



Technical Report

# **SnapMirror synchronous configuration and best practices**

Tony Ansley, NetApp  
February 2024 | TR-4733

## Abstract

This document contains information and best practices for configuring NetApp® SnapMirror® synchronous (SM-S) replication in NetApp ONTAP® including features and updates introduced in ONTAP 9.14.1.

## TABLE OF CONTENTS

<b>Executive summary</b> .....	<b>4</b>
Purpose and intended audience .....	4
<b>SM-S overview</b> .....	<b>5</b>
Terminology .....	5
<b>What's new for SM-S in ONTAP 9?</b> .....	<b>6</b>
<b>Use cases for SM-S</b> .....	<b>6</b>
Single or multiple volume file protection (Network Attached Storage) using two data centers – short distance .....	7
Database application protection across two data centers – long distance .....	7
Database application protection across three data centers .....	<b>Error! Bookmark not defined.</b>
<b>SM-S concepts</b> .....	<b>8</b>
Licensing .....	8
Relationships .....	8
Common snapshot .....	8
Sync versus StrictSync policies .....	10
Operational states of SM-S .....	11
SM-S behavior when state is OutOfSync .....	11
<b>Implementing SM-S</b> .....	<b>12</b>
Prerequisites .....	12
Scalability .....	13
Limitations .....	13
Creating SM-S volume relationships .....	14
Monitoring and managing SM-S .....	14
<b>Advanced SM-S topics</b> .....	<b>15</b>
Restoring an SM-S source volume from an NDMP backup .....	15
SM-S support for increased FlexVol volume sizes .....	16
SM-S support for non-disruptive operations .....	17
Support for NFS 4.2 sparse files and extended attributes .....	17
Accessing the secondary volume .....	18
Fan-out support .....	19
Performing accessibility testing of DR site data .....	19
Performing DR rehearsals .....	20
Converting to or from SM-S .....	21
Conversions to or from an SM-S StrictSync policy .....	22

Migration of 7-Mode synchronous SnapMirror to SM-S.....	23
SM-S replication interoperability .....	23
<b>Performance considerations .....</b>	<b>24</b>
Quality of service .....	24
Mixing array models .....	25
<b>Reversion to previous ONTAP versions.....</b>	<b>25</b>
<b>Where to find additional information .....</b>	<b>25</b>
<b>Version history .....</b>	<b>26</b>

## LIST OF TABLES

Table 1) Required time to create a common Snapshot per ONTAP version. ....	9
Table 2) Relationship status .....	11
Table 3) Concurrent replication operations allowed per high-availability (HA) pair.....	13
Table 4) SM-S relationship interoperability .....	24
Table 5) QoS support for SM-S. ....	25

## LIST OF FIGURES

Figure 1) DR between two local data centers .....	5
Figure 2) SM-S data protection between two short distance datacenters .....	7
Figure 3) Combining synchronous and Asynchronous SnapMirror relationships. ....	<b>Error! Bookmark not defined.</b>
Figure 4) Three data center enterprise-application use case. ....	<b>Error! Bookmark not defined.</b>
Figure 5) Monitoring SM-S using System Manager.....	15
Figure 6) Monitoring SM-S using Active IQ Unified Manager. ....	15
Figure 7) Fan-out support with SM-S. ....	19
Figure 8) Changing SnapMirror relationship between asynchronous and synchronous modes .....	22

## Executive summary

Businesses can use several approaches to increase data availability in the face of hardware, software, or site failures. Data protection (DP) is one of the most critical aspects because any loss of data translates directly into lost money and time. Data protection is the process of taking data located in one location and making a copy of it in a different location to serve two use cases:

- **Backup.** The objective is to restore from the secondary to the primary with no intention of failing over to the secondary. This implies that the primary purpose of the secondary is archival storage. Therefore, you might have more data in the secondary than in the primary.
- **Disaster recovery (DR).** An exact replica or copy is maintained in the secondary and used for failover from the primary to the secondary if there is a failure at the primary site.

Although backups allow you to recover lost data from an archival medium (tape, disk, or the cloud), mirroring is the most popular data availability mechanism for business continuity and DR because it supports minimizing data loss and downtime. NetApp SnapMirror® technology offers a fast and flexible enterprise solution for mirroring or replicating data over LANs and WANs.

SnapMirror synchronous (SM-S) replication enhances SnapMirror asynchronous data protection to the next level by replicating in real-time every write or modify request from the application to two different volumes simultaneously. The main advantages of using SM-S include:

- **Robust enterprise technology.** SM-S is a mature feature of ONTAP storage systems that has been enhanced and improved over time. SM-S can recover from update failures, use concurrent processes for replication processing, and much more.
- **Speed and efficiency.** Logical block-level, real-time data transfer makes sure that only the data that has changed is sent to the destination replica.
- **Testability.** SM-S supports cloning of destination volumes, using the latest common snapshot, as writable volumes using NetApp FlexClone® technology, irrespective of their size, and in a space-efficient manner, without needing to stop data replication from the source. These FlexClone volumes can be used for a wide variety of secondary data use cases.
- **Failover and failback.** If DR systems must be brought online, SM-S relationships can be broken, which stops replication from the primary to secondary volumes and makes the destination volumes writeable and ready for application restart using the secondary volume. SM-S allows you to resynchronize the source with the changes made at the destination and then re-establish the original SM-S relationship.
- **Ease of use.** Using either ONTAP System Manager, REST APIs, or command-line interfaces (CLIs), you can perform data protection operations including monitoring and managing all SM-S replication relationships in one place.
- **Secure.** SnapMirror relationships can be encrypted natively end-to-end using open standard TLS.
- **Zero RPO.** SM-S technology ensures that any data change on a protected volume is not lost in the event of a disaster by replicating every data change operation to the secondary volume before acknowledging the successful write operation to the client operating system and application.

## Purpose and intended audience

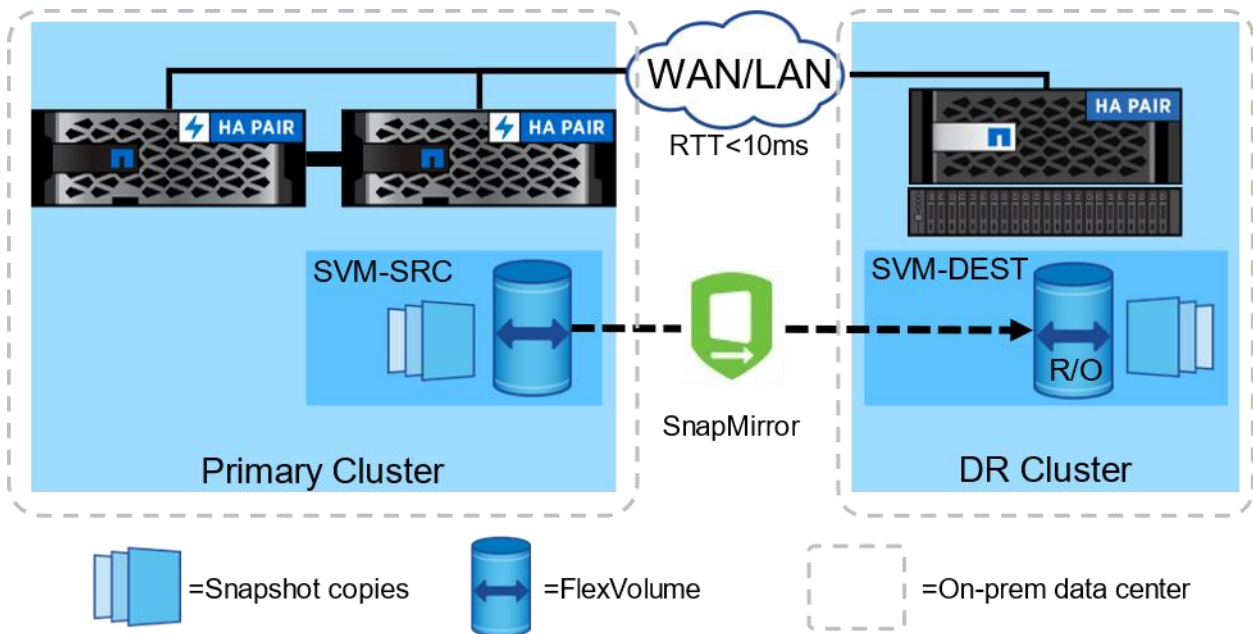
This document is intended for people who administer, install, or support ONTAP storage systems and who expect to configure and use SM-S technology for data replication. It assumes that the reader understands the following processes and technologies:

- A working knowledge of ONTAP operational processes
- A working knowledge of NetApp features such as Snapshot copy technology, NetApp FlexVol® volumes, and NetApp FlexClone volumes
- General knowledge of data protection, DR, and data replication solutions and methodologies

## SM-S overview

SM-S provides volume-granular synchronous data replication that enterprises depend on for backup and DR. SM-S enables customers to achieve zero RPO by replicating data from FlexVol volumes on one ONTAP storage system to a second ONTAP storage system situated within a data center or between two separate data centers. The distance between these data centers is limited by the SM-S requirement for network RTT of less than 10 milliseconds (Figure 1). SM-S provides the flexibility to protect any number of FlexVol volumes within a cluster to other, remote ONTAP storage systems. This functionality addresses regulatory or industry-mandated requirements for real-time data protection in industries such as finance, healthcare, or any other industries that depend on zero data loss.

Figure 1) DR between two local data centers.



### Terminology

- **Primary or source.** The originating volume from which SM-S retrieves data.
- **Secondary or destination.** The targeted volume to which SM-S writes.
- **Recovery point objective (RPO).** The amount of data loss your business application can tolerate.
- **Recovery time objective (RTO).** The amount of time required to restart a failed application.
- **SnapMirror unified replication.** A DR technology designed for failover from primary storage to secondary storage at a geographically remote site by creating a replica or mirror of your working data in secondary storage. You can continue to serve data from the secondary storage in the event of a catastrophe at the primary site. Generally, SnapMirror replication technology can operate in several different modes and data repository types. This term refers to the asynchronous mode of SnapMirror replication for FlexVol volumes and storage virtual machines (SVMs)
- **SnapMirror synchronous (SM-S) replication.** SnapMirror replication mode enables the synchronous mirroring of FlexVol volumes by mirroring all data modifications and write operations to a secondary volume. This ensures that the volume data is not lost (zero RPO) in the event of a disaster event.
- **Common snapshot.** SM-S uses ONTAP Snapshot copy technology to optimize performance, reduce RPO and RTO, and support read-only access to the secondary volume data. At regular intervals (default is 6 hours), SM-S will create a snapshot at the primary and secondary sites of each protected volume. These snapshots are referred to as the common snapshot for that volume.

- **Active file system (AFS).** The AFS is the real-time state of the volume on the source or destination as data is written to the volume by SM-S. This may differ from the data presented in a snapshot, including the common snapshot.
- **SnapMirror sync (Sync) mode.** A mode of SM-S that provides zero RPO replication while tolerating temporary replication interruptions or failures.
- **SnapMirror strict synchronous (StrictSync) mode.** A mode of SM-S that provides zero RPO replication but stops primary application I/O if a replication failure occurs. This ensures that both primary and secondary volumes are always identical.
- **InSync.** An SM-S state where it is actively replicating each application write or update I/O to the secondary storage system.
- **OutOfSync.** An SM-S state where the application I/O is not replicating to the secondary storage system. This state will generate an application I/O failure response for volumes protected by SM-S strict sync mode.
- **Round trip time (RTT).** The duration in milliseconds (ms) it takes for a network request (such as a network ping) to go from a source network port to a destination network port and back again to the source.
- **Logical Interface (LIF).** A LIF is a virtual network port used for various types of communication within a NetApp ONTAP array. SnapMirror communicates with client applications using a data LIF and replicates data between the source and destination volumes using an intercluster LIF.
- **Storage virtual machine (SVM).** An SVM is a logical storage server that provides data access to LUNs or network-attached storage (NAS) namespace using one or more data LIFs. Each SVM enforces data visibility and security critical for multi-tenant environments.

## What's new for SM-S in ONTAP 9?

In an ongoing effort to provide better ONTAP features and capabilities, the following items are new for SM-S in ONTAP 9.14.1:

- Replication of all application-generated Snapshot copies regardless of SnapMirror label value.
- Improved cascade behavior for SnapMirror asynchronous replication of SM-S DP volumes
- Support for new storage efficiency enabled volumes reporting capacity up to 600TB
- Immediate availability of application created snapshots for replication in cascade or fan-out cascade relationships.

A full list of changes since ONTAP 9 was first released can be found in the [ONTAP 9 release notes](#).

## Use cases for SM-S

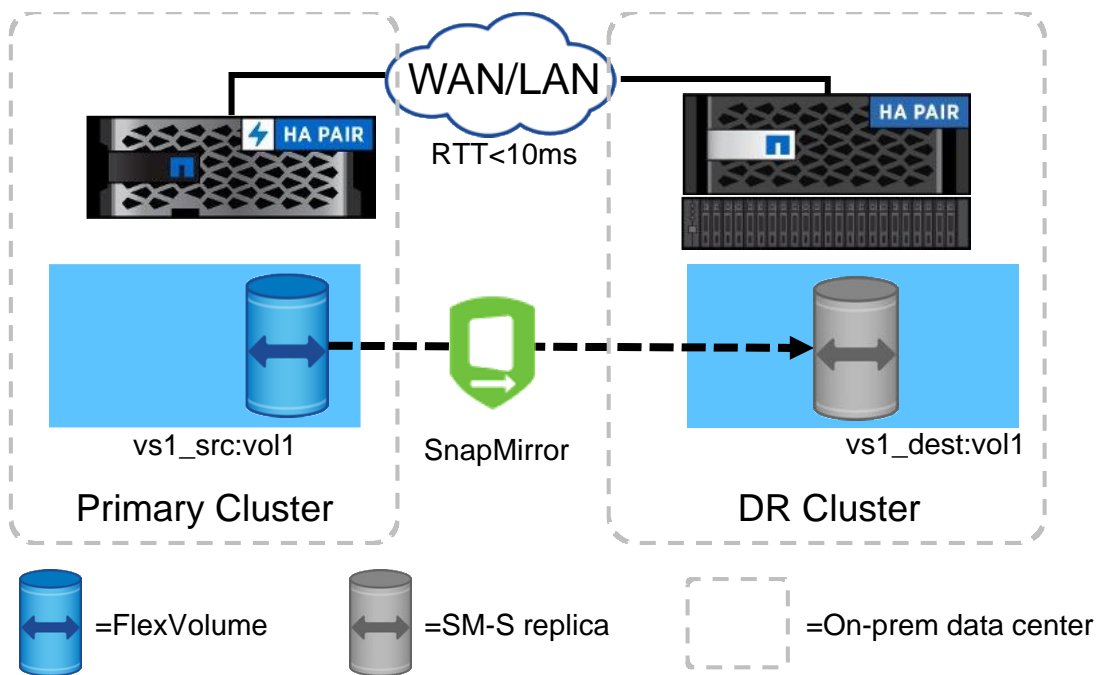
The primary use case for SM-S is DR. In today's digital world, a loss of vital business data can cripple a company, and any downtime can have serious repercussions for the organization's financial performance and competitive advantage. SM-S replication software can safeguard data from loss due to natural disasters, fire, application failure, user error, or software malfunction.

DR solutions are very dependent on the requirements of applications using protected data. Applications can be as simple as a general-purpose file system for storing users' working files, or as complex as transactional or data analytics applications that require data to be managed with one or more databases. The following three examples are offered to highlight the capabilities and flexibility of SM-S as a stand-alone data protection solution or in conjunction with other ONTAP features such as SnapMirror asynchronous.

## Single or multiple volume file protection (Network Attached Storage) using two data centers – short distance

This is the most basic data protection use case for SM-S. SM-S provides a zero-RPO solution for mission-critical file protection for enterprises using NFS or SMB file system sharing such as those under regulatory restrictions. As shown in Figure 2, one or more FlexVol volumes (vs1\_src:vol1) can have an SM-S relationship to mirror all data and metadata changes from a primary site to a DR volume (vs1\_dest:vol1) at a secondary site in real-time. This solution is typically used with network-attached storage (NAS) environments that use NFS or SMB protocols that provide file shares to end-users or other NAS-compatible applications such as VMware ESX. The distance between the two data centers is considered a “short distance” concerning the round-trip time (RTT) latency. RTT latency must be less than 10ms for SM-S.

Figure 2) SM-S data protection between two short distance datacenters



## Application protection across multiple data centers

For applications that depend on coordinated, consistent writes and updates to their data, SM-S can be an important component – along with SnapMirror asynchronous – of a comprehensive data protection plan. In many cases, these applications depend on transaction logs for data consistency. Any replication of these transaction logs must ensure that data replicated to secondary locations is identical to the primary source transaction log and if there is any situation that prevents these copies from being identical, the application cannot continue. Other data stores may not require real-time protection as they can easily be recovered using the transaction logs if there is a failover event.

This different treatment for replication based on the data role within the application suggests a variety of application data architectures that can be leveraged for protection of the application’s data. These architectures will typically involve not only different replication policies (synchronous or asynchronous), but also employ more than one data repository in multiple disparate sites. To learn more details about these various use cases, please refer to [TR-4832: SnapMirror data protection using multiple data centers](#).

# SM-S concepts

The following sections will provide a deeper discussion of basic SM-S concepts that will need to be understood as planning and implementation proceed.

## Licensing

An SM-S license is required on all nodes on the source and destination clusters. How this license is obtained will depend on the ONTAP version:

- Starting with NetApp ONTAP 9.6 and later, an SM-S license is included with the Data Protection or Premium bundle.
- Starting in ONTAP 9.9.1, SM-S is automatically enabled on systems where the Data Protection or Premium bundle has been installed.
- Systems purchased with ONTAP 9.6 or later are shipped with the license installed.
- Systems purchased with ONTAP 9.5 or earlier, or before June 2019, require the [master license key](#) to be installed regardless of the current version.
- Starting on June 2023, 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 synchronous is included as part of the ONTAP One license.

**Note:** ONTAP One is also available for deployed clusters running older ONTAP version:

- For customers who already have the Data Protection bundle installed, an upgrade to ONTAP One is available at no additional cost.
- For customers who have not purchased the Data Protection bundle, ONTAP One is available for an additional fee as a substitute for purchasing the Data Protection bundle.

## Relationships

SM-S requires that a set of relationships be created before volumes can be protected. These relationships provide information to each participating cluster about accessibility, security, and health monitoring. There are two relationship types that are required:

- **Cluster peer relationship.** Creating a cluster peer relationship is a one-time operation between any two ONTAP clusters. By peering cluster, each cluster has critical information about the remote cluster to enable replication. Read [Create a cluster peer relationship \(netapp.com\)](#) and the [cluster peer create command reference](#) for details.
- **SVM peer relationship.** Before a volume can be protected using SM-S, the SVM that the volume is hosted by must also have a peer relationship with an SVM on the remote cluster that will host the replica volume. This is done once for every SVM relationship regardless of the number of protected volumes within the SVM. Read [Create an intercluster SVM peer relationship \(netapp.com\)](#) and the [vserver peer create command reference](#) for details.

### Best practice

Name an SVM with a unique fully qualified domain name (FQDN): for example, `dataVserver.HQ` or `mirrorVserver.Offsite`. SVM peering requires unique SVM names, and using the FQDN naming style makes it much easier to establish uniqueness.

## Common snapshot

SnapMirror synchronous periodically creates a point-in-time snapshot of the active file system (AFS) on both the source and destination volumes. These snapshots are used as a common base to support faster resynchronization after any temporary out-of-sync state.



The process ONTAP uses to create a common snapshot includes the following steps:

- Pause application I/O.
- Flush all write queues to the source and destination volumes.
- Create a snapshot of the quiesced volumes on both primary and secondary volumes.
- Verify that the common snapshot was successfully created on both primary and secondary volumes.
- Restart application I/O.

The advantage of this process is:

- To enable quick and easy resynchronization of volumes after any out-of-sync state.
- Allow for secondary uses of the data from the DR location for applications that do not need access to the latest data such as backup, software development, or data analytics.
- Efficient use of network resources by not requiring physical replication of snapshots from source volume to destination volume.

By default, a common snapshot is created every six hours. The interval between creation of each common snapshot can be modified by creating a custom schedule with a minimum interval of 30 minute and maximum interval of 24 hours. This custom schedule is then used in a custom SM-S policy or applied to an existing SM-S policy.

```
Cluster::*> cron create -name 3hourly_schedule -hour 01,04,07,10,13,16,19,22 -minute 03
(job schedule cron create)

Cluster::*> snapmirror policy modify -policy Sync -common-snapshot-schedule 3hourly_schedule -
vserver vs1
```

The time required to create a common snapshot for each volume will depend on several factors such as outstanding I/O transactions, array model, disk technology, etc. Table 1 illustrates improvements made based on lab tests in common snapshot creation based on ONTAP version. Note that duration is provided as a range of possible times to take the various factors into account.

**Table 1) Improvements in the time to create a common snapshot per ONTAP version.**

ONTAP version	Required time to create a common snapshot
ONTAP 9.9.1 and earlier	250ms-5s (where most operations take 2-4 seconds)
ONTAP 9.10.1	10ms-5s (where most operations finish within 1 second)
ONTAP 9.11.1	5ms-512ms (where most operations finish within 10ms -256ms)

NetApp recommends evaluating the effect of any change to the default schedule, taking these factors into consideration:

- A maximum of two common snapshots are retained on each protected volume on the source and destination.
- The `snapmirror update` command provides a manual method of creating a common snapshot by capturing the latest state of AFS and creating an on-demand common snapshot.

**Note:** The SnapMirror policies support only one rule for common snapshot creation. Common snapshots will have a snapshot label of `sm_created`.

### Best practice

NetApp recommends not making changes by adding rules to default policies, but create custom policies when your needs deviate from the pre-defined settings. Making changes to default, built-in policies have a global effect, which might not be desirable.

## Sync versus StrictSync policies

SM-S operates in one of two modes based on the SM-S policy selected during relationship creation. These SM-S policies and their operational differences are listed below.

- **Sync**

In synchronous mode, application I/O operations are sent in parallel to the primary and secondary storage systems. If the write to the secondary storage system does not complete for any reason, the application is allowed to continue writing to the primary storage system. When the error condition is corrected, SM-S technology automatically resynchronizes with the secondary storage system and then resumes replicating synchronously from the primary storage to the secondary storage system. This process provides an RPO of zero while the volumes are in the InSync state.

```
Primary::snapmirror policy*> show -vserver vs1 -policy Sync
Vserver: cluster3
SnapMirror Policy Name: Sync
SnapMirror Policy Type: sync-mirror
Policy Owner: cluster-admin
Tries Limit: 8
Transfer Priority: normal
Ignore accesstime Enabled: false
Transfer Restartability: always
Network Compression Enabled: false
Common Snapshot Schedule: hourly
Max Average Replication Latency Threshold: 10ms
Transition to Out of Sync by Latency: false
Comment: Policy for SnapMirror Synchronous where client access will not be disrupted on
replication failure
Total Number of Rules: 1
Total Keep: 2
Rules:
SnapMirror Label Keep Preserve Warn Schedule Prefix
-----
sm_created      2    false    0    -    -
```

- **StrictSync**

In StrictSync mode, application I/O operations are sent in parallel to primary and secondary storage systems. If the I/O to the secondary storage system does not complete for any reason (ONTAP, storage, network, and so on), then the application I/O fails, and synchronous replication is terminated. This event helps to make the primary and the secondary volumes identical with zero data loss. In this case, SnapMirror also tries to bring the relationship back in synchronization automatically.

If the primary storage system becomes inoperable, the application I/O can be failed over to and resumed on the secondary storage system with manual or scripted actions. Depending on the necessary failover actions, this process makes sure that the volume has an RPO equal to zero.

```
Primary::snapmirror policy*> show -vserver vs1 -policy StrictSync
Vserver: cluster3
SnapMirror Policy Name: StrictSync
SnapMirror Policy Type: strict-sync-mirror
Policy Owner: cluster-admin
Tries Limit: 8
Transfer Priority: normal
Ignore accesstime Enabled: false
Transfer Restartability: always
Network Compression Enabled: false
Common Snapshot Schedule: hourly
Max Average Replication Latency Threshold: 10ms
Transition to Out of Sync by Latency: false
Comment: Policy for SnapMirror Synchronous where client access will not be disrupted on
replication failure
```

```

Total Number of Rules: 1
Total Keep: 2

Rules:
SnapMirror Label Keep Preserve Warn Schedule Prefix
-----
sm_created      2    false    0    -    -

```

## Operational states of SM-S

SM-S relationship will be in one of many possible operational states. These states provide insight into the operational status of SM-S. Table 2 describes the different relationship states that can occur between the source and destination volumes.

**Table 2) Relationship status.**

Relationship status	Details
Idle	The idle state has the following characteristics: <ul style="list-style-type: none"> <li>No transfer is in progress.</li> <li>If Healthy is true, this is a newly created relationship.</li> <li>If Healthy is false, auto-resync has terminated after all default attempts have failed.</li> </ul>
Quiescing	Synchronous replication suspension is in progress.
Quiesced	Synchronous replication is suspended.
Transferring	Asynchronous transmission phase is in progress when a synchronous relationship is being established or reestablished.
Transitioning	The transitioning state has the following characteristics: <ul style="list-style-type: none"> <li>Transformation from asynchronous to synchronous replication has started.</li> <li>Incoming ops are executed on primary and applied to secondary after the last asynchronous transfer.</li> </ul>
InSync	Incoming I/O successfully applied to primary and secondary, and the replication path is active.
OutOfSync	The destination volume is not in sync with the source volume because SnapMirror replication is not happening due to some failure in the data replication path. If the mirror state is SnapMirrored, this indicates a transfer failure or failure due to an unsupported operation. <p>The following can result in a change to the OutOfSync state:</p> <ul style="list-style-type: none"> <li>Unplanned non-disruptive operations</li> <li>Network errors due to congestion or equipment failure which results in packet loss</li> <li>Not receiving replication partner "heartbeat"</li> <li>Destination cluster failure</li> <li>Destination cluster IO failure notification</li> <li>User-initiated management operations such as abort, quiesce, release, etc..</li> </ul>

## SM-S behavior when state is OutOfSync

If SM-S is operating in StrictSync mode, an IO error is sent to the requesting application server to manage as required for the host file system and application.

If SM-S is operating in Sync mode, host application server IO will continue to be processed by the source volume.

Table 3 provides a list of source volume operations allowed or blocked when volume is in OutOfSync state.

**Table 3) Operations allowed when volume is in OutOfSync state**

Operation	StrictSync mode	Sync mode
Write to source volume	Blocked	Allowed
Read from source volume	Blocked	Allowed
Read from snapshot	Allowed	Allowed
Snapshot creation	Blocked	Allowed
Snapshot deletion	Blocked	Allowed
Clone creation from existing snapshot	Allowed	Allowed
Clone creation from new snapshot	Blocked	Allowed
Sis-clone creation	Blocked	Allowed
Volume move	Allowed	Allowed

Once SM-S transitions to an OutOfSync state, the destination cluster will enter an auto-resync mode. If the initial auto-resync attempt fails, retries will be attempted as below:

- Prior to ONTAP 9.9.1, auto-resync will be retried with a fixed interval of 5 mins after the previous attempt for a maximum of 5 attempts. If after five consecutive unsuccessful attempts to resynchronize the volumes, SM-S transitions to an IDLE state with Healthy set to False indicating that manual intervention is required to correct any issues.
- Starting with ONTAP 9.9.1, auto-resync will be retried with progressively increasing intervals starting with 5 mins from the initial attempt and adding an additional five minutes to the previous retry interval until it is either successful or has reached a maximum interval of 30 minutes. SM-S continues to retry auto-resynchronization every 30 minutes until it is successful or manual intervention resolves the issue.

## Implementing SM-S

### Prerequisites

Before a volume can be protected using SM-S, the following prerequisites are required:

- Each node that will participate in the SM-S relationship must have at least 16GB of DRAM.
- Each node that will participate is supported for SM-S:
  - NetApp AFF A-Series
  - NetApp AFF C-Series
  - NetApp AFF All-SAN Array (ASA)
  - NetApp ASA A-Series
  - NetApp ASA C-Series
  - NetApp FAS
  - NetApp ONTAP Select
- All nodes are running ONTAP 9.5 or later.

- A cluster peer relationship must be created between the participating clusters.
- SVMs on the secondary cluster must be created.
- SVM peer relationships must be created between the SVMs that will host source and destination volumes for a given relationship. This SVM peer relationship will support protecting multiple volumes within the source SVM.
- Volumes on the destination SVM must be created with the same or greater capacity as the source volumes being protected. These volumes must be of type DP.
- Network:
  - At least one intercluster LIF must be created on each node in both the source and destination clusters.
  - NetApp recommends a dedicated, high-bandwidth, low-latency intercluster network path.
  - The intercluster network can be physical or logical ports (port group, ifgroup, or VLAN).
  - Round trip time (RTT) must be 10ms or lower.

## Scalability

The number of concurrent SM-S relationships supported per HA pair varies depending on the ONTAP version and the array model type. Table 4 provides detailed scalability information for SM-S.

**Table 4) Concurrent replication relationships allowed per high-availability (HA) pair.**

	ONTAP 9.8 or earlier	ONTAP 9.9.1	ONTAP 9.10.1	ONTAP 9.11.1	ONTAP 9.12.1	ONTAP 9.13.1 or later
AFF A-Series	80	160	200	400	400	400
AFF C-Series	N/A	N/A	200	400	400	400
AFF ASA	80	160	200	400	400	400
ASA A-Series	N/A	N/A	N/A	N/A	N/A	400
ASA C-Series	N/A	N/A	N/A	N/A	N/A	400
FAS	40	80	80	80	80	80
ONTAP Select	20	40	40	40	40	40

**Note:** These limits represent the total number of relationships per HA pair.

When replicating between HA pairs of different families (e.g. AFF source volume to FAS destination volume) the maximum number of concurrent operations will be the limit imposed by the least scalable HA pair (e.g. 80 instead of 400 in the hypothetical AFF → FAS example).

The concurrent relationship scale is applied only if all the nodes are running the indicated ONTAP version or later. Where the source and destination clusters have different ONTAP versions, the limit will be the lower of the two endpoints.

These limits apply to both the surviving node and the failed-over node of an HA pair. If more than the maximum number of SnapMirror volume replications are scheduled to run concurrently, each additional transfer generates an error message stating that resource limits have been reached.

Each transfer beyond the maximum is retried once per minute until either it succeeds, SnapMirror is turned off, or the update is terminated.

If the source and destination volume are hosted on the same HA pair, the relationship counts as two for the purposes of determining the maximum concurrent relationships.

## Limitations

The following limitations must be considered during SM-S planning:

- The following ONTAP features are not supported with SM-S:
  - NetApp SnapLock® volumes
  - FlexGroup volumes
  - NetApp FlexCache® volumes
    - SM-S cannot replicate from or to a FlexCache volume.
    - The SM-S source and target volumes cannot be an origin volume for FlexCache.
- An SM-S primary volume cannot be a destination (DP) volume for any SnapMirror relationship (cascade). SM-S must be the first SnapMirror relationship in any cascade configuration.
- Automatic failover is not supported. All failover operations must be manually started using the `snapmirror break` command.
- Fanout:
  - A source volume can have only one SM-S relationship.
  - Secondary relationships must be SnapMirror asynchronous and should use `MirrorAllSnapshots` or `MirrorAndVault` policies.
- I/O performance:
  - Flash and spinning media are supported in the same volume relationship.
  - The performance will be throttled by the lowest-performing media type.

## Creating SM-S volume relationships

SM-S replication relationships can be created using the following ONTAP management tools:

- NetApp ONTAP System Manager
- NetApp ONTAP CLI
- NetApp ONTAP REST APIs
- NetApp BlueXP (on-premises to on-premises only)

To create an SM-S relationship using the ONTAP CLI, use the `snapmirror create` command with the `-policy` of `Sync` or `StrictSync`.

```
Destination:> snapmirror create -source-path vs1:vol1 -destination-path vs1_sync_dr:vol1 -policy
StrictSync
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:vol1.
```

```
Destination:> snapmirror create -source-path vs1:vol1 -destination-path vs1_sync_dr:vol1 -policy
Sync
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:vol1.
```

## Monitoring and managing SM-S

After the SM-S relationship is active, it will be a necessity to monitor the health of that relationship. There are several ways that an administrator can monitor SM-S relationship health, such as:

- **CLI and REST API.** Through scripting or manual execution, the ONTAP CLI or REST APIs can be used to periodically provide SM-S status and health. For example, using the ONTAP CLI, you can execute the `snapmirror show` command to show the fields `state`, `status`, `health` and `last-transfer-error`.

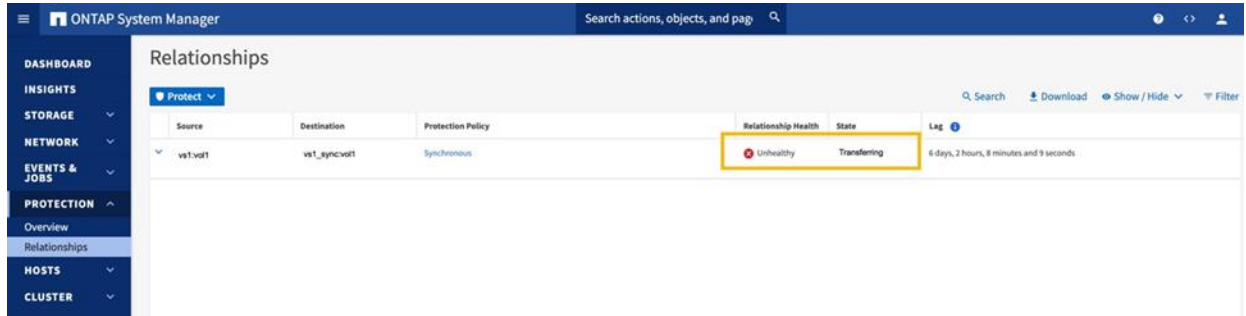
```
Source:> snapmirror show -fields state, status, health, last-transfer-error
source-path destination-path state          status          healthy last-transfer-error
```

```
vs1:vol1 vs1_sync:vol1 Snapmirrored Transferring false Prechecks on source volume failed.
(CSM: A get-session operation failed because no (local) transport address was registered for the
node.)
```

Refer to the [Operational states of SM-S](#) in this document.

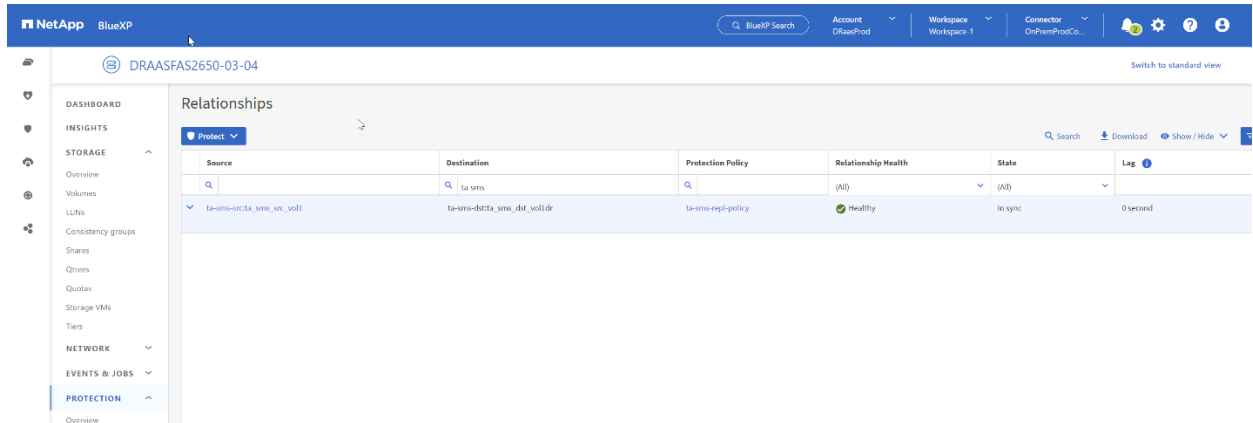
- **System Manager.** System Manager provides a real-time view of the health status of the SM-S volume relationship on the Protection > Relationships page of the destination cluster (Figure 3).

**Figure 3) Monitoring SM-S using System Manager.**



- **BlueXP.** BlueXP in advanced mode provides a view of the health status of the SM-S volume relationship on the Protection > Relationships page (Figure 4).

**Figure 4) Monitoring SM-S using BlueXP advanced mode.**



## Advanced SM-S topics

### Deploying an intra-cluster SM-S deployment

Starting with ONTAP 9.14.1, NetApp has qualified SM-S for replication between two volumes within the same cluster. These volumes may be in the same SVM or two different SVMs within the same cluster. These volumes can be hosted on different HA pairs or the same HA pair.

Hosting the source and destination SM-S volumes will not exhibit the same level of availability that hosting the source and destination volumes on different clusters, in different locations would provide due to reducing the impact of redundant hardware provided by different clusters. To minimize the impact on availability when implementing an intra-cluster SM-S deployment, the following items are recommended:

- If possible, place the source and destination volumes on different HA pairs to enhance HA.

- If possible, consider placing the source and destination volumes on different aggregates to enhance HA.

**Note:** If the entire cluster becomes unavailable due to a site or catastrophic power failure, access to both source and destination volumes will not be possible until the entire cluster is recovered. In this scenario, data is not lost, it is simply inaccessible.

**Note:** Since both the source and destination volumes are on the same cluster, the relationship will count against each HA pair's SM-S relationship limit. If the source and destination volume are hosted on the same HA pair, this relationship will count twice for purposes of determining the HA pair's SM-S relationship limit.

## Restoring an SM-S source volume from an NDMP backup

Starting with ONTAP 9.13.1, an SM-S source volume, can be restored from an NDMP backup. To support this feature both the source and destination clusters must be running ONTAP 9.13.1 or later. SM-S will replicate changes made during the NDMP restore operation to the SM-S volume's secondary destination.

The restore supports the following NDMP restore types:

- NDMPcopy
  - Full volume restore
  - Single File
  - Online LUN or Namespace
  - Folders and sub-folder trees
  - Files with streams and ACLs
- NDMP restore
  - Baseline (full volume)
  - Incremental
  - DAR (direct access recovery)
  - Folders and sub-folder trees
  - Individual files with rename
  - Files with streams and ACLs
  - Extract restore
  - Audit (restore with NoWrite)
  - Restore with list

## SM-S support for increased FlexVol volume sizes

Starting with ONTAP 9.12.1P2, FlexVol volume maximum sizes increased from 100TB to 300TB and increased the maximum size for individual files and LUNs from 16TB to 128TB for AFF A-Series, AFF C-Series, ASA, FAS, and NetApp Cloud Volumes ONTAP platforms.

Starting with ONTAP 9.14.1, ONTAP supports new storage efficiency modes for FlexVol volumes and changes in capacity reporting. SnapMirror synchronous supports volumes reporting capacity up to 600TB.

Large FlexVol volumes are enabled on a per-volume basis using the `-is-large-size-enabled true` parameter when creating or modifying a volume.

SM-S supports the replication of large FlexVol volumes with the following limitations:

- Both source and destination nodes must be running ONTAP 9.12.1P2 or later.
- Both source and destination FlexVol volumes must have the `-is-large-size-enabled true` parameter configured.



- If a destination node is ever reverted to an ONTAP version earlier than 9.12.1P2, SM-S replication will fail until both source and destination volumes have the same volume size support capabilities or large file support is disabled on the source volume.

## SM-S support for non-disruptive operations

### ONTAP 9.12.1 and later

Starting with ONTAP 9.12.1, SM-S will continue to operate without error during planned and unplanned nondisruptive operations (NDO) events. SM-S will continue uninterrupted during the following operations:

- Planned NDO on source and destination:
  - Moving FlexVol volumes between aggregates
  - Aggregate relocation
  - HA takeover and give-back operations.
- Unplanned NDO on source and destination:
  - Controller disruptions
  - Node failures

Execution of ONTAP NDO will no longer result in I/O disruption for both Sync and StrictSync SM-S policies.

**Note:** Both source and destination nodes hosting a protected volume must be AFF (A-Series or C-Series) or ASA models for SM-S resiliency during NDO.

**Note:** Both source and destination clusters must be running ONTAP 9.12.1 or later for SM-S resiliency during NDO.

During upgrades to 9.12.1, SM-S will still result in a temporary SM-S disruption, as described in the next section.

### ONTAP 9.11.1 and earlier

Execution of ONTAP NDO could result in I/O disruption depending on the SM-S policy type:

- For Sync relationships, the volume would show a status of OutOfSync, but application I/O will continue.
- For StrictSync relationships, NDO execution results in a replication failure by SM-S causing disruption of application I/O. After completion of such operations, a resynchronization of protected volumes from the last exported common snapshot starts resulting in several minutes of disruption.

Upgrading from and to any ONTAP version prior to ONTAP 9.12.1 will result in the following temporary degradation of SM-S operation:

- For Sync relationships, the volumes protected with SM-S will show Out of Sync when moving the volume between the nodes of an HA pair.
- For SM-S relationships of type `StrictSync`, I/O failure will result as the source and destination volume ownership is transitioned to the alternate node in the HA pair.

## Support for NFS 4.2 sparse files and extended attributes

Starting with ONTAP 9.12.1 and later, SM-S supports replication of NFS 4.2 volumes using extended attributes and sparse files.

## NFS extended attributes

Extended attributes (xattrs) are a means to associate opaque metadata with file system objects, organized as key/value pairs. They are especially useful when they add information that is not or cannot be, present in the associated object itself. Extended attributes are defined in [IETF RFS 8276](#).

All major operating systems provide facilities to access and modify extended attributes and many advanced file systems require xattrs to operate efficiently. Many user-space tools allow xattrs to be included together with regular attributes that need to be preserved when objects are updated, moved, or copied.

Starting with ONTAP 9.12.1, ONTAP supports NFS 4.2 xattrs of 2KB or smaller. In conjunction with the expanded NFS support for xattrs, SM-S will replicate xattrs associated with NFS 4.2 enabled FlexVol volumes.

## NFS sparse file and space reservation support

NFS 4.2 introduces the concept of sparse files through the implementation of space reservation guarantees without storing data within a file.

- **Sparse files** are files that have regions of unallocated or uninitialized data blocks that look like “holes” in the file. When clients do read operations on a sparse file, the data in the holes is represented as zeros.
- **Space reservations** guarantee that space is available on disk for future writes to the file.

Space reservations guarantee that a sparse file has disk blocks reserved, so no future writes into the sparse file can fail.

Examples of sparse files include virtual machine (VM) disk images, database files, log files, and even checkpoint recovery files used by the high-performance computing (HPC) community. Sparse files and space reservations are complementary features because such applications expect space reservations to work in conjunction with sparse files to ensure enough data blocks are available for future writes on the file.

For details on NFS sparse files, see [RFC 7862: Network File System \(NFS\) Version 4 Minor Version 2 Protocol](#).

## Accessing the secondary volume

SM-S allows access to the DP volume on the destination cluster in a read-only mode at any time. This view of the protected volume is not a real-time view of the active filesystem, but a view of the volume based on the latest exported common snapshot (See [Common snapshots](#) for details). Depending on the configured common snapshot schedule, this view of the volume data could be between 30 minutes and 24 hours old.

FlexClone volumes can be created from the latest exported common snapshot to provide read-write access to data in that most recent snapshot. This FlexClone volume is not based on the current state of the volume, but only that of the last exported common snapshot.

If a more up-to-date view of the protected volume’s data is required, manually create a new common snapshot using the `snapmirror update` command then create a new FlexClone volume using the `-parent-snapshot` parameter designating this latest common snapshot as the clone source.

```
Destination::> snapmirror update -destination-path vs1:vol1
Operation is queued: snapmirror update of destination "vs1:vol1".

Destination::> volume clone create -vserver vs1 -flexclone flexvol1 -parent-vserver vs1 -parent-
volume vol1 -parent-snapshot snapmirror.8e36712a-6073-11ed-8a4f-00a098d41def_2157573437.2022-11-
16_190358
[Job 11524] Job succeeded: Successful
```

## Replicating application created snapshots

SM-S will create replica snapshots located on the source volume that are either manually created (e.g. created using the `volume snapshot create` CLI command) or application created (e.g. created using external tools such as NetApp SnapCenter®) snapshots. Which snapshots are replicated will depend on the ONTAP version currently deployed on the source and destination:

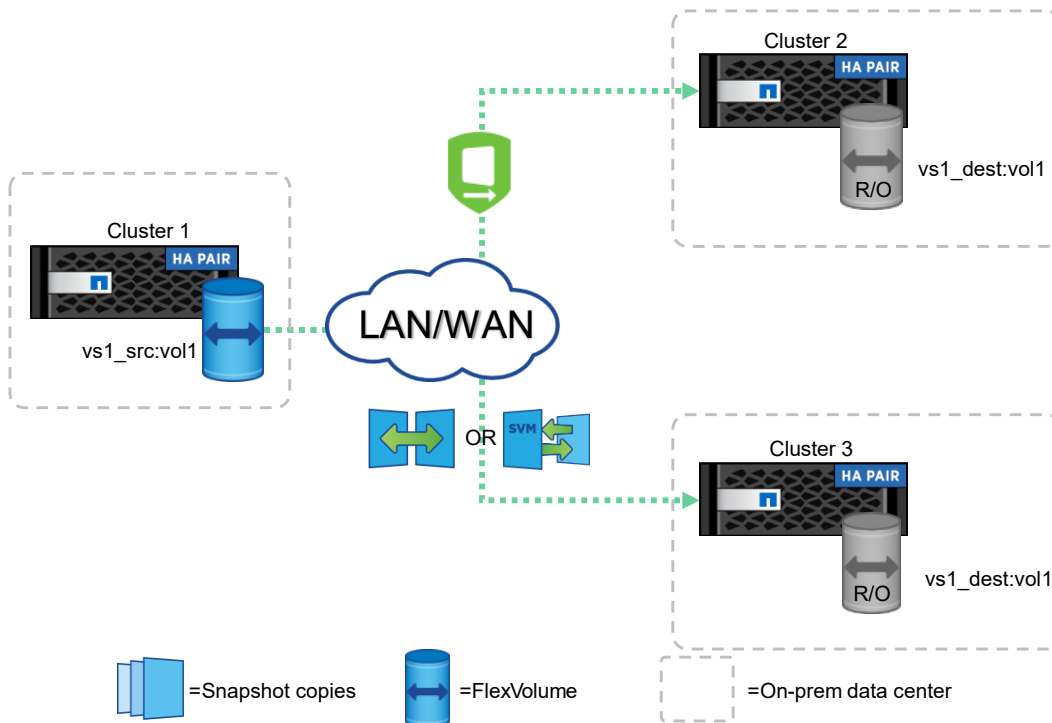
- ONTAP 9.13.1 and earlier SM-S will replicate only the application created snapshots that have a `snapmirror` label that matches a SnapMirror policy rule.
- ONTAP 9.14.1 and later SM-S will replicate all application created snapshots regardless of `snapmirror` label.

**Note:** To disable replication of application created snapshots to the DR site, change the SnapMirror policy parameter `-enable-acs-replication` to `false` available in advanced mode.

## Fan-out support

A source volume that is protected using SM-S can also be asynchronously replicated to another destination SVM. This asynchronous destination SVM can reside on the same destination cluster or a different destination cluster. If both SM-S and asynchronous SnapMirror relationships are using the same target cluster, each relationship must target a different SVM. This can be either a volume-scoped SnapMirror asynchronous relationship or an SVM-scoped (SVM-DR) relationship (Figure 5).

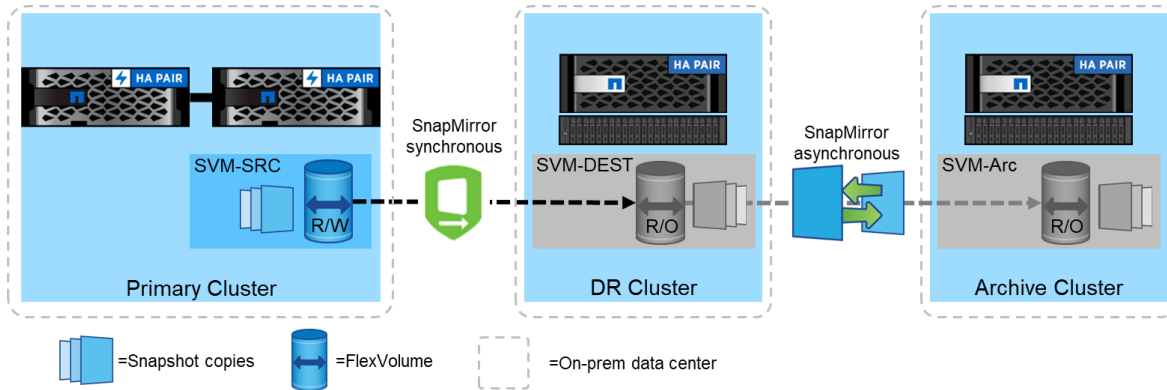
Figure 5) Fan-out support with SM-S.



## Cascade support

SM-S can be a part of a cascaded SnapMirror relationship but must be the initial relationship for the source volume. Downstream relationships must be asynchronous as shown in Figure 6. The asynchronous relationship will only replicate the latest exported snapshot of the source volume which does not reflect the real-time state of the active file system.

Figure 6) SM-S role in a SnapMirror cascade relationship



Also, the active file system (AFS) and any application-created snapshots (ACS) from the SM-S source relationship may not be immediately visible to the secondary SnapMirror relationship for cascade depending on ONTAP version:

- **ONTAP 9.13.1 and earlier:** The cascade relationship will not have visibility to the SM-S replicated ACS snapshots until the SM-S relationship has been updated and a new common exported snapshot is available. The result of this behavior is that an ACS snapshot created on the SM-S source volume that is duplicated on the secondary will be delayed in being replicated to the tertiary volume.
- **ONTAP 9.14.1 and later:** The asynchronous SnapMirror cascade relationship no longer must wait for SM-S to create a new common exported snapshot. The result of this change in behavior is that ACS snapshots created on the SM-S source will be visible and replicated based on the cascade SnapMirror relationship to the tertiary site.

## Performing accessibility testing of DR site data

One of the key activities of any DR plan is to periodically perform a DR rehearsal by failing over to the DR site, restarting applications, and then recovering back to the primary site. This section provides guidance on successfully performing a DR rehearsal while minimizing downtime due to recovery to primary site. The process described here maintains the production application environment while allowing a DR test to be run. This maintains RPO=0 on production systems.

As with all SnapMirror data protection solutions, this process involves the following steps from the destination cluster:

1. Perform an update operation to update the common snapshot:

```
Dest::> snapmirror update -destination-path dest_SVM:voll-dr
```

2. Create a FlexClone volume of the DR volume.

```
Dest::> volume clone create -vserver dest_SVM -flexclone temp-voll-dr -parent-volume voll-dr -junction active true -foreground true -comment "DR test volume for voll"
```

3. Perform application-level tests using test client servers using the FlexClone volume.
4. Delete FlexClone volume when finished tests

```
Dest::> volume delete -vserver dest_SVM -volume temp-voll-dr
```

## Performing DR rehearsals

One of the key activities of any DR plan is to periodically perform a DR rehearsal by failing over to the DR site, restarting applications, and then recovering back to the primary site. This section provides guidance on successfully performing a DR rehearsal while minimizing downtime due to recovery to the primary site.

**Note:** Performing this procedure does not maintain the production volume's RPO=0 state as the SM-S relationship is modified to activate the DR volume and resync any changes made on the DR volume back to the primary application volume once the rehearsal is complete.

As with all SnapMirror data protection solutions, this process involves the following steps:

1. Failover to DR site volume:

```
Dest::> snapmirror quiesce -destination-path dest_SVM:voll-dr
Dest::> snapmirror break -destination-path dest_SVM:voll-dr
```

2. Perform application-level tests using the volume on the DR site.

3. Resync any changes back to the production site after testing is completed:

```
Source::> snapmirror create -destination-path source_SVM:voll -source-path dest_SVM:voll-dr -
policy Sync
Source::> snapmirror resync -destination-path source_SVM:voll
```

4. Delete the reverse relationship and reenble the primary SM-S relationship.

```
Source::> snapmirror quiesce -destination-path source_SVM:voll
Source::> snapmirror break -source-path dest_SVM:voll-dr
Source::> snapmirror delete -destination-path source_SVM:voll
```

```
Dest::> snapmirror release -destination-path source_SVM:voll -relationship-info-only true
Dest::> snapmirror resync -destination-path dest_SVM:voll-dr
```

5. Reconnect applications to the reactivated production volume.

When performing a rehearsal, it is imperative that the following requirements be understood:

- When executing the `snapmirror release` command, you must use the `-relationship-info-only true` parameter to ensure that a common snapshot on both volumes is retained. This common snapshot is used to determine the amount of changed data to restore to the production site after the rehearsal is finished. If this parameter is not supplied, a full baseline replication of the volume from the DR site to the production site will be required before the production site can be restarted.
- The `snapmirror release` command deletes the relationship information from the source and performs a resource cleanup. For SM-S relationships using the StrictSync policy, this process removes any existing fences on the primary volume to prevent I/O disruption on the primary.
- When performing a reverse resync, all snapshots on the production volume newer than the latest common snapshot will be deleted, and all the snapshots on the DR site volume newer than the common snapshot will be transferred to the production site.
- When performing a reverse resync, some snapshots may be left behind after the resync operation finishes, and the reverse relationship is deleted. These snapshots will have to be manually deleted as they will continue to consume storage capacity on the DR cluster.

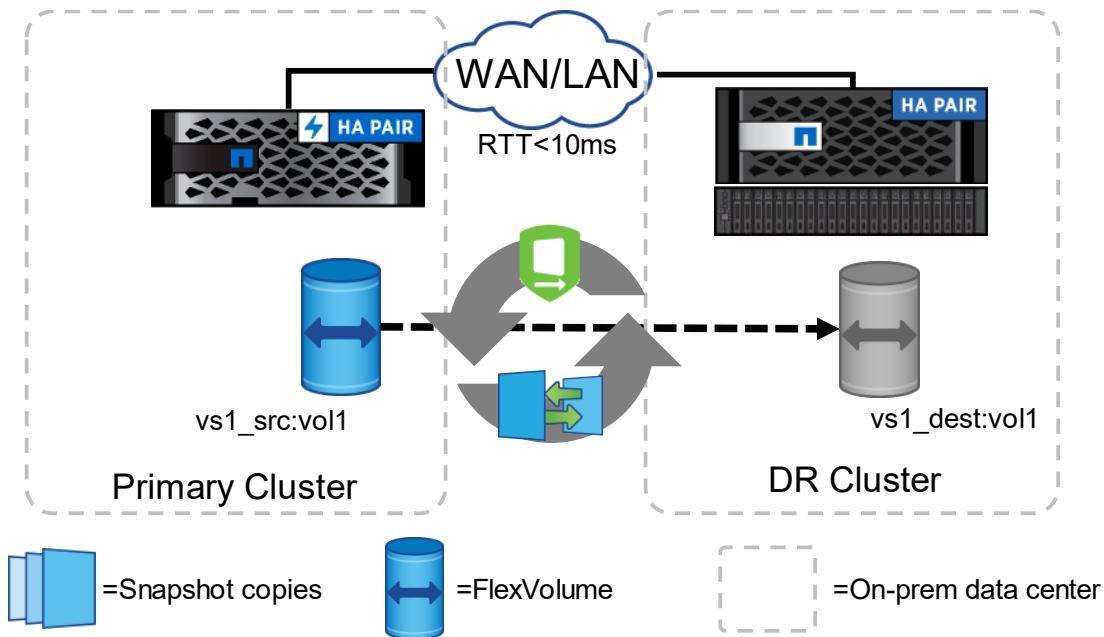
## Converting to or from SM-S

Existing SnapMirror relationships can be changed between asynchronous and synchronous relationships. Direct conversions are not supported (Figure 7). The original asynchronous or synchronous relationship must be deleted and released, and a new asynchronous or synchronous relationship created.

**Note:** When releasing the existing relationship, you must use the `-relationship-info-only true` parameter to the `snapmirror release` command. This allows the new relationship to be

resynchronized with minimal data transfer between source and destination as the common snapshots are retained.

Figure 7) Changing SnapMirror relationship between asynchronous and synchronous modes



To convert a volume with a SnapMirror Asynchronous relationship to an SM-S relationship, the current asynchronous relationship will need consider the following:

- All SM-S [prerequisites](#) (see section Prerequisites above) must be met before converting an asynchronous relationship to SM-S
- SM-S only supports two fan-out destinations, one synchronous and one asynchronous (volume or SVM scoped). If the current volume is a source of a fan-out with greater than two relationships, the conversion will not be supported until the additional asynchronous relationships are removed.
- SM-s does not support SVM-scoped relationships (SVM-DR) but individual volumes within an SVM-DR protected SVM can have SM-S relationships. There are two options available for enabling SM-S on individual volumes within an SVM-DR protected SVM:
  - Remove the SVM-DR relationships and create SM-S relationships for each volume that needs protection. This will remove the benefits of SVM identity and configuration replication.
  - Create individual SM-S relationships for each volume within the SVM. These relationships must have a different destination SVM from that used by the SVM-DR relationship. These SM-S SVMs can be on the same or different destination cluster as the SVM-DR relationship. SVM identity is not protected by the SM-S relationship.

Read [Convert the type of a SnapMirror relationship \(netapp.com\)](#) for more details.

## Conversions to or from an SM-S StrictSync policy

SM-S relationships can be changed between Sync and StrictSync operational modes. Direct conversions are not supported. The original SM-S relationship must be deleted and released, and a new SM-S relationship created with the new operational mode.

**Note:** When releasing the existing relationship, you must use the `-relationship-info-only true` parameter to the `snapmirror release` command. This allows the new relationship to be resynchronized with minimal data transfer between source and destination.

Read [Convert the mode of a SnapMirror synchronous relationship \(netapp.com\)](https://netapp.com) for more details.

## Migration of 7-Mode synchronous SnapMirror to ONTAP 9 SnapMirror synchronous

SM-S does not support using a 7-mode source volume, but ONTAP supports converting a 7-mode SnapMirror relationship to an SM-S relationship without requiring a new baseline copy of the volume data. NetApp provides two methods for converting a 7-mode relationship to SM-S:

- 7-Mode Transition Tool (7MTT)
- ONTAP CLI using the `vserver peer transition create` command.

Consider the following items when converting a 7-mode synchronous relationship to SM-S:

- Protocols, networking, and other configuration information will need to be configured on ONTAP 9 after migration.
- Consider using 7MTT to perform conversions because it provides additional prechecks before migrating data and relationships.
- Review the list of features that are not supported for transition for [NAS](#) and [SAN](#) deployment before performing migrations.
- Be sure your 7-mode transition is supported by the 7MTT by reviewing the [interoperability list](#) in the 7MTT documentation.
- The maximum number of concurrent SnapMirror transfers may be limited based on the storage system model. Visit [Hardware Universe](#) for specific model limits.



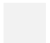
## SnapMirror synchronous replication interoperability

SM-S supports volume replication relationships between clusters running different ONTAP versions.

Starting with ONTAP 9.12.1, NetApp supports replication of a source volume running ONTAP 9.12.1 or later with a second cluster running an ONTAP version that was released no more than three years before or after the source cluster's ONTAP version. For SM-S relationships that do not involve at least one replication endpoint running ONTAP 9.12.1 or later, SM-S replication interoperability is limited to two ONTAP versions older or newer. Table 5 shows specific relationship interoperability for ONTAP versions starting with ONTAP 9.5.

**Table 5) SM-S relationship interoperability**

	9.14.1	9.13.1	9.12.1	9.11.1	9.10.1	9.9.1	9.8	9.7	9.6	9.5
9.14.1	Support	Support	Support	Support	Support	Support	Support	No support	No support	No support
9.13.1	Support	Support	Support	Support	Support	Support	Support	Support	No support	No support
9.12.1	Support	Support	Support	Support	Support	Support	Support	Support	No support	No support
9.11.1	Support	Support	Support	Support	Support	Support	No support	No support	No support	No support
9.10.1	Support	Support	Support	Support	Support	Support	Support	No support	No support	No support
9.9.1	Support	Support	Support	Support	Support	Support	Support	Support	No support	No support
9.8	Support	Support	Support	No support	Support	Support	Support	Support	Support	No support
9.7	No support	Support	Support	No support	No support	Support	Support	Support	Support	Support
9.6	No support	No support	No support	No support	No support	No support	Support	Support	Support	Support
9.5	No support	No support	No support	No support	No support	No support	No support	Support	Support	Support

 &  = Support between versions  
 = No support between versions

## Performance considerations

### Network and storage infrastructure

There are many design decisions that must be considered when building a real-time replication solution for high-performance applications. These considerations include array capabilities such as memory capacity, CPU performance, balancing storage capacity versus IO performance, network performance, and data change rate generated by the application. The following provides some guidance on sizing each of these components.

To ensure adequate CPU and memory, review [NetApp Hardware Universe](#) to identify the appropriate ONTAP controller models to fit the expected number of concurrent SM-S relationships.

Networking will most likely be the most impactful design choice for SM-S. Performing real-time replication between two clusters requires a very low latency Ethernet network infrastructure with adequate available bandwidth to support all volumes being protected by SM-S.

As already stated, any network used for intercluster SM-S replication must exhibit a stable RTT (round-trip time) of less than 10ms between the source and destination clusters. This ensures that application IO is not unduly hindered as it awaits IO write verification from the destination ONTAP system.

Net effective throughput – the throughput available to SM-S – of the intercluster network infrastructure must be adequate to support transmitting all the data writes and updates between the source and destination as well as any ONTAP metadata that must be sent with the replicated data.

For example, if your application generates 100 megabytes of peak data writes or updates, then the minimum net effective bandwidth required should be no less than 1Gbps and preferably 2Gbps. If your application's peak data writes and updates generate 1 gigabytes of peak activity, then it would be recommended to provide 10Gbps of network throughput. If you are placing intercluster LIFs on a port shared by other LIF types, you need to account for the non-SM-S traffic when selecting a port for SM-S inter-cluster replication.



If possible, NetApp recommends that applications be tested for peak throughput and latency behavior prior to designing an SM-S infrastructure. Depending on the number of volumes being protected and their data change rate, you may want to consider deploying intercluster LIFs to dedicated Ethernet ports to ensure that each application can access its data without any additional latency due to lack of available bandwidth.

## Quality of service

Quality of service (QoS) is a method of providing predictable behavior for storage I/O operations. Prior to ONTAP 9.12.1, SM-S did not support QoS settings for ONTAP. Starting with ONTAP 9.12.1, SM-S supports setting a ceiling value as well as supports adaptive QoS peak operations settings. Table 6 lists the supported QoS policy limits.

**Table 6) QoS support for SM-S.**

QoS type	ONTAP 9.11.1 and earlier	ONTAP 9.12.1
Floor	No	No
Ceiling	No	Yes. QoS policies that have maximum throughput settings are supported on the source and destination clusters.
Adaptive QoS	No	Yes. Adaptive QoS settings honor peak operations only.

**Note:** Both the primary and secondary clusters must run ONTAP 9.12.1 to support QoS.

## Mixing array models

SM-S is supported between storage clusters consisting of differing node types and storage types. When creating an SM-S volume relationship where the source and target cluster nodes have different technology – such as flash and spinning storage disks, or NVMe and SATA storage disks – the performance of SM-S will be restricted by the slower technology. In the examples above, if a source volume is hosted on flash and the destination volume is hosted on SATA spinning disks, then SM-S performance characteristics will be dictated by the cluster nodes using the slower SATA spinning disks.

## Reversion to previous ONTAP versions

If the SM-S relationship is set up, you must perform the following actions before you can revert to a previous version of ONTAP:

1. Release the SM-S relationship on the source.
2. Break and delete the SM-S relationship on the destination.
3. Verify that all newly created SM-S policies are deleted before you revert to previous version of ONTAP.

**Note:** When reverting the ONTAP software to a previous version, keep in mind that there are certain features introduced in newer ONTAP versions that are not supported by older versions.

## Where to find additional information

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

- ONTAP 9 documentation  
<https://docs.netapp.com/us-en/ontap/index.html>

- ONTAP and ONTAP System Manager documentation resources  
<https://www.netapp.com/data-management/oncommand-system-documentation/>
- NetApp product documentation  
<https://www.netapp.com/support-and-training/documentation/>

## Version history

Version	Date	Document version history
Version 3.0	February 2024	ONTAP 9.14.1 release
Version 2.0	June 2023	ONTAP 9.13.1 release
Version 1.7	February 2023	ONTAP 9.12.1 release
Version 1.6	May 2022	ONTAP 9.11.1 — Minor update
Version 1.5	December 2021	9.10.1 release
Version 1.4	November 2021	Minor update: Addition of “Performance” section.
Version 1.3	June 2021	9.9.1 release
Version 1.2	December 2020	9.8 release
Version 1.1	May 2019	9.6 release
Version 1.0	December 2018	Initial 9.5 release

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 © 2023 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.

