

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

SnapMirror Synchronous Configuration and Best Practices
ONTAP 9.11.1



Table of Contents

1. Executive Summary	7
Purpose and Scope	7
Terminology	8
2. SnapMirror Synchronous Overview	9
Use Cases	9
Zero Data Loss	9
Disaster Recovery Between Two Local Data Centers	9
Combining Synchronous and Asynchronous Replication Between Two Data Centers	10
Application Disaster Recovery Among Three Data Centers	11
Integrated Data Protection	12
SnapMirror Synchronous Modes	13
3. SnapMirror Synchronous Data Protection Relationships	14
Prerequisites	14
Considerations Before Using SnapMirror Synchronous	15
SnapMirror Synchronous Licensing	15
Maximum Number of SnapMirror Synchronous Replication Operations	16
Restrictions Using SnapMirror Synchronous	16
SnapMirror Policies	17
Common Snapshot Copies	18
Relationship Status	19
Cluster and SVM Peering	19
Populate Source Volume with Data	21
Set Up a SnapMirror Synchronous Relationship	22
Set Up a SnapMirror StrictSync Relationship	24
View Destination Volume on the Secondary Storage System	25
Interrupt the Network Between Primary and Remote (Prevents Replication) and Observe the Behavior of the SnapMirror StrictSync Relationship	26
Interrupt the Network Between Source and Destination (Prevents Replication) and Observe the Behavior of SnapMirror Synchronous Mode	27
Failover from Primary to Secondary Data Center	28

Failover and Failback	29
Convert a SnapMirror Synchronous Relationship from StrictSync Mode to Synchronous	30
Convert a SnapMirror Asynchronous Relationship to Synchronous	32
Convert a Synchronous SnapMirror Relationship to Asynchronous	36
Flip Resync	37
Mirror State and Relationship Status	39
4. SnapMirror Synchronous and Storage Efficiency	40
5. Revert to Previous ONTAP Versions	41
6. Performance.....	42
7. Troubleshooting.....	43
Operations That Cause SnapMirror Relationships to Go OutOfSync	43
Unhealthy Cluster Peering	43
Nondisruptive Operation	44
Aggregate or Volume Goes Offline	46
Insufficient Space on Volume or Aggregate	48
SnapMirror Abort	49
Resync Failures	50
Unable to Meet Transition Criteria Because Resync Takes a Long Time	50
No Common Snapshot Copy Found	51
Transition Failures	52
Common Snapshot Failures	52
Transition Failure Due to High Metadata	53
Transition Failure	53
Blacklisted Operations	54
Interoperability	55
8. Limitations.....	56

List of Figures

Figure 1	Disaster recovery between two local data centers	9
Figure 2	Combining Synchronous and Asynchronous SnapMirror relationships.....	10
Figure 3	Three data center enterprise-application use case	11
Figure 4	Integrated data protection with SM-S.....	12

List of Tables

Table 1	Concurrent replication operations allowed	16
Table 2	Required time to create a common Snapshot per ONTAP version	18
Table 3	Relationship status.....	19
Table 4	SM-S relationship: source and destination values	23
Table 5	Mirror state and relationship status	39

Preface

This document contains information and best practices for configuring SnapMirror Synchronous (SM-S) replication in ONTAP 9.11.1.

Copyright 2022 Fujitsu Limited

Second Edition
December 2022

Trademarks

Third-party trademark information related to this product is available at:

<https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html>

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

About This Manual

Intended Audience

This manual is intended for system administrators who configure and manage operations of the ETERNUS AX/HX, or field engineers who perform maintenance. Refer to this manual as required.

Related Information and Documents

The latest information for the ETERNUS AX/HX is available at:

<https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html>

Document Conventions

■ Notice Symbols

The following notice symbols are used in this manual:

Caution

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

Note

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

1. 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.

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/writable.

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 features such as Snapshot technology, FlexVol volumes, and FlexClone volumes
- General knowledge of data protection, disaster recovery, and data replication solutions

Terminology

- **Primary or source**
The originating volume from which SM-S retrieves data.
- **Secondary or destination**
The targeted volume to which SM-S writes.
- **Recovery point objective (RPO)**
The amount of data loss your business application can tolerate.
- **Recovery time objective (RTO)**
The amount of time required to restart a failed application.
- **SnapMirror**
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 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.
- **SM-S mode**
Provides zero RPO replication with no primary I/O restriction if a replication failure occurs.
- **SnapMirror Strict Synchronous (StrictSync) mode**
Provides continuous primary I/O even if 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.
- **Round trip time (RTT)**
The duration in milliseconds (ms) it takes for a network request to go from a source to a destination and back again to the source.

2. SnapMirror Synchronous Overview

SM-S is volume-granular synchronous data replication that enterprises depend on for backup, disaster recovery, and data mobility. SM-S enables customers to achieve zero RPO by replicating data from FlexVol volumes between ONTAP storage systems situated across your data center or metro region. SM-S is targeted at workloads 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. 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.

Use Cases

Zero Data Loss

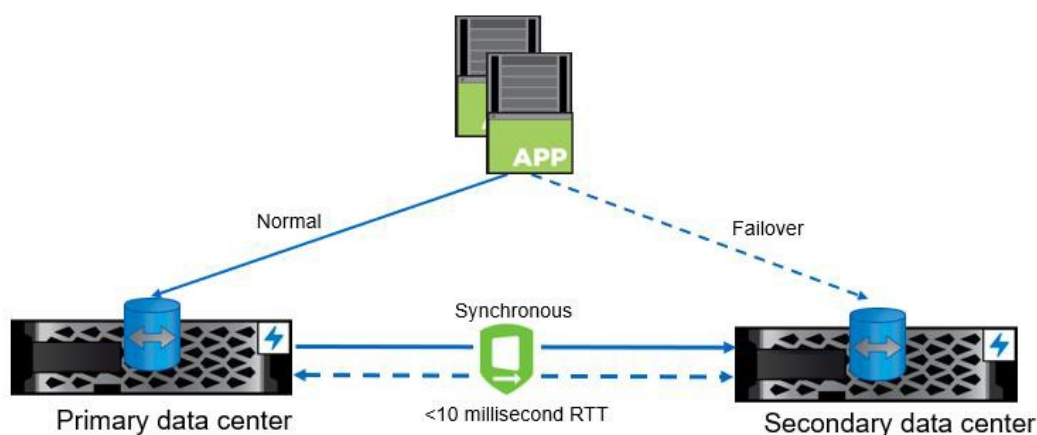
In today's digital world, a loss of vital business data can cripple a company, and any downtime can have serious repercussions for the organization's financial performance and competitive advantage. You can use SM-S replication software to safeguard data from loss due to natural disaster, fire, application 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.

Disaster Recovery Between Two Local Data Centers

This use case includes two data centers for local disaster recovery with one synchronous relationship of a volume to a relatively local secondary data center (see [Figure 1](#)). If the primary site fails, the enterprise application fails over to the replica volume in the secondary data center.

Figure 1 Disaster recovery between two local data centers



Combining Synchronous and Asynchronous Replication Between Two Data Centers

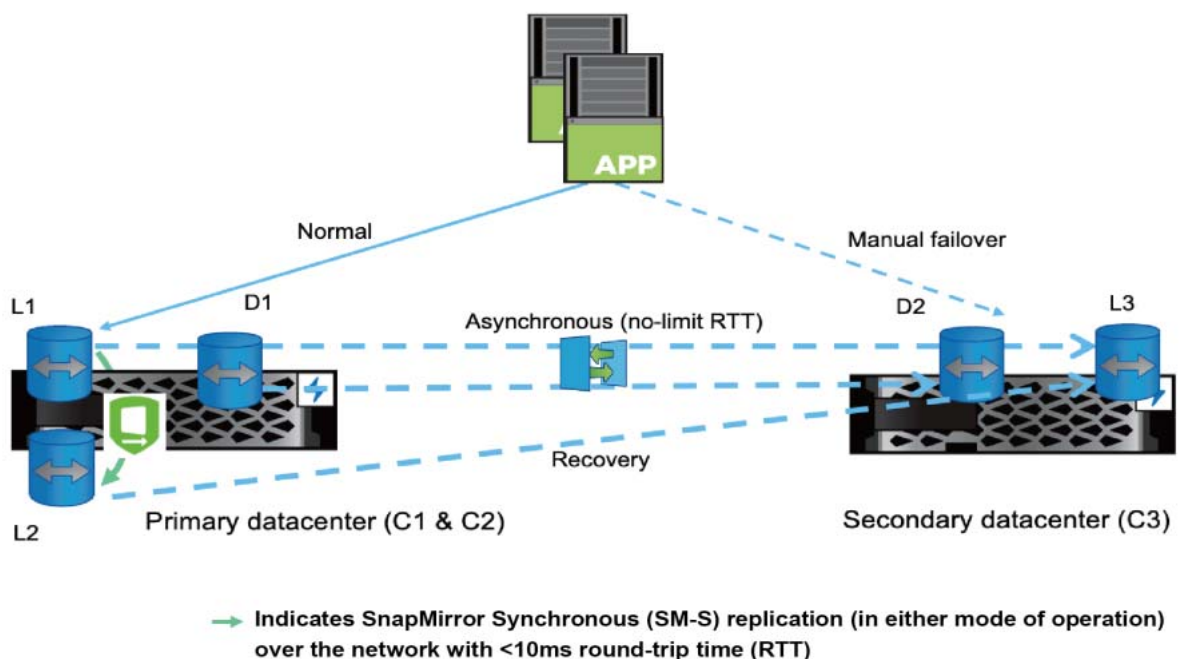
This use case includes two data centers for an enterprise application with one synchronous relationship for transaction logs to another cluster or volume in the primary data center, and one asynchronous relationship for regular data to the secondary data center. If a primary cluster fails, the application fails over to the secondary data center, which recovers and synchronizes with data from the transaction logs in the primary data center, as shown in [Figure 2](#).

- **C1**
Primary cluster where D1 (data) and L1 (log) are hosted
- **C2**
Primary cluster hosting the log volume's synchronous replica (L2)
- **C3**
Secondary data centre that hosts L3 and D2, the asynchronous replicas of log volume and data volume, respectively.
- When there is a crash at C1, SnapMirror break is issued on L2 and L3, respectively. L2 is mounted at the surviving site. The difference between L2 and L3 is applied onto L3. Run a log recovery onto D2 by using L3. After D2 is ready, the application can be served by using <D2, L3>.

Caution

While deploying SM-S for databases, the logs should be in one volume whereas the data files can be spread across multiple volumes.

Figure 2 Combining Synchronous and Asynchronous SnapMirror relationships



Application Disaster Recovery Among Three Data Centers

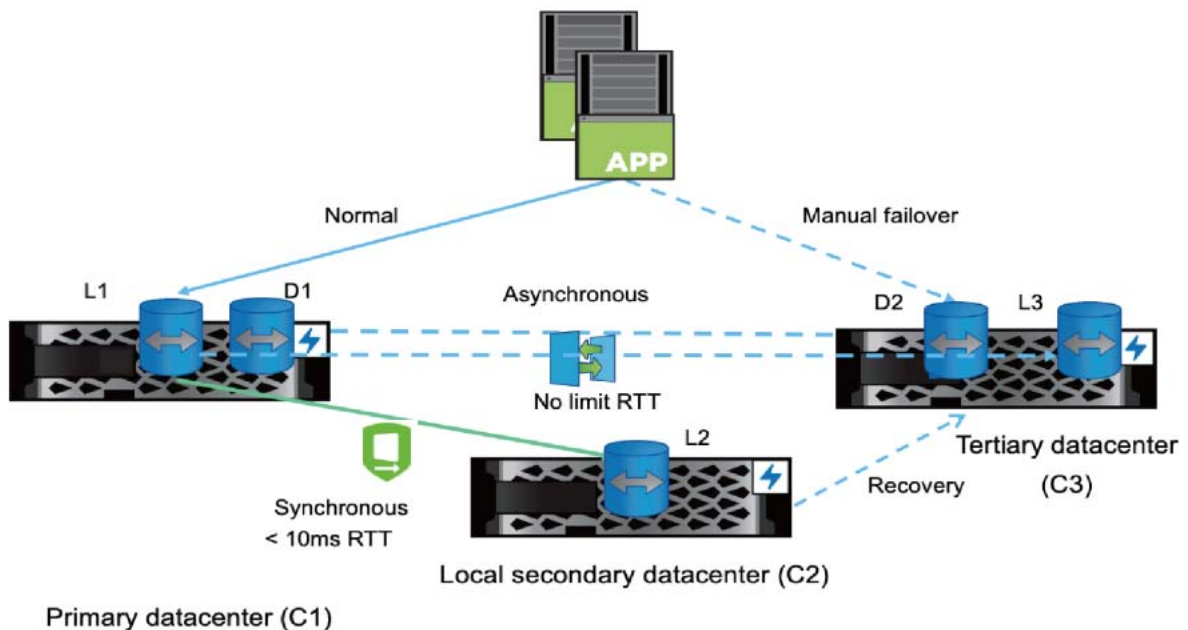
This use case includes three data centers for an enterprise application with one synchronous relationship for transaction logs in a relatively local secondary data center, and one asynchronous relationship for regular data to a more remote tertiary data center. If the primary site fails, the application fails over to the remote tertiary data center, which recovers and synchronizes with data from the transaction logs in the local secondary data center, as shown in [Figure 3](#).

- **C1**
Primary cluster where D1 (data) and L1 (log) are hosted
- **C2**
Secondary cluster hosting the log volume's synchronous replica (L2).
- **C3**
Tertiary data centre that hosts L3 and D2, the asynchronous replicas of log volume and data volume, respectively.
- When there is a crash at C1, SnapMirror break is issued on L2 and L3, respectively. L2 is mounted at the surviving site. The difference between L2 and L3 is applied onto L3. Run a log recovery onto D2, by using L3. After D2 is ready, the application can be served by using <D2, L3>.

Caution

While deploying SM-S for databases, the logs should be in one volume whereas the data files can be spread across multiple volumes.

Figure 3 Three data center enterprise-application use case



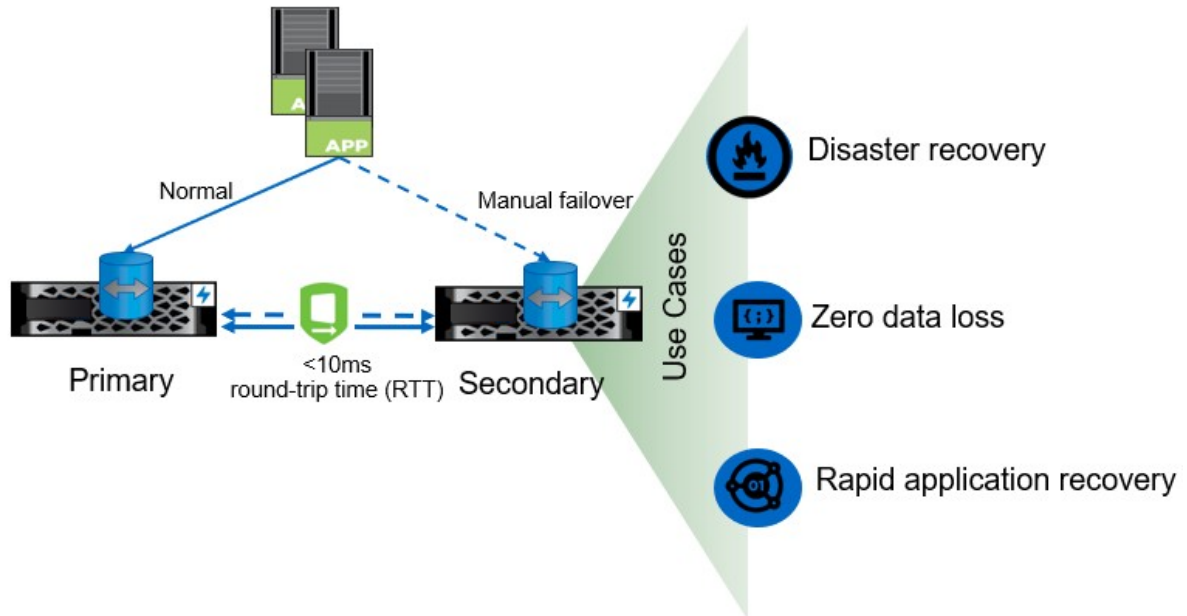
■ Best Practice

To make sure that a common Snapshot copy is available between the synchronous and asynchronous relationships, set up a SnapMirror Asynchronous relationship with the MirrorAllSnapshots policy and a schedule more frequent than the synchronous relationships' common Snapshot schedule.

Integrated Data Protection

SM-S tightly integrates with Snapshot technology to create common Snapshot copies, enabling quick and easy resynchronization between data centers. SM-S provides efficient data protection across Fujitsu platforms with native data format of secondary data copies. This kind of protection provides multiuse for disaster recovery, zero data loss, and rapid application recovery as illustrated in [Figure 4](#).

Figure 4 Integrated data protection with SM-S



Application-consistent Snapshot support starting with ONTAP 9.7: if a Snapshot copy is tagged with the appropriate label at the time of the snapshot create operation by using the CLI or the ONTAP API, SnapMirror-Synchronous replicates the Snapshot copies (user created or created with external scripts after quiescing the applications).

■ Best Practice

Fujitsu recommends that you do not make changes by adding rules to default policies and create custom policies. Making changes to default policies have a global effect, which might not be desirable.

SnapMirror Synchronous Modes

You can configure SM-S in either of the following two modes.

- **Synchronous mode**

In Synchronous mode, application I/O operations are sent in parallel to the primary and secondary storage systems. If the write to the secondary storage system does not complete for any reason, the application is allowed to continue writing to the primary storage system. When the error condition is corrected, SM-S technology automatically resynchronizes with the secondary storage system and then resumes replicating synchronously from the primary storage to the secondary storage system.

This process provides an RPO equal to zero.

- **StrictSync mode**

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

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

3. SnapMirror Synchronous Data Protection Relationships

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

Prerequisites

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

- SM-S technology is supported on all ETERNUS AX/HX series and ETERNUS AX series All SAN Array (ASA) systems that have at least 16GB of memory.

- To check the system memory configuration, use the following commands:

```
cluster::*> set -privilege advanced
cluster::*> run local sysconfig
      slot 0: System Board
      Model Name:      ETERNUS AX4100
      Processors:      20
      Processor type:   Intel(R) Xeon(R) Silver 4210 CPU @ 2.20GHz
      Memory Size:     147328 MB
```

- Each node in the cluster must be running ONTAP 9.7 or later.

- To check the system version, use this command:

```
cluster::> version
Data ONTAP Release 9.11.1P1: Tue Aug 09 13:13:19 UTC 2022 (Fujitsu)
```

- The functional network to transfer data between two different storage systems should have an RTT latency of less than 10ms.

- The following is the CLI command to verify network latency between source C1 and destination C2 clusters.

```
C1_stiA400-355_cluster::*> cluster peer ping -originating-node stiA400-355
  -destination-cluster C2_stiA400-357_cluster -destination-node stiA400-357

Node: stiA400-355          Destination Cluster: C2_stiA400-357_cluster

Destination Node IP Address      Count   TTL RTT(ms) Status
-----
stiA400-357      172.20.37.240      1    64    0.075 interface_reachable
```

- Cluster and storage virtual machine (SVM) peer relationships must be configured.

- See [configure peer relationship](#) document for more information.

- Storage system names should be correctly resolved by DNS.

- To check DNS entry, use the following command:

```
C1_sti92-vsim-ucs532m_cluster::*> getXXbyyy gethostbyname
-node sti92-vsim-ucs532m -vserver C1_sti92-vsim-ucs532m_cluster
-hostname sti92-vsim-ucs532m

(vserver services name-service getxxbyyy gethostbyname)

Host name: sti92-vsim-ucs532m
Canonical name: sti92-vsim-ucs532m.ctl.gdl.englab.netapp.com
IPv4: 10.237.140.174
```

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.
- The slower among the primary and secondary will impact the latency.
 - While the drive technology (FC or SATA, flash) underlying the source and destination volumes can vary, the least performant drive technology determines the I/O characteristics.
 - The kind of system, the CPU headroom, and the drive configuration on the secondary system affect the latency of the workloads. Therefore, the destination system should be sized appropriately when volumes are SM-S protected.
- The network transport should be optimized for SM-S replication. Fujitsu recommends a dedicated, high-bandwidth, low-latency network between the source and destination systems. Otherwise, replication performance and hence user workload performance could be affected.
- The source volume can reside in a FabricPool-enabled aggregate with any tiering policy, including None, Snapshot, or Auto. However, the destination volume in a FabricPool aggregate cannot be set to All tiering policy.

SnapMirror Synchronous Licensing

- With the ETERNUS AX/HX, an SM-S license is included with the Data Protection bundle.
- Starting in ONTAP 9.9.1, SM-S is automatically enabled on systems where the Data Protection bundle has been installed.

Caution

The license must be present on both source and destination.

Maximum Number of SnapMirror Synchronous Replication Operations

Table 1 Concurrent replication operations allowed

SM-S limits per platform	ONTAP 9.8 or earlier (per HA pair)	ONTAP 9.9.1 (per HA pair)	ONTAP 9.10.1 (per HA pair)	ONTAP 9.11.1 (per HA pair)
ETERNUS AX series	80	160	200	400
ETERNUS HX series	40	80	80	80

Caution

Above limits are total operations per HA pair - all could be source, or destination or in any combination of source or destination.

The updated scale is applied only if all the nodes are running ONTAP 9.11.1 or later.

These limits apply to both the surviving node and the failed-over node of an HA pair. If a node failover occurs during data backup, the transfers running at the time of the failure on the failed node are aborted and are rescheduled by the surviving node. Conversely, if a giveback occurs during data backup, the transfers running at the time of the giveback on behalf of the surviving node are aborted and are rescheduled by the restored 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.

Restrictions Using SnapMirror Synchronous

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

- The source and destination volumes must be online.
- The capacity of the destination volume must be greater than or equal to the total space used on the source volume.
- The 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.
- RTT should be less than 10milliseconds.

SnapMirror Policies

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

- Sync
- StrictSync

```
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      -      -
```

Common Snapshot Copies

SnapMirror-Synchronous periodically drains the I/O, which creates a common Snapshot copy on both the primary and the secondary storage systems during in-flight between these systems. This Snapshot copy is used to resynchronize the relationship if it fell out of sync. At any given point in time, a maximum of two common Snapshot copies are maintained.

Beginning in ONTAP 9.9.1, the following features are available:

- By default, a common Snapshot copy is created every 6 hours.
- You can choose to modify the common Snapshot copy frequency to anywhere between 30 minutes and 24 hours. To modify the common Snapshot copy frequency, you must create a cron schedule and modify the Sync policy and attach the newly created schedule. For example:

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

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

Note

The common Snapshot copy workflow time is optimized in ONTAP 9.11.1. The test results are shown below:

Table 2 Required time to create a common Snapshot per ONTAP version

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

Fujitsu 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          6-hourly
Primary::*> snapmirror policy show -vserver vs1 -policy StrictSync -fields -common-snapshot-
schedule
vserver policy common-snapshot-schedule
-----
vs1          StrictSync    6-hourly
```

Caution

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
```

Relationship Status

[Table 3](#) describes the different relationship states that can occur between the source and destination volumes.

Table 3 Relationship status

Relationship Status	Details
Idle	<ul style="list-style-type: none">• No transfer is in progress.• If health is true, this is a newly created relationship.• If health is false, auto-resync has terminated after all default attempts have failed.
Quiescing	Synchronous replication suspension is in progress.
Quiesced	Synchronous replication is suspended.
Transferring	Asynchronous transmission phase is in progress when a synchronous relationship is being established or reestablished.
Transitioning	<ul style="list-style-type: none">• Transformation from asynchronous to synchronous replication has started.• Incoming ops are executed on primary and applied to secondary after the last asynchronous transfer.
InSync	Incoming I/O successfully applied to primary and secondary, and the replication path is active.
OutOfSync	<ul style="list-style-type: none">• Destination volume is not in sync with the source volume because SnapMirror replication is not happening due to some failure like network. If the mirror state is Snapmirrored, this indicates a transfer failure or failure due to an unsupported operation.• SM-S monitors the relationship every 5 minutes and attempts an autoresync operation up to 5 times.

Cluster and SVM Peering

You must perform the one-time operation of creating [cluster peer and SVM peer relationships](#) 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.

To view cluster details, complete the following steps:

Procedure ►►►

1 View the source cluster details.

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

2 View the destination cluster details.

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

3 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
```

Caution

In the Availability field, the cluster shows as Unavailable if the cluster is not reachable or if there is network congestion between the source and destination clusters.

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

4 Calculate the round-trip time (RTT) for SM-S by 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
```



Populate Source Volume with Data

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

Procedure ►►►

- 1 Note that this example talks about NAS; however, SM-S supports other protocols as well - NVMe/FC, NVMe/TCP, FC, iSCSI, CIFS/SMB 2.0+, NFSv3, NFSv4.0, and NFSv4.1. Use the `volume create` command to create a volume on the primary cluster.

```
Primary::> vol create -vserver vs1 -volume voll -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
```

- 2 Populate the source volume with data.

```
[root@client~]# mount | grep 10.236.130.54
10.236.130.54:/voll on /root/mount type nfs
(rw,relatime,vers=3,rsize=65536,wsiz=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:/voll 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 voll -aggregate aggr1_Remote_01
-size 10GB -type DP
[Job 81] Job succeeded: Successful
```



Set Up a SnapMirror Synchronous Relationship

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

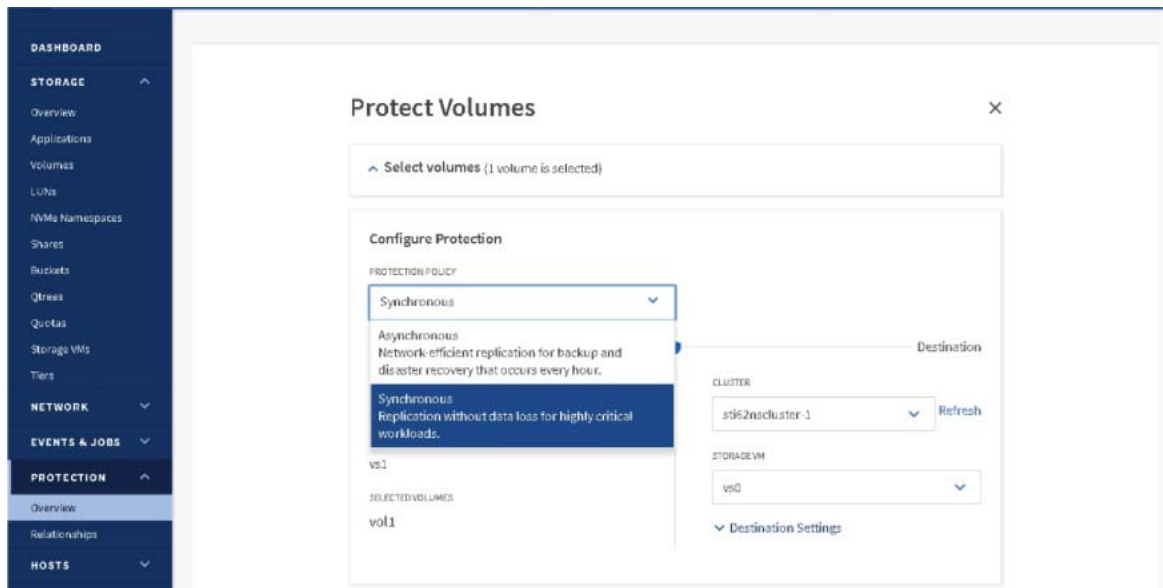
Procedure ►►►

- 1 Create an SM-S relationship between volumes on the primary and remote clusters.

Caution

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

In System Manager (for ONTAP 9.8 and later), the Synchronous listed in volume protect workflows is a Sync policy. To protect by using the StrictSync policy, navigate to Protection > Overview > Local Policy Settings > Protection Policies > Add.



An SM-S 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 SM-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.
```

Caution

You can set up an SM-S relationship between the two nodes of an HA pair for local protection of your data.

- 2 Initialize the SnapMirror relationship. The first time you do this, there is a baseline transfer from the source to the destination volume. 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.

3. SnapMirror Synchronous Data Protection Relationships Set Up a SnapMirror Synchronous Relationship

- 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.

```
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
```

- Do 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.

3 Monitor the SM-S relationship. You can monitor the data transfer status of an SM-S relationship on the source or destination by using the following commands, which provide details of any possible errors or unhealthy states.

Table 4 SM-S relationship: source and destination values

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 error 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-d046-11e8-9a21-0050568edde
72 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))
```



Set Up a SnapMirror StrictSync Relationship

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

Procedure ►►►

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 StrictSync 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.
```

Caution

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.
```

Caution

If the I/O to the secondary storage system 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 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- 005056aca11e_2157376758.2018-
09-10_034839
```


6 You can choose to assign a higher priority for Synchronous replication traffic by 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)
```



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. This situation occurs as an SM-S destination volume is redirected to the latest exported Snapshot copy, which is updated only once every six hours by default. However, if you wish to access the changed data immediately at destination for test/verification purposes, you can manually initiate the SnapMirror update.

How often an SM-S relationship updates the exported Snapshot copy is controlled by the SnapMirror policy's `-common-snapshot-schedule` parameter. You can modify the `-common-snapshot-schedule` parameter to a minimum of 30 minutes and a maximum of 24 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 mode.

```
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 SM-S relationship:

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

Interrupt the Network Between Primary and Remote (Prevents Replication) and Observe the Behavior of the SnapMirror StrictSync Relationship

Procedure ▶▶▶

1 Bring down the intercluster LIFs.

```
Primary::> net int show -role intercluster
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
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)
Vserver      Logical      Status      Network      Current      Current      Is
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 StrictSync.

```
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      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
```

Caution

In StrictSync mode, when you cannot perform replication, client application read or modification of a source might fail to provide zero data loss for a zero RPO.

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

Procedure ►►►

- 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 SM-S.

```
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 -
```

- 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@cli-
ent ~]# more /root/mount/realxed_outage.foo
sync_behavior_during_network_outage
```

Caution

Because SM-S is configured in Synchronous mode, client application modification of the source does not fail. However, replication to the destination volume does not occur and zero RPO cannot be achieved.

Failover from Primary to Secondary Data Center

When the source volume of a SnapMirror relationship is disabled, you must permanently end a SnapMirror relationship between a source and destination volumes.

To failover the application, follow these steps.

Procedure ►►►

1 Run the `snapmirror quiesce` command.

The quiesce 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.

Caution

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.

Caution

Recovery in a three-site data center scenario includes SnapMirror quiesce and break-off:

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

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



Failover and Failback

To perform failover and failback, follow these steps.

Procedure ►►►

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.

3 Create a reverse relationship.

```
Primary::> snapmirror create -source-path vs1_sync_dr:vol1 -destination-path vs1:vol1 -policy Strict-Sync  
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 like `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  -
```

Caution

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 vol1 false
```

Convert a SnapMirror Synchronous Relationship from StrictSync Mode to Synchronous

To convert a SnapMirror relationship from StrictSync to Synchronous, 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.

Caution

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

3. SnapMirror Synchronous Data Protection Relationships

Convert a SnapMirror Synchronous Relationship from StrictSync Mode to Synchronous

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

Procedure ►►►

1 Run the `snapmirror quiesce` command.

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

2 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.
```

3 Run the `snapmirror release` command with `-relationship-info-only true` option.

```
Remote:*> snapmirror release -relationship-info-only true -destination-path vs1_dr:vol1
[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:vol1 -destination-path vs1_dr:vol1 -policy Sync
Operation succeeded: snapmirror create the relationship with destination vs1_sync_dr:vol1.
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 Idle      -      false      -

Remote:> snapmirror show -fields policy-type,policy
source-path destination-path policy-type      policy
-----
vs1:vol1 vs1_sync_dr:vol1 sync-mirror      Sync
```

5 Run the `snapmirror resync` command.

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

Warning:      All data newer than Snapshot copy snapmirror.3fd9730b-8192-11e2-9caa-
123478563412_2147484699.2013-02-28_1 10732 on volume vs1_sync_dr:vol1 will be deleted.

Verify there is no XDP relationship whose source volume is
"vs1_sync_dr:vol1 ". If such a relationship exists then you are creating an unsupported XDP to XDP cas-
cade.

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      -
```

Convert a SnapMirror Asynchronous Relationship to Synchronous

You cannot convert a SnapMirror Asynchronous 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.

Procedure ►►►

1 Create a SnapMirror Asynchronous 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.
```


3. SnapMirror Synchronous Data Protection Relationships Convert a SnapMirror Asynchronous Relationship to Synchronous

2 View the progress of the SnapMirror Asynchronous 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 SnapMirror Asynchronous relationship.

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

3. SnapMirror Synchronous Data Protection Relationships

Convert a SnapMirror Asynchronous Relationship to Synchronous

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
```



Convert a Synchronous SnapMirror Relationship to Asynchronous

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

Procedure ►►►

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.
```

Caution

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
```

Flip Resync

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

Procedure ►►►

1 Before deletion, quiesce the existing relationship.

```
Remote::~> snapmirror quiesce -destination-path cnodes1_vsl:dstvol_1
Operation succeeded: snapmirror quiesce for destination "cnodes1_vsl:dstvol_1".

Remote::~> snapmirror show
```

Source Path	Destination Type Path	Mirror State	Relationship Status	Total Progress	Quiesced	Last Healthy	Updated
cnodes0_vsl:srcvol_1	XDP	cnodes1_vsl:dstvol_1	Snapmirrored	Quiesced	- false	-	-
cnodes0_vsl:srcvol_2	XDP	cnodes1_vsl:dstvol_2	Snapmirrored	Insync	- true	-	-

2 entries were displayed.

```
Remote::~> snapmirror quiesce -destination-path cnodes1_vsl:dstvol_2
Operation succeeded: snapmirror quiesce for destination "cnodes1_vsl:dstvol_2".

Remote::~> snapmirror show
```

Source Path	Destination Type Path	Mirror State	Relationship Status	Total Progress	Quiesced	Last Healthy	Updated
cnodes0_vsl:srcvol_1	XDP	cnodes1_vsl:dstvol_1	Snapmirrored	Quiesced	- false	-	-
cnodes0_vsl:srcvol_2	XDP	cnodes1_vsl:dstvol_2	Snapmirrored	Quiesced	- false	-	-

2 entries were displayed.

2 Delete the existing relationship.

```
Remote::~> snapmirror delete -destination-path cnodes1_vsl:dstvol_1
Operation succeeded: snapmirror delete for the relationship with destination "cnodes1_vsl:dstvol_1".

Remote::~> snapmirror delete -destination-path cnodes1_vsl:dstvol_2
Operation succeeded: snapmirror delete for the relationship with destination "cnodes1_vsl:dstvol_2".

Remote::~> snapmirror show
This table is currently empty.
```

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

```
Remote::~> snapmirror release -destination-path cnodes1_vsl:dstvol_1 -relationship-info-only true [Job 60] Job succeeded: SnapMirror Release Succeeded

Primary::~> snapmirror release -destination-path cnodes1_vsl:dstvol_2 -relationship-info-only true [Job 61] Job succeeded: SnapMirror Release Succeeded
```

4 Create the reverse relationship.

```
Primary:.*> snapmirror create -source-path cnodes1_vsl:dstvol_1 -destination-path  
cnodes0_vsl:srcvol_1 -policy StrictSync
```

```
Warning: You are creating 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_vsl:srcvol_1"
```

```
Primary:.*> snapmirror create -source-path cnodes1_vsl:dstvol_2 -destination-path  
cnodes0_vsl:srcvol_2 -policy Sync
```

```
Operation succeeded: snapmirror create for the relationship with destination  
"cnodes0_vsl:srcvol_2"
```

```
Warning: You are creating 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  
application such as databases, VMWare, etc.
```

5 Resync the relationship.

```
Primary:.*> snapmirror resync -destination-path cnodes0_vsl:srcvol_1
```

```
Warning: All data newer than Snapshot copy  
snapmirror. 3811aa03-d894-11e8-aef9-00a0985a6a80_2147894256.2018-10-26_0  
22311 on volume cnodes0_vsl:srcvol_1 will be deleted.
```

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

```
Operation is queued: initiate snapmirror resync to destination "cnodes0_vsl:srcvol_1".
```

```
Primary:.*> snapmirror resync -destination-path cnodes0_vsl:srcvol_2
```

```
Warning: All data newer than Snapshot copy  
snapmirror. 3811aa03-d894-11e8-aef9-00a0985a6a80_2147894257.2018-10-26_0  
22839 on volume cnodes0_vsl:srcvol_2 will be deleted.
```

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

```
Operation is queued: initiate snapmirror resync to destination "cnodes0_vsl:srcvol_2".
```



Mirror State and Relationship Status

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

Table 5 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
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. SnapMirror Synchronous and Storage Efficiency

SM-S does not maintain storage efficiency over the wire; however, storage efficiency is maintained for data at rest on both the primary and the secondary storage systems.

5. Revert to Previous ONTAP Versions

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

Procedure ►►► ---

- 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.



Caution

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.

6. Performance

If data loss is something your business cannot tolerate, then SM-S is a solution you must consider. As with any zero RPO solution, SM-S also increases the latency of the application it is protecting. The factors affecting the latency/performance of SM-S protected workloads include network RTT, storage platforms, percentage of modify operations, I/O profile of the application, data amount, and frequency of common Snapshot copies.

7. Troubleshooting

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-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 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-vsim-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-Uid=72a3c715-173b-40b5-9c2c-f20f1a68dd1d 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-Uid=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-Uid=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-Uid=72a3c715-173b-40b5-9c2c-f20f1a68dd1d 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-Uid=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-Uid=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-vsim-ucs542i_cluster:> cluster peer show
Peer Cluster Name Cluster Serial Number Availability Authentication
-----
C2_sti97-vsim-ucs542k_cluster 1-80-000011 Partial ok

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

■ Resolution

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

Nondisruptive Operation

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.

7. Troubleshooting

Operations That Cause SnapMirror Relationships to Go OutOfSync

```
C2_sti97-vsim-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-vsim-ucs542i_cluster::> volume move start -volume Vol1 -destination-aggregate
sti97_vsim_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_vsim_ucs542j_aggr1". Use the
"volume move show -vserver vs0 -volume Vol1" command to view the status of this operation.

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

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

Vserver Name: vs0
Volume Name: Vol1
Actual Completion Time: -
Bytes Remaining: 148KB
Destination Aggregate: sti97_vsim_ucs542j_aggr1
Detailed Status: Volume move job locking down volumes for final transfer - This where mutual exclu-
sive comes in and relation goes OutOfSync.
Estimated Time of Completion: Wed Oct 24 06:16:55 2018
Managing Node: sti97-vsim-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.
```

Caution

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-Uid= 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-Uid=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-Uid=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-Uid=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-Uid=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-Uid=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

Autoresync 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_PriVol1" and
destination volume "vs0:SE_PriVol1_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 desti-
nation does not generate an OutOfSync EMS this could be a transient failure during a resync opera-
tion.))".
```

■ 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

Procedure ►►►

- 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 autosync to complete to bring the relationship status back to InSync.

- 3 If the autoresync attempts fail because wafiron 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:vol3 vs2:vol3 Prechecks on source volume failed. (Wafiron has fixed inconsistencies on the
volume. Run "snapmirror resync". Ensure that snapmirror checker tool is run on the common snapshots.)
```

- 4 If wafiron 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

Procedure ►►►

- 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 drives to the aggregate by running the following command:

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

- 3 Wait for autosync to recover the SnapMirror relationship to status InSync.



SnapMirror Abort

An Event Management System (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_PriVol1" and destination volume "df811f1b-cd8e-11e8-b2c3- 00a098b9251d:SE_PriVol1_dst" with relationship UUID "cblb3a18-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.

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-Uid=bd49c7d7-6f62-4efa-96d3-c4bclc02ebf6 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-Uid=bd49c7d7-6f62-4efa-96d3-c4bclc02ebf6 Group=none Operation-Cookie=0 action=Info Autore-
sync 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-Uid=bd49c7d7-6f62-4efa-96d3-c4bclc02ebf6 Group=none Operation-Cookie=0 action=Info sm_rp-
m_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-Uid=bd49c7d7-6f62-4efa-96d3-c4bclc02ebf6 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-Uid=bd49c7d7-6f62-4efa-96d3-c4bclc02ebf6 Group=none Operation-Cookie=0 action=Info Autore-
sync 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

Procedure ►►►

- 1 Size the environment based on best practices and use quality of service (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-8060-03bINFORMATIONAL 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-8060-03bINFORMATIONAL 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-8060-03bINFORMATIONAL 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-8060-03bINFORMATIONAL 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

Procedure ►►►

- 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] WAFL_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] WAFL_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] |  
[0xfffffffff93ab8d40] msg:WAFL_SPINNP_REPL_SEND_HOLESstatus:45result:14op:29
```

The transition completes because 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 sup-
port.)
```

■ 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 SnapMir-
ror
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 tran-
sition 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-vsimg-ucs520i% grep -i dnl 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 1733334210432295 [1:0] REPL_0: repl_granular::SplitterDrls::getDRL(): | failed to
get DRL
2018-10-26T04:12:36Z 1733334210453541 [1:0] REPL_0:
granular::Splitter::doDrlLoggingIfReq(): | [0xffffffff96b1cdeb8] result: [status: 10 failure_msg:
5898505 failure_msg_detail: 0] fileId=96 Failed to get DRL
2018-10-26T04:12:36Z 1733334210473508 [1:0] REPL_0:
granular::Splitter::doDrlLoggingIfReq(): | [0xffffffff96b1cdeb8] 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 1733334210610453 [1:0] REPL_0:
granular::Splitter::setCGOutOfSync(): | [0xffffffff96b1cdeb8] result: [status: 10 failure_msg: 5898505
failure_msg_detail: 0] fileId=0 Setting CG out-of-sync. state: virtual const char *repl_granu-
lar::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 many files.

■ Resolution

Procedure ►►►

- 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 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. (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 per-  
formed 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.

8. Limitations

- Snapshot symmetry: User or Scheduled Snapshot copies are not replicated. However, Application created snapshot labeled as app-consistent will get replicated.
- A SnapMirror Synchronous primary cannot be a destination of a synchronous or asynchronous relationship.
- A fan-out and cascade configuration of SM-S is not supported. A SnapMirror Synchronous source cannot be a source of another synchronous relationship.
- 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.
- SnapLock, FlexGroup, and FlexCache are not supported.

Fujitsu Storage
ETERNUS AX series All-Flash Arrays,
ETERNUS HX series Hybrid Arrays
SnapMirror Synchronous Configuration and Best Practices
ONTAP 9.11.1

P3AG-5702-02ENZ0

Date of issuance: December 2022
Issuance responsibility: Fujitsu Limited

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

FUJITSU