

FUJITSU Storage
ETERNUS AX series All-Flash Arrays,
ETERNUS HX series Hybrid Arrays

Data Protection and Backup for ONTAP FlexGroup Volumes




Table of Contents

1. Overview	8
Document Scope	8
Data Protection Terminology	8
Data Protection Features with FlexGroup Volumes	10
2. Snapshot Copies with FlexGroup Volumes	11
How Snapshot Copies Work with FlexGroup	12
Snapshot Creation Failures	12
FlexGroup Snapshot Guidelines	13
3. SnapMirror and SnapVault with FlexGroup Volumes.....	14
SnapMirror or SnapVault?	15
How SnapMirror Works with FlexGroup	16
Creating a SnapMirror and SnapVault Relationship for a FlexGroup Volume	16
Basic General Steps for Creating a SnapMirror and SnapVault Relationship with FlexGroup (ONTAP 9.4 and Later)	16
Using "Protect" with FlexGroup Volumes in ONTAP System Manager	17
Peering a Cluster	17
Protecting a New FlexGroup Volume	20
Protecting an Existing FlexGroup Volume	22
SnapMirror in ONTAP 9.7	24
SnapMirror and SnapVault Considerations	25
SnapMirror Fan-Out and Cascading SnapMirror	25
SnapMirror Resync Considerations	25
SnapMirror Restore Considerations	26
Expanding a FlexGroup Volume/Adding New Member Volumes	26
FlexGroup SnapMirror Guidelines	27
How SVM Disaster Recovery Behaves When FlexGroup Volumes Are Present	27
4. MetroCluster	28

5. FlexVol to FlexGroup Conversion: Data Protection Considerations	29
Why to Convert a Volume to a FlexGroup Volume	29
When Not to Convert a FlexVol Volume	30
FlexVol Conversion: ONTAP Snapshot Considerations	31
FlexVol Conversion: SnapMirror Considerations	32
Steps for Converting a FlexVol Volume in a SnapMirror Relationship	32
Expanding the FlexGroup Volume in SnapMirror Relationships	33
Converting a FlexVol Volume in a SnapMirror Relationship: Example	33
Data Protection Feature Parity in FlexVol and FlexGroup Volumes: Considerations	37
6. Backing Up FlexGroup Volumes.....	38
7. FlexGroup Volumes as a Backup Repository: Customer Success Stories.....	40
Use Case 1: Backing Up Oracle RMAN Databases	40
Backup, Conversion, and Migration Steps	40
Use Case 2: Backing Up SQL Server Databases	42
Results	42
Data Protection	43
Scale-Out Performance	43
Conclusion	43

List of Figures

Figure 1	FlexGroup volumes	9
Figure 2	Snapshot copies in FlexGroup volumes.....	12
Figure 3	SnapMirror versus SnapVault with FlexGroup volumes	15
Figure 4	Protecting a FlexGroup volume during deployment	21
Figure 5	Protecting a FlexGroup volume during deployment (Advanced Options)	21
Figure 6	Viewing unprotected volumes	22
Figure 7	Creating manual filters for protection relationships in System Manager	22
Figure 8	Protecting an existing FlexGroup volume	23
Figure 9	Protecting an existing FlexGroup volume: protect volumes	23
Figure 10	Protecting an existing FlexGroup volume: validate	24
Figure 11	ONTAP System Manager 9.7 SnapMirror relationship setup.....	24
Figure 12	SnapMirror cascade and fan-out.....	25
Figure 13	Workflow for Oracle RMAN backup to FlexGroup, conversion, migration to cloud.....	41
Figure 14	FlexGroup volume design	41
Figure 15	SQL Server backup environment.....	42
Figure 16	Throughput and total operations during test runs.....	42
Figure 17	CPOC scale-out throughput results	43

List of Tables

Table 1	Data protection features: FlexGroup volumes.....	10
Table 2	Snapshot feature support with FlexGroup volumes	11
Table 3	SnapMirror feature support with FlexGroup volumes.....	14
Table 4	FlexGroup data protection minimums	14
Table 5	Member volume count considerations for FlexGroup SnapMirror relationships	27
Table 6	Data protection features: FlexVol volumes versus FlexGroup volumes.....	37

Preface

This document covers data protection and backup for ONTAP FlexGroup volumes. Topics include Snapshot copies, SnapMirror, and other data protection and backup solutions. For data protection best practices pertaining to FlexVol volumes, see "FUJITSU Storage ETERNUS AX/HX series SnapMirror Configuration and Best Practices Guide for ONTAP 9". For general FlexGroup volume information, see "FUJITSU Storage ETERNUS AX/HX series Technical Overview of ONTAP FlexGroup Volumes". For all other general data protection information, see "FUJITSU Storage ETERNUS AX/HX series Data Protection Power Guide" on the Fujitsu manual site.

Copyright 2021 FUJITSU LIMITED

First Edition
October 2021

Trademarks

Third-party trademark information related to this product is available at:
<https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html>

Trademark symbols such as ™ and ® are omitted in this document.

About This Manual

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, contact Fujitsu support.

Related Information and Documents

The latest information for the ETERNUS AX/HX is available at:
<https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html>

Document Conventions

■ Notice Symbols

The following notice symbols are used in this document:

Caution

Indicates information that you need to observe when using the ETERNUS AX/HX. Make sure to read the information.

Note

Indicates information and suggestions that supplement the descriptions included in this manual.

1. Overview

Document Scope

This document intends to cover data protection best practices, considerations, and other related items as they pertain to ONTAP FlexGroup volumes. This document is not intended to be used for FlexVol volume configurations, nor is it aimed at audiences using Data ONTAP operating in 7-Mode. Visit the [Fujitsu manual site](#) for links to documents covering FlexVol volume configurations.

Data Protection Terminology

This section defines key terminology that's 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 logical interfaces, 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 logical interfaces 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**

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. 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.

Figure 1 FlexGroup volumes



For more information regarding FlexGroup volumes, see ["FUJITSU Storage ETERNUS AX/HX series Technical Overview of ONTAP FlexGroup Volumes"](#).

- **Logical Directory Replication/Unified Replication**

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

- **MirrorVault**

A MirrorVault is a SnapMirror relationship that integrates SnapVault functionality into the replication, allowing asymmetric Snapshot copy counts on the source and destination. This replication is controlled by the MirrorAndVault policy in ONTAP.

- **SnapMirror**

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

- **SnapRestore**

SnapRestore is a licensable feature that lets you 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**

Volumes can be copied 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, allows 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

Data Protection Feature/Software	First Supported ONTAP Version
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
MetroCluster	ONTAP 9.6
NDMP	ONTAP 9.7
XCP Migration Tool	All versions of ONTAP
Qtree SnapMirror	N/A
SnapProtect	N/A
SnapManager	N/A
SnapCenter	N/A
SVM DR	N/A (for details, see "How SVM Disaster Recovery Behaves When FlexGroup Volumes Are Present" (page 27))
SnapMirror to Cloud Backup	N/A
SnapLock	N/A
SnapMirror Synchronous	N/A
SnapMirror (data protection/DP style)	N/A
SnapMirror fan-out	N/A
Cascading SnapMirror	N/A
Single-file SnapRestore (through ONTAP CLI)	N/A
Data Availability Services (DAS)	N/A
VMware vSphere/Site Recovery Manager	N/A
1,023 Snapshot support	N/A
SnapMirror to tape (SMTape)	N/A

2. Snapshot Copies with FlexGroup Volumes

Snapshot copies are point-in-time copies of a file system. ONTAP has supported Snapshot copies for decades.

When a 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 FlexVol volumes. Also, you can restore volumes or aggregates together by using the `snap restore` command (with a valid SnapRestore license). The `snap restore` command is available only at the diagnostic privilege level for FlexGroup 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

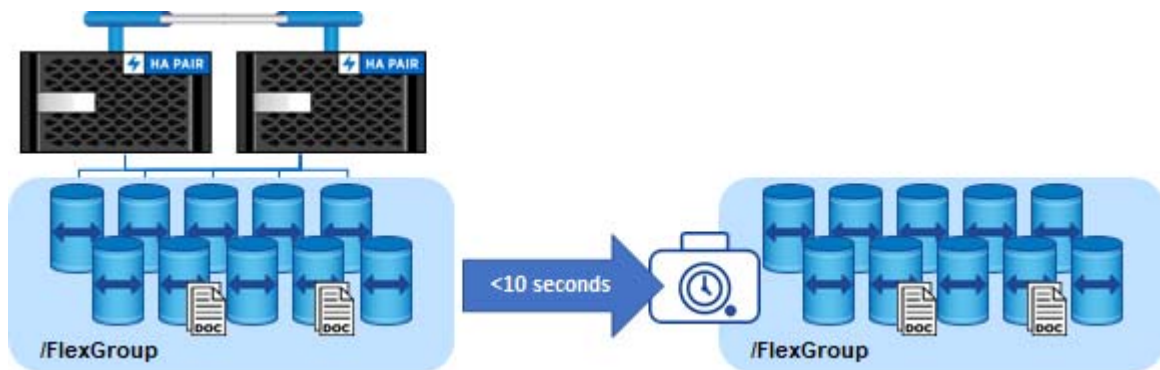
Snapshot Feature/Functionality	Supported?	ONTAP Version First Supported
Snapshot create	Yes	ONTAP 9.1
Snapshot restore (diagnostic level only)	Yes	ONTAP 9.1
Single file SnapRestore from CLI	No	N/A
Previous Versions tab	Yes	ONTAP 9.1
<code>.snapshot</code> directory access	Yes	ONTAP 9.1
Support for 1,023 Snapshot copies	No	N/A (FlexVol support in 9.4)

How Snapshot Copies Work with FlexGroup

ONTAP FlexGroup volumes are made up of a group of FlexVol member volumes. Snapshot copies are still made at the FlexVol volume level. When a FlexGroup Snapshot copy must be created, the member volumes must be coordinated to create a consistent Snapshot copy of the file system. Keep in mind that a FlexGroup volume makes generous use of remote hard links. Therefore, if a Snapshot copy is made of a member volume and the hard link isn't captured in flight, or if a hard link is captured without the associated file, the Snapshot copy is essentially no good.

As a result, 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 around one second, depending on load on the system. 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:

```
wafl.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.
- Make an attempt at keeping 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 can’t 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. Member FlexVol volumes cannot be restored 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 schedules should be in intervals of 30 minutes or more.

3. SnapMirror and SnapVault with FlexGroup Volumes

SnapMirror support for ONTAP FlexGroup volumes is present in ONTAP 9.1 and later. 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 rebaseline	Yes	ONTAP 9.3
SnapVault	Yes	ONTAP 9.3
Unified SnapMirror and MirrorVault	Yes	ONTAP 9.3
SnapMirror to Cloud Volumes ONTAP	Yes	ONTAP 9.6
SnapLock	No	N/A
Storage virtual machine disaster recovery (SVM DR) ("How SVM Disaster Recovery Behaves When FlexGroup Volumes Are Present" (page 27))	No	N/A
Load-sharing mirrors (LS) (data I/O)	No	N/A
SnapMirror (block/DP)	No	N/A
SnapMirror to AltaVault (deprecated)	No	N/A
SolidFire to ONTAP SnapMirror	No	N/A
Cascading SnapMirror relationships	No	N/A
SnapMirror Synchronous	No	N/A
SnapMirror fan-out	No	N/A
SnapMirror to tape (SMTape)	No	N/A
Qtree SnapMirror	No	Supported only in Data ONTAP operating in 7-Mode
Snapshot naming/autodelete	No	N/A
SnapMirror from FlexVol to FlexGroup	No	N/A
SnapMirror from FlexGroup to FlexVol	No	N/A

Table 4 FlexGroup data protection minimums


	Value	Value Type
SnapMirror schedule	30 minutes	Tested/recommended
Snapshot schedule	30 minutes	Tested/recommended

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 what exist on the source volume. SnapMirror and SnapVault use cases can be broken down as follows:

- **SnapMirror relationships** are intended for disaster recovery and provide an exact replica of the source volume, including the number of Snapshot copies on the source.
- **SnapVault relationships** are intended for backup and archive use cases by providing a path to Snapshot copies older than what might exist on the source volume.

Figure 3 SnapMirror versus SnapVault with FlexGroup volumes



How SnapMirror Works with FlexGroup

SnapMirror with FlexGroup operates 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. Before ONTAP 9.4, you primarily created such a relationship by using the command line. ONTAP 9.4 brought a plethora of simplicity enhancements for deploying and protecting FlexGroup volumes.

Basic General Steps for Creating a SnapMirror and SnapVault Relationship with FlexGroup (ONTAP 9.4 and Later)

Procedure ►►►

- 1 Make sure that the SVM and/or clusters are peered. See the SnapMirror documentation for more information about peering.
- 2 Create a destination volume with type "data protection" (DP) that has the same number of member volumes as the source volume and equal or greater capacity. For example, an eight-member source volume needs an eight-member destination.
- 3 Select a SnapMirror and SnapVault policy.
- 4 Create the SnapMirror relationship.
- 5 Initialize SnapMirror and SnapVault.
- 6 Select a Snapshot policy and label.
- 7 Create a schedule, a Snapshot policy, and Snapshot schedule.
- 8 Attach the schedule to the SnapMirror relationship.



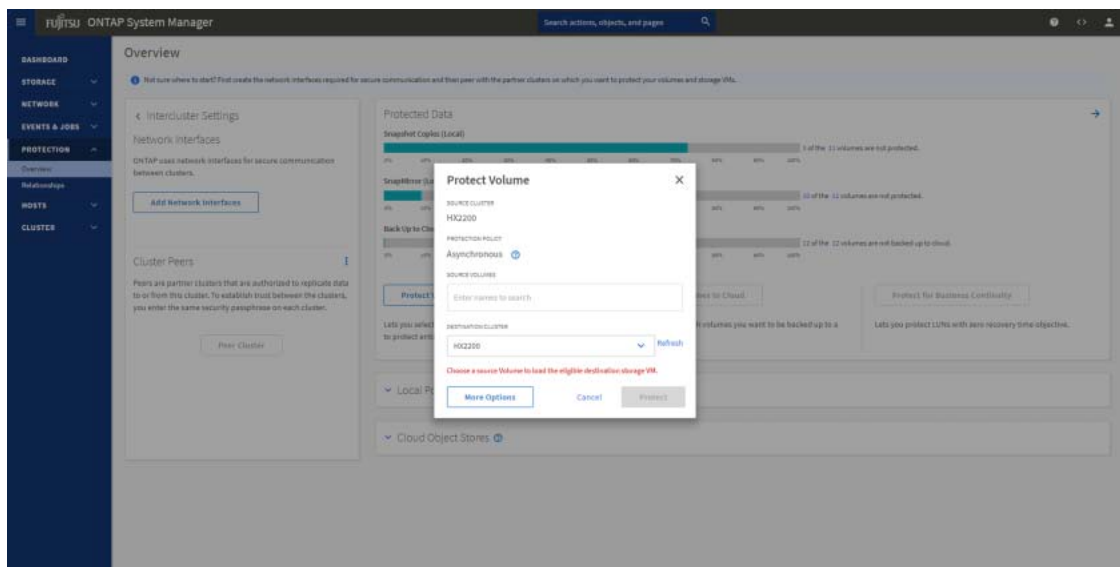
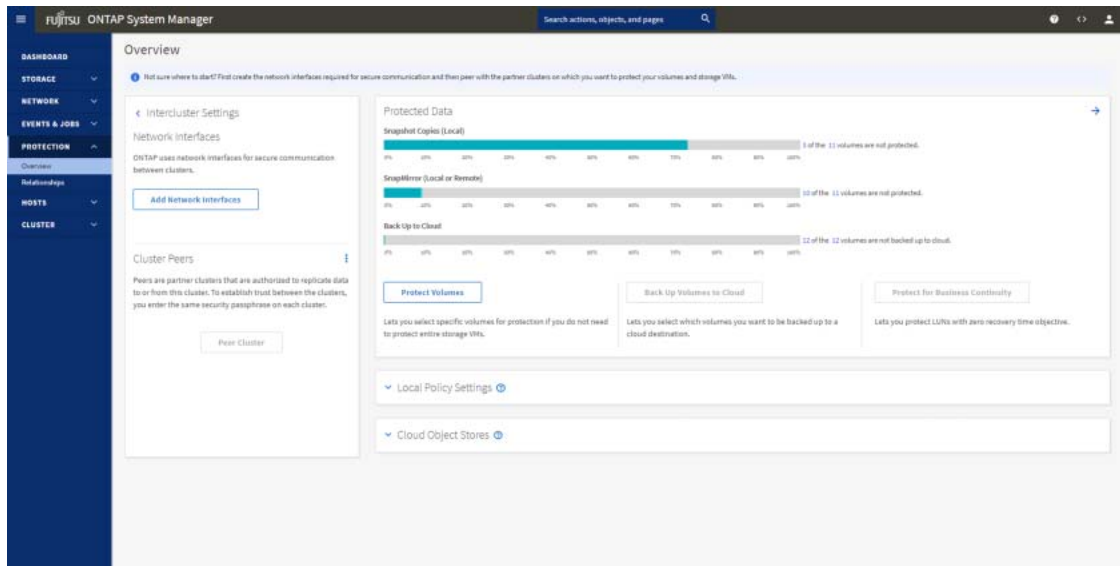
Using "Protect" with FlexGroup Volumes in ONTAP System Manager

ONTAP 9.4 introduced the ability to use the robust "protect" functionality of ONTAP System Manager (formerly OnCommand System Manager) for FlexGroup volumes. This functionality can be used at the time of the FlexGroup volume's deployment or after the FlexGroup volume has already been deployed.

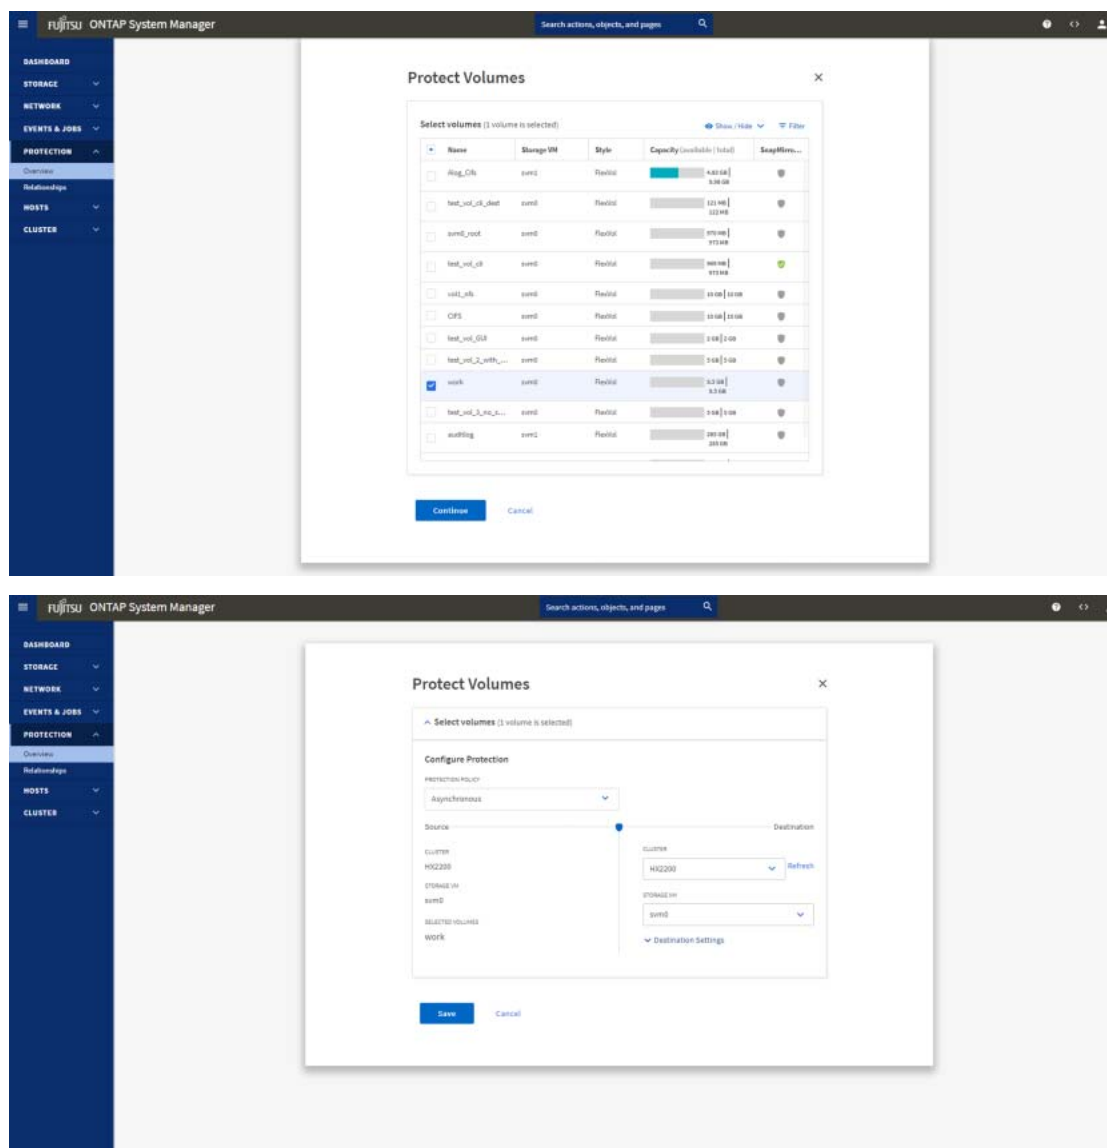
Peering a Cluster

Procedure ►►►

- 1 Click Configuration in the left menu, then SVM Peers (for local SnapMirror) or Cluster Peers (for remote SnapMirror) if you haven't already peered the source and destination. In this example, we're peering SVMs to do a local SnapMirror configuration.



3. SnapMirror and SnapVault with FlexGroup Volumes Using "Protect" with FlexGroup Volumes in ONTAP System Manager



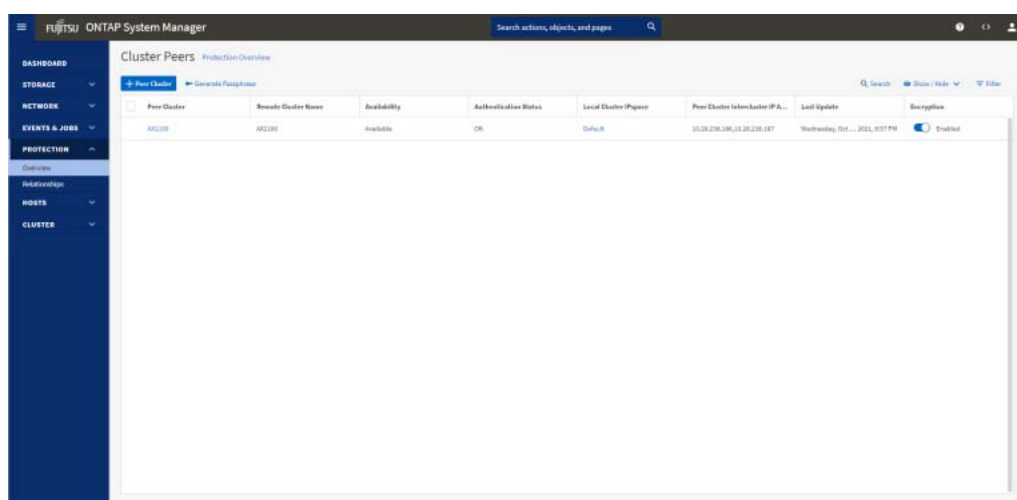
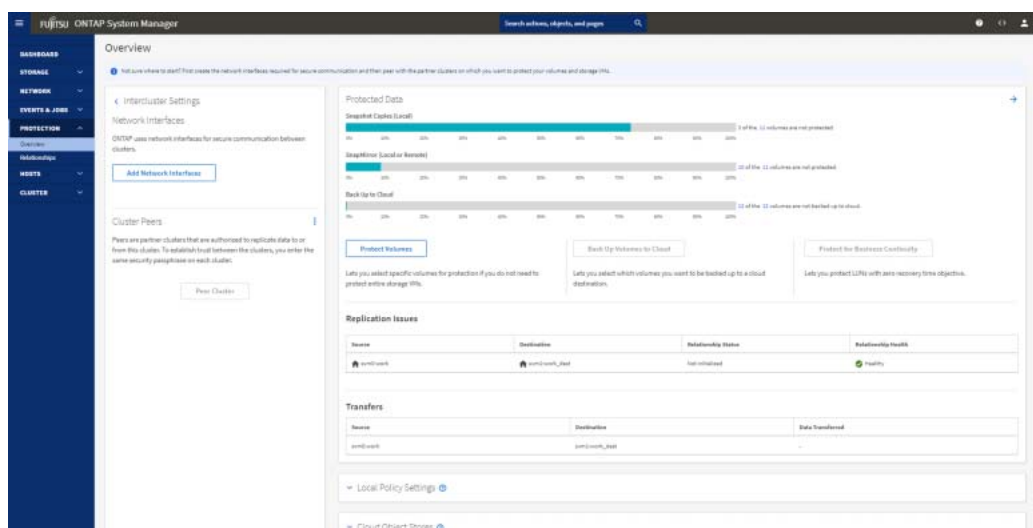
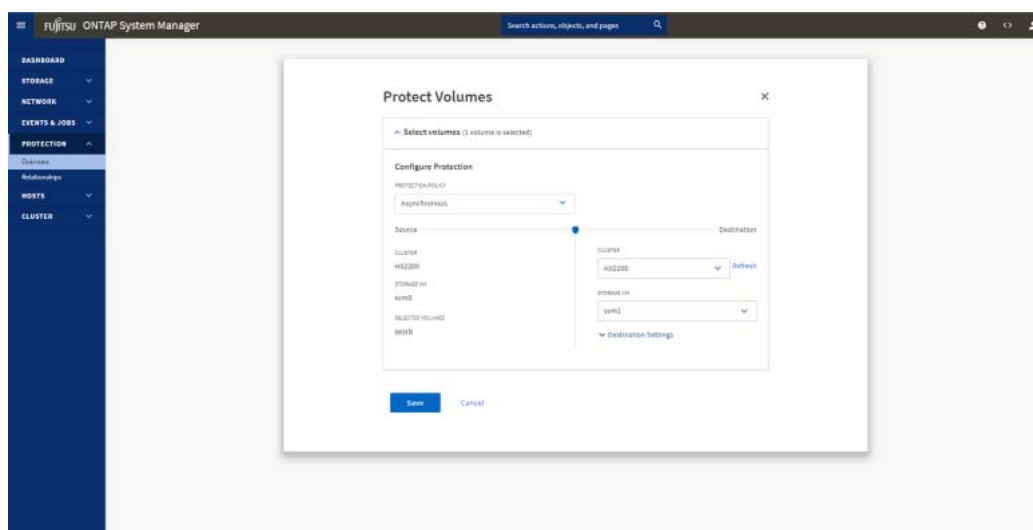
2 Peer the SVM or cluster.

■ SVM Peering

- (1) Click Create and choose your SVMs.
- (2) Click Initiate SVM Peering.
In a few seconds, you should see "SVM peering successful."

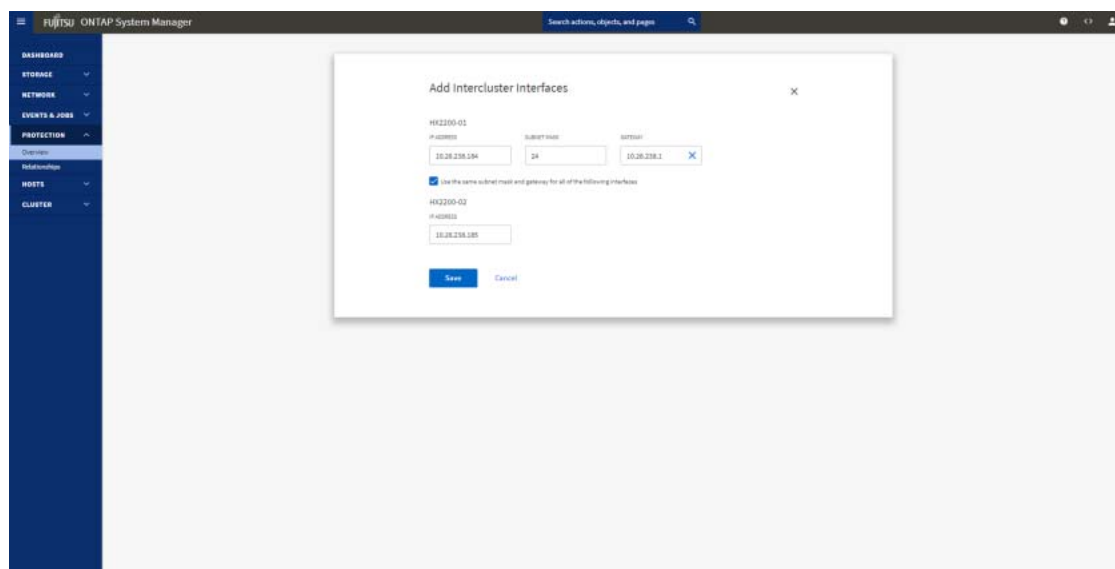
3. SnapMirror and SnapVault with FlexGroup Volumes Using "Protect" with FlexGroup Volumes in ONTAP System Manager

(3) Click Save.



■ Cluster Peering

If you plan to implement intercluster SnapMirror, cluster peering is needed. In System Manager for ONTAP 9.3 and later, this peering is done with "cluster peers." Doing this peering also allows you to peer SVMs in the same configuration steps.



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, provided the cluster has been peered successfully with another cluster.

Caution

Both clusters must have a SnapMirror license.

After you click Create, ONTAP System Manager creates the source and destination volumes with the correct size and geometries. If you leave the Initialize Protection field at the default value, the relationship also initializes automatically.

3. SnapMirror and SnapVault with FlexGroup Volumes Using "Protect" with FlexGroup Volumes in ONTAP System Manager

Figure 4 Protecting a FlexGroup volume during deployment

The screenshot shows the 'Add Volume' dialog in the ONTAP System Manager interface. The left sidebar contains navigation links for Dashboard, Storage, Applications, Volumes, LUNs, Shares, Buckets, Qtree, Quotas, Storage Vols, Tiers, Network, Events & Jobs, Protection, Hosts, and Cluster. The 'Add Volume' dialog has the following fields and options:

- Name:** Snap_Test
- Storage VSI:** svml
- ☐ Add as a cache for a remote volume. Simplifies file distribution, reduces WAN latency, and lowers WAN bandwidth costs.
- Storage and Optimization:**
 - Capacity:** 100 GB
 - Performance Service Level:** Performance
 - ☐ Distribute volume data across the cluster
- Access Permissions:**
 - ☒ Export via NFS

Figure 5 Protecting a FlexGroup volume during deployment (Advanced Options)

The screenshot shows the 'Add Volume' dialog in the ONTAP System Manager interface, displaying advanced protection options. The left sidebar is the same as in Figure 4. The 'Add Volume' dialog has the following fields and options:

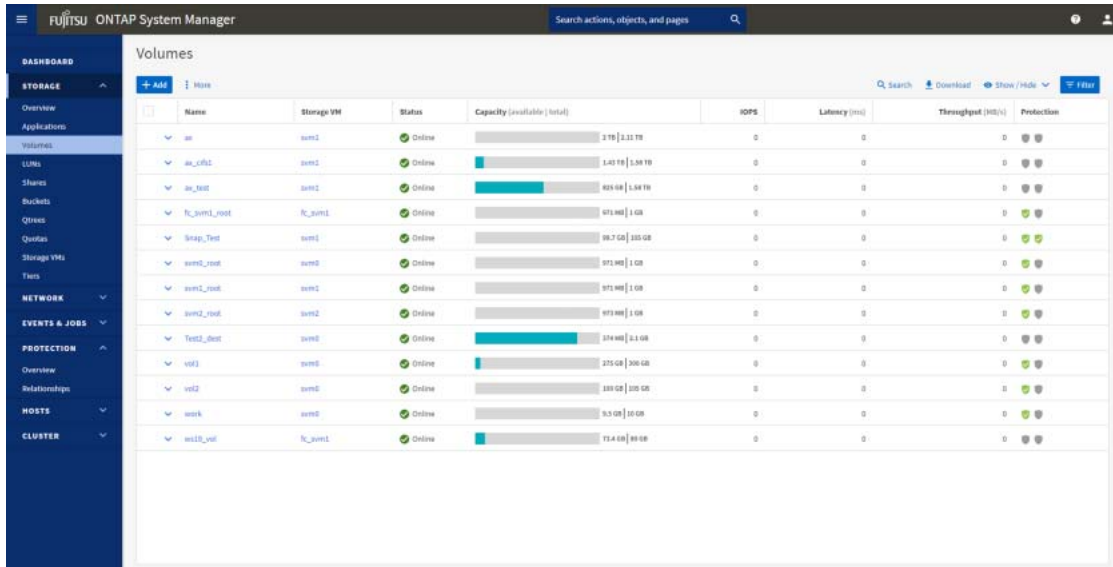
- ☒ Enable SnapMirror (Local or Remote)
- Protection Policy:** Asynchronous
- Source:**
 - Cluster:** AXC2100
 - Storage VSI:** svml
- Destination:**
 - Cluster:** HQ2100
 - Storage VSI:** svml
- Destination Settings:**
 - Volume Name:** vol_ <SourceVolumeName> _dest
 - Performance Service Level:** Auto
 - ☒ Enforce performance limit
- Configuration Details:**
 - ☒ Initialize relationship

3. SnapMirror and SnapVault with FlexGroup Volumes Using "Protect" with FlexGroup Volumes in ONTAP System Manager

Protecting an Existing FlexGroup Volume

If a FlexGroup volume has already been created without a SnapMirror or SnapVault relationship, you can use System Manager to query for unprotected volumes and then retroactively protect them. To query for unprotected volumes, click View Missing Protection Relationships on the volumes screen.

Figure 6 Viewing unprotected volumes

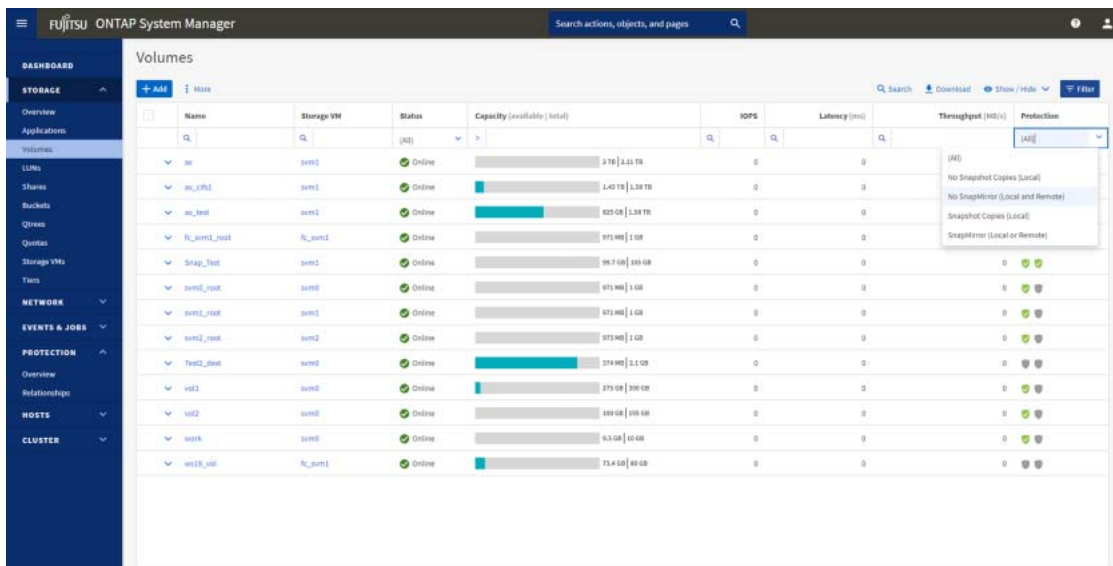


The screenshot shows the 'Volumes' page in the ONTAP System Manager. The table lists various volumes, including 'all', 'all_cgfs', 'all_test', 'fc_svm1_root', 'Snap_Test', 'svm2_root', 'svm1_root', 'svm2_root', 'Test2_root', 'vol1', 'vol2', 'work', and 'vol19_vol'. The 'Protection' column shows icons for protection status, with some volumes having a green shield icon and others having a red shield icon.

Name	Storage VM	Status	Capacity (available / total)	IOPS	Latency (ms)	Throughput (MB/s)	Protection
all	svm1	Online	3 TB / 3.33 TB	0	0	0	🛡️
all_cgfs	svm1	Online	1.43 TB / 1.38 TB	0	0	0	🛡️
all_test	svm1	Online	825 GB / 1.38 TB	0	0	0	🛡️
fc_svm1_root	fc_svm1	Online	975 MB / 1 GB	0	0	0	🛡️
Snap_Test	svm1	Online	99.7 GB / 100 GB	0	0	0	🛡️
svm2_root	svm2	Online	975 MB / 1 GB	0	0	0	🛡️
svm1_root	svm1	Online	975 MB / 1 GB	0	0	0	🛡️
svm2_root	svm2	Online	975 MB / 1 GB	0	0	0	🛡️
Test2_root	svm2	Online	374 MB / 1.1 GB	0	0	0	🛡️
vol1	svm2	Online	375 GB / 300 GB	0	0	0	🛡️
vol2	svm2	Online	389 GB / 300 GB	0	0	0	🛡️
work	svm2	Online	6.5 GB / 10 GB	0	0	0	🛡️
vol19_vol	fc_svm1	Online	75.4 GB / 80 GB	0	0	0	🛡️

This query filters out all volumes that show "Protection Relationship = yes." You can also accomplish this result by manually selecting No in the menu.

Figure 7 Creating manual filters for protection relationships in System Manager



The screenshot shows the 'Volumes' page with a dropdown menu open for the 'Protection' column. The menu options are: (All), No Snapshot Copies (Local), No SnapMirror (Local and Remote), Snapshot Copies (Local), and SnapMirror (Local or Remote).

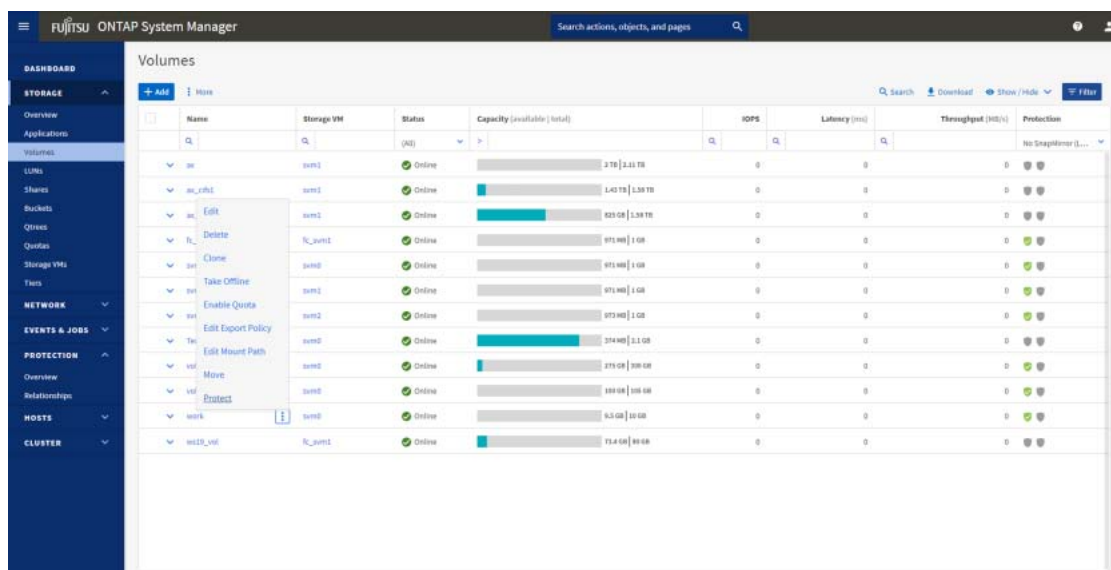
Name	Storage VM	Status	Capacity (available / total)	IOPS	Latency (ms)	Throughput (MB/s)	Protection
all	svm1	Online	3 TB / 3.33 TB	0	0	0	🛡️
all_cgfs	svm1	Online	1.43 TB / 1.38 TB	0	0	0	🛡️
all_test	svm1	Online	825 GB / 1.38 TB	0	0	0	🛡️
fc_svm1_root	fc_svm1	Online	975 MB / 1 GB	0	0	0	🛡️
Snap_Test	svm1	Online	99.7 GB / 100 GB	0	0	0	🛡️
svm2_root	svm2	Online	975 MB / 1 GB	0	0	0	🛡️
svm1_root	svm1	Online	975 MB / 1 GB	0	0	0	🛡️
svm2_root	svm2	Online	975 MB / 1 GB	0	0	0	🛡️
Test2_root	svm2	Online	374 MB / 1.1 GB	0	0	0	🛡️
vol1	svm2	Online	375 GB / 300 GB	0	0	0	🛡️
vol2	svm2	Online	389 GB / 300 GB	0	0	0	🛡️
work	svm2	Online	6.5 GB / 10 GB	0	0	0	🛡️
vol19_vol	fc_svm1	Online	75.4 GB / 80 GB	0	0	0	🛡️

To protect an existing FlexGroup volume, simply click the volume to highlight it and click More Actions. Then select Protect.

3. SnapMirror and SnapVault with FlexGroup Volumes

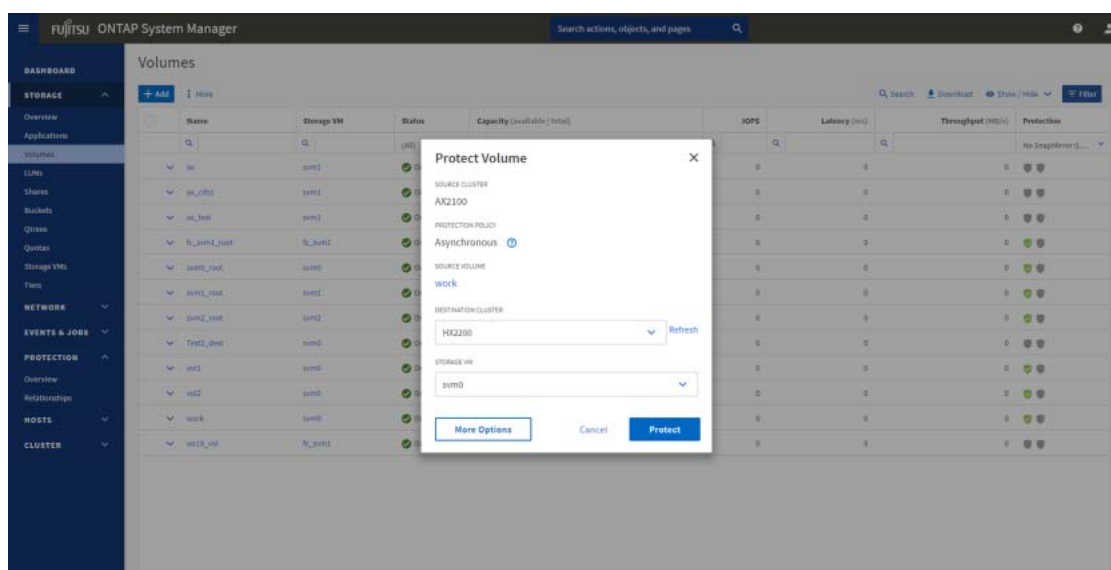
Using "Protect" with FlexGroup Volumes in ONTAP System Manager

Figure 8 Protecting an existing FlexGroup volume



This selection brings you to the Protect Volumes wizard. From here, specify your destination cluster and desired destination volume suffix name.

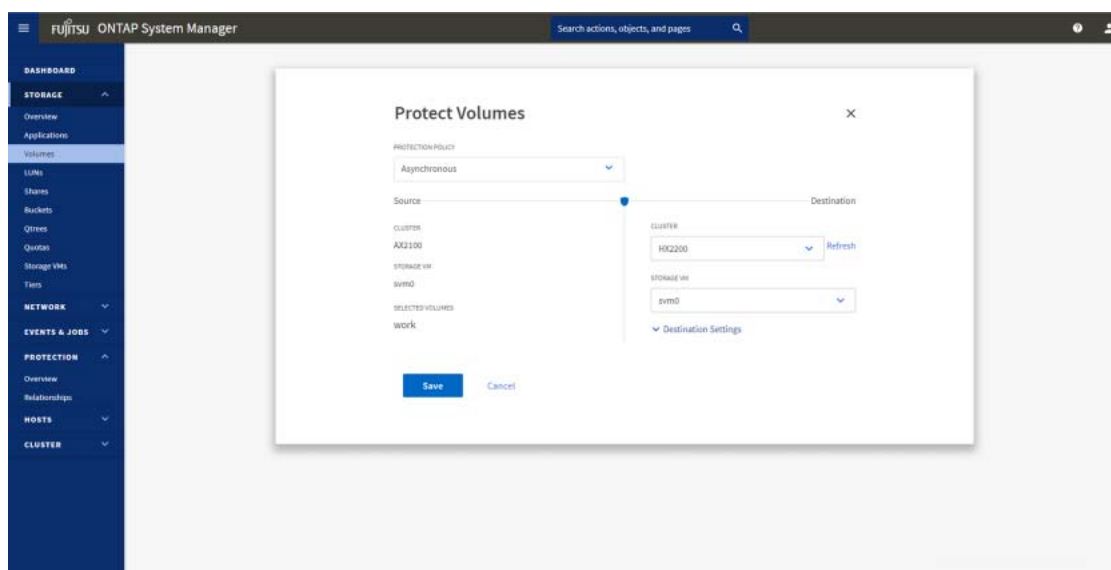
Figure 9 Protecting an existing FlexGroup volume: protect volumes



3. SnapMirror and SnapVault with FlexGroup Volumes Using "Protect" with FlexGroup Volumes in ONTAP System Manager

Then click Validate to check for labels.

Figure 10 Protecting an existing FlexGroup volume: validate



If all looks good, click Save to complete the operation.

SnapMirror in ONTAP 9.7

ONTAP 9.7 introduces a new look and feel for System Manager, designed to simplify day-to-day operations by reducing the number of clicks needed to perform tasks, as well as presenting a more logical layout.

In ONTAP 9.7's new view, the SnapMirror setup is in the same place as before—FlexGroup volume creation—but now the fields don't populate until the checkbox is selected.

Figure 11 ONTAP System Manager 9.7 SnapMirror relationship setup

Unchecked:

This screenshot shows the 'Protection' settings dialog. It has a title 'Protection'. Below it are two checkboxes: 'Enable Snapshot Copies (Local)' which is checked, and 'Enable SnapMirror (Local or Remote)' which is unchecked. At the bottom, there are 'Save' and 'Cancel' buttons.

Checked:

This screenshot shows the 'Protection' settings dialog with both checkboxes checked. The 'Enable SnapMirror (Local or Remote)' checkbox is now checked, and the fields below it are populated. The 'PROTECTION POLICY' dropdown is set to 'FullSync'. The 'Source' section shows 'CLUSTER' as 'ontap9-tme-8040' and 'STORAGE VM' as 'DEMO'. The 'Destination' section shows 'CLUSTER' as 'ontap9-tme-8040' and 'STORAGE VM' as 'DEMO'. There is a 'Refresh' button next to the destination cluster field. At the bottom, there is a 'Destination Settings' dropdown. 'Save' and 'Cancel' buttons are at the bottom.


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 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).

Figure 12 SnapMirror cascade and fan-out



These SnapMirror options are not currently supported with FlexGroup volumes.

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, which fences off read/write access 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 GUI) is not supported, but you can restore single files from clients by using the Previous Versions tab in CIFS/SMB or the `.snapshot` directory in NFS.
- 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

Volume expansion can be used 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.

FlexGroup SnapMirror Guidelines

- The guidelines for Snapshot copies also apply to SnapMirror. 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 using SnapMirror must have an equal number of member volumes on source and destination.
- 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 AX4100 all-flash storage system cluster can be mirrored to a single-node HX2000 series node as long as it has the same number of member volumes.
- For optimal results, 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 (*1)	N/A	50
Member volume count per node (multiple FlexGroup volumes) (*2)	N/A	500
Member volume count per cluster (all FlexGroup volumes) (*2)	100	6,000

*1: Exceeding this limit can affect RPO.

*2: Same limits as FlexVol volumes.

Caution

To request a higher limit, contact Fujitsu support.

How SVM Disaster Recovery Behaves When FlexGroup Volumes Are Present

If you are currently using SVM disaster recovery in your cluster, 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, the command fails with an error. If you want to use FlexGroup volumes with a cluster using SVM disaster recovery, create an SVM for the FlexGroup volumes.

4. MetroCluster

ONTAP 9.6 introduced support for FlexGroup volumes on MetroCluster deployments (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 ["FUJITSU Storage ETERNUS AX/HX series MetroCluster® IP Installation and Configuration Guide"](#).

5. FlexVol to FlexGroup Conversion: Data Protection Considerations

In ONTAP 9.7, it is now possible to convert a single FlexVol volume to a FlexGroup volume with a single member volume, in place, with less than 40 seconds disruption. This can be done 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 any other modifications that could create a maintenance window. After the FlexVol volume is converted to a FlexGroup volume, you can add new member volumes to expand the capacity.

Why to 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, perhaps you have a workload that is growing rapidly and you don't want to have to migrate the data, but still want to provide more capacity. Or perhaps a workload's performance isn't good enough on a FlexVol volume, so you want to provide better performance handling with a FlexGroup volume. Converting can help here.

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 aren't available in FlexGroup volumes, then you should hold off. For example, SVM-DR and cascading SnapMirror relationships aren't supported in ONTAP 9.7, so if you need them, you should stay with FlexVol volumes.

Also, if you have a FlexVol volume that's already very large (80–100TB) and already very full (80–90%), you might want to copy the data rather than convert, because the converted FlexGroup volume would have a very large, very full member volume. This could create performance issues and doesn't 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 would look like this:



If you were to convert this 90% full volume to a FlexGroup volume, you would have a 90% full member volume. If you add new member volumes, they would be 100TB each and 0% full, so they would take on a majority of new workloads. The data would not rebalance and if the original files grew over time, you could still run out of space with nowhere to go (because 100TB is the maximum member volume size).



The following sections pertain to FlexVol conversion only as it applies to data protection and disaster recovery ONTAP functionality.

FlexVol Conversion: ONTAP Snapshot Considerations

When a FlexVol volume has 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 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.

If a FlexVol volume has more than 255 Snapshot copies present, then FlexVol to FlexGroup conversion is blocked, because FlexGroup volumes do not support more than 255 Snapshot copies. Snapshot copies would need to be deleted to get below 255.

Additionally, 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 -vservers DEMO -volume flexvol -physical-snap-id >255 -fields physical-snap-id
```

vservers	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.

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

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

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

Procedure ►►►

- 1 Optional: Create a FlexClone copy of the volume to be converted, and split it to test FlexVol to FlexGroup conversion.
- 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: ONTAP Snapshot Considerations" \(page 31\)](#) and ["FlexVol Conversion: SnapMirror Considerations" \(page 32\)](#).
- 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

Once 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, so in most cases, 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". Additionally, 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. Here 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 Flex-
Group. 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 would need to convert the destination first. To do that, you must 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 "snap-
mirror quiesce -destination-path data_dst" and then try the conversion.
```

Now convert the volume.

```
cluster::*> snapmirror quiesce -destination-path DEMO:data_dst Operation suc-
ceeded: 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 lets you know 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-con-
version". 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 suc-
cessfully converted to a FlexGroup volume. You must now convert the relationship's
source volume, "DEMO:data", to a FlexGroup. Then, re-establish the SnapMirror rela-
tionship using the "snapmirror resync" command.
```

Now 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-conver-
sion". Pre-conversion Snapshot copies cannot be restored.
Do you want to continue? {y|n}: y
[Job 23712] Job succeeded: success
```

And 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
```

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

First, 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 is 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
```

5. FlexVol to FlexGroup Conversion: Data Protection Considerations

FlexVol Conversion: SnapMirror Considerations

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
```

And 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
```

Now, you 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 Flex-
Group "data": 4. Expanding the FlexGroup will cause the state of all Snapshot cop-
ies 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.
```

5. FlexVol to FlexGroup Conversion: Data Protection Considerations

FlexVol Conversion: SnapMirror Considerations

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: Considerations

Before converting a FlexVol volume to a FlexGroup volume, it is important to evaluate any and all ONTAP features that exist in FlexVol volumes but are currently unavailable in FlexGroup volumes. The following table lists data protection features, their availability across volume types, and the release they appeared in. If your FlexVol volume uses any feature that FlexGroup volumes do not offer in ONTAP 9.7, carefully decide whether that feature is necessary.

Table 6 Data protection features: FlexVol volumes versus FlexGroup volumes

Data Protection Feature	FlexVol Availability	FlexGroup Availability
Snapshot copies	Y	Y (ONTAP 9.1)
SnapMirror (XDP/logical)	Y	Y (ONTAP 9.1)
SnapVault/Unified SnapMirror	Y	Y (ONTAP 9.3)
NDMP	Y	Y (ONTAP 9.7)
NDMPCopy	Y	Y (ONTAP 9.7)
MetroCluster	Y	Y (ONTAP 9.6)
SnapMirror Synchronous (SM-S)	Y (ONTAP 9.5)	N
SnapLock	Y (ONTAP 9.4)	N
Storage virtual machine disaster recovery (SVM DR)	Y (Data ONTAP 8.3.1)	N
Load-sharing mirrors (LS) (data I/O)	Y (deprecated in ONTAP 9.5; use FlexCache instead)	N
SnapMirror (block/DP)	Y	N
SnapMirror to AltaVault (deprecated)	Y	N
SolidFire to ONTAP SnapMirror	Y	N
Cascading SnapMirror relationships	Y	N
SnapMirror fan-out	Y	N
SnapMirror to tape (SMTape)	Y	N
Qtree SnapMirror	7-Mode only	N
1,023 Snapshot copies	Y (ONTAP 9.4)	N
Snapshot names/autodelete	Y	N

6. Backing Up FlexGroup Volumes

ONTAP 9.7 introduced NDMP support for FlexGroup volumes. As a result, a new way to back up FlexGroup volumes is available. The following are 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 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. SnapVault is used to keep asynchronous 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 backups**

If you're running ONTAP 9.7 or later, you can perform backups of a FlexGroup volume over NDMP. You can expect the same general performance as for FlexVol volumes, but keep in mind 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, to adhere to the backup window.

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

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 ndm-
puser: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
Ndmpcopy: 10.x.x.x: Notify: restore successful
Ndmpcopy: Transfer successful [ 0 hours, 50 minutes, 53 seconds ]
Ndmpcopy: Done
```

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

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

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

Using the XCP Migration Tool, that dataset took just under 20 minutes—or around 60% faster than `ndmcopy`:

```
# xcp copy 10.193.67.219:/flexvol/nfs 10.193.67.219:/flexgroup_16/xcp Sending sta-
tistics...
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.
```

Caution

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

7. FlexGroup Volumes as a Backup Repository: Customer Success Stories

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 PostGRES 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 around 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.

- 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 13 Workflow for Oracle RMAN backup to FlexGroup, conversion, migration to cloud

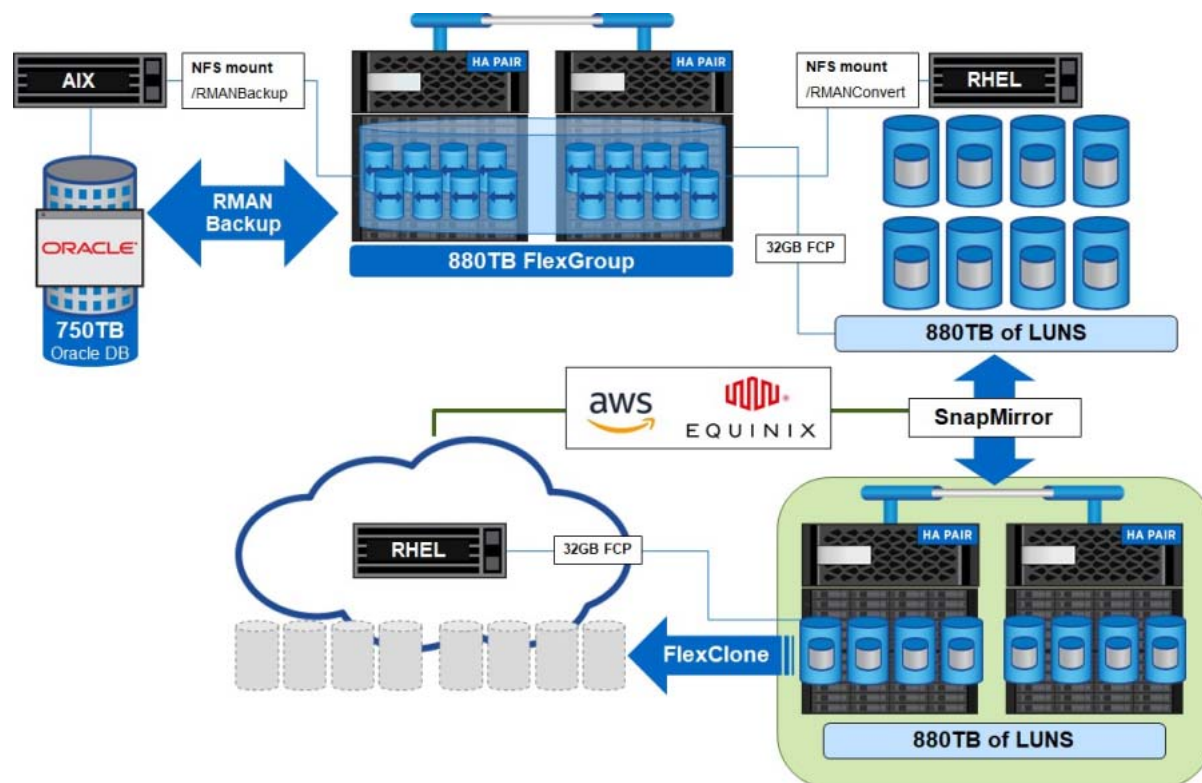



Figure 14 FlexGroup volume design



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 SnapMirror for extra data protection.

Each site uses 6TB near-line SAS (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 15 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's about 12TB/hour at about 3.5GB per second.

During testing, 222 servers were used 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, which is about 2.4 times the amount of throughput the job needed. That 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. When more clients are added to this workload, we expect that the throughput maxes out at around 9GB per second.

Figure 16 Throughput and total operations during test runs

avg		ops	total		fsocache		total		data		cluster		cluster		disk		pkts	
ops	busy	avg	nfs-ops	cifs-ops	ops	spin-ops	recv	sent	busy	recv	sent	busy	recv	sent	read	write	recv	sent
56%	81%	54530	0	54530	0	54420	6.16GB	2.65GB	44%	3.34GB	26.3MB	22%	2.82GB	3.42GB	12MB	3.31GB	968237	898917
65%	78%	70482	0	70482	0	70407	8.03GB	3.44GB	47%	4.39GB	30.5MB	24%	3.70GB	3.41GB	11MB	4.79GB	1178768	1102912
74%	87%	88725	0	88725	0	88105	10.2GB	4.30GB	49%	5.44GB	37.1MB	36%	4.78GB	4.26GB	15MB	5.56GB	1389743	1324559
86%	92%	111577	0	111577	0	110569	12.8GB	5.88GB	53%	6.84GB	41.9MB	31%	6.00GB	5.84GB	15MB	6.77GB	1726469	1679506
88%	92%	115036	0	115036	0	113509	13.2GB	6.44GB	54%	7.06GB	45.9MB	49%	6.14GB	6.40GB	14MB	7.65GB	1845740	1814549
92%	95%	118148	0	118148	0	117104	13.4GB	6.11GB	45%	7.26GB	48.5MB	42%	6.34GB	6.07GB	14MB	8.11GB	1802939	1769902
95%	98%	122953	0	122953	0	122123	14.3GB	7.10GB	47%	7.54GB	45.5MB	43%	6.75GB	7.04GB	13MB	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	13MB	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	10MB	8.30GB	1995908	1967271
95%	98%	123079	0	123079	0	121113	13.9GB	5.71GB	41%	7.56GB	49.0MB	38%	6.37GB	5.64GB	12MB	8.40GB	1741037	1712061
95%	97%	120567	0	120567	0	120493	13.7GB	7.01GB	42%	7.41GB	47.4MB	36%	6.34GB	6.96GB	11MB	8.49GB	1888934	1882711
95%	98%	119573	0	119573	0	119458	13.6GB	5.74GB	37%	7.39GB	44.4MB	35%	6.28GB	5.49GB	11MB	8.19GB	1702969	1671163
95%	97%	119538	0	119538	0	119829	13.5GB	6.88GB	41%	7.34GB	46.2MB	35%	6.17GB	6.93GB	12MB	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	11MB	8.42GB	1666066	1630785
95%	98%	118462	0	118462	0	118327	13.4GB	6.29GB	38%	7.29GB	47.1MB	33%	6.30GB	6.24GB	11MB	8.31GB	1784134	1759246
96%	99%	121039	0	121039	0	121136	13.7GB	6.47GB	38%	7.44GB	44.5MB	34%	6.21GB	6.43GB	12MB	8.35GB	1832520	1827158
96%	98%	120852	0	120852	0	120920	13.7GB	5.77GB	39%	7.42GB	47.8MB	33%	6.24GB	5.72GB	11MB	8.51GB	1706939	1678778
94%	97%	119819	0	119819	0	120129	13.7GB	7.05GB	41%	7.34GB	42.6MB	35%	6.29GB	7.01GB	11MB	8.49GB	1882656	1877381

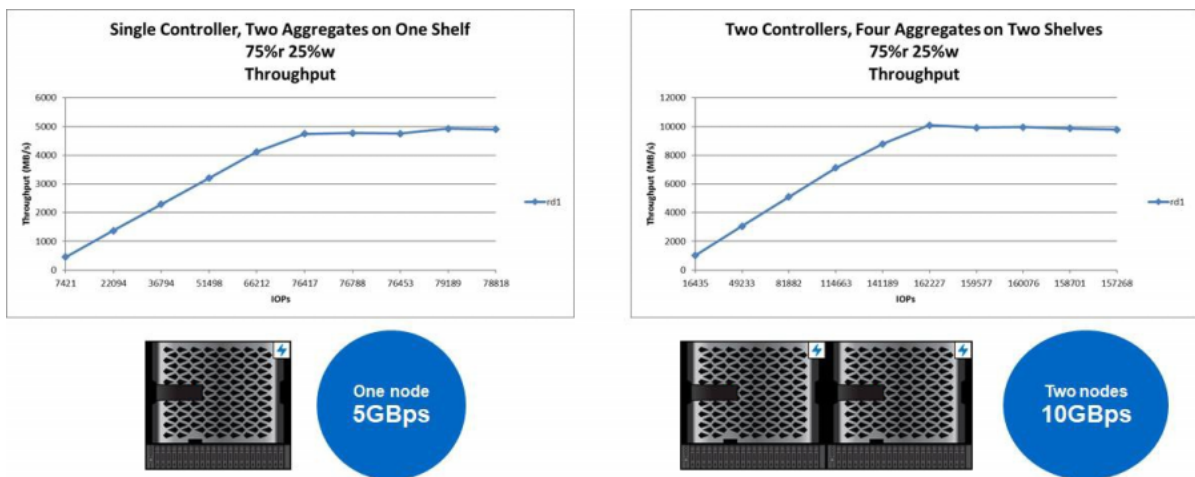
Data Protection

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—8.4GB per second for the SnapMirror transfer. This means that the replication window for a 150TB dataset would be 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've seen near-linear performance gains by adding nodes to a cluster. The following graphs show throughput results for a single-node AX4100 and a two-node AX4100.

Figure 17 CPOC scale-out throughput results



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

Conclusion

Not only is a FlexGroup volume great for small or high-file-count workloads such as EDA and software builds, but it also can handle 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 disk.

FUJITSU Storage
ETERNUS AX series All-Flash Arrays,
ETERNUS HX series Hybrid Arrays
Data Protection and Backup for ONTAP FlexGroup Volumes

P3AG-5682-01ENZ0

Date of issuance: October 2021
Issuance responsibility: FUJITSU LIMITED

- The content of this manual is subject to change without notice.
- This manual was prepared with the utmost attention to detail.
However, Fujitsu shall assume no responsibility for any operational problems as the result of errors, omissions, or the use of information in this manual.
- Fujitsu assumes no liability for damages to third party copyrights or other rights arising from the use of any information in this manual.
- The content of this manual may not be reproduced or distributed in part or in its entirety without prior permission from Fujitsu.

FUJITSU