



Data protection for FlexGroup volumes

ONTAP 9

NetApp
March 08, 2024

Table of Contents


- Data protection for FlexGroup volumes 1
 - Data protection workflow for FlexGroup volumes 1
 - Create a SnapMirror relationship for FlexGroup volumes 2
 - Create a SnapVault relationship for FlexGroup volumes 5
 - Create a unified data protection relationship for FlexGroup volumes 8
 - Create an SVM disaster recovery relationship for FlexGroup volumes 11
 - Transition an existing FlexGroup SnapMirror relationship to SVM DR 13
 - Convert a FlexVol volume to a FlexGroup volume within an SVM-DR relationship. 15
 - Considerations for creating SnapMirror cascade and fanout relationships for FlexGroups. 16
 - Considerations for creating a SnapVault backup relationship and a unified data protection relationship for FlexGroup volumes 17
 - Monitor SnapMirror data transfers for FlexGroup volumes 18

Data protection for FlexGroup volumes

Data protection workflow for FlexGroup volumes

You can create SnapMirror disaster recovery (DR) relationships for FlexGroup volumes. Beginning with ONTAP 9.3, 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.

The data protection workflow consists of verifying the cluster and SVM peer relationships, creating a destination volume, creating a job schedule, specifying a policy, creating a data protection relationship, and initializing the relationship.



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. The following table shows the default policy types and supported custom policy types for different types of data protection relationships.

Relationship type	Default Policy	Custom policy type
SnapMirror DR	MirrorAllSnapshots	async-mirror
SnapVault backup	XDPDefault	vault
Unified data protection	MirrorAndVault	mirror-vault

The MirrorLatest policy is not supported with FlexGroup volumes.

Create a SnapMirror relationship for 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.

What you'll need

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

Cluster and SVM peering

About this task

- 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, you must 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
```

```
cluster1::> volume show -volume srcFG* -is-constituent true
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
vss	srcFG	-	online	RW	400TB
172.86GB	56%				
vss	srcFG__0001	Aggr_cmode	online	RW	25GB
10.86TB	56%				
vss	srcFG__0002	aggr1	online	RW	25TB
10.86TB	56%				
vss	srcFG__0003	Aggr_cmode	online	RW	25TB
10.72TB	57%				
vss	srcFG__0004	aggr1	online	RW	25TB
10.73TB	57%				
vss	srcFG__0005	Aggr_cmode	online	RW	25TB
10.67TB	57%				
vss	srcFG__0006	aggr1	online	RW	25TB
10.64TB	57%				
vss	srcFG__0007	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
```

```
cluster2::> volume show -volume dstFG* -is-constituent true
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
-----	-----	-----	-----	----	-----
-----	-----				
vsd	dstFG	-	online	DP	400TB
172.86GB	56%				
vsd	dstFG__0001	Aggr_cmode	online	DP	25GB
10.86TB	56%				
vsd	dstFG__0002	aggr1	online	DP	25TB
10.86TB	56%				
vsd	dstFG__0003	Aggr_cmode	online	DP	25TB
10.72TB	57%				
vsd	dstFG__0004	aggr1	online	DP	25TB
10.73TB	57%				
vsd	dstFG__0005	Aggr_cmode	online	DP	25TB
10.67TB	57%				
vsd	dstFG__0006	aggr1	online	DP	25TB
10.64TB	57%				
vsd	dstFG__0007	Aggr_cmode	online	DP	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 `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 Snapshot copies for FlexGroup volumes.

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 -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 Snapshot copies that are enabled with the qtree functionality. You must also perform these steps before reverting from ONTAP 9.3 to an earlier version of ONTAP, if you have the qtree functionality enabled on the FlexGroup volumes. [Disable qtree functionality in FlexGroup volumes before reverting](#)

After you finish

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

Create a SnapVault relationship for FlexGroup volumes

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

What you'll need

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
```

```
cluster1::> volume show -volume src* -is-constituent true
Vserver    Volume          Aggregate      State      Type      Size
Available Used%
-----
vss        src              -              online     RW         400TB
172.86GB   56%
vss        src__0001        Aggr_cmode     online     RW         25GB
10.86TB    56%
vss        src__0002        aggr1          online     RW         25TB
10.86TB    56%
vss        src__0003        Aggr_cmode     online     RW         25TB
10.72TB    57%
vss        src__0004        aggr1          online     RW         25TB
10.73TB    57%
vss        src__0005        Aggr_cmode     online     RW         25TB
10.67TB    57%
vss        src__0006        aggr1          online     RW         25TB
10.64TB    57%
vss        src__0007        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 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
```



```
cluster2::> volume show -volume dst* -is-constituent true
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
vsd	dst	-	online	RW	400TB
172.86GB	56%				
vsd	dst__0001	Aggr_cmode	online	RW	25GB
10.86TB	56%				
vsd	dst__0002	aggr1	online	RW	25TB
10.86TB	56%				
vsd	dst__0003	Aggr_cmode	online	RW	25TB
10.72TB	57%				
vsd	dst__0004	aggr1	online	RW	25TB
10.73TB	57%				
vsd	dst__0005	Aggr_cmode	online	RW	25TB
10.67TB	57%				
vsd	dst__0006	aggr1	online	RW	25TB
10.64TB	57%				
vsd	dst__0007	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 `-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 Snapshot copies 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
```

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".
```

Create a unified data protection relationship for 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.

What you'll need

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 srcFG* -is-constituent true
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
vss	srcFG	-	online	RW	400TB
172.86GB	56%				
vss	srcFG__0001	Aggr_cmode	online	RW	25GB
10.86TB	56%				
vss	srcFG__0002	aggr1	online	RW	25TB
10.86TB	56%				
vss	srcFG__0003	Aggr_cmode	online	RW	25TB
10.72TB	57%				
vss	srcFG__0004	aggr1	online	RW	25TB
10.73TB	57%				
vss	srcFG__0005	Aggr_cmode	online	RW	25TB
10.67TB	57%				
vss	srcFG__0006	aggr1	online	RW	25TB
10.64TB	57%				
vss	srcFG__0007	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
```

```
cluster2::> volume show -volume dstFG* -is-constituent true
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
-----	-----	-----	-----	-----	-----
vsd	dstFG	-	online	RW	400TB
172.86GB	56%				
vsd	dstFG__0001	Aggr_cmode	online	RW	25GB
10.86TB	56%				
vsd	dstFG__0002	aggr1	online	RW	25TB
10.86TB	56%				
vsd	dstFG__0003	Aggr_cmode	online	RW	25TB
10.72TB	57%				
vsd	dstFG__0004	aggr1	online	RW	25TB
10.73TB	57%				
vsd	dstFG__0005	Aggr_cmode	online	RW	25TB
10.67TB	57%				
vsd	dstFG__0006	aggr1	online	RW	25TB
10.64TB	57%				
vsd	dstFG__0007	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 Snapshot copies 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 `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 MirrorAndVault
```

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".
```

Create an SVM disaster recovery relationship for 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 FlexGroups 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`

```
cluster2::> snapmirror show
```

Progress							
Source		Destination	Mirror	Relationship	Total		
Last							
Path	Type	Path	State	Status	Progress	Healthy	
Updated							
-----	----	-----	-----	-----	-----	-----	

vs1:	XDP	vs1_dst:	Snapmirrored				
			Idle		-	true	-

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

```
cluster2::> snapmirror show -expand
```

Progress	Source	Destination	Mirror	Relationship	Total		
Last	Path	Type	Path	State	Status	Progress	Healthy
Updated							
-----	----	-----	-----	-----	-----	-----	-----
	vs1:	XDP	vs1_dst:	Snapmirrored			
				Idle		-	true -
	vs1:fg_src	XDP	vs1_dst:fg_src	Snapmirrored			
				Idle		-	true -
	vs1:fg_src__0001						
		XDP	vs1_dst:fg_src__0001	Snapmirrored			
				Idle		-	true -
	vs1:fg_src__0002						
		XDP	vs1_dst:fg_src__0002	Snapmirrored			
				Idle		-	true -
	vs1:fg_src__0003						
		XDP	vs1_dst:fg_src__0003	Snapmirrored			
				Idle		-	true -
	vs1:fg_src__0004						
		XDP	vs1_dst:fg_src__0004	Snapmirrored			
				Idle		-	true -

6 entries were displayed.

Transition an existing FlexGroup SnapMirror relationship to SVM DR

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

What you'll need

- 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.

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

```
snapmirror show -destination-path fg_vs_renamed: -source-path fg_vs:
```

Progress

Source	Destination	Mirror	Relationship	Total		
Last						
Path	Type	Path	State	Status	Progress	Healthy
Updated						
-----	----	-----	-----	-----	-----	-----

fg_vs:	XDP	fg_vs1_renamed:	Broken-off			
			Idle	-	true	-

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`

Convert a FlexVol volume to a FlexGroup volume 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.

What you'll need

- 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:
```

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 SVM-DR 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:
```

6. 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:
```

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

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

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

Considerations for creating SnapMirror cascade and fanout relationships for FlexGroups

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.
- 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 supported 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 in ONTAP 9.9.1.
- 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 FlexGroups 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 FlexGroups.
- All fanout configurations with policy types supported by rest are also supported for REST APIs in FlexGroup fanout configurations.

Considerations for creating a SnapVault backup relationship and a unified data protection relationship for 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 Snapshot copies on the destination volume that are

newer than the latest common Snapshot copy.

- Long-term retention is not supported with FlexGroup volumes.

Long-term retention enables creating Snapshot copies directly on the destination volume without requiring to store the Snapshot copies 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 Snapshot copies 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 copy on the source FlexGroup volume and re-create a Snapshot copy with the same name, the next update transfer to the destination FlexGroup volume fails if the destination volume has a Snapshot copy of the same name.

This is because Snapshot copies cannot be renamed for FlexGroup volumes.

Monitor SnapMirror data transfers for 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::> snapmirror show -relationship-group-type flexgroup
```

Progress

Source Destination Mirror Relationship Total

Last

Path Type Path State Status Progress Healthy

Updated

vss:s	XDP	vsd:d	Snapmirrored				
			Idle		-	true	-
vss:s2	XDP	vsd:d2	Uninitialized				
			Idle		-	true	-

2 entries were displayed.

Copyright information

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

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

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

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

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

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

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

Trademark information

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