



Technical Report

## SnapMirror Synchronous for ONTAP 9.7

Cheryl George, NetApp  
November 2019 | TR-4733

### Abstract

This document contains information and best practices for configuring NetApp® SnapMirror® Synchronous (SM-S) replication in NetApp ONTAP® 9.7.

## TABLE OF CONTENTS

<b>1</b>	<b>Executive Summary</b>	<b>4</b>
1.1	Purpose and Scope	4
1.2	Terminology	4
<b>2</b>	<b>SnapMirror Synchronous Overview</b>	<b>5</b>
2.1	Use Cases	5
2.2	Integrated Data Protection	7
2.3	SnapMirror Synchronous Modes	8
<b>3</b>	<b>SnapMirror Synchronous Data Protection Relationships</b>	<b>9</b>
3.1	Prerequisites	9
3.2	Considerations Before Using SnapMirror Synchronous	9
3.3	Maximum Number of SnapMirror Synchronous Replication Operations	10
3.4	Restrictions Using SnapMirror Synchronous	10
3.5	SnapMirror Policies	10
3.6	Common Snapshot Copies	12
3.7	Relationship Status	12
3.8	Cluster and SVM Peering	13
3.9	Populate Source Volume with Data	14
3.10	Set Up a SnapMirror Synchronous Relationship	15
3.11	Set Up a SnapMirror Strict Synchronous Relationship	17
3.12	View Destination Volume on the Secondary Storage System	18
3.13	Interrupt Network Between Primary and Remote (Prevents Replication) and Observe the Behavior of the SnapMirror Strict Synchronous Relationship	19
3.14	Interrupt Network Between Source and Destination (Prevents Replication) and Observe the Behavior of SnapMirror Synchronous Mode	20
3.15	Failover from Primary to Secondary Data Center	20
3.16	Failover and Failback	21
3.17	Convert a SnapMirror Synchronous Relationship from Strict Synchronous Mode to Synchronous	23
3.18	Convert an Asynchronous SnapMirror Relationship to Synchronous	24
3.19	Convert a Synchronous SnapMirror Relationship to Asynchronous	27
3.20	Flip Resync	27
3.21	Migrate Existing Synchronous SnapMirror Relationships from 7-Mode to ONTAP	29
3.22	Mirror State and Relationship Status	31
<b>4</b>	<b>Manage SnapMirror Synchronous Operations Through ONTAP System Manager</b>	<b>32</b>
<b>5</b>	<b>SnapMirror Synchronous and NetApp Storage Efficiency</b>	<b>33</b>

<b>6 SnapMirror Synchronous Calculator</b> .....	<b>33</b>
<b>7 Revert to Previous ONTAP Versions</b> .....	<b>33</b>
<b>8 Troubleshooting</b> .....	<b>34</b>
8.1 Operations That Cause SnapMirror Relationships to Go OutOfSync.....	34
8.2 Resync Failures .....	40
<b>9 Limitations</b> .....	<b>44</b>
<b>Where to Find Additional Information</b> .....	<b>44</b>
<b>Version History</b> .....	<b>44</b>

**LIST OF TABLES**

Table 1) Concurrent replication operations allowed. ....	10
Table 2) Relationship status. ....	12
Table 3) Mirror state and relationship status. ....	31
Table 4) Workload metrics.....	33

**LIST OF FIGURES**

Figure 1) SM-S to a relatively local secondary data center. ....	5
Figure 2) Combining synchronous and asynchronous SnapMirror relationships.....	6
Figure 3) Three data center enterprise-application use case. ....	7
Figure 4) Integrated data protection with SM-S. ....	8
Figure 5) SnapMirror policies in ONTAP System Manager. ....	11

# 1 Executive Summary

In today's constantly connected global business environment, you must be able to protect your data from damaging natural or human-made events and recover it quickly when necessary. At the same time, you must maximize your investments to get the most out of your IT infrastructure. An effective data protection strategy is vital to prevent your operations from stalling, which could result in lost productivity and revenue, and damage to your company's reputation. Furthermore, if you reuse your secondary facility for business intelligence or development and testing, you can turn your backup and disaster recovery solution into a business accelerator.

NetApp SnapMirror Synchronous (SM-S) is a cost-effective, easy-to-use disaster recovery solution that replicates data synchronously at high speeds over LAN or metro area network (MAN). It provides high data availability and fast disaster recovery for your business-critical applications such as Microsoft Exchange Server, Microsoft SQL Server, and Oracle, in both virtual and traditional environments.

SM-S reduces management overhead and simplifies data protection to provide zero-data-loss data replication. SM-S offers a zero recovery point objective (RPO), along with the flexibility to protect a subset of volumes in the cluster. Replication can also occur between ONTAP storage systems of different models, thereby lowering your TCO and accelerating your business. While SM-S facilitates business continuity, you can also choose to clone the secondary data and make it readable.

## 1.1 Purpose and Scope

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 NetApp Snapshot™ technology, NetApp FlexVol® volumes, and NetApp FlexClone® volumes
- General knowledge of disaster recovery and data replication solutions
- Familiarity with the [ONTAP 9 Data Protection Power Guide](#) on the NetApp support site

## 1.2 Terminology

- **Primary or source.** The source volume of SM-S.
- **Secondary or destination.** The destination volume of SM-S.
- **SnapMirror.** A disaster recovery 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.
- **SnapMirror Synchronous (SM-S).** The NetApp SnapMirror technology that enables you to synchronously mirror your data and continually update the mirrored data, so that your data is kept current and is available whenever you need it.
- **SnapMirror Synchronous mode.** Provides zero RPO replication with no primary I/O restriction if a replication failure occurs.
- **SnapMirror Strict Synchronous (StrictSync) mode.** provides zero RPO replication but stops primary I/O if a replication failure occurs.
- **InSync.** An SM-S relationship is actively replicating each application I/O to the secondary storage system.
- **OutOfSync.** The application I/O is not replicating to the secondary storage system.
- **Data protection relationship.** The relationship between the source volume in primary storage and the destination volume in secondary storage.

- **Recovery point objective (RPO).** How much data loss your business application can tolerate.
- **Recovery time objective (RTO).** How fast you can restart a failed application.

## 2 SnapMirror Synchronous Overview

Beginning with ONTAP 9.5, SM-S is volume-granular synchronous data replication that enterprises depend on for backup, disaster recovery, and data mobility. SM-S replicates data on FlexVol volumes between ONTAP storage systems situated across your data center or metro region to achieve zero RPO. SM-S is targeted at relatively short distances with a <10 millisecond round trip time (RTT) to provide exact replicas that can be used in any kind of failure. SM-S gives you the flexibility to protect a subset of volumes in the cluster with replication between ONTAP storage systems that can be of different controller models. This functionality addresses the national, regulatory, and industry-mandated need for synchronous replication in industries such as finance and healthcare or any other industries that depend on zero data loss. SM-S does not require any new or special hardware, software, or networking capabilities, which significantly lowers the overall cost of the solution.

### 2.1 Use Cases

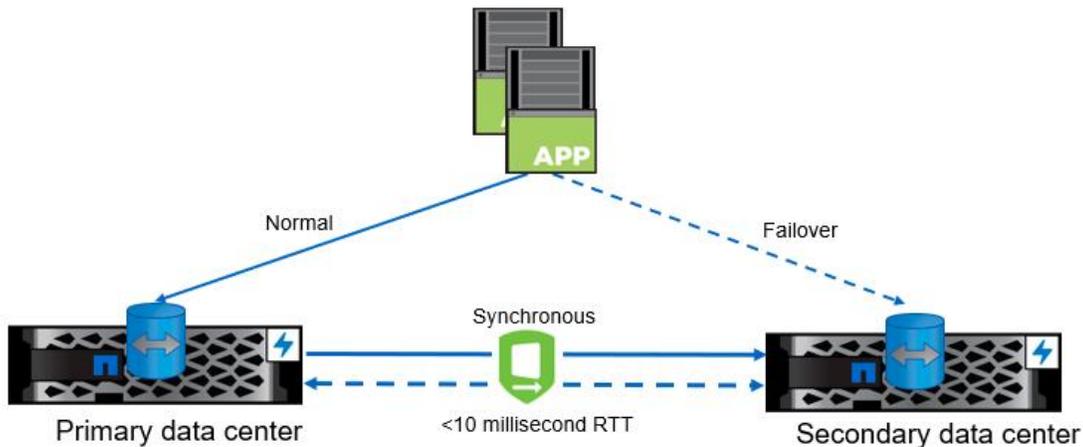
This section describes two examples of solutions using SM-S.

#### Disaster Recovery Between Two Local Data Centers

When an application's workload is suitable and a secondary data center is available within a relatively short distance, it is possible to synchronously replicate all the application's data (log files as well as data) to the secondary data center. Synchronous replication provides zero data loss and very low recovery times from failovers.

If the primary site fails, the enterprise application can failover to the replica volumes in the local secondary data center, which typically is in the same metro area as the primary data center.

Figure 1) SM-S to a relatively local secondary data center.

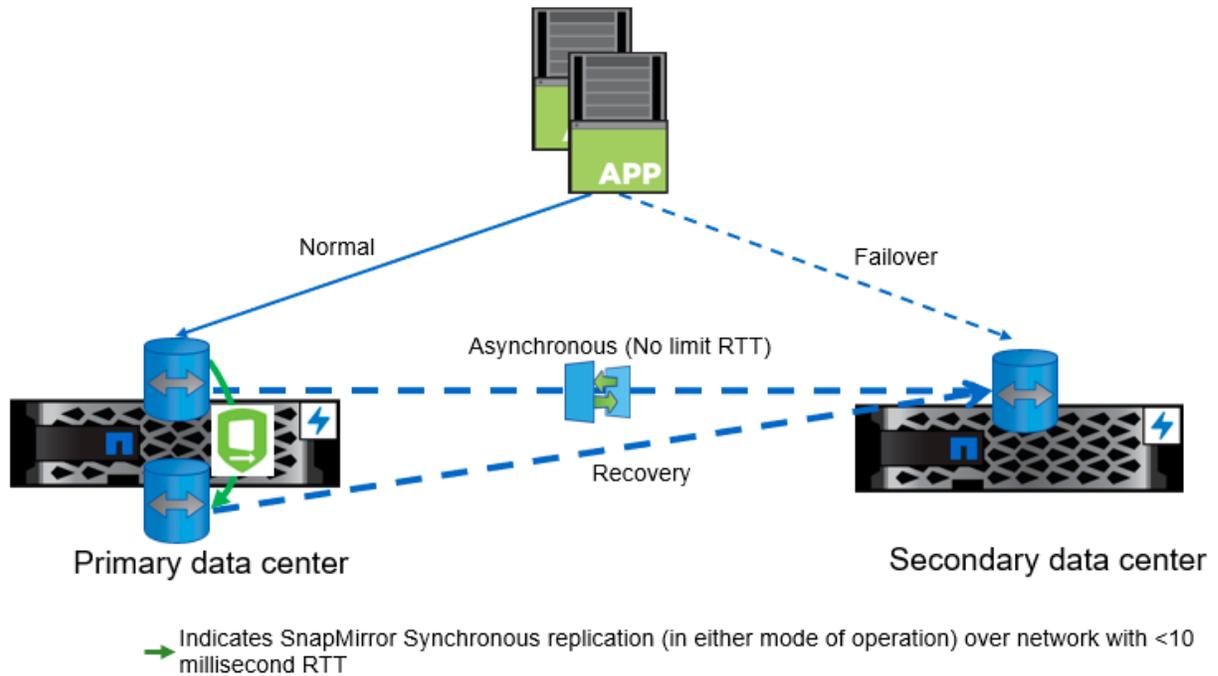


#### Combining Synchronous and Asynchronous Replication Between Two Data Centers

If failure of the primary data center is not considered a risk and you want to limit the effect of RTT on application performance, then an alternative is to synchronously replicate the log volume in the primary data center. Then, at the same time, you can asynchronously replicate the data volume to the secondary

data center. If the primary cluster fails, then the application fails over to the secondary data center, which recovers the data files and rolls forward the log data from the primary data center to achieve zero RPO.

Figure 2) Combining synchronous and asynchronous SnapMirror relationships.

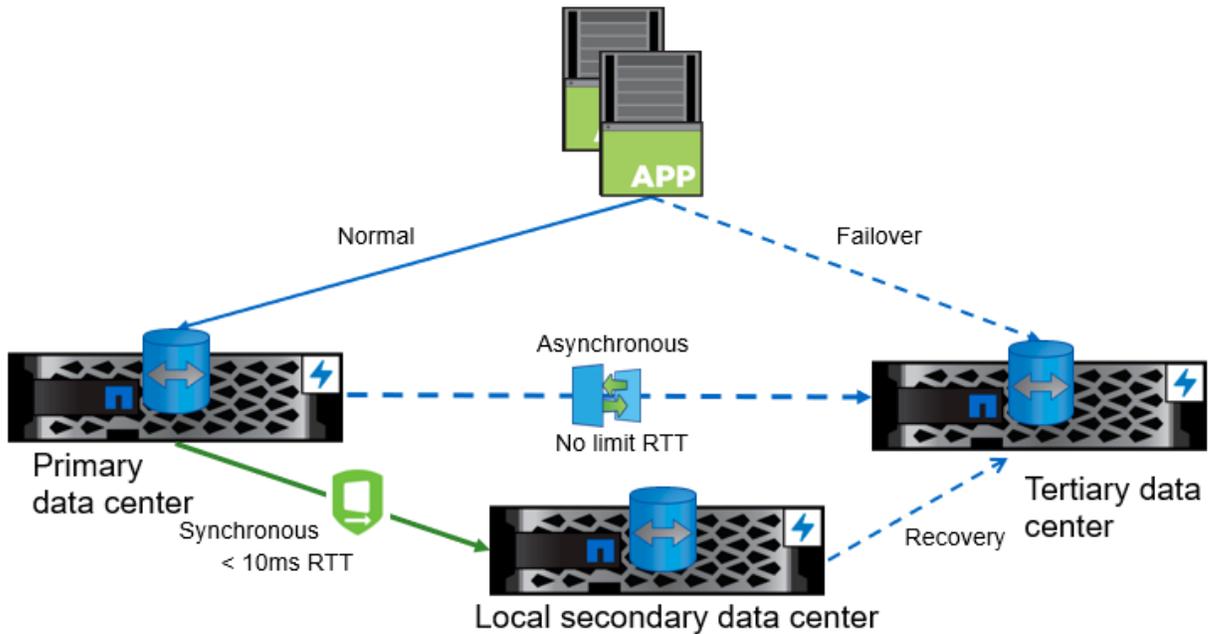


## Application Disaster Recovery Among Three Data Centers

To achieve a zero RPO over a longer distance, set up one synchronous relationship for log volumes to a relatively local secondary data center that is in a different fault zone than the primary. Simultaneously, set up an asynchronous relationship for regular data from the primary to the more remote tertiary data center. This architecture is popular with banks for regional regulatory compliance where the local secondary data center is typically in the same metro area as the primary and is not used for failover purposes. The tertiary data center is a time zone away.

Typical regulations dictate that the near disaster recovery synchronous destination should not be used for failover purposes. If the primary site fails, the enterprise application is failed over to the tertiary data center and the logs are rolled forward from the secondary location, bringing the data to the same level as the source. In this case, a zero RTO is not as critical as the RPO, allowing the enterprise application to operate from the far disaster recovery site as quickly as possible and fail the application back to primary at the first available opportunity.

Figure 3) Three data center enterprise-application use case.



#### Best Practice

- Set up an asynchronous SnapMirror relationship with the MirrorAllSnapshots policy and a schedule more frequent than the synchronous relationships common snapshot schedule. This configuration ensures that a common snapshot is available between the synchronous and asynchronous relationships.

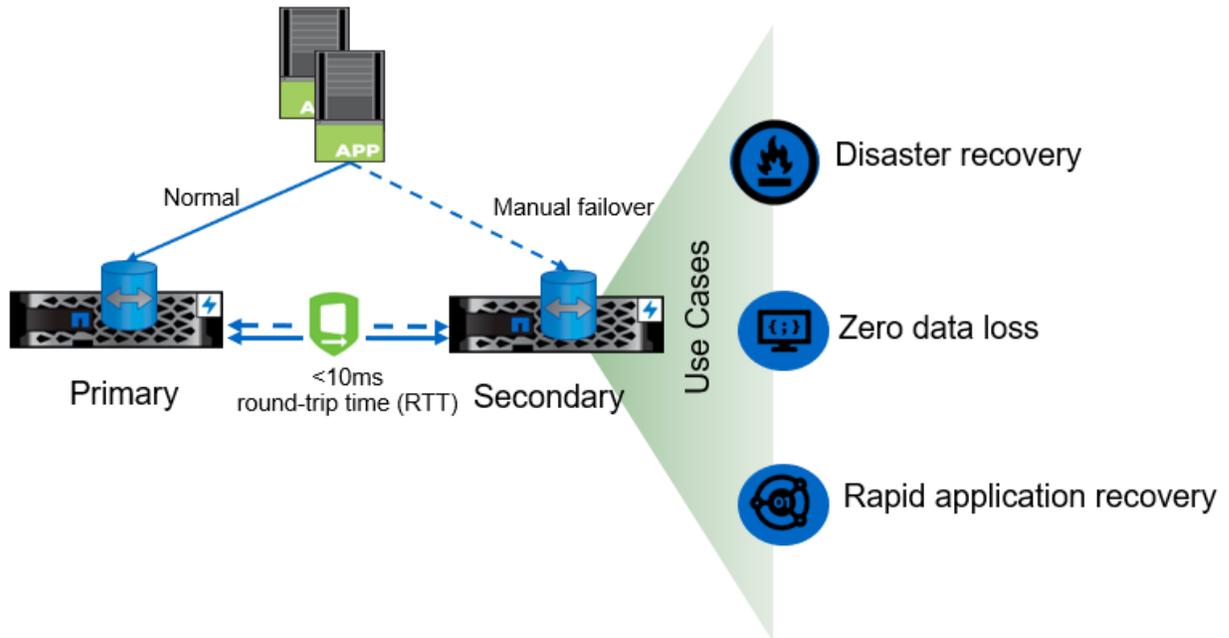
#### Benefits

- Minimal compute is required at the local disaster recovery site because the application is never actually failed over to it.
- Lower-end hardware can be used at the distant disaster recovery site, with degraded post-disaster application performance.
- Software licensing costs are reduced for the local and distant disaster recovery sites.

## 2.2 Integrated Data Protection

NetApp SM-S tightly integrates with NetApp Snapshot technology to create cost-effective, common Snapshot copies, enabling quick and easy resynchronization between data centers. SM-S in ONTAP 9.5 is completely different in architecture from 7-mode Sync SnapMirror. SM-S intercepts application writes before they reach primary WAFL. It then splits the write to secondary WAFL so that WAFL processes it using NVRAM logging and consistency checkpoint cycles on source and destination storage systems separately. Synchronous replication is based on individual I/O and is not snapshot based.

Figure 4) Integrated data protection with SM-S.



## SnapMirror for Disaster Recovery

You can synchronously replicate critical log data to a second location in the same metro area to achieve zero RPO by simply rolling forward the logs to the data file volume that has been asynchronously replicated. In this configuration, a failure in hardware or software does not result in data being unavailable for extended periods, because the business-critical enterprise application can be switched over to the disaster recovery site. When the production site is back online, SnapMirror can resynchronize data efficiently back to the production site.

## Zero Data Loss

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. You can use SM-S replication software to safeguard data from loss due to natural disaster, fire, application failure, or software malfunction.

IT demands are changing rapidly, so you must rapidly repurpose and reconfigure data centers to cater to the business drivers of zero RPO and zero RTO.

## 2.3 SnapMirror Synchronous Modes

SM-S can be configured in either of the two modes, specified by the SnapMirror policies:

- Synchronous
- Strict Synchronous (StrictSync)

### SnapMirror Synchronous

In Synchronous mode, application I/O is sent in parallel to 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 primary storage to secondary storage systems.

RPO=0 and RTO is very low until a secondary replication failure occurs, at which time RPO and RTO become indeterminate but are equal to the time required to repair the issue that caused secondary replication to fail and for the resync to complete.

### **SnapMirror Strict Synchronous**

In SnapMirror StrictSync mode, application I/O is also 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. Therefore, the primary and secondary volumes remain identical (that is, zero data loss). In this case, SM-S tries to bring the relationship back in sync automatically.

If 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. This provides RPO=0 and very low RTO, depending on the failover actions necessary.

## **3 SnapMirror Synchronous Data Protection Relationships**

This section describes the different types of SM-S relationships and how to initialize and manage them.

### **3.1 Prerequisites**

The following prerequisites must be in place before you can use SM-S:

- SM-S technology is supported on all shipped FAS and AFF systems that have at least 16GB of memory and support ONTAP 9.5. SM-S is also supported on ONTAP Select platforms.
- Each node in the cluster must be running ONTAP 9.5; however, the model can vary.
- The functional network to transfer data between two different storage systems should have RTT latency of no more than 10 milliseconds. Essentially, this means LAN or MAN for distances of up to 150km.
- The cluster and storage virtual machine (SVM) peer relationships must be configured.
- The storage systems names should correctly resolve using the domain naming system (DNS).

### **3.2 Considerations Before Using SnapMirror Synchronous**

You must also consider the following when planning to use SM-S replication:

- One source volume cannot have SM-S relationships to multiple destination volumes. But one source volume can have one synchronous and one asynchronous relationship to two different destination volumes.
- You must verify that the source and destination have ONTAP 9.5 or later installed.
- The disk technology (FC or SATA, flash or rotating) underlying the source and destination volumes can vary, but the least performant disk technology determines the I/O characteristics.
- The type of system and the disk configuration on the destination system affect the performance of the source system. Therefore, the destination system should have the bandwidth required for the increased traffic and for message logging.
- The network transport should be optimized for SM-S replication. NetApp recommends a dedicated, high-bandwidth, low-latency network between the source and destination systems. Otherwise performance could be affected.

### 3.3 Maximum Number of SnapMirror Synchronous Replication Operations

The limit on the number of SM-S replication operations per node depends on the controller model. The system resources are shared between SnapMirror and SM-S replication operations. Therefore, the limit for a particular type of replication is reduced if any other types of replication operations are being performed concurrently.

Table 1) Concurrent replication operations allowed.

Platform	SnapMirror Synchronous operations per HA
AFF	80
FAS	40
ONTAP Select	20

**Note:** These limits apply to the surviving system, the failed-over system, or both. If a failover occurs during data backup, all of the concurrent transfers happening at the time of the failure on the failed node are aborted and are rescheduled by the partner node. Conversely, if a giveback occurs during data backup, all of the concurrent transfers happening at the time of the giveback on behalf of the partner node are aborted and are rescheduled by the partner node. 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.

```
Remote::> snapmirror initialize -source-path vs0:ssm_test_452985__245_RW1 -type P -policy
StrictSync -destination-path vs1:ssm_test_452991__249_DF1
Error: command failed: The number of endpoints in SnapMirror Synchronous relationships on node
sti2650-249 and its failover partner cannot exceed 40.
```

### 3.4 Restrictions Using SnapMirror Synchronous

When planning the configuration of SM-S replication, you need to consider the following restrictions:

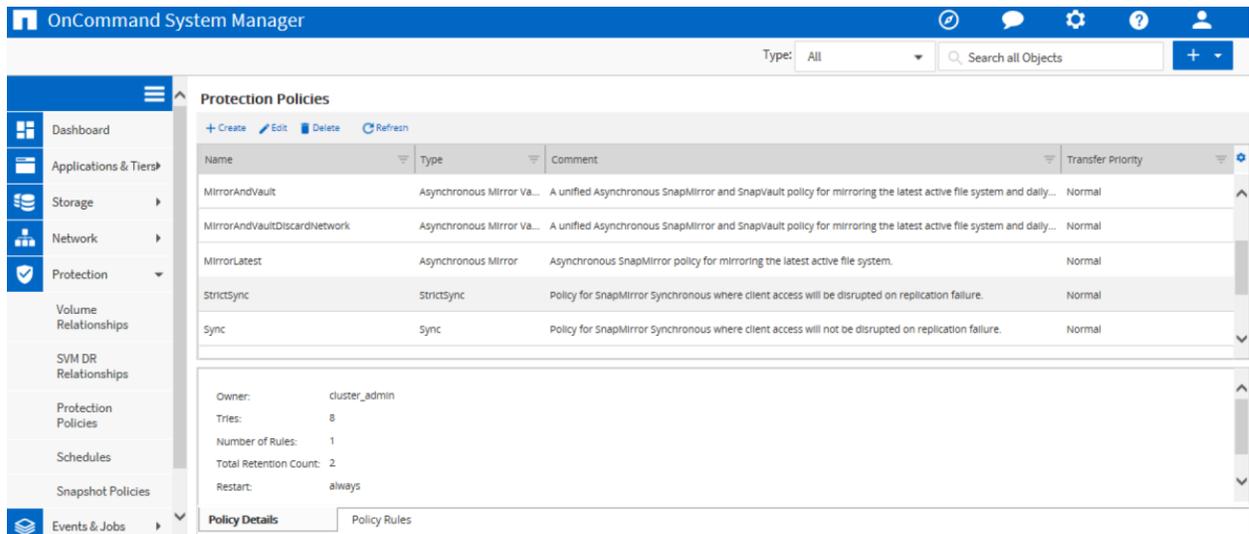
- The source volume must be online.
- The capacity of the destination volume must be greater than or equal to the capacity of the source volume.
- The SnapMirror source or destination volume cannot be the root volume of a storage system.
- There must be a functional network to transfer data between two different storage systems.

### 3.5 SnapMirror Policies

SM-S has two new predefined policies available in ONTAP System Manager and the CLI:

- Sync
- StrictSync

Figure 5) SnapMirror policies in ONTAP System Manager.



```
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   -   -
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   -   -
```

### 3.6 Common Snapshot Copies

Snapshot copies stored on both the source and destination volumes enable the rapid resynchronization of the volumes in the data protection relationship. Recovery from a replication failure is asynchronous and uses the last common snapshot copy as a starting point to avoid a re-baseline. When using advanced privileges in the CLI, SnapMirror policy has a new optional field, `-common-snapshot-schedule`, with a default hourly schedule that specifies when common Snapshot copies are created on source and destination volumes. The `-common-snapshot-schedule` can be modified from a minimum of 30 minutes to a maximum of 2 hours.

A shorter schedule uses more system resources and involves more latency but decreases the time it takes to perform a resync operation after a failure. A longer schedule uses fewer system resources and uses network latency more effectively but increases the time it takes perform a resync operation after a failure.

NetApp recommends that you evaluate and understand the effect of any change to the default schedule, taking these factors into consideration:

- A maximum of two common Snapshot copies are retained on the volumes.
- The `snapmirror update` command captures the latest state of AFS and creates an on-demand common Snapshot copy.

```
Primary::*> snapmirror policy show -vserver vs1 -policy Sync -fields -common-snapshot-schedule
vserver policy common-snapshot-schedule
-----
vs1          Sync          hourly
Primary::*> snapmirror policy show -vserver vs1 -policy StrictSync -fields -common-snapshot-
schedule
vserver policy common-snapshot-schedule
-----
vs1          StrictSync      hourly
```

**Note:** The SnapMirror policies support only one rule for common Snapshot copy creation based on the SnapMirror label `sm_created`.

```
Primary::*> snapmirror policy modify -vserver vs1 -policy Sync -common-snapshot-schedule ?
(snapmirror policy modify)
<text>                               Common Snapshot Copy Creation Schedule for SnapMirror Synchronous
```

### 3.7 Relationship Status

Table 2 describes the different relationship states that can occur between the source and destination volumes.

Table 2) Relationship status.

Relationship Status	Details
Idle	<ul style="list-style-type: none"> <li>• No transfer is in progress.</li> <li>• If health is true, this is a newly created relationship.</li> <li>• If health 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 transfer phase is in progress when a synchronous relationship is being established or reestablished.

Relationship Status	Details
Transitioning	<ul style="list-style-type: none"> <li>Transition 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	<ul style="list-style-type: none"> <li>Destination volume is not in sync with the source volume because SnapMirror replication is not happening. If the mirror state is Snapmirrored, this indicates a transfer failure or failure due to an unsupported operation.</li> <li>SM-S monitors the relationship every 5 minutes and attempts an autoresync operation.</li> </ul> <p><b>Note:</b> SM-S attempts a maximum of five back-to-back resync operations.</p>

### 3.8 Cluster and SVM Peering

You must perform the one-time operation of creating cluster peer and SVM peer relationships in order to facilitate SM-S relationships between the two clusters. Also, both the source and destination SVMs must have the same language type setting to successfully replicate between them. Intercluster SnapMirror relationships are primarily used to provide disaster recovery capability in another site or location. If you have replicated all necessary volumes to a disaster recovery site with SnapMirror, then you can perform a recovery to restore operations from the disaster recovery site.

A peer relationship can be created in one of two ways:

- A peer relationship can be created by a cluster administrator who has security credentials (a cluster admin login and password) for the other cluster.
- Suppose that two administrators do not want to exchange cluster admin passwords to peer their clusters. In this method, each administrator enters the `cluster peer create` command specifying the intercluster IP addresses of the other cluster.

To view cluster details, complete the following steps:

- View the source cluster details.

```
Primary::> cluster show
Node                Health  Eligibility
-----
PrimaryNode1       true   true
PrimaryNode2       true   true
2 entries were displayed.
```

- View the destination cluster details.

```
Remote::> cluster show
Node                Health  Eligibility
-----
RemoteNode1        true   true
RemoteNode2        true   true
2 entries were displayed.
```

- View the cluster and SVM peer status.

```
Remote::> cluster peer show
Peer Cluster Name  Cluster Serial Number  Availability  Authentication
-----
Primary           1-80-000011           Available    ok
```

**Note:** In the Availability field, the cluster becomes Unavailable if the cluster is not reachable or if there is network congestion between the source and destination clusters.

```
Remote::> vserver peer show
Peer Peer Peering Remote
Vserver Vserver State Peer Cluster Applications Vserver
-----
vs1 vs1_sync_dr peered Primary snapmirror vs1_sync_dr
```

4. Calculate the round-trip time (RTT) for SM-S using the `cluster peer ping` command. The originating node hosts the source volume of the relationship and the destination node hosts the destination volume of the relationship.

```
Primary::> cluster peer ping -originating-node PrimaryNode1 -destination-node RemoteNode1
Destination Cluster: Remote

Destination Node IP Address Count TTL RTT(ms) Status
-----
RemoteNode1 172.26.145.10 1 64 0.452 interface_reachable
RemoteNode1 172.26.145.12 1 64 0.352 interface_reachable
```

For intracluster relationships, you can run the `network ping` command in a similar way.

```
Primary::> network ping -node RemoteNode1 -destination PrimaryNode1 -show-detail
```

### 3.9 Populate Source Volume with Data

To populate the source volume with data, complete the following steps:

1. Use the `volume create` command to create a volume on the primary cluster.

```
Primary::> vol create -vserver vs1 -volume vol1 -aggregate aggr1_Primary_01 -size 10GB -state
online -policy default -junction-path /dp_src -unix-permissions ---rwxr-xr-x -type RW
[Job 81] Job succeeded: Successful
```

#### Best Practice

- The source volume can reside in a FabricPool aggregate with any tiering policy, including None, Snapshot, or Auto. However, the destination volume can be in a FabricPool aggregate only if the tiering policy is not of type Backup.

```
Primary::> vol create -volume vol_src -vserver vs0 -tiering-policy auto -aggregate aggr5 -size 1g
-tiering-minimum-cooling-days 2 -type RW
```

2. Populate the source volume with data.

```
[root@client~]# mount | grep 10.236.130.54
10.236.130.54:/vol1 on /root/mount type nfs
(rw,relatime,vers=3,rsize=65536,wsz=65536,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys
,mountaddr=10.236.130.54,mountvers=3,mountport=635,mountproto=udp,local_lock=none,addr=10.236.130
.54)

[root@client mount]# df -ha /root/mount/
Filesystem Size Used Avail Use% Mounted on
10.236.130.54:/vol1 973M 95M 879M 10% /root/mount
```

3. Create a data protection volume on the remote cluster. The destination volume should be the same as or greater in size than the source volume and of `-type dp`.

```
Remote:> vol create -vserver vs1_sync_dr -volume vol1 -aggregate aggr1_Remote_01
-size 10GB -type DP
[Job 81] Job succeeded: Successful
```

#### 4. The destination volume is in a Fabric Pool aggregate.

```
Remote:> vol create -volume vol_dst -vserver vs1 -tiering-policy auto -aggregate aggr5 -size 1g -
tiering-minimum-cooling-days 2 -type DP
```

**Note:** Creation of the SM-S relationship creates a destination volume of type `dp`, so this step is not mandatory.

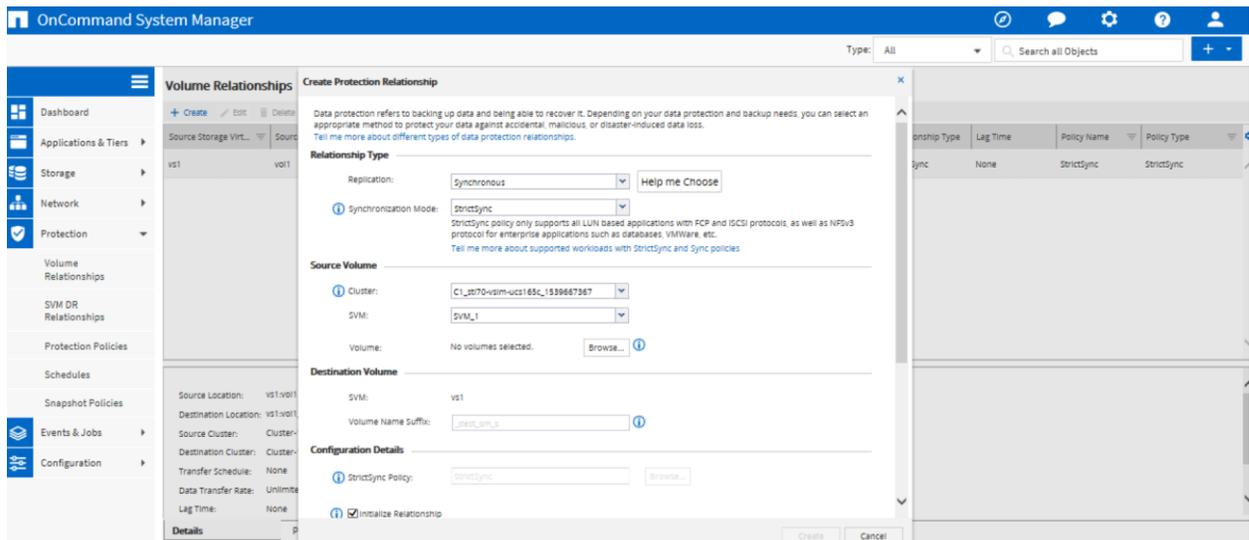
### 3.10 Set Up a SnapMirror Synchronous Relationship

To set up the SM-S relationship, complete the following steps:

1. Create a SnapMirror Synchronous relationship between volumes on the primary and remote clusters.

**Note:** Creating an SM-S relationship between two volumes validates that the destination volume is not already the source for another SnapMirror relationship. During this time, client access to the source volume is disrupted and the destination volume is read-only.

**Note:** Fan-out or cascade from an SM-S destination is not supported.



2. Click Help Me Choose.

## Help Me Choose

What is your expected recovery point objective? **Zero** Minutes - Hours Hours  
Synchronous mirror, copying data without a delay

Do you want disaster recovery?

Do you want backups?

If it is not possible to keep the mirror synchronized, what are your priorities? **Maintain Protection** Continues I/Os  
Fail primary I/O to maintain strict synchronization

**Protection Result**

- Synchronous
- Mirror
- 
- StrictSync

Cancel Apply

A SnapMirror Synchronous relationship can also be created by using the `snapmirror create` command and specifying a suitable SnapMirror policy such as `Sync` or `StrictSync`. This creates a configuration entry for a relationship. After you have created it, the `snapmirror show` command can be used to observe the SMS-S relationship.

```
Remote::> 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.
```

**Note:** You can set up a SnapMirror Synchronous relationship between the two nodes of an HA pair.

3. Initialize the SnapMirror relationship. The first time you do this, there is a baseline transfer from the source to the destination volume. The SnapMirror policy for the relationship defines the content of the baseline and any updates. Initializing the destination triggers a set of checks for interoperability and limits, and then the following phases occur during the baseline transfer:
  - a. In the initial asynchronous transfer phase, back-to-back asynchronous transfers occur, updating the destination volume with data from the source volume. During this time, the replication status is shown as `Transferring`, and the destination volume is marked as `invalid` in the output of the `vol status` command.
  - b. In the transition phase, the last asynchronous transfer of all incoming I/O occurs before switching to synchronous replication with the relationship status of `Transitioning`. After this phase, the destination volume is `online` in a `read-only` state. After synchronous replication kicks in, the active I/O is written to both source and destination volumes.

You should consider the following issues:

- The destination volume should already be created with the same language as the source volume.
- SnapMirror monitors the SM-S relationship every 5 minutes and attempts `autoresync` if the relationship is `OutOfSync`. A maximum of five back-to-back `autoresync` attempts occur to try to bring the relationship to `InSync`.

```
Remote::> snapmirror initialize -destination-path vs1_sync_dr:vol1 Operation is queued:
snapmirror initialize of destination vs1_sync_dr:vol1.
```

```
Remote::> snapmirror status
Snapmirror is on.
Source          Destination          State          Status
-----
Primary:vol1    vs1_sync_dr:vol1    Snapmirrored  Idle
```

- You should not delete Snapshot copies that SnapMirror creates in the source volume. The most recent SnapMirror Snapshot copy is referred to as the newest common Snapshot copy (NCS). The resync operation depends on the NCS. If SnapMirror cannot find the required Snapshot copy on the source, it cannot perform the resync operation to the destination.
  - You should not restrict or take the destination volume offline while SnapMirror is configured to transfer. Taking the destination offline prevents SnapMirror from making changes to the destination. Also, any attempt to bring this volume online manually will succeed only after the initial transfer is complete.
4. Monitor the SnapMirror Synchronous relationship. You can monitor the data transfer status of the SM-S relationship on the source or destination by using the following commands, which provide details of any possible errors or unhealthy states.

Source	Destination
snapmirror list-destination	snapmirror show
snapmirror list-destination -instance	snapmirror show -instance
	snapmirror show -history

```
Remote::~*> snapmirror show -fields last-transfer-error, policy,policy-type,status
source-path destination-path policy-type          policy erroer status last-transfer-error
-----
vs1:vol1     vs1_sync_dr:vol1 sync-mirror          Sync          InSync -
vs1:volX     vs1:volY         strict-sync-mirror  StrictSync InSync -
2 entries were displayed.
Primary::~*> snapmirror list-destinations
Source          Destination          Transfer    Last          Relationship
Path           Type    Path           Status Progress  Updated      Id
-----
vs1:vol1       XDP     vs1_sync_dr:vol1 InSync -        -            603f951a-cf83-11e8-9a21-
0050568edde7
vs1:volX       XDP     vs1:volY         InSync -        -            078c2397-d046-11e8-
9a21-0050568edde72 entries were displayed. Primary::~*> snapmirror list-destinations -destination-
path vs2:volY -fields out-of-sync-reasonsource-path destination-path out-of-sync-reason
-----
vs2:volX       vs2:volY          Transfer failed.(Replication engine error(No space left on
device))
```

The Status field reflects the status of the data transfer:

- InSync indicates that the secondary copy is in sync with the primary and that the replication path is active.
- OutOfSync indicates that the replication from primary to secondary volume is not occurring.
- Transitioning indicates the transition from asynchronous to synchronous transfer.

### 3.11 Set Up a SnapMirror Strict Synchronous Relationship

To set up a SnapMirror StrictSync relationship, complete the following steps:

1. Create a volume on the primary cluster.

```
Primary::> vol create -vserver vs1 -volume voll -aggregate aggr1_Primary_01 -size 10GB (volume create)
```

```
[Job 81] Job succeeded: Successful
```

## 2. Create a data protection volume on the remote cluster.

```
Remote::> vol create -vserver vs1_sync_dr -volume voll -aggregate aggr1_Remote_01 -size 10GB -type DP (volume create)
```

```
[Job 81] Job succeeded: Successful
```

## 3. Create a SnapMirror Strict Synchronous relationship between volumes on the primary and remote clusters.

```
Remote::> snapmirror create -source-path vs1:voll -destination-path vs1_sync_dr:voll -policy StrictSync
```

```
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:voll.
```

**Note:** Verify that the destination volume is not already a source of another SnapMirror relationship.

## 4. Initialize the SnapMirror relationship.

```
Remote::> snapmirror initialize -destination-path vs1_sync_dr:voll Operation is queued: snapmirror initialize of destination vs1_sync_dr:voll.
```

**Note:** If the I/O to the secondary storage does not complete for any reason, the application I/O fails (that is, zero data loss) and synchronous replication is terminated in StrictSync mode. This keeps the primary and secondary volumes identical. You can check whether or not a volume is fenced for protocol access and who the owner of this fence is.

```
Primary:*> volume show -fields is-protocol-access-fenced, protocol-access-fenced-by -volume voll vserver volume is-protocol-access-fenced protocol-access-fenced-by
```

```
-----  
vs1 voll true snapmirror_synchronous_strict_sync
```

```
Primary:*> snapmirror update -destination-path vs1_sync_dr:voll
```

```
Operation is queued: snapmirror update of destination " vs1_sync_dr:voll".
```

## 5. To check for the exported snapshot copy, check the following field in `snapmirror show`:

```
Remote::> snapmirror show -fields exported-snapshot source-path destination-path exported-snapshot
```

```
-----  
vs1:voll vs1_sync_dr:voll snapmirror.2e09b182-b4c3-11e8-80c2-005056acalle_2157376758.2018-09-10_034839
```

## 6. You can choose to assign a higher priority for Synchronous replication traffic using the Differentiated Services Code Point (DSCP) packet header value.

```
Primary:*> network qos-marking modify -ip-space Default -protocol SnapMirror-Sync -dscp 10 -is-enabled true
```

```
Primary:*> network qos-marking show
```

```
(A DSCP value of 10 indicates a relatively high priority)
```

## 3.12 View Destination Volume on the Secondary Storage System

When viewing a read-only SM-S destination volume (type 'dp'), changes made to the source volume are not always reflected immediately on the destination volume and might only appear later. However, the data is still there. This situation occurs when a SnapMirror Synchronous (SM-S) destination volume is redirected to the latest exported Snapshot copy, which is updated only once an hour by default. How often a SnapMirror Synchronous relationship updates the exported Snapshot copy is controlled by the SnapMirror policy's `-common-snapshot-schedule` parameter. The `-common-snapshot-schedule` parameter can be modified to a minimum of 30 minutes and a maximum of 2 hours. Note that changing this parameter can affect controller performance.

You can modify the `-common-snapshot-schedule` parameter with the `snapmirror policy modify` command in advanced privilege.

```
Remote::> set advanced
Remote::*> snapmirror policy modify -vserver <dest_svm> -policy <sync_policy> -common-snapshot-schedule <schedule>
```

You can also update the exported Snapshot copy on-demand by issuing the `snapmirror update` command against the SnapMirror Synchronous relationship:

```
Remote::> snapmirror update -destination-path <destination_svm>:<dest_volume>
```

### 3.13 Interrupt Network Between Primary and Remote (Prevents Replication) and Observe the Behavior of the SnapMirror Strict Synchronous Relationship

#### 1. Bring down the intercluster LIFs.

```
Primary::> net int show -role intercluster
(network interface show)
Logical Status Network Current Current Is
Vserver Interface Admin/Oper Address/Mask Node Port Home
-----
Primary
PrimaryNode1_inet4_intercluster1 up/up 10.236.130.76/20 PrimaryNode1 e0d true
PrimaryNode2_inet4_intercluster1 up/up 10.236.130.84/20 PrimaryNode2 e0d true
2 entries were displayed.
Primary::> net int modify -vserver vs1 -lif PrimaryNode1_inet4_intercluster1 -status-admin down
(network interface modify)
Primary::> net int modify -vserver vs1 -lif PrimaryNode2_inet4_intercluster1 -status-admin down
(network interface modify)
Primary::> net int show -role intercluster
(network interface show)
Logical Status Network Current Current Is
Vserver Interface Admin/Oper Address/Mask Node Port Home
-----
Primary
PrimaryNode1_inet4_intercluster1 down/down 10.236.130.76/20 PrimaryNode1 e0d true
PrimaryNode2_inet4_intercluster1 down/down 10.236.130.84/20 PrimaryNode2 e0d true
2 entries were displayed.
```

```
Primary::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability Authentication
-----
Primary 1-80-000011 Unavailable ok
```

#### 2. Observe the fence set by SnapMirror Strict Synchronous.

```
Remote::*> vol show -vserver vs1 -volume voll -fields protocol-access-fenced-by,is-protocol-access-fenced
vserver volume is-protocol-access-fenced protocol-access-fenced-by
-----
vs1 voll true snapmirror_synchronous_strict_sync
```

#### 3. Observe the application behavior, try to modify the source, and verify that the operation fails.

```
[root@client ~]# cd /root/mount
-bash: cd: /root/mount: Stale file handle
[root@client ~]#

[root@client ~]# echo Test_strict_sync_application_behavior_when_replication_fails >
/root/mount/failure.foo
-bash: /root/mount/failure.foo: Remote I/O error
```

**Note:** In Strict Synchronous mode when replication cannot be performed, client application read or modification of a source might fail to provide zero data loss for a zero RPO.

### 3.14 Interrupt Network Between Source and Destination (Prevents Replication) and Observe the Behavior of SnapMirror Synchronous Mode

1. To interrupt the network between the source and destination, bring down the intercluster LIFs.

```
Primary::> net int show -role intercluster
(network interface show)
Logical          Status Network          Current Current Is
Vserver         Interface      Admin/Oper Address/Mask Node Port Home
-----
Primary
PrimaryNode1_inet4_intercluster1 up/up 10.236.130.76/20 PrimaryNode1 e0d true
PrimaryNode2_inet4_intercluster1 up/up 10.236.130.84/20 PrimaryNode2 e0d true
2 entries were displayed.
Primary::> net int modify -vserver vs1 -lif PrimaryNode1_inet4_intercluster1 -status-admin down
(network interface modify)
Primary::> net int modify -vserver vs1 -lif PrimaryNode2_inet4_intercluster1 -status-admin down
(network interface modify)
```

```
Primary::> net int show -role intercluster
(network interface show)
Logical Status Network Current Current Is
Vserver Interface Admin/Oper Address/Mask Node Port Home
-----
Primary
PrimaryNode1_inet4_intercluster1 down/down 10.236.130.76/20 PrimaryNode1 e0d true
PrimaryNode2_inet4_intercluster1 down/down 10.236.130.84/20 PrimaryNode2 e0d true
2 entries were displayed.
Primary::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability Authentication
-----
Primary                               1-80-000011 Unavailable ok
```

```
Primary::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_sync_dr:vol1 Snapmirrored OutofSync 0B false 09/11 02:23:05
```

2. Observe the fence set by SnapMirror Synchronous.

```
Remote::~*> vol show -vserver vs1 -volume vol1 -fields protocol-access-fenced-by,is-protocol-
access-fenced
vserver volume is-protocol-access-fenced protocol-access-fenced-by
-----
vs1          vol1         false                               snapmirror_synchronous_sync
```

3. Observe application behavior, try to modify the source, and verify that the operation succeeds.

```
[root@client ~]# echo sync_behavior_during_network_outage > /root/mount/realxed_outage.foo
[root@client ~]# more /root/mount/realxed_outage.foo
sync_behavior_during_network_outage
```

**Note:** Because SM-S is configured in Synchronous mode, client application modification of the source does not fail, but zero RPO cannot be met because this data wasn't replicated to the remote storage system.

### 3.15 Failover from Primary to Secondary Data Center

When disaster disables the source of a SnapMirror relationship, you must permanently end a SnapMirror relationship between a source and destination volumes.

To failover the application, follow these steps.

1. Run the `snapmirror quiesce` command.

The quiescence of a synchronous relationship aborts or stops ongoing transfers to secondary until a resume command is issued. To avoid client I/O disruption, a StrictSync relationship switches to Sync mode with a relationship status that is marked Unhealthy until a resume command is issued. A confirmation message warns about zero RPO compromise. If the relationship is InSync, the quiesce operation creates a common Snapshot copy to reduce the effect on RPO during the resume operation.

**Note:** The SnapMirror Synchronous abort command might lead to primary I/O disruption; you are prompted with the following warning message.

```
Warning: It is recommended to quiesce the relationship using the "snapmirror quiesce" command instead of aborting the SnapMirror Synchronous transfer. For relationships with a policy of type "strict-sync-mirror" this will fail client I/O on source volume if the status is InSync. Do you want to continue? {y|n}:
```

```
Remote::> snapmirror quiesce -destination-path vs1_sync_dr:vol1
Operation succeeded: snapmirror quiesce the relationship with destination vs1_sync_dr:vol1.
```

2. Create a SnapMirror break.

Break terminates the synchronous protection on the volume and converts the data protection destination volume to read/write, with the mirror state showing as Broken-Off and the relationship status as OutOfSync at the source and Idle at the destination. A relationship that is Broken-Off does not autoresync.

```
Remote::> snapmirror break -destination-path vs1_sync_dr:vol1
```

**Best Practice**

- Quiesce the SnapMirror relationship before breaking it.

**Note:** Break is not allowed on an Uninitialized or Broken-Off relationship, and a force option is not supported.

**Note:** Recovery in a three-site datacenter scenario includes SnapMirror quiesce and break-off:

- An asynchronous data volume at the far DR site.
- A synchronous log volume at the near DR site.

3. Use external third-party tools to replay the logs from the near DR site to the far DR site.

### 3.16 Failover and Failback

To perform failover and failback, follow these steps.

1. Run the `snapmirror delete` command.

```
Remote::> snapmirror delete -destination-path vs1_sync_dr:vol1
Operation succeeded: snapmirror delete the relationship with destination vs1_sync_dr:vol1.
```

**Best Practice**

- Quiesce the SnapMirror relationship before deleting it.

2. Run the `snapmirror release` command.

The `snapmirror release` command deletes the relationship information from the source and performs a resource cleanup. For StrictSync relationships, release first unfences the primary volume

so that it does not cause I/O disruption on the primary. A common Snapshot copy must be retained. Therefore, you must use the `-relationship-info-only` parameter with the `snapmirror release` command; otherwise it is not possible to reestablish a sync relationship between the two copies.

```
Primary::*> snapmirror release -relationship-info-only true -destination-path vs1_sync_dr:vol1
[Job 51] Job succeeded: SnapMirror Release Succeeded\
```

### Best Practice

- Verify that a SnapMirror transfer is not in progress when you attempt a release.
- Quiesce and delete the SnapMirror relationship before releasing it.

### 3. Create a reverse relationship.

```
Primary::> snapmirror create -source-path vs1_sync_dr:vol1 -destination-path vs1:vol1 -policy StrictSync
```

```
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:vol1.
```

### 4. Run the `snapmirror resync` command.

To perform a resync operation by using the `snapmirror resync` command, a common Snapshot copy must exist between the source and destination volumes. The resync operation creates a new SnapMirror Snapshot copy containing the changes that were made since the common Snapshot copy, and it replicates only those changes between the volumes. `snapmirror resync` is similar to `snapmirror initialize`, but without the baseline transfer. Each transfer also deletes the Snapshot copy taken by the previous transfer as part of cleanup.

```
Primary::> snapmirror resync -destination-path vs1:vol1
Warning: All data newer than Snapshot copy snapmirror.3fd9730b-8192-11e2-9caa-123478563412_2147484699.2013-02-28_1 10732 on volume vs1:vol1 will be deleted.

Verify there is no XDP relationship whose source volume is
" vs1:vol1". If such a relationship exists then you are creating an unsupported XDP to XDP
cascade.

Do you want to continue? {y|n}: y
[Job 133] Job succeeded: SnapMirror Resync Transfer Queued
Primary::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_sync_dr:vol1 Snapmirrored Finalizing 2.30KB false 09/11 02:03:22
Primary::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_sync_dr:vol1 Snapmirrored InSync - true -
```

**Note:** Autoresync is driven from the destination and is attempted every 5 minutes for five tries to bring back the relationship to InSync. If the failure persists, the relationship state transfers to Idle, which then requires manual intervention because the relationship is marked as Unhealthy.

### 5. Check whether the fence from SnapMirror is deactivated.

```
Remote::*> vol show -vserver vs1 -volume vol1 -fields is-protocol-access-fenced
vserver volume is-protocol-access-fenced
-----
```

```
vs1      voll  false
```

### 3.17 Convert a SnapMirror Synchronous Relationship from Strict Synchronous Mode to Synchronous

You cannot convert a SnapMirror relationship from Strict Synchronous to Synchronous by simply modifying the policy from StrictSync to Sync (and vice versa). You must quiesce, delete, and release (with the parameter `-relationship-info-only`) the existing SnapMirror relationship, then create a new SnapMirror relationship with the intended policy type and resynchronize.

**Note:** Using the `snapmirror modify` command, you can modify SnapMirror relationship attributes. However, conversion of the policy type on the relationship is not supported.

To convert a SnapMirror relationship from Strict Synchronous to Synchronous, follow these steps.

1. Run the `snapmirror quiesce` command.

```
Remote:> snapmirror quiesce -destination-path vs1_sync_dr:voll
Operation succeeded: snapmirror quiesce the relationship with destination vs1_sync_dr:voll.
Remote:> snapmirror show
Progress
Source  Destination      Mirror      Relationship Total      Last
Path   Type Path           State       Status      Progress Healthy Updated
-----
vs1:voll XDP vs1_sync_dr:voll Snapmirrored Quiesced      -      false -
```

2. Run the `snapmirror delete` command.

```
Remote:> snapmirror delete -destination-path vs1_sync_dr:voll
Operation succeeded: snapmirror delete the relationship with destination vs1_sync_dr:voll.
```

3. Run the `snapmirror release` command.

```
Remote:*> snapmirror release -relationship-info-only true -destination-path vs1_dr:voll
[Job 52] Job succeeded: SnapMirror Release Succeeded
Remote:*> snapmirror list-destinations
This table is currently empty.
```

4. Create a new SnapMirror relationship for Synchronous mode.

```
Remote:> snapmirror create -source-path vs1:voll -destination-path vs1_dr:voll -policy Sync
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:voll.
Remote:> snapmirror show
Progress
Source  Destination      Mirror      Relationship Total      Last
Path   Type Path           State       Status      Progress Healthy Updated
-----
vs1:voll XDP vs1_sync_dr:voll Snapmirrored Idle      -      false -

Remote:> snapmirror show -fields policy-type,policy
source-path destination-path policy-type      policy
-----
vs1:voll    vs1_sync_dr:voll sync-mirror      Sync
```

5. Run the `snapmirror resync` command.

```
Remote:> snapmirror resync -destination-path vs1_sync_dr:voll

Warning:      All data newer than Snapshot copy snapmirror.3fd9730b-8192-11e2-9caa-
123478563412_2147484699.2013-02-28_1 10732 on volume vs1_sync_dr:voll will be deleted.

Verify there is no XDP relationship whose source volume is
"vs1_sync_dr:voll ". If such a relationship exists then you are creating an unsupported XDP to
XDP cascade.

Do you want to continue? {y|n}: y
[Job 133] Job succeeded: SnapMirror Resync Transfer Queued
```

```

Remote::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_sync_dr:vol1 Snapmirrored Transferring 2.30KB false 09/11 02:18:43
Remote::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_sync_dr:vol1 Snapmirrored InSync - true -

```

### 3.18 Convert an Asynchronous SnapMirror Relationship to Synchronous

Most customers in industries like finance and healthcare cannot tolerate a single point of failure in their systems that could cause data loss.

For example, consider the use case of Site A synchronously replicating to Site B and Site B asynchronously replicating to Site C. Site B must replicate synchronously to Site C only when Site A fails over to Site B. In this case, you have the option to convert a SnapMirror relationship between Site B to Site C from asynchronous to synchronous as follows:

1. Create an asynchronous SnapMirror relationship and perform a baseline transfer.

```

Primary::> snapmirror create -source-path vs1:vol1 -destination-path vs1_dr:vol1 -type XDP -
policy MirrorAllSnapshots

```

Operation succeeded: snapmirror create the relationship with destination vs1\_sync\_dr:vol1.

2. View the progress of the asynchronous SnapMirror relationship.

```

Primary::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_dr:vol1 Uninitialized Idle - true -

Primary::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_dr:vol1 Snapmirrored Finalizing 0B true 09/10 03:48:41

Primary::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_dr:vol1 Snapmirrored Idle - true -

Primary::> snapmirror show -instance
Source Path: vs1:vol1
Destination Path: vs1_dr:vol1
Relationship Type: XDP
Relationship Group Type: none
SnapMirror Schedule: -
SnapMirror Policy Type: async-mirror
SnapMirror Policy: MirrorAllSnapshots
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.2e09b182-b4c3-11e8-80c2-005056aca11e_2157376758.2018-09-10_034839
Newest Snapshot Timestamp: 09/10 03:48:39
Exported Snapshot: snapmirror.2e09b182-b4c3-11e8-80c2-005056aca11e_2157376758.2018-09-10_034839
Exported Snapshot Timestamp: 09/10 03:48:39
Healthy: true
Unhealthy Reason: -
Destination Volume Node: Primary
Relationship ID: d512efeb-b4cd-11e8-886b-005056ac1de0
Current Operation ID: -
Transfer Type: -
Transfer Error: -
Current Throttle: -
Current Transfer Priority: -
Last Transfer Type: update
Last Transfer Error: -
Last Transfer Size: 94.03MB
Last Transfer Network Compression Ratio: 1:1
Last Transfer Duration: 0:0:15
Last Transfer From: vs1:vol1
Last Transfer End Timestamp: 09/10 03:48:56
Progress Last Updated: -
Relationship Capability: 8.2 and above
Lag Time: 0:0:48
Identity Preserve Vserver DR: -
Volume MSIDs Preserved: -
Is Auto Expand Enabled: -
Number of Successful Updates: 1
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: 98610000
Total Transfer Time in Seconds: 17

```

### 3. Run the `snapmirror delete` command to delete the asynchronous SnapMirror relationship.

```

Primary::> snapmirror delete -destination-path vs1_dr:vol1
Operation succeeded: snapmirror delete the relationship with destination vs1_dr:vol1.

```

### 4. Run the `snapmirror release` command.

```

Remote::> snapmirror release -relationship-info-only true -destination-path vs1_dr:vol1
Operation succeeded: snapmirror release the relationship with destination vs1_dr:vol1.

```

### 5. Create a new SnapMirror Strict Synchronous relationship.

```

Primary::> snapmirror create -source-path vs1:vol1 -destination-path vs1_dr:vol1 -policy
StrictSync
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:vol1.

```

```

Primary::> snapmirror show
Progress
Source      Destination  Mirror      Relationship  Total   Last
Path        Type Path      State        Status      Progress Healthy Updated
-----
vs1:vol1    XDP vs1_dr:vol1 Snapmirrored Idle        -        true   -

```

### 6. Run the `snapmirror resync` command.

```

Primary::> snapmirror resync -destination-path vs1_dr:vol1
Warning: All data newer than Snapshot copy snapmirror.3fd9730b-8192-11e2-9caa-
123478563412_2147484699.2013-02-28_1 10732 on volume vs1_dr:vol1 will be deleted.

```

Verify there is no XDP relationship whose source volume is

" vs1\_dr:vol1". If such a relationship exists then you are creating an unsupported XDP to XDP cascade.

```

Do you want to continue? {y|n}: y
[Job 133] Job succeeded: SnapMirror Resync Transfer Queued
Primary::> snapmirror show
Progress
Source          Destination  Mirror      Relationship  Total      Last
Path            Type Path      State        Status       Progress Healthy Updated
-----
vs1:vol1        XDP vs1_dr:vol1 Snapmirrored Transferring  0B         true   09/10 04:29:45
Primary::> snapmirror show
Progress
Source          Destination  Mirror      Relationship  Total      Last
Path            Type Path      State        Status       Progress Healthy Updated
-----
vs1:vol1        XDP vs1_dr:vol1 Snapmirrored InSync      -         true   -
Primary::> snapmirror show -destination-path vs1_dr:vol1
Source Path: vs1:vol1
Destination Path: vs1_dr:vol1
Relationship Type: XDP
Relationship Group Type: none
SnapMirror Schedule: -
SnapMirror Policy Type: strict-sync-mirror
SnapMirror Policy: StrictSync
Tries Limit: -
Throttle (KB/sec): unlimited
Mirror State: Snapmirrored
Relationship Status: InSync
File Restore File Count: -
File Restore File List: -
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: snapmirror.2e09b182-b4c3-11e8-80c2-005056aca11e_2157376758.2018-09-10_060505
Newest Snapshot Timestamp: 09/10 06:05:05
Exported Snapshot: snapmirror.2e09b182-b4c3-11e8-80c2-005056aca11e_2157376758.2018-09-10_060505
Exported Snapshot Timestamp: 09/10 06:05:05
Healthy: true
Unhealthy Reason: -
Destination Volume Node: PrimaryNode1
Relationship ID: 96486a06-b4d3-11e8-886b-005056ac1de0
Current Operation ID: -
Transfer Type: -
Transfer Error: -
Current Throttle: -
Current Transfer Priority: -
Last Transfer Type: resync
Last Transfer Error: -
Last Transfer Size: 845.8KB
Last Transfer Network Compression Ratio: 1:1
Last Transfer Duration: 0:1:17
Last Transfer From: vs1:vol1
Last Transfer End Timestamp: 09/10 04:31:02
Progress Last Updated: -
Relationship Capability: 8.2 and above
Lag Time: 0:0:0
Identity Preserve Vserver DR: -
Volume MSIDs Preserved: -
Is Auto Expand Enabled: -
Number of Successful Updates: 3
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: 868498
Total Transfer Time in Seconds: 77

```

**Note:** You cannot convert an asynchronous SnapMirror relationship to synchronous (and vice versa) by simply modifying the SnapMirror policy. You must delete and release the existing

SnapMirror relationship with the parameter `-relationship-info-only` and then create a new relationship with the intended policy and resynchronize.

**Note:** A new baseline transfer is not required if common Snapshot copies exist. Also, you have the option to revert back to asynchronous mode if necessary.

### 3.19 Convert a Synchronous SnapMirror Relationship to Asynchronous

To convert an SM-S relationship to asynchronous, follow these steps.

1. Run the `snapmirror quiesce` command.

```
Remote::> snapmirror quiesce -destination-path vs1_dr:vol1
Operation succeeded: snapmirror quiesce the relationship with destination vs1_dr:vol1.
Remote::> snapmirror show
Progress           Mirror Relationship Total           Last
Source Destination
Path  Type Path           State  Status Progress Healthy Updated
-----
vs1:vol1 XDP vs1_dr:vol1 Snapmirrored Quiesced - false -
```

2. Run the `snapmirror delete` command.

```
Remote::> snapmirror delete -destination-path vs1_dr:vol1
Operation succeeded: snapmirror delete the relationship with destination vs1_dr:vol1.
```

3. Run the `snapmirror release` command.

```
Primary::> snapmirror release -relationship-info-only true -destination-path vs1_dr:vol1
Operation succeeded: snapmirror release the relationship with destination vs1_dr:vol1.
```

**Note:** The default release operation (without the parameter `relationship-info-only` set to true) deletes the Snapshot copies created by this relationship. This option does not allow a resync operation for a new relationship created for these volumes.

4. Run the `snapmirror create` command to create a new asynchronous SnapMirror relationship.

```
Remote::> snapmirror create -source-path vs1:vol1 -destination-path vs1_dr:vol1 -policy
MirrorAllSnapshots
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:vol1.
```

5. Run the `snapmirror resync` command to resync the newly created asynchronous relationship.

```
Remote::> snapmirror resync -destination-path vs1_dr:vol1

Warning:           All data newer than Snapshot copy snapmirror.3fd9730b-8192-11e2-9caa-
123478563412_2147484699.2013-02-28_1 10732 on volume vs1_dr:vol1 will be deleted.

Verify there is no XDP relationship whose source volume is
" vs1_dr:vol1". If such a relationship exists then you are creating an unsupported XDP to XDP
cascade.

Do you want to continue? {y|n}: y
[Job 133] Job succeeded: SnapMirror Resync Transfer Queued
```

### 3.20 Flip Resync

To establish a relationship in the reverse direction, complete the following steps:

1. Before deletion, quiesce the existing relationship.

```

Remote::*> snapmirror quiesce -destination-path cnodes1_vs1:dstvol_1
Operation succeeded: snapmirror quiesce for destination "cnodes1_vs1:dstvol_1".

Remote::*> snapmirror show

Source          Destination Mirror Relationship Total          Progress
Path           Type   Path      State   Status   Progress  Healthy  Last Updated
-----
cnodes0_vs1:srcvol_1 XDP  cnodes1_vs1:dstvol_1 Snapmirrored Quiesced - false -
cnodes0_vs1:srcvol_2 XDP  cnodes1_vs1:dstvol_2 Snapmirrored InSync - true -
2 entries were displayed.

Remote::*> snapmirror quiesce -destination-path cnodes1_vs1:dstvol_2
Operation succeeded: snapmirror quiesce for destination "cnodes1_vs1:dstvol_2".

Remote::*> snapmirror show

Source          Destination Mirror Relationship Total          Progress
Path           Type   Path      State   Status   Progress  Healthy  Last Updated
-----
cnodes0_vs1:srcvol_1 XDP  cnodes1_vs1:dstvol_1 Snapmirrored Quiesced - false -
cnodes0_vs1:srcvol_2 XDP  cnodes1_vs1:dstvol_2 Snapmirrored Quiesced - false -
2 entries were displayed.

```

## 2. Delete the existing relationship.

```

Remote::*> snapmirror delete -destination-path cnodes1_vs1:dstvol_1
Operation succeeded: snapmirror delete for the relationship with destination "cnodes1_vs1:dstvol_1".

Remote::*> snapmirror delete -destination-path cnodes1_vs1:dstvol_2
Operation succeeded: snapmirror delete for the relationship with destination "cnodes1_vs1:dstvol_2".

Remote::*> snapmirror show
This table is currently empty.

```

## 3. Release the existing relationship but retain the common Snapshot copies.

```

Primary::*> snapmirror release -destination-path cnodes1_vs1:dstvol_1 -relationship-info-only true
[Job 60] Job succeeded: SnapMirror Release Succeeded

Primary::*> snapmirror release -destination-path cnodes1_vs1:dstvol_2 -relationship-info-only true
[Job 61] Job succeeded: SnapMirror Release Succeeded

```

## 4. Create the reverse relationship.

```

Primary::*> snapmirror create -source-path cnodes1_vs1:dstvol_1 -destination-path cnodes0_vs1:srcvol_1 -policy StrictSync

Warning: You are creating a SnapMirror relationship with a policy of type
"strict-sync-mirror" that only supports all LUN based applications
with FCP and iSCSI protocols, as well as NFSv3 protocol for enterprise
applications such as databases, VMWare, etc.

Warning: For a SnapMirror relationship with a policy of type
"strict-sync-mirror", client I/O will fail in order to maintain strict
synchronization when the secondary is inaccessible.
Do you want to continue? {y|n}: y
Operation succeeded: snapmirror create for the relationship with destination "cnodes0_vs1:srcvol_1".

Primary::*> snapmirror create -source-path cnodes1_vs1:dstvol_2 -destination-path cnodes0_vs1:srcvol_2 -policy Sync
Operation succeeded: snapmirror create for the relationship with destination "cnodes0_vs1:srcvol_2".

Warning: You are creating a SnapMirror relationship with a policy of type
"sync-mirror" that only supports all LUN based applications with FCP
and iSCSI protocols, as well as NFSv3 protocol for enterprise
applications such as databases, VMWare, etc.

```

## 5. Resync the relationship.

```

Primary::*> snapmirror resync -destination-path cnodes0_vs1:srcvol_1

Warning: All data newer than Snapshot copy
snapmirror.3811aa03-d894-11e8-aef9-00a0985a6a80_2147894256.2018-10-26_0
22311 on volume cnodes0_vs1:srcvol_1 will be deleted.
Do you want to continue? {y|n}: y
Operation is queued: initiate snapmirror resync to destination "cnodes0_vs1:srcvol_1".

Primary::*> snapmirror resync -destination-path cnodes0_vs1:srcvol_2

Warning: All data newer than Snapshot copy
snapmirror.3811aa03-d894-11e8-aef9-00a0985a6a80_2147894257.2018-10-26_0
22839 on volume cnodes0_vs1:srcvol_2 will be deleted.
Do you want to continue? {y|n}: y
Operation is queued: initiate snapmirror resync to destination "cnodes0_vs1:srcvol_2".

```

## 3.21 Migrate Existing Synchronous SnapMirror Relationships from 7-Mode to ONTAP

SM-S, introduced in ONTAP 9.5, does not support using a source volume from a 7-mode storage system. However, you can use the ONTAP CLI or the 7-Mode Transition Tool (7MTT) to convert existing 7-Mode SnapMirror relationships to SnapMirror Synchronous relationships without another baseline. To do so complete the following steps:

1. Enable transition of Snapshot copies from 7-Mode to ONTAP storage systems in the primary and secondary sites.
  - a. Check the status of SnapMirror on 7-mode storage systems.

```

7m-dst:> snapmirror status
Snapmirror is on.
Source Destination State Lag Status
7m-src:ssm_src 7m-dst:ssm_dst Snapmirrored 00:00:00 In-sync

```

- b. Enable the DP replication type using the cluster-wide option on the primary and secondary ONTAP storage system clusters.

```
options replication.create_data_protection_rels.enable on
```

- c. Create peer transitions between the following systems:

- The 7-Mode primary storage system and an ONTAP primary cluster
- The 7-Mode secondary storage system and an ONTAP secondary cluster

```
cm_primary::> vserver peer transition create -local-vserver vs0 -src-filer-name 10.63.38.63
Transition peering created
```

- d. Create and initialize the relationship between the 7-mode primary storage system and the ONTAP primary storage system.

```
cm_primary::> snapmirror create -source-path 10.63.38.63:ssm_src -destination-path vs0:ssm_src7m
-type TDP
Operation succeeded: snapmirror create for the relationship with destination "vs0:ssm_src7m".
cm_primary::> snapmirror initialize vs0:ssm_src7m
Operation is queued: snapmirror initialize of destination "vs0:ssm_src7m".
```

- e. Create and initialize the relationship between the 7-mode secondary storage system and the ONTAP secondary storage system.

```
cm_secondary::> snapmirror create -source-path 10.63.39.5:ssm_dst -destination-path vs1:ssm_dst7m
-type TDP
Operation succeeded: snapmirror create for the relationship with destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror initialize vs1:ssm_dst7m
Operation is queued: snapmirror initialize of destination "vs1:ssm_dst7m".
```

- f. Plan to cutover, quiesce, and break the 7-mode to ONTAP relationships for both the primary and secondary ONTAP storage systems.

```
cm_primary::> snapmirror quiesce vs0:ssm_src7m
Operation succeeded: snapmirror quiesce for destination "vs0:ssm_src7m".
cm_primary::> snapmirror break vs0:ssm_src7m
Operation succeeded: snapmirror break for destination "vs0:ssm_src7m".
cm_secondary::> snapmirror quiesce vs1:ssm_dst7m
Operation succeeded: snapmirror quiesce for destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror break vs1:ssm_dst7m
Operation succeeded: snapmirror break for destination "vs1:ssm_dst7m".
```

- g. Delete the 7-mode to ONTAP SnapMirror relationships.

```
cm_secondary::> snapmirror delete vs1:ssm_dst7m
Operation succeeded: snapmirror delete for the relationship with destination "vs1:ssm_dst7m".
```

2. Create and resync the primary and secondary volumes on the ONTAP storage systems using `-type DP`.

```
cm_secondary::> snapmirror create -source-path vs0:ssm_src7m -destination-path vs1:ssm_dst7m -
type DP
Operation succeeded: snapmirror create for the relationship with destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror resync vs1:ssm_dst7m
Warning: All data newer than Snapshot copy vs0(4143489260)_ssm_src7m.2 on
volume vs1:ssm_dst7m will be deleted.
Do you want to continue? {y|n}: y
Operation is queued: initiate snapmirror resync to destination "vs1:ssm_dst7m".
```

3. Complete the resynchronization, perform a quiesce operation, and break and delete the DP relationship.

```
cm_secondary::> snapmirror quiesce vs1:ssm_dst7m
Operation succeeded: snapmirror quiesce for destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror break vs1:ssm_dst7m
Operation succeeded: snapmirror break for destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror delete vs1:ssm_dst7m
Operation succeeded: snapmirror delete for the relationship with destination "vs1:ssm_dst7m".
```

4. Create a SnapMirror relation between the primary and secondary volumes of the ONTAP storage systems using `-type XDP`, and complete the resynchronization.

```
cm_secondary::> snapmirror create -source-path vs0:ssm_src7m -destination-path vs1:ssm_dst7m -
type XDP
Operation succeeded: snapmirror create for the relationship with destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror resync vs1:ssm_dst7m
```

```

Operation is queued: initiate snapmirror resync to destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror show Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs0:ssm_src7m XDP vs1:ssm_dst7m Snapmirrored Idle - true -

```

5. Convert the SnapMirror relationship from asynchronous to synchronous by completing the following steps:

a. Delete the asynchronous relationship.

```

cm_secondary::> snapmirror delete vs1:ssm_dst7m
Operation succeeded: snapmirror delete for the relationship with destination "vs1:ssm_dst7m".

```

b. Release the asynchronous relationship with the parameter `-relationship-info-only`.

```

cm_secondary::> snapmirror release vs1:ssm_dst7m -relationship-info-only
[Job 36] Job succeeded: SnapMirror Release Succeeded

```

c. Create a SM-S relationship with the policy type of `Sync` or `Strict Sync`.

```

cm_secondary::> snapmirror create -source-path vs0:ssm_src7m -destination-path vs1:ssm_dst7m -
policy Sync
Operation succeeded: snapmirror create for the relationship with destination "vs1:ssm_dst7m".

```

d. Perform resynchronization. The SnapMirror relationship then enters the `InSync` status.

```

cm_secondary::> snapmirror resync vs1:ssm_dst7m
Operation is queued: initiate snapmirror resync to destination "vs1:ssm_dst7m".
cm_secondary::> snapmirror show Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs0:ssm_src7m XDP vs1:ssm_dst7m Snapmirrored InSync - true -

```

### 3.22 Mirror State and Relationship Status

A SnapMirror relationship can have its mirror state as Uninitialized until the baseline transfer completes, after which the relationship status changes to Snapmirrored. The destination dictates the mirror state and the source dictates the relationship status.

Table 3) Mirror state and relationship status.

Operation	Mirror State	Relationship Status	Healthy	Autoresync Enabled
Create new SnapMirror relationship	Uninitialized	Idle	Yes	No
Initialize – baseline phase	Uninitialized	Transferring	Yes	No
Initialize – baseline failure	Uninitialized	Idle	No	No
Initialize - baseline successful - back-to-back asynchronous transfer phase	Snapmirrored	Transferring	Yes	Yes
Initialize - back-to-back asynchronous transfer failure	Snapmirrored	OutOfSync	No	Yes

Operation	Mirror State	Relationship Status	Healthy	Autoresync Enabled
Initialize/Resync - back-to-back transfer successful - Transition Phase	Snapmirrored	Transitioning	Yes	Yes
Initialize/Resync - transition failure	Snapmirrored	OutOfSync	No	Yes
Initialize/Resync - transition successful	Snapmirrored	InSync	Yes	Yes
Synchronous replication failure - autoresync enabled	Snapmirrored	OutOfSync	No	Yes
Synchronous replication failure - autoresync gives up	Snapmirrored	Idle	No	No
Quiesce - transfer in progress	Snapmirrored	Quiescing	Yes	No
Quiesce - in progress transfer failure	Snapmirrored	Quiesced	No (last transfer failed)	No
Quiesce - in progress transfer completed	Snapmirrored	Quiesced	No (schedule missed)	No
Resume	Snapmirrored	Idle	Yes/No	No
Break - successful	Broken-Off	Idle	Yes	No
Break - failed	Snapmirrored	Idle	No	No
Resync - asynchronous resynchronization failure	Broken-Off	Idle	No	No
Resync - asynchronous resynchronization successful – back-to-back asynchronous transfer phase	Snapmirrored	Transferring	Yes	Yes

## 4 Manage SnapMirror Synchronous Operations Through ONTAP System Manager

You can use ONTAP System Manager to perform various SnapMirror operations, including create, delete, and manage SM-S relationships.

You can perform the following tasks:

- Initialize SnapMirror destinations
- Update SnapMirror relationships
- Quiesce SnapMirror relationships

- Resume quiesced SnapMirror relationships
- Break SnapMirror relationships
- Resynchronize and reverse resynchronize SnapMirror relationships
- Abort SnapMirror data transfers
- Delete SnapMirror relationships

## 5 SnapMirror Synchronous and NetApp Storage Efficiency

Because SnapMirror Synchronous writes data to both primary and secondary storage volumes simultaneously, there is no question of storage efficiency being maintained over the wire, as is the case for asynchronous SnapMirror relationships. Deduplication takes place only after data is written to the primary and secondary storage systems. This helps maintain storage efficiency benefits in replicated volumes.

## 6 SnapMirror Synchronous Calculator

The [SM-S Replication Advisor](#) allows you to input application workload metrics along with model types and network characteristics to determine whether synchronous replication is feasible for your configuration. Table 4 shows the inputs to the sizer.

Table 4) Workload metrics.

Configuration	Inputs
Application:	<ul style="list-style-type: none"> <li>• IOPS:</li> <li>• Write %:</li> </ul>
Source node:	<ul style="list-style-type: none"> <li>• CPU %:</li> <li>• Volumes:</li> <li>• Model:</li> </ul>
Destination node:	<ul style="list-style-type: none"> <li>• CPU %:</li> <li>• Model:</li> </ul>
Network:	<ul style="list-style-type: none"> <li>• Bandwidth:</li> <li>• RTT:</li> </ul>

**Note:** You determine the maximum IOPS that the source and destination nodes can support based on the model.

## 7 Revert to Previous ONTAP Versions

If an 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.

## 8 Troubleshooting

### 8.1 Operations That Cause SnapMirror Relationships to Go OutOfSync

#### Unhealthy Cluster Peering

The SnapMirror relationship goes to a persistent failure state under any of the following circumstances:

- If a cluster peer is not available
- If the source and destination clusters are not reachable after five autoresync attempts
- If there is network congestion between the source and destination clusters

```
C2_sti97-vsिम-ucs542k_cluster::> snapmirror show -fields state, status, health, last-transfer-error
source-path destination-path state status healthy last-transfer-error
-----
vs0:Vol1 vs1: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.)
vs0:Vol2 vs1:Vol2 Snapmirrored Transferring false Prechecks on source volume failed. (CSM: A get-
session operation failed because no (local) transport address was registered for the node.)
vs0:Vol3 vs1:Vol3 Snapmirrored Transferring false Prechecks on source volume failed. (CSM: A get-
session operation failed because no (local) transport address was registered for the node.)
vs0:Vol4 vs1:Vol4 Snapmirrored Transferring false Prechecks on source volume failed. (CSM: A get-
session operation failed because no (local) transport address was registered for the node.)
4 entries were displayed.

sti97-vsिम-ucs542k% tail -f /mroot/etc/log/snapmirror_audit
Wed Oct 24 05:00:03 EDT 2018 ResyncTransfer[Oct 24 05:00:03]:21adcfb5-d755-11e8-85a7-005056a7ecc7
Operation-Uuid=72a3c715-173b-40b5-9c2c-f20fla68ddld Group=none Operation-Cookie=0 action=Start
source=vs0:Vol4 destination=vs1:Vol4
Wed Oct 24 05:00:03 EDT 2018 ResyncTransfer[Oct 24 05:00:03]:633eeb95-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=8d673383-c9ab-4efe-9296-09e935ebf1e8 Group=none Operation-Cookie=0 action=Start
source=vs0:Vol2 destination=vs1:Vol2
Wed Oct 24 05:01:03 EDT 2018 ResyncTransfer[Oct 24 05:00:03]:633eeb95-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=8d673383-c9ab-4efe-9296-09e935ebf1e8 Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol2 destination=vs1:Vol2 status=Failure message=Prechecks on source volume
failed.(CSM: Connection aborted.)
Wed Oct 24 05:01:03 EDT 2018 ResyncTransfer[Oct 24 05:00:03]:21adcfb5-d755-11e8-85a7-005056a7ecc7
Operation-Uuid=72a3c715-173b-40b5-9c2c-f20fla68ddld Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol4 destination=vs1:Vol4 status=Failure message=Prechecks on source volume
failed.(CSM: Connection aborted.)
Wed Oct 24 05:01:03 EDT 2018 ResyncTransfer[Oct 24 05:00:03]:79919e97-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=58b86c96-48bf-4f05-b211-d0b49132d687 Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol3 destination=vs1:Vol3 status=Failure message=Prechecks on source volume
failed.(CSM: Connection aborted.)
Wed Oct 24 05:01:03 EDT 2018 ResyncTransfer[Oct 24 05:00:03]:7542a8ff-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=29ae8a4e-110f-4e00-80e6-c1236cf20166 Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol1 destination=vs1:Vol1 status=Failure message=Prechecks on source volume
failed.(CSM: Connection aborted.)

l_sti97-vsिम-ucs542i_cluster::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability Authentication
-----
C2_sti97-vsिम-ucs542k_cluster 1-80-000011 Partial ok

C1_sti97-vsिम-ucs542i_cluster::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability Authentication
-----
C2_sti97-vsिम-ucs542k_cluster 1-80-000011 Unavailable ok
```

#### Resolution

Wait for autoresync to complete to bring the relationship status back to InSync.

## Nondisruptive Operation

There might be situations in which a particular node in the cluster has higher usage (for example, I/O) than others. Or there might be a need for a planned decommission of an HA pair that requires data evacuation from those particular nodes. Nondisruptive operations (NDOs) such as storage failover (SFO) and aggregate relocation (ARL) allow aggregates to be relocated from one node to another in a cluster. Volume moves can be used to level-down the effect on the performance of each node and nondisruptively move volumes across aggregates.

Typically, during the cutover phase of NDO operations, client access is temporarily blocked during the short period when a final replication from source to destination is being made. An identity swap changes the destination to the source location, where the system can then start routing client traffic to the new source location, resuming client access and disposing of the old data location. However, performing a volume move on a volume that is synchronously replicating data affects zero RPO, so there is automatic conversion from StrictSync to Sync mode during this cutover phase to avoid client I/O disruption or application failure. For SFO or ARL, you are prompted with a warning to choose to convert from StrictSync to Sync to avoid application I/O failure.

```
C2_sti97-vsिम-ucs542k_cluster::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
-----
vs0:Vol1 XDP vs1:Vol1 Snapmirrored InSync - true -
vs0:Vol2 XDP vs1:Vol2 Snapmirrored InSync - true -
vs0:Vol3 XDP vs1:Vol3 Snapmirrored InSync - true -
vs0:Vol4 XDP vs1:Vol4 Snapmirrored InSync - true -

C1_sti97-vsिम-ucs542i_cluster::> volume move start -volume Vol1 -destination-aggregate
sti97_vsिम_ucs542j_aggr1

Warning: Volume "Vol1" in Vserver "vs0" is currently part of a SnapMirror
Synchronous relationship with a policy of type "strict-sync-mirror". A
volume move operation will fail the zero RPO protection. To prevent
primary IO failure, this relationship will behave as if had a policy
of type "sync-mirror" from the beginning of volume move cutover until
the volume move completes and the relationship status is "InSync".
Do you want to continue? {y|n}: y
[Job 48] Job is queued: Move "Vol1" in Vserver "vs0" to aggregate "sti97_vsिम_ucs542j_aggr1". Use
the "volume move show -vserver vs0 -volume Vol1" command to view the status of this operation.

C1_sti97-vsिम-ucs542i_cluster::> volume move show
Vserver Volume State Move Phase Percent-Complete Time-To-Complete
-----
vs0 Vol1 healthy replicating
--

C1_sti97-vsिम-ucs542i_cluster::> volume move show
Vserver Volume State Move Phase Percent-Complete Time-To-Complete
-----
vs0 Vol1 healthy replicating
75% Wed Oct 24 06:16:04 2018

C1_sti97-vsिम-ucs542i_cluster::> volume move show -instance

Vserver Name: vs0
Volume Name: Vol1
Actual Completion Time: -
Bytes Remaining: 148KB
Destination Aggregate: sti97_vsिम_ucs542j_aggr1
Detailed Status: Volume move job locking down volumes for final transfer - This where mutual
exclusive comes in and relation goes OutofSync.
Estimated Time of Completion: Wed Oct 24 06:16:55 2018
Managing Node: sti97-vsिम-ucs542i
Percentage Complete: 98%
Move Phase: replicating
Estimated Remaining Duration: 00:00:30
```

```

Replication Throughput: 8KB/s
Duration of Move: 00:01:03
Source Aggregate: sti97_vsim_ucs542i_aggr1
Start Time of Move: Wed Oct 24 06:15:22 2018
Move State: healthy
Is Source Volume Encrypted: false
Encryption Key ID of Source Volume:
Is Destination Volume Encrypted: false
Encryption Key ID of Destination Volume:

C2_sti97-vsim-ucs542k_cluster::> snapmirror show -fields state, status, health, last-transfer-
error
source-path destination-path state status healthy last-transfer-error
-----
vs0:Vol1 vs1:Vol1 Snapmirrored OutOfSync false -
vs0:Vol2 vs1:Vol2 Snapmirrored InSync true -
vs0:Vol3 vs1:Vol3 Snapmirrored InSync true -
vs0:Vol4 vs1:Vol4 Snapmirrored InSync true -
4 entries were displayed.

C2_sti97-vsim-ucs542k_cluster::> snapmirror show -fields state, status, health, last-transfer-
error
source-path destination-path state status healthy last-transfer-error
-----
vs0:Vol1 vs1:Vol1 Snapmirrored Transferring false -
vs0:Vol2 vs1:Vol2 Snapmirrored InSync true -
vs0:Vol3 vs1:Vol3 Snapmirrored InSync true -
vs0:Vol4 vs1:Vol4 Snapmirrored InSync true -
4 entries were displayed.

```

**Note:** Audit log entries confirm replication failure at the destination.

```

C2_sti97-vsim-ucs542k_cluster::> tail -f /mroot/etc/log/snapmirror_audit

Wed Oct 24 06:15:42 EDT 2018 InSyncTransfer[Oct 24 05:21:03]:7542a8ff-d754-11e8-85a7-005056a7ecc7
Operation-Uuid= Group=none Operation-Cookie=0 action=Info Transfer failed..
Wed Oct 24 06:16:02 EDT 2018 ResyncTransfer[Oct 24 06:16:02]:7542a8ff-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=d34d8b85-9017-405e-a4e5-830258937d64 Group=none Operation-Cookie=0 action=Start
source=vs0:Vol1 destination=vs1:Vol1
Wed Oct 24 06:16:02 EDT 2018 ResyncTransfer[Oct 24 06:16:02]:7542a8ff-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=d34d8b85-9017-405e-a4e5-830258937d64 Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol1 destination=vs1:Vol1 status=Failure message=Prechecks on source volume
failed.(Volume move of the source volume is in progress.)
Wed Oct 24 06:17:05 EDT 2018 ResyncTransfer[Oct 24 06:16:02]:7542a8ff-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=d34d8b85-9017-405e-a4e5-830258937d64 Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol1 destination=vs1:Vol1 status=Failure message=Prechecks on source volume
failed.(Volume move of the source volume is in progress.)
Wed Oct 24 06:18:08 EDT 2018 ResyncTransfer[Oct 24 06:16:02]:7542a8ff-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=d34d8b85-9017-405e-a4e5-830258937d64 Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol1 destination=vs1:Vol1 status=Failure message=Prechecks on source volume
failed.(Volume move of the source volume is in progress.)
Wed Oct 24 06:19:11 EDT 2018 ResyncTransfer[Oct 24 06:16:02]:7542a8ff-d754-11e8-85a7-005056a7ecc7
Operation-Uuid=d34d8b85-9017-405e-a4e5-830258937d64 Group=none Operation-Cookie=0 action=Defer
source=vs0:Vol1 destination=vs1:Vol1 status=Failure message=Prechecks on source volume
failed.(Failed to get volume attributes for 8589e0de-d74e-11e8-8e4e-005056a73476:Vol1.(Volume is
not known or has been moved))

```

## Resolution

Autosync or manual resync after an NDO event helps to recover the SnapMirror relationship with the status InSync.

## Aggregate or Volume Goes Offline

When a primary volume or aggregate goes offline, the following event occurs on the destination volume:

```

C2_stiA800-245_cluster:*> event log show -event *sms*
Time Node Severity Event

```

```
-----  
10/24/2018 01:42:29 stiA800-245 ERROR sms.status.out.of.sync: Source volume "vs0:ssm_c1nlvs0src1"  
and destination volume "vs1:ssm_c2nlvs1dst1" with relationship UUID "15f584d5-d67f-11e8-95f2-  
00a098b9251d" is in "out-of-sync" status due to the following reason: "Transfer failed."
```

A corresponding event occurs on the source volume in case I/O did not happen, because the source volume is offline:

```
C1_stiA800-243_cluster::*> event log show -event *sms*  
Time Node Severity Event  
-----  
10/24/2018 01:42:29 stiA800-243 ERROR sms.status.out.of.sync: Source volume "e9d4bab4-cd8e-11e8-  
96ae-00a098ba7d2b:ssm_c1nlvs0src1" and destination volume "df811f1b-cd8e-11e8-b2c3-  
00a098b9251d:ssm_c2nlvs1dst1" with relationship UUID "15f584d5-d67f-11e8-95f2-00a098b9251d" is in  
"out-of-sync" status due to the following reason: "Transfer failed.(Job initiated abort(Refer to  
the corresponding sms.status.out.of.sync EMS generated on the SnapMirror destination cluster for  
more details on the OutOfSync reason. If the SnapMirror destination does not generate an  
OutOfSync EMS this could be a transient failure during a resync operation.))".
```

A corresponding event occurs on the source volume in case I/O writes were in progress when the volume went offline:

```
sti8080-447-0011536251499::> event log show -event sms*  
Time Node Severity Event  
-----  
10/24/2018 01:47:45 stiA800-243 ERROR sms.status.out.of.sync: Source volume  
"alias1:SE_PriVoll1" and destination volume "vs0:SE_PriVoll1_dst" with relationship UUID "c08c0250-  
c34e-11e8-b009-00a0986102e7" is in "out-of-sync" status due to the following reason: "Transfer  
failed.(Granular Sync Op encountered a failure while performing Quick Reconciliation.)".
```

When a secondary volume or aggregate goes offline, the following event occurs on the destination:

```
C2_stiA800-245_cluster::*> event log show -event *sms*  
Time Node Severity Event  
-----  
10/24/2018 01:38:19 stiA800-245 ERROR sms.status.out.of.sync: Source volume "vs0:ssm_c1nlvs0src1"  
and destination volume "vs1:ssm_c2nlvs1dst1" with relationship UUID "15f584d5-d67f-11e8-95f2-  
00a098b9251d" is in "out-of-sync" status due to the following reason: "Transfer failed.(Job  
initiated abort)".
```

A corresponding event occurs on the source volume:

```
C1_stiA800-243_cluster::*> event log show -event *sms*  
Time Node Severity Event  
-----  
10/24/2018 01:38:24 stiA800-243 ERROR sms.status.out.of.sync: Source volume "e9d4bab4-cd8e-11e8-  
96ae-00a098ba7d2b:ssm_c1nlvs0src1" and destination volume "df811f1b-cd8e-11e8-b2c3-  
00a098b9251d:ssm_c2nlvs1dst1" with relationship UUID "15f584d5-d67f-11e8-95f2-00a098b9251d" is in  
"out-of-sync" status due to the following reason: "Transfer failed.(The destination replication  
engine could not be reached by the source(Refer to the corresponding sms.status.out.of.sync EMS  
generated on the SnapMirror destination cluster for more details on the OutOfSync reason. If the  
SnapMirror destination does not generate an OutOfSync EMS this could be a transient failure  
during a resync operation.))".
```

## Probable Cause

Check the volume and aggregate status on both the source and the destination by running the following commands:

```
volume show  
aggregate show
```

Check whether an event that took the source volume offline was triggered during the time of failure by running the following command:

```
C1_stiA800-243_cluster::*> event log show -event *offline*
```

```
Time Node Severity Event
```

```
-----  
9/9/2018 13:02:18 sti8060-htp-001 INFORMATIONAL waf1.vvol.offline: Volume  
'sm_c1nlvs0src1@vserver:d3bcb7ee-b1f5-11e8-a547-00a098540637' has been set temporarily offline
```

## Resolution

1. Bring the volume or aggregate online by using the following command.

```
volume online <volume name> -vserver <vserver name>  
aggregate online <aggregate name>
```

2. Wait for autoresync to complete to bring the relationship status back to InSync.
3. If the autoresync attempts fail because waf1ron was run on the source volume, you get the following error message:

```
cluster-2::*> sn show -fields last-transfer-error  
source-path destination-path last-transfer-error  
-----  
vs1:vol13 vs2:vol13 Prechecks on source volume failed. (Waf1ron has fixed inconsistencies on the  
volume. Run "snapmirror resync". Ensure that snapmirror checker tool is run on the common  
snapshots.)
```

4. If waf1ron was run on the destination volume, you get the following error message:

```
cluster-1::*> snapmirror show -fields status,last-transfer-error  
source-path destination-path status last-transfer-error  
-----  
vs1:srcvol vs1:dstvol OutOfSync Destination vs1:dstvol must be a data-protection volume.
```

5. Run the SnapMirror checker tool to verify consistency.
6. Run a manual resync operation.

## Insufficient Space on Volume or Aggregate

If there is insufficient space on the destination volume, the following event occurs:

```
C2_stiA800-245_cluster::*> event log show -event *sms*  
Time Node Severity Event  
-----  
10/24/2018 01:42:29 stiA800-245 ERROR sms.status.out.of.sync: Source volume "vs0:ssm_c1nlvs0src1"  
and destination volume "vs1:ssm_c2nlvs1dst1" with relationship UUID "15f584d5-d67f-11e8-95f2-  
00a098b9251d" is in "out-of-sync" status due to the following reason: "Transfer failed."
```

If the primary volume or aggregate is out of space, a corresponding source event occurs:

```
C1_stiA800-243_cluster::*> event log show -event *sms*  
Time Node Severity Event  
-----  
10/24/2018 01:42:29 stiA800-243 ERROR sms.status.out.of.sync: Source volume "e9d4bab4-cd8e-11e8-  
96ae-00a098ba7d2b:ssm_c1nlvs0src1" and destination volume "df811f1b-cd8e-11e8-b2c3-  
00a098b9251d:ssm_c2nlvs1dst1" with relationship UUID "15f584d5-d67f-11e8-95f2-00a098b9251d" is in  
"out-of-sync" status due to the following reason: "Transfer failed.(Granular Sync Op encountered  
a failure while performing Quick Reconciliation.)".
```

If the secondary volume or aggregate is out of space, a corresponding source event occurs:

```
C1_stiA800-243_cluster::*> event log show -event *sms*  
Time Node Severity Event  
-----  
10/24/2018 01:42:29 stiA800-243 ERROR sms.status.out.of.sync: Source volume "e9d4bab4-cd8e-11e8-  
96ae-00a098ba7d2b:ssm_c1nlvs0src1" and destination volume "df811f1b-cd8e-11e8-b2c3-  
00a098b9251d:ssm_c2nlvs1dst1" with relationship UUID "15f584d5-d67f-11e8-95f2-00a098b9251d" is in  
"out-of-sync" status due to the following reason: "Transfer failed.(Granular Sync failed to  
replicate Op to Secondary.)".
```

## Probable Cause

If the primary volume or aggregate runs out of space, the I/O writes to primary fail and quick reconciliation tries to read from secondary and write to primary. However, quick reconciliation might fail.

If the secondary volume or aggregate runs out of space, write to secondary might fail.

Check the volume or aggregate size and its availability by running the following commands:

```
volume show
aggregate show
```

## Resolution

1. Increase the size of the volume by running the following command:

```
volume size -volume <volume name> -vserver <vserver name> -newsize <size>
```

2. If necessary, add additional disks to the aggregate by running the following command:

```
Aggregate -adddisk -aggregate <aggregate name> diskcount <number>
```

3. Wait for autoresync to recover the SnapMirror relationship to status InSync.

## SnapMirror Abort

An EMS notification is displayed on the source cluster when `snapmirror abort` operations are performed.

```
C1_stiA800-243_cluster::*> event log show -event *sms*
Time Node Severity Event
-----
10/25/2018 07:44:23 stiA800-243 ERROR sms.status.out.of.sync: Source volume "e9d4bab4-cd8e-11e8-96ae-00a098ba7d2b:SE_PriVol11" and destination volume "df811f1b-cd8e-11e8-b2c3-00a098b9251d:SE_PriVol11_dst" with relationship UUID "cb1b3a18-d831-11e8-b165-00a098ba7cef" is in "out-of-sync" status due to the following reason: "Transfer failed.(Job initiated abort(Refer to the corresponding sms.status.out.of.sync EMS generated on the SnapMirror destination cluster for more details on the OutOfSync reason. If the SnapMirror destination does not generate an OutOfSync EMS this could be a transient failure during a resync operation.))".
```

Verify that a `snapmirror abort` operation has been attempted by running the `snapmirror show-history` command:

```
C2_stiA800-245_cluster::*> snapmirror show-history -destination-path vs1:SE_PriVol3_dst
Destination Source Start End
Path Path Operation Time Time Result
-----
vs1:SE_PriVol3_dst vs0:SE_PriVol3 scheduled-update 10/25/2018 07:05:00 10/25/2018 07:23:32
success
vs1:SE_PriVol3_dst vs0:SE_PriVol3 automatic-resync 10/25/2018 06:23:00 10/25/2018 07:03:28
success
vs1:SE_PriVol3_dst vs0:SE_PriVol3 automatic-resync 10/25/2018 03:00:00 10/25/2018 06:22:17
failure
vs1:SE_PriVol3_dst vs0:SE_PriVol3 automatic-resync 10/25/2018 02:17:00 10/25/2018 02:59:58
failure
vs1:SE_PriVol3_dst vs0:SE_PriVol3 abort 10/25/2018 02:55:48 10/25/2018 02:55:48 success
vs1:SE_PriVol3_dst vs0:SE_PriVol3 quiesce 10/25/2018 02:55:28 10/25/2018 02:55:28 failure
vs1:SE_PriVol3_dst vs0:SE_PriVol3 initialize 12/31/1969 19:00:00 10/25/2018 02:11:45 success
vs1:SE_PriVol3_dst vs0:SE_PriVol3 create 10/25/2018 02:11:43 10/25/2018 02:11:43 success
```

## Resolution

Autoresync or manual `resync` helps to recover the SnapMirror relationship to status InSync.

## 8.2 Resync Failures

### Unable to Meet Transition Criteria Because Resync Takes a Long Time

Auto or manual resync fails with an error in the `snapmirror show` output.

```
path,type,state,status,last-transfer-error
source-path destination-path type policy state status last-transfer-error
-----
vs0:vol2 vs1:vol2 XDP StrictSync Snapmirrored OutOfSync Failed to transition SnapMirror
Synchronous relationship "vs1:vol2" into InSync. Reason: Asynchronous transfers are consistently
above acceptable duration threshold making transition to InSync unsuitable.
```

Check each resync duration from the destination `snapmirror_audit` log, available at `/mroot/etc/log`.

```
Thu Jun 28 06:37:30 EDT 2018 ResyncTransfer[Jun 28 06:29:24]:d7800405-74b9-11e8-8288-00a0985f663d
Operation-Uuid=bd49c7d7-6f62-4efa-96d3-c4bc1c02ebf6 Group=none Operation-Cookie=0
action=Transfer_snapshots (1 of 1 log entries): snapmirror.cc42a2cd-65ca-11e8-bbaa-
00a0985f663d_2161757991.2018-06-28_063728
Thu Jun 28 06:37:59 EDT 2018 ResyncTransfer[Jun 28 06:29:24]:d7800405-74b9-11e8-8288-00a0985f663d
Operation-Uuid=bd49c7d7-6f62-4efa-96d3-c4bc1c02ebf6 Group=none Operation-Cookie=0
Autoresync B2B Transfer: dest: vs1:vol2. Total: 19, low: 0, high: 19, duration: 3601
Thu Jun 28 06:38:02 EDT 2018 ResyncTransfer[Jun 28 06:29:24]:d7800405-74b9-11e8-8288-00a0985f663d
Operation-Uuid=bd49c7d7-6f62-4efa-96d3-c4bc1c02ebf6 Group=none Operation-Cookie=0 action=Info
sm_rpm_snaplist_delete([86]): Snapshot snapmirror.cc42a2cd-65ca-11e8-bbaa-
00a0985f663d_2161757991.2018-06-28_063643 deleted.
Thu Jun 28 06:38:02 EDT 2018 ResyncTransfer[Jun 28 06:29:24]:d7800405-74b9-11e8-8288-00a0985f663d
Operation-Uuid=bd49c7d7-6f62-4efa-96d3-c4bc1c02ebf6 Group=none Operation-Cookie=0
action=Transfer_snapshots (1 of 1 log entries): snapmirror.cc42a2cd-65ca-11e8-bbaa-
00a0985f663d_2161757991.2018-06-28_063759
Thu Jun 28 06:38:36 EDT 2018 ResyncTransfer[Jun 28 06:29:24]:d7800405-74b9-11e8-8288-00a0985f663d
Operation-Uuid=bd49c7d7-6f62-4efa-96d3-c4bc1c02ebf6 Group=none Operation-Cookie=0 action=Info
Autoresync B2B Transfer: dest: vs1:vol2. Total: 20, low: 0, high: 20, duration: 3601
```

### Probable Cause

The change rate might be higher than the available network can support. Check IOPS and CPU statistics, or the difference since the last common Snapshot copy is too great, contributing to the extended resync.

Verify that consistency point (CP) statistics are using node-level commands in diagnostic mode:

```
run local waf1 cpstats -s -a <aggregate name>.
```

An initialize or resync operation occurred while converting from an asynchronous to a synchronous SnapMirror relationship.

Common Snapshot copy creation failure or transition aborts due to blacklisted operations or a breach of limits.

### Resolution

1. Size the environment based on best practices and use QoS for workload optimization.
2. Perform a manual resync to recover the SnapMirror relationship to status InSync.

### No Common Snapshot Copy Found

Auto or manual resync fails with an error in the `snapmirror show` output.

```
cluster-1::*> snapmirror show -fields status,last-transfer-error
source-path destination-path status last-transfer-error
-----
vs1:srcvol vs1:dstvol OutOfSync No common Snapshot copy found between vs1:srcvol and vs1:dstvol
```

## Probable Cause

The Snapshot copy autodelete feature is enabled, causing deletion of the common Snapshot copies on these volumes.

```
ssan-8060-03::*> event log show -message-name waf1.volume.snap.autoDelete
Time                Node                Severity            Event
-----
7/12/2018 05:46:33  ssan-8060-03b      INFORMATIONAL      waf1.volume.snap.autoDelete: Deleting Snapshot
copy 'snapmirror.b61aad3d-7b80-11e8-888d-00a098a31ddd_2159361896.2018-07-12_054007' in volume
'SSMSrc@vserver:b61aad3d-7b80-11e8-888d-00a098a31ddd' to recover storage.
7/12/2018 05:46:20  ssan-8060-03b      INFORMATIONAL      waf1.volume.snap.autoDelete: Deleting Snapshot
copy 'snapmirror.b61aad3d-7b80-11e8-888d-00a098a31ddd_2159361896.2018-07-12_054000' in volume
'SSMSrc@vserver:b61aad3d-7b80-11e8-888d-00a098a31ddd' to recover storage.
7/12/2018 05:45:56  ssan-8060-03b      INFORMATIONAL      waf1.volume.snap.autoDelete: Deleting Snapshot
copy 'S3' in volume 'SSMSrc@vserver:b61aad3d-7b80-11e8-888d-00a098a31ddd' to recover storage.
7/12/2018 05:45:40  ssan-8060-03b      INFORMATIONAL      waf1.volume.snap.autoDelete: Deleting Snapshot
copy 'S2' in volume 'SSMSrc@vserver:b61aad3d-7b80-11e8-888d-00a098a31ddd' to recover storage.
```

The asynchronous relationship schedule might be higher than the synchronous relationship in a three-site disaster recovery scenario, with an async fan-out relationship and a new sync between the disaster recovery volumes.

## Resolution

1. Disable the autodelete parameter or provision sufficient space on the source and destination volumes.
2. In a site disaster recovery scenario, set the asynchronous schedule to be one hour or less.

## Transition Failures

Auto or manual resync fails with an error in the `snapmirror show` output.

```
sti8060-2671527862493::*> snapmirror show -fields source-path,destination-
path,type,state,status,last-transfer-error
source-path destination-path type policy state status last-transfer-error
-----
vs0:vol2 vs1:vol2 XDP StrictSync Snapmirrored OutOfSync Failed to start replication on source.
```

From the destination `snapmirror_error` log available at `/mroot/etc/log`, check for an error signature at the time of failure.

```
> 2018-07-24T15:59:39Z 1171725196136205 [8:0] REPL_0:
repl_writer::PushFileUpdateHoleReplOpHandler::replXlateOut(): | [254b8fa9-8f4b-11e8-b4c4-
00a098a46c68] | [0xfffff80f9e743058] WAF1_SPINNP_REPL_SEND_HOLES error response 45
> 2018-07-24T15:59:39Z 1171725196151878 [8:0] REPL_0:
repl_writer::PushFileUpdateHoleReplOpHandler::replXlateOut(): | [254b8fa9-8f4b-11e8-b4c4-
00a098a46c68] | [0xfffff80f9e743058] WAF1_SPINNP_REPL_SEND_HOLES failed on Sync SM dst due to
Punch hole beyond EOF/on internal files
> 2018-07-24T15:59:39Z 1171725196167961 [8:0] REPL_0:
repl_writer::ReplOpHandler::replOpHandlerCallback(): | [254b8fa9-8f4b-11e8-b4c4-00a098a46c68] |
[0xfffffff93ab8d40] msg:WAF1_SPINNP_REPL_SEND_HOLESstatus:45result:14op:29
```

The transition completes as a result of multiple autoresync attempts to bring the relationship status to InSync. If the transition still fails, avoid this workload until the relationship status is InSync.

## Common Snapshot Failures

Auto or manual resync fails with an error in the `snapmirror show` output.

```
sti8060-2671527862493::*> snapmirror show -fields source-path,destination-
path,type,state,status,last-transfer-error
source-path destination-path type policy state status last-transfer-error
-----
```

```
vs0:vol2 vs1:vol2 XDP StrictSync Snapmirrored OutOfSync SnapMirror Common Snapshot creation
operation failed when creating Snapshot copies "snapmirror.9b3cb6d0-959c-11e8-b04b-
00a098aa7373_2151051109.2018-08-13_032751" on source "vs0:vol2" and destination "vs1:vol2".
```

## Probable Cause

Verify that CP statistics are using a node-level command in diagnostic mode: `run local wafl cpstats -s -a <aggregate name>`.

## Resolution

- If this is a transient problem, wait for autoresync attempts to put the relationship status InSync.
- If this issue persists, verify that sizing is done appropriately to optimize the environment and workload. Perform manual resync to bring the relationship to InSync.

## Transition Failure Due to High Metadata

```
::*> snapmirror show -fields last-transfer-error
source-path destination-path last-transfer-error
-----
vs0:vol1 vs0:vol1dp Failed to start replication on source. (Ran out of resources to
queue metadata ops. Aborting transition. Retry the operation. If the error persists contact
technical support.)

::~*> snapmirror show-history -fields additional-info
destination-path operation-id additional-info
-----
vs0:vol1dp 8081726d-d8cc-11e8-8917-005056a7cdb4 Failed to start replication on source. (Ran
out of resources to queue metadata ops. Aborting transition. Retry the operation. If the error
persists contact technical support.)
```

## Resolution

Attempt the transition again by performing a manual resync.

## Transition Failure

```
::*> snapmirror show-history -fields out-of-sync-reason
source-path destination-path out-of-sync-reason
-----
vs0:vol1 vs0:vol1dp Transfer failed. (Out of memory)

::~*> snapmirror show -fields last-transfer-error
source-path destination-path last-transfer-error
-----
vs0:vol1 vs0:vol1dp Failed to configure source. (Failed to transition SnapMirror
Synchronous relationship to the "InSync" status. Use the "snapmirror resync" command to resync
the relationship to recover from the failure. (Replication engine error))

::~*> snapmirror show-history -fields additional-info
destination-path operation-id additional-info
-----
vs0:vol1dp 125275b5-d8d4-11e8-aed4-005056a717a1 Failed to configure source. (Failed to
transition SnapMirror Synchronous relationship to the "InSync" status. Use the "snapmirror
resync" command to resync the relationship to recover from the failure. (Replication engine
error))

sti46-vsिम-ucs520i% grep -i drl sktrace.log | grep REPL_0
2018-10-26T04:12:35Z 173333821799639 [0:0] REPL_0: repl_granular::SyncCache::startDRL(): |
[0xffffffff96abbc30] result: [status: 0 failure_msg: 0 failure_msg_detail: 0] Start DRL log
called for CG with UUID: b53225c5-d8cb-11e8-aed4-005056a717a1, Transfer ID: 30eb4bc4-d8ce-11e8-
8917-005056a7cdb4
2018-10-26T04:12:35Z 173333828263613 [1:0] REPL_0:
repl_granular::SyncCache::startDRLCompleted(): | [0xffffffff96abbc30] result: [status: 0
failure_msg: 0 failure_msg_detail: 0] Notifying d-control about start DRL completion for CG with
UUID: b53225c5-d8cb-11e8-aed4-005056a717a1, Transfer ID: 30eb4bc4-d8ce-11e8-8917-005056a7cdb4
```

```

2018-10-26T04:12:36Z 173334210432295 [1:0] REPL_0: repl_granular::SplitterDrls::getDRL():
| failed to get DRL
2018-10-26T04:12:36Z 173334210453541 [1:0] REPL_0:
granular::Splitter::_doDrlLoggingIfReq(): | [0xffffffff96blcdb8] result: [status: 10
failure_msg: 5898505 failure_msg_detail: 0] fileId=96 Failed to get DRL
2018-10-26T04:12:36Z 173334210473508 [1:0] REPL_0:
granular::Splitter::_doDrlLoggingIfReq(): | [0xffffffff96blcdb8] result: [status: 10
failure_msg: 5898505 failure_msg_detail: 0] fileId=96 DRL logging write failed for Offset: 21360
Length: 3
2018-10-26T04:12:36Z 173334210610453 [1:0] REPL_0: granular::Splitter::setCGOutOfSync(): |
[0xffffffff96blcdb8] result: [status: 10 failure_msg: 5898505 failure_msg_detail: 0] fileId=0
Setting CG out-of-sync. state: virtual const char
*repl_granular::Splitter::SplitterStateLogDrl::name() const

```

## Probable Cause

- DRL has a per-node limit of 60K file handles.
- There was an I/O burst during the last async transfer or a lengthy last async transfer consisting of I/O on a large number of files.

## Resolution

1. Check the type of workload.
2. If the number of active file handles is within the limit, perform a manual resync.

## Blacklisted Operations

Auto or manual resync fails with an error in `snapmirror show output`.

```

sti8060-2671527862493:.*> snapmirror show -fields source-path,destination-
path,type,state,status,last-transfer-error
source-path destination-path type policy state status last-transfer-error
-----
vs0:vol2 vs1:vol2 XDP StrictSync Snapmirrored OutOfSync Failed to start replication on source.
(Resync completion deferred due to a conflicting metadata operation.)
Respective source EMS event,
sti8060-2671527862493:.*> event log show -event *sms*
8/23/2018 17:25:58 sti8060-267 ERROR sms.operation.blacklisted: SnapMirror Sync operation
encountered a failure when relationship status is "Transitioning" for relationship UUID
"3548f219-a6cb-11e8-9548-0050568e99eb" because blacklisted operation "volume qtree security" was
performed in "SplitterStateLogDrl" state.

```

## Probable Cause

A volume or qtree security modify operation was attempted during the Transition state; this is not supported. Check the audit or message logs.

## Resolution

Perform an autoresync or manual resync to try to bring the relationship to InSync. Otherwise, do not modify the security style during a resync operation.

## Interoperability

Auto or manual resync fails with an error in the `snapmirror show output`.

```

cluster1:.*> snapmirror show -fields source-path,destination-path,type,state,status,last-
transfer-error
source-path destination-path type policy state status last-transfer-error
-----
vs0:vol2 vs1:vol2 XDP StrictSync Snapmirrored OutOfSync Prechecks on source volume failed.
Reason: "SnapMirror Synchronous operations are not supported on volume " vs0:vol2" because a file
or LUN copy operation is in progress on the volume.".)

```

## Probable Cause

A disallowed interoperability operation was performed on the volume involved in the SnapMirror relationship.

## Resolution

Remove or stop the disallowed interoperability configuration or operation and then retry the resync operation.

## 9 Limitations

- Snapshot symmetry: User, scheduled, or application Snapshot copies are not replicated.
- A SnapMirror Synchronous primary cannot be a destination of a synchronous or asynchronous relationship.
- Synchronous fan-out: A SnapMirror Synchronous source cannot be a source of another synchronous relationship.
- Systems with less than 16GB of memory are not supported
- DP\_Optimized (DPO) is not supported.
- Automatic failover is not supported. You must manually use SnapMirror commands to fail over the application from primary to secondary after break of the relationship.
- QoS throughput floors cannot be set for SnapMirror Synchronous source volumes.
- NetApp SnapLock®, FlexGroup, and FlexCache® are not supported.

## 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 Center  
<http://docs.netapp.com/ontap-9/index.jsp>
- SM-S Replication Advisor  
<https://fieldportal.netapp.com/content/835440>

## Version History

Version	Date	Document Version History
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 © 2019 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

### **Trademark Information**

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