of a SnapMirror relationship must match.

If the number of constituents in the volumes does not match, the SnapMirror transfers fail.

- You should not perform any SnapMirror operation when the expansion process is in progress.
- If a disaster strikes before the expansion process is complete, you must break the SnapMirror relationship and wait until the operation succeeds.



You should break the SnapMirror relationship when the expansion process is in progress only in the case of a disaster. In the case of a disaster, the break operation can take some time to complete. You should wait for the break operation to get completed successfully before performing a resync operation. If the break operation fails, you must retry the break operation. If the break operation fails, some of the new constituents might remain in the destination FlexGroup volume after the break operation. It is best to delete these constituents manually before proceeding further.

Expand the source ONTAP FlexGroup volume of a SnapMirror relationship

Beginning with ONTAP 9.3, you can expand the source FlexGroup volume of a SnapMirror relationship by adding new constituents to the source volume. You can expand the source volume in the same way that you expand a regular FlexGroup volume (read-write volume).

Steps

1. Expand the source FlexGroup volume: volume expand -vserver vserver_name -volume fg_src -aggr-list aggregate name,... [-aggr-list-multiplier constituents_per_aggr]

```
cluster1::> volume expand -volume src_fg -aggr-list aggr1 -aggr-list -multiplier 2 -vserver vs_src

Warning: The following number of constituents of size 50GB will be added to FlexGroup "src_fg": 2.

Expanding the FlexGroup will cause the state of all Snapshot copies to be set to "partial".

Partial Snapshot copies cannot be restored.

Do you want to continue? {y|n}: Y

[Job 146] Job succeeded: Successful
```

The state of all of the snapshots that are taken before the volume is expanded changes to partial.

Expand the destination ONTAP FlexGroup volume of a SnapMirror relationship

You can expand the destination FlexGroup volume and reestablish the SnapMirror relationship either automatically or manually. By default, the SnapMirror relationship is set for automatic expansion, and the destination FlexGroup volume expands automatically if the source volume expands.

Before you begin

- The source FlexGroup volume must have been expanded.
- The SnapMirror relationship must be in the SnapMirrored state.

The SnapMirror relationship must not be broken or deleted.

About this task

• When the destination FlexGroup volume is created, the volume is set up for automatic expansion by default.

You can modify the destination FlexGroup volume for manual expansion, if required.



The best practice is to expand the destination FlexGroup volume automatically.

- All SnapMirror operations fail until both the source FlexGroup volume and destination FlexGroup volume have expanded and have the same number of constituents.
- If you expand the destination FlexGroup volume after the SnapMirror relationship is broken or deleted, you cannot resync the original relationship again.

If you intend to reuse the destination FlexGroup volume, do not expand the volume after deleting the SnapMirror relationship.

Choices

- Perform an update transfer to expand the destination FlexGroup volume automatically:
 - a. Perform a SnapMirror update transfer: snapmirror update -destination-path svm:vol name
 - b. Verify that the status of the SnapMirror relationship is in the SnapMirrored state: snapmirror show

Based on the size and availability of aggregates, the aggregates are automatically selected, and new constituents that match the constituents of the source FlexGroup volume are added to the destination FlexGroup volume. After expansion, a resynchronization operation is automatically triggered.

- Expand the destination FlexGroup volume manually:
 - a. If the SnapMirror relationship is in the auto-expand mode, set the SnapMirror relationship to the manual expand mode: snapmirror modify -destination-path svm:vol_name -is-auto-expand -enabled false

```
cluster2::> snapmirror modify -destination-path vs_dst:dst_fg -is
-auto-expand-enabled false
Operation succeeded: snapmirror modify for the relationship with
destination "vs_dst:dst_fg".
```

b. Quiesce the SnapMirror relationship: snapmirror quiesce -destination-path svm:vol name

```
cluster2::> snapmirror quiesce -destination-path vs_dst:dst_fg
Operation succeeded: snapmirror quiesce for destination
"vs_dst:dst_fg".
```

c. Expand the destination FlexGroup volume: volume expand -vserver vserver_name -volume fg_name -aggr-list aggregate name,... [-aggr-list-multiplier constituents per aggr]

```
cluster2::> volume expand -volume dst_fg -aggr-list aggr1 -aggr-list
-multiplier 2 -vserver vs_dst

Warning: The following number of constituents of size 50GB will be
added to FlexGroup "dst_fg": 2.
Do you want to continue? {y|n}: y
[Job 68] Job succeeded: Successful
```

d. Resynchronize the SnapMirror relationship: snapmirror resync -destination-path svm:vol name

```
cluster2::> snapmirror resync -destination-path vs_dst:dst_fg
Operation is queued: snapmirror resync to destination
"vs_dst:dst_fg".
```

e. Verify that the status of the SnapMirror relationship is SnapMirrored: snapmirror show

Related information

- snapmirror quiesce
- · snapmirror resync
- snapmirror show

Perform a SnapMirror single file restore from an ONTAP FlexGroup volume

Beginning with ONTAP 9.8, you can restore a single file from a FlexGroup SnapMirror vault or from a UDP destination.

About this task

- You can restore from a FlexGroup volume of any geometry to FlexGroup volume of any geometry.
- Only one file per restore operation is supported.
- You can restore to either the original source FlexGroup volume or to a new FlexGroup volume.
- Remote fenced file lookup is not supported.

Single file restore fails if the source file is fenced.

- You can restart or clean up an aborted single file restore.
- You should clean up a failed single file restore transfer by using the clean-up-failure option of the snapmirror restore command.

Learn more about snapmirror restore in the ONTAP command reference.

• Expansion of FlexGroup volumes is supported when a FlexGroup single file restore is in progress or in an aborted state.

Steps

 Restore a file from a FlexGroup volume: snapmirror restore -destination-path destination_path -source-path source_path -file-list /f1 -throttle throttle -source-snapshot snapshot

The following is an example of a FlexGroup volume single file restore operation.

```
vserverA::> snapmirror restore -destination-path vs0:fg2 -source-path
vs0:fgd -file-list /f1 -throttle 5 -source-snapshot snapmirror.81072ce1-
d57b-11e9-94c0-005056a7e422 2159190496.2019-09-19 062631
[Job 135] Job is queued: snapmirror restore from source "vs0:fqd" for
the snapshot snapmirror.81072ce1-d57b-11e9-94c0-
005056a7e422 2159190496.2019-09-19 062631.
vserverA::> snapmirror show
               Destination Mirror Relationship
Source
Total Last
Path Type Path State Status Progress
Healthy Updated
_____
vs0:vld RST vs0:v2 - Transferring Idle 83.12KB
true 09/19 11:38:42
vserverA::*> snapmirror show vs0:fg2
Source Path: vs0:fqd
Source Cluster: -
Source Vserver: vs0
Source Volume: fgd
Destination Path: vs0:fg2
```

```
Destination Cluster: -
Destination Vserver: vs0
Destination Volume: fg2
Relationship Type: RST
Relationship Group Type: none
Managing Vserver: vs0
SnapMirror Schedule: -
SnapMirror Policy Type: -
SnapMirror Policy: -
Tries Limit: -
Throttle (KB/sec): unlimited
Current Transfer Throttle (KB/sec): 2
Mirror State: -
Relationship Status: Transferring
File Restore File Count: 1
File Restore File List: f1
Transfer Snapshot: snapmirror.81072ce1-d57b-11e9-94c0-
005056a7e422 2159190496.2019-09-19 062631
Snapshot Progress: 2.87MB
Total Progress: 2.87MB
Network Compression Ratio: 1:1
Snapshot Checkpoint: 2.97KB
Newest Snapshot: -
Newest Snapshot Timestamp: -
Exported Snapshot: -
Exported Snapshot Timestamp: -
Healthy: true
Physical Replica: -
Relationship ID: e6081667-dacb-11e9-94c0-005056a7e422
Source Vserver UUID: 81072ce1-d57b-11e9-94c0-005056a7e422
Destination Vserver UUID: 81072ce1-d57b-11e9-94c0-005056a7e422
Current Operation ID: 138f12e6-dacc-11e9-94c0-005056a7e422
Transfer Type: cg file restore
Transfer Error: -
Last Transfer Type: -
Last Transfer Error: -
Last Transfer Error Codes: -
Last Transfer Size: -
Last Transfer Network Compression Ratio: -
Last Transfer Duration: -
Last Transfer From: -
Last Transfer End Timestamp: -
Unhealthy Reason: -
Progress Last Updated: 09/19 07:07:36
Relationship Capability: 8.2 and above
Lag Time: -
```

```
Current Transfer Priority: normal
SMTape Operation: -
Constituent Relationship: false
Destination Volume Node Name: vserverA
Identity Preserve Vserver DR: -
Number of Successful Updates: 0
Number of Failed Updates: 0
Number of Successful Resyncs: 0
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
Total Transfer Bytes: 0
Total Transfer Time in Seconds: 0
Source Volume MSIDs Preserved: -
OpMask: fffffffffffffff
Is Auto Expand Enabled: -
Source Endpoint UUID: -
Destination Endpoint UUID: -
Is Catalog Enabled: false
```

Related information

· snapmirror show

Restore ONTAP FlexGroup volumes from SnapVault backups

You can perform a full-volume restore operation of FlexGroup volumes from a snapshot in the SnapVault secondary volume. You can restore the FlexGroup volume either to the original source volume or to a new FlexGroup volume.

Before you begin

You must be aware of certain considerations when you restore from SnapVault backups for FlexGroup volumes.

- Only baseline restore is supported with partial snapshots from a SnapVault backup.

 The number of constituents in the destination volume must match the number of constituents in the source volume when the snapshot was taken.
- If a restore operation fails, no other operations are allowed until the restore operation is complete. You can either retry the restore operation or run the restore operation with the cleanup parameter.
- A FlexGroup volume can be the source volume of only one backup relationship or restore relationship.
 A FlexGroup volume cannot be the source of two SnapVault relationships, two restore relationships, or a SnapVault relationship and a restore relationship.
- SnapVault backup and restore operations cannot run in parallel.
 When either a baseline restore operation or an incremental restore operation is in progress, you should quiesce the backup operations.
- You must abort a restore operation of a partial snapshot from the destination FlexGroup volume. You cannot abort the restore operation of a partial snapshot from the source volume.

• If you abort a restore operation, you must restart the restore operation with the same snapshot that was used for the previous restore operation.

About this task

Any active quota rules on the destination FlexGroup volume are deactivated before the restore is performed.

You can use the volume quota modify command to reactivate quota rules after the restore operation is complete.

Steps

1. Restore the FlexGroup volume: snapmirror restore -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -snapshot snapshot_name snapshot_name is the snapshot that is to be restored from the source volume to the destination volume. If the snapshot is not specified, the destination volume is restored from the latest snapshot.

```
vserverA::> snapmirror restore -source-path vserverB:dstFG -destination -path vserverA:newFG -snapshot daily.2016-07-15_0010 Warning: This is a disruptive operation and the volume vserverA:newFG will be read-only until the operation completes Do you want to continue? \{y|n\}: y
```

Related information

· snapmirror restore

Disable SVM protection on ONTAP FlexGroup volumes

When the SVM DR flag is set to protected on a FlexGroup volume, you can set the flag to unprotected to disable SVM DR protection on a FlexGroup volume.

Before you begin

- The SVM DR relationship between the primary and secondary is healthy.
- SVM DR protection parameter is set to protected.

Steps

1. Disable protection by using the volume modify command to change the vserver-dr-protection parameter for the FlexGroup volume to unprotected.

```
cluster2::> volume modify -vserver vs1 -volume fg_src -vserver-dr
-protection unprotected
[Job 5384] Job is queued: Modify fg_src.
[Job 5384] Steps completed: 4 of 4.
cluster2::>
```

2. Update the SVM at the secondary site: snapmirror update -destination-path destination_svm_name: -source-path Source_svm_name:

- 3. Verify that the SnapMirror relationship is healthy: snapmirror show
- 4. Verify that the FlexGroup SnapMirror relationship has been removed: snapmirror show -expand

Related information

- snapmirror show
- · snapmirror update

Enable SVM protection on ONTAP FlexGroup volumes

When the SVM DR protection flag is set to unprotected on a FlexGroup volume, you can set the flag to protected to enable SVM DR protection.

Before you begin

- The SVM DR relationship between the primary and secondary is healthy.
- SVM DR protection parameter is set to unprotected.

Steps

1. Enable protection by using the volume modify to change the vserver-dr-protection parameter for the FlexGroup volume to protected.

```
cluster2::> volume modify -vserver vs1 -volume fg_src -vserver-dr
-protection protected
[Job 5384] Job is queued: Modify fg_src.
[Job 5384] Steps completed: 4 of 4.
cluster2::>
```

2. Update the SVM at the secondary site: snapmirror update -destination-path destination svm name -source-path source svm name

```
snapmirror update -destination-path vs1_dst: -source-path vs1:
```

3. Verify that the SnapMirror relationship is healthy: snapmirror show

cluster2::> snapmirror show								
Progress								
Source		Destination	Mirror	Relationship	Total			
Last								
Path	Type	Path	State	Status	Progress	Healthy		
Updated								
vs1:	XDP	vs1_dst:	Snapmir					
				Idle	_	true -		

4. Verify that the FlexGroup SnapMirror relationship is healthy: snapmirror show -expand

cluster2::>	snapr	mirror show -	expand				
Progress Source Last		Destination	Mirror	Relationship	Total		
Path	Type	Path	State	Status	Progress	Healthy	
Updated							
vs1:	XDP	vs1_dst:	Snapmir				
wal fa ara	AUD	vsl dst:fg s	ıra	Idle	_	true	-
vs1.1g_s1c	ADE	vs1_usc.1g_s	Snapmir	rored			
			-	Idle	_	true	-
vs1:fg_src_	_	1 1	0.001				
	XDP	vs1_dst:fg_s	rc0001 Snapmir				
			2110[21112	Idle	_	true	_
vsl:fg_src_	_						
	XDP	vs1_dst:fg_s	rc0002 Snapmir				
			SHAPIHII	Idle	_	true	_
vs1:fg_src_	_						
	XDP	vs1_dst:fg_s	rc0003 Snapmir				
			Shapilitt	Idle	_	true	_
vsl:fg_src_	_0004						
	XDP	vs1_dst:fg_s					
			Snapmir	rored Idle	_	true	_
6 entries we	ere di	isplayed.				0140	

Related information

• snapmirror show

Convert FlexVol volumes to FlexGroup volumes

Learn about converting ONTAP FlexVol volumes to FlexGroup volumes

If you want to expand a FlexVol volume beyond its space limit, you can convert the FlexVol volume to a FlexGroup volume. Beginning with ONTAP 9.7, you can convert standalone FlexVol volumes or FlexVol volumes that are in a SnapMirror relationship to FlexGroup volumes.

Considerations for converting FlexVol volumes to FlexGroup volumes

You should be aware of the features and operations that are supported before you decide to convert FlexVol volumes to FlexGroup volumes.

Operations not supported during conversion

The following operations are not allowed when volume conversion is in progress:

- Volume move
- · Aggregate relocation
- · Planned takeover and giveback in a high-availability configuration
- · Manual and automatic giveback in an high-availability configuration
- · Cluster upgrade and revert
- · FlexClone volume split
- Volume rehost
- · Volume modify and autosize
- · Volume rename
- · Attaching an object store to an aggregate
- Negotiated switchover in MetroCluster configuration
- SnapMirror operations
- Restoring from a snapshot
- Quota operations
- · Storage efficiency operations

You can perform these operations on the FlexGroup volume after successful conversion.

Configurations that are not supported with FlexGroup volumes

- · Offline or restricted volume
- SVM root volume
- SAN
- SMB 1.0
- NVMe namespaces
- Remote Volume Shadow Copy Service (VSS)

Convert ONTAP FlexVol volumes to ONTAP FlexGroup volumes

Beginning with ONTAP 9.7, you can perform an in-place conversion of a FlexVol volume to a FlexGroup volume without requiring a data copy or additional disk space.

Before you begin

- Transitioned volumes can be converted to FlexGroup volumes beginning with ONTAP 9.8.
- The FlexVol volume that is being converted must be online.

• The operations and configurations on the FlexVol volume must be compatible with the conversion process.

Check for the following conditions that can prevent the conversion from succeeding:

• A FlexVol volume was transitioned from 7-Mode using 7MTT (ONTAP 9.7).

Transitioned volumes can be converted beginning with ONTAP 9.8.

- Something is enabled on the volume that is not yet supported with FlexGroup volume; for example, SAN LUNs, Windows
 NFS, SMB1, snapshot naming/autodelete, vmalign set, SnapLock with releases earlier than ONTAP 9.11.1 (SnapLock is supported beginning with ONTAP 9.11.1), space SLO, or logical space
 - 9.11.1 (SnapLock is supported beginning with ONTAP 9.11.1), space SLO, or logical space enforcement/reporting. For more information see Supported and unsupported configurations for FlexGroup volumes.
- The SVM where the FlexVol volume to be converted is located is currently using SVM DR.
- NetApp FlexClone volumes are present, and the FlexVol volume is the parent volume. The volume being converted cannot be a parent or a clone.
- The volume is a NetApp FlexCache origin volume.
- For ONTAP 9.7 and earlier, NetApp snapshots must not exceed 255. For ONTAP 9.8 and later, 1023 snapshots are supported.
- Storage efficiencies are enabled. These must be disabled and can be reenabled after conversion.
- The volume is a source of a SnapMirror relationship, and the destination has not yet been converted.
- The volume is part of an active (not quiesced) SnapMirror relationship.
- Quotas are enabled. These must be disabled and can be reenabled after conversion.
- Volume names are longer than 197 characters.
- The volume is associated with an application.

This is applicable to ONTAP 9.7 only. The limitation is removed in ONTAP 9.8.

- ONTAP processes are running, such as mirroring, jobs, wafliron, NDMP backup, and inode conversion in process.
- The volume is an SVM root volume.
- The volume is too full.

If any of these incompatibilities exist, an error message is generated if the FlexVol volume, and the volume conversion is aborted. You can take corrective actions and retry the conversion.

 If a FlexVol volume is currently at 80% or greater maximum capacity, consider copying the data to a newly created FlexGroup volume instead of performing an in-place conversion. Although FlexGroup member volumes will naturally rebalance over time, converting a high-capacity FlexVol volume to a FlexGroup volume may create performance or balance issues that will not quickly be rebalanced across member volumes.



Converting a very large FlexGroup volume results in a very full FlexGroup volume member constituent, which can create performance issues. For more information, see the section called "When not to create a FlexGroup volume" in the TR FlexGroup volumes - Best Practices and Implementation Guide.

Steps

 Verify that the FlexVol volume is online: volume show -fields vol_name volume-styleextended, state

```
cluster-1::> volume show my_volume -fields volume-style-extended, state
vserver volume    state volume-style-extended
-----
vs0    my_volume online flexvol
```

- 2. Verify whether the FlexVol volume can be converted without issues:
 - a. Log in to the advance privilege mode: set -privilege advanced
 - b. Verify the conversion process: volume conversion start -vserver vs1 -volume flexvol -check-only true

You must rectify all errors before converting the volume.



You cannot convert a FlexGroup volume back to a FlexVol volume.

3. Start the conversion: volume conversion start -vserver svm name -volume vol name

```
cluster-1::*> volume conversion start -vserver vs0 -volume my_volume

Warning: Converting flexible volume "my_volume" in Vserver "vs0" to a
FlexGroup

will cause the state of all Snapshot copies from the volume to
be set

to "pre-conversion". Pre-conversion Snapshot copies cannot be
restored.

Do you want to continue? {y|n}: y
[Job 57] Job succeeded: success
```

4. Verify that the conversion is successful: volume show vol_name -fields volume-style-extended, state

Results

The FlexVol volume is converted to a single-member FlexGroup volume.

After you finish

You can expand the FlexGroup volume, as required.

Convert ONTAP FlexVol volume SnapMirror relationships to ONTAP FlexGroup volume SnapMirror relationships

To convert a FlexVol volume SnapMirror relationship to a FlexGroup volume SnapMirror relationship in ONTAP, you must first convert the destination FlexVol volume followed by the source FlexVol volume.

About this task

- Mixing FlexGroups with FlexVol volumes in SnapMirror/SnapVault relationships is not supported beyond the conversion process.
- FlexGroup conversion is supported only for SnapMirror asynchronous relationships.
- FlexGroup conversion is not supported in SnapMirror cloud relationships.
- Conversion time depends on several variables. Some of the variables include:
 - CPU of the controller
 - Utilization of CPU by other applications
 - Amount of data in the initial snapshot
 - Network bandwidth
 - Bandwidth used by other applications

Before you begin

- The FlexVol volume that is being converted must be online.
- The source FlexVol volume in the SnapMirror relationship must not be the source volume for multiple SnapMirror relationships.

Beginning with ONTAP 9.9.1, fanout SnapMirror relationships are supported for FlexGroup volumes. For more information, see Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes.

• The operations and configurations on the FlexVol volume must be compatible with the conversion process.

An error message is generated if the FlexVol volume has any incompatibility and the volume conversion is aborted. You can take corrective actions and retry the conversion.

Steps

1. Verify that the SnapMirror relationship is healthy:

snapmirror show

Only XDP type mirror relationships can be converted.

Example:

```
Progress
Source Destination Mirror Relationship Total
Last
Path Type Path State Status Progress Healthy
Updated
------
vs0:src_dpv DP vs2:dst_dpv Snapmirrored
Idle - true -
vs0:src_xdp XDP vs2:dst_xdp Snapmirrored
Idle - true -
```

- 2. Verify whether the source volume is compatible for conversion:
 - a. Log in to the advance privilege mode:

```
set -privilege advanced
```

b. Verify the conversion process:

```
volume conversion start -vserver <src_svm_name> -volume <src_vol>
-check-only true
```

Example:

```
volume conversion start -vserver vs1 -volume src_vol -check-only true
```

You must rectify all errors before converting the volume.

- 3. Convert the destination FlexVol volume to FlexGroup volume.
 - a. Quiesce the FlexVol SnapMirror relationship:

```
snapmirror quiesce -destination-path <dest_svm:dest_volume>
```

Example:

```
cluster2::> snapmirror quiesce -destination-path vs2:dst_xdp
```

b. Start the conversion:

volume conversion start -vserver <dest_svm> -volume <dest_volume>

Example:

```
Cluster-1::> volume conversion start -vserver vs2 -volume dst_xdp

Warning: After the volume is converted to a FlexGroup, it will not be possible to change it back to a flexible volume.

Do you want to continue? {y|n}: y

[Job 510] Job succeeded: SnapMirror destination volume "dst_xdp" has been successfully converted to a FlexGroup volume.

You must now convert the relationship's source volume, "vs0:src_xdp", to a FlexGroup.

Then, re-establish the SnapMirror relationship using the "snapmirror resync" command.
```

4. Convert the source FlexVol volume to FlexGroup volume: `

```
volume conversion start -vserver <src_svm_name> -volume <src_vol_name>
```

Example:

```
Cluster-1::> volume conversion start -vserver vs0 -volume src_xdp

Warning: Converting flexible volume "src_xdp" in Vserver "vs0" to a

FlexGroup

will cause the state of all Snapshot copies from the volume to
be set

to "pre-conversion". Pre-conversion snapshots cannot be
restored.

Do you want to continue? {y|n}: y

[Job 57] Job succeeded: success
```

5. Resync the relationship:

```
snapmirror resync -destination-path dest_svm_name:dest_volume
```

Example:

cluster2::> snapmirror resync -destination-path vs2:dst_xdp

After you finish

You must ensure that when the source FlexGroup volume is expanded to include more constituents, the destination volume is also expanded.

Related information

- snapmirror quiesce
- snapmirror resync
- snapmirror show

Copyright information

Copyright © 2025 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.