



Technical Report

SnapMirror configuration and best practices guide for ONTAP 9

Tony Ansley, NetApp
February 2026 | TR-4015

Abstract

This document describes information and best practices related to configuring NetApp® SnapMirror® replication and has been updated to include changes introduced in NetApp ONTAP® 9.18.1.

TABLE OF CONTENTS

Purpose and intended audience	6
Solution overview	6
What is SnapMirror?	6
Why Choose NetApp SnapMirror?	7
Use case summary	8
Unified architecture flexibility	9
ONTAP networking basics	10
Common SnapMirror networking terms	10
SnapMirror networking requirements	12
Intercluster networking requirements	12
Intercluster multipathing	12
How SnapMirror uses intercluster LIFs to replicate traffic	16
Firewall requirements	17
NetApp replication basics	17
Licensing	17
NetApp SnapMirror Asynchronous technology	17
SnapMirror Synchronous	21
SnapMirror for NetApp cloud volume platforms	22
SnapMirror configuration	22
Cluster peering	22
SVM peering	23
SnapMirror data protection relationship	23
Protection policies	29
SnapMirror schedules	36
Create a SnapMirror relationship	37
Baseline transfer during initialization of SnapMirror relationship	39
Manual update to the SnapMirror relationship	39
Long-term retention of snapshots	40
Converting a legacy DP SnapMirror relationship to an XDP SnapMirror relationship	41
SnapMirror relationship conversion details	41
SnapMirror Asynchronous ONTAP feature interoperability	47
SnapMirror and Snapshot copies	47
SnapMirror and qtrees	48

SnapMirror and FlexGroup volumes.....	48
SnapMirror and FlexClone technologies.....	49
SnapMirror and NetApp storage efficiency	50
SnapMirror and volume move.....	52
SnapMirror for disk shelf failure protection	52
SnapMirror and volume autosize.....	52
SnapMirror and NDMP	53
SnapMirror and Fabric Pool.....	53
SnapMirror and consistency groups	54
SnapMirror for CGs interoperability with other ONTAP features.....	59
SnapMirror for SVM DR.....	59
Defining what SVM DR will protect.....	60
Creating an SVM DR relationship.....	63
SVM DR scalability.....	63
SVM DR advanced topics.....	64
SVM DR interoperability	66
Performance	70
Calculate SnapMirror throughput for performance.....	71
SnapMirror and network compression	72
SnapMirror throttling.....	74
How to change TCP receive buffer size	75
Concurrent replication operations.....	76
Network sizing requirements	77
S3 SnapMirror	77
SVM data mobility	78
Hardware interoperability.....	78
Troubleshooting tips	78
Troubleshooting cluster peer relationships	78
Troubleshooting SVM peer relationships.....	79
Understanding SnapMirror relationship status.....	80
Troubleshooting SnapMirror relationships	80
Best practices for DR configurations	82
Configuration and failover for DR.....	83

Environment failover requirements and assumptions	83
Preparing the destination for failover	84
Performing a failover	85
Post failover volume configuration	86
Where to find additional information	86
Version history	87

LIST OF TABLES

Table 1) ONTAP networking terminology.	10
Table 2) Shared versus dedicated interface groups.	15
Table 3) Minimum RPO for SnapMirror replication of FlexVol volumes and FlexGroup volumes.....	25
Table 4) Maximum fan out for a SnapMirror source volume.....	26
Table 5) Cascade relationships types between any three volumes in a SnapMirror relationship	27
Table 6) SnapMirror policy types.....	30
Table 7) Replication behavior when source volume is on a non-TSSE capable controller and destination volume is on a TSSE enabled controller.....	51
Table 8) Compression behavior when source volume has TSSE enabled, and destination volume controller is not on a TSSE capable.	51
Table 9) Compression behavior when source volume has TSSE enabled, and destination volume controller is TSSE capable.....	52
Table 10) What snapshots get replicated as part of a SnapMirror replication of a consistency group.....	57
Table 11) SnapMirror interoperability with other ONTAP features.....	59
Table 12) Differences between SVM DR and SnapMirror.	60
Table 13) SVM DR scalability.....	63
Table 14) SVM DR limitations when SVMs contain FlexGroup volumes or consistency groups	63
Table 15) TCP receive buffer windows.....	75
Table 16) Maximum number of concurrent SnapMirror transfers per node by ONTAP version and controller model.	76

LIST OF FIGURES

Figure 1) NetApp SnapMirror replication overview.	7
Figure 2) Unified architecture flexibility.....	10
Figure 3) Intercluster LIFs and relationship to other ONTAP networks.	12
Figure 4) Failover group used in failover mode.	14
Figure 5) Failover group used in multiplex mode.....	15
Figure 6) TCP stream distribution examples.	16
Figure 7) SnapMirror custom protection policy create.	20
Figure 8) Add custom SnapMirror protection policy.....	21
Figure 9) SnapMirror Asynchronous replication in action	24

Figure 10) SnapMirror fan-out and fan-in.	26
Figure 11) SnapMirror cascade.	27
Figure 12) Cascade relationships involving SnapMirror synchronous.	28
Figure 13) Bypassing a leg in a SnapMirror cascade infrastructure.	29
Figure 14) SnapMirror Asynchronous policy definition.	30
Figure 15) DailyBackup asynchronous SnapMirror policy definition.	31
Figure 16) DPDefault asynchronous SnapMirror policy definition.	32
Figure 17) MirrorAllSnapshots asynchronous SnapMirror policy definition.	33
Figure 18) MirrorLatest asynchronous SnapMirror policy definition.	33
Figure 19) MirrorAndVault asynchronous SnapMirror policy definition.	34
Figure 20) Unified7year asynchronous SnapMirror policy definition.	35
Figure 21) XDPDefault policy definition.	35
Figure 22) SnapMirror schedules can be listed and created in System Manager.	36
Figure 23) SnapMirror schedules can be listed and created using CLI.	36
Figure 24) Start a SnapMirror relationship update.	39
Figure 25) Relationship update dialog.	40
Figure 26) FlexGroup volumes used in SnapMirror cascade and fan-out configurations.	48
Figure 27) SnapMirror behavior during FlexGroup rebalance operation.	49
Figure 28) Creating a FlexClone volume at a SnapMirror destination.	49
Figure 29: Comparing ONTAP CGs with SnapMirror CG support.	54
Figure 30: SnapMirror for consistency groups.	55
Figure 31) Replicated information for SVM DR using <code>-identity-preserve true</code>	61
Figure 32) Replicated information for SVM DR using <code>-identity-preserve true</code> with <code>-discard-configs network policy</code>	62
Figure 33) Replicated information for SVM DR using <code>-identity-preserve false</code>	62
Figure 34) SVM DR fan-out limitations for SVMs.	65
Figure 35) SVM DR cascade support.	66
Figure 36) SVM DR support for FlexCache origin volumes.	67
Figure 37) SVM DR with MetroCluster.	68
Figure 38) SVM DR used to migrate an SVM from an HA cluster to a MetroCluster sync-source cluster.	68
Figure 39) SVM DR support for FlexGroup volumes.	69
Figure 40) SVM DR support for consistency groups.	70
Figure 41) SnapMirror network compression functional diagram.	73
Figure 42) ONTAP S3 SnapMirror overview.	78
Figure 43) Volume layout for DR.	84

Purpose and intended audience

This document is intended for individuals who administer, install, or support ONTAP systems and who intend to configure and use SnapMirror technology for data replication.

This document assumes that the reader understands the following processes and technologies:

- A working knowledge of ONTAP operations
- A working knowledge of NetApp features such as NetApp Snapshot™ copy technology, NetApp FlexVol® or FlexGroup® volumes, and NetApp FlexClone technology
- General knowledge of DR and data replication solutions
- Familiarity with the [ONTAP Data Protection Guide](#) on the NetApp Support site

Solution overview

Businesses need robust strategies to maintain data availability during hardware, software, or site failures. Data protection is critical because any data loss translates directly into lost revenue and productivity.

Historically, organizations relied on recovering data from archival media such as tape, disk, or cloud storage. Today, **mirroring data to another location** has become the preferred approach because it minimizes data loss and enables faster recovery compared to traditional archival methods.

What is SnapMirror?

NetApp SnapMirror technology is an ONTAP integrated replication capability designed for **backup, disaster recovery (DR), and archival use cases**. It establishes a data protection relationship between:

- FlexVol or FlexGroup volumes
- Consistency groups (CGs) for application-consistent replication
- Storage Virtual Machines for protection of complete NAS services

Replication is **policy-driven** and scheduled to keep secondary copies synchronized with the primary. SnapMirror's policy-driven architecture maximizes flexibility by offering the following features:

- **Deployment Options**
 - **Same Cluster:** For local protection and quick recovery.
 - **Remote Cluster:** For site-level DR across LAN/WAN.
 - **Cloud Targets:** Using **Cloud Volumes ONTAP** or **Amazon FSx for NetApp ONTAP** for hybrid and public cloud deployments.
- **Failover and Failback**

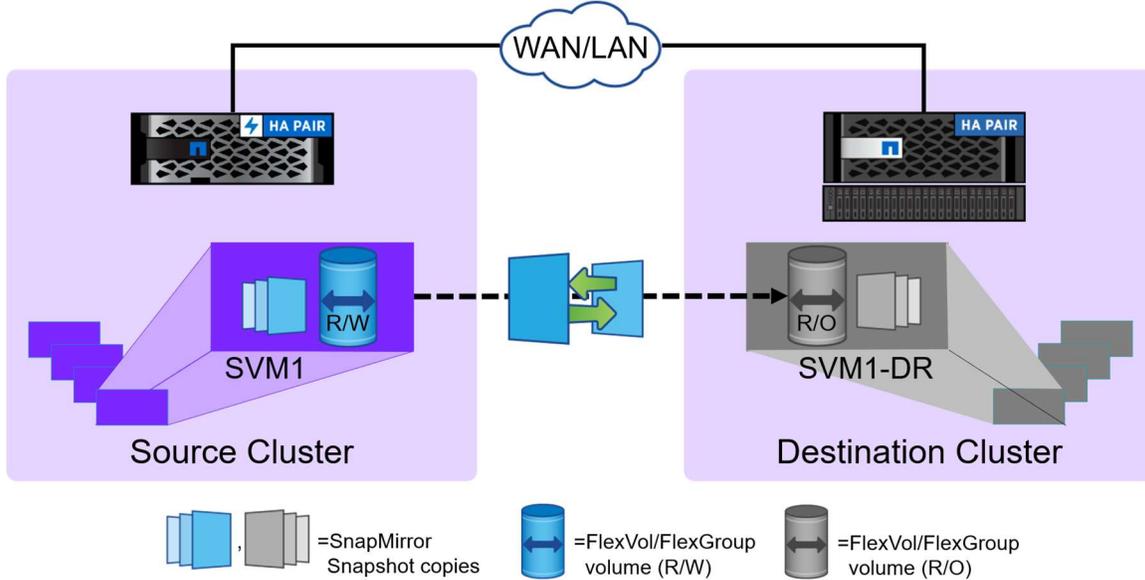
In the event of a primary site outage:

 - Storage administrators can break the SnapMirror relationship and serve data from the secondary site.
 - After recovery, SnapMirror supports resync operations to replicate changes back to the primary without requiring a full baseline transfer.
- **Business Value**
 - **Lower TCO:** Leverage DR infrastructure for secondary workloads (e.g., analytics, testing).
 - **Flexible Topologies:** Supports many-to-one (fan-in), one-to-many (fan-out) and cascade replication for multi-site protection.

- **Integrated Efficiency:** Preserves ONTAP storage efficiencies (compaction, compression, and deduplication) during replication.

For an overview of NetApp SnapMirror replication, see Figure 1.

Figure 1) NetApp SnapMirror replication overview.



Why Choose NetApp SnapMirror?

NetApp SnapMirror technology offers a fast, flexible enterprise solution for replicating data over LAN or WAN networks. Its key advantages include:

- **Robust Enterprise Technology**

Mature ONTAP features with capabilities such as recovery from update failures, concurrent replication, and bandwidth throttling.

Example: A financial services firm uses SnapMirror to maintain data replication across multiple data centers, ensuring compliance with strict regulatory requirements for data integrity.

- **Speed and Efficiency**

Transfers only changed blocks, preserves storage efficiencies, and uses network compression for faster performance.

Example: A healthcare provider replicates electronic health records between sites in minutes instead of hours, reducing downtime during maintenance windows.

- **Flexibility**

Supports custom schedules, reverse synchronization, and advanced topologies like fan-out and cascade.

Example: A global retailer uses SnapMirror fan-out to replicate inventory data from a primary site to multiple regional sites for faster local access.

- **Testability**

Enables instant writable clones of replicated volumes using ONTAP FlexClone technology for disaster recovery (DR) testing.

Example: A SaaS company creates test environments from SnapMirror destinations without impacting production, accelerating application updates and QA cycles.

- **Ease of Use**

Integrated with ONTAP System Manager for simplified workflows and monitoring.

Example: An IT team at a mid-sized enterprise manages SnapMirror relationships through a graphical interface, reducing administrative overhead and training costs.

- **Security**

Provides end-to-end encryption with TLS 1.2 and AES-256.

Example: A government agency replicates sensitive data across sites while meeting strict security compliance standards for encrypted data in transit.

- **Cloud-Enabled**

Supports replication to NetApp cloud solutions such as Amazon FSx and Cloud Volumes ONTAP.

Example: A media company replicates large video archives to AWS using SnapMirror for cost-effective cloud storage and global content distribution.

Use case summary

SnapMirror replication technology supports a broad set of use cases including:

- **Near-line backup**

One of the primary use cases for SnapMirror is data backup. Since the early days of enterprise data storage, data backup has been the purview of tape. Tape backup created some challenges for responsive data recovery, foremost of which is that restore operation simply failed and at the best of times it would take hours to recover from a disaster scenario.

SnapMirror can be used as a primary backup tool by replicating data within the same cluster or to remote targets. Using SnapMirror, storage administrators can restore single files or an entire storage configuration quickly.

- **Disaster recovery**

SnapMirror technology also enables comprehensive disaster recovery (DR) plans. If critical data is replicated to a different physical location, a serious disaster does not have to cause extended periods of unavailable data for business-critical applications. Clients can access replicated data across the network until the recovery of the production site.

In the case of failback to the primary site, SnapMirror provides an efficient means of resynchronizing the DR site with the primary site by transferring only changed or new stored data back to the primary site by simply reversing the SnapMirror relationship. After all changes have been restored, the primary production site resumes normal application operations, and SnapMirror continues the transfer to the DR site without requiring another baseline transfer.

- **DR testing and application testing and development**

NetApp FlexClone® technology quickly creates a copy of a SnapMirror destination volume enabling read-write access of the secondary copy to allow administrators to confirm if all the production data is available and verify that applications can operate normally from the DR site.

- **Data distribution and remote data access**

SnapMirror technology can be used to distribute large amounts of data throughout an enterprise, enabling access to data at remote locations. Remote data access provides faster access by clients in

remote locations. It also allows more efficient and predictable use of expensive network and server resources because storage administrators can replicate production data at a specific time to minimize overall network utilization.

- **Backup offloading and remote tape archiving**

SnapMirror technology can also be used for backup consolidation and for offloading tape backup overhead from production servers. This approach facilitates centralized backup operations and reduces backup administrative requirements at remote locations. NetApp Snapshot technology eliminates the traditional backup window on the primary storage system.

Unified architecture flexibility

SnapMirror provides data protection across a broad set of platforms to meet many different user requirements. SnapMirror started as a native feature of ONTAP for data center environments but has expanded in recent updates by embracing new cloud-centric platforms for private, hybrid, and public cloud deployments, as shown in Figure 2.

- **Data center deployments**

SnapMirror can be used on any ONTAP platform in the data center to meet a broad set of performance and scalability needs from NetApp ONTAP arrays to ONTAP Select software-defined storage within VMware vSphere infrastructures. These deployments can be managed by using CLI, REST APIs, or the web-based System Manager graphical user interface.

- **Private cloud deployments**

SnapMirror private cloud deployments are architected similarly to data center deployments, with all the same platform flexibility, with the added management flexibility offered by REST APIs to support cloud-centric management and operational platforms such as Kubernetes, Ansible, and VMware vRealize.

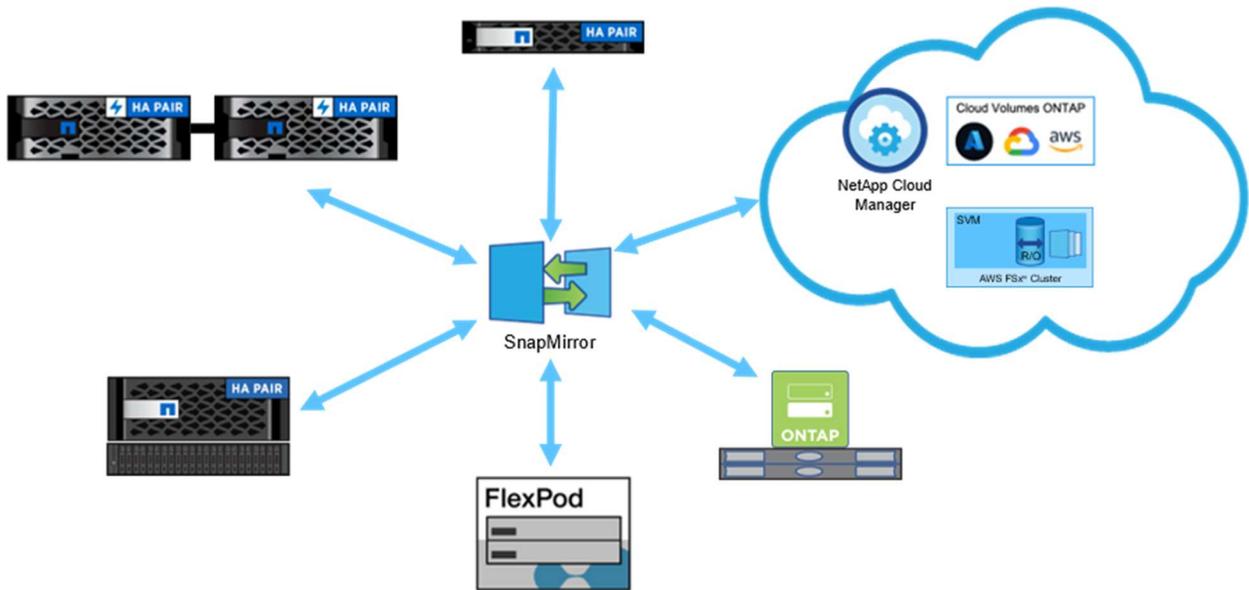
- **Hybrid cloud deployments**

SnapMirror provides support for data replication between ONTAP on-premises and ONTAP public-cloud implementations such as NetApp Cloud Volumes ONTAP deployments on major cloud provider environments using the NetApp Cloud Manager service and Amazon FSx for ONTAP service.

- **Public cloud deployments**

SnapMirror can support ONTAP deployments completely hosted in the cloud through solutions such as Cloud Volumes ONTAP deployments on major cloud provider environments, and Amazon FSx for NetApp ONTAP (Figure 2). These ONTAP solutions provide access to SnapMirror functionality to support primary and DR cloud-based deployments and can use System Manager, CLI, or REST APIs for management. In addition, through NetApp Cloud Manager, SnapMirror can be configured to replicate data between Cloud Volumes for ONTAP clusters within the same geographical region or in different regions for added protection.

Figure 2) Unified architecture flexibility.



ONTAP networking basics

For a basic ONTAP networking overview, see [TR-4182-Ethernet Storage Design Considerations and Best Practices for Clustered Data ONTAP Configurations](#). This section addresses only SnapMirror specific networking requirements.

Common SnapMirror networking terms

Table 1 lists the basic terminology used in NetApp ONTAP for SnapMirror replication.

Table 1) ONTAP networking terminology.

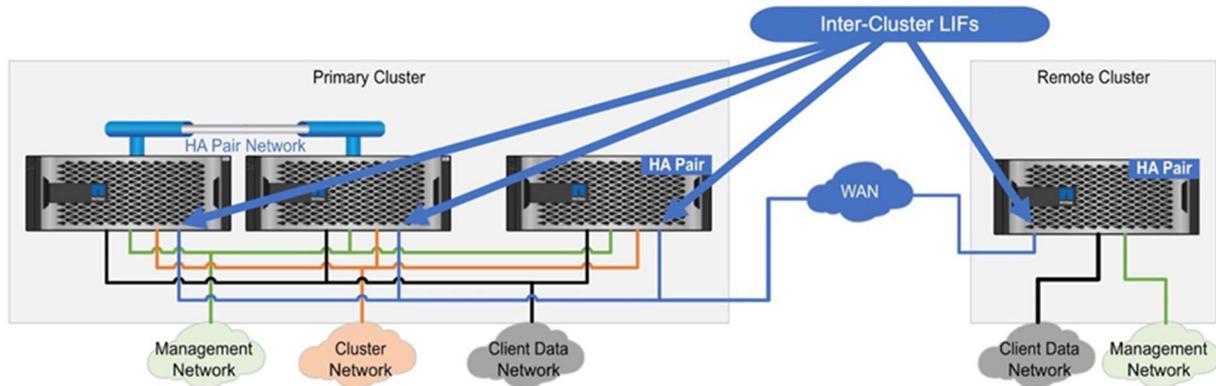
Term	Definition
Node	A single device offering ONTAP storage services. A node can be stand-alone (NetApp AFF, FAS, and ASA chassis, ONTAP Select, or Cloud Volumes ONTAP) or integrated into the same physical chassis as an HA pair. Typically, ONTAP storage clusters consist of one or more 2-node HA pairs, though some cloud deployments can be single-node deployments.
High-availability (HA) pair	Two ONTAP storage nodes configured in a pair for high availability. Each node can take ownership of storage resources from its paired node in the event of a node failure.
Cluster	One or more HA pairs that are interconnected and managed as a single storage solution. ONTAP clusters can consist of up to 6 SAN or 12 NAS HA pairs.
IPspace	An IPspace defines a distinct IP address space in which ONTAP arrays can participate. IPspaces offer a distinct routing table and enable ONTAP support for multitenancy deployments where tenants might use a common IP address range.
Broadcast domain	A broadcast domain is a group of physical network ports in the same IPspace that can communicate through a single layer-2 network protocol.

Term	Definition
Intercluster network	The network used for communication and replication between different clusters. Intercluster networks are used by SnapMirror and other data protection solutions within the ONTAP environment.
Physical port	A physical network port such as e0e or e0f Ethernet or 0c or 0e FC ports. Physical ports can support Ethernet, FC, or provide unified protocol support. Physical ports host virtual ports and/or logical interfaces (LIFs).
Interface group (ifgrp)	A collection of physical ports combined to create one logical port used for link aggregation. An interface group can offer expanded throughput, redundancy or both. Interface groups comply with the IEEE 802.3ad and 802.1AX network standards.
Virtual LAN (VLAN)	A virtual LAN is an IEEE 802.1Q standard protocol that subdivides a physical network into distinct broadcast domains. As a result, traffic is completely isolated between VLANs unless a router (layer 3) is used to connect VLANs. In ONTAP, VLANs subdivide a physical port into several separate virtual ports, allowing for one of the key components of our secure multitenant messaging—isolation of data.
Virtual port	A virtual port is a logical network interface that can come in various forms: <ul style="list-style-type: none"> • Interface group • VLAN
Logical interface (LIF)	A LIF is an IP address or a worldwide port name (WWPN) that is associated with a port. LIFs can be attached to physical ports, interface groups or VLANs. LIFs have associated attributes such as failover rules, role and firewall rules.
Intercluster LIF	A LIF used to connect to an intercluster network. An intercluster LIF must be created on each cluster node before a cluster peering relationship can be established. Intercluster LIFs can only fail over to ports in the same node.
Failover group	A failover group is a cluster-scoped list of physical ports within a cluster and in the same IPspace and same broadcast domain that can host a logical interface in the event of a physical port failure.
Cluster peer	Cluster peers are participants in an intercluster relationship. An intercluster peer relationship must be created before any data movement through ONTAP data protection services can be performed. The act of creating this intercluster relationship is called cluster peering.
Storage virtual machine (SVM)	An SVM is a logical storage server that provides data access to LUNs and/or a network-attached storage (NAS) namespace from one or more data LIFs.
SVM peer	SVM peers are participants in an SnapMirror relationship. An inter-SVM peer relationship must be created before any data movement through ONTAP data protection services can be performed. The act of creating this inter-SVM relationship is called SVM peering.
Volume	A data storage construct that refers to ONTAP native FlexVol volumes or FlexGroup volumes.
Consistency group	An ONTAP container that associates multiple volumes such that they can be simultaneously protected in an application consistent way. SnapMirror uses consistency groups to take application consistent snapshots for replication purposes.

SnapMirror networking requirements

There are multiple types of networks in ONTAP, as shown in Figure 3. Cluster management, inter-node communications and client data access each have dedicated networks and provide critical access to an ONTAP cluster for everyday use. Additionally, ONTAP uses another network type called the intercluster network that is used by SnapMirror for data replication between SnapMirror endpoints. These endpoints are typically associated with different, physically remote ONTAP clusters. For this technical report, focus will be on this intercluster network.

Figure 3) Intercluster LIFs and relationship to other ONTAP networks.



Intercluster networking requirements

An intercluster network consists of multiple Intercluster LIFs within each cluster that must meet the following requirements:

- All intercluster LIFs for a given SnapMirror relationship must be in the same IPspace. Different IPspaces can be used to peer with different clusters by using different intercluster LIFs.
- At least one intercluster LIF must be configured on every node in the source cluster and every node in the destination cluster participating in the SnapMirror relationship.
- The IP addresses assigned to intercluster LIFs can reside in the same subnet as a cluster's data LIFs or in a different subnet.
- Intercluster LIFs can be assigned to ports, interface groups, or VLANs.
- Intercluster LIFs can reside on the same physical ports as data LIFs.
- Intercluster LIFs are node scoped. Therefore, when the port hosting an intercluster LIF fails, the LIF can only fail over to another intercluster-capable port on the same node, as defined by the LIF's failover policy.
- All intercluster LIFs must have consistent settings (the same maximum transmission units [MTUs], flow control, TCP options, and so on) using the same IPspace.
- SnapMirror replication over an FC network is not available in ONTAP.

For additional information regarding intercluster networking, see [TR-4182-Ethernet Storage Design Considerations and Best Practices for Clustered Data ONTAP Configurations](#) and the [ONTAP 9 Cluster and SVM peering with CLI](#).

Intercluster multipathing

It is recommended to configure multipathing for intercluster LIFs used for a SnapMirror relationship. Multipathing provides data path redundancy, bandwidth aggregation, or both depending on the multipath

configuration used. This section discusses the use of interface groups (interface groups) and failover groups by highlighting the benefits of each and when to use each to meet the solution requirements.

The decision on which type of multi-pathing solution to use will depend on the capabilities provided by the network switches to which the cluster nodes connect. Interface groups require switches that support the IEEE 802.3ad link aggregation protocol. Failover groups are completely managed by ONTAP and do not require any additional switch features to implement.

Intercluster multipathing using interface groups

Interface groups use the [IEEE 802.3ad](#) link aggregation (LAG) protocol and the 802.1AX link aggregation control protocol (LACP) standards to provide logical network interfaces that consist of two or more physical ports. Interface groups can host LIFs or VLANs to provide path redundancy or path redundancy with bandwidth aggregation depending on selected interface group type. It is important that all ports added to an interface group have the same link-speeds, duplex settings, allowed VLANs (if it is layer-2), native VLAN, and so on settings.

To create an interface group, use the `network port ifgrp create` command. This section will reference various parameters to this command. To learn more about interface groups and how to create and manage them, see [Combine physical ports to create interface groups](#).

The following recommendations are suggested when using interface groups for SnapMirror intercluster replication:

- Use a dedicated broadcast domain (VLAN) for SnapMirror to simplify identification of upstream links in the network for troubleshooting and added network security.
- If the switch infrastructure supports LACP, use interface groups that are configured as type `-mode multimode_lacp`; otherwise, configure the interface group as `-mode multi-mode`
- If switches are capable, always configuring inter-switch links as multi-chassis control (vPC, MLAG, and so on) mode and connect ports in the interface group residing on different nodes to different physical switches.
- Select constituent ports from different network interface controllers (NICs) that have the same physical network characteristics such as throughput, frame size, duplex settings, etc. and are in the same broadcast domain.
- Create interface groups with either 2, 4, or 8 constituent physical ports. This ensures the link distribution algorithm has an opportunity to evenly distribution connections between physical ports. The algorithm always returns a result (bucket) between 0 and 7 with each bucket assigned to one physical link in the interface group at interface group initialization. Sizing the interface group with a number of physical links that are evenly divisible by eight ensures that each link will get the same chance of assignment. Assuming there is variability in the input values used by the hash algorithm, network sessions should get evenly distributed.
- If source or destination clusters are small or single node (as may be possible in a cloud deployment) consider creating multiple intercluster LIFs per interface group to allow the hash algorithm to more evenly distribute connections across all physical ports. A maximum of eight intercluster LIFs per node can be configured.
- Use the `-distr-func port` parameter for interface groups used for intercluster data transfer. This provides the hash algorithm with the broadest set of values to use resulting in a more even distribution of connections mapped to the eight hash result buckets.

Best practice

When using interface groups, use the following parameters to the `network port ifgrp create` command:

- Use `-mode multimode_lacp`
- Use `-distr-func port`.

Intercluster multipathing using failover groups

Failover groups provide a non-switch dependent method of link failure protection. By default, every NetApp cluster creates a failover group for each broadcast domain. Custom failover groups can be defined to optimize the failover behavior of an intercluster LIF.

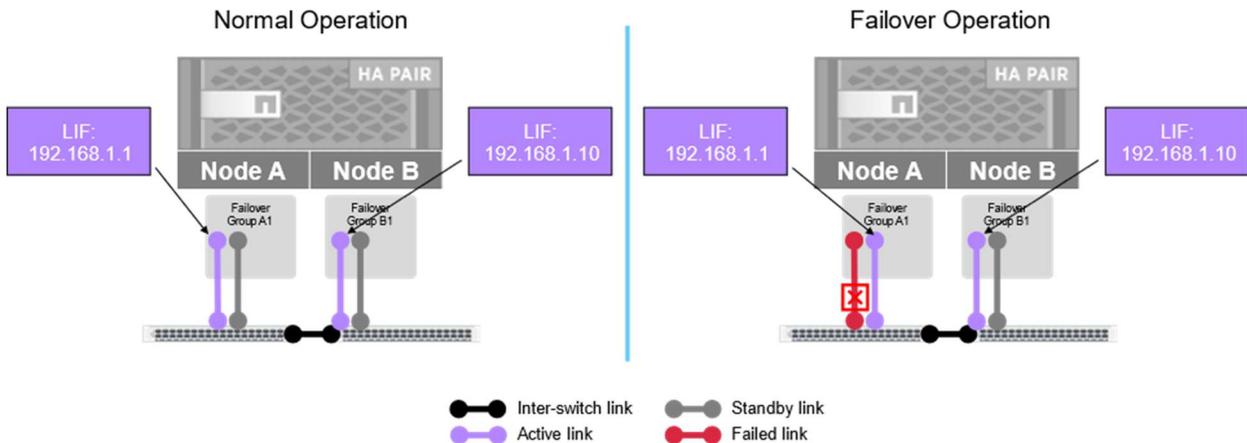
Intercluster LIFs can be assigned to failover groups that span cluster nodes, but intercluster LIFs will only failover to ports within the local node of the home port for that LIF. Intercluster LIFs are defined with the `-role intercluster` or `-service-policy default-intercluster` parameters to the `interface create` command.

Depending on how LIFs are assigned to the ports in a failover group, the failover group can operate in one of two modes:

- Failover mode
- Multiplexing mode

Once the failover group is created, LIFs are assigned to individual ports in the failover group. The port a LIF is initially assigned to is considered that LIF's home port. All other ports in the failover group are considered standby ports for that LIF. If a single intercluster LIF is assigned to a failover group, then the failover group operates in failover mode, as shown in Figure 4.

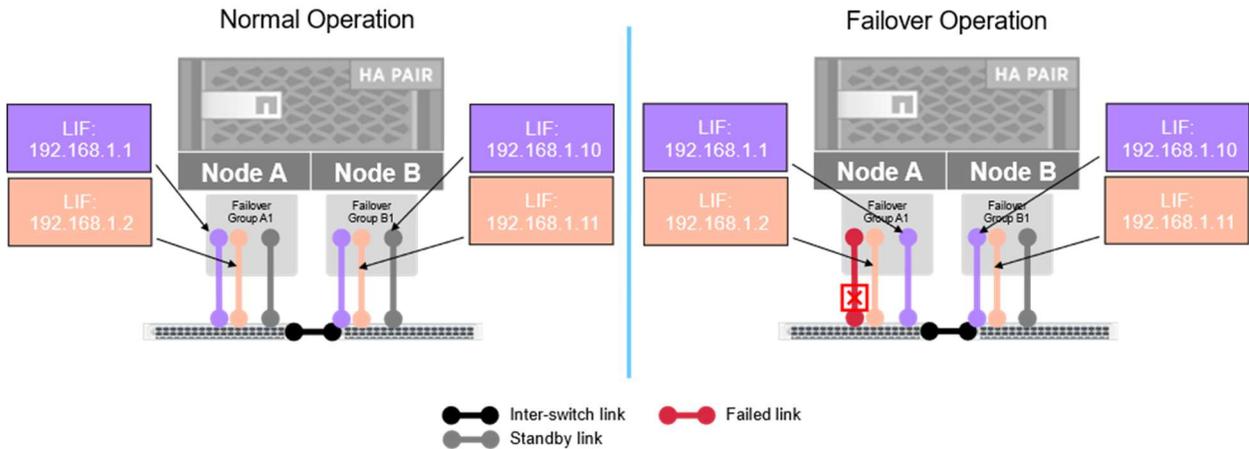
Figure 4) Failover group used in failover mode.



If more than one intercluster LIF is assigned to ports within the failover group, then those LIFs enable users of those LIFs – such as SnapMirror – to multiplex across the active LIFs within that failover group. NetApp recommends that these LIFs be assigned to different ports within the failover group to provide additional bandwidth for SnapMirror replication, as shown in Figure 5.

For each active LIF in the failover group, the port initially assigned to that LIF is considered its home port and all other ports – even if hosting other LIFs – is considered that LIF's standby ports. If a LIF's active link goes down, it is possible that a port in the failover group could host multiple LIFs. There is no guarantee that the failed over LIF will be assigned to a port not hosting any other LIFs.

Figure 5) Failover group used in multiplex mode.



Best practice

- Create custom failover groups on each node specifically for intercluster LIFs.
- Use a dedicated broadcast domain for SnapMirror to simplify identification of upstream links in the network for troubleshooting and added network security.
- Create at least two intercluster LIFs and assign one LIF to different port in the failover group.
- Use the following parameters when creating intercluster LIFs:
 - `-service-policy default-intercluster`
 - `-data-protocol none`
 - `-failover-group <custom failover group name>`

Share or dedicate failover groups?

Intercluster LIFs can be assigned to failover groups that contain other types of LIFs or can be assigned to ports in failover groups specifically defined for intercluster communications. Table 2 highlights some considerations for each.

Table 2) Shared versus dedicated interface groups.

Shared interface group	Dedicated interface group
<ul style="list-style-type: none"> • Replication performance is not critical. • Higher bandwidth (25–100GbE) ports are available and SnapMirror replication workload is not enough to saturate links. • Node interface available bandwidth is not the bottleneck...the WAN link is. • Replication requirements support a high RPO (for example, I only need to replicate once a day). • Data change rate on the source volume is low. • Port availability on each node is low. 	<ul style="list-style-type: none"> • DR is a primary goal for SnapMirror. • Lower bandwidth node ports that might easily get saturated with SnapMirror replication workload. • Replication requirements need to support a very small RPO (dictating the need for different MTU). • Data change rate on the source volume is very high (dictating the need for different maximum MTU). • Port availability on each node supports having a dedicated failover group.

Best practice recommendations

Although it is not required, the following are recommendations for using failover groups for SnapMirror:

- Configure at least two intercluster LIFs per node using a failover group containing at least two physical ports with the same network characteristics (speed, duplex mode, and so on)
- If the network utilization generated by the data protocols (SMB, NFS, iSCSI, or NVMe) is above 50%, then a dedicated failover group for intercluster communication is recommended.
- When using dedicated failover groups for intercluster communication, consider the following:
 - Remove the physical ports from other failover groups, including the default failover group associated with the cluster broadcast domain.
 - Use a dedicated broadcast domain for additional security.
 - Consider setting MTU size to jumbo frames to maximize network efficiency (consult network engineer for viability of this option)
- Use a standardized naming convention for intercluster LIFs. For example, `node_name_icl#` or `node-name-ic#`, depending on preference.

How SnapMirror uses intercluster LIFs to replicate traffic

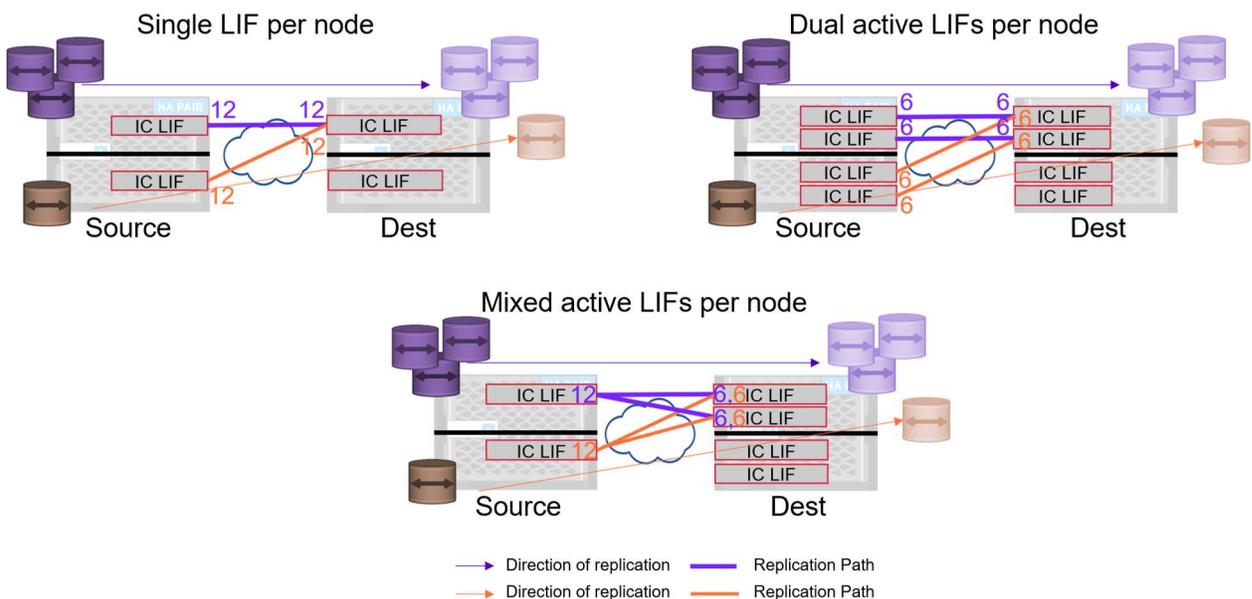
ONTAP provides a centralized network session manager to provide optimized network traffic management for various functions involving intercluster communications such as SnapMirror. For SnapMirror, this cluster session manager provides up to 12 TCP send streams and 12 TCP receive streams for each source-destination node peer relationship. For discussion purposes, only the 12 send sessions will be used, but there are also 12 receive sessions that are managed similarly.

These 12 sessions are controlled using the following rules:

- A set of 12 sessions is created between each source-destination node interaction between clusters that have been peered.
- On each node, the 12 sessions are distributed across the available IC LIFs for that source-destination node relationship.
- All volumes replicating between a given source-destination node pair share the 12 sessions.

Using these rules, Figure 6 illustrates how these 12 TCP sessions are distributed.

Figure 6) TCP stream distribution examples.



Note: It is not possible to select a specific LIF pair to use for a replication event. All session management is performed automatically by the ONTAP session manager.

Firewall requirements

SnapMirror uses the typical socket, bind, listen, and accept sequence on a TCP port. The firewall and the intercluster firewall policy must allow the following protocols:

- TCP ports 11104 and 11105 for intercluster control and data respectively
- TCP port 10000 for NDMP backup services
- Optionally, TCP port 443 (HTTPS) in each direction between the intercluster LIFs for managing ONTAP arrays using System Manager or REST APIs

Note: HTTPS is not required to set up cluster peering using the CLI.

NetApp replication basics

Licensing

A SnapMirror license is required for each cluster that will participate in SnapMirror replication. This license is included in the Data Protection Bundle and ONTAP One license. If the SnapMirror source and destination are on different clusters, a SnapMirror license must be enabled on each cluster. All nodes in each cluster must have a license.

Starting with ONTAP 9.12.1P1 for NetApp AFF C-Series controllers and ONTAP 9.13.1 for all other NetApp FAS, ASA, and AFF A-Series controllers, SnapMirror is included as part of ONTAP One license sold with every NetApp ONTAP controller.

Note: The license must be present on both source and destination.

Note: ONTAP One is available for clusters operating older versions of ONTAP. Contact your NetApp representative for more information.

NetApp SnapMirror Asynchronous technology

SnapMirror replicates data from a source volume or consistency group in a cluster to an equivalent volume or consistency group in a destination cluster by using snapshots. SnapMirror performs the following operations:

1. A snapshot of the data on the source is created. This snapshot will have a SnapMirror label of `sm_created`.
2. The snapshot is copied to the destination during baseline synchronization. This process creates a destination that is online, read-only, and contains the same data as the source at the time of the most recent common snapshot.
3. Ongoing replication is performed based on the SnapMirror policy schedule. What gets replicated is based on the SnapMirror policy type:
 - a. **Async-Mirror:** A new `sm_created` snapshot is created and replicated to the destination cluster. If `mirror-all-snapshots` is part of the `async-mirror` policy rules, all snapshots – manually created and policy created – are replicated to the destination cluster regardless of SnapMirror label.
 - b. **Vault:** Snapshots created manually or using a snapshot policy that have a SnapMirror label matching any SnapMirror policy rules will be replicated to the destination cluster

- c. **Mirror-Vault:** A new `sm_created` snapshot is created and replicated along with any snapshots that have a SnapMirror label matching any SnapMirror policy rules will be replicated to the destination cluster.

Note: See **Error! Reference source not found.** for details on policy types and creating SnapMirror replication policies.

When a SnapMirror data protection relationship is established (`snapmirror initialize`), the destination volumes are identical replicas of the source volumes, including snapshots, volume settings, and ONTAP space efficiency features that are accessible as a read-only volumes. Breaking the SnapMirror relationship (`snapmirror break`) makes the destination volumes writable and is used to perform a failover for `async-mirror` and `mirror-vault` relationships to recover from a failure on the source volume. SnapMirror is sophisticated enough to identify the data changed at the failover site and replicating that changed data back to the primary system. The original SnapMirror relationship can then be reestablished (`snapmirror resync`).

SnapMirror may use either of two different replication engines to create replicas – LRSE and BRE. Although both engines operate at the volume level, they have different characteristics:

- **Logical replication with storage efficiency (LRSE).** LRSE uses block-level metadata and knowledge of the file system to determine differences between snapshots at the indirect pointer level. LRSE organizes the transfer of data from the source to the destination in two streams.
 - The data stream consists of data blocks that are transferred with specific volume block number (`vvbn#`) within the destination volume. This `vvbn#` helps identify the block number at which the data is stored on the source FlexVol volume, but without specifying a file context. On the destination, the data is written to the data warehouse (DW) file with a file block number (`fbn#`) which corresponds to the `vvbn#`.
 - The user files are transferred by reference using the user file inodes, which share blocks with the data warehouse file and do not use buffer trees that require parsing to reach a specific object. LRSE makes explicit requests to the block-sharing infrastructure of the DW blocks (the donors) with user files (recipients) while replication transfer is in progress.

The mirror has a structure of logical block pointers to the original data set that has a completely different on-disk physical layout relative to the source.

LRSE preserves space efficiency over the wire and on the destination when replicating data in storage-efficient source volumes. Storage efficiency is an important part of LRSE because features such as block sharing and compression allow a volume to effectively hold far more data than the space used. This efficiency must be preserved during replication to avoid the replica growing to an intolerably large size, not to mention the time needed to transfer it. LRSE also allows enablement of storage efficiency features on the secondary, independent of the primary storage settings. For more information, see [Use deduplication, data compression and data compaction to increase storage efficiency – overview](#).

In addition to asymmetric storage efficiency on primary and secondary storage, LRSE enables version flexibility where the destination version can be different than the source. It also supports asymmetric snapshots where the destination can support a greater number of snapshots than the source. All the files and directories in the source file system are created in the destination file system. Therefore, SnapMirror can replicate data between a storage system running an older version of ONTAP and a storage system running a newer version. This approach allows reduced downtime because the controllers on either side can be upgraded at any time while reducing overhead and managing complex topologies (fan-in, fan-out, and cascade).

The SnapMirror relationship is created with `-type XDP` using the SnapMirror policy type `async-mirror`, `vault` or `mirror-vault`.

- **(DEPRECATED) Block Replication Engine (BRE).** BRE replicates the on-disk layout from a source volume to a destination volume either as a whole or as an incremental update with 4K blocks. The BRE uses knowledge of the file system to determine differences between snapshots at the block-allocation level and replicates only those changed blocks. Therefore, the copy of data created on the

destination has an identical structure of physical block pointers to the original data set on the source. The BRE replicates volumes using volume block (vbn#) read and write operations.

This SnapMirror relationship is created with `-type DP` using the SnapMirror policy type `async-mirror`. BRE only supports volume mirroring use cases and does not support vault use cases.

Note: As of 9.11.1, BRE SnapMirror relationships are used for legacy data protection policies only. All newly created SnapMirror relationships default to using LRSE (see section Logical replication with storage efficiency).

Note: As of ONTAP 9.12.1, BRE is no longer supported and systems that still host BRE SnapMirror relationships will not be able to upgrade to ONTAP 9.12.1 until those BRE relationships are converted to LRSE (XDP). For more information, see [Convert an existing DP-type relationship to XDP](#).

The performance characteristics of LRSE are similar to those of BRE because the replication engine only transfers the difference between two snapshots from the primary to the secondary. This incremental-only transfer leads to savings in terms of storage and network bandwidth. Starting with ONTAP 9.3, SnapMirror XDP mode replaces SnapMirror DP mode as the SnapMirror default. More details can be found at [XDP replaces DP as the SnapMirror default](#).

SnapMirror can also be integrated with NetApp SnapCenter® to replicate application consistent snapshots, such as those used for enterprise database applications. snapshots are created in coordination with the application to guarantee that no in-flight I/O operations cause inconsistencies in the snapshot. After creating an application consistent snapshot, SnapCenter can then trigger a SnapMirror replication of these application consistent snapshots to the secondary storage system.

Unified data protection

The XDP relationship type enables SnapMirror to provide a single, unified replication engine to perform both mirroring (for disaster recovery) and vaulting (for disk-based backup and archiving) use cases. Moving forward, the use cases previously referred to as SnapVault will be subsumed into the SnapMirror standard nomenclature as SnapMirror with a policy type of `vault` and `mirror-vault`.

Overall, unified replication with SnapMirror provides powerful data management capabilities for virtualization, protecting critical data while providing the flexibility to move data between locations and storage tiers, including cloud service provider hosted SnapMirror instances. The relationship is created with relationship type of XDP, and a policy type of `mirror-vault`. NetApp provides a pre-defined SnapMirror policy called `Asynchronous` that uses these settings. The policy can always be modified to include custom rules for backing up specific snapshots. In addition, this functionality reduces the number of secondary snapshots needed on the destination.

The major benefits for SnapMirror unified replication are as follows:

- Only one baseline copy of a volume is needed to perform mirroring and archiving to the secondary volume.
- Less network traffic is required between the primary and secondary (a single baseline plus fewer snapshots over time).
- The flexibility to replicate between storage systems running different ONTAP releases. XDP relationships are version independent, while DP relationships are not.
- To avoid corrupting replication from primary to secondary, unified replication makes it possible to recover the primary volume from available snapshots.
- An XDP relationship removes the limitation of the destination controller requiring an ONTAP major version number equal to or later than the major version of the source controller.

In System Manager for ONTAP 9.8 and later, SnapMirror Asynchronous uses the `Asynchronous (mirror-vault)` protection policy by default. A custom policy is required to modify any policy parameters.

The following example shows how unified replication can be configured with the Asynchronous policy from the CLI:

```
cluster02::> snapmirror create -source-path snap_src1:Source -destination-path  
svm_dst1:Source_dest -type XDP -policy Asynchronous
```

Best practice

Use a mirror-vault relationship whenever possible. There is no performance or space impact in using this policy type even when archive is not a requirement of your data protection plan. By doing so, it is easy to modify your data protection plan to include archive in the future without requiring additional volume baseline images.

To create a custom policy, navigate to Protection > Overview > Local Policy Settings > Protection Policies > Add (Figure 7 and Figure 8).

Figure 7) SnapMirror custom protection policy create.

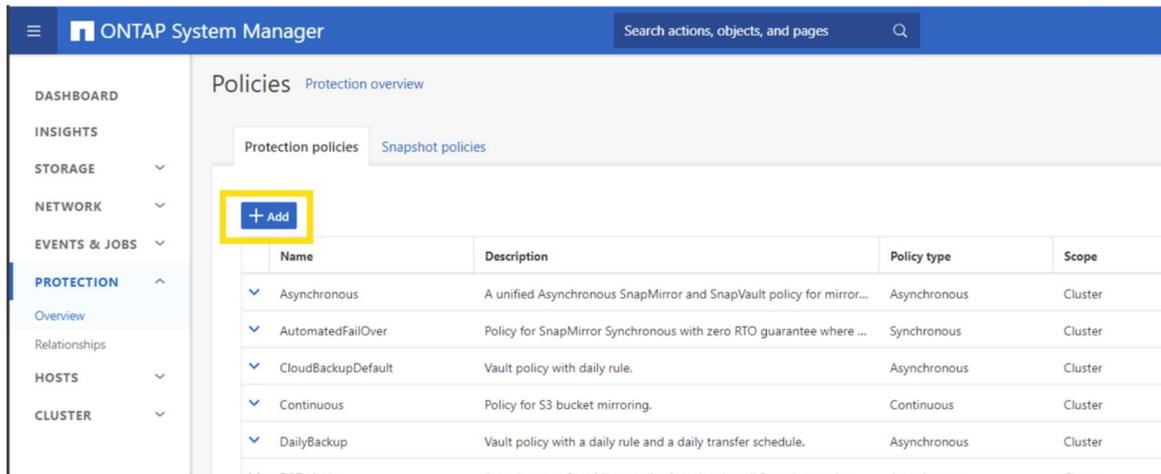
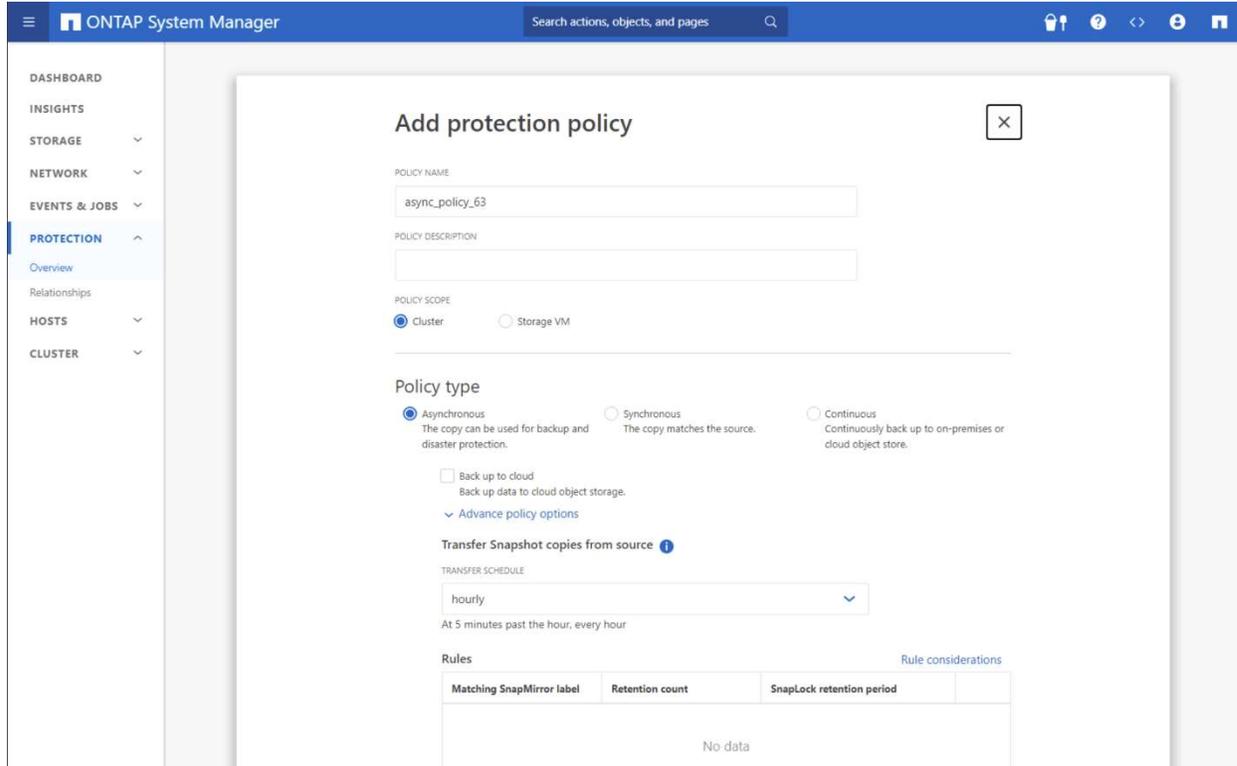


Figure 8) Add custom SnapMirror protection policy.



Load-sharing mirror

Every SVM in a NAS environment has a unique namespace. The SVM root volume is the entry point to this namespace hierarchy. For clusters consisting of two or more HA pairs, load-sharing mirrors (LSMs) of SVM root volume should be considered to ensure the namespace remains accessible to clients in the event of both nodes of an HA pair failing. Load-sharing mirrors are not suitable for clusters consisting of a single HA pair and are not suitable for a MetroCluster environment.

Note: Starting with ONTAP 9.1, load-sharing mirrors have been deprecated for data volumes and are only supported for SVM root volumes.

Note: SnapMirror load-sharing mirrors are only capable of supporting NAS (CIFS/NFSv3). Load-sharing mirrors do not support NFSv4 clients or SAN client protocol connections (FC, FCoE, or iSCSI). However, NFSv4 and load-sharing mirrors can coexist in the same environment.

Best practice

LSMs are of limited value in modern ONTAP cluster solutions due to the multiple hardware redundancies integrated into the hardware and cluster software. For most deployments, LSMs are not recommended, rather it is recommended to use SVM disaster recovery as an alternative solution.

For additional information concerning load-sharing mirrors, see [Managing SnapMirror Root Volume Replication](#).

SnapMirror Synchronous

SnapMirror Synchronous (SM-S) is an easy-to-use solution that replicates data synchronously between a source volume and a destination volume over a LAN or metropolitan area network (MAN). SM-S provides

high data availability and rapid DR for business-critical applications without loss of data due to primary site or cluster failure.

SM-S uses a different mechanism than SnapMirror Asynchronous for transferring volume data and is covered in [TR-4733 SnapMirror Synchronous configuration and best practices for ONTAP 9](#).

SnapMirror for NetApp cloud volume platforms

SnapMirror can be used to provide asynchronous data protection of on-premises or cloud-hosted ONTAP volumes and SVMs. Working in cooperation with Google, Amazon, and Microsoft, NetApp offers a broad range of ONTAP cloud offerings that can be used as an overall data management fabric. The following sections describe each offering that is supported for use with SnapMirror.

NetApp Cloud Volumes ONTAP

Cloud Volumes ONTAP enables customers to use performance and cost-optimized cloud storage in conjunction with on-premises ONTAP data management. Cloud Volumes ONTAP offers a full array of ONTAP APIs and data management services built upon the native cloud-provider compute, storage, and network products. Cloud Volumes ONTAP offerings are available from Amazon Web Services (AWS), Microsoft Azure, and Google Cloud.

To use a Cloud Volumes ONTAP instance as a SnapMirror endpoint, it must first be created using the cloud provider management tools or NetApp Cloud Manager. Once created, a SnapMirror asynchronous relationship can be created between an on-premises ONTAP cluster and the Cloud Volumes ONTAP instance using any NetApp ONTAP management interface – CLI, System Manager, REST APIs or NetApp BlueXP.

Amazon FSx for NetApp

Amazon FSx for NetApp ONTAP is a native AWS storage service that offers fully managed ONTAP as a service that can be administered through NetApp BlueXB or the AWS suite of management tools such as AWS Management Console, AWS CLI, or AWS REST APIs. After an Amazon FSx for NetApp ONTAP cluster is created, it can be used as a source or destination cluster in a SnapMirror Asynchronous relationship with other Amazon FSx for NetApp ONTAP instances, Cloud Volumes ONTAP instances, or on-premises ONTAP clusters.

SnapMirror configuration

Cluster peering

The cluster peering feature of ONTAP allows administrators of independent clusters to establish a peer relationship between them. SnapMirror leverages the ONTAP core networking infrastructure and depends on intercluster LIFs and cluster peering to transfer replication data between volumes or SVMs. For complete details on the cluster peering process, please review [ONTAP Cluster and SVM peering overview](#) in the online documentation.

- Before setting up cluster peering, confirm that the connectivity, port, IP address, subnet, firewall, and cluster-naming requirements are met. Cluster peering requirements include the following:
- The time on the clusters must be in sync to within 300 seconds (five minutes) for peering to be successful. Cluster peers can be in different time zones.
- At least one intercluster LIF must be created on every node in the cluster.
- Every intercluster LIF in the local cluster IPspace must be able to communicate with every intercluster LIF in the remote cluster IPspace.

- Every intercluster LIF requires an IP address dedicated to intercluster replication.
- The MTU settings of the ports must be consistent. The default value of 1,500 is correct for most environments.
- All paths on a node used for intercluster replication should have equal performance characteristics.

Note: SnapMirror does not support network address translation (NAT).

Starting with ONTAP 9.6, cluster peering uses encrypted communication, which means any SnapMirror relationship that is created uses an additional layer of security through TLS encryption.

Best practice

- If using host files instead of DNS for name resolution of cluster nodes, the name and IP address of the source system must be in the `hosts` file of the destination system and vice versa.
- Use at least one intercluster IP address from each node in the remote cluster, so that the peer relationship remains available in the event of a single node failure.

SVM peering

SVM peering connects two SVMs to allow replication to occur between them, which requires cluster peering first. SVM peering enables granularity of access or the delegation of various replication operations to the SVM admin.

Note: A peer relationship is not required to mirror data between two SVMs in the same cluster or between two volumes in the same SVM.

For more information about SVM peering, see [ONTAP Cluster and SVM peering overview](#) in the online documentation.

SnapMirror data protection relationship

A relationship created between the source object (for example, a FlexVol or FlexGroup volume, or a consistency group) in primary storage and the destination object in secondary storage is called a data protection relationship.

SnapMirror relationships have the following characteristics:

- SnapMirror relationships are created and managed on the destination cluster.
- SnapMirror relationship transfers are triggered by the scheduler on the destination cluster.
- Destination volumes must be created with the volume type of DP (`-type DP`) for SnapMirror initialization to succeed. After volumes have been created, the volumes' type cannot be changed.
- Destination volumes in a SnapMirror relationship are read-only until a manual failover is initiated by a storage administrator. Use the `snapmirror break` command to initiate a failover to the secondary copy and make the destination volume writable. The `snapmirror break` command must be performed separately for each volume or consistency group.
- The destination volumes can be mounted into an SVM namespace as read-only, but only after the initial transfer is complete.
- Destination volumes in a SnapMirror relationship between two separate clusters cannot be mounted in the same NAS namespace as the source volumes. However, the destination volumes in a SnapMirror relationship configured within a cluster can be mounted in the same namespace as the source volume if both the source and destination volumes exist in the same SVM. However, they cannot be mounted to the same mount point.
- LUNs contained in mirror destination volumes can be mapped to initiator groups (igroups) and connected to clients. However, the client must be able to support connection to a read-only LUN.
- SnapMirror relationships can be managed using the ONTAP CLI, ONTAP System Manager, REST API, Active IQ Unified Manager (formerly OnCommand Unified Manager), or NetApp BlueXP.

- If an in-progress transfer is interrupted by a network outage or aborted by an administrator, a subsequent restart of that transfer can automatically continue from a saved restart checkpoint.
- Both the source and destination SVM can have different language types, but the source and destination volumes must have the same language type.
- The destination aggregate must have adequate space to host the replicated volumes and any vault snapshots retained per the configured protection policy.
- If the SnapMirror source is a consistency group, then the destination cluster must be running an ONTAP version of 9.13.1 or later.

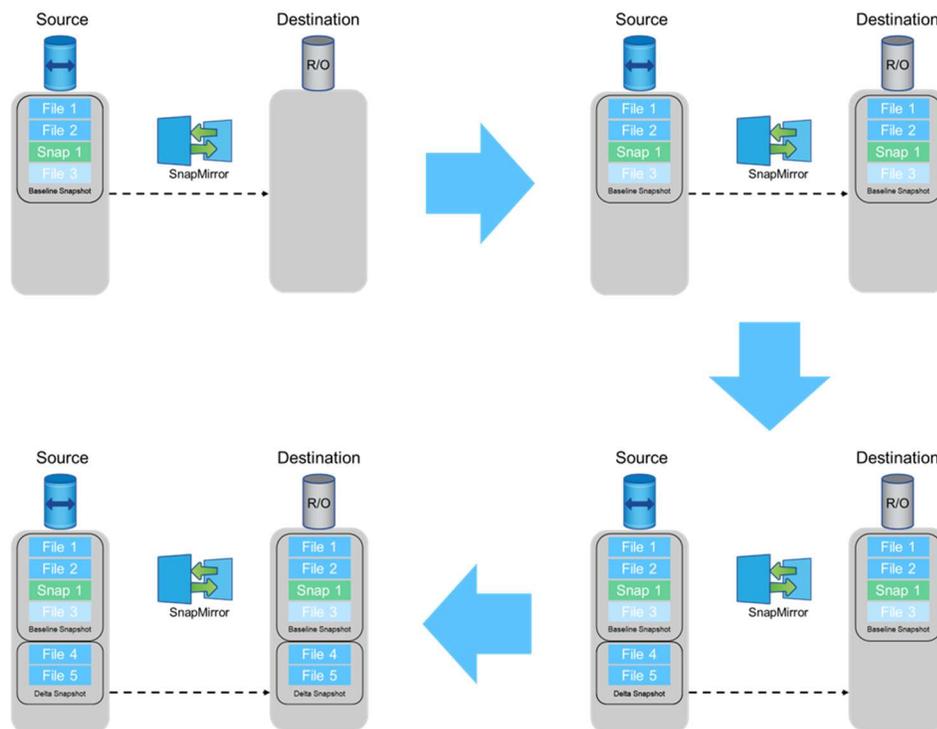
How SnapMirror replicates volumes and volume changes

SnapMirror Asynchronous replicates data using native ONTAP snapshots to create a static image of the volume and/or changes since the last SnapMirror snapshot. This ensures that only data not previously replicated to the DR volume gets replicated.

After a SnapMirror replication relationship is created, the relationship must be initialized. During this initialization phase, an initial snapshot is created of the volume. Since this snapshot is the very first SnapMirror related snapshot created on the volume, this snapshot captures the entire volume's contents. This snapshot – and thus the entire volume's contents – are replicated to the DR volume. This “baseline” is used as the starting point for scheduled SnapMirror snapshots subsequently created as part of its replication policy schedule.

Once the baseline data has been replicated, SnapMirror goes into a replication mode where it regularly creates new snapshots – based on the requirements described in the SnapMirror replication policy – on a regular schedule defined by the `-schedule` parameter supplied during the SnapMirror creation. This is illustrated in Figure 9.

Figure 9) SnapMirror Asynchronous replication in action



When a snapshot is replicated to the DR volume, that latest SnapMirror snapshot will be located on both the source volume and the DR volume. This is called the newest common snapshot (NCS). The NCS is used for the next scheduled SnapMirror replication, and this continues as required by the supplied schedule.

When a failure event occurs on the source volume, any data written to the DR site volume during the event will need to be replicated to the production site volume once the failure event is rectified. The NCS is used as a reference point to how far back data must be resynchronized after a failure event once the production site is back in operation.

Best practices

- Do not reuse a destination volume from a previously existing SnapMirror relationship. Always use a newly created volume to start a new SnapMirror relationship.
- Do not delete snapshots that SnapMirror creates in the source volume before copying the data to the destination. The most recent SnapMirror snapshot is referred to as the newest common snapshot (NCS). Incremental changes to the destination depend on the NCS. If SnapMirror cannot find the required snapshot on the source, it cannot perform incremental changes to the destination.
- To avoid unnecessary autosize changes for the data protection destination FlexGroup volume, make sure to specify that the total size of the FlexGroup volume at time of volume creation is the same as the primary FlexGroup volume.
- Do not restrict or take the destination volume offline while SnapMirror is configured to transfer. Taking the destination offline prevents SnapMirror from performing updates to the destination.
- The number of constituents on the primary FlexGroup directly relates to the number of aggregate entries that need to be specified in the `-aggr-list` parameter. When choosing which aggregate is specified in the `-aggr-list`, make sure that the aggregates have enough space for the constituents.
- To ensure an efficient SnapMirror operation of FlexGroup volumes, make sure that for each FlexGroup hosted by the same set of aggregates, use a different order in the `-aggr-list` parameter. One recommendation is to rotate the aggregates in a round-robin fashion.
- Make sure that the size of each destination constituent is such that it can ingest data from the primary constituent. Otherwise, SnapMirror operations fail when they run out of space.

Replication intervals

SnapMirror updates must establish a communication session between the source and destination nodes, creating and deleting snapshots, and determining which blocks of data to send to the destination.

SnapMirror uses the built-in ONTAP scheduler functionality to manage how often data is replicated. The period between replication events should be considered the volume's or consistency group's recovery point objective (RPO), or the acceptable amount of data loss in the event of a catastrophic failure of the primary volume or consistency group. Although the ONTAP scheduler supports creating schedules that run every minute, SnapMirror has minimum supported replication intervals listed in Table 3.

Table 3) Minimum RPO for SnapMirror replication of FlexVol volumes and FlexGroup volumes.

SnapMirror replication source	Minimum supported RPO
FlexVol volumes	5 minutes
FlexGroup volumes	30 minutes
Consistency groups	30 minutes

Fan in and fan out

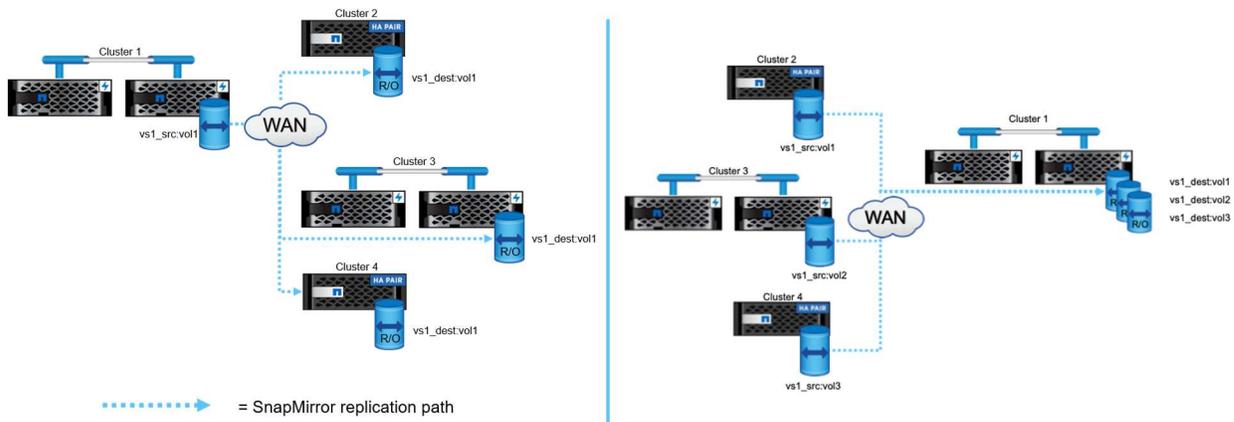
It is possible for a volume or consistency group in the source cluster to be replicated to multiple different destinations (fan-out) or volumes or consistency groups from different source SVMs can be replicated to separate SVMs in a single destination cluster (fan-in), as shown in Figure 10. Table 4 lists the fan-out limits for SnapMirror.

Table 4) Maximum fan out for a SnapMirror source volume.

ONTAP Version	Maximum fan-out
ONTAP 9.11.1 and earlier	8
ONTAP 9.12.1 and later	8 or 16
ONTAP 9.14.1 and later	8 or 20

Note: Fan out to greater than eight destination volumes from a single source volume depends on type of source and destination controllers. See [Hardware Universe](#) for fan-out limits for specific controller models.

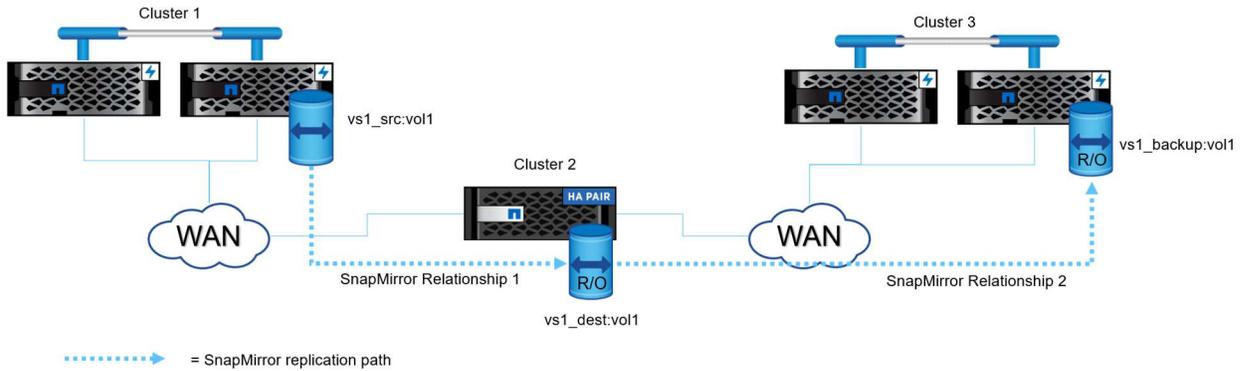
Figure 10) SnapMirror fan-out and fan-in.



Cascade relationship

SnapMirror can replicate data from a SnapMirror destination to another destination system. Therefore, a system that is a destination for one SnapMirror relationship can act as the source for another SnapMirror relationship. This is useful for distributing data from one site to multiple sites. This is referred to as cascading. In a cascade topology, intercluster networks need to be created between the primary and secondary clusters and between the secondary and tertiary clusters. An intercluster network between the primary and the tertiary cluster is not needed. While there is no limit on the number of “hops” for a single source volume, there may be practical limits in the ability to schedule each downstream replication without being interfered with by an upstream replication. An example cascade configuration with two hops is shown in Figure 11.

Figure 11) SnapMirror cascade.



The function of this deployment is to make a uniform set of data available on a read-only basis to users from various locations throughout a network and to enable the updating of that data uniformly at regular intervals.

Snapshots behave in the following ways:

- SnapMirror creates a soft lock on the snapshot of the source volume (tag) or consistency group.
- The destination system carries an extra snapshot.

SnapMirror supports a different relationship for each leg of the cascade. Cascade relationships can also be created using either SnapMirror asynchronous or synchronous relationships. For cascade relationships that include SnapMirror synchronous, the synchronous relationship must be the first relationship in the cascade chain. All other SnapMirror relationships in the cascade chain must be asynchronous.

Cascade relationships that use only SnapMirror asynchronous relationships.

For any relationship in the cascade chain, that relationship can be a mirror or a vault. Therefore, in long-chain cascades, the following combinations of data protection relationships can be used between any three clusters in the chain as described in Table 5:

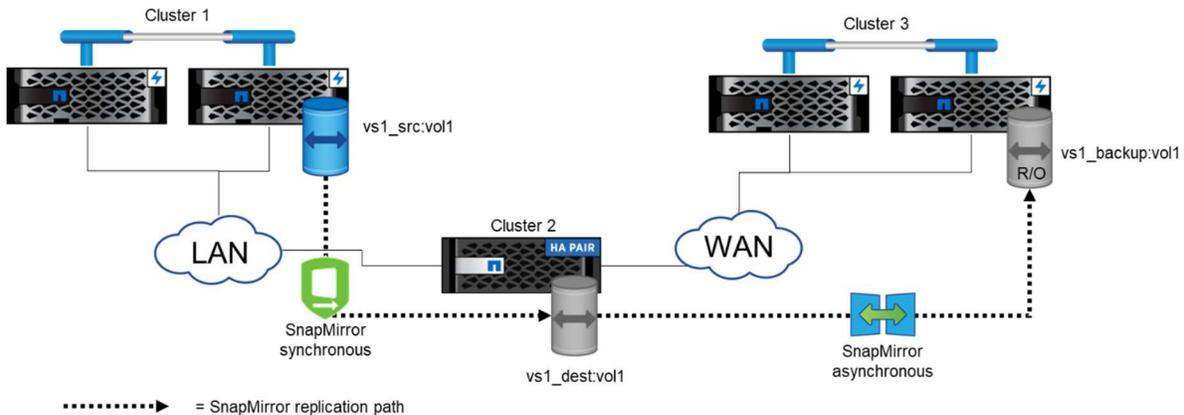
Table 5) Cascade relationships types between any three volumes in a SnapMirror relationship

Cluster 1-Cluster 2	Cluster 2-Cluster 3	Description
Mirror	Mirror	One cluster has a mirror relationship with a second cluster and that second cluster has a mirror relationship with a third cluster.
Mirror	Vault	One cluster has a mirror relationship with a second cluster and that second cluster has a vault relationship with a third cluster.
Vault	Vault	One cluster has a vault relationship with a second cluster and that second cluster has a vault relationship with a third cluster.

Cascade relationships that use SnapMirror asynchronous and synchronous relationships

For cascade relationships can be configured with the root source volume relationship that uses a SnapMirror synchronous relationship and all subsequent SnapMirror relationships in the cascade chain using SnapMirror asynchronous as shown in Figure 12.

Figure 12) Cascade relationships involving SnapMirror synchronous



The asynchronous relationship that uses the SnapMirror synchronous DP volume as the source volume (the secondary relationships in any long-chain cascade that includes SnapMirror synchronous) should be configured as an `async-mirror` with `MirrorAllSnapshots`. The resulting volume on the tertiary site (Cluster 3 in Figure 12) will depend on the ONTAP version used in the clusters involved in the SnapMirror cascade.

ONTAP 9.13.1 and older The replica volume on Cluster 3 will have snapshots available that have been exported by the SnapMirror synchronous common snapshot (exported snapshot). This means that while the DP volume on Cluster 2 will have all application created snapshots from Cluster 1 on Cluster 2, those snapshots that were created/replicated since the last exported snapshot will not be visible. The asynchronous relationship between Cluster 2 and Cluster 3 will not see these snapshots until the next scheduled SnapMirror synchronous relationship's common snapshot creation.

ONTAP 9.14.1 and later The replica volume on Cluster 3 will reflect all application created snapshots that have been created/replicated from Cluster 1 to Cluster 2 regardless of the current exported common snapshot.

Any cascade relationships downstream from Cluster 3 can be created as described for those cascades not involving synchronous relationships.

Best practice

- Keep the number of cascades as small as possible to ensure that each leg has the time to perform its scheduled replication without interfering with upstream or downstream relationships. Make sure that all the legs of the cascaded relationships complete successfully to make sure that subsequent SnapMirror updates do not fail with a `snapmirror busy` error.

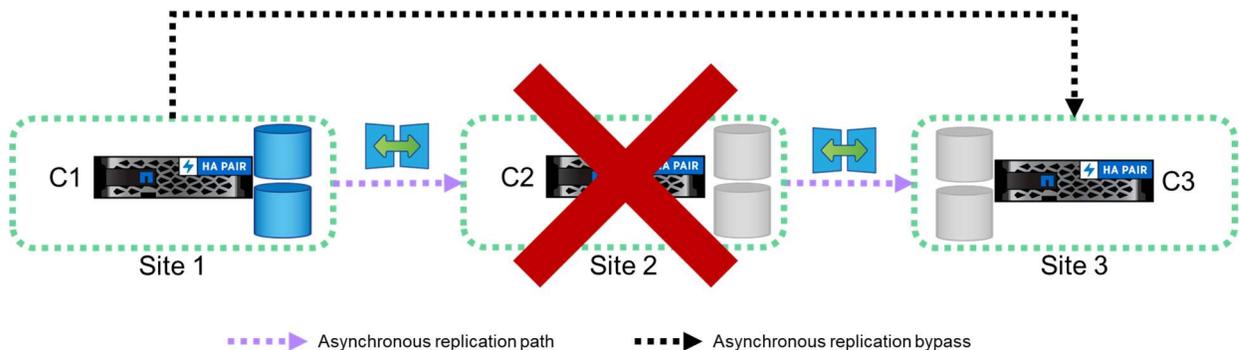
- If using a combination mirror-vault, fan-out, or cascade deployment, keep in mind that updates fail if a common snapshot does not exist on the source and destination volumes or consistency groups. Use the [snapmirror snapshot-owner create](#) command to preserve a labeled snapshot on the secondary in a mirror-vault deployment. Doing so provides a common snapshot for the update of the vault relationship.
- If the cascade includes an initial SnapMirror synchronous relationship and the clusters are running ONTAP 9.13.1 or earlier, ensure that the synchronous relationship's common snapshot schedule matches the schedule duration of the asynchronous relationship between Cluster 2 and Cluster 3. This will ensure that both relationships will result in all snapshots being replicated to all three clusters.

Note: Creation of SnapMirror cascade relationships is not supported by ONTAP System Manager and must be created using the ONTAP CLI or REST APIs.

Dual-hop volume SnapMirror

This configuration involves volume or consistency group SnapMirror replication among three clusters, which consists of a chain of relationships in which a source volume or consistency group is mirrored to a secondary volume or consistency group, and the secondary volume or consistency group is mirrored to a tertiary volume or consistency group. If the secondary volume or consistency group becomes unavailable, SnapMirror can synchronize the relationship between the primary and tertiary volumes or consistency groups without performing a new baseline transfer if there can be a common snapshot copy identified on primary and tertiary clusters (Figure 13).

Figure 13) Bypassing a leg in a SnapMirror cascade infrastructure.



Protection policies

ONTAP relies on policies to dictate when to create snapshots and how many copies are retained and/or replicated as a part of the relationship. Additionally, the policy helps determine the type of relationship that exists between the source and destination. SnapMirror replication limits the contents of the baseline transfer to the snapshot created by SnapMirror at initialization. At each update, SnapMirror creates another snapshot of the source. It then transfers the changes from this snapshot and the previously transferred snapshot (mirror) along with any new snapshots that have labels matching the labels defined in the snapshot rules defined on the SnapMirror policy if the policy is a vault policy. ONTAP comes with several predefined protection policies.

SnapMirror Asynchronous policy types

Each SnapMirror protection policy (standard or custom) is one of several different policy types. The policy types are described in Table 6.

Table 6) SnapMirror policy types.

Policy type	Definition
Async-mirror	<p>The async-mirror policy type is used by SnapMirror to transfer snapshots of two types:</p> <ul style="list-style-type: none">• Those source snapshots created by the SnapMirror engine (rule label = <code>sm_created</code>)• All source snapshots created on the volume from other Snapshot copy policies or manually created on the source volume. (rule label = <code>all_source_snapshots</code>) <p>Note: No other rules can be applied to a protection policy of type <code>async-mirror</code>.</p>
Vault	<p>The vault policy type is used by SnapMirror to copy only source volume snapshots that match the labels provided with each rule.</p> <p>The Vault policy type replaces SnapVault functionality.</p> <p>This policy does not replicate snapshots created by the SnapMirror relationship (label = <code>sm_created</code>)</p>
Mirror-vault	<p>The mirror-vault policy type is used by SnapMirror to transfer snapshots created by the SnapMirror replication engine (snapshot label = <code>sm_created</code>) and any desired source snapshots that match the snapshot label defined in each rule.</p> <p>The mirror-vault protection policy can have multiple rules defined to match the desired data protection requirements.</p>

Standard asynchronous protection policies

The following protection policy variants are available for the SnapMirror Asynchronous relationship creation:

- **Asynchronous.** This is an asynchronous SnapMirror policy of the mirror-vault policy type. As such, the asynchronous protection policy is for mirroring the latest file system on an hourly schedule and retaining seven snapshots with the daily label and retaining 52 snapshots with the weekly label from the source volume or consistency group. This is the default policy for the SnapMirror relationship creation and replaces DPDefault used as the default protection policy in previous versions of ONTAP (Figure 14).

This policy consists of the following settings:

- Policy type is `mirror-vault`.
- Create snapshot is set to `true`.
- There are three rules:
- `sm_created` replicates the changes on the source volume since the last SnapMirror generated snapshot.
- `daily` keeps seven daily snapshots.
- `weekly` and 52 weekly snapshots.

Figure 14) SnapMirror Asynchronous policy definition.

```
cluster_dst:> snapmirror policy show -policy Asynchronous -instance
```

```

Vserver: vs0
SnapMirror Policy Name: Asynchronous
SnapMirror Policy Type: mirror-vault
Policy Owner: cluster-admin
Tries Limit: 8
Transfer Priority: normal
Ignore accesstime Enabled: false
Transfer Restartability: always
Network Compression Enabled: false
Create Snapshot: true
Comment: A unified Asynchronous SnapMirror and SnapVault policy for
mirroring the latest active file system and daily and weekly Snapshot copies with an hourly
transfer schedule.
Total Number of Rules: 3
Total Keep: 60
Transfer Schedule Name: hourly
Throttle: unlimited

Rules:
SnapMirror Label      Keep Preserve Warn Schedule Prefix
-----
sm_created            1 false      0 -         -
daily                 7 false      0 -         -
weekly               52 false      0 -         -

```

- **DailyBackup.** This policy is an asynchronous SnapMirror policy of the vault policy type. It can be used to create an archive of the source volume snapshots that have a label = daily and will retain the last seven snapshots on the data protection volume. The Create Snapshot field is set to False, which indicates that this policy does not create any SnapMirror-related snapshots on the source volume.

The included protection policies listed here are considered legacy policies but can be very useful in supporting additional data protection strategies. In System Manager, these policies are only shown if the Show Legacy Policies option is selected when creating or editing a protection relationship (Figure 15).

This policy consists of the following settings:

- The policy type is set to `vault`.
- The Create Snapshot value is set to `false`, which means the policy does not create a snapshot when an update is triggered.
- There is one rule:
- Keep seven daily snapshots.

Figure 15) DailyBackup asynchronous SnapMirror policy definition.

```

cluster_dst::> snapmirror policy show -policy DailyBackup -instance

Vserver: vs0
SnapMirror Policy Name: DailyBackup
SnapMirror Policy Type: vault
Policy Owner: cluster-admin
Tries Limit: 8
Transfer Priority: normal
Ignore accesstime Enabled: false
Transfer Restartability: always
Network Compression Enabled: false
Create Snapshot: false
Comment: Vault policy with a daily rule and a daily transfer schedule.
Transfer Schedule Name: daily
Throttle: unlimited
Total Number of Rules: 1
Total Keep: 7
Transfer Schedule Name: daily
Throttle: unlimited

Rules:
SnapMirror Label      Keep Preserve Warn Schedule Prefix
-----

```

- **DPDefault.** This is an asynchronous SnapMirror policy for mirroring all snapshots and the latest active file system from the source to the destination.

Note: This policy is available for legacy SnapMirror relationships. Administrators should use the newer MirrorAllSnapshots policy for new SnapMirror relationships.

With this configuration, the SnapMirror engine creates a snapshot and then replicates the difference between the new SnapMirror snapshot and the previous one and all the other snapshots. If the relationship is being initialized, then a snapshot is taken and everything before it is replicated. After the update is complete, the older snapshot is deleted, leaving just one common SnapMirror snapshot in place (Figure 16).

This policy consists of the following settings:

- Policy type is `async-mirror`.
- Create snapshot is set to `true`.
- There are two rules:
- `sm_created` replicates the changes on the source volume since the last SnapMirror-generated snapshot.
- `all_source_snapshots` keep one copy of each unique snapshot from the source volume.

Figure 16) DPDefault asynchronous SnapMirror policy definition.

```
cluster_dst:> snapmirror policy show -policy DPDefault -instance

                Vserver: vs0
SnapMirror Policy Name: DPDefault
SnapMirror Policy Type: async-mirror
                Policy Owner: cluster-admin
                Tries Limit: 8
                Transfer Priority: normal
Ignore accesstime Enabled: false
                Transfer Restartability: always
Network Compression Enabled: false
                Create Snapshot: true
                Comment: Asynchronous SnapMirror policy for mirroring all
Snapshot copies and the latest active file system.
                Total Number of Rules: 2
                Total Keep: 2
                Transfer Schedule Name: -
                Throttle: unlimited

Rules:
SnapMirror Label          Keep Preserve Warn Schedule Prefix
-----
sm_created                1 false      0 -         -
all_source_snapshots      1 false      0 -         -
```

- **MirrorAllSnapshots.** This also is an asynchronous policy for mirroring all snapshots and the latest active file system from the primary to the secondary. This policy is similar to DPDefault (Figure 17).

This policy consists of the following settings:

- Policy type is `async-mirror`.
- Create snapshot is set to `true`.
- There are two rules:
- `sm_created` replicates the changes on the source volume since the last SnapMirror-generated snapshot.

- `all_source_snapshots` keep one copy of each unique snapshot from the source volume.

Figure 17) MirrorAllSnapshots asynchronous SnapMirror policy definition.

```
cluster_dst:> snapmirror policy show -policy MirrorAllSnapshots -instance

          Vserver: vs0
SnapMirror Policy Name: MirrorAllSnapshots
SnapMirror Policy Type: async-mirror
          Policy Owner: cluster-admin
            Tries Limit: 8
          Transfer Priority: normal
Ignore accesstime Enabled: false
          Transfer Restartability: always
Network Compression Enabled: false
          Create Snapshot: true
          Comment: Asynchronous SnapMirror policy for mirroring all snapshots
                  and the latest active file system.
          Total Number of Rules: 2
            Total Keep: 2
          Transfer Schedule Name: -
            Throttle: unlimited

Rules:
SnapMirror Label      Keep  Preserve Warn Schedule Prefix
-----
sm_created            1  false      0 -      -
all_source_snapshots 1  false      0 -      -
```

- **MirrorLatest.** This is an asynchronous policy for mirroring the latest active file system from the primary to the secondary. Using this policy, the SnapMirror engine creates a snapshot and then replicates the difference between the new SnapMirror snapshot and the previous one. If the relationship is being initialized, then a snapshot is taken and everything before it is replicated. After the update is complete, the older snapshot is deleted, leaving just one common SnapMirror snapshot in place (Figure 18).

This policy consists of the following settings:

- Policy type is `async-mirror`.
- Create snapshot is set to `true`.
- There is one rule:
- `sm_created` replicates the changes on the source volume since the last SnapMirror-generated snapshot.

Figure 18) MirrorLatest asynchronous SnapMirror policy definition.

```
cluster_dst:> snapmirror policy show -policy MirrorLatest -instance

          Vserver: vs0
SnapMirror Policy Name: MirrorLatest
SnapMirror Policy Type: async-mirror
          Policy Owner: cluster-admin
            Tries Limit: 8
          Transfer Priority: normal
Ignore accesstime Enabled: false
          Transfer Restartability: always
Network Compression Enabled: false
          Create Snapshot: true
          Comment: Asynchronous SnapMirror policy for mirroring the latest active
file system.
          Total Number of Rules: 1
            Total Keep: 1
```

```

Transfer Schedule Name: -
                        Throttle: unlimited
Rules:
SnapMirror Label      Keep Preserve Warn Schedule Prefix
-----
sm_created            1 false      0 -      -

```

- **MirrorAndVault.** This is a unified SnapMirror policy for mirroring the latest active file system and daily and weekly snapshots. Starting with ONTAP 9.5, MirrorAndVault is the new default policy when no data protection mode is specified or when XDP mode is specified as the relationship type (Figure 19).

This policy consists of the following settings:

- Policy type is `mirror-vault`.
- Create snapshot is set to `true`.
- There are three rules:
- `sm_created` replicates the changes on the source volume since the last SnapMirror-generated snapshot.
- `daily` keeps seven daily Snapshot copies.
- `weekly` keeps 52 weekly snapshots.

Figure 19) MirrorAndVault asynchronous SnapMirror policy definition.

```

cluster_dst::> snapmirror policy show -policy MirrorAndVault -instance

Vserver: vs0
SnapMirror Policy Name: MirrorAndVault
SnapMirror Policy Type: mirror-vault
Policy Owner: cluster-admin
Tries Limit: 8
Transfer Priority: normal
Ignore accesstime Enabled: false
Transfer Restartability: always
Network Compression Enabled: false
Create Snapshot: true
Comment: A unified Asynchronous SnapMirror and SnapVault policy for
mirroring the latest active file system and daily and weekly Snapshot copies.
Total Number of Rules: 3
Total Keep: 60
Transfer Schedule Name: -
Throttle: unlimited
Rules:
SnapMirror Label      Keep Preserve Warn Schedule Prefix
-----
sm_created            1 false      0 -      -
daily                 7 false      0 -      -
weekly                52 false     0 -      -

```

- **Unified7year.** This policy provides a `monthly` rule to transfer monthly snapshots and retain them for seven years (Figure 20).

This policy consists of the following settings:

- Policy type is `mirror-vault`.
- Create snapshot is set to `true`.
- There are four rules:
- `sm_created` replicates the changes on the source volume since the last SnapMirror-generated snapshot.

- daily keeps seven daily snapshots.
- weekly keeps 52 weekly snapshots.
- monthly keeps 84 monthly (7 years) snapshots.

Figure 20) Unified7year asynchronous SnapMirror policy definition.

```
cluster_dst:> snapmirror policy show -policy Unified7year -instance
      Vserver: vs0
SnapMirror Policy Name: Unified7year
SnapMirror Policy Type: mirror-vault
      Policy Owner: cluster-admin
      Tries Limit: 8
      Transfer Priority: normal
Ignore accesstime Enabled: false
      Transfer Restartability: always
Network Compression Enabled: false
      Create Snapshot: true
      Comment: Unified SnapMirror policy with 7year retention.
      Total Number of Rules: 4
      Total Keep: 144
      Transfer Schedule Name: -
      Throttle: unlimited

Rules:
SnapMirror Label      Keep Preserve Warn Schedule Prefix
-----
sm_created            1 false      0 -         -
daily                 7 false      0 -         -
weekly               52 false      0 -         -
monthly              84 false      0 monthly  -
```

- **XDPDefault.** This is an asynchronous vault policy for mirroring all snapshots with labels of Daily and Weekly. This is a legacy vault-only policy (Figure 21).

This policy consists of the following settings:

- The policy type is set to `vault`.
- The `Create Snapshot` value is set to `false`, which means the policy does not create a snapshot when an update is triggered.
- There are two rules:
 - `daily` keeps seven daily snapshots.
 - `weekly` keeps 52 weekly Snapshot copies.

Figure 21) XDPDefault policy definition.

```
cluster_dst:> snapmirror policy show -policy XDPDefault -instance
      Vserver: vs0
SnapMirror Policy Name: XDPDefault
SnapMirror Policy Type: vault
      Policy Owner: cluster-admin
      Tries Limit: 8
      Transfer Priority: normal
Ignore accesstime Enabled: false
      Transfer Restartability: always
Network Compression Enabled: false
      Create Snapshot: false
      Comment: Vault policy with daily and weekly rules.
      Total Number of Rules: 2
      Total Keep: 59
```

```

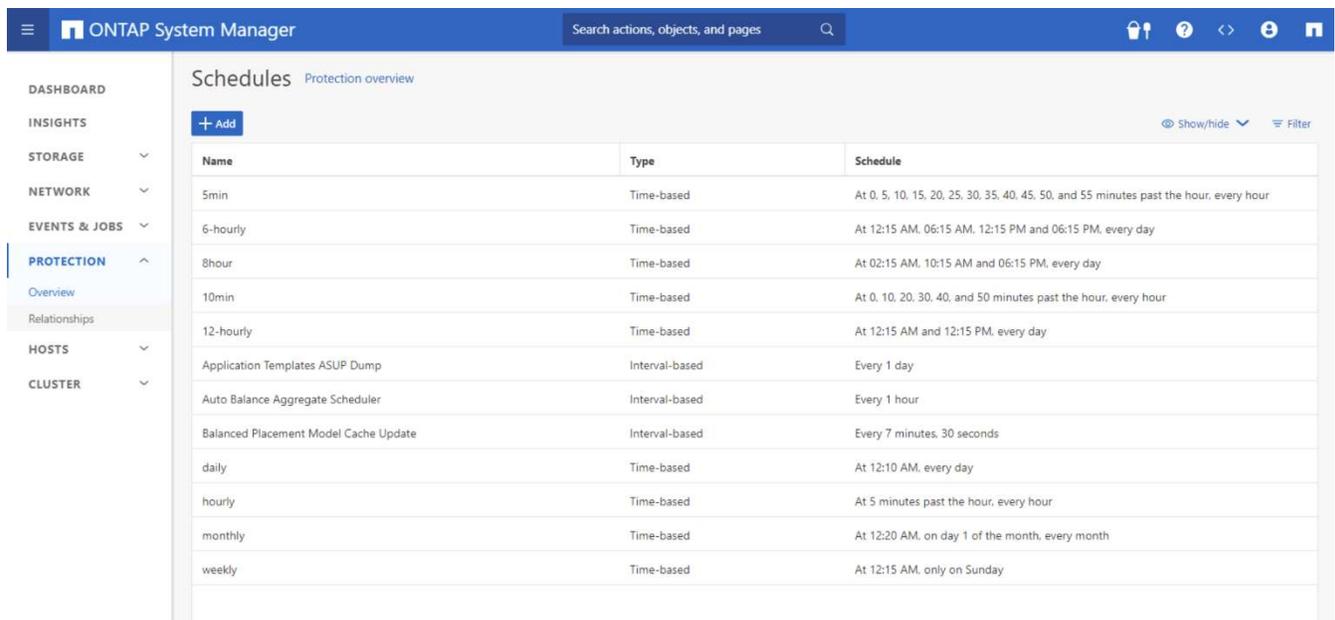
Transfer Schedule Name: -
                        Throttle: unlimited
Rules:
SnapMirror Label      Keep Preserve Warn Schedule Prefix
-----
daily                 7 false      0 -      -
weekly               52 false      0 -      -

```

SnapMirror schedules

A SnapMirror policy requires at least one SnapMirror job schedule be assigned to the SnapMirror relationship. The schedule can be a pre-defined schedule, or a custom schedule can be created to run periodic asynchronous replication by assigning a schedule to a SnapMirror relationship in the destination cluster. Figure 22 shows the pre-defined job schedules in ONTAP System Manager.

Figure 22) SnapMirror schedules can be listed and created in System Manager.



Alternatively, use the `job schedule cron create` command to create a schedule through the command line or create a new schedule by clicking the +Add button in System Manager. The following example demonstrates the creation of a schedule called `Hourly_SnapMirror` that runs at the top of every hour (on the zero minute of every hour) as shown in Figure 23.

Figure 23) SnapMirror schedules can be listed and created using CLI.

```

cluster02::> job schedule cron create Hourly_SnapMirror -minute 0
cluster02::> job schedule cron show
Name          Description
-----
5min          @:00, :05, :10, :15, :20, :25, :30, :35, :40, :45, :50, :55
8hour        @2:15,10:15,18:15
Hourly_SnapMirror  @:00
avUpdateSchedule @2:00
daily        @0:10
hourly       @:05
weekly       Sun@0:15

```

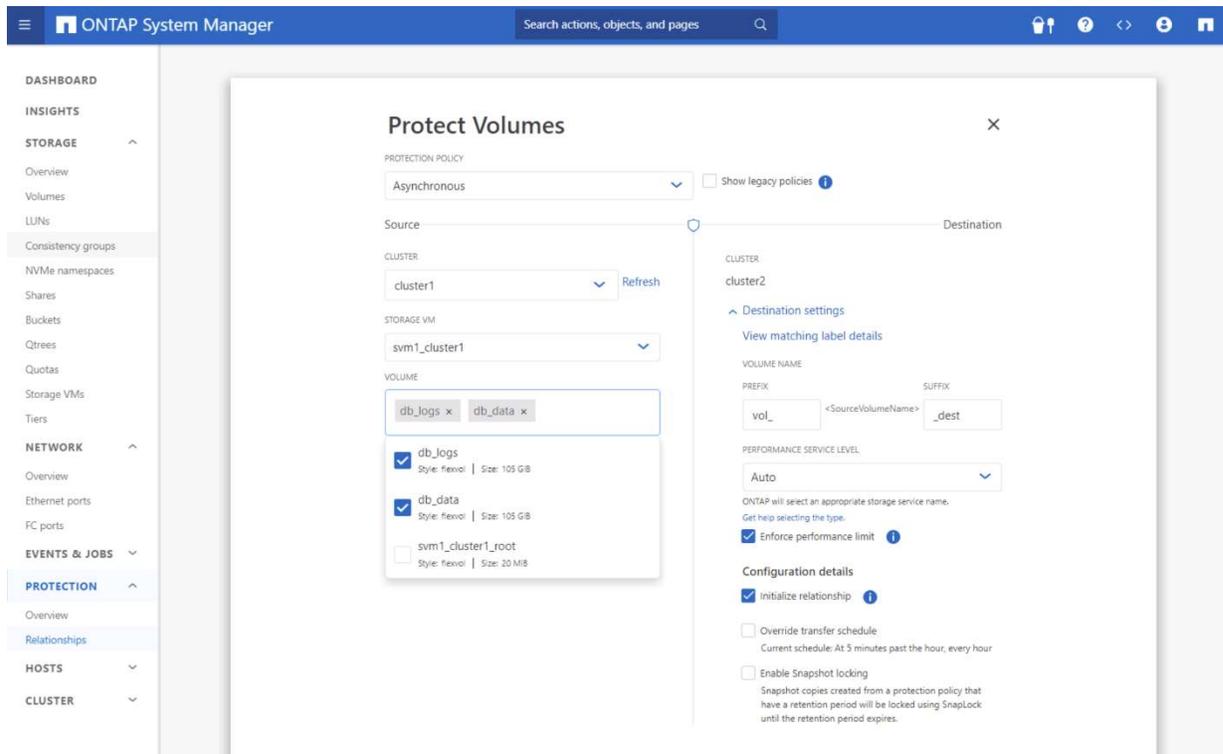
The schedule can then be applied to a SnapMirror relationship at the time of creation using the `-schedule` option or added to an existing relationship using the `snapmirror modify` command and the `-schedule` option. In this example, the `Hourly_SnapMirror` schedule is applied to an existing relationship.

```
cluster02::> snapmirror modify -destination-path cluster02://vs1/vol1 -schedule Hourly_SnapMirror
```

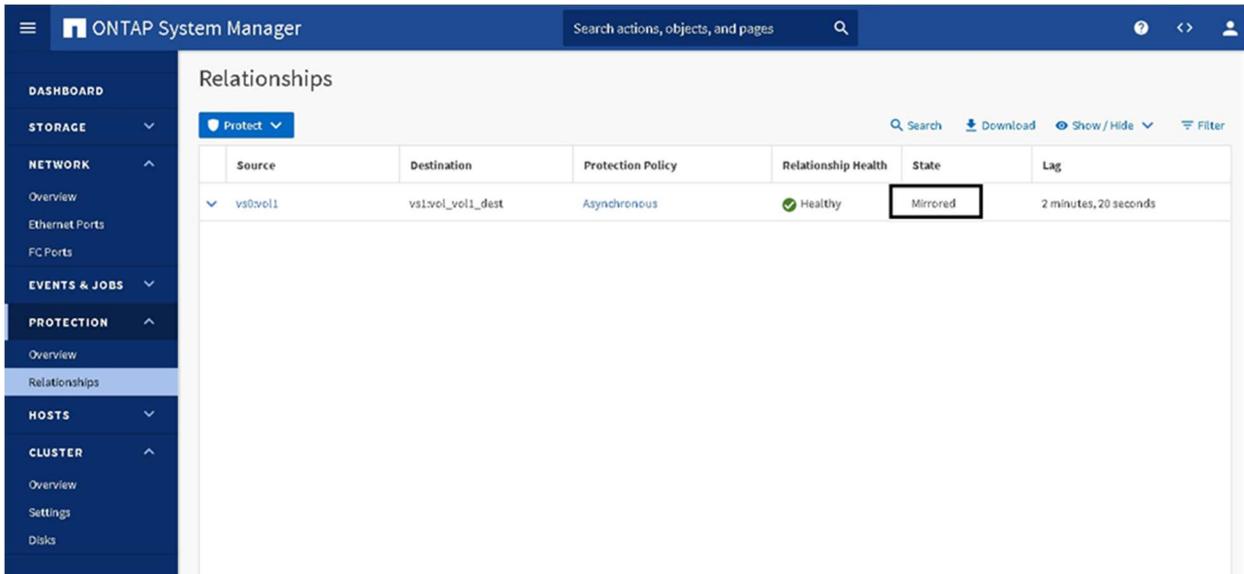
Create a SnapMirror relationship

SnapMirror relationships are created between the source volume on one cluster and a destination volume on the peered cluster enabling data protection. From ONTAP System Manager on the source cluster, complete the following steps:

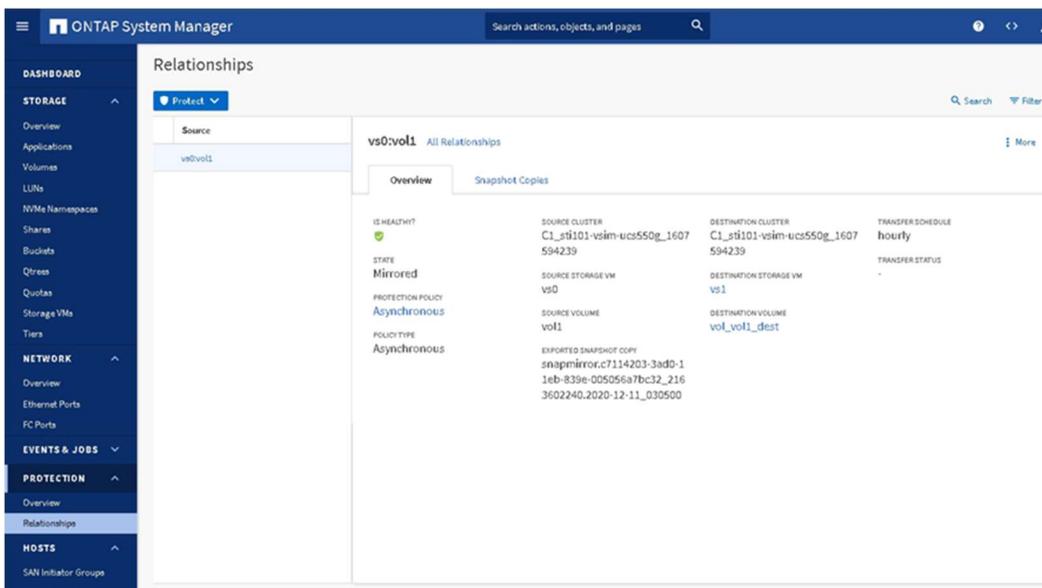
1. Click Relationships.
2. Select the volume or consistency group on which to create a SnapMirror relationship, and then click Save.



3. After the relationship is created and initialized, verify that the relationship status of the SnapMirror relationship is in the Mirrored state.



4. Select the SnapMirror relationship between the source and the destination volumes, and then verify the status in the Details tab.
5. The Details tab displays the health status of the SnapMirror relationship and shows the transfer errors and lag time.
6. The Is Healthy field must display Yes.
7. For most SnapMirror data transfer failures, the field displays No. In some failure cases, however, the field continues to display Yes. This might require checking the transfer errors in the Details section to verify that no data transfer failure occurred.
8. The State field must display Mirrored.
9. The lag time must be no more than the transfer schedule interval. For example, if the transfer schedule is hourly, then the lag time must not be more than an hour.
10. Also, navigate to the Volumes window, and then select the desired volume. Double-click the volume to view the volume details and the data protection status.



Baseline transfer during initialization of SnapMirror relationship

When a new SnapMirror relationship is created, it establishes the relationship and the metadata that defines it. Optionally, select Initialize the Relationship to perform a baseline transfer from the source to the destination based on the SnapMirror policy that defines the content of the baseline and any updates. SnapMirror uses the following process to initialize a relationship:

1. Make a snapshot of the source.
2. Transfer the snapshot to the destination.
3. Depending on the SnapMirror policy attached to the relationship, it also transfers other snapshots from the source to the destination.
4. After baseline transfer, updates to this relationship occur according to the schedule assigned to the SnapMirror relationship.

The destination is a volume or consistency group that is already created and marked restricted. After SnapMirror finishes transferring the data, it brings the destination volumes online in a read-only state. While the initial data transfer is taking place, the destination is marked invalid in the output of a `vol status` command. The volumes become valid and go online after the initial transfer is complete. Now, the files and snapshots in the source volumes should be available on the destination.

Manual update to the SnapMirror relationship

The SnapMirror relationship might need a manual update either from the newest snapshot or from a specific Snapshot to prevent data loss due to an upcoming power outage, scheduled maintenance, or data migration (Figure 24 and Figure 25).

Figure 24) Start a SnapMirror relationship update.

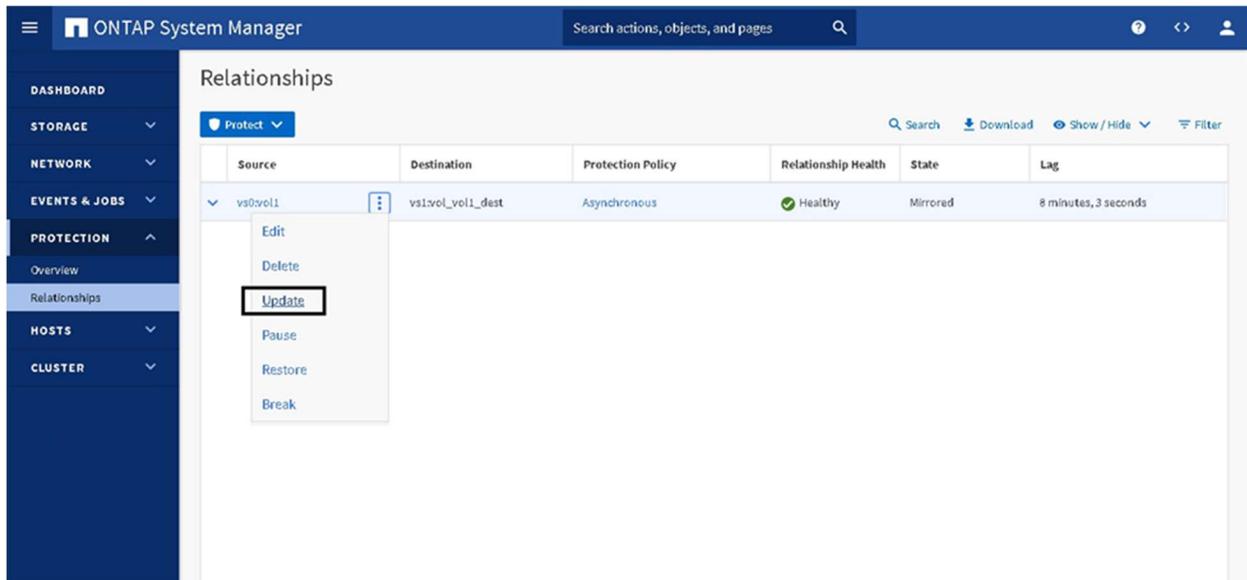
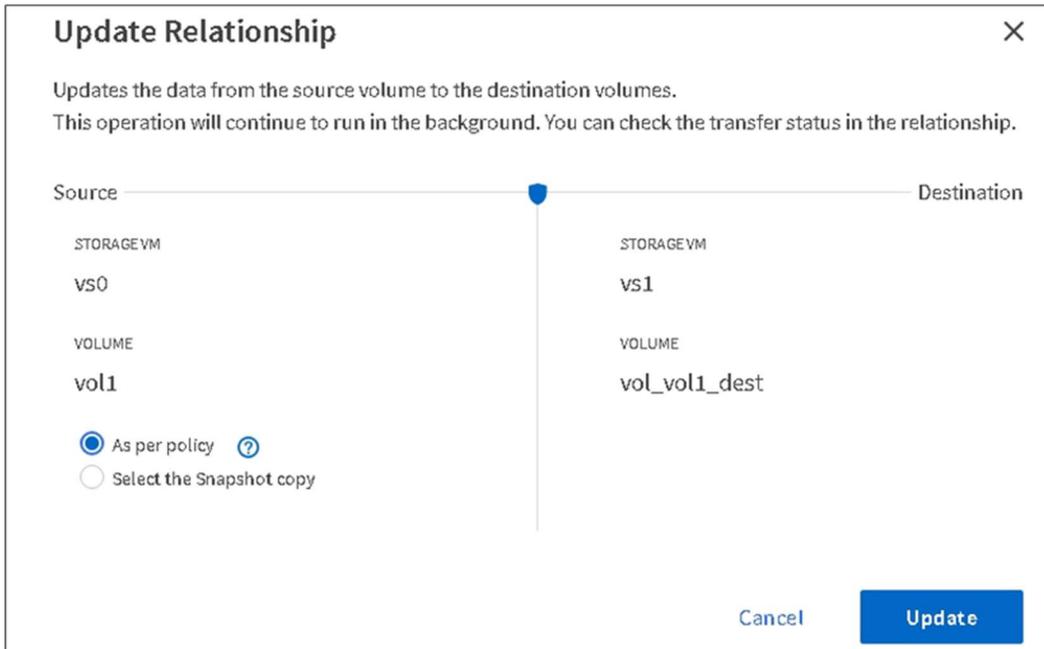


Figure 25) Relationship update dialog.



After the update completes, the Transfer State field changes from Transferring to Idle.

Long-term retention of snapshots

Snapshot policies on primary systems are designed to create scheduled snapshots and retain a desired number of historical snapshots on the primary system. Using a vault or mirror-vault SnapMirror policy, these snapshots can be replicated by SnapMirror replication to the secondary system. The SnapMirror vault-type policy uses a series of replication rules that inform SnapMirror which primary system snapshots should be replicated and how many snapshots to retain on the secondary ONTAP system. Also, additional historical snapshots can be maintained on the secondary ONTAP system...one of the primary benefits of a vault-type policy. Understand though that SnapMirror does not delete these snapshots from the primary system once they are replicated, which must be controlled by the primary system's snapshot policy. This can result in unintended consumption of volume capacity on the primary system. It is also inefficient to create lots of snapshots on a primary volume, replicate those snapshots to the secondary volume and then remove them from the primary system.

It would be more efficient if you could create those snapshots directly on the replicated volume located on the secondary ONTAP cluster. This methodology of snapshot management is called Long-Term Retention snapshots (LTR snapshots). Using this methodology, you could maintain only your most often created snapshot policy on primary and replicated to secondary and maintaining a limited number of historical snapshots on each. For example, you would use the SnapMirror vault policy to keep seven daily snapshots on both sides of the SnapMirror relationship using SnapMirror replication.

For longer-term retention, you would then use LTR to create longer-term retention snapshots on the secondary ONTAP cluster only. For example, you could tell SnapMirror to create local snapshots on the DP volumes on a monthly, quarterly, and yearly basis and keep 3 monthly, 4 quarterly and 7 yearly snapshots on the secondary cluster.

Snapshots cannot be created on the secondary DP volume using the local cluster's snapshot functionality using the standard ONTAP `volume snapshot create` command. To create LTR snapshots, we must use an extension of the SnapMirror policy for relationships using the vault relationship type. Just like we create a SnapMirror policy rule to tell SnapMirror to replicate certain snapshots from the source volume to

the destination volume, we can also use policy rules to tell SnapMirror to create local snapshots on the DP volume managed by the SnapMirror relationship.

To tell SnapMirror to create local snapshots on the DP volume, the added or modified SnapMirror policy rule must have the parameter `-schedule <text>` supplied with an appropriate ONTAP cron schedule for each scheduled LTR snapshot you wish to create. Using the example of creating a monthly and quarterly LTR snapshot described earlier, the following rules could be used:

```
DEST::> snapmirror policy add-rule -vserver SVM1 -policy MyVaultPolicy -snapmirror-label monthly
-keep 12 -schedule monthly

DEST::> snapmirror policy add-rule -vserver SVM1 -policy MyVaultPolicy -snapmirror-label
quarterly -keep 12 -schedule my3monthschedule
```

There are some considerations to discuss when using LTR snapshots:

- LTR snapshots on an upstream cluster in a cascade SnapMirror use case are not supported. So, LTR cannot be used on cluster “B” in an A→B→C cascade architecture, but can be used on cluster “C.”
- Using LTR snapshots on a SnapMirror relationship that is type mirror-vault can result in LTR snapshots being deleted in the event of a DR scenario where a snapmirror resync operation is required. This is because a resync operation will result in any snapshots created after the common snapshot used to perform the resync are deleted by ONTAP as part of the resync operation. Therefore, it is strongly recommended that LTR snapshots be used only on SnapMirror relationships with a policy of type vault.

If you need data retention for regulatory reasons or for extremely long periods of time, it is recommended to use SnapLock® volumes rather than LTR snapshots. SnapLock volumes provide an additional layer of protection preventing snapshots from being modified or deleted. LTR snapshots only guarantee that the data within a snapshot is unmodified, but does not guarantee that snapshots will not be deleted.

Converting a legacy DP SnapMirror relationship to an XDP SnapMirror relationship

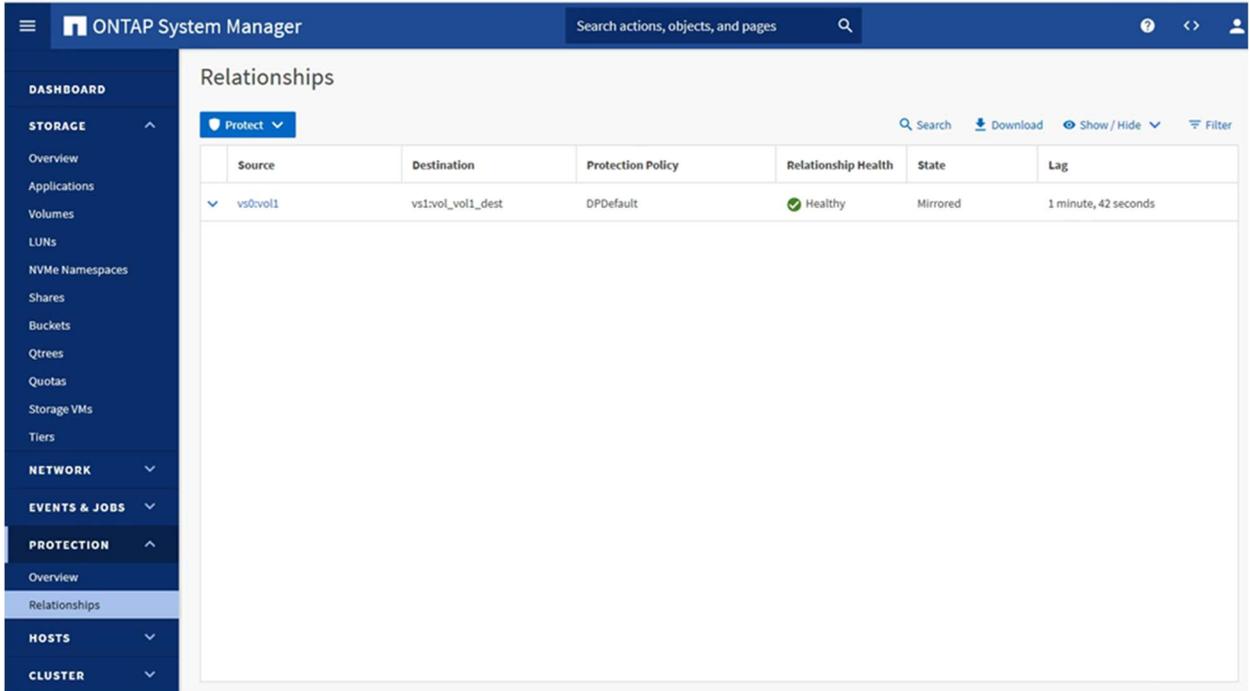
As of ONTAP 9.11.1, SnapMirror relationships of type DP are supported only for pre-existing SnapMirror relationships within the cluster, and as of ONTAP 9.12.1, SnapMirror relationships of type DP are no longer supported, and NetApp requires that existing DP relationships be converted to new XDP relationships. A summary of the process includes:

1. Break the SnapMirror volume relationship from the destination cluster.
2. Delete the SnapMirror volume relationship.
3. Create a new SnapMirror relationship using one of the built-in policies (all built-in policies use XDP relationship type), or create a custom policy with the `-type XDP` parameter, between the same endpoints with one of the default SnapMirror unified replication policies.
4. Perform a resync operation between the endpoints. This resync converts the relationship to a SnapMirror unified replication configuration without having to perform a re-baseline.

SnapMirror relationship conversion details

The following steps present this process in detail:

1. View the status of the SnapMirror relationship. The SnapMirror relationship state is shown as Mirrored.

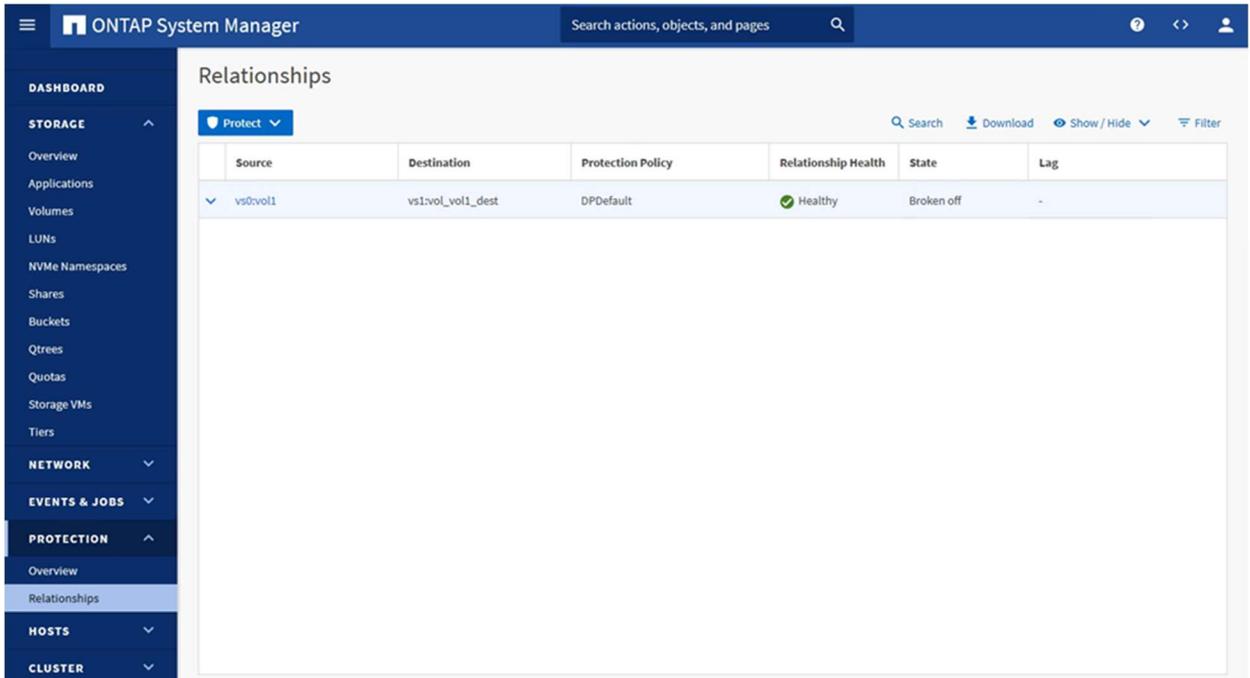


2. Perform a SnapMirror Break operation:

d. Using the CLI, run the following commands:

```
Remote::> snapmirror break -destination-path vs1:vol_vol1_dr
[Job 128] Job succeeded: SnapMirror Break Succeeded
```

e. Using System Manager, click the desired relationship then click the More > Break command. Upon completion, System Manager shows the relationship to be Broken Off in the State column.

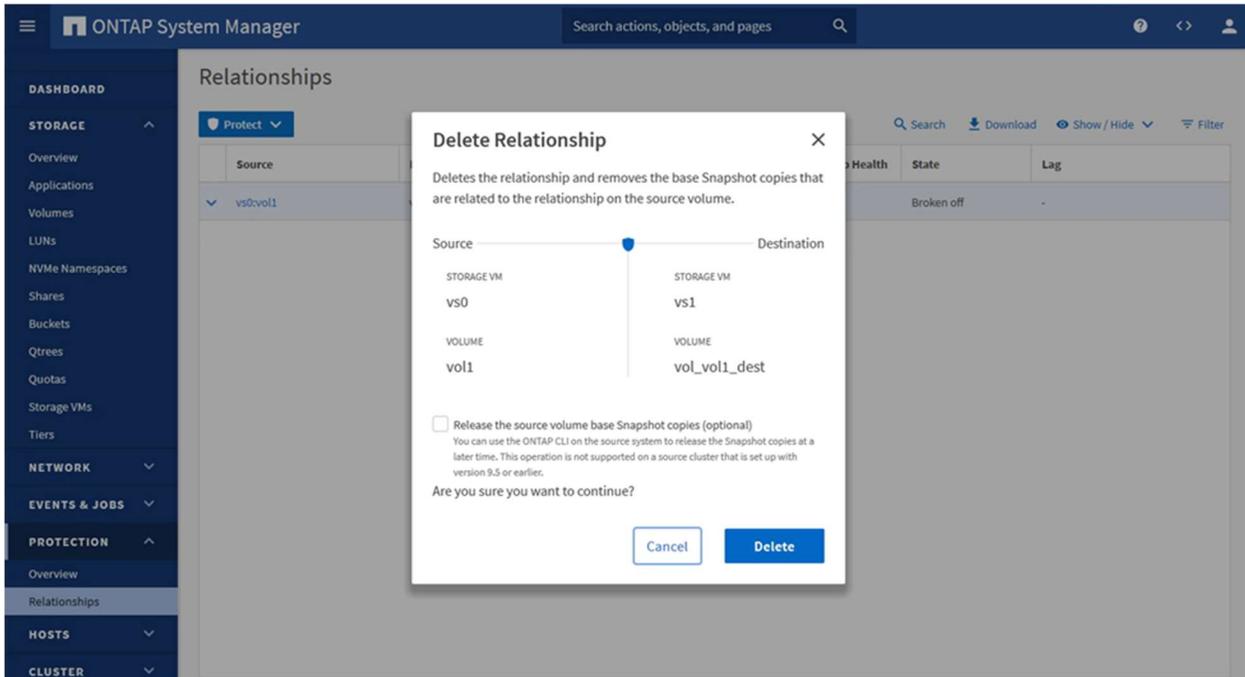


3. Delete the SnapMirror relationship:

- a. Using the CLI: Run the `snapmirror delete` command.

```
Remote:::> snapmirror delete -destination-path vs1:vol_voll_dr -relationship-info-only true
Operation succeeded: snapmirror delete the relationship with destination vs1:vol_voll_dr.
```

- b. Using System Manager, click the More > Delete command. Ensure that the Release the Source Volume Base Snapshot Copies field is unchecked before clicking Delete.



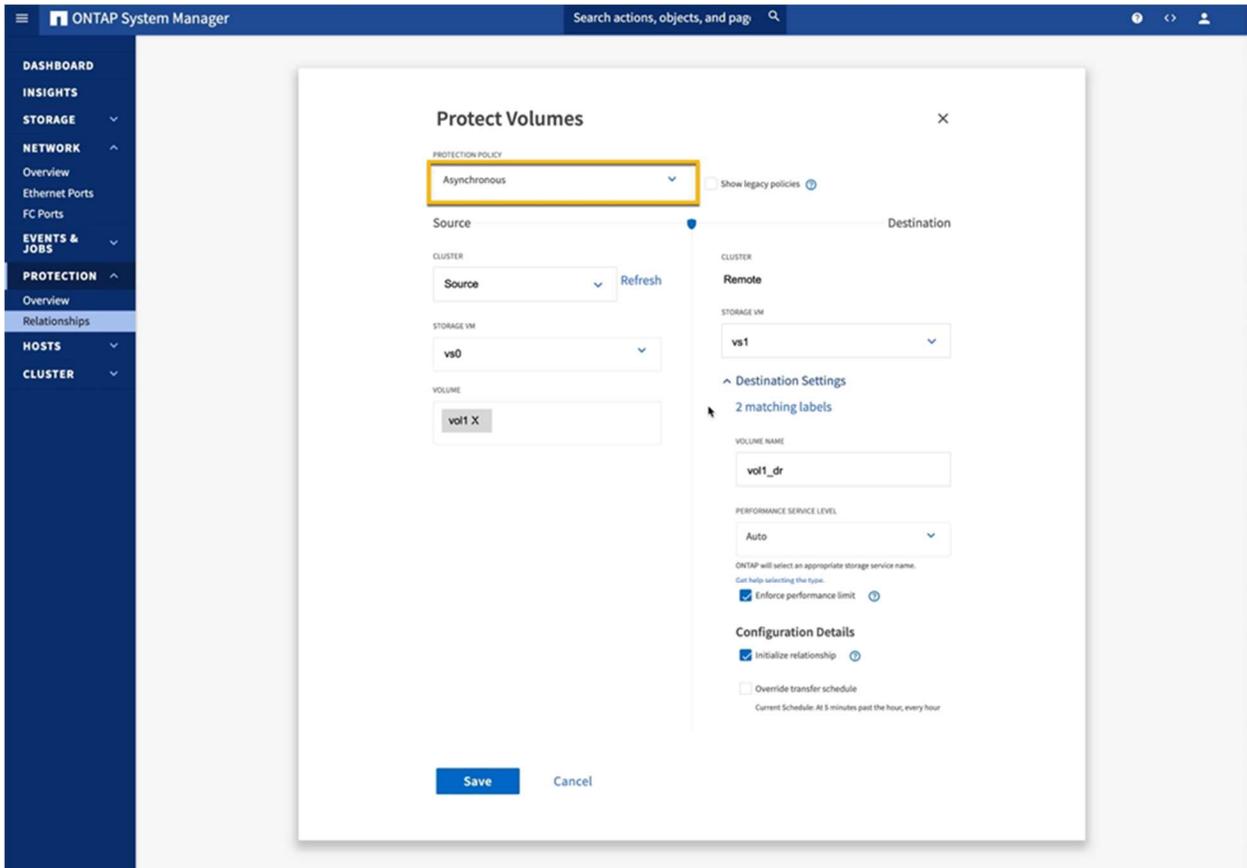
This relationship disappears from the destination cluster under Protection > Relationships.

4. Create the unified replication relationship:

- a. Using the CLI, run the `snapmirror create` command with a `-policy Asynchronous` parameter.

```
Remote:::> snapmirror create -source-path vs0:vol1 -destination-path vs1:vol_voll_dr -type XDP
-policy Asynchronous
Operation succeeded: snapmirror create the relationship with destination svm_dst1:Source_dest.
```

- b. Using System Manager, click the Protect > Volumes option and select the Protection Policy Asynchronous from the drop-down menu.

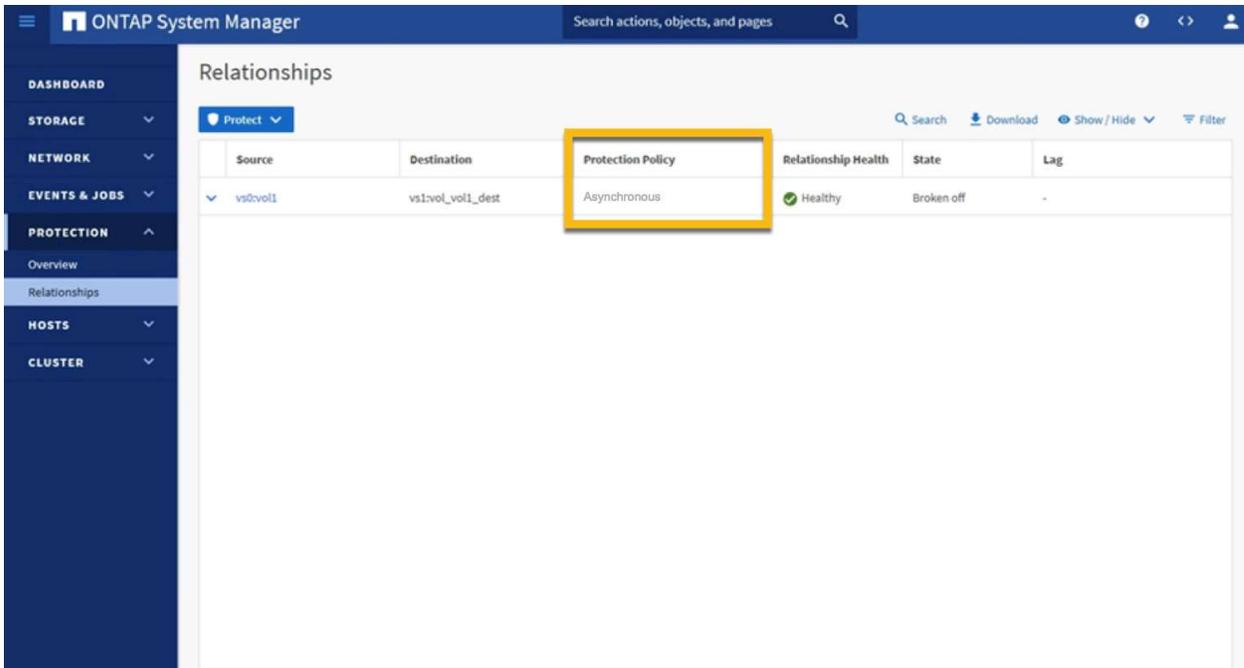


5. Verify that the new relationship is in place:
 - a. Using the CLI, run the `snapmirror show` command.

```
Remote:::> snapmirror show
```

Source Path	Destination Type Path	Mirror State	Relationship Status	Total Progress	Healthy	Progress Last Updated
vs0:vol1	XDP vs1:vol_voll_dr	Broken-off	Idle	-	true	-

- b. Using System Manager, verify that the Protection Policy is listed as Asynchronous.

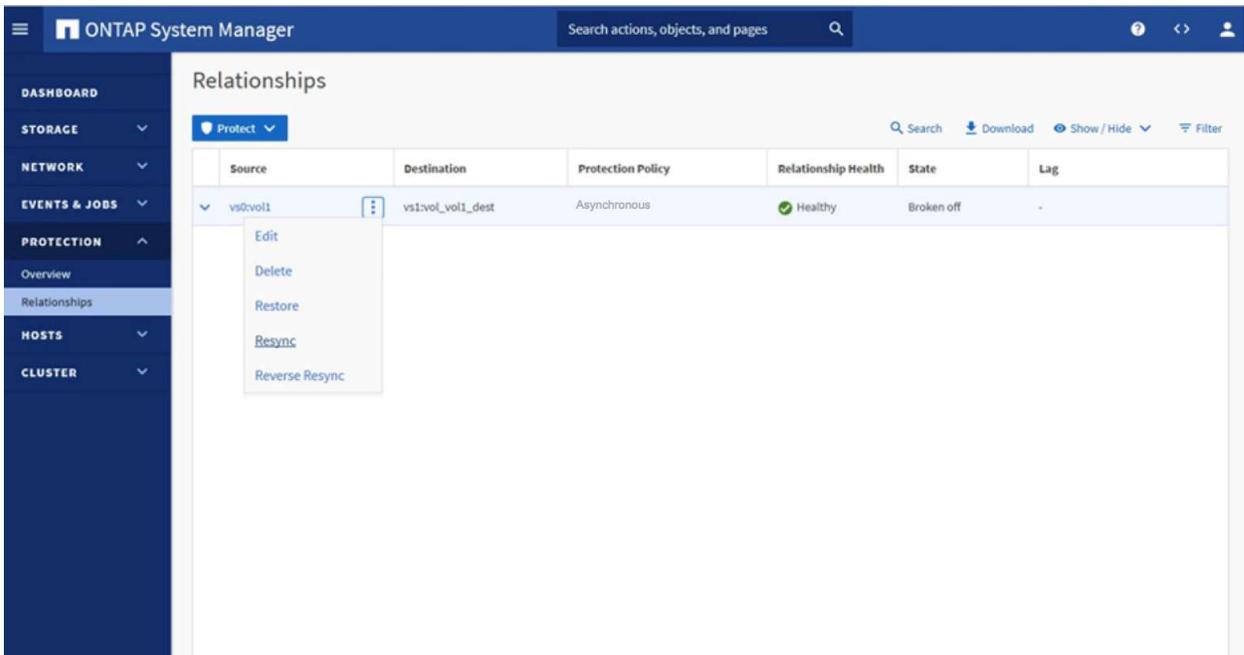


6. Resync the SnapMirror relationship:

- a. Using the CLI, run the `snapmirror resync` command Remote::> `snapmirror resync -destination-path vs1:vol_voll_dr`.

```
Warning: All data newer than Snapshot copy snapmirror.12ceb7f0-b078-11e8-baec-005056b013db_2160175149.2020-01-24_091316 on volume vs1:vol_voll_dr will be deleted.
Do you want to continue? {y/n}: y
Operation is queued: initiate snapmirror resync to destination "vs1:vol_voll_dr".
```

- b. Using System Manager, select More > Resync.



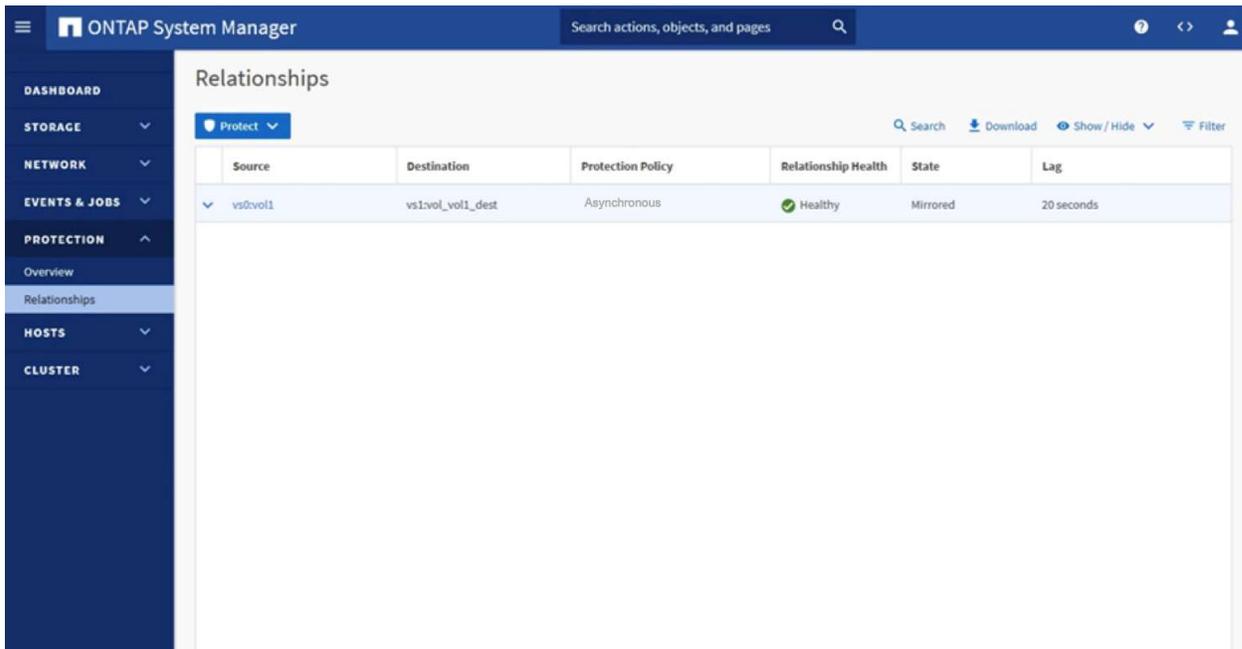
7. Verify the relationship status:

- a. Using the CLI, run the `snapmirror show` command and verify the Relationship Status column shows Idle.

```
Remote:::> snapmirror show
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Relationship Healthy	Progress Last Updated
vs0:vol1	XDP	vs1:vol_vol1_dr	Snapmirrored	Idle	-	true	-

- b. Using System Manager, verify the State column shows Mirrored and the Relationship Health column shows Healthy.



8. Verify the volume changes from type `rw` to type `dp`:

- a. Using the CLI, run the `volume show -volume vol1` command and verify that the volume type is listed as `DP`.

```
Remote:::> volume show -volume vol1
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used%
vs1	vol1_dr	data_01	online	DP	250GB	217.8GB	12%

- b. Using System Manager, verify the Type field shows Data Protection.

The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation menus for Dashboard, Storage, Network, Events & Jobs, Protection, Hosts, and Cluster. The main area displays a table of volumes. The table has columns for Name, Storage VM, Status, Capacity (available | total), Throughput (MB/s), Protection, and Type. The volume 'vol_vol1_dest' is highlighted with a red box.

Name	Storage VM	Status	Capacity (available total)	Throughput (MB/s)	Protection	Type
root_vs0	vs0	Online	94.7 MB 100 MB	0	Read/Write	Read/Write
vol1	vs0	Online	17.5 MB 20 MB	0	Read/Write	Read/Write
vol1_dr	vs1	Online	17.7 MB 20 MB	0	Read/Write	Read/Write
vol_dr	vs1	Online	19.8 MB 20 MB	0	Data Protection	Data Protection
vol_vol1_dest	vs1	Online	17.3 MB 20 MB	0	Data Protection	Data Protection
vs1_root	vs1	Online	17.7 MB 20 MB	0	Read/Write	Read/Write

SnapMirror Asynchronous ONTAP feature interoperability

SnapMirror and Snapshot copies

SnapMirror creates a snapshot before it performs a replication update. A SnapMirror snapshot is created on the source volume and a snapshot label of `sm_created` is applied. The new snapshot is then compared to the previous SnapMirror snapshot that was replicated to the data protection volume. Any data changes between the new SnapMirror snapshot and the previous one (including all snapshots on the volume between the two SnapMirror snapshots and all data in those snapshots) is replicated to the destination volume. After the SnapMirror update is complete, the new SnapMirror snapshot is exported on the destination system. SnapMirror maintains a history of one SnapMirror snapshot on the source volume and two on the destination volume.

Best practice

Verify that SnapMirror updates are not scheduled to occur on the source volume at the same time as other snapshots.

ONTAP maintains locks on snapshots created by SnapMirror to prevent them from being deleted by mistake because these snapshots are required to perform scheduled updates. If the snapshots created by SnapMirror must be deleted, the volumes can still be resynchronized. A full baseline is not required if other common snapshots between the two volumes still exist on the volumes.

In the following example, a SnapMirror resync is performed on a volume where all snapshots created by SnapMirror were deleted and uses the hourly snapshot as the base for the resync.

```
remote::> snapmirror resync -source-path cluster01://vs1/vol1 -destination-path remote://vs2/vol1
Warning: All data newer than Snapshot copy hourly.2011-12-06_1805 on volume remote://vs2/vol1
will be deleted.
Do you want to continue? {y|n}: y
```

SnapMirror and locked Snapshot copies

Introduced in ONTAP 9.12.1, locked Snapshot copies can provide indelible protection for individual snapshots without the need for a dedicated SnapLock volume. SnapMirror replicates snapshots that are marked as tamper-proof along with other snapshots based on the SnapMirror replication policy rules. The replicated locked snapshots will be replicated with either:

- The same retention time of the source tamper-proof snapshot, or
- The default retention time of the destination cluster,

whichever is longer.

Note: Though tamper-proof snapshots do not require the configuration of a SnapLock volume, it does require a SnapLock license on each cluster that implements this feature. Contact a NetApp representative for details.

SnapMirror and qtrees

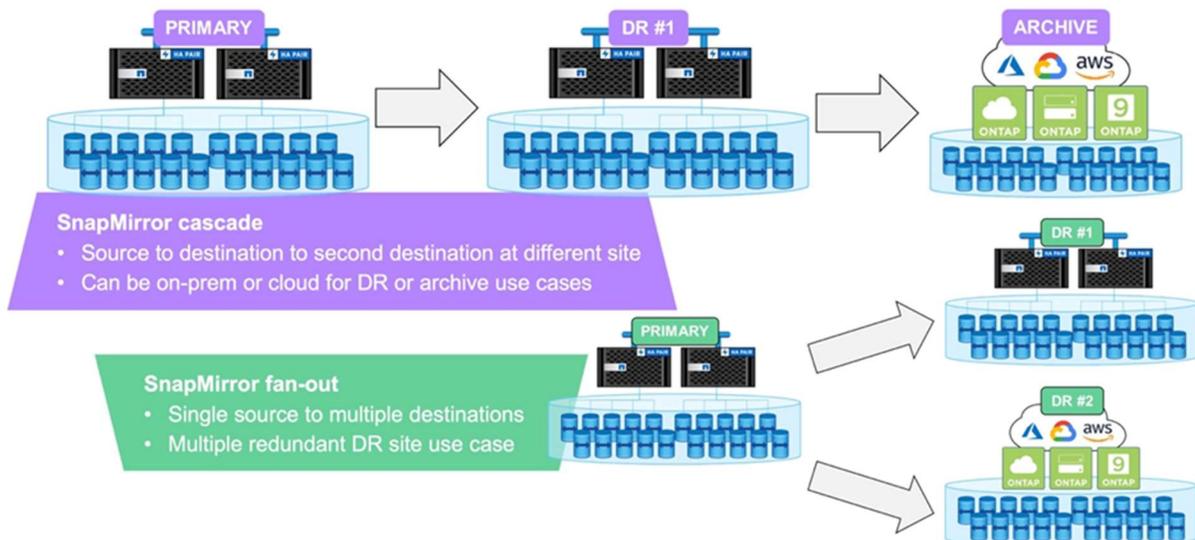
Qtrees are special directories that allow the application of file system quotas for NAS. ONTAP allows creation of qtrees, and qtrees can exist in volumes that are replicated with SnapMirror. However, SnapMirror does not allow replication of individual qtrees or qtree-level replication because SnapMirror replications operate only at the volume level.

SnapMirror and FlexGroup volumes

Starting with ONTAP 9.9.1, SnapMirror supports FlexGroup volumes as source and destination in cascade and fan-out configurations (Figure 26). Destinations can be on-premises or cloud-hosted Cloud Volumes ONTAP clusters.

Note: Both the source and destination volumes in a SnapMirror relationship involving FlexGroup volumes must be FlexGroup volumes. It is not supported to replicate a FlexVol volume source to a FlexGroup volume destination, or vice versa. If you want to convert the volumes in an existing SnapMirror relationship from a FlexVol volume to a FlexGroup volume, please review [Convert ONTAP FlexVol volume SnapMirror relationships to ONTAP FlexGroup volume SnapMirror relationships](#). It is not possible to convert a FlexGroup volume to a FlexVol volume.

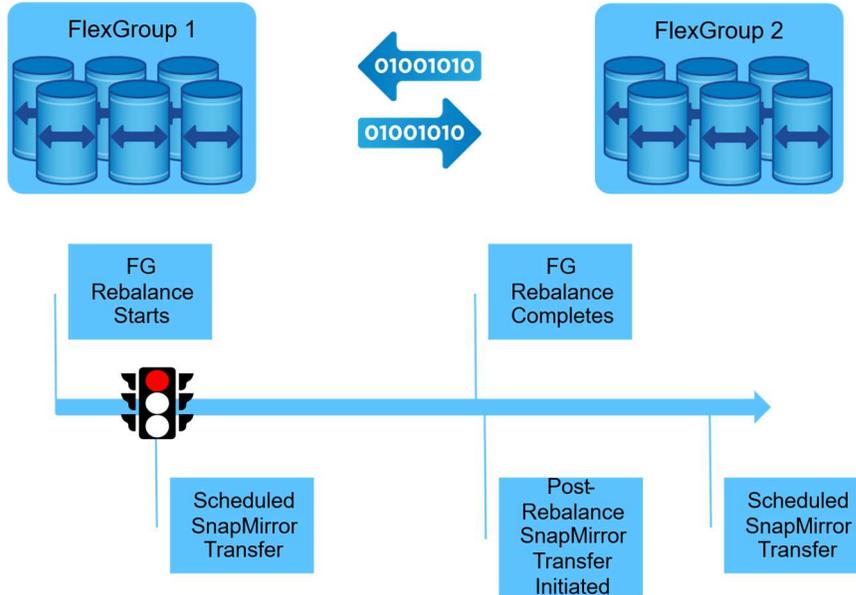
Figure 26) FlexGroup volumes used in SnapMirror cascade and fan-out configurations.



SnapMirror behavior during FlexGroup rebalance operations

Beginning with ONTAP 9.12.1, SnapMirror behavior is changed to execute the last scheduled SnapMirror operation immediately following the completion of the FlexGroup volume rebalance operation. This ensures that the RPO is minimized due to the rebalance operation preventing a scheduled SnapMirror operation from completing. This updated process is illustrated in Figure 27.

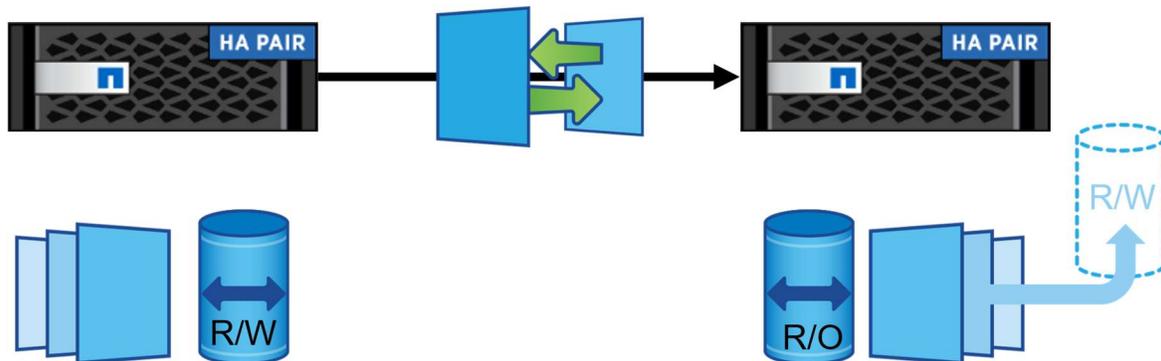
Figure 27) SnapMirror behavior during FlexGroup rebalance operation.



SnapMirror and FlexClone technologies

FlexClone technology allows storage administrators to create a writable volume from a read-only SnapMirror destination volume without interrupting the SnapMirror replication process. Although a SnapMirror relationship can be created using a FlexClone volume as the source, the SnapMirror destination volume cannot be a FlexClone volume. Figure 28 illustrates the creation of a FlexClone volume at the SnapMirror destination.

Figure 28) Creating a FlexClone volume at a SnapMirror destination.



SnapMirror replicates the snapshot history from a source to a destination volume. If a snapshot is removed from the source volume, the next SnapMirror update removes that snapshot from the destination volume. If that snapshot is the basis for a FlexClone volume, then the SnapMirror update will fail. The

only way for a SnapMirror update to proceed is to delete the FlexClone volume or split it to remove the snapshot dependency.

To avoid this issue when creating FlexClone volumes on SnapMirror destination, manually create a base snapshot required by the FlexClone volume on the source system, then replicate that snapshot to the destination system and use that snapshot as the base for the FlexClone volume, as shown in Figure 28. Using a snapshot specifically created for the FlexClone volume in this manner prevents the SnapMirror update from failing due to an automatically created snapshot being removed from the source system. The label associated with this snapshot should also not use any labels that are associated with any SnapMirror replication policy rules in effect. This ensures that the retention policies for snapshots using the preexisting labels do not attempt to delete the snapshot associated with the FlexClone volume.

SnapMirror and NetApp storage efficiency

In general, SnapMirror will always replicate data to the destination in the same state as the source data with respect to storage efficiency (SE) regardless of source and destination default SE settings. SnapMirror maintains all storage efficiency benefits in replicated volumes including deduplication, compression, and compaction. This ensures that ONTAP maximizes resource utilization for data at rest as well as over the network.

SnapMirror creates a snapshot before performing an update transfer. Any blocks in the snapshot are locked and cannot be deduplicated. Therefore, if maximum space savings from deduplication are required, run the dedupe process before performing SnapMirror updates.

Best practice

Make sure that deduplication and SnapMirror operations do not run at the same time. Start SnapMirror transfers of a deduplicated volume after the deduplication operation is complete. This prevents any effects on replication performance while deduplication is in progress and sending of non-deduplicated data and additional temporary deduplication metadata files over the network.

One aspect of ONTAP SE that can change depending on the source and destination array models is compression. Several modern array controller models support enhanced compression techniques that may not be available on the older array controller model. Once a DP volume is converted to read/write state using the `snapmirror break` command, SE for new updates and writes may change to the default for the destination array model depending on the ONTAP version for all type XDP relationships. All-flash array controllers can have temperature-sensitive SE (TSSE) enabled (AFF C-Series arrays have TSSE enabled by default).

For more information about ONTAP storage efficiency, please review [TR-3966 NetApp data compression and deduplication deployment and implementation guide](#).

Storage compression for XDP SnapMirror relationships

Table 7 provides guidance on SE behavior for SnapMirror protected volumes when the source does not have TSSE enabled, but the destination volume is on a controller that has TSSE enabled as a default (for example, A-Series with TSSE enabled or a C-Series controller).

Table 7) Replication behavior when source volume is on a non-TSSE capable controller and destination volume is on a TSSE enabled controller.

Destination ONTAP versions	Source volume compression	Replicated data (Pre-Break)	Post-break		
			Replicated data	Hot data (new data only)	Cold data (new and replicated data)
←ONTAP 9.11.1P10 ←ONTAP 9.12.1P3	No compression	No compression	No compression	standard compression	No TSSE-enabled cold data compression
	Compression enabled	Source compression	Source compression		
ONTAP 9.11.1P11 → ONTAP 9.12.1P4 → ONTAP 9.13.1 →	No compression	No compression	No compression	TSSE compression	
	Compression enabled	Source compression	Source compression		

Table 8 provides guidance on SE behavior for SnapMirror protected volumes hosted on an AFF C-Series array or AFF A-Series array with TSSE enabled on the protected volume replicating to a non-TSSE capable array model.

Table 8) Compression behavior when source volume has TSSE enabled, and destination volume controller is not on a TSSE capable.

Destination ONTAP versions	Source volume compression	Replicated data (Pre-Break)	Post-Break		
			Replicated data	Hot data (new data only)	Cold data (new and replicated data)
All	TSSE	TSSE	TSSE	FAS: Uncompressed AFF A-Series (TSSE not enabled): File compression	NA

Table 9 provides guidance on SE behavior for SnapMirror protected volumes hosted on an AFF C-Series array or AFF A-Series array with TSSE enabled on the protected volume when replicating to another TSSE capable array controller (AFF C-Series or AFF A-Series).

Table 9) Compression behavior when source volume has TSSE enabled, and destination volume controller is TSSE capable.

Destination ONTAP versions	Source volume compression	Replicated data (Pre-Break)	Post-Break		
			Replicated data	Hot data (new data only)	Cold data (new and replicated data)
All	TSSE	TSSE compression	TSSE compression	TSSE compression	

Storage compression for DP SnapMirror relationships

If SnapMirror is using legacy (`-type dp`) relationships, storage efficiency configurations cannot differ between the source and destination volumes. For example, it is not possible to compress or deduplicate the SnapMirror destination volume alone without enabling compression or deduplication on the SnapMirror source volume.

SnapMirror and volume move

The volume-move capability allows volumes to be moved nondisruptively between nodes in the cluster using the `volume move` command. The SnapMirror relationship does not have to be reconfigured or modified on the source or destination when a volume move is performed. If a volume in an intercluster SnapMirror relationship is moved, the node to which the volume is moved must have an intercluster LIF and be connected to the intercluster network to successfully perform SnapMirror updates.

The effect a volume move has on a SnapMirror relationship depends on whether the source volume or the destination volume is being moved. If a SnapMirror transfer is currently in progress and the SnapMirror source volume is being moved, then both the SnapMirror transfer, and the volume move transfer can run simultaneously. However, when the volume move cutover occurs (the moment ONTAP redirects I/O to the new volume), the active SnapMirror transfer is then momentarily interrupted and automatically continues from the source volume's new location.

For more information about volume move, see the [ONTAP documentation on this subject](#).

SnapMirror for disk shelf failure protection

SnapMirror can mirror volumes to nodes in a different HA pair on the same cluster. Mirroring a volume to a different HA pair ensures that the other volume is always placed on a different disk shelf. If mirroring to a different disk shelf on the same node, then the mirror must be on a different aggregate. There is still a risk that an aggregate might have a constituent disk from any disk shelf due to drive failure and having a spare assigned. This configuration avoids having a single point of failure and provides protection against disk shelf failure.

One caveat is that the configuration does not failover automatically. The storage administrator must manually break the SnapMirror relationship, unmount the clients, remount the clients on the destination volumes, and change the NFS export policies.

SnapMirror and volume autosize

When using SnapMirror XDP relationships, it is possible to mirror a larger volume from a source to a smaller volume on the destination due to the integrated data efficiencies native to the LRSE replication process used by XDP relationships. It is recommended that destination volumes be similar in size (or larger) than the source volume and enable the Autosize option.

Best practice

Keep the source and destination volumes the same size or slightly larger with the Auto Grow option enabled on the destination volume.

If Autosize is enabled on the source, NetApp recommends that Autosize be enabled on the destination to ensure adequate capacity for the SnapMirror transfers.

When autosize increases the size of the source volume of a SnapMirror relationship, the destination volume also automatically increases in size.

SnapMirror and NDMP

NDMP backups can be performed from either a SnapMirror source or destination volume. When a SnapMirror destination is backed up to tape with the dump engine, only the data in the volume is backed up. However, if a SnapMirror destination is backed up to tape using SMTape, then the metadata is also backed up. The SnapMirror relationships and the associated metadata are not backed up to tape. Therefore, during restore, only the data on that volume is restored, but the associated SnapMirror relationships are not restored. There are advantages to performing NDMP backups from SnapMirror destination volumes rather than from source volumes, including the following:

- SnapMirror transfers can happen quickly and with less effect on the source system. Use NetApp snapshots and perform SnapMirror replication from a primary system as a first stage of backup to significantly shorten or eliminate backup windows. Then perform NDMP backup to tape from the secondary system.
- SnapMirror source volumes are more likely to be moved to optimize the production environment than the DR volumes on the destination.

SnapMirror and Fabric Pool

SnapMirror supports replication of volumes hosted on FabricPool-enabled aggregates. When replicating volumes on a Fabric Pool aggregate, the replication interval should be set to a lower value than the FabricPool tiering policy to make sure that all data is protected.

When a SnapMirror relationship is created for a Fabric Pool enabled volume, ONTAP will replicate both the performance tier data and the capacity tier data without requiring the cool data in the capacity tier to be repositioned onto the performance tier of the source volume. Data in the capacity tier is streamed from the source volume's object store through the source cluster node directly to the IC LIF network. That cool data will temporarily be placed on the destination volume's performance tier data before immediately being queued for transfer to the destination volume's capacity tier.

While this happens very quickly, there are scenarios where cool data from the source capacity tier will arrive at the destination performance tier faster than the destination cluster can tier those blocks to the destination capacity tier. This can result in a higher-than-expected consumption of the destination aggregate hosting the Fabric Pool volume. Testing may be required to determine how your data change rate – impacting the size of each snapshot transferred – may impact the amount of Fabric Pool aggregate space needed to maintain your RPO. For an initial baseline transfer, it is recommended that you have enough reserve capacity on the aggregate to host all of the volume's data in the capacity tier.

For additional information related to Fabric Pool refer to [Technical Report 4598](#).

Note: FabricPool alone does not represent a data protection strategy.

Note: Starting with ONTAP 9.12.1, SnapMirror SVM DR supports replication of SVMs hosting both FlexGroup volumes and FabricPool volumes.

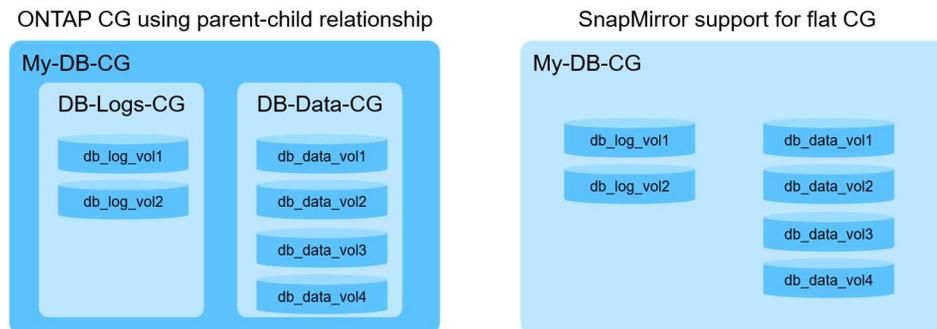
SnapMirror and consistency groups

ONTAP supports the use of consistency groups (CGs) to group multiple volumes such that operations can be applied to the group of volumes in a coherent way. This ensures that all volumes in a CG are write consistent with each other. CGs are typically used for backup of more advanced applications – such as relational databases – that utilize multiple volumes to store data for varying use cases (such as logs and data table storage) but must be in lockstep. The ONTAP implementation of CGs provides the ability to create CGs that contain CG that contain volumes...a parent-child CG configuration. This feature is used primarily for snapshot creation and SnapMirror Business Continuity configurations.

Starting with ONTAP 9.13.1, SnapMirror provides support for replicating volumes in consistency groups. For this release, only single-tier – what is referred to as flat – consistency groups are supported for SnapMirror asynchronous replication. Creating SnapMirror relationships that target CGs that contain other CGs (parent-child CGs) will result in an error. Other than this parent-child CG restriction, SnapMirror relationships targeting CG sources will support the same functionality currently offered for SnapMirror relationships created against individual volume sources unless otherwise noted in this section.

Figure 29 compares an example of a parent-child CG to a flat CG supported by SnapMirror for asynchronous replication.

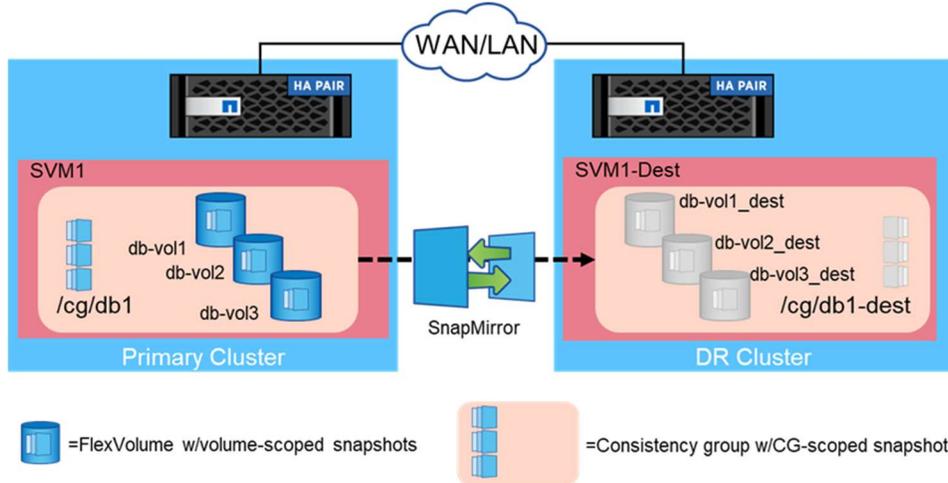
Figure 29: Comparing ONTAP CGs with SnapMirror CG support



Overview of SnapMirror with CGs

The establishment of SnapMirror relationships that target CGs is like that for creating SnapMirror relationships for individual volumes. The difference is associated with referencing source and destination CGs (instead of individual volumes) and providing a volume map between source CG component volumes and destination CG component volumes (Figure 30).

Figure 30: SnapMirror for consistency groups



SnapMirror relationships targeting CGs support all SnapMirror policy types. What snapshots will be replicated will depend on the ONTAP version and the SnapMirror replication policy. In all ONTAP versions, SnapMirror supports all SnapMirror replication policies (async-mirror, vault, and mirror-vault). For vault, mirror-vault relationships and async-mirror relationships using MirrorAllSnapshots, what snapshots get replicated will vary depending on ONTAP version.

For ONTAP 9.13.1, SnapMirror will replicate only those snapshots associated with the CG that meet the policy type and any SnapMirror snapshot label rules. It will not replicate any member volume-scoped snapshots that may meet the SnapMirror CG policy definitions.

Starting with ONTAP 9.14.1, SnapMirror will replicate all snapshots associated with the CG and the member volumes that have snapshots matching snapshot label rules that are part of the SnapMirror CG policy definition.

Creating a SnapMirror CG relationship

To create a SnapMirror relationship between two CGs, the source and destination CGs must be identified along with a mapping of the volumes in the source CG to the volumes in the destination CG. Prior to creating the SnapMirror relationship, make sure all volumes and CGs are created on both the source and destination clusters.

Note: CGs cannot be created, managed, or deleted using ONTAP CLI. Acceptable interfaces for CG management are ONTAP REST API and ONTAP System Manager.

Creating SnapMirror relationships for CGs using ONTAP CLI

Using Figure 30 as a reference, the following CLI command executed from the destination cluster will create an hourly SnapMirror replication of the source CG to the destination CG.

```
Dest:> snapmirror create -source-path SVM1:/cg/db1 -destination-path SVM1-Dest:/cg/db1-dest -cg-  
item-mappings db-vol1:@db-vol1_dest, db-vol2:@db-vol2_dest, db-vol3:@db-vol3_dest -policy  
MirroAndVault -schedule hourly
```

Creating SnapMirror relationships for CGs using ONTAP REST APIs

ONTAP REST APIs can be used to create a SnapMirror relationship from either the source cluster or the destination cluster. REST API commands can be executed from the source or destination cluster.

If the destination CG does not exist, the REST API operation will automatically create the CG. As part of provisioning the destination ONTAP CG, destination constituent volumes matching the source CG

constituent volumes are created on the destination cluster and then the destination CG is created using these new constituent volumes. The operation also establishes the SVM peering relationship, if it does not already exist, then creates the CG Async SnapMirror relationship and then optionally initializes the relationship. This operation can be initiated either on the source or on the destination cluster.

Using Figure 30 as a reference, the following REST command executed from the source cluster will create a SnapMirror replication of the source CG to the destination CG.

```
POST /api/snapmirror/relationships/
'{"source":      { "path": "SVM1:/cg/db1", "consistency_group_volumes": "db-
voll1, db-vol2, db-vol3"},
  "destination": { "cluster.name": "dest_cluster", "path": "SVM1-
Dest:/cg/db1-dest", "consistency_group_volumes": "db-voll1_dest, db-vol2_dest,
db-vol3_dest"},
  "policy":      "MirrorAllSnapshots",
"create_destination": { "enabled": "true", "storage_service": { "enabled":
"true", "name": "extreme", "enforce_performance": "true" } }'
```

Creating SnapMirror relationships for CGs using ONTAP System Manager

ONTAP System Manager can be used to create SnapMirror relationships for CG. The process for creating a SnapMirror relationship for a CG using System Manager will be the similar to the process for creating a volume scoped SnapMirror relationship except you start creating the relationship on consistency group object rather than a volume object. When System Manager detects the source of the relationship is a CG, the user will be offered a full set of pre-created or custom asynchronous SnapMirror protection policies alongside any SnapMirror Business Continuity (SM-BC) protection policies. SnapMirror will create the appropriate relationship type based on the selected protection policy.

Once the clusters and SVMs are peered, the user can protect a CG through the System Manager Consistency Group page by selecting Protect consistency group menu option from the destination cluster.

Scalability of SnapMirror relationships for CGs

SnapMirror can scale up to 50 CGs per cluster with a maximum of 80 FlexVol member volumes per CG. The maximum number of total member volumes supported per cluster is 400 FlexVol volumes. Also, the total number of CGs and FlexGroup volumes cannot exceed 50 for any cluster.

Managing SnapMirror protection of CGs

CGs add an additional layer of complexity to the data protection process since it must coordinate the replication of multiple member volumes with a single SnapMirror relationship. Each of the following sections describes how SnapMirror behaves for creation and recovery scenarios.

SnapMirror replication

When a SnapMirror scheduled replication executes, the following actions take place:

1. All constituent volumes within the CG will be fenced from all I/O operations.
2. A CG snapshot of each member volume is created.
3. The fence is then removed to permit continued I/O operations.
4. SnapMirror will replicate the changes captured in the member volumes' CG snapshots based on the SnapMirror policy as shown in xxx.

Table 10 lists which type of snapshots are replicated based on the SnapMirror policy when replicating CGs.

Table 10) What snapshots get replicated as part of a SnapMirror replication of a consistency group

SnapMirror policy type	Replicated snapshots
Async-mirror (only <code>sm_created</code> rule)	<ul style="list-style-type: none">The scheduled <code>sm_created</code> CG snapshot on each member volume
Async-mirror (<code>sm_created</code> + <code>all_source_snapshots</code> rules)	<ul style="list-style-type: none">The scheduled <code>sm_created</code> CG snapshot on each member volumeAll CG-scoped snapshots since the previous CG-scoped SnapMirror scheduled replication.All volume-scoped snapshots since previous CG-scoped SnapMirror scheduled replication.
Mirror-vault	<ul style="list-style-type: none">The scheduled <code>sm_created</code> CG snapshot on each member volumeThe CG-scoped snapshots since the previous CG-scoped SnapMirror scheduled replication that have <code>snapmirror-label</code> values matching the SnapMirror policy rules.The volume-scoped snapshots since the previous CG-scoped SnapMirror scheduled replication that have <code>snapmirror-label</code> values matching the SnapMirror policy rules.
Vault	<ul style="list-style-type: none">The CG-scoped snapshots since the previous CG-scoped SnapMirror scheduled replication that have <code>snapmirror-label</code> values matching the SnapMirror policy rules.The volume-scoped snapshots since the previous CG-scoped SnapMirror scheduled replication that have <code>snapmirror-label</code> values matching the SnapMirror policy rules.

If the snapshot schedules are very short when compared to SnapMirror replication schedules, it may be possible that not all snapshots will be available on the secondary due to `-keep` parameter settings. The SnapMirror policy's `-keep` setting for any rule is applied to the CG-scoped snapshots and the volume-scoped snapshots independently.

Note: Prior to ONTAP 9.14.1, snapshots independently created on each member volume, using a member volume's snapshot policy, are not replicated as part of the SnapMirror CG relationship. Only snapshots created as part of the SnapMirror CG relationship policy are replicated.

SnapMirror restore

The following behavior is true for restore operations to the original source CG or for a restore to a new CG (redirection).

- During a SnapMirror restore operation, all CG member volumes must be restored using the same CG snapshot.
- If this is not possible, the member volumes that cannot be restored will remain in a fenced state until all volumes are restored from the same CG snapshot.
- This fenced state for each failed member volume restore will remain until all member volumes get restored from the requested CG snapshot.

Converting existing volume relationships to CG relationships

Consider a scenario where a set of individual volumes each has a SnapMirror asynchronous relationship. You now wish to combine these volumes into a CG and then protect the entire CG using a SnapMirror asynchronous relationship.

Converting these individual relationships into a single CG relationship is possible without requiring a new baseline replication. The following high-level steps can be used to convert the SnapMirror configuration:

1. Delete each volume SnapMirror relationship.
2. Release each volume SnapMirror relationship with the `-relationship-info-only true` parameter.
3. If the source volumes are not already in a CG, then create a new CG with the desired volumes.
4. If the destination volumes are not already in a CG, create a new destination CG with the desired volumes.
5. Create the new SnapMirror relationship using the source and destination CGs
6. Resync the new SnapMirror relationship.

Changing the composition of a CG involved in a SnapMirror relationship.

In general, SnapMirror relationships for CGs supports the following lifecycle functions:

- Modifying CG snapshot policies.
- Adding volumes to source CG.
- Adding or deleting LUNs on protected constituent volumes.
- Adding or deleting NVMe namespaces on protected constituent volumes.
- Retrieving information about the SnapMirror relationship or the protected CG.
- Modifying volume attributes while being part of the protected CG.
- Resizing volumes that are constituent members of a protected CG.

Interchanging or removing constituent volumes

Consider the scenario where it may be necessary to remove member volumes or exchange (delete and add) member volumes in an existing CG that has a SnapMirror relationship in place. The following steps can be used to accomplish this task:

1. Delete the current SnapMirror relationship for the CG on the destination cluster.
2. Release the SnapMirror relationship with the `-relationship-info-only true` parameter on the source cluster.
3. Add and/or remove the desired constituent volumes from the CG
4. Recreate the new SnapMirror relationship using the source and destination CGs
5. Resync the new SnapMirror relationship.

Note: The success of the resync operation depends on the availability of a common, valid CG snapshot between the new source and destination CGs.

Adding constituent volumes to a CG

Consider the scenario where additional volumes need to be added to an existing CG with an existing SnapMirror relationship. Volumes can be added to the source CG without interrupting the SnapMirror replication relationship on the source CG. Once the new volumes have been added to the source CG, SnapMirror will perform the following activities automatically in the next scheduled or manual SnapMirror update operation:

1. SnapMirror detects the difference in CG composition between the source and the destination.
2. Destination cluster creates new volume(s) to match the added source CG volumes.
3. Adds the newly created volumes to the destination CG.
4. Execute a resync operation which will:

- a. Perform a baseline replication on new volumes.
- b. Resync original volumes to get them all to the same consistency point.

In case there is a failover to the destination before the non-disruptive expand resync completes, the CG snapshot prior to the ONTAP CG expand will be available. The new CG snapshots on the expanded ONTAP CG will be available after the successful resync operation.

SnapMirror for CGs interoperability with other ONTAP features

Table 11 defines how SnapMirror relationships for CGs interoperates with other ONTAP features

Table 11) SnapMirror interoperability with other ONTAP features

Category	Feature	Definition
Mobility	Volume move	Constituent volumes can be moved to another aggregate while in protected CG.
	FlexClone volumes	Constituent volumes can be cloned while in a protected CG.
	Volume rehost	Constituent volumes cannot be rehosted to another SVM while in a protected CG
	LUN move	Supported
	LUN clone	Supported
Data Protection	Snapshot copies	Snapshots associated with the SnapMirror CG are replicated. Starting with ONTAP 9.14.1, Snapshots created on constituent volumes independently of the CG are replicated if they match the SnapMirror policy rules.
	Volume SnapMirror	Constituent volumes can have separate volume-scoped asynchronous SnapMirror relationships while in a protected CG.
	SVM DR	Starting with ONTAP 9.14.1, SVM DR relationships can be defined for SVMs containing SnapMirror compliant CGs.
	MetroCluster	Supported
	SnapLock®	SnapLock volumes cannot be members of a protected CG
Lifecycle Management	FlexGroup volumes	FlexGroup volumes cannot be members of a protected CG
	FabricPool	Supported
	FlexCache	FlexCache is not compatible with protected CGs
	Volume encryption	Member volumes can be encrypted at the volume level
	Aggregate encryption	Member volume can reside on encrypted aggregates

SnapMirror for SVM DR

SnapMirror can be used to protect entire SVMs and is called SnapMirror SVM DR. SVM DR replicates FlexVol volumes, FlexGroup volumes, and consistency groups owned by the SVM as well as the SVM configuration and identity information to a remote destination for NAS protocols. Although SVM DR uses

the same SnapMirror replication technology as SnapMirror for volumes, there are differences (listed in Table 12).

Table 12) Differences between SVM DR and SnapMirror.

SVM DR	SnapMirror
Works at SVM-level granularity	Works at volume-level granularity
Protects FlexVol, FlexGroup, and CG data containers as well as SVM configuration	Protects only data stored in FlexVol, CG, and FlexGroup volumes
Minimum supported RPO: <ul style="list-style-type: none"> SVMs containing FlexVol volumes: 15 min SVMs containing FlexGroup volumes: 30 min SVMs containing consistency groups: 30 min 	Minimum supported RPO: <ul style="list-style-type: none"> FlexVol volumes: 5 min FlexGroup volumes: 15 min Consistency groups: 30 min

Common use cases for SVM DR include, but are not limited to, the following:

- DR of entire SVMs in case of loss of access to a cluster, site, or data center.
- Provide load balancing of data traffic between various clusters for content consumption.
- Development and testing purposes on an idle secondary site.
- Ensuring all data volumes are available for a specific tenant or application.
- Automatic protection of new FlexVol volumes, consistency groups, or FlexGroup volumes created in the protected SVM.

Note: SVM DR, when used for DR use cases, supports replication between clusters running the same version of ONTAP software. SVM DR used for use cases that do not have an expectation of recovery back to the original source cluster, such as using SVM DR for migration of SVMs and contained data volumes as part of a technology refresh, supports replication of SVMs from a source cluster up to two major versions lower than the destination cluster.

Best practice

Use [SVM data mobility](#) instead of SVM DR for migrating an SVM between clusters for reasons other than ongoing data protection. SVM data mobility is the preferred tool for the following use cases:

- Manual workload balancing for resource management
- One-time migration of an SVM for technology refresh

Defining what SVM DR will protect

SVM data replication

SVM DR can be used to protect volume data organized as FlexVol volumes, FlexGroup volumes, or multiple volumes contained within consistency groups as described in SnapMirror and consistency groups section of this technical report.

There are some additional limitations associated with SVMs containing CG:

- CG-scoped asynchronous SnapMirror relationships cannot be converted to SVM DR relationships as is possible for FlexVol volume relationships, but it can be accomplished by performing some manual steps using the ONTAP CLI:
 1. Delete the CG relationship and release the CG async relationship with relationship-info-only.
 2. Rename the destination CG volumes to that of the source SVM CG volume names.
 3. Create and initialize/resync the SVMDR relationship.
 4. Verify the SVMDR relationship is in SnapMirrored state, Idle with healthy set as true.

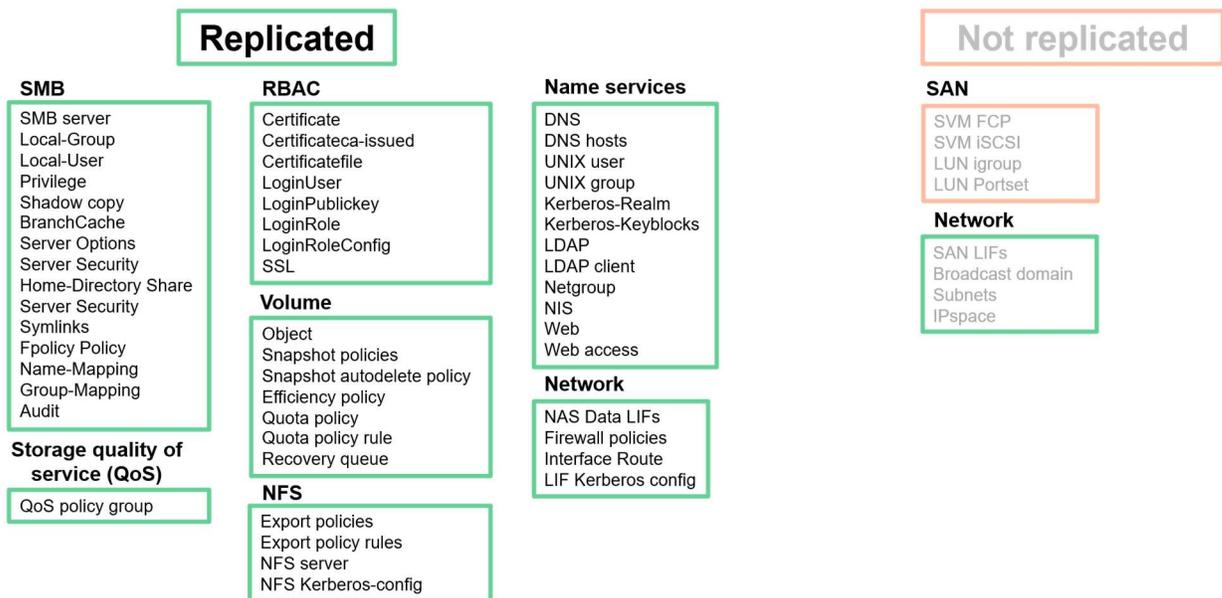
- SVM DR fan-out relationships are not supported if the source SVM contains a CG.
- SVM DR relationships from source SVMs residing on a sync-source SVM in a MetroCluster infrastructure are not supported.
- CGs cannot be excluded from SVM DR protection.
- The CGs in an SVM protected by SVM DR cannot have member volumes removed from the protected CGs.
- SVM DR will automatically create required DP member volumes and CG on the destination cluster. These volumes, and their containing CG, will have the same name as the source volumes and CG respectively.

SVM configuration information

When creating an SVM DR relationship, SnapMirror can protect just volume or CG data, or it can protect both the data and the SVM NAS configuration information. By default, SVM configuration information is not replicated.

For DR use cases, where it is desirable to enable NAS clients to reconnect to the NFS exports after the DR failover event, SnapMirror must be configured to replicate the SVM configuration information along with the SVM's volumes. To configure SnapMirror to replicate the SVM configuration information, use the `-identity-preserve true` parameter to the `snapmirror create` command. Figure 31 lists what is and is not replicated when `-identity-preserve` is true.

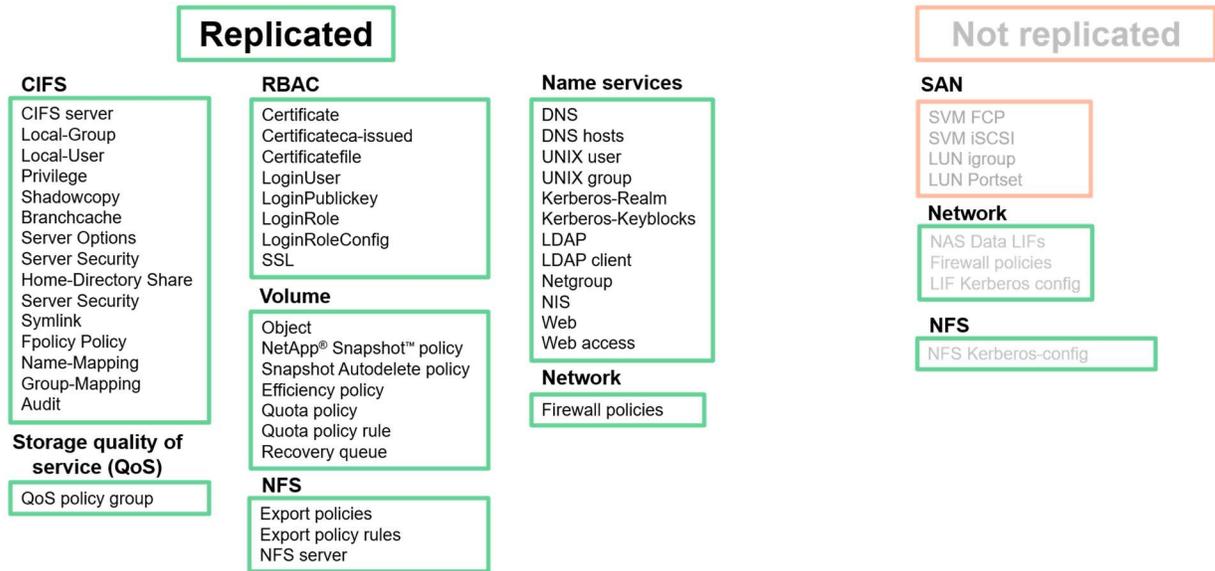
Figure 31) Replicated information for SVM DR using `-identity-preserve true`



It is also possible to replicate SVM configuration information without any of the network configuration information. This is desirable when the destination SVM is in a different IP subnet than the source SVM. To replicate SVM configuration information without network settings, a custom SnapMirror policy must be created that uses the `-discard-configs network` parameter. This parameter is only supported with SnapMirror policy type `async-mirror`. Figure 32 lists the additional information not replicated when the replication policy uses `-discard-configs network`.

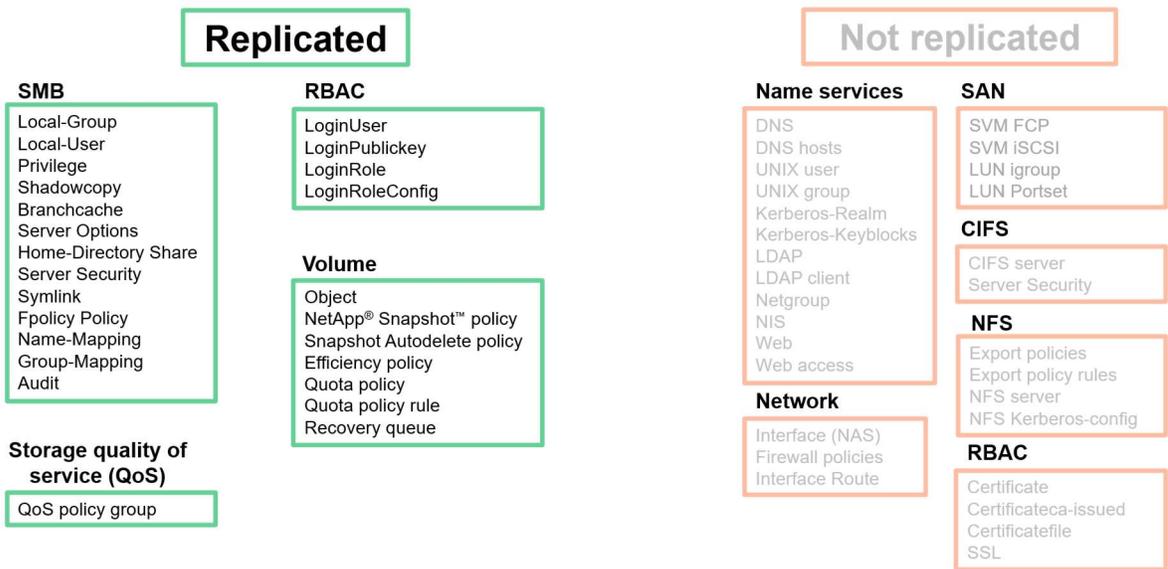
```
remote::>snapmirror policy modify -vserver vs0 -policy MySVMDRPolicy -type async-mirror -discard-configs network
```

Figure 32) Replicated information for SVM DR using `-identity-preserve true` with `-discard-configs network` policy



By default, SVM DR does not replicate the SVM configuration information. This is the equivalent of using `-identity-preserve false` during the SVM DR relationship creation. Figure 33 lists what is and is not replicated by SVM DR when `-identity-preserve` is false.

Figure 33) Replicated information for SVM DR using `-identity-preserve false`



For details regarding which SVM configuration information is replicated, see [About SnapMirror SVM replication](#).

Creating an SVM DR relationship

Creating an SVM DR relationship uses the same `snapmirror create` commands as volume scoped SnapMirror. Instead of providing source and destination volume parameters, SVM DR requires only providing the source and destination SVM names. This example demonstrates how to create an SVM DR relationship that replicates data volumes and SVM configuration information.

```
remote::>snapmirror create -source-cluster clust1 -source-vserver vs0 -destination-cluster remote
-destination-vserver vs1 -policy Asynchronous -schedule daily -type XDP
```

Prior to creating the SVM DR relationship, a target SVM must be created and designated as a data protection SVM. To designate an SVM as a data protection SVM, use the `-subtype dp-destination` parameter of the `vserver create` command.

```
remote::>vserver create -source-cluster clust1 -vserver vs1 -subtype dp-destination
```

SVM DR scalability

SVM DR has different scalability limits than volume scoped SnapMirror relationships. Table 13 and Table 14 lists the scalability limits of SVM DR data protection relationships.

Table 13) SVM DR scalability.

Parameter	Source SVM contains...	Limits
Number of SVM DR relationships per HA pair	FlexVol volumes only	Physical arrays: ONTAP 9.9.1 and earlier: 32 ONTAP 9.10.1: 64 ONTAP 9.11.1 and later: 128
		ONTAP Select (all versions): 16
	FlexGroup volumes	Physical arrays (ONTAP 9.9.1 and later): 32* ONTAP Select (all versions): 2
	Consistency groups	ONTAP 9.14.1 and later: 32*

* The sum of FlexGroup volumes and consistency groups in an SVM DR protected SVM cannot exceed 32

Table 14) SVM DR limitations when SVMs contain FlexGroup volumes or consistency groups

Parameter	Limits
Number of volumes per SVM (FlexVol volumes + FlexGroup volume constituent volumes)	300
Number of CG constituent volumes	16
Number of FlexGroup constituent volumes	20

Note: When using FlexGroup volumes and/or CGs the total number combined cannot exceed 32.

SVM DR advanced topics

Limiting the volumes replicated as part of SVM DR relationship

By default, SVM DR replicates all data volumes contained within the source SVM. This has several advantages for SVM data and identity protection. Because SVM DR has a primary use case of ensuring access to the SVM's data in the case of a source SVM failure event, it might be important to ensure that all the data volumes owned by the SVM are replicated. This default behavior also enables automatic replication of any new volumes created in the protected SVM without any further administrative steps.

There might be a scenario where it is not desirable to replicate all volumes contained within an SVM. In these cases, ONTAP supports excluding one or more data volumes from the SVM DR replication relationship.

To exclude a volume from SVM DR replication, use the `volume modify` command with the `-vserver-dr-protection unprotected` parameter for each volume to be excluded. See [Exclude volumes from SVM replication](#) for more information.

Note: CGs and any member volumes cannot be excluded from replication.

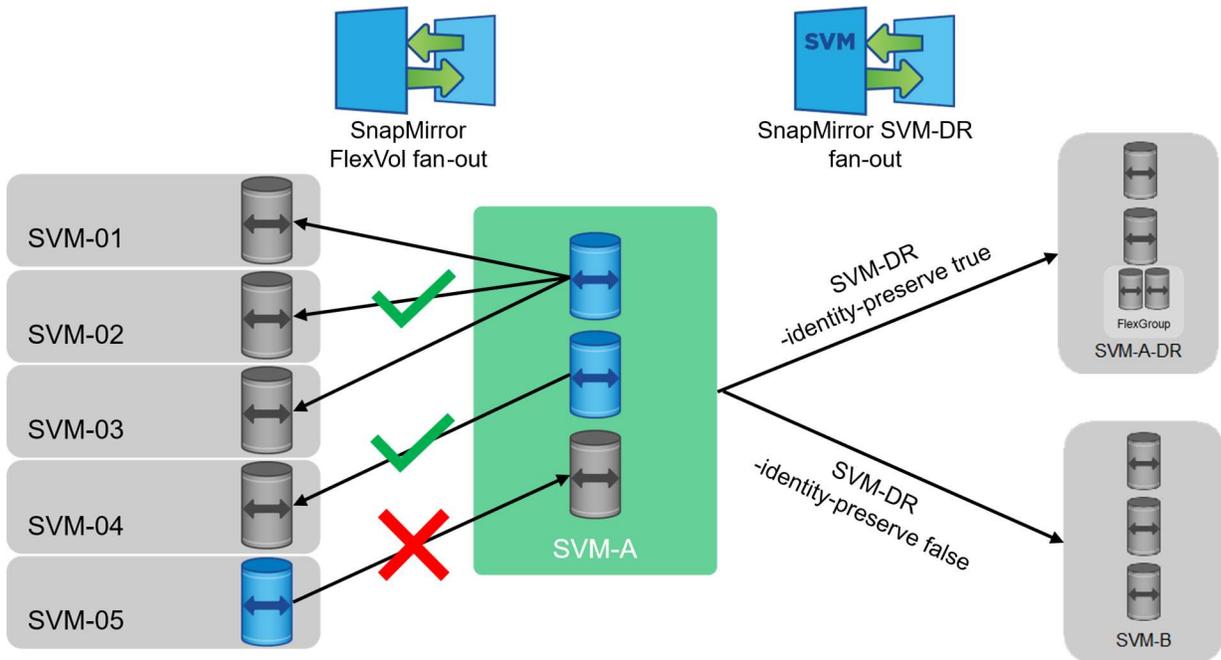
SVM DR fan-in and fan-out

SVM DR supports limited fan-out capabilities. In general, a single source SVM can be replicated to up to two destination SVMs. In this configuration one of the SVM DR relationships can be configured with the `-identity-preserve true` parameter, and the second SVM DR relationship for the source SVM must use the `-identity-preserve false` parameter as shown in Figure 34.

Additionally, the following replication rules apply:

- Individual FlexVol volumes that are contained in an SVM DR source SVM can be a source volume for additional SnapMirror relationships (e.g., volume in SVM-A to volume in SVM-04).
- Individual FlexVol volume relationships can be fan-out relationships (e.g., volume in SVM-A to volumes SVM-01/02/03).
- SVM DR does not support replication of DP volumes (e.g., volume in SVM-05 to volume in SVM-A).
- There are additional restrictions for FlexGroup volumes and consistency groups.

Figure 34) SVM DR fan-out limitations for SVMs.



SVM DR does not support fan-in of multiple source SVMs to the same destination SVM. An ONTAP cluster can be the destination for multiple SVM DR relationships assuming all SVMs have different names and configurations.

SVM DR fan-out for FlexGroup volumes

Starting with ONTAP 9.13.1, SVM DR supports SVM replication fan-out for SVMs containing FlexGroup volumes. Both the source and destination clusters must be running ONTAP 9.13.1 or later.

SVM DR fan-out for consistency groups

Fan-out of SVMs using SVM DR is not supported when the source SVM contains consistency groups.

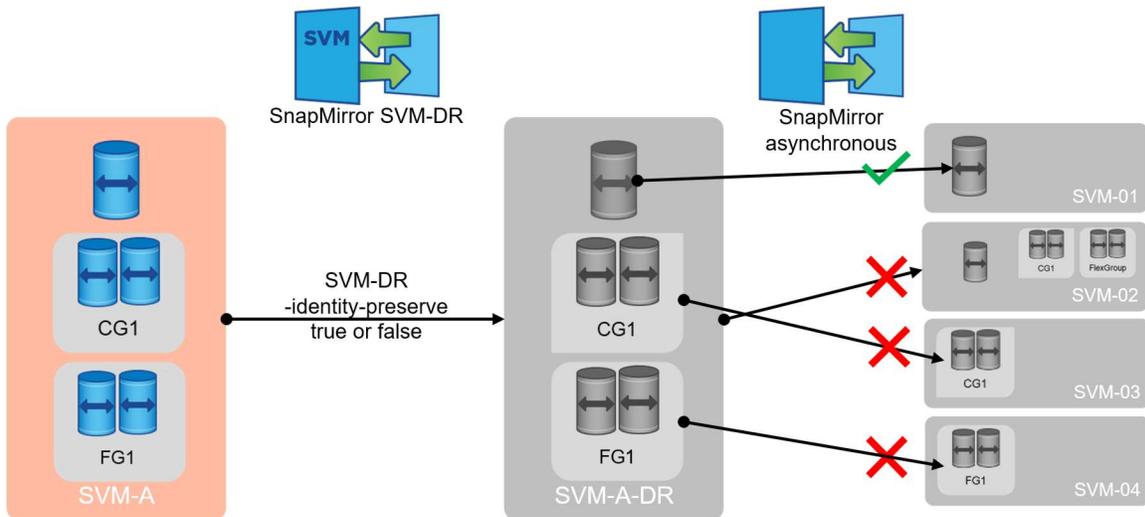
SVM DR cascade

Configuring a cascade relationship of a target SVM in an SVM DR relationship using SVM DR is not supported. Individual DP FlexVol volumes in an SVM DR destination SVM can be replicated SnapMirror asynchronous relationships. DP CGs and DP FlexGroup volumes are not allowed as sources for SnapMirror asynchronous relationships (Figure 35).

The following restrictions are imposed for secondary SnapMirror relationships from a dp-destination SVM containing CGs:

- Individual volume-scoped SnapMirror relationships are supported (e.g., single volume in SVM-A-DR to dp volume in SVM-01).
- SVM DR does not support cascade (e.g., SVM-A-DR to SVM-02).
- CG scoped SnapMirror relationships are not supported (e.g., CG1 in SVM-A-DR to CG1 in SVM-03).
- Flexgroup volume SnapMirror relationships are not supported (e.g., FG1 in SVM-A-DR to FG1 in SVM-04).

Figure 35) SVM DR cascade support



Volumes contained in the destination SVM of an SVM DR relationship that is using `-identity-preserve true` are not accessible because the destination SVM is not running as its configuration would conflict with the source SVM on the network (due to the SVM configuration data being the same on both the source and destination SVMs). This prevents SnapMirror from accessing the volumes as a source for a downstream replication relationship.

SVM DR interoperability

SVM DR and Fabric Pool

Starting with ONTAP 9.12.1, SnapMirror SVM DR supports replication of SVMs hosting FabricPool volumes.

SnapMirror supports replication of SVMs containing volumes hosted on FabricPool-enabled aggregates. When replicating these SVMs, the replication interval should be set to a lower value than the FabricPool tiering policy to make sure that all data is protected.

While replication of capacity tier data happens very quickly, there are scenarios where cool data from the source volumes' capacity tiers will arrive at the destination SVM's volumes' performance tier faster than the destination cluster can tier those blocks to the destination capacity tier. This can result in a higher-than-expected consumption of the destination aggregate(s) hosting the SVM's Fabric Pool volume(s). Testing may be required to determine how your data change rate – impacting the size of each volume's snapshot transferred – may impact the amount of aggregate space needed to maintain your RPO. For an initial baseline transfer, it is recommended that you have enough reserve capacity on the aggregate to host all the SVM's Fabric Pool volume's data in the capacity tier.

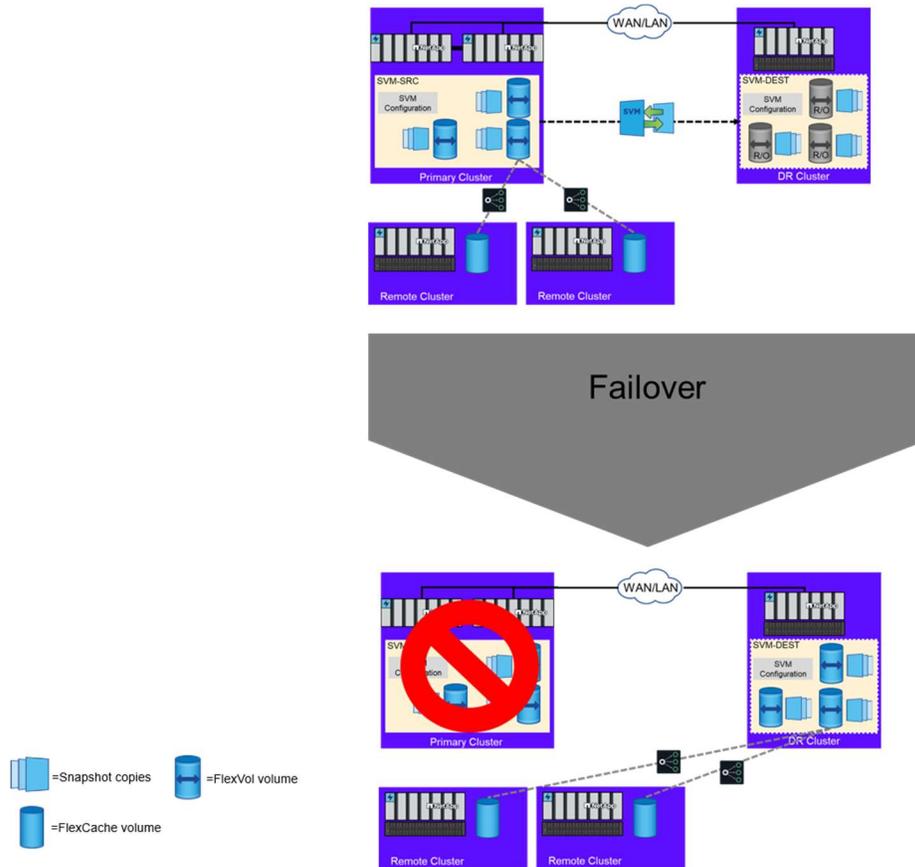
For additional information related to Fabric Pool refer to [Technical Report 4598](#).

SVM DR and FlexCache

Starting with ONTAP 9.18.1, SnapMirror SVM DR relationships support protection of SVMs containing FlexCache origin volumes. To use this feature, all ONTAP clusters participating in the SnapMirror SVM DR relationship and the FlexCache relationship must be running ONTAP 9.18.1.

Upon execution of a `snapmirror break` command of the SVM hosting the FlexCache origin volume, each dependent FlexCache relationship will be redirected to the newly activated FlexCache origin volume on the destination DR SVM (Figure 36).

Figure 36) SVM DR support for FlexCache origin volumes



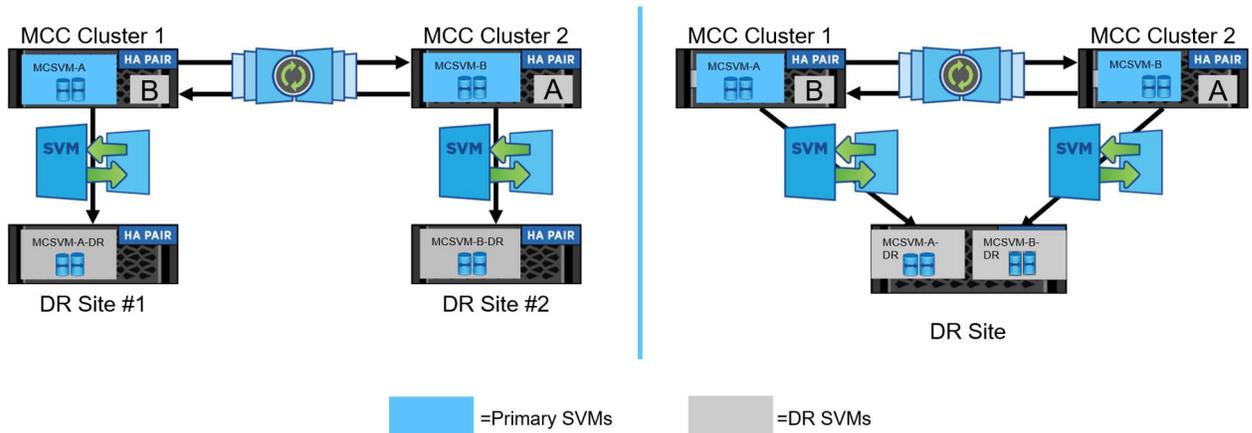
SVM DR and SAN protocols

SVM DR supports replication of iSCSI and FC LUNs and NVMe namespaces hosted on replicated volumes. SAN configuration information (initiator IDs, LUN IDs, igroups or SAN LIFs) are not replicated as part of the replicated SVM configuration data set. Additional configuration on the destination cluster, and potentially on clients accessing the namespaces or LUNs, is required.

SVM DR and MetroCluster

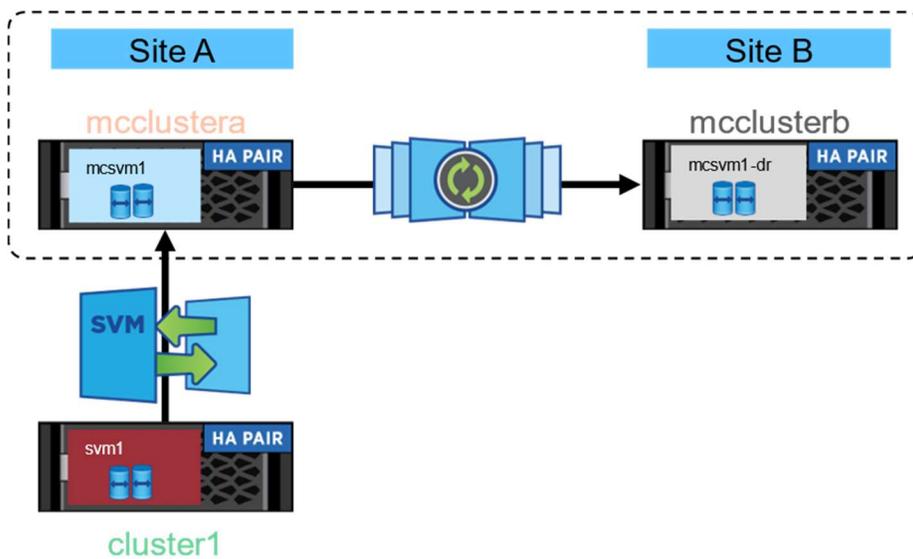
As of ONTAP 9.11.1, SVMs residing on either end of a MetroCluster relationship can be used as a source SVM for SVM DR. The combination of MetroCluster and SVM DR can be used to build a three-site DR topology that provides continuous availability at local or metropolitan distances and DR capabilities at longer distances with the ability maintain the SVM identity across all three sites, as shown in Figure 37.

Figure 37) SVM DR with MetroCluster.



Starting with ONTAP 9.16.1, SVM DR can also be used to migrate an SVM from a stand-alone HA cluster to a MetroCluster infrastructure sync-source cluster (Figure 38).

Figure 38) SVM DR used to migrate an SVM from an HA cluster to a MetroCluster sync-source cluster



Best practice

If all clusters are running ONTAP 9.16.1 or later, and the SVM contains only FlexVol volumes, the SVM data mobility tool (`vserver migrate`) can now be used to move an SVM in the following scenarios:

- HA cluster to MetroCluster
- MetroCluster to HA cluster
- MetroCluster to MetroCluster

Best practice

Use SVM DR for migration of SVMs when the new destination is part of a MetroCluster source cluster and the ONTAP version is older than 9.16.1 on any cluster involved in the migration. See [TR-4966](#) for details on how to implement this migration using SVM DR.

SVM DR support for FlexGroup volumes

Starting with ONTAP 9.9.1, SVM DR can be used with FlexGroup volumes. With this support, SnapMirror SVM relationships transferred to the destination cluster will retain full awareness of FlexGroup volumes and properly mount those volumes in a DR scenario (Figure 39).

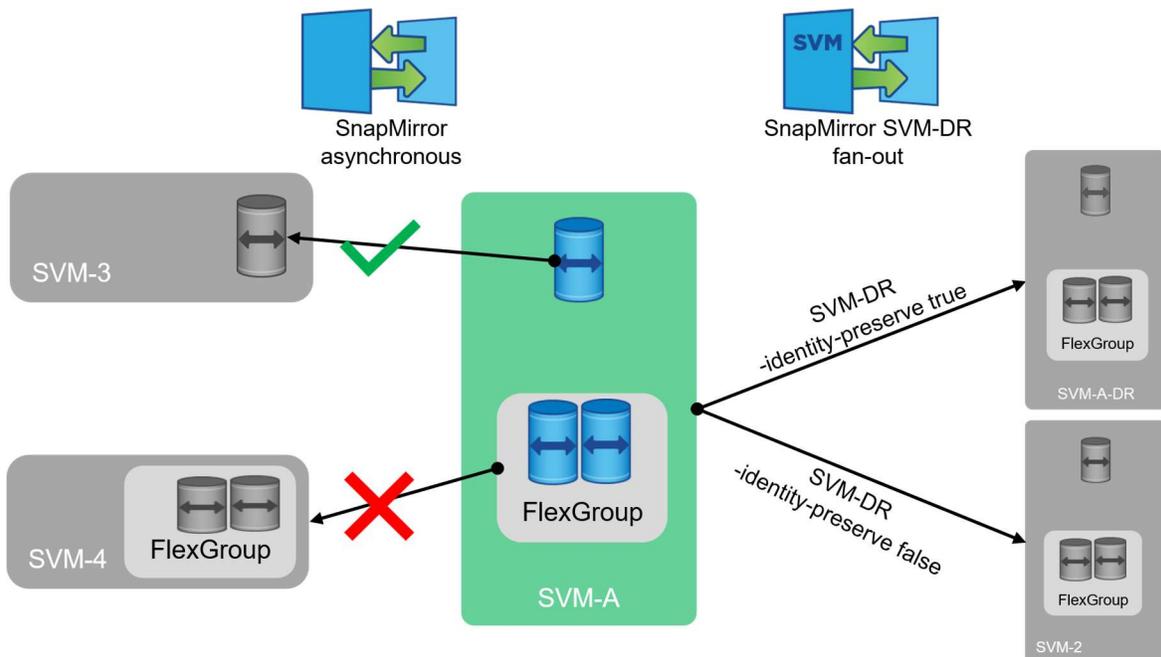
The following features are not supported for SVM DR using FlexGroup volumes:

- FlexClone volumes on source or destination clusters
- SnapMirror cascade configurations
- Converting FlexVol volumes to FlexGroup volumes

Starting with ONTAP 9.12.1, SnapMirror SVM DR supports replication of SVMs hosting both FlexGroup volumes and FabricPool volumes.

Starting with ONTAP 9.13.1, SnapMirror supports SVM DR fan-out configurations for source SVMs containing FlexGroup volumes.

Figure 39) SVM DR support for FlexGroup volumes



When a FlexGroup volume in a source SVM is modified, such as when a FlexGroup volume is expanded or contracted (add or remove constituent volumes), the changes are replicated to the destination SVM by the SVM DR relationship.

SVM DR uses a centralized ONTAP volume balanced placement algorithm to determine the location of replicated FlexVol volumes and FlexGroup volumes. The constituent volumes of a FlexGroup are distributed across all nodes and available suitable aggregates within the destination cluster.

FlexGroup volumes in a source SVM can be flagged for non-protection similar to FlexVol volumes using the `-vserverdr-protection` parameter of the FlexGroup volume.

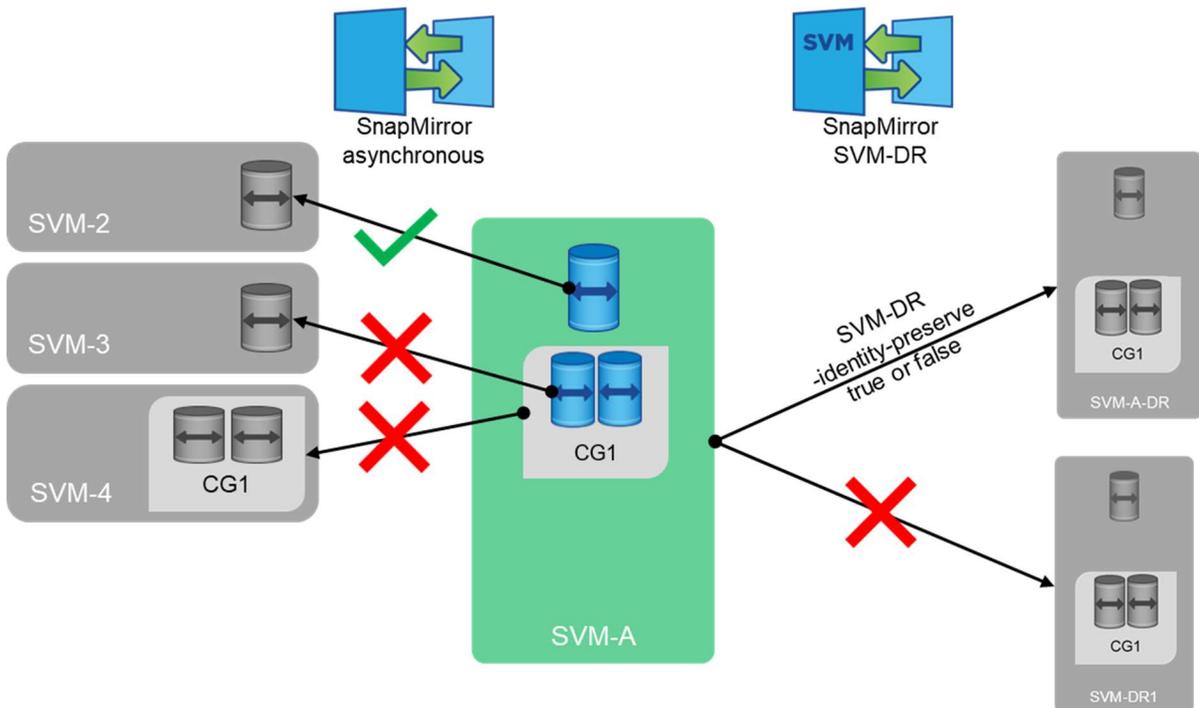
SVM DR support for consistency groups

Starting with ONTAP 9.14.1, SVM DR can be used to protect SVMs containing consistency groups (CGs).

The following limitations apply to CG-scoped SnapMirror relationships (Figure 40):

- SVM DR fan-out is not supported for SVMs containing CGs (e.g., SVM-A to SVM-DR1).
- CGs in an SVM DR protected source SVM cannot be replicated using a separate CG-scoped SnapMirror asynchronous relationships (e.g. volume in SVM-A replicating to volume in SVM-2).
- The member volumes in a CG that is in an SVM DR protected source SVM cannot be replicated using a separate volume-scoped SnapMirror relationship (member volume of CG1 to volume in SVM-3).

Figure 40) SVM DR support for consistency groups



SVM DR and cloud target interoperability

SVMs and SnapMirror are supported to some extent by NetApp Cloud Volumes ONTAP implementations offered by major cloud providers, the extent of SVM DR support for NetApp volume services offered by various cloud providers might be limited or not supported. It is imperative that customers work with their cloud provider to understand the full extent of support for SVMs, SnapMirror, and SVM DR.

SVM DR is not supported by the SnapMirror Cloud API and the NetApp Backup and Recovery service.

Performance

There are multiple factors that can affect the performance of replication:

- Node CPU utilization
CPUs will be shared between various data operations such as application data access and data protection operations.
- Number of concurrent SnapMirror operations

Each transfer operation takes additional CPU cycles and network bandwidth to move data. The fewer concurrent transfers occurring at a given time, the faster each transfer operation will complete. The supported number of concurrent transfers supported will depend on node model and ONTAP version. For more information, see [NetApp Hardware Universe](#).

- Type of transfer: initialization or update

A new SnapMirror relationship requires a baseline snapshot to be transferred that include all data in the volume at the time of relationship initialization. Subsequent updates will only transfer the differential data changes since the previous SnapMirror snapshot creation.

- Node hardware type.

The node model, configuration (including disk types, number of disks in the aggregate, number of volumes in the aggregate, and intercluster network physical port type) will directly affect SnapMirror performance.

Calculate SnapMirror throughput for performance

Throughput for a relationship can be determined based on the amount of data moved over a set period. To determine throughput, the fields to note are the Transfer Size and Transfer Duration. To find the transfer throughput, divide the transfer size by the transfer duration.

```
cluster::> snapmirror show -destination-path vs3:dst -instance

                Source Path: vs1:src_test
                Destination Path: vs3:dst
                Relationship Type: DP
                Relationship Group Type: none
                SnapMirror Schedule: -
                SnapMirror Policy Type: async-mirror
                SnapMirror Policy: DPDefault
                Tries Limit: -
                Throttle (KB/sec): unlimited
                Mirror State: Snapmirrored
                Relationship Status: Transferring
                File Restore File Count: -
                File Restore File List: -
                Transfer Snapshot: snapmirror.89659724-bd35-11e4-9f11-
000c299bf0b8_2147484674.2015-03-02_134417
                Snapshot Progress: 0B
                Total Progress: 0B
                Network Compression Ratio: 2:1
                Snapshot Checkpoint: 0B
                Newest Snapshot: snapmirror.89659724-bd35-11e4-9f11-
000c299bf0b8_2147484674.2015-02-25_134212
                Newest Snapshot Timestamp: 02/25 13:22:08
                Exported Snapshot: snapmirror.89659724-bd35-11e4-9f11-
000c299bf0b8_2147484674.2015-02-25_134212
                Exported Snapshot Timestamp: 02/25 13:22:08
                Healthy: true
                Unhealthy Reason: -
                Constituent Relationship: false
                Destination Volume Node: vsim
                Relationship ID: d8b4cbc8-bd36-11e4-9f11-000c299bf0b8
                Current Operation ID: 46da2fc6-c125-11e4-9f1a-000c299bf0b8
                Transfer Type: update
                Transfer Error: -
                Current Throttle: unlimited
                Current Transfer Priority: normal
                Last Transfer Type: initialize
                Last Transfer Error: -
                Last Transfer Size: 240GB
                Last Transfer Network Compression Ratio: 3.1:1
                Last Transfer Duration: 02:13:32
                Last Transfer From: vs1:src_test
                Last Transfer End Timestamp: 02/25 13:42:15
                Progress Last Updated: 03/02 13:44:17
                Relationship Capability: 8.2 and above
```

```
Lag Time: 120:22:10
Number of Successful Updates: 0
Number of Failed Updates: 0
Number of Successful Resyncs: 0
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
Total Transfer Bytes: 245760
Total Transfer Time in Seconds: 3
```

SnapMirror and network compression

With increasing network bandwidth costs and increasing data growth, customers must do more with less. As the amount of data to be protected increases, more network bandwidth is needed to maintain the same RPO. Otherwise, replication times increase as the amount of data sent over the network to the DR site increases.

The SnapMirror native network compression feature can cut down on the amount of data replicated over the network. It also offers more flexibility and choices, as described in the following section.

Maintaining the same RPO level

- **Challenge.** Data replication needs are growing and additional bandwidth is needed to maintain the same level of RPO.
- **Solution.** By using network compression, it is possible to maintain the same RPO without purchasing additional network bandwidth.

Improve Your RPO without buying additional bandwidth

- **Challenge.** Network bandwidth is fully utilized. However, the customer wants to reduce their exposure to data loss and improve their RPO.
- **Solution.** Enabling network compression improves the RPO without purchasing more network bandwidth.

Use the network bandwidth for other purposes

- **Challenge.** Data replication is consuming all available network bandwidth, but there are other applications and services being introduced that also require network bandwidth.
- **Solution.** By using network compression, it is possible to reduce the bandwidth consumed by SnapMirror without sacrificing RPO, thereby freeing up network bandwidth for other purposes.

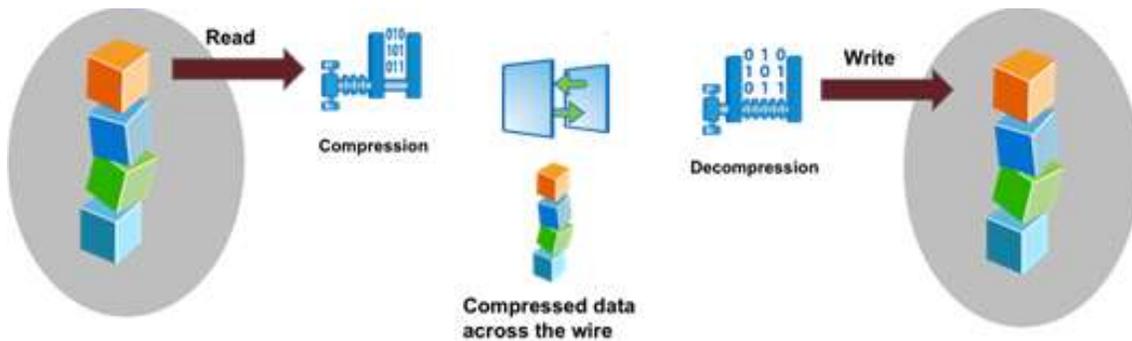
Speeding up the initial transfers

- **Challenge.** Initial SnapMirror transfers can be large and therefore can take a long time to complete under bandwidth constraints.
- **Solution.** Network compression can speed up the initial SnapMirror transfers.

What Is SnapMirror network compression?

Network compression is natively built into SnapMirror to enable increased efficiency over the network for SnapMirror transfers. It does not, however, compress data at rest. SnapMirror network compression is not the same as volume compression. Figure 41 shows a very high-level flow of SnapMirror network compression.

Figure 41) SnapMirror network compression functional diagram.



On the source system, the data blocks that must be sent to the destination system are handed off to the compression engine, which compresses the data blocks. The compression engine on the source system creates several threads, depending on the number of CPUs available on the storage system. These compression threads help to compress data in parallel. The compressed blocks are then sent over the network.

On the destination system, the compressed blocks are received and decompressed in parallel by using multiple threads. Decompressed data is then written to the appropriate volume.

Enable and disable network compression

SnapMirror network compression can be enabled or disabled by the `-is-network-compression-enabled` option in SnapMirror policy. It cannot be enabled for an active transfer. To enable compression for an existing transfer, first abort outstanding transfers, set the `-is-network-compression-enabled` option to `true` in the SnapMirror policy, and then resume the transfer.

Best practice

SnapMirror network compression increases resource utilization on both the SnapMirror source and destination systems. Therefore, evaluate the resource usage and benefits before deploying compression. For example, compression might not be useful for a high-bandwidth, low-latency connection. But it can be useful for connections that have relatively low bandwidth, such as WAN connections.

Reporting the compression ratio

The SnapMirror network compression ratio is reported in the `snapmirror show -instance` output.

```
cluster::> snapmirror show -destination-path vs3:dst -instance
      Source Path: vs1:src_test
      Destination Path: vs3:dst
      Relationship Type: DP
      Relationship Group Type: none
      SnapMirror Schedule: -
      SnapMirror Policy Type: async-mirror
      SnapMirror Policy: DPDefault
      Tries Limit: -
      Throttle (KB/sec): unlimited
      Mirror State: Snapmirrored
      Relationship Status: Transferring
      File Restore File Count: -
      File Restore File List: -
```

Compression ratio is only shown in transferring state.

```

Transfer Snapshot: snapmirror.89659724-bd35-11e4-9f11-
000c299bf0b8_2147484674.2015-03-02_134417
Snapshot Progress: 0B
Total Progress: 0B
Network Compression Ratio: 2:1
Snapshot Checkpoint: 0B
Newest Snapshot: snapmirror.89659724-bd35-11e4-9f11-
000c299bf0b8_2147484674.2015-02-25_134212
Newest Snapshot Timestamp: 02/25 13:22:08
Exported Snapshot: snapmirror.89659724-bd35-11e4-9f11-
000c299bf0b8_2147484674.2015-02-25_134212
Exported Snapshot Timestamp: 02/25 13:22:08
Healthy: true
Unhealthy Reason: -
Constituent Relationship: false
Destination Volume Node: vsim
Relationship ID: d8b4cbc8-bd36-11e4-9f11-000c299bf0b8
Current Operation ID: 46da2fc6-c125-11e4-9f1a-000c299bf0b8
Transfer Type: update
Transfer Error: -
Current Throttle: unlimited
Current Transfer Priority: normal
Last Transfer Type: initialize
Last Transfer Error: -
Last Transfer Size: 240KB
Last Transfer Network Compression Ratio: 1:1
Last Transfer Duration: 0:0:3
Last Transfer From: vs1:src_test
Last Transfer End Timestamp: 02/25 13:42:15
Progress Last Updated: 03/02 13:44:17
Relationship Capability: 8.2 and above
Lag Time: 120:22:10
Number of Successful Updates: 0
Number of Failed Updates: 0
Number of Successful Resyncs: 0
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
Total Transfer Bytes: 245760
Total Transfer Time in Seconds: 3

```

SnapMirror throttling

The SnapMirror throttle setting is used to throttle the network bandwidth consumed, which limits the amount of bandwidth used by intercluster SnapMirror. In other words, SnapMirror throttle does not control network bandwidth. Rather, it works by limiting the blocks that WAFL can use for SnapMirror transfers.

Note: All replication throttles in ONTAP are in kilobytes per second.

SnapMirror throttle can be set on a per relationship basis when creating a new relationship by using the `-throttle` option and by modifying an existing relationship with the `snapmirror modify` command. In this example, a 10MB throttle is applied to an existing relationship by using the `snapmirror modify` command.

```
cluster02::> snapmirror modify -destination-path vs1:vol1 -throttle 10240
```

Note: To change the throttle of an active SnapMirror relationship, terminate the existing transfer and restart it to use the new value. SnapMirror restarts the transfer from the last restart checkpoint by using the new throttle value, rather than restarting from the beginning.

Note: Intracluster throttle is supported, and it works the same way as intercluster throttle.

ONTAP 9 introduces global SnapMirror throttling available for each node in a cluster to perform SnapMirror transfer at a fixed maximum bandwidth for outgoing and incoming transfers. SnapMirror global throttling restricts the bandwidth used by incoming and/or outgoing SnapMirror transfers. The restriction is enforced cluster wide on all nodes in the cluster. This capability is in addition to the throttle for each

SnapMirror relationship as described earlier. Each node has a global throttle for sender-side (outgoing) transfers as well as receiver-side (incoming) transfers and an option to enable or disable this throttling. The per-transfer throttle is capped at the node-level throttle if it exceeds the global node throttle value. Otherwise, the transfers take place at the specified value.

Global throttling works with the per-relationship throttle feature for SnapMirror transfers. The per-relationship throttle is enforced until the combined bandwidth of per-relationship transfers exceeds the value of the global throttle, after which the global throttle is enforced. A throttle value 0 implies that global throttling is disabled.

Note: Global throttling should not be enabled on clusters that have SnapMirror Synchronous relationships.

The minimum throttle bandwidth should be 4KBps, and the maximum can be up to 2TBps. A throttle bandwidth of 0 implies that the transfer is not throttled or that bandwidth is unlimited.

A new cluster-wide option to control throttling is as follows:

```
cluster::> options replication*

cluster
  replication.throttle.enable      on          -
  replication.throttle.incoming.max_kbs 4000      -
  replication.throttle.outgoing.max_kbs 2000      -
3 entries were displayed.
```

Each entry can be edited individually. The enable option either enables or disables both the outgoing and incoming throttle.

```
cluster::> options replication.throttle.enable on
1 entry was modified.
```

Changing the outgoing and incoming throttle is reflected in the actual transfer only if the enable option is on. The outgoing and incoming throttle values can be changed irrespective of the enable option value.

```
cluster::> options replication.throttle.outgoing.max_kbs 8000
1 entry was modified.

cluster::> options replication.throttle.incoming.max_kbs 5000
1 entry was modified.
```

How to change TCP receive buffer size

SnapMirror uses the network service `ctlopccp` with a tunable TCP receive buffer window for both intercluster (WAN network) and intra-cluster (LAN network) replication. The TCP receive buffer window is configured per cluster, and an increase in the TCP receive buffer size takes effect immediately with no requirement to reboot (Table 15).

Table 15) TCP receive buffer windows.

	Default	Minimum	Maximum
Intercluster TCP receive buffer window	2MB	256KB	7MB
Intracluster TCP receive buffer window	256KB	256KB	7MB

Note: The intercluster TCP receive buffer window autotunes by default. The window starts at 64KB per TCP stream and grow to a maximum of the configured value (default 2MB) to accommodate intercluster replications. Once grown, the receive buffer window does not shrink if the TCP stream remains open.

Concurrent replication operations

The number of supported simultaneous SnapMirror operations is limited. This limit is per node and varies depending on the platform and version of ONTAP shown in Table 16. For information about the number of concurrent SnapMirror operations allowed per node model, see [NetApp Hardware Universe](#).

Table 16) Maximum number of concurrent SnapMirror transfers per node by ONTAP version and controller model.

ONTAP Version	FAS								FAS Legacy									
	2720	2750	2820	500F	8300	8700	9000	9500	2520	2552	2554	2620	2650	8020	8040	8060	8080EX	8200
9.7.x	100	100			100	100	100		100	100	100	100	100	100	100	100	100	100
9.8.x	100	100		100	100	100	250		100	100	100	100	100	100	100	100	100	100
9.9.x	100	100		100	100	100	250				100	100						100
9.10.x	100	100		100	100	100	250				100	100						100
9.11.x	100	100		100	100	100	250	100				100						100
9.12.x	100	100		100	100	100	250	100				100						100
9.13.x	100	100	100	100	100	100	250	100										100
9.14.x	100	100	100	100	100	100	250	100										100
9.15.x	100	100	100	100	100	100	250	100										100
9.16.x	100	100	100	100	100	100	250	100										100

ONTAP Version	ASA Axxx-Series						Legacy ASA AFF						ONTAP Version	C30		C60			C250			C400			C800		
	A150	A250	A400	A700	A800	A900	A20	A30	A50	A70	A90	A1K		C30	C60	C250	C400	C800									
9.7.x			100	100	100								9.7.x														
9.8.x		100	100	250	250								9.8.x														
9.9.x		100	100	250	250								9.9.x														
9.10.x	100	100	100	250	250	250							9.10.x														
9.11.x	100	100	100	250	250	250							9.11.x				100	100	250								
9.12.x	100	100	100	250	250	250							9.12.x				100	100	250								
9.13.x	100	100	100	250	250	250							9.13.x				100	100	250								
9.14.x	100	100	100	250	250	250							9.14.x				100	100	250								
9.15.x	100	100	100	250	250	250	250	250	250	250	250		9.15.x				100	100	250								
9.16.x	100	100	100	250	250	250	100	100	100	100	250	250	9.16.x	100	100		100	100	250								

ONTAP Version	ASA Axxx-Series					Legacy ASA AFF				
	A150	A250	A400	A800	A900	A220	A250	A400	A700	A800
9.7.x						100			100	
9.8.x						100	100	100	250	250
9.9.x						100	100	100	250	250
9.10.x						100	100	100	250	250
9.11.x						100	100	100	250	250
9.12.x						100	100	100	250	250
9.13.x	100	100	100	250	250	100	100	100	250	250
9.14.x	100	100	100	250	250	100	100	100	250	250
9.15.x	100	100	100	250	250	100	100	100	250	250
9.16.x	100	100	100	250	250	100	100	100	250	250

Best practices

- When planning concurrent operations, it is a best practice to consider the frequency of volume move and volume copy operations in the environment in addition to SnapMirror replications.
- Size the system correctly with enough CPU headroom to allow the CPU workload to execute.

ONTAP provides a greater level of scalability by allowing expansion of a NetApp cluster beyond two nodes. Each node in the cluster provides CPU and memory resources that are used for replication of volumes owned by that node.

Best practice

To optimize replication, distribute replicated volumes across different nodes in the clusters rather than placing all volumes requiring replication on a single node. This best practice allows all nodes in the cluster to share replication activity.

Network sizing requirements

When deploying SnapMirror, consider the round-trip travel time of a packet from the source to the destination storage system, because network distance causes write latency. A network with the appropriate bandwidth available to transfer the system data is required to support the desired replication interval, so that application performance is not affected. There are limitations on the network characteristics that are supported for intercluster replication.

Network sizing requirements for intercluster replication

The intercluster network must be sized appropriately depending on the data change rate and the update interval to meet the RPO of the solution and individual node performance characteristics. Intercluster SnapMirror is supported across networks that have the following characteristics:

- A minimum bandwidth of 0.5Mbps
- A packet loss of 1%

Best practice

All paths used for intercluster replication must have equal performance characteristics. Configuring multipathing in such a way that a node has one intercluster LIF on a slow path and another intercluster LIF on a fast path degrades performance, because data is multiplexed across both paths simultaneously.

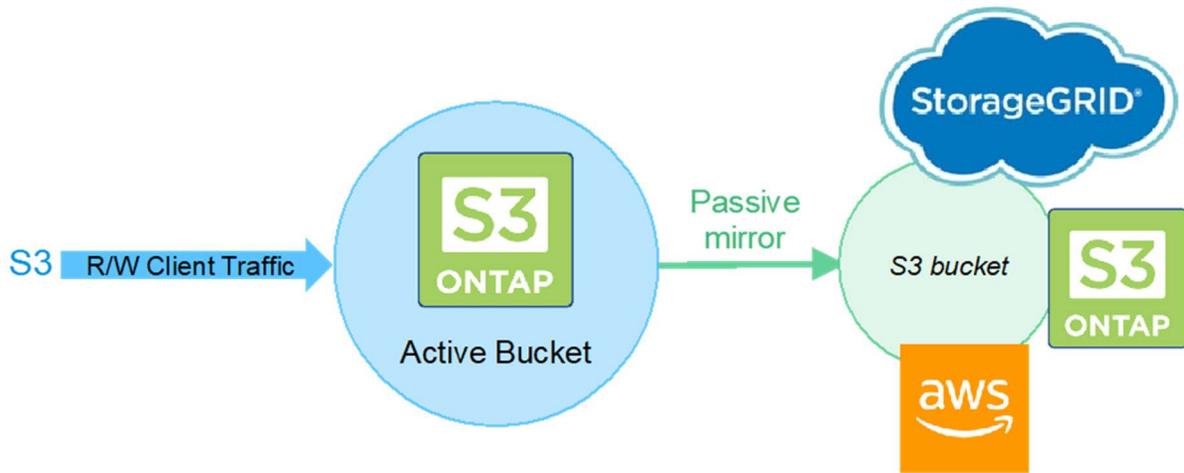
Network sizing requirements for intracluster replication

All intracluster transfers, including SnapMirror, volume move, and volume copy operations, use the private cluster interconnect between nodes in the same cluster whose bandwidth is not configurable.

S3 SnapMirror

ONTAP 9.10.1 introduced ONTAP S3 SnapMirror. S3 SnapMirror provides customers with a native replication and backup solution for their ONTAP S3 object stores. S3 SnapMirror supports a wide variety of S3 targets for data protection and DR including another ONTAP S3 bucket, StorageGRID, ONTAP Cloud Volumes ONTAP S3 buckets and native cloud-provided S3 buckets such as AWS S3 (Figure 42).

Figure 42) ONTAP S3 SnapMirror overview.



S3 SnapMirror uses a specialized replication engine different from the standard LRSE replication engine used by SnapMirror for FlexVolume and FlexGroup asynchronous replication. When setting up S3 SnapMirror relationships, the replication protection policy used is Continuous (`-type continuous`).

For more information on S3 SnapMirror, see [S3 SnapMirror overview](#).

SVM data mobility

SVM data mobility (SVM Migrate) is an ONTAP feature that allows a cluster administrator to move an SVM (including data and SVM configuration information) from one cluster to another. This feature is not dependent on having any previous SnapMirror volume or SVM DR relationship configured for the SVM being moved.

To learn more about SVM Migrate, please visit [SVM data mobility overview - NetApp Documentation](#)

Hardware interoperability

Refer to [NetApp Hardware Universe](#) to understand details about hardware and limits based on each platform. Also, 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. 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 the published specifications. Verify that the source and destination volumes are running [compatible ONTAP versions before creating a SnapMirror data protection relationship](#).

Troubleshooting tips

Troubleshooting cluster peer relationships

1. Run the `cluster peer show` command to verify the availability of the cluster peer relationship. This command displays all existing configured cluster peer relationships.

```
cluster01::> cluster peer show
Peer Cluster Name      Cluster Serial Number Availability
-----
```

2. Add `-instance` to the command to view more detailed information about the cluster peers. Include `-cluster <cluster_name>` to view results for a specific cluster. The `-instance` option displays the remote addresses that are used for intercluster communication.

```
cluster01::> cluster peer show -cluster cluster02 -instance
Peer Cluster Name: cluster02
Remote Intercluster Addresses: 10.12.12.3,10.12.12.4
Availability: Available
Remote Cluster Name: cluster02
Active IP Addresses: 10.12.12.3,10.12.12.4
Cluster Serial Number: 1-80-000013
```

3. Run the `cluster peer ping` command to view information about connectivity between each intercluster address, including RTT response times. For multiple configured cluster peers, use the `-cluster <cluster_name>` option to perform the ping for one specific peer relationship. The `cluster peer ping` command displays the results of a ping between intercluster interfaces. As mentioned earlier, when performing intercluster SnapMirror mirroring over multiple paths between the local and remote clusters, each path must have the same performance characteristics. In this example, the ping response times (RTTs) are comparatively equal to the pings to nodes where the destination cluster displays as `cluster02`.

```
cluster01::> cluster peer ping cluster02

Node: cluster01-01      Destination Cluster: cluster01
Destination Node IP Address      Count TTL RTT(ms) Status
-----
cluster01-01      10.12.12.1      1    255  0.186  interface_reachable
cluster01-02      10.12.12.2      1    255  1.156  interface_reachable

Node: cluster01-01      Destination Cluster: cluster02
Destination Node IP Address      Count TTL RTT(ms) Status
-----
cluster02-01      10.12.12.3      1    255  7.164  interface_reachable
cluster02-02      10.12.12.4      1    255  7.065  interface_reachable

Node: cluster01-02      Destination Cluster: cluster01
Destination Node IP Address      Count TTL RTT(ms) Status
-----
cluster01-01      10.12.12.1      1    255  1.324  interface_reachable
cluster01-02      10.12.12.2      1    255  0.809  interface_reachable

Node: cluster01-02      Destination Cluster: cluster02
Destination Node IP Address      Count TTL RTT(ms) Status
-----
cluster02-01      10.12.12.3      1    255  7.279  interface_reachable
cluster02-02      10.12.12.4      1    255  7.282  interface_reachable
```

Troubleshooting SVM peer relationships

Here is a list of common issues and how to troubleshoot them:

- SVM peer action failure for the intercluster environment:
 - Verify that the peer cluster is reachable.
 - Verify that both clusters support ONTAP versions with SVM peering capability enabled.
 - Verify that the peer SVM name is not associated with another cluster from peer SVM names in the SVM peering table.
 - Check `mgwd.log` and the console logs for error messages.
- SVM peer action failure for the intra-cluster or intercluster environment:
 - Verify that both clusters supported ONTAP versions, with SVM peering capability enabled. Verify that local and peer SVM names are not the same.

- Check `mgwd.log` and the console logs for error messages.
- Run the `vserver peer show` command to verify the SVM peer relationship. This command displays all existing configured SVM peer relationships.

```
cluster02::> vserver peer show
Vserver      Peer      Peer
Vserver      Vserver   State
-----
vs1_dest     vs1_backup peered
vs1_dest     vs1_src   peered
```

- Check for any notifications with the `vserver peer show-all` command.

```
cluster02::> vserver peer show-all
Vserver      Peer      Peer      Peering
Vserver      Vserver   State      Peer Cluster Applications
-----
vs1_dest     vs1_backup peered     cluster03  snapmirror
vs1_dest     vs1_src   peered     cluster01  snapmirror
```

Understanding SnapMirror relationship status

The `Healthy` column indicates the SnapMirror relationship status. This column is shown in the output of the `snapmirror show` command on the CLI and as the `Healthy` column in the displayed status of SnapMirror relationships in System Manager.

```
cluster02::> snapmirror show
Source      Destination  Mirror  Relationship  Total      Progress
Path        Type  Path      State  Status  Progress  Healthy  Last Updated
-----
vs1_src:vol1
           XDP  vs1_dest:vol1
                        Snapmirrored
                        Transferring  128KB  true  02/25 15:43:53
```

The `Mirror State` column also displays if the destination volume is offline or if it cannot be reached.

Troubleshooting SnapMirror relationships

To determine when the last SnapMirror transfer for a specific relationship completed, see the `Exported Snapshot Timestamp` field for instance information.

```
cluster02::> snapmirror show -instance
Source Path: snap_src1:SMSsource
Destination Path: svm_dst1:SMSsource_dest
Relationship Type: XDP
Relationship Group Type: none
SnapMirror Schedule: -
SnapMirror Policy Type: vault
SnapMirror Policy: XDPDefault
Tries Limit: -
Throttle (KB/sec): unlimited
Mirror State: Snapmirrored
Relationship Status: Idle
File Restore File Count: -
File Restore File List: -
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: snapmirror.12ceb7f0-b078-11e8-baec-0050
56b013db_2160175147.2020-01-24_043858
```

```

Newest Snapshot Timestamp: 01/24 04:38:59
Exported Snapshot: snapmirror.12ceb7f0-b078-11e8-baec-0050
56b013db_2160175147.2020-01-24_043858
Exported Snapshot Timestamp: 01/24 04:38:59
    Healthy: true
    Unhealthy Reason: -
    Constituent Relationship: false
    Destination Volume Node: cluster2-01
    Relationship ID: 1a46a611-3e64-11ea-86bf-005056b013db
    Current Operation ID: -
    Transfer Type: -
    Transfer Error: -
    Current Throttle: -
    Current Transfer Priority: -
    Last Transfer Type: resync
    Last Transfer Error: -
    Last Transfer Size: 0B
Last Transfer Network Compression Ratio: 1:1
    Last Transfer Duration: 0:0:1
    Last Transfer From: snap_src1:SMSSource
    Last Transfer End Timestamp: 01/24 04:45:16
    Progress Last Updated: -
    Relationship Capability: 8.2 and above
    Lag Time: 5:27:1
Identity Preserve Vserver DR: -
    Volume MSIDs Preserved: -
    Is Auto Expand Enabled: -
Number of Successful Updates: 0
Number of Failed Updates: 0
Number of Successful Resyncs: 1
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
    Total Transfer Bytes: 0
Total Transfer Time in Seconds: 1

```

For SnapMirror relationship issues, review information about relationships in the event log. Use the `-messagemname` option with the `event log show` command to filter the event log for messages related to SnapMirror, as shown in the following example. Specify the `mgmt.snapmir*` message name to filter the output and find only messages related to SnapMirror.

```

cluster01::> event log show -messagemname mgmt.snapmir*
Time          Node          Severity      Event
-----
12/6/2011 17:35 cluster02-01 ERROR         mgmt.snapmir.update.fail: Update from source
volume 'cluster01://vs1/vol03' to destination volume(s) 'cluster02://vs2/vol03' failed with error
'Failed to setup transfer. (Duplicate transfer specified. (Other error.))'. Job ID 1322.
12/6/2011 17:34:35 cluster02-01 DEBUG         mgmt.snapmir.abnormal.abort: Source Path
cluster01://vs1/vol01, Destination Path cluster02://vs2/vol01, Error Transfer failed.
(Destination volume cluster02://vs2/vol01 is smaller than the source volume.), Function
copySnapshot, line 5030, job ID 1355.
12/5/2011 05:15:45 cluster02-01 DEBUG         mgmt.snapmir.abnormal.abort: Source Path
cluster01://vs2/vol12, Destination Path cluster02://vs8/vol12, Error Failed to delete Snapshot
copy weekly.2011-12-04_0015 on volume cluster02://vs8/vol12. (Snapshot is in use.), Function
deleteSnapshot, line 4285, job ID 1215.

```

To find an error message about a specific volume, filter the message list further by specifying the name of the volume, enclosed in asterisks, with the `-event` option, as shown in the following example.

```

cluster01::> event log show -messagemname mgmt.snapmir* -event *vol01*
Time          Node          Severity      Event
-----
12/6/2011 17:34:35 cluster02-01 DEBUG         mgmt.snapmir.abnormal.abort: Source Path
cluster01://vs1/vol01, Destination Path cluster02://vs2/vol01, Error Transfer failed.
(Destination volume cluster02://vs2/vol01 is smaller than the source volume.), Function
copySnapshot, line 5030, job ID 1355.

```

All SnapMirror events are logged to the `SnapMirror_audit.log` and `SnapMirror_error.log` files on the node where the destination volume resides. This node might be different from the one where the command was issued. The node running the operation can be determined by running the `snapmirror show -fields destination-volume-node` command. System Manager allows viewing of the SnapMirror log files.

Best practices for DR configurations

Best practices

- **Replicate volumes hosted by a single SVM on the source to a single SVM on the destination.**

An SVM hosts the root of a NAS namespace for NAS clients and a single storage target in SAN environments. If some NAS volumes are replicated from one SVM into different SVMs at the destination, then those volumes cannot be recovered into the same namespace. The same is true of volumes containing LUNs. If the volumes are replicated into different SVMs at the destination, then all the LUNs are not presented under the same SAN target.

- **Ensure that the source and destination SVMs are hosted within the same Active Directory, LDAP, or NIS domain.**

This configuration is required so that access control lists (ACLs) stored in NAS files are not broken if a NAS volume is recovered into an SVM that cannot authenticate those ACLs. The process of changing file-level ACLs to reenforce access from a different domain can be extremely difficult and time consuming. It is also important for authentication of tools running in SAN clients such as NetApp SnapCenter Plug-in for Windows.

- **Consider using DNS aliases to associate source and destination SVMs for volume-based DR use cases.**

Volume-based SnapMirror replication requires that source and destination SVMs use different names to support the ability to access both SVMs as needed for secondary activities such as backup to tape or cloud from the destination SVM. In a DR failover scenario, clients must access the volumes from the secondary site. Although changing the DNS records for the destination SVM to reflect the source SVM's name is an option, creating DNS alias records is recommended. This practice makes sure that SMB shares are still accessible using the same Uniform Naming Convention (UNC) path name or NFS exports they did prior to the DR event.

- **Consider using SnapMirror SVM DR when an entire SMB or NFS namespace needs to be protected.**

SVM DR using the `-identity-preserve true` parameter enables client access to the namespace by using the same client configuration in the event of a DR failover scenario. This also ensures that the namespace is identical because all replicated volumes within the SVM will have the same names.

- **When using volume replication, consider using the same volume names in both the source and destination SVMs.**

Although using destination volume names that are the same as the source volume names is not required, this practice can make mounting destination volumes into the destination simpler to manage if the junction path where the volume is mounted also has the same name as the volume.

- **Many SAN clients cannot access a LUN that resides in a read-only container, such as a SnapMirror destination volume.**

Generally, LUNs should be mapped to igroups and mounted by SAN clients after the SnapMirror break operation is performed.

- **Configure the destination SVMs ahead of time as described in the following section.**

This approach can greatly speed up the storage system DR process, possibly reducing it to a few SnapMirror break operations and the update of some DNS aliases.

Best practices

- **If not using SVM DR, as new volumes are created at the source site, SnapMirror volume relationships must be created to replicate those volumes.**

Make configuration settings pertaining to those volumes in the DR site after the volumes are created and replicated so they can be ready in the event of a disaster.

Configuration and failover for DR

Configuration and failover for DR is presented here in an overview of the DR process for intracluster SnapMirror data protection (async-mirror) replication. The process is presented in two sections. The first section provides steps that must be completed before a failover is required to prepare the destination for failover. These steps should be completed to prepare the DR site for a DR scenario. The second section provides the steps necessary to perform a failover.

Every environment has its own unique characteristics. Each environment can affect a DR plan. Depending on the type of DR solutions deployed, each organization's DR situation can be very different. To enable success, proper planning, documentation, and a realistic walkthrough of a DR scenario are required.

Environment failover requirements and assumptions

To provide a successful DR experience, consider some general requirements and assumptions. The following is not an all-inclusive list:

- System administrator access to a workstation or server desktop session from which to administer the DR site and perform the failover.
- System administrators have all appropriate credentials, accounts, passwords, and so on required to access the systems.
- Connectivity to the DR network is available from wherever operations are performed.
- Certain infrastructure servers already exist in the DR site and are accessible. These systems provide basic services necessary for the administrators to work in the environment and execute the recovery plan.
 - DR site Active Directory or LDAP services to provide authentication.
 - DR site DNS services to provide name resolution.
 - DR site license servers to provide licensing services for all applications that require them.

Note: A server must be available at the DR site to perform the necessary Active Directory FS MO roles. For information regarding transferring roles to a surviving Active Directory server or seizing these roles from a failed server, see [Microsoft KB 255504](#).

- The DR site has time synchronized to the same source as the primary site or a source in sync with the primary site.
- All required NetApp volumes are replicated using SnapMirror to the DR site.
- SnapMirror operations have been monitored and are up to date with respect to the designed RPO.
- The required capacity exists on the DR NetApp controller. This refers to the capacity required to support day-to-day operations that have been planned for in the DR environment.
- All DR site application servers have the proper connectivity configured to be able to connect to the DR storage arrays.
- A method exists to isolate or fence the failed primary network from the DR site. This approach is necessary if the event causing the disaster is temporary or intermittent in nature, such as an extended power outage. When the primary site systems restart, services might conflict with the recovered operations that are then running at the DR site.

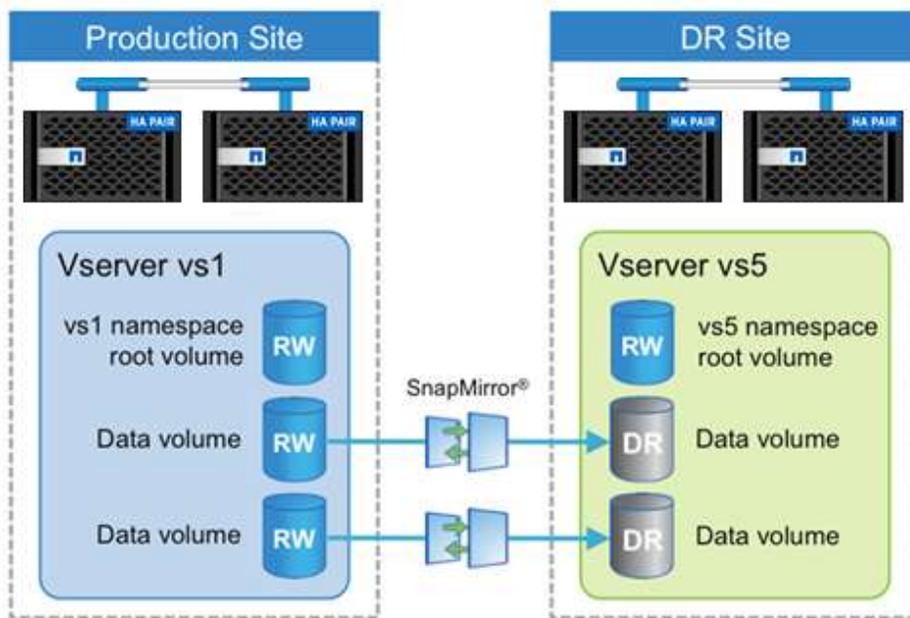
- Plans have been made for providing users and applications access to the data and services at the DR site. For example, updating DNS records such that home directory mount requests to the primary site SVM are directed to the DR site SVM instead.

Preparing the destination for failover

Many parts of the DR process can be prepared for prior to a DR event. For example, mounting volumes into the namespace, creating SMB shares, and assigning NFS export policies, can all be performed ahead of time. SnapMirror volume replication cannot be used to replicate configuration information that could be independent in the destination SVMs. These configurations include SVM domain membership, SMB configuration, NFS policies, Snapshot policy schedules, or NetApp storage efficiency policies.

Figure 43 illustrates volume layout for DR.

Figure 43) Volume layout for DR.



After volumes have been replicated, complete the following steps to prepare the destination system for failover.

NAS and SAN environments

1. Configure the destination SVM membership into the appropriate Active Directory, LDAP, or NIS domain.
2. Determine that the destination SVM is a member of the same domain as the source SVM so that authentication is not broken for tools such as NetApp SnapCenter. This configuration also ensures that the same users can be authenticated against file-level ACLs that are replicated by SnapMirror.
3. Create any nondefault Snapshot copy policies needed in the destination cluster.

Note: NetApp recommends configuring Snapshot copy policies in the destination cluster with the same schedules as those in the source. Snapshot copy policies must be applied after failover (`snapmirror break`) has completed.

4. Create NetApp storage efficiency policies in the destination SVM.

Note: If NetApp storage efficiency policies are assigned to the volumes in the source SVM, an identical policy must be created in the destination SVM in order to schedule the dedupe

process after failover at the DR site. These storage efficiency policies must be applied after failover (`snapmirror break`) has completed.

NAS only environments

1. Verify that all necessary volumes in the source SVM are being replicated to the destination SVM. Volumes can be mounted in subfolders or inside other volumes in the namespace. If this condition exists, it is important to make sure that all the volumes required to properly reconstruct the namespace at the destination are being replicated.
2. Verify the security style and permissions on the destination SVM root volume. The security style and permissions of the root of the destination SVM namespace must be set correctly, or the NAS namespace might be inaccessible after failover.
3. Mount the destination NAS volumes into the destination SVM namespace.
SnapMirror does not replicate the SVM namespace junction path information. NAS volumes have no junction path, so they are not accessible after a SnapMirror break occurs unless they are premounted before failover or until they are mounted after failover.
When mounting the volumes, mount them into the namespace by using the same junction path into which the source volume was mounted in the source SVM. This configuration is important so that paths in the recovered namespace are not different than paths that existed at the primary site. If the paths are different, then client mount points, links, shortcuts, and aliases might not be able to find the correct paths.
Note: Volumes cannot be mounted inside (nested in) other volumes that are still in a DP state. After using the `snapmirror break` command, any volume that has a mount point nested inside a replicated volume must be mounted, and any CIFS shares must be created.
4. Create CIFS shares on the destination SVM by using the same share names that were used at the source. Clients can access the CIFS shares. However, all data is read-only until the volume is failed over.
5. Apply the proper ACLs to the CIFS shares at the destination.
6. Create appropriate NFS export policies for the destination SVM.
7. Assign the NFS export policies to the destination volumes. Although clients can access the NFS exports, all data is read-only until the volume is failed over.

SAN only environments

1. If the destination SVMs use portsets, they can be configured as required before failover.
2. Configure igroups on the destination SVM.

Typically, there are different application servers that connect to the recovered storage at the DR site. The initiators from these servers can be preconfigured into appropriate igroups in the destination SVM.

Note: Because many host operating systems do not support connecting to LUNs in read-only containers, map LUNs to igroups after failover (`snapmirror break`) has completed.

Performing a failover

With most of the configuration necessary for DR performed prior to a failover, the actual steps required to fail over during a DR scenario are greatly reduced. They are as follows.

NAS environment

1. Perform a SnapMirror break operation to fail over each volume. In ONTAP, wildcards can be used to perform a SnapMirror operation on multiple volumes with one command. The following example performs failover for all volumes in the destination SVM called `vs5`. It can be restricted to certain volumes by using part of the volume name in the command.

```
cluster02::> snapmirror break -destination-path cluster02://vs5/*
```

2. If the volumes have been mounted in the namespace and CIFS shares and NFS export policies have been created and applied, clients then have read-write access to the NAS data.
3. Redirect clients to the recovered storage.

It is a common practice to have a DR system with a different name than the source system. In DR failover scenarios, it is typical to change DNS name resolution or use DNS aliases to redirect clients to the name of the recovered storage systems. This approach enables CIFS share access using the same UNC path name, and NFS clients can also access the expected path. Alternatively, the failed source storage system can be removed from Active Directory. The recovery storage system can then be removed and added again to Active Directory by using the same name as the source system. However, it can take time for this change to propagate through a large Active Directory environment.

SAN environment

1. Perform a SnapMirror break operation to fail over each volume. Wildcards can be used to perform a SnapMirror operation on multiple volumes with one command. The following example performs failover for all volumes in the destination SVM called `vs5`. It can be restricted to certain volumes by using part of the volume name in the command.

```
cluster02::> snapmirror break -destination-path cluster02://vs5/*
```

2. Make the LUNs in the volume available to the SAN clients at the DR site by mapping the LUN into the appropriate igroup.
3. On the SAN client, perform a storage rescan to detect the connected LUN.

Post failover volume configuration

Snapshot copy policies and NetApp storage efficiency policies cannot be assigned to volumes in a DP state, so they must be assigned after failover.

1. If using an ONTAP snapshot schedule, assign a Snapshot copy policy to the recovered volumes. In SAN environments, snapshots are typically scheduled in the client.
2. If using NetApp storage efficiency technology, assign a storage efficiency policy to the recovered volumes.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- ONTAP 9 Data protection and disaster recover documentation
<https://docs.netapp.com/us-en/ontap/data-protection-disaster-recovery/index.html>
- NetApp support website
<https://mysupport.netapp.com/site/global/dashboard>
- ONTAP and ONTAP System Manager documentation resources
<https://www.netapp.com/data-management/oncommand-system-documentation/>
- ONTAP command reference
<https://docs.netapp.com/us-en/ontap-cli-9161/>

Version history

Version	Date	ONTAP version
Version 11	January 2026	ONTAP 9.18.1 updates
Version 10	March 2025	ONTAP 9.16.1 updates
Version 9.1	February 2024	Errors and omissions
Version 9.0	January 2024	ONTAP 9.14.1 updates
Version 8.0	June 2023	ONTAP 9.13.1 updates
Version 7.0	February 2023	ONTAP 9.12.1 updates
Version 6.0	October 2022	ONTAP 9.11.1 updates
Version 5.2	August 2021	ONTAP 9.10.1
Version 5.1	December 2020	ONTAP 9.8
Version 5.0	March 2020	ONTAP 9.7
Version 4.1	November 2019	ONTAP 9.6
Version 4.0	November 2017	ONTAP 9.3
Version 3.1	February 2016	ONTAP 9.0
Version 3.0	April 2015	ONTAP 8.3
Version 2.2	November 2013	ONTAP 8.2 updates
Version 2.1	July 2013	ONTAP 8.2 updates
Version 2.0	April 2013	ONTAP 8.2
Version 1.0	February 2012	ONTAP 8.1

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 © 2025 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, 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) (FEB 2014).

Trademark information

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

