



Technical Report

Data protection and backup

NetApp ONTAP FlexGroup volumes

Justin Parisi, NetApp
October 2021 | TR-4678

Abstract

This document covers data protection and backup for NetApp® ONTAP® FlexGroup volumes. Topics include NetApp Snapshot™ copies, NetApp SnapMirror®, and other data protection and backup solutions. For general FlexGroup volume best practices, see [TR-4571: NetApp ONTAP FlexGroup volumes](#). For all other data protection information, see the [data protection guide for your version of ONTAP 9.x](#) on the NetApp product documentation pages.

TABLE OF CONTENTS

Overview	4
Document scope.....	4
Intended audience	4
Data protection terminology.....	4
Data protection features with FlexGroup volumes	6
Snapshot copies with FlexGroup volumes	7
SnapMirror and SnapVault with FlexGroup volumes.....	9
MetroCluster	17
NetApp SnapDiff support	18
FlexVol to FlexGroup conversion: Data protection considerations	18
Backing up FlexGroup volumes	25
NDMP with FlexGroup volumes.....	25
Customer success stories: Backup repositories	29
Conclusion	33
Where to find additional information	33
Version history.....	33
Contact us	33

LIST OF TABLES

Table 1) Data protection features: FlexGroup volumes.	6
Table 2) Snapshot feature support with FlexGroup volumes.....	7
Table 3) SnapMirror feature support with FlexGroup volumes.	9
Table 4) FlexGroup data protection minimums.....	10
Table 5) Member volume count considerations for FlexGroup SnapMirror relationships.	15
Table 6) Data protection features: FlexVol volumes versus FlexGroup volumes.	24
Table 7) NDMP dump performance — ONTAP 9.7 versus ONTAP 9.8.	27
Table 8) NDMP restore performance — ONTAP 9.7 versus ONTAP 9.8.	27
Table 9) NDMP dump performance — ONTAP 9.8: FlexGroup versus FlexVol.....	27
Table 10) NDMP Restore Performance — ONTAP 9.8: FlexGroup versus FlexVol.....	27

LIST OF FIGURES

Figure 1) FlexGroup volume.....	5
---------------------------------	---

Figure 2) Snapshot copies in FlexGroup volumes	8
Figure 3) SnapMirror versus SnapVault with FlexGroup volumes.	11
Figure 4) ONTAP System Manager Data Protection Overview dashboard.	12
Figure 5) ONTAP System Manager — Protect Volume.....	12
Figure 6) SnapMirror cascade and fan-out.	13
Figure 7) 100TB capacity FlexVol volume with 90TB capacity used.	19
Figure 8) FlexVol volume with a 90% full member volume.	19
Figure 9) Workflow for Oracle RMAN backup to FlexGroup volume, conversion, migration to cloud.	30
Figure 10) FlexGroup volume design.	30
Figure 11) SQL Server backup environment.	31
Figure 12) Throughput and total operations during test runs.	32
Figure 13) CPOC scale-out throughput results.....	32

Overview

Document scope

This document intends to cover data protection best practices, considerations, and other related items as they pertain to NetApp ONTAP FlexGroup volumes. This document is not intended to be used for NetApp FlexVol volume configurations, nor is it aimed at audiences using NetApp Data ONTAP operating in 7-Mode.

Intended audience

This document's intended audience includes, but is not limited to, the following roles:

- Storage administrators
- Storage architects
- Field resources
- Business decision makers

If there are any questions about the content of this document, see the "Contact us" section of this document.

Data protection terminology

This section defines key terminology that is used when describing data protection.

Storage virtual machine

A storage virtual machine (SVM) is a logical file system namespace capable of spanning beyond the boundaries of physical nodes in a cluster:

- Clients can access virtual servers from any node in the cluster but only through the associated LIFs.
- Each SVM has a root volume under which more volumes are mounted, extending the namespace.
- An SVM can span several physical nodes.
- It is associated with one or more LIFs; clients access the data on the virtual server through the LIFs, which can live on any node in the cluster.

LIF

A LIF is essentially an IP address with associated characteristics, such as a home port, failover ports, a firewall policy, and a routing group:

- Client network data access is through LIFs dedicated to the SVM.
- An SVM can have more than one LIF. You can have many clients mounting one LIF or one client mounting several LIFs. This fact means that IP addresses are no longer tied to a single physical interface.

Data protection

Data protection is the process of safeguarding data from corruption or loss due to sitewide outages, ransomware, or other unforeseen circumstances that can cost businesses valuable time and money.

FlexClone


NetApp FlexClone® technology enables near-zero-space, exact, writable virtual copies of datasets, including volumes, files, and LUNs. It offers rapid, space-efficient creation of extra data copies ideally suited for disaster recovery testing and test and development environments.

FlexGroup volumes

FlexGroup volumes were introduced in ONTAP 9.1. FlexGroup volumes take the concept of the FlexVol volume and use ONTAP to create a single large container composed of multiple FlexVol volume members as shown in Figure 1. This approach enables a true scale-out NAS file system that can use all resources in a cluster while blending capacity, performance, and simplicity in its deployment.

For more information regarding FlexGroup volumes, see [TR-4571: NetApp ONTAP FlexGroup Volumes—Best practices and implementation guide](#).

Figure 1) FlexGroup volume.



Logical directory replication/unified replication

SnapMirror unified replication refers to the use of NetApp SnapMirror® software with the same (unified) logical replication engine as in NetApp SnapVault® technology. This unified relationship type is designated extended data protection (XDP) and provides single baseline functionality at the volume level. This functionality drastically reduces storage and network bandwidth, which translates immediately into cost savings.

MirrorVault

A MirrorVault is a SnapMirror relationship that integrates SnapVault functionality into the replication, enabling asymmetric NetApp Snapshot™ copy counts on the source and destination. The MirrorAndVault policy in ONTAP controls this replication.

SnapMirror

SnapMirror provides asynchronous replication of volumes, independent of protocol. This replication occurs either in the cluster or to another ONTAP system for data protection and disaster recovery.

SnapRestore

NetApp SnapRestore® is a licensable feature that enables you to restore data from Snapshot copies in ONTAP.

Snapshot copies

Snapshot copies are automatically scheduled point-in-time copies that take up no space and incur no performance overhead when created. Over time, Snapshot copies consume minimal storage space, because only changes to the active file system are written. Individual files and directories can be easily recovered from any Snapshot copy, and the entire volume can be restored back to any Snapshot state in seconds. Snapshot copies can also be triggered manually.

SnapVault

You can copy volumes for space-efficient, read-only, disk-to-disk backup either in the cluster or to another ONTAP system. SnapVault, when used with version-independent SnapMirror, enables a single destination volume to serve as both a backup and disaster recovery copy.

Recovery point objective

Recovery point objective (RPO) refers to the amount of data at risk. This value determines the acceptable loss in data recovery scenarios because the more data you back up, the more cost is allocated to data protection infrastructure and management.

Recovery time objective

Recovery time objective (RTO) refers to the amount of downtime that is deemed as acceptable in disaster recovery scenarios.

Version independence/version flexibility with data protection

SnapMirror XDP offers the ability to replicate between ONTAP clusters that differ in ONTAP versions on source and destination clusters.

Data protection features with FlexGroup volumes

Table 1 lists data protection features and the ONTAP versions that support them with FlexGroup volumes. Check back for each new ONTAP release for more feature support for FlexGroup volumes.

Table 1) Data protection features: FlexGroup volumes.

Supported feature	Version of ONTAP first supported
Snapshot copies	ONTAP 9.1
SnapRestore	ONTAP 9.1 (diagnostic privilege only)
Microsoft Windows Previous Versions tab	ONTAP 9.1
Logical directory replication (XDP)	ONTAP 9.1
Version independence/version flexibility	ONTAP 9.1
SMB/NFS backup	ONTAP 9.1
SnapVault	ONTAP 9.3
MirrorVault	ONTAP 9.3
SnapDiff (2.0 and later)	ONTAP 9.4
NetApp MetroCluster	ONTAP 9.6
NDMP	ONTAP 9.7 (basic) ONTAP 9.8 (Restartable Backup Extension/RBE, EXCLUDE, MULTI_SUBTREE_NAMES, IGNORE_CTIME_MTIME, per qtree exclude)
NetApp ONTAP tools for VMware vSphere (formerly VMware vSphere/Site Recovery Manager)	ONTAP 9.8
1,023 NetApp Snapshot support	ONTAP 9.8
Single-file SnapRestore (through ONTAP CLI)	ONTAP 9.8 (ONTAP tools for VMware vSphere and SnapMirror restore commands only)
NetApp SnapCenter®	ONTAP 9.8 (virtualization only through ONTAP tools for VMware vSphere)
SVM DR	ONTAP 9.9.1*

Supported feature	Version of ONTAP first supported
	* See “How SVM disaster recovery behaves (ONTAP 9.8 and earlier) when FlexGroup volumes are present” for limitations. FlexClone support is added in ONTAP 9.10.1
SnapMirror fan-out	ONTAP 9.9.1
Cascading SnapMirror	ONTAP 9.9.1
Single File SnapRestore	ONTAP 9.10.1
Snapshot rename	ONTAP 9.10.1
NetApp® XCP	All versions of ONTAP
Qtree SnapMirror	n/a
NetApp SnapProtect®	n/a
NetApp SnapManager®	n/a
SnapMirror to NetApp Cloud Backup (formerly NetApp AltaVault™)	n/a
NetApp SnapLock®	n/a
SnapMirror Synchronous	n/a
SnapMirror (data protection/DP style)	n/a
NetApp Snapshot automatic deletion	n/a
SnapMirror to Simple Storage Service (S3)	n/a
SnapMirror to tape (SMTape)	n/a
SnapMirror Business Continuity	n/a (SAN only)

Snapshot copies with FlexGroup volumes

NetApp Snapshot copies are point-in-time copies of a file system. NetApp ONTAP has supported Snapshot copies for decades, and NetApp is one of the pioneers of the technology.

When a NetApp Snapshot copy is created in ONTAP, inode pointers to the data in the active file system are created and are pointed to new locations in the storage. These pointers are read only from a client's perspective. As data is deleted from the active file system, it stays locked in a Snapshot copy until the Snapshot copy is deleted. Thus, space is not freed up when data is deleted until a Snapshot copy is deleted.

You can restore files individually from Snapshot copies from clients by navigating to the `.snapshot` directory in NFS or the `~snapshot` directory in CIFS/SMB or by using the Previous Versions tab in Windows. Snapshot copies can also restore individual files through `snap restore` commands with NetApp FlexVol volumes.

ONTAP supports Snapshot copies with FlexGroup volumes, with most of the functionality that a Snapshot copy provides with regular FlexVol volumes. Table 2 shows which features are supported for Snapshot copies in ONTAP with FlexGroup volumes.

Table 2) Snapshot feature support with FlexGroup volumes.

Snapshot feature/functionality	Supported?	ONTAP version first supported
Snapshot create	Yes	ONTAP 9.1
Snapshot restore (diagnostic level only)	Yes	ONTAP 9.1
Previous Versions tab	Yes	ONTAP 9.1
<code>.snapshot</code> directory access	Yes	ONTAP 9.1

Snapshot feature/functionality	Supported?	ONTAP version first supported
Support for 1,023 Snapshot copies	Yes	ONTAP 9.8
Single file SnapRestore from CLI (see the following Note)	Yes	ONTAP 9.8 (SnapMirror restore only) ONTAP 9.10.1 (full support)
Snapshot rename	No	ONTAP 9.10.1
Snapshot automatic deletion	No	n/a
Snapshot ordinal naming	No	n/a
Snapshot reclaimable space calculation	No	n/a
Snapshot comments	No	n/a
Snapshot remove owner	No	n/a

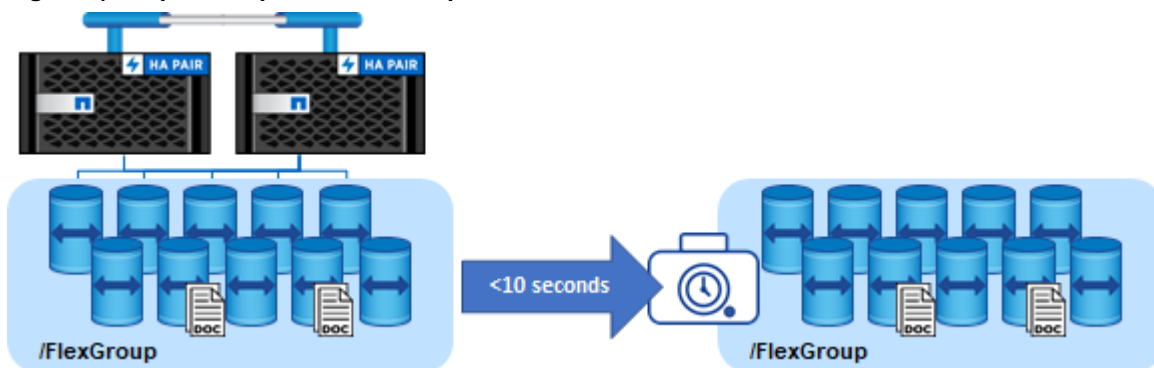
Note: Single File Snapshot Restore (SFSR) is available for virtual machines (VMs) in VMware datastores through the ONTAP tools for VMware vSphere UI and for use with SnapMirror restore in ONTAP 9.8. Single File SnapRestore through the CLI is available in ONTAP 9.10.1 and later.

How Snapshot copies work with FlexGroup volumes

NetApp ONTAP FlexGroup volumes are made up of a group of FlexVol member volumes. Snapshot copies are still made at the FlexVol volume level. When you need to create a FlexGroup Snapshot copy, the member volumes must be coordinated to create a consistent Snapshot copy of the file system. Remember that a FlexGroup volume makes generous use of remote hard links. Therefore, if a Snapshot copy is made of a member volume and the remote hard link is captured in flight, the Snapshot copy is essentially no good. Or, if a hard link is captured without the associated file, the Snapshot copy is no good.

To prevent these scenarios, a FlexGroup volume fences off data access during a Snapshot job and flushes cache entries so that the Snapshot copy is consistent. This process is similar to the Snapshot process that is used for SAN LUNs or for crash-consistent or application-consistent Snapshot copies, such as with VMware virtual machines (VMs). During this access fencing, reads and writes are paused for no more than 10 seconds in a worst-case scenario. Generally, this process completes in about one second, depending on load on the system, and is depicted in Figure 2. If a Snapshot copy takes more than 10 seconds to complete, it fails.

Figure 2) Snapshot copies in FlexGroup volumes.



Snapshot creation failures

If a Snapshot creation fails, ONTAP logs an event management system (EMS) error message:

```
waf1.snap.create.skip.reason: volume X skipping creation of daily.Y snapshot copy (snapshot creation could not be initiated within ten seconds).
```


This error generally occurs when a Snapshot copy cannot finish in the 10-second timeout period. The 10-second value is not configurable.

To avoid scenarios that can cause Snapshot copies to fail because of timeouts, consider the following best practices:

- Avoid creating Snapshot copies during scheduled RAID scrub windows.
- Avoid creating Snapshot copies on a FlexGroup volume when other volumes are experiencing peak workloads.
- If a cluster has multiple FlexGroup volumes, stagger the scheduled Snapshot copy creation not to occur simultaneously.
- Attempt to keep the node's CPU level from 40% to 70%.

If the issue persists, contact technical support for assistance.

FlexGroup Snapshot guidelines

- If any FlexVol member volume cannot make a Snapshot copy because of space or performance issues, then the FlexGroup Snapshot copy is marked as “invalid” and is automatically cleaned up by ONTAP. Invalid Snapshot copies are not viewable by default through the CLI. The process is transparent to storage administrators.
- If a Snapshot copy is considered “partial,” it cannot be used in SnapRestore operations. However, partial Snapshot copies can be used to restore individual files from `.snapshot` directories or the Previous Versions tab.
- SnapRestore is an all-or-nothing proposition. When you restore a FlexGroup volume, you restore the entire container. You cannot restore member FlexVol volumes individually.
- If a FlexGroup volume is modified to include more members, previously created Snapshot copies are considered “partial” and are available only for `.snapshot` directory or Previous Versions access from clients.
- Because of the effect of fencing the FlexGroup volume access during Snapshot copy creation, Snapshot copy schedules must be in intervals of 30 minutes or more.

SnapMirror and SnapVault with FlexGroup volumes

NetApp SnapMirror support for NetApp ONTAP FlexGroup volumes is present in ONTAP 9.1 and later. NetApp SnapVault support is present in ONTAP 9.3 and later.

The current support includes only the logical replication engine (logical replication with storage efficiency [LRSE]/extended data protection [XDP] relationships). Table 3 shows a list of SnapMirror features supported with FlexGroup and the version of ONTAP in which the feature was first introduced.

Table 3) SnapMirror feature support with FlexGroup volumes.

SnapMirror feature/functionality	Supported?	ONTAP version first supported
SnapMirror (logical/XDP)	Yes	ONTAP 9.1
Version-flexible SnapMirror	Yes	ONTAP 9.1
FlexGroup expansion without SnapMirror return to baseline	Yes	ONTAP 9.3
SnapVault	Yes	ONTAP 9.3
Unified SnapMirror and MirrorVault	Yes	ONTAP 9.3
SnapMirror to NetApp Cloud Volumes ONTAP	Yes	ONTAP 9.6
Storage virtual machine disaster recovery (SVM DR)	Yes	ONTAP 9.9.1

SnapMirror feature/functionality	Supported?	ONTAP version first supported
Cascading SnapMirror relationships	Yes	ONTAP 9.9.1
SnapMirror fan-out	Yes	ONTAP 9.9.1
SnapLock	No	n/a
Load-sharing mirrors for data I/O (LSM)	No	n/a (deprecated for all volume types; use NetApp FlexCache volumes instead)
SnapMirror (block/DP)	No	n/a
SnapMirror to AltaVault (deprecated, replaced by Cloud Backup in 2018)	No	n/a
NetApp SolidFire® to ONTAP SnapMirror	No	n/a
SnapMirror Synchronous	No	n/a
SnapMirror to tape (SMTape)	No	n/a
Qtree SnapMirror	No	Supported only in Data ONTAP operating in 7-Mode
NetApp Snapshot naming/autodelete See: Table 2) Snapshot feature support with FlexGroup volumes.	No	n/a
SnapMirror from FlexVol to FlexGroup	No	n/a
SnapMirror from FlexGroup to FlexVol	No	n/a
SnapMirror to S3	No	ONTAP 9.10.1 (S3 to S3 only)
SnapMirror Business Continuity	No	n/a (SAN only)

A NetApp FlexGroup volume can span multiple nodes and there is coordination needed across member FlexVol volumes when a Snapshot or SnapMirror update is performed. Thus, there are limitations to how frequently these operations can take place before job failures might start to be regularly seen. These limits are shown in Table 4 and although they aren't hard limits they should be adhered to for the best possible results.

Table 4) FlexGroup data protection minimums.


Schedule	Supported interval	Hard limit?
SnapMirror schedule	30 minutes	No
Snapshot schedule	30 minutes	No

Should I use SnapMirror or SnapVault?

ONTAP 9.3 adds support for SnapVault for FlexGroup volumes. This support enables storage administrators to replicate FlexGroup volumes asynchronously to a destination volume and preserve more Snapshot copies than exist on the source volume. Figure 3 shows the layout for SnapMirror versus SnapVault with FlexGroup volumes, and SnapMirror and SnapVault use cases can be broken down as follows:

- **SnapMirror relationships.** Intended for disaster recovery and provide an exact replica of the source volume, including the number of Snapshot copies on the source.
- **SnapVault relationships.** Intended for backup and archive use cases by providing a path to Snapshot copies older than what might exist on the source volume. ONTAP 9.8 and later supports up to 1023 Snapshot copies per FlexVol member volume in a FlexGroup.

Figure 3) SnapMirror versus SnapVault with FlexGroup volumes.



How SnapMirror works with FlexGroup volumes

SnapMirror with FlexGroup volumes operate much like Snapshot copies do, because the basis of a SnapMirror copy is the Snapshot copy. Access is fenced, and all volumes must make Snapshot copies at the same time to promote consistency. After that similarity, however, SnapMirror operates by applying concurrent transfers of the member volume constituents to the disaster recovery site. All members transfer at the same time; there is no way to mirror individual member volumes with SnapMirror. If Snapshot copies fail on the source for any member, the SnapMirror transfer does not initiate.

Creating a SnapMirror and SnapVault relationship for a FlexGroup volume

The following section covers basic steps for creating a SnapMirror and SnapVault relationship for a FlexGroup volume. ONTAP has introduced numerous simplicity enhancements for deploying and protecting FlexGroup volumes, including simplified data protection workflows in ONTAP System Manager.

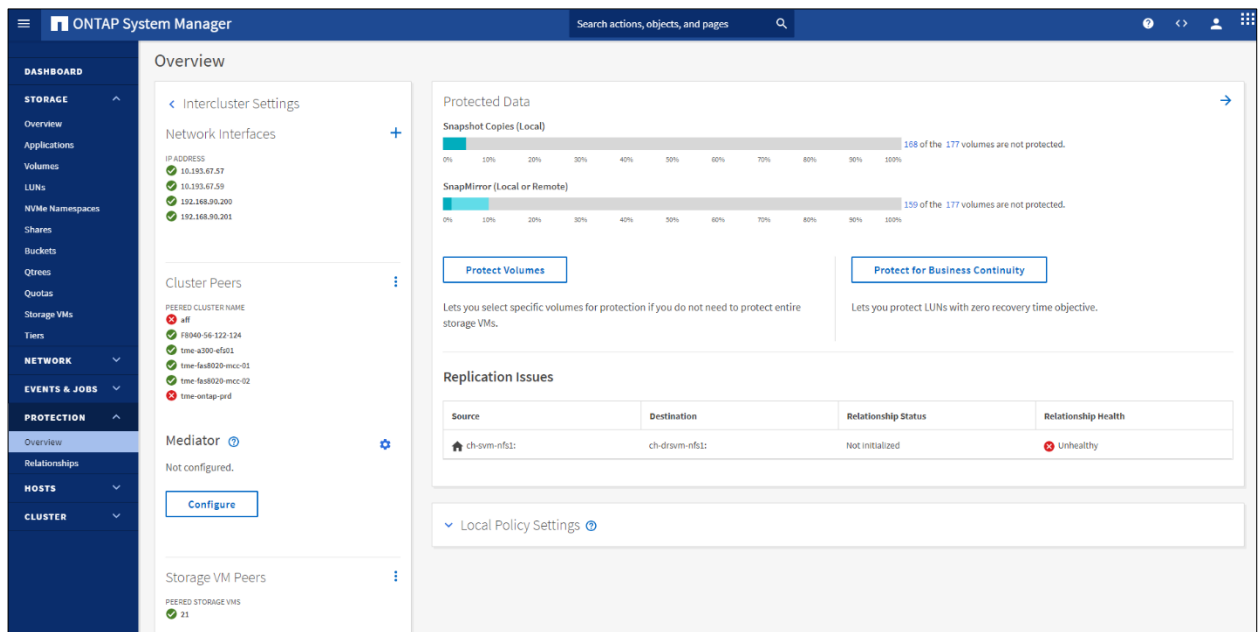
You can view which volumes have data protection enabled in the main volume view or you can use the Protection Overview dashboard to show which clusters have been peered. You can also view how many volumes are and are not protected as well as which SnapMirror relationships might be unhealthy.

In general, the same rules for FlexVol volumes and SnapMirror apply to FlexGroup volumes.

The high-level requirements are:

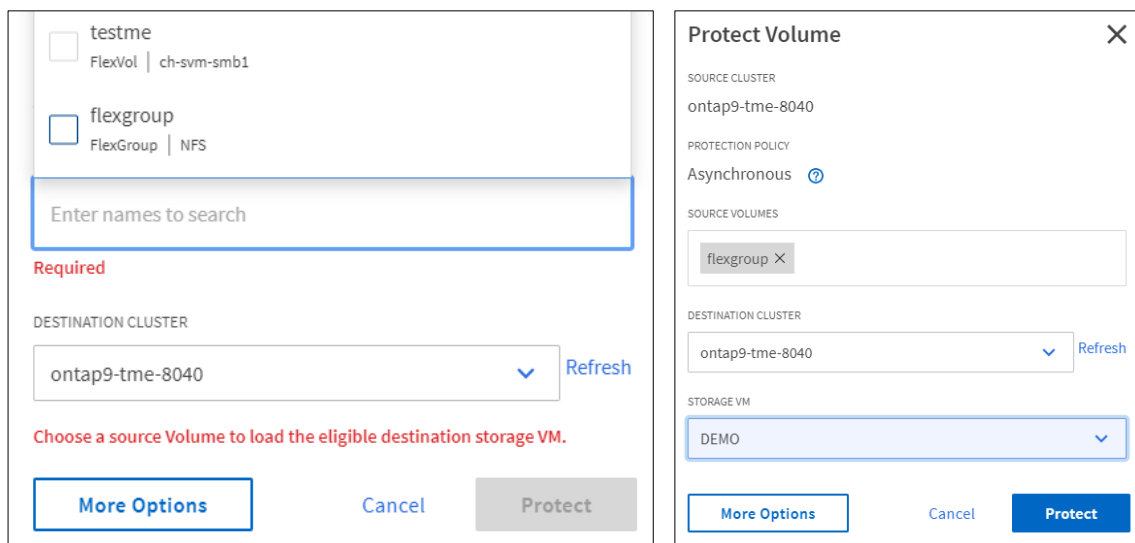
- Valid SnapMirror licenses on source and destination clusters.
- Source and destination clusters within the supported ONTAP version range.
- Cluster and SVM peers created.
- Available space and volume count on the destination to enable volume creation.

Figure 4) ONTAP System Manager Data Protection Overview dashboard.



As shown in Figure 4 if you click “Protect Volumes”, you can select one FlexGroup to configure a relationship. ONTAP System Manager allows you to select multiple FlexVol volumes, but not multiple FlexGroup volumes. Then, you select a destination cluster and destination SVM.

Figure 5) ONTAP System Manager — Protect Volume.



After you have identified the volume that you want to protect and selected the destination cluster and SVM as shown in Figure 5, the rest is simple. Click Protect and ONTAP System Manager creates the destination FlexGroup, SnapMirror relationship and then initializes the first transfer in a single click.

If you do not see your desired destination cluster or SVM in the list, you need to create cluster and SVM peers.

Protecting a new FlexGroup volume

When a new FlexGroup volume is created, storage administrators can easily protect the volume with a SnapMirror or SnapVault relationship during the volume creation. This is provided that the cluster has been peered successfully with another cluster. Simply click “More Options” during volume creation and do the following:

1. To make a FlexGroup volume, select the Distribute Volume Data Across the Cluster checkbox.

Storage and Optimization

CAPACITY

200 TB

PERFORMANCE SERVICE LEVEL

Extreme

Not sure? [Get help selecting type](#)

OPTIMIZATION OPTIONS

☒ Distribute volume data across the cluster

Protection

☒ Enable Snapshot Copies (Local)

☒ Enable SnapMirror (Local or Remote)

PROTECTION POLICY

Asynchronous

Source

CLUSTER

ontap9-tme-8040

STORAGE VM

DEMO

Destination

CLUSTER

ontap9-tme-8040

STORAGE VM

DEMO

Destination Settings

Save Cancel

2. Scroll down and select the desired protection options.
3. Click Save.

ONTAP does the rest.

Note: Both clusters must have a SnapMirror license.

SnapMirror and SnapVault considerations


The following section covers some considerations about SnapMirror and SnapVault use with FlexGroup volumes.


SnapMirror fan-out and cascading SnapMirror

With NetApp FlexVol volumes, it is possible to set up SnapMirror relationships that either span source → destination → secondary destination (cascade) or source A → destination A and source A → destination B (fan-out).

These SnapMirror options, as shown in Figure 6, are available in ONTAP 9.9.1 for use with FlexGroup volumes.

Figure 6) SnapMirror cascade and fan-out.





SnapMirror resync considerations

A `snapmirror resync` operation restores or redefines a SnapMirror source or destination relationship that was broken with the `snapmirror break` command. When a resync occurs, a common Snapshot copy between the primary and secondary FlexGroup volumes is used to restore the FlexGroup volume to a common checkpoint. With this approach, any data accumulated on the secondary volume after a SnapMirror break is lost. An appropriate confirmation occurs when `snapmirror resync` is run. After you run the resync, all member volumes in the FlexGroup volume are set to a data protection status. This setting fences off read/write permission to the volumes and maintains consistency in the volumes. Resyncs might report failures if a FlexGroup volume's source or destination has been expanded (more member volumes added) between the previous `snapmirror break` and `snapmirror resync` commands.

SnapMirror restore considerations

SnapMirror restore operations restore the entire contents of a Snapshot copy from one volume to another volume. A SnapMirror relationship of type RST is created from the source volume to the destination volume by the `snapmirror restore` command. This relationship lasts during the restore operation and is deleted when the command completes successfully.

When you use `snapmirror restore` to recover from a backup, the entire FlexGroup volume is restored, similar to Snapshot restore. Again, keep the following in mind:

- Individual member volumes cannot be restored.
- Single File SnapRestore (using the ONTAP CLI or UI) is supported in ONTAP 9.10.1 and later, and you can restore single files from clients by using the Previous Versions tab in CIFS/SMB or the `.snapshot` directory in NFS. [Single-file SnapMirror restore](#) is supported by the CLI in ONTAP 9.8 and later.
- When a read/write volume is converted to a data protection volume and reverted to read/write, any data accumulated beyond the latest Snapshot copy is lost. Before restoring a SnapMirror volume, create a new Snapshot copy on the primary volume to avoid data loss.

Expanding a FlexGroup volume/adding new member volumes

You can use volume expansion to increase the number of member volumes in a FlexGroup volume to add capacity and scale out a volume across multiple nodes in a cluster.

The `volume expand` command does not work natively with FlexGroup volumes participating in SnapMirror relationships earlier than ONTAP 9.3 because those commands require a return to baseline of the SnapMirror relationship. ONTAP 9.3 introduces the enhancement to enable volume expansion on FlexGroup volumes participating in a SnapMirror relationship without the need to return to baseline. As of ONTAP 9.3, ONTAP automatically adjusts the FlexGroup member volume count on the next SnapMirror update.

Note: If you use SnapMirror with FlexGroup volumes, use ONTAP 9.3 or later.

Expanding FlexGroup volumes in SnapMirror relationships earlier than ONTAP 9.3

To expand a volume (to add more members) in a SnapMirror relationship earlier than ONTAP 9.3, perform the following steps:

1. Perform `snapmirror delete` of the existing relationship on the destination.
2. Perform `snapmirror release` on the source.
3. Perform `volume delete` of the destination FlexGroup data protection volume.
4. Perform `volume expand` of the source FlexGroup volume.
5. Perform `volume create` of a new destination FlexGroup data protection volume with same size and constituent count as the source FlexGroup volume.
6. Perform `snapmirror initialize` of the new relationship (return to baseline).

Growing the member volumes without needing to return to baseline the relationship is supported with SnapMirror and FlexGroup as of ONTAP 9.3.

FlexGroup SnapMirror guidelines

- The guidelines for Snapshot copies also apply to SnapMirror and are summarized in Table 5. Partial Snapshot copies fail the entire Snapshot operation. Snapshot copies are created as a group. Access is fenced off for no longer than 10 seconds.
- FlexGroup volumes that are being copied by using SnapMirror must have an **equal number** of member volumes on the source and destination.
- In versions earlier than ONTAP 9.3, if a FlexGroup volume is expanded to include more members, you must return to baseline the SnapMirror relationship. This step is completed by using a new secondary FlexGroup volume with the correct number of member volumes. After ONTAP 9.3, ONTAP manages the SnapMirror relationship adjustment.
- Destination FlexGroup volumes can be larger than source volumes but cannot be smaller.
- Destination FlexGroup volumes can reside on any supported ONTAP architecture, regardless of hardware, provided the capacity and member volume requirements are met. For example, a FlexGroup volume that lives on a four-node NetApp AFF A700 all-flash storage system cluster can be mirrored to a single-node NetApp FAS2600 series node if it has the same number of member volumes.
- SnapMirror schedule intervals should not be less than 30 minutes.

Table 5) Member volume count considerations for FlexGroup SnapMirror relationships.

Member volume count limits for FlexGroup volumes participating in SnapMirror relationships	ONTAP 9.4 and earlier	ONTAP 9.5 and later
Member volume count	32	200
Member volume count per node*	n/a	50
Member volume count per node (multiple FlexGroup volumes)**	n/a	500
Member volume count per cluster (all FlexGroup volumes)**	100	6,000

*Exceeding this limit can affect recovery point objectives (RPOs).

**Same limits as FlexVol volumes.

Creating SnapMirror relationships when NetApp FabricPool is involved

When you create a FlexGroup volume on aggregates with FabricPool enabled, each aggregate where member volumes are created must be FabricPool aggregates. If there are non-FabricPool aggregates, then creation fails. These aggregates include FlexGroup volumes that are part of a SnapMirror relationship. Automated tools like ONTAP System Manager and the CLI option `-auto-provision-as`

try to use all available aggregates in the system. If you have a mixture of aggregates or if the FlexGroup creation fails due to FabricPool aggregates, use the manual creation methods covered in [TR-4571](#).

How SVM disaster recovery behaves (ONTAP 9.8 and earlier) when FlexGroup volumes are present

If you are using SVM disaster recovery in your cluster in ONTAP releases earlier than ONTAP 9.9.1, you cannot use FlexGroup volumes in the same SVM. Also, if you try to create an SVM disaster recovery relationship in an SVM where FlexGroup volumes are present before ONTAP 9.9.1, the command fails with an error. ONTAP 9.9.1 adds support for SVM-DR for FlexGroup volumes, but with the following caveats:

When SVM-DR is in use, there is no current support in ONTAP 9.10.1 for the following:

- FabricPool
- Fan-out
- Cascading SnapMirror

Note: Support for FlexClone volumes and FlexVol to FlexGroup volume conversion without a SnapMirror return to baseline is available in ONTAP 9.10.1.

SnapMirror throttle behavior

When you set a [SnapMirror throttle](#) on a relationship with a FlexGroup volume, the throttle is not divided across member volumes. Instead, the throttle is set the same to each member volume in the FlexGroup.

For example, if you set a throttle of 100Mbps on a FlexGroup with 16 members, the throttle is not 100Mbps/16 members. Instead, it is 100Mbps * 16 members, for a total throttle of 1600Mbps.

SnapMirror restore of a single file from the CLI

In releases earlier than ONTAP 9.10.1, you cannot restore single files from the CLI by using the `snapshot restore` command, but you can use the `snapmirror restore` command with the `-file-list` flag to restore single files in a FlexGroup volume from a SnapMirror destination volume in ONTAP 9.8 and later.

```
cluster::> snapmirror restore -file-list [/folder1/folder2/file.name] -source-snapshot
[snapname.202x-xx-xx_xxxx] -source-path [SVM_TO_RESTORE_FROM:vol] -destination-path
{SVM_TO_RESTORE_TO:vol}
```

Note: In ONTAP 9.8 and earlier, SnapMirror cascade and fan-out is not supported, which means you cannot restore a file by using SnapMirror restore without first breaking the SnapMirror relationship. In ONTAP 9.9.1 and later, you can restore single files by using SnapMirror restore without needing to break the mirror relationship.

Note: In ONTAP 9.10.1, Single File SnapRestore is supported.

For example, in the Tech_ONTAP volume, there is a file named `ILoveNetApp.mp3`.

```
# pwd
/ToTarchive/TechONTAP
# ls -la | grep Love
-rwxr-xr-x 1 host games      41086 Sep 21  2017 ILoveNetApp.mp3
```

This example is the SnapMirror relationship:

```
cluster::*> snapmirror show
```

Source	Destination	Mirror	Relationship	Total	Progress
Path	Type	Path	State	Status	Progress
					Healthy
DEMO:Tech_ONTAP					


```
XDP  COMPANYB:Tech_ONTAP_mirror
      Snapmirrored
      Idle          -          true  -
```

We delete the file:

```
# rm ILoveNetApp.mp3
rm: remove regular file 'ILoveNetApp.mp3'? y
# ls -la | grep Love
#
```

Then, we restore it by using `snapmirror restore` using the snapshot `daily.2021-01-19_0010`.

```
cluster::> snapmirror restore -file-list /TechONTAP/ILoveNetApp.mp3 -source-snapshot daily.2021-01-19_0010 -source-path COMPANYB:Tech_ONTAP_mirror -destination-path DEMO:Tech_ONTAP
```

Warning: This command will overwrite any file on destination "DEMO:Tech_ONTAP" that has the same path as any of the files to be restored.

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

[Job 35659] Job is queued: snapmirror restore from source "COMPANYB:Tech_ONTAP_mirror" for the snapshot daily.2021-01-19_0010.

Within seconds, the file is restored:

```
cluster::> job show -id 35659 -instance
```

```
          Job ID: 35659
      Owning Vserver: cluster
            Name: Snapmirror FG Restore
      Description: snapmirror restore from source "COMPANYB:Tech_ONTAP_mirror" for the
snapshot daily.2021-01-19_0010
            Priority: High
              Node: cluster-01
            Affinity: Cluster
            Schedule: @now
      Queue Time: 03/11 13:32:53
      Start Time: 03/11 13:32:54
      End Time: 03/11 13:33:09
      Drop-dead Time: -
      Restarted?: false
            State: Success
      Status Code: 0
      Completion String: SnapMirror FG Restore Succeeded
            Job Type: FG RestoreV2
      Job Category: FG SnapMirror
      Execution Progress: Complete: SnapMirror FG Restore Succeeded [0]
            User Name: admin
Restart Is Delayed by Module: -
```

This is the output from the client:

```
# ls -la | grep Love
-rwxr-xr-x 1 host games 41086 Mar 11 13:33 ILoveNetApp.mp3
```

MetroCluster

NetApp ONTAP 9.6 introduces support for FlexGroup volumes on NetApp MetroCluster deployments (FC and IP).

MetroCluster software is a solution that combines array-based clustering with synchronous replication to deliver continuous availability and zero data loss at the lowest cost. There are no stated limitations or caveats for FlexGroup volumes with MetroCluster.

For more information about MetroCluster, see [TR-4705: NetApp MetroCluster: Solution Architecture and Design](#).

NetApp SnapDiff support

ONTAP 9.4 and later introduces support for NetApp SnapDiff® for FlexGroup volumes. SnapDiff is described in the kilobyte article, [FAQ: SnapDiff Support in ONTAP](#).

NetApp SnapDiff is only used by licensed backup partners. The APIs are not available for public consumption but only through supported backup software. FlexGroup volumes support only NetApp SnapDiff 2.0 and later, so contact your backup vendor to find out what SnapDiff support they currently offer.

For more information about NetApp SnapDiff, see the following resources:

- [Tech ONTAP Podcast Episode 264: NetApp ONTAP SnapDiff](#)
- [A New Backup Architecture – SnapDiff V3](#)

FlexVol to FlexGroup conversion: Data protection considerations

In NetApp ONTAP 9.7 and later, you can convert a single NetApp FlexVol volume to a FlexGroup volume with a single member volume, in place, with less than 40 seconds disruption. You can perform this conversion regardless of the amount of data capacity or number of files in the volume. There is no need to remount clients, copy data, or make other modifications that could create a maintenance window. After you convert the FlexVol volume to a FlexGroup volume, you can add new member volumes to expand the capacity.

Why convert a volume to a FlexGroup volume

FlexGroup volumes offer a few advantages over FlexVol volumes, such as:

- Ability to expand beyond 100TB and 2 billion files in a single volume
- Ability to scale out capacity or performance nondisruptively
- Multi-threaded performance for high ingest workloads
- Simplification of volume management and deployment

For example, if you have a workload that is growing rapidly but you do not want to migrate the data (but you still want to provide more capacity), you can convert to a FlexGroup volume. Or perhaps a workload's performance is not good enough on a FlexVol volume, so you want to provide better performance handling with a FlexGroup volume. Converting can help here, too.


When not to convert a FlexVol volume

Converting a FlexVol volume to a FlexGroup volume might not always be the best option. If you require FlexVol features that are not available in FlexGroup volumes, then you should hold off. For example, SnapMirror Synchronous relationships are currently not supported, so if you need these features, you should stay with FlexVol volumes.

Also, if you have a FlexVol volume that is already large (80–100TB) and already full (80–90%), you might want to copy the data rather than convert it. This option might be preferable because the converted FlexGroup volume has a large, full member volume. Such a large volume can create performance issues and does not fully resolve your capacity issues, particularly if that dataset contains files that grow over time.


For example, if you have a FlexVol volume that is 100TB in capacity and 90TB used, it looks as shown in Figure 7:

Figure 7) 100TB capacity FlexVol volume with 90TB capacity used.



If you convert this 90% full volume to a FlexGroup volume, you have a 90% full member volume. If you add new member volumes, as shown in Figure 8, they are 100TB each and 0% full, so they take on most new workloads. The data does not automatically rebalance, and if the original files grow over time, you can still run out of space with nowhere to go (because 100TB is the maximum member volume size).

Figure 8) FlexVol volume with a 90% full member volume.



For more detailed information about FlexVol to FlexGroup conversion, including commands, caveats, functionality, and examples, see [TR-4571: NetApp ONTAP FlexGroup volumes](#). The following sections pertain to FlexVol conversion only as it applies to data protection and disaster recovery ONTAP functionality.

FlexVol conversion: NetApp ONTAP Snapshot considerations

When a FlexVol volume has NetApp Snapshot copies present and is converted to a FlexGroup volume, those Snapshot copies remain in place and can be used to do client-side restores of files and folders through the `.snapshot` or `~snapshot` directory or through the Previous Versions tab in Windows. However, the leftover Snapshot copies are unavailable for full NetApp SnapRestore operations. These Snapshot copies age out eventually and are deleted by ONTAP. New FlexGroup Snapshot copies can be used to perform SnapRestore operations. Snapshot copies being tiered by FabricPool continue to be tiered with no disruption or rehydration of data until file access is requested from those Snapshot copies.

In ONTAP 9.7, if a FlexVol volume has more than 255 Snapshot copies present, then FlexVol to FlexGroup volume conversion is blocked because FlexGroup volumes do not support more than 255 Snapshot copies in that release. Snapshot copies must be deleted to get below 255.

Also, in ONTAP 9.7, each Snapshot copy per volume in ONTAP has a unique ID number (`physical-snap-id`), up to 1,023. These IDs cycle back to 1 after the maximum ID number has been reached. When you attempt to convert a FlexVol volume to a FlexGroup volume, if `physical-snap-id` values greater than 255 exist, conversion is blocked until those IDs are cleared. You can view `physical-snap-id` values with the advanced privilege command `snapshot show -fields physical-snap-id`. To view only offending Snapshot copies, use the following command:

```
cluster::*> snapshot show -vserver DEMO -volume flexvol -physical-snap-id >255 -fields physical-snap-id
```

vserver	volume	snapshot	physical-snap-id
DEMO	flexvol	weekly.2020-01-05_0015	471
DEMO	flexvol	weekly.2020-01-12_0015	648
DEMO	flexvol	daily.2020-01-14_0010	698
DEMO	flexvol	daily.2020-01-15_0010	723
DEMO	flexvol	hourly.2020-01-15_1405	737
DEMO	flexvol	hourly.2020-01-15_1505	738
DEMO	flexvol	hourly.2020-01-15_1605	739
DEMO	flexvol	hourly.2020-01-15_1705	740
DEMO	flexvol	hourly.2020-01-15_1805	741
DEMO	flexvol	hourly.2020-01-15_1905	742

To clear the volume of `physical-snap-id` IDs greater than 255, you have the following options:

- Delete the offending Snapshot copies.
- Wait for the offending Snapshot copies to roll off.
- Run a script that creates and deletes Snapshot copies until the `physical-snap-id` IDs cycle back below 255.

ONTAP 9.8 lifts this restriction by adding support for 1023 Snapshot copies. If you don't want to delete Snapshot copies when converting, upgrade to ONTAP 9.8.

FlexVol conversion: SnapMirror considerations

To convert FlexVol volumes participating in SnapMirror relationships, follow the same considerations required for Snapshot copies in the previous section.

If you transitioned a FlexVol volume from a system running Data ONTAP operating in 7-Mode by using the 7-Mode Transition Tool (7MTT), conversion is prevented in ONTAP 9.7 but allowed in ONTAP 9.8.

To check the volume's transition status, use the following command:

```
volume show -volume [volname] -fields is-transitioned
```

Also, there are considerations needed for both source and destination volumes to preserve the existing SnapMirror relationship and to help avoid the need to return to baseline the SnapMirror relationship.

Steps for converting a FlexVol volume in a SnapMirror relationship

When converting a FlexVol volume participating in a SnapMirror relationship, use the following steps.

1. Optional: Create a NetApp FlexClone copy of the volume to be converted and split it to test FlexVol to FlexGroup volume conversion.
For more information, see [TR-4571: NetApp ONTAP FlexGroup volumes](#).
2. Before you convert a FlexVol volume, it is a best practice to first run the command with the `-check-only` flag on both source and destination volumes to see if any preparation steps are required and if there are any conversion blockers present. For details about possible conversion blockers, see "FlexVol conversion: NetApp ONTAP Snapshot considerations" and "FlexVol conversion: SnapMirror considerations."

3. Clear any blockers on the FlexVol volume or wait until the blockers can be cleared (for example, if Snapshot copies must be deleted).
4. Quiesce the SnapMirror relationship.
5. Convert the SnapMirror destination volume first.
6. After the destination volume is converted, test access and functionality of the new FlexGroup volume if desired.
7. Convert the source volume.
8. Resync the SnapMirror relationship.
9. Test access and file restore functionality as desired.

Expanding the FlexGroup volume in SnapMirror relationships

When a FlexVol volume is converted to a FlexGroup volume, the result is a FlexGroup volume with a single member volume. A FlexGroup volume can scale out to hundreds of member volumes for optimal capacity and performance considerations. Usually, a newly converted FlexGroup volume must then be expanded to add more member volumes to realize the FlexGroup volume's full potential.

Because a SnapMirror relationship with FlexGroup volumes requires the same number of member volumes on the source and the destination, expanding the source volume causes a temporary disruption in SnapMirror operations. ONTAP automatically expands the destination FlexGroup volume for you on the next SnapMirror update, as covered in the section "Expanding a FlexGroup volume/adding new member volumes". Also, see the example in the next section.

Converting a FlexVol volume in a SnapMirror relationship: Example

You can also convert FlexVol volumes that are part of existing SnapMirror relationships without disruption. The following is a volume in a SnapMirror relationship:

```
cluster::*> snapmirror show -destination-path data_dst -fields state
source-path destination-path state
-----
DEMO:data    DEMO:data_dst    Snapmirrored
```

If you try to convert the source, you get an error.

```
cluster::*> vol conversion start -vserver DEMO -volume data -check-only true

Error: command failed: Cannot convert volume "data" in Vserver "DEMO" to a FlexGroup. Correct the
following issues and retry the command:
* Cannot convert source volume "data" because destination volume "data_dst" of the
SnapMirror relationship with "data" as the source is not converted. First check if the source
can be converted to a FlexGroup volume using "vol conversion start -volume data -convert-to
flexgroup -check-only true". If the conversion of the source can proceed then first convert the
destination and then convert the source.
```

So, you must convert the destination first.

1. To convert the destination, quiesce the SnapMirror relationship:

```
cluster::*> vol conversion start -vserver DEMO -volume data_dst -check-only true

Error: command failed: Cannot convert volume "data_dst" in Vserver "DEMO" to a FlexGroup. Correct
the following issues and retry the command:
* The relationship was not quiesced. Quiesce SnapMirror relationship using "snapmirror quiesce -
destination-path data_dst" and then try the conversion.
```

2. Next, convert the volume.

```
cluster::*> snapmirror quiesce -destination-path DEMO:data_dst
Operation succeeded: snapmirror quiesce for destination "DEMO:data_dst".

cluster::*> vol conversion start -vserver DEMO -volume data_dst -check-only true
```

Conversion of volume "data_dst" in Vserver "DEMO" to a FlexGroup can proceed with the following warnings:

- * After the volume is converted to a FlexGroup, it will not be possible to change it back to a flexible volume.
- * Converting flexible volume "data_dst" in Vserver "DEMO" 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.

When you convert the volume, it tells you your next steps.

```
cluster::*> vol conversion start -vserver DEMO -volume data_dst
```

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

Warning: Converting flexible volume "data_dst" in Vserver "DEMO" 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 23710] Job succeeded: SnapMirror destination volume "data_dst" has been successfully converted to a FlexGroup volume. **You must now convert the relationship's source volume, "DEMO:data", to a FlexGroup. Then, re-establish the SnapMirror relationship using the "snapmirror resync" command.**

3. Convert the source volume:

```
cluster::*> vol conversion start -vserver DEMO -volume data
```

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

Warning: Converting flexible volume "data" in Vserver "DEMO" 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 23712] Job succeeded: success

4. Resync the mirror:

```
cluster::*> snapmirror resync -destination-path DEMO:data_dst
```

Operation is queued: snapmirror resync to destination "DEMO:data_dst".

```
cluster::*> snapmirror show -destination-path DEMO:data_dst -fields state
```

```
source-path destination-path state
```

```
-----
```

```
DEMO:data DEMO:data_dst Snapmirrored
```

The conversion works, but the most important part of a SnapMirror relationship is the restore operation. Therefore, you must see if you can access files from the destination volume's Snapshot copy.

5. Mount the source and destination and compare the `ls` output:

```
# mount -o nfsvers=3 DEMO:/data_dst /dst
```

```
# mount -o nfsvers=3 DEMO:/data /data
```

This following output shows the content of the source volume:

```
# ls -lah /data
```

```
total 14G
```

```
drwxrwxrwx 6 root root 4.0K Nov 14 11:57 .
```

```
dr-xr-xr-x. 54 root root 4.0K Nov 15 10:08 ..
```

```
drwxrwxrwx 2 root root 4.0K Sep 14 2018 cifslink
```

```
drwxr-xr-x 12 root root 4.0K Nov 16 2018 nas
```

```
-rwxrwxrwx 1 prof1 ProfGroup 0 Oct 3 14:32 newfile
```

```
drwxrwxrwx 5 root root 4.0K Nov 15 10:06 .snapshot
```

```
lrwxrwxrwx 1 root root 23 Sep 14 2018 symlink -> /shared/unix/linkedfile
```

```
drwxrwxrwx 2 root bin 4.0K Jan 31 2019 test
```

```
drwxrwxrwx 3 root root 4.0K Sep 14 2018 unix
```

```
-rwxrwxrwx 1 newuser1 ProfGroup 0 Jan 14 2019 userfile
```

```
-rwxrwxrwx 1 root root 6.7G Nov 14 11:58 Windows2.iso
```

```
-rwxrwxrwx 1 root root 6.7G Nov 14 11:37 Windows.iso
```

The destination volume matches exactly, as it should.

```
# ls -lah /dst
total 14G
drwxrwxrwx 6 root root 4.0K Nov 14 11:57 .
dr-xr-xr-x 54 root root 4.0K Nov 15 10:08 ..
drwxrwxrwx 2 root root 4.0K Sep 14 2018 cifslink
dr-xr-xr-x 2 root root 0 Nov 15 2018 nas
-rwxrwxrwx 1 prof1 ProfGroup 0 Oct 3 14:32 newfile
drwxrwxrwx 4 root root 4.0K Nov 15 10:05 .snapshot
lrwxrwxrwx 1 root root 23 Sep 14 2018 symlink -> /shared/unix/linkedfile
drwxrwxrwx 2 root bin 4.0K Jan 31 2019 test
drwxrwxrwx 3 root root 4.0K Sep 14 2018 unix
-rwxrwxrwx 1 newuser1 ProfGroup 0 Jan 14 2019 userfile
-rwxrwxrwx 1 root root 6.7G Nov 14 11:58 Windows2.iso
-rwxrwxrwx 1 root root 6.7G Nov 14 11:37 Windows.iso
```

If you `ls` to the Snapshot copy in the destination volume, you see the expected files.

```
# ls -lah /dst/.snapshot/snapmirror.7e3cc08e-d9b3-11e6-85e2-00a0986b1210_2163227795.2019-11-15_100555/
total 14G
drwxrwxrwx 6 root root 4.0K Nov 14 11:57 .
drwxrwxrwx 4 root root 4.0K Nov 15 10:05 ..
drwxrwxrwx 2 root root 4.0K Sep 14 2018 cifslink
dr-xr-xr-x 2 root root 0 Nov 15 2018 nas
-rwxrwxrwx 1 prof1 ProfGroup 0 Oct 3 14:32 newfile
lrwxrwxrwx 1 root root 23 Sep 14 2018 symlink -> /shared/unix/linkedfile
drwxrwxrwx 2 root bin 4.0K Jan 31 2019 test
drwxrwxrwx 3 root root 4.0K Sep 14 2018 unix
-rwxrwxrwx 1 newuser1 ProfGroup 0 Jan 14 2019 userfile
-rwxrwxrwx 1 root root 6.7G Nov 14 11:58 Windows2.iso
-rwxrwxrwx 1 root root 6.7G Nov 14 11:37 Windows.iso
```

6. Next, expand the FlexGroup source to provide more capacity.

```
cluster::*> volume expand -vserver DEMO -volume data -aggr-list aggr1_node1,aggr1_node2 -aggr-list-multiplier
```

Warning: The following number of constituents of size 30TB will be added to FlexGroup "data": 4.
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 23720] Job succeeded: Successful

The source volume now has five member volumes. The destination volume has only one.

```
cluster::*> vol show -vserver DEMO -volume data*
Vserver Volume Aggregate State Type Size Available Used%
-----
DEMO data - online RW 150TB 14.89TB 0%
DEMO data__0001 aggr1_node2 online RW 30TB 7.57TB 0%
DEMO data__0002 aggr1_node1 online RW 30TB 7.32TB 0%
DEMO data__0003 aggr1_node2 online RW 30TB 7.57TB 0%
DEMO data__0004 aggr1_node1 online RW 30TB 7.32TB 0%
DEMO data__0005 aggr1_node2 online RW 30TB 7.57TB 0%
DEMO data_dst - online DP 30TB 7.32TB 0%
DEMO data_dst__0001
aggr1_node1 online DP 30TB 7.32TB 0%
8 entries were displayed.
```

7. Update the mirror, and ONTAP fixes it for you:

```
cluster::*> snapmirror update -destination-path DEMO:data_dst
Operation is queued: snapmirror update of destination "DEMO:data_dst".
```

The update initially fails with the following error message:

```
Last Transfer Error: A SnapMirror transfer for the relationship with destination FlexGroup
"DEMO:data_dst" was aborted because the source FlexGroup was expanded. A SnapMirror AutoExpand
job with id "23727" was created to expand the destination FlexGroup and to trigger a SnapMirror
transfer for the SnapMirror relationship. After the SnapMirror transfer is successful, the
```

"healthy" field of the SnapMirror relationship will be set to "true". The job can be monitored using either the "job show -id 23727" or "job history show -id 23727" commands.

The job expands the volume, and you can then update again.

```
cluster::*> job show -id 23727
Owning
Job ID Name Vserver Node State
-----
23727 Snapmirror Expand cluster
node1
Success
Description: SnapMirror FG Expand data_dst

cluster::*> snapmirror show -destination-path DEMO:data_dst -fields state
source-path destination-path state
-----
DEMO:data DEMO:data_dst Snapmirrored
```

Now both FlexGroup volumes have the same number of member volumes.

```
cluster::*> vol show -vserver DEMO -volume data*
Vserver Volume Aggregate State Type Size Available Used%
-----
DEMO data - online RW 150TB 14.88TB 0%
DEMO data__0001 aggr1_node2 online RW 30TB 7.57TB 0%
DEMO data__0002 aggr1_node1 online RW 30TB 7.32TB 0%
DEMO data__0003 aggr1_node2 online RW 30TB 7.57TB 0%
DEMO data__0004 aggr1_node1 online RW 30TB 7.32TB 0%
DEMO data__0005 aggr1_node2 online RW 30TB 7.57TB 0%
DEMO data_dst - online DP 150TB 14.88TB 0%
DEMO data_dst__0001
aggr1_node1 online DP 30TB 7.32TB 0%
DEMO data_dst__0002
aggr1_node1 online DP 30TB 7.32TB 0%
DEMO data_dst__0003
aggr1_node2 online DP 30TB 7.57TB 0%
DEMO data_dst__0004
aggr1_node1 online DP 30TB 7.32TB 0%
DEMO data_dst__0005
aggr1_node2 online DP 30TB 7.57TB 0%
```

Data protection feature parity in FlexVol and FlexGroup volumes

Before converting a FlexVol volume to a FlexGroup volume, it is important to evaluate all ONTAP features that exist in FlexVol volumes but are currently unavailable in FlexGroup volumes. Table 6 lists data protection features, their availability across volume types, and the release in which they appeared. If your FlexVol volume uses any feature that FlexGroup volumes do not offer in ONTAP 9.7, carefully decide whether that feature is necessary. For more supported feature lists, see [TR-4571: NetApp ONTAP FlexGroup volumes](#).

Table 6) Data protection features: FlexVol volumes versus FlexGroup volumes.

Data protection feature	FlexVol availability	FlexGroup availability
Snapshot copies	Yes	Yes (ONTAP 9.1)
SnapMirror (XDP/logical)	Yes	Yes (ONTAP 9.1)
SnapVault/Unified SnapMirror	Yes	Yes (ONTAP 9.3)
NDMP	Yes	Yes (ONTAP 9.7)
NDMPcopy	Yes	Yes (ONTAP 9.7)
NetApp MetroCluster	Yes	Yes (ONTAP 9.6)
SnapMirror Synchronous (SM-S)	Yes (ONTAP 9.5)	No

Data protection feature	FlexVol availability	FlexGroup availability
SnapLock	Yes (ONTAP 9.4)	No
Storage virtual machine disaster recovery (SVM DR)	Yes (Data ONTAP 8.3.1)	Yes (ONTAP 9.9.1)
Load-sharing mirrors (LS) (data I/O)	Yes (deprecated in ONTAP 9.5; use NetApp FlexCache® instead)	No
SnapMirror (block/DP)	Yes	No
SnapMirror to AltaVault (deprecated)	Yes	No
NetApp SolidFire to ONTAP SnapMirror	Yes	No
Cascading SnapMirror relationships	Yes	Yes (ONTAP 9.9.1)
SnapMirror fan-out	Yes	Yes (ONTAP 9.9.1)
SnapMirror to tape (SMTape)	Yes	No
Qtree SnapMirror	7-Mode only	No
1,023 Snapshot copies	Yes (ONTAP 9.4)	Yes (ONTAP 9.8)
Snapshot names/autodelete	Yes	No
SnapMirror to Simple Storage Service (S3)	Yes (ONTAP 9.8)	Yes (ONTAP 9.10.1; S3 to S3 only)
SnapMirror Business Continuity	Yes (ONTAP 9.8)	No

Backing up FlexGroup volumes

NetApp ONTAP 9.7 introduces basic NDMP support for FlexGroup volumes. As a result, a new way to back up FlexGroup volumes is available.

The following are the supported backup methods for FlexGroup volumes.

- **NAS-based backups.**

One way to back up a FlexGroup volume is to use CIFS/SMB or NFS protocols to copy files. However, in high-file-count environments, this approach can take a long time and can add undue stress on a cluster as the backup utility crawls the files and metadata. Therefore, if you are using a NAS-based backup solution, either back up the file system during off-peak hours or else run the backups on a replica of the file system. For example, run the backups on a NetApp SnapMirror and SnapVault destination. Some backup vendors provide interaction with ONTAP APIs that can make backups faster and more efficient. Check with your backup vendor for information about the level of support provided with ONTAP systems and FlexGroup volumes.

- **SnapMirror and SnapVault.**

As previously mentioned, FlexGroup volumes support both SnapMirror and SnapVault replication technologies. SnapMirror is more suitable for disaster recovery. NetApp SnapVault is used to keep asynchronous NetApp Snapshot copies on a destination system and is a better fit for a backup solution. ONTAP uses the same license and replication engine for both technologies, making management of relationships simpler.

- **NDMP-based backups.**

NDMP with FlexGroup volumes

If you are running ONTAP 9.7 or later, you can perform backups of a FlexGroup volume over NDMP.

FlexGroup volumes support the same NDMP versions and topologies (such as local, three-way, and remote) as FlexVol volumes, including NDMP copy.

Supported NDMP functionality

ONTAP 9.7 does not support all NDMP features. ONTAP 9.8 introduces support for the following NDMP features for use with FlexGroup volumes and can be considered to have feature parity with FlexVol volumes (with a few [exceptions](#)):

- Restartable Backup Extensions (RBE)
- EXCLUDE
- MULTI_SUBTREE_NAMES
- IGNORE_CTIME_MTIME
- Single file restores

In addition, support for per-qtree exclude lists is added in ONTAP 9.8. This support is a configurable option in ONTAP in the `vserver services ndmp modify` command.

Listed following are the available options for NDMP in advanced privilege level.

```
cluster::*> vserver services ndmp modify
Usage:
  [-vserver] <vserver name>                                Vserver
  [-ignore-ctime-enabled {true|false} ]                    Ignore Ctime
  [-offset-map-enabled {true|false} ]                       Enable Offset Map
  [-tcpnodelay {true|false} ]                               Enable TCP Nodelay
  [-tcpwinsize <integer> ]                                  TCP Window Size
  [-data-port-range <text> ]                                 Data Port Range
  [-backup-log-enabled {true|false} ]                       Enable Backup Log
  [-per-qtree-exclude-enabled {true|false} ]                Enable per Qtree Exclusion
  [-authtype <NDMP Authentication types>, ... ]            Authentication Type
  [-debug-enabled {true|false} ]                            *Enable Debug
  [-debug-filter <text> ]                                    *Debug Filter
  [-dump-logical-find <text> ]                               *Enable Logical Find for Dump
  [-abort-on-disk-error {true|false} ]                      *Enable Abort on Disk Error
  [-fh-dir-retry-interval <integer> ]                      *FH Throttle Value for Dir
  [-fh-node-retry-interval <integer> ]                      *FH Throttle Value for Node
  [-restore-vm-cache-size <integer> ]                      *Restore VM File Cache Size
  [-enable {true|false} ]                                   Enable NDMP on Vserver
  [-preferred-interface-role {cluster|data|node-mgmt|intercluster|cluster-mgmt}, ... ] Preferred Interface Role
  [-secondary-debug-filter <text> ]                        *Secondary Debug Filter
  [-is-secure-control-connection-enabled {true|false} ]    Is Secure Control Connection Enabled
```

NDMP multistream support

Both FlexGroup and FlexVol volumes support multiple NDMP backup streams. However, some backup vendors implement multiple streams in nonstandard or unconventional ways. For example, if the NDMP DUMP_DATE is interpreted in a nonstandard way by the backup application, then multistream backups fail for FlexGroup volumes. However, multistream backups work for FlexVol volumes due to the difference in the way the volumes handle DUMP_DATE. Be sure to check with your backup vendor for details and support with FlexGroup volumes and NDMP multistream backups.

Incremental NDMP backup support

FlexGroup volumes support incremental backups by using NDMP but there is a limit of 31 incrementals allowed per FlexGroup volume before another regular full backup is needed. This limit is based on the desire to keep NDMP restore windows smaller and more manageable and is the same limit as FlexVol volumes.

NDMP performance

With NDMP, you can expect the same general performance as for NetApp FlexVol volumes (depending on file size/count) but remember that NDMP itself can be a slow process—particularly with high file count environments. If you plan to use NDMP with a FlexGroup volume, it might be more prudent to back up at the folder or qtree level rather than backing up an entire FlexGroup volume, to adhere to the backup window.

ONTAP 9.8 improves NDMP performance over ONTAP 9.7. Table 7, Table 8, Table 9, and Table 10 each show performance comparisons between FlexVol volumes and FlexGroup volumes, as well as between ONTAP 9.7 and ONTAP 9.8 releases.

Table 7) NDMP dump performance — ONTAP 9.7 versus ONTAP 9.8.

Dataset	FlexGroup — L0 ONTAP 9.7 (MBps)	FlexGroup — L0 ONTAP 9.8 (MBps)	Delta	FlexGroup — L1 ONTAP 9.7 (MBps)	FlexGroup — L1 ONTAP 9.8 (MBps)	Delta
4K (20m)	250	375	+50%	63	81	+28.6%
64K (10m)	472	584	+23.7%	319	353	+10.6%
512K (0.5m)	488	593	+21.5%	434	513	+18.2%
1G (500)	480	593	+23.5%	462	582	+25.9%

Table 8) NDMP restore performance — ONTAP 9.7 versus ONTAP 9.8.

Dataset	FlexGroup — L0 ONTAP 9.7 (MBps)	FlexGroup — L0 ONTAP 9.8 (MBps)	Delta	FlexGroup — L1 ONTAP 9.7 (MBps)	FlexGroup — L1 ONTAP 9.8 (MBps)	Delta
4K (20m)	36	37	+3%	1	2	+100%
64K (10m)	131	148	+12%	15	19	+27%
512K (0.5m)	514	540	+5%	99	127	+28%
1G (500)	159	327	105%	154	304	+97%

Table 9) NDMP dump performance — ONTAP 9.8: FlexGroup versus FlexVol.

Dataset	FlexVol — L0 (MBps)	FlexGroup — L0 (MBps)	Delta	FlexVol — L1 (MBps)	FlexGroup — L1 (MBps)	Delta
4K (20m)	286	250	-13%	113	63	-45%
64K (10m)	596	472	-21%	439	319	-28%
512K (0.5m)	618	488	-21%	584	434	-26%
1G (500)	555	480	-14%	545	462	-15%

Table 10) NDMP Restore Performance — ONTAP 9.8: FlexGroup versus FlexVol.

Dataset	FlexVol — L0 (MBps)	FlexGroup — L0 (MBps)	Delta	FlexVol — L1 (MBps)	FlexGroup — L1 (MBps)	Delta
4K (20m)	19	36	+89%	1	1	-

Dataset	FlexVol — L0 (MBps)	FlexGroup — L0 (MBps)	Delta	FlexVol — L1 (MBps)	FlexGroup — L1 (MBps)	Delta
64K (10m)	114	131	+15%	14	15	+7%
512K (0.5m)	428	514	+20%	97	99	+2%
1G (500)	250	159	-36%	243	154	+97%

NDMPCopy example

In the following example, `ndmcopy` was used to migrate about 5 million folders and files from a FlexVol volume to a FlexGroup volume. The process took about 51 minutes:

```
cluster::*> system node run -node ontap9-tme-8040-01 ndmcopy -sa ndmpuser:AcDjtsU827tputjN -da
ndmpuser:AcDjtsU827tputjN 10.x.x.x:/DEMO/flexvol/nfs 10.x.x.x:/DEMO/flexgroup_16/ndmcopy
Ndmpcopy: Starting copy [ 2 ] ...
Ndmpcopy: 10.x.x.x: Notify: Connection established
Ndmpcopy: 10.x.x.x: Notify: Connection established
Ndmpcopy: 10.x.x.x: Connect: Authentication successful
Ndmpcopy: 10.x.x.x: Connect: Authentication successful
Ndmpcopy: 10.x.x.x: Log: Session identifier: 12584
Ndmpcopy: 10.x.x.x: Log: Session identifier: 12589
Ndmpcopy: 10.x.x.x: Log: Session identifier for Restore : 12589
Ndmpcopy: 10.x.x.x: Log: Session identifier for Backup : 12584
Ndmpcopy: 10.x.x.x: Log: DUMP: creating "/DEMO/flexvol/./snapshot_for_backup.1" snapshot.
Ndmpcopy: 10.x.x.x: Log: DUMP: Using subtree dump
Ndmpcopy: 10.x.x.x: Log: DUMP: Using snapshot_for_backup.1 snapshot
Ndmpcopy: 10.x.x.x: Log: DUMP: Date of this level_0 dump snapshot: Thu Jan  9 11:53:18 2020.
Ndmpcopy: 10.x.x.x: Log: DUMP: Date of last level 0 dump: the epoch.
Ndmpcopy: 10.x.x.x: Log: DUMP: Dumping /DEMO/flexvol/nfs to NDMP connection
Ndmpcopy: 10.x.x.x: Log: DUMP: mapping (Pass I)[regular files]
Ndmpcopy: 10.x.x.x: Log: DUMP: Reference time for next incremental dump is : Fri Jun 21 16:41:27
2019
Ndmpcopy: 10.x.x.x: Log: DUMP: mapping (Pass II)[directories]
Ndmpcopy: 10.x.x.x: Log: DUMP: estimated 12524018 KB.
Ndmpcopy: 10.x.x.x: Log: DUMP: dumping (Pass III) [directories]
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:05:07 2020: Begin level 0 restore
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:05:09 2020: Reading directories from the backup
Ndmpcopy: 10.x.x.x: Log: DUMP: dumping (Pass IV) [regular files]
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:09:37 2020: Creating files and directories.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:10:04 2020 : We have processed 58223 files and
directories.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:15:04 2020 : We have processed 850477 files and
directories.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:20:04 2020 : We have processed 1821373 files and
directories.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:25:04 2020 : We have processed 2810141 files and
directories.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:30:04 2020 : We have processed 3807403 files and
directories.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:35:04 2020 : We have processed 4814787 files and
directories.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:38:41 2020: Writing data to files.
Ndmpcopy: 10.x.x.x: Log: DUMP: Thu Jan  9 12:38:41 2020 : We have written 1597813 KB.
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:40:04 2020 : We have read 4215061 KB from the
backup.
Ndmpcopy: 10.x.x.x: Log: DUMP: Thu Jan  9 12:43:41 2020 : We have written 10995860 KB.
Ndmpcopy: 10.x.x.x: Log: ACL_START is '11842836480'
Ndmpcopy: 10.x.x.x: Log: RESTORE: Thu Jan  9 12:44:00 2020: Restoring NT ACLs.
Ndmpcopy: 10.x.x.x: Log: DUMP: dumping (Pass V) [ACLs]
Ndmpcopy: 10.x.x.x: Log: DUMP: Debug: 11566072 KB
Ndmpcopy: 10.x.x.x: Log: DUMP: DUMP IS DONE
Ndmpcopy: 10.x.x.x: Log: DUMP: Deleting "/DEMO/flexvol/./snapshot_for_backup.1" snapshot.
Ndmpcopy: 10.x.x.x: Log: DUMP_DATE is '5856116983'
Ndmpcopy: 10.x.x.x: Notify: dump successful
Ndmpcopy: 10.x.x.x: Log: RESTORE: RESTORE IS DONE
```

```
Ndmppcopy: 10.x.x.x: Notify: restore successful
Ndmppcopy: Transfer successful [ 0 hours, 50 minutes, 53 seconds ]
Ndmppcopy: Done
```

The same dataset using `cp` over NFS took 316 minutes—six times as long as `ndmppcopy`:

```
# time cp -R /flexvol/nfs/* /flexgroup/nfscp/

real    316m26.531s
user    0m35.327s
sys     14m8.927s
```

Using the NetApp XCP, that dataset took just under 20 minutes—or approximately 60% faster than `ndmppcopy`:

```
# xcp copy 10.193.67.219:/flexvol/nfs 10.193.67.219:/flexgroup_16/xcp
Sending statistics...
5.49M scanned, 5.49M copied, 5.49M indexed, 5.60 GiB in (4.81 MiB/s), 4.55 GiB out (3.91 MiB/s),
19m52s.
```

Note: This XCP copy was performed on a VM with a 1GB network and not much RAM or CPU; more robust servers perform even better.

Customer success stories: Backup repositories

FlexGroup volumes are being used in various ways. A common use case for FlexGroup volumes is using one as a backup repository. Because of its ability to spread performance and capacity across cluster nodes and to provide a viable single namespace, a FlexGroup volume provides an excellent way to store archival data.

The following sections describe two customer environments in which FlexGroup volumes are playing a pivotal role in providing a high-performing, resilient backup target.

Use case 1: Backing up Oracle RMAN databases

This customer is running a large Oracle database on AIX and wants to migrate the database to PostgreSQL SQL Server in the cloud to reduce maintenance and storage costs. Some of the challenges this customer is facing for this migration are:

- **Immense capacity needs.** The database is 890TB (and growing about 10TB per month) and must be staged to a single namespace.
- **Converting from a legacy application to a modern application.** The source database is running on AIX, which is [big endian](#). The target application runs little endian. The database must be dumped, converted, and then moved to the cloud.
- **Must remain online during migration.** The database must keep running during the migration, with a small cutover window. Incremental updates of changes are needed as the conversion/migration occurs.

Backup, conversion, and migration steps

The following steps occur when completing the migration task of a database that is larger than 800TB. They are also shown in Figure 9 and Figure 10.

- The AIX server performs an Oracle Recovery Manager (RMAN) backup of the database to an NFS file system. The NFS file system is a FlexGroup volume. You must use a FlexGroup volume because RMAN needs one mount point to direct the RMAN dump, and it needs to support more than 100TB. FlexGroup works well because RMAN can generate output in parallel to several target files that are spread across four aggregates on two controllers. Transfer speed is important, and FlexGroup helps by using parallel operations and load balancing across nodes and aggregates.

- A Linux server performs the big endian to little endian conversion; an NFS mount is required. The conversion ends up as a little-endian database on an FC LUN.
- The converted database is then replicated to a second NetApp array (array 2). After the initial dump of the database has been confirmed and fully replicated to array 2, array 2 is shipped to an Equinix data center and direct-connected to AWS.
- The process repeats several times, except it is RMAN level-1 dumps (incrementals) that are dumped, converted, and replicated to array 2 in Equinix. A final incremental is processed, and the database is shut down on the AIX system and brought up on the PostGRES SQL Server system in AWS.

Figure 9) Workflow for Oracle RMAN backup to FlexGroup volume, conversion, migration to cloud.

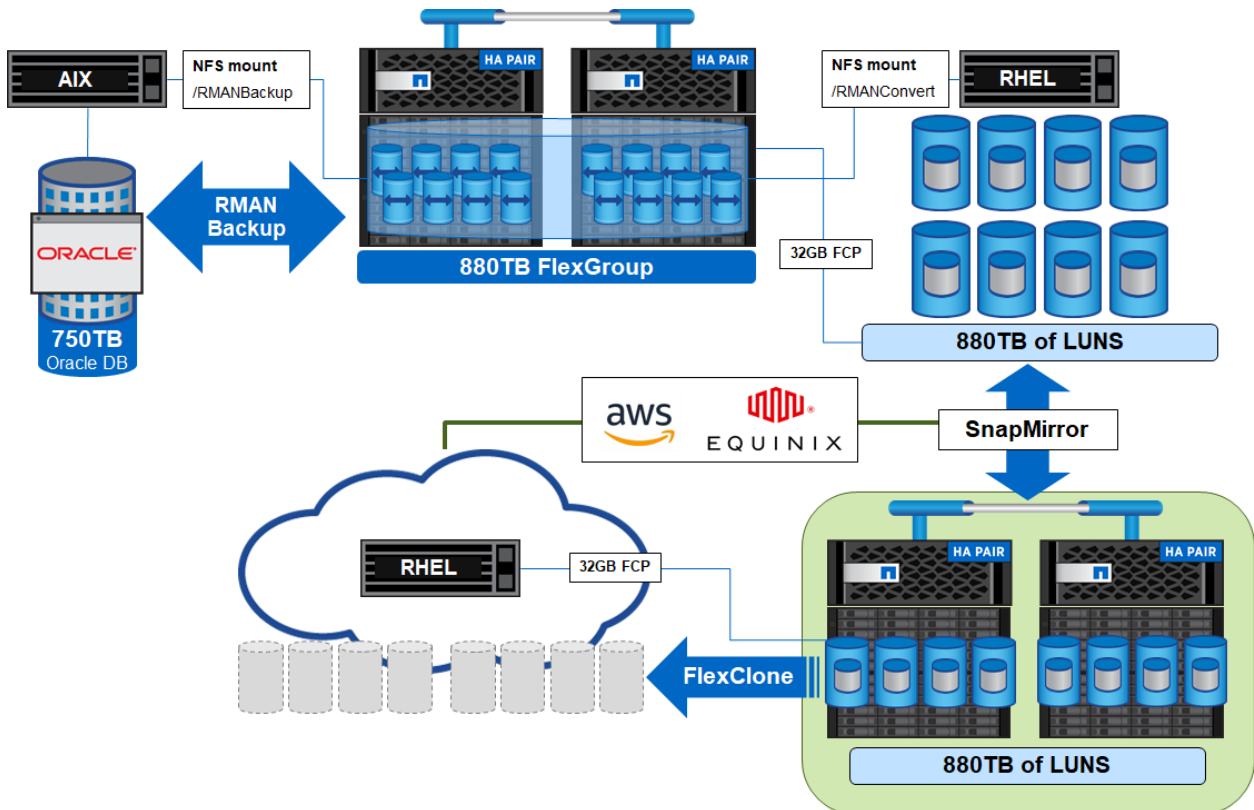



Figure 10) FlexGroup volume design.



Use case 2: Backing up SQL Server databases


A customer wanted to perform compressed backups of 5,000 Microsoft SQL Servers over SMB. This test was done with approximately 200 servers to vet out the solution, with a slow ramp-up over the course of a few months.

The customer not only wanted the database to be a backup target, but also wanted to replicate it to a disaster recovery site by using NetApp SnapMirror for extra data protection.

Each site has a six-node NetApp FAS8200 cluster (as seen in Figure 11) running ONTAP 9.4 using 6TB NL-SAS encrypted drives. Each cluster holds 3PB of usable capacity. The clusters use 30 FlexGroup volumes and use qtrees within the volumes for data organization.

The FlexGroup volumes themselves are 64TB each and the member volumes are 2.6TB apiece, with four members per node across six nodes (24 total members per FlexGroup volume).

Figure 11) SQL Server backup environment.



Results

This customer needed a single namespace that could collect about 150TB worth of SQL backup data over a 12-hour period. That rate is about 12TB/hour at about 3.5GB per second.

During testing, NetApp used 222 servers at site A and 171 servers at site B. During the test, each cluster's CPU was at 95% utilization and the backup jobs (sequential writes) were able to accomplish about 8.4GB per second. This amount was about 2.4 times the amount of throughput the job needed, which means the backups would complete in about 5 hours, rather than the 12-hour window. Also, this SMB workload was able to perform about 120,000 IOPS. Throughput and total operations during this test run can be seen in Figure 12. When more clients are added to this workload, we expect that the throughput maxes out at approximately 9GB per second.

Figure 12) Throughput and total operations during test runs.

cpu avg	cpu busy	total ops	nfs-ops	cifs-ops	foache ops	spin-ops	total recv	total sent	data busy	data recv	data sent	cluster busy	cluster recv	cluster sent	disk read	disk write	pkts recv	pkts sent
56%	81%	54530	0	54530	0	54420	6.16GB	2.65GB	44%	3.34GB	28.3MB	22%	2.82GB	2.62GB	128MB	3.31GB	968237	898917
65%	78%	70482	0	70482	0	70407	8.03GB	3.44GB	47%	4.33GB	30.9MB	24%	3.70GB	3.41GB	114MB	4.79GB	1178768	1102912
76%	87%	89725	0	89725	0	89105	10.2GB	4.30GB	49%	5.44GB	37.1MB	36%	4.76GB	4.26GB	157MB	5.54GB	1389743	1324559
86%	92%	111577	0	111577	0	110569	12.8GB	5.88GB	53%	6.84GB	41.5MB	31%	6.00GB	5.84GB	153MB	6.77GB	1724469	1679506
88%	92%	115036	0	115036	0	113509	13.2GB	6.44GB	51%	7.06GB	45.9MB	49%	6.14GB	6.40GB	142MB	7.65GB	1845760	1814640
92%	95%	118148	0	118148	0	117104	13.6GB	6.11GB	45%	7.26GB	49.9MB	42%	6.34GB	6.07GB	149MB	8.11GB	1802929	1769902
95%	98%	122953	0	122953	0	122123	14.3GB	7.10GB	47%	7.54GB	45.5MB	43%	6.75GB	7.06GB	134MB	8.29GB	1978205	1952416
96%	99%	126241	0	126241	0	125104	14.6GB	6.43GB	53%	7.75GB	54.3MB	44%	6.80GB	6.37GB	133MB	8.28GB	1865375	1849777
95%	97%	121948	0	121948	0	120719	13.9GB	7.25GB	44%	7.47GB	47.3MB	40%	6.41GB	7.20GB	108MB	8.30GB	1955908	1967271
95%	98%	123079	0	123079	0	121113	13.9GB	5.71GB	41%	7.56GB	49.0MB	38%	6.37GB	5.66GB	129MB	8.40GB	1761097	1712061
95%	97%	120567	0	120567	0	120493	13.7GB	7.01GB	42%	7.41GB	47.4MB	36%	6.34GB	6.96GB	114MB	8.48GB	1888934	1882711
95%	98%	119573	0	119573	0	119458	13.6GB	5.74GB	37%	7.33GB	44.4MB	35%	6.28GB	5.69GB	111MB	8.19GB	1702969	1671363
95%	97%	119538	0	119538	0	119829	13.5GB	6.98GB	41%	7.34GB	46.2MB	35%	6.17GB	6.93GB	120MB	8.44GB	1880298	1873821
95%	98%	118119	0	118119	0	118373	13.4GB	5.56GB	37%	7.25GB	45.4MB	37%	6.17GB	5.52GB	118MB	8.42GB	1666066	1630785
95%	98%	118862	0	118862	0	118327	13.6GB	6.29GB	39%	7.29GB	47.1MB	33%	6.30GB	6.24GB	114MB	8.31GB	1784134	1759266
96%	99%	121039	0	121039	0	121136	13.7GB	6.67GB	38%	7.44GB	44.5MB	34%	6.21GB	6.63GB	120MB	8.35GB	1832520	1827158
96%	99%	120852	0	120852	0	120920	13.7GB	5.77GB	39%	7.42GB	47.8MB	33%	6.24GB	5.72GB	111MB	8.51GB	1706939	1678778
94%	97%	119819	0	119819	0	120129	13.7GB	7.05GB	41%	7.36GB	42.4MB	35%	6.29GB	7.01GB	118MB	8.49GB	1882656	1877381

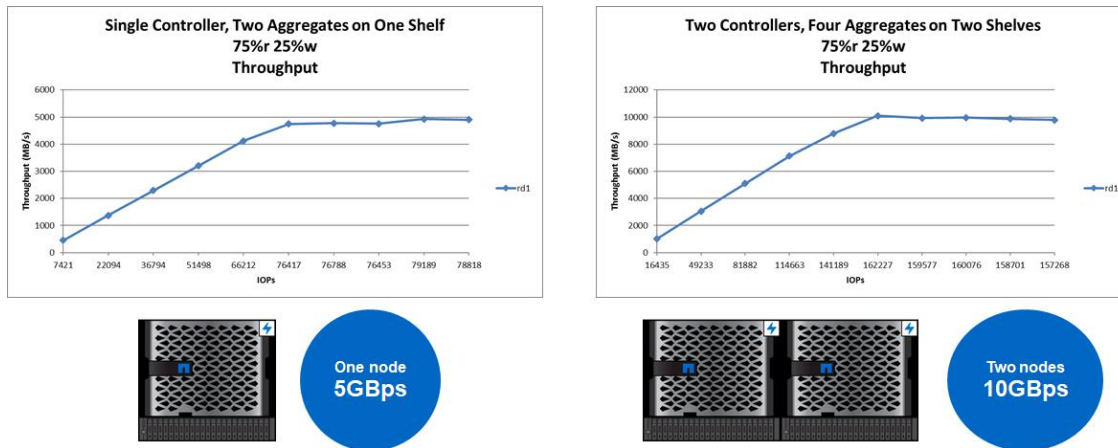
Data protection

In addition to the performance seen on the FlexGroup volume for the production workload, this customer was also able to achieve a high rate of transfer for the SnapMirror relationships between sites. This transfer rate of 8.4GB per second means that the replication window for a 150TB dataset is about 5.5 hours for the initial transfer. After that, the deltas should be able to complete well within the required transfer window, providing a solid disaster recovery plan for these SQL Server backups.

Scale-out performance

This six-node cluster was able to push over 8.4GB per second to a FlexGroup volume. In our customer proof of concept (CPOC) labs, we have seen near-linear performance gains by adding nodes to a cluster. The following graphs in Figure 13 show throughput results for a single-node NetApp AFF A700 and a two-node AFF A700.

Figure 13) CPOC scale-out throughput results.



If we wanted to add even more performance to our backup workload, we could add more nodes.

Lessons learned

During this backup architecture, some valuable lessons were learned that can help improve the experience even further.

- Stagger the workloads – Some NAS clients/protocol versions are unable to handle longer transfer times when thousands of other jobs are competing for resources. Adjust jobs to run in batches at different times and leverage ONTAP quality of service (QoS) for performance throttling.

- When possible, have the database file dumps break up over multiple smaller files for best data balance results across the FlexGroup volume.
- Leverage qtrees, when possible, for their reporting, quotas, and QoS functionality. (Qtree QoS is available in ONTAP 9.8 and later)
- Deploy the FlexGroup across as many homogenous nodes in the cluster as possible to leverage more of the cluster's hardware resources.

Conclusion

Not only is a FlexGroup volume great for small or high-file-count workloads such as electronic design automation (EDA) and software builds, it also can manage high throughput requirements for larger streaming files. It also reduces backup windows by scaling out storage across multiple nodes and applies all your cluster resources while maintaining performance, even with spinning disks.

Where to find additional information

To learn more about the information described in this document, refer to the following documents:

- TR-4015: SnapMirror configuration and best practices guide
www.netapp.com/us/media/tr-4015.pdf
- TR-4571: NetApp ONTAP FlexGroup volumes: Best practices and implementation guide
www.netapp.com/us/media/tr-4571.pdf

Version history

Version	Date	Document version history
Version 1.0	April 2018	First edition
Version 2.0	November 2018	ONTAP 9.5
Version 2.1	June 2019	ONTAP 9.6
Version 2.2	January 2020	ONTAP 9.7
Version 2.3	January 2021	ONTAP 9.8
Version 2.4	June 2021	ONTAP 9.9.1
Version 2.5	October 2021	ONTAP 9.10.1

Contact us

Let us know how we can improve this technical report.

Contact us at doccomments@netapp.com.

Include TECHNICAL REPORT 4678 in the subject line.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2018-2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

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.

TR-4678-1021