



# SnapMirror active sync

ONTAP 9

NetApp  
May 18, 2024

# Table of Contents

- SnapMirror active sync ..... 1
- Introduction ..... 1
- Plan ..... 9
- Configure ..... 16
- Manage SnapMirror active sync and protect data ..... 31
- Troubleshoot ..... 45

# SnapMirror active sync

## Introduction

### SnapMirror active sync overview

SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. Manual intervention nor custom scripting are required to trigger a failover with SnapMirror active sync.

SnapMirror active sync is available beginning with ONTAP 9.9.1. SnapMirror active sync is supported on AFF clusters, All-Flash SAN Array (ASA) clusters, C-Series (AFF or ASA). Primary and secondary clusters must be of the same type: either ASA or AFF. SnapMirror active sync protects applications with iSCSI or FCP LUNs.

Beginning with ONTAP 9.15.1, SnapMirror active sync supports a [symmetric active/active capability](#), enabling read and write I/O operations from both copies of a protected LUN with bidirectional synchronous replication, enabling both LUN copies to serve I/O operations locally.



Beginning July 2024, content from technical reports previously published as PDFs has been integrated with ONTAP product documentation. The ONTAP SnapMirror active sync documentation now includes content from TR-4878: SnapMirror active sync.

### Benefits

SnapMirror active sync provides the following benefits:

- Continuous availability for business-critical applications.
- Ability to host critical applications alternately from primary and secondary site.
- Simplified application management using consistency groups for dependent write-order consistency.
- The ability to test failover for each application.
- Instantaneous creation of mirror clones without impacting application availability.
- The ability to deploy protected and non-protected workloads in the same ONTAP cluster.
- LUN identity remains the same, so the application sees them as a shared virtual device.
- The ability to reuse secondary clusters with flexibility to create instantaneous clones for application usage for dev-test, UAT or reporting purposes without impacting application performance or availability.

SnapMirror active sync allows you to protect your data LUNs, which enables applications to fail over transparently for the purpose of business continuity in the event of a disaster. For more information, see [Use cases](#).

### Key concepts

SnapMirror active sync utilizes consistency groups and the ONTAP Mediator to ensure your data is replicated and served even in the event of a disaster scenario. When planning your SnapMirror active sync deployment, it is important to understand the essential concepts in SnapMirror active sync and its architecture.

## Asymmetry and symmetry

SnapMirror active sync supports asymmetric and, beginning with ONTAP 9.15.1, symmetric active/active solutions. These options refer to how hosts access storage paths and write data. In an asymmetric configuration, data on the secondary site is proxied to a LUN. In a symmetric active/active configuration, both sites are able to access local storage for active I/O.

Symmetric active/active is optimized for clustered applications including VMware VMSc, Windows Failover Cluster with SQL, and Oracle RAC.

For more information, see [SnapMirror active sync architecture](#).

## Consistency group

A [consistency group](#) is a collection of FlexVol volumes that provide a consistency guarantee for the application workload that must be protected for business continuity.

The purpose of a consistency group is to take simultaneous snapshot images of multiple volumes, thus ensuring crash-consistent copies of a collection of volumes at a point in time. A consistency group ensures all volumes of a dataset are quiesced and then snapped at precisely the same point in time. This provides a data-consistent restore point across volumes supporting the dataset. A consistency group thereby maintains dependent write-order consistency. If you decide to protect applications for business continuity, the group of volumes corresponding to this application must be added to a consistency group so a data protection relationship is established between a source and a destination consistency group. The source and destination consistency must contain the same number and type of volumes.

## Constituent

An individual volume or LUN that is part of the consistency group protected in the SnapMirror active sync relationship.

## ONTAP Mediator

The [ONTAP Mediator](#) receives health information about peered ONTAP clusters and nodes, orchestrating between the two and determining if each node/cluster is healthy and running. ONTAP Mediator provides the health information about:

- Peer ONTAP clusters
- Peer ONTAP cluster nodes
- Consistency groups (which is the failover unit); for each consistency group, the following information is provided:
  - Replication state: Uninitialized, In Sync, or Out of Sync
  - Which cluster hosts the primary copy
  - Operation context (used for planned failover)

With this ONTAP Mediator health information, clusters can differentiate between distinct types of failures and determine whether to perform an automated failover. ONTAP Mediator is one of the three parties in the SnapMirror active sync quorum along with both ONTAP clusters (primary and secondary). To reach consensus, at least two parties in the quorum must agree to a certain operation.



Beginning in ONTAP 9.15.1, System Manager displays the status of your SnapMirror active sync relationship from either cluster. You can also monitor the ONTAP Mediator's status from either cluster in System Manager. In earlier releases of ONTAP, System Manager displays the status of SnapMirror active sync relationships from the source cluster.

## Planned failover

A manual operation to change the roles of copies in a SnapMirror active sync relationship. The primary sites becomes the secondary, and the secondary becomes the primary.

## Primary-first and primary bias

SnapMirror active sync uses a primary-first principle that gives preference to primary copy to serve I/O in case of a network partition.

Primary-bias is a special quorum implementation that improves availability of a SnapMirror active sync protected dataset. If the primary copy is available, primary-bias comes into effect when the ONTAP Mediator is not reachable from both clusters.

Primary-first and primary bias are supported in SnapMirror active sync beginning with ONTAP 9.15.1. Primary copies are designated in System Manager and output with the REST API and CLI.

## Automatic unplanned failover (AUFO)

An automatic operation to perform a failover to the mirror copy. The operation requires assistance from the ONTAP Mediator to detect that the primary copy is unavailable.

## Out of Sync (OOS)

When the application I/O is not replicating to the secondary storage system, it will be reported as **out of sync**. An out of sync status means the secondary volumes are not synchronized with the primary (source) and that SnapMirror replication is not occurring.

If the mirror state is `Snapmirrored`, this indicates a transfer failure or failure due to an unsupported operation.

SnapMirror active sync supports automatic resync, enabling copies to return to an InSync state.

Beginning with ONTAP 9.15.1, SnapMirror active sync supports [automatic reconfiguration in fan-out configurations](#).

## Uniform and non-uniform configuration

When creating protection with *symmetric active/active*, the process depends on your configuration:

- **Uniform host access** means that hosts from both sites are connected to all paths to storage clusters on both sites. Cross site paths are stretched across distance.
- **Non-uniform host access** means hosts in each site are connected only to the cluster in the same site. Cross-site paths and stretched paths aren't connected.



Uniform host access is supported for any SnapMirror active sync deployment; non-uniform host access is only supported for symmetric active/active deployments.

## Zero RPO

RPO stands for recovery point objective, which is the amount of data loss deemed acceptable during a given time period. Zero RPO signifies that no data loss is acceptable.


## Zero RTO

RTO stands for recovery time objective, which is the amount of time that is deemed acceptable for an application to return to normal operations non-disruptively following an outage, failure, or other data loss event. Zero RTO signifies that no amount of downtime is acceptable.

## SnapMirror active sync architecture

The SnapMirror active sync architecture enables active workloads on both clusters, where primary workloads can be served simultaneously from both clusters. Regulations for financial institutions in some countries require businesses to be periodically serviceable from their secondary data centers as well, called “Tick-Tock” deployments, which SnapMirror active sync enables.

The data protection relationship to protect for business continuity is created between the source storage system and destination storage system, by adding the application specific LUNs from different volumes within a storage virtual machine (SVM) to the consistency group. Under normal operations, the enterprise application writes to the primary consistency group, which synchronously replicates this I/O to the mirror consistency group.



Even though two separate copies of the data exist in the data protection relationship, because SnapMirror active sync maintains the same LUN identity, the application host sees this as a shared virtual device with multiple paths while only one LUN copy is being written to at a time. When a failure renders the primary storage system offline, ONTAP detects this failure and uses the Mediator for re-confirmation; if neither ONTAP nor the Mediator are able to ping the primary site, ONTAP performs the automatic failover operation. This process results in failing over only a specific application without the need for the manual intervention or

scripting which was previously required for the purpose of failover.

Other points to consider:

- Unmirrored volumes which exist outside of protection for business continuity are supported.
- Only one other SnapMirror asynchronous relationship is supported for volumes being protected for business continuity.
- Cascade topologies are not supported with protection for business continuity.

## ONTAP Mediator

ONTAP Mediator is installed in a third failure domain, distinct from the two ONTAP clusters. Its key role is to act as a passive witness to SnapMirror active sync copies. In the event of a network partition or unavailability of one copy, SnapMirror active sync uses Mediator to determine which copy continues to serve I/O, while discontinuing I/O on the other copy. There are three key components in this setup:

- Primary ONTAP cluster hosting the SnapMirror active sync primary CG
- Secondary ONTAP cluster hosting the mirror CG
- ONTAP Mediator

The ONTAP Mediator plays a crucial role in SnapMirror active sync configurations as a passive quorum witness, ensuring quorum maintenance and facilitating data access during failures. It acts as a ping proxy for controllers to determine liveness of peer controllers. Although the Mediator does not actively trigger switchover operations, it provides a vital function by allowing the surviving node to check its partner's status during network communication issues. In its role as a quorum witness, the ONTAP Mediator provides an alternate path (effectively serving as a proxy) to the peer cluster.

Furthermore, it allows clusters to get this information as part of the quorum process. It utilizes the node management LIF and cluster management LIF for communication purposes. It establishes redundant connections through multiple paths to differentiate between site failure and InterSwitch Link (ISL) failure. When a cluster loses connection with the ONTAP Mediator software and all its nodes due to an event, it is considered not reachable. This triggers an alert and enables automated failover to the mirror Consistency Group (CG) in the secondary site, ensuring uninterrupted I/O for the client. The replication data path relies on a heartbeat mechanism, and if a network glitch or event persists beyond a certain period, it can result in heartbeat failures, causing the relationship to go out-of-sync. However, the presence of redundant paths, such as LIF failover to another port, can sustain the heartbeat and prevent such disruptions.

To summarize, ONTAP Mediator is used for the following purposes:

- Establish a quorum
- Continuous availability via automatic failover (AUFO)
- Planned failovers (PFO)



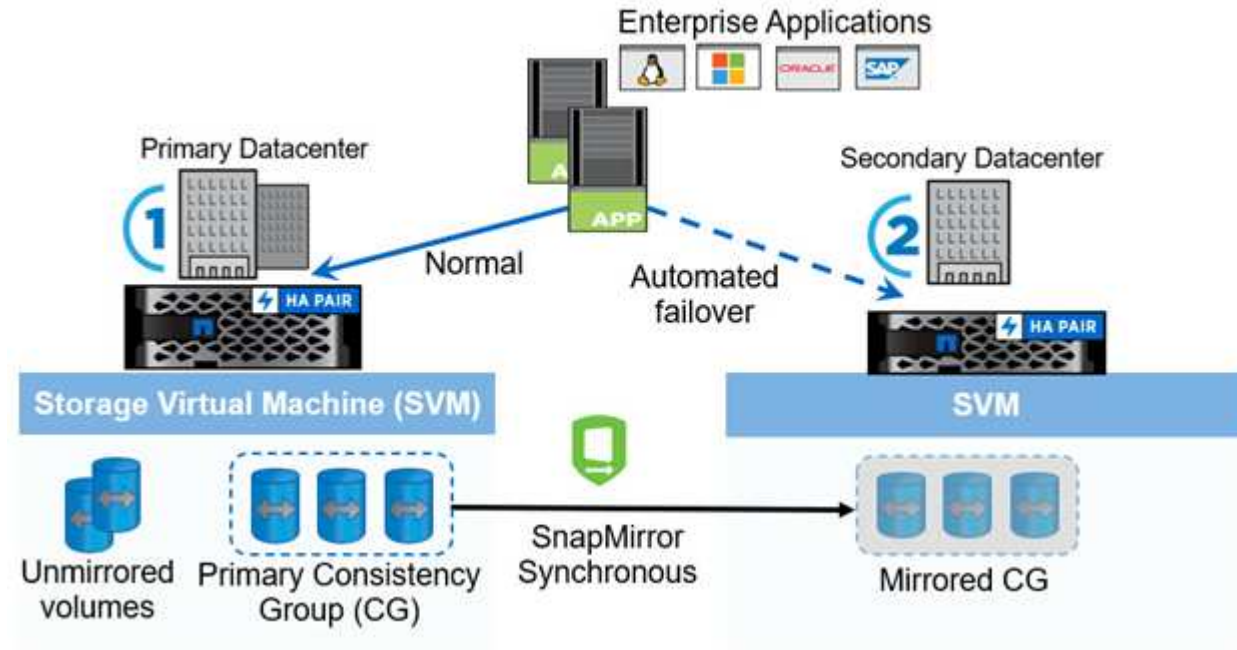
ONTAP Mediator 1.7 can manage ten cluster pairs for the purpose of business continuity.



When the ONTAP Mediator is not available, you cannot perform planned or automated failovers. The application data continues to synchronously replicate without any interruption to for zero data loss.

## Operations

The following figure illustrates the design of SnapMirror active sync at a high level.



The diagram shows an enterprise application that is hosted on a storage VM (SVM) at the primary data center. The SVM contains five volumes, three of which are part of a consistency group. The three volumes in the consistency group are mirrored to a secondary data center. In normal circumstances, all write operations are performed to the primary data center; in effect, this data center serves as the source for I/O operations, while the secondary data center serves as a destination.

In the event of a disaster scenario at the primary data center, ONTAP directs the secondary data center to act as the primary, serving all I/O operations. Only the volumes that are mirrored in the consistency group are served. Any operations pertaining to the other two volumes on the SVM is be affected by the disaster event.

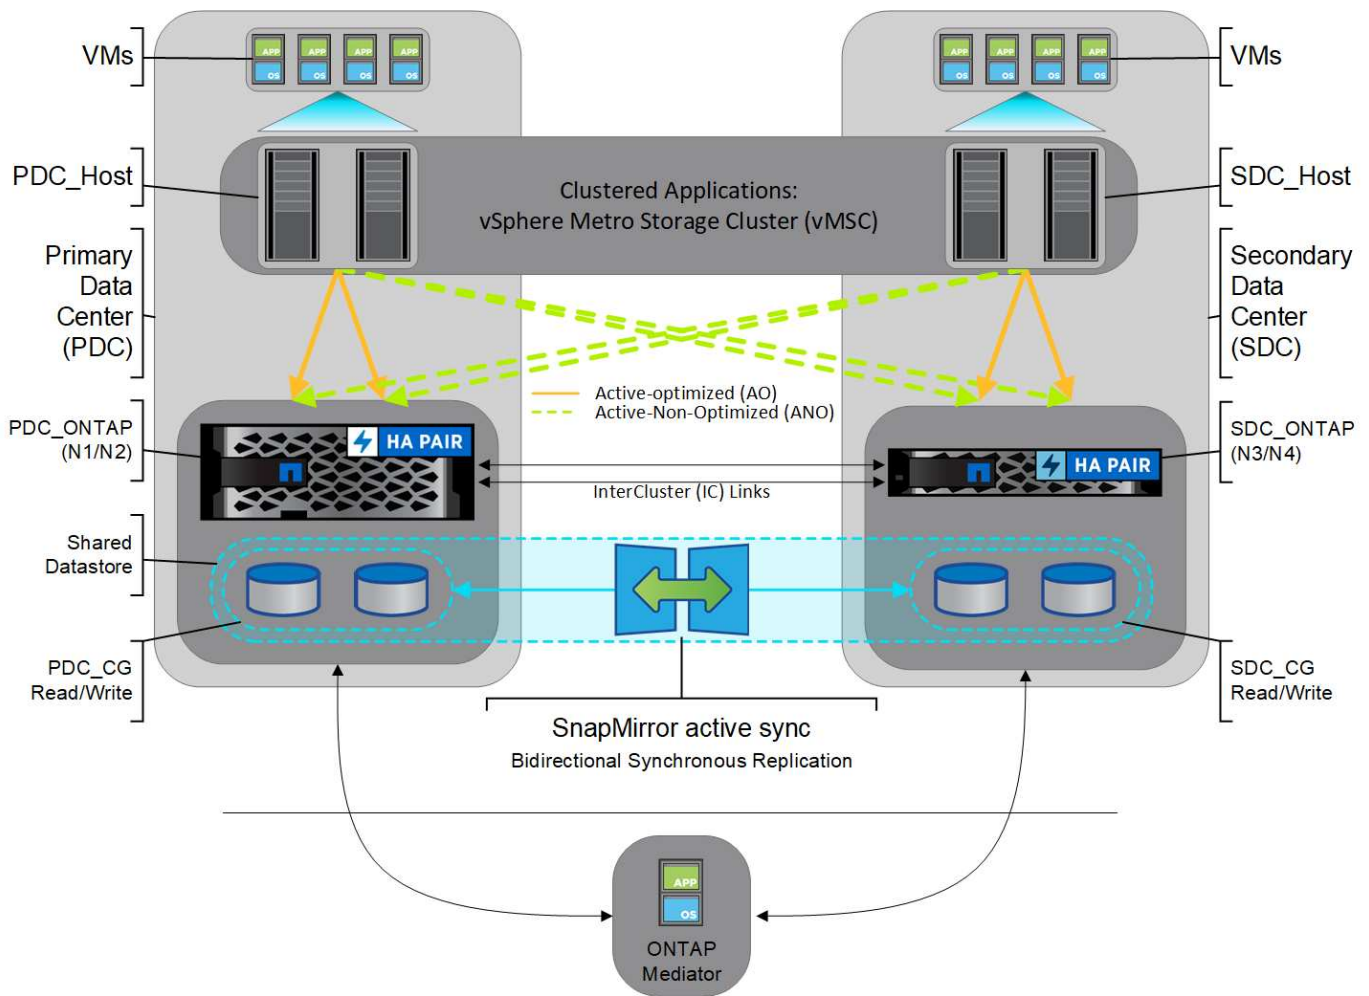
### Symmetric active/active

SnapMirror active sync offers asymmetric and symmetric solutions.

In *asymmetric configurations*, the primary storage copy exposes an active-optimized path and actively serves client I/O. The secondary site uses a remote path for I/O. The storage paths for the secondary site are considered active-non-optimized. Access to the write LUN is proxied from the secondary site.

In *symmetric active/active configurations*, active-optimized paths are exposed on both sites, are host specific, and are configurable, meaning hosts on either side are able to access local storage for active I/O.





Symmetric active/active is targeted for clustered applications including VMware Metro Storage Cluster, Oracle RAC, and Windows Failover Clustering with SQL.

## Use cases for SnapMirror active sync

The demands of a globally connected business environment demand rapid recovery of business-critical application data with zero data loss in the event of a disruption such as a cyber attack, power outage, or natural disaster. These demands are heightened in arenas such as finance and those adhering to regulatory mandates such as the General Data Protection Regulation (GDPR).

SnapMirror active sync provides the following use cases:

### Application deployment for zero recovery time object (RTO)

In a SnapMirror active sync deployment, you have a primary and secondary cluster. A LUN in the primary cluster (L1P) has a mirror (L1S) on the secondary; both LUNs share the same serial ID and are reported as read-write LUNs to the host. Read and write operations, however, are only serviced to the primary LUN, L1P. Any writes to the mirror L1S are served by proxy.

### Application deployment for zero RTO or TAF

TAF is based on host MPIO software-based path failover to achieve non-disruptive access to the storage. Both LUN copies—for example, primary (L1P) and mirror copy (L1S)—have the same identity (serial number) and are reported as read-writable to the host. However, reads and writes are serviced only by the primary volume.

I/Os issued to the mirror copy are proxied to the primary copy. The host's preferred path to L1 is VS1:N1 based on asymmetric logical unit access (ALUA) access state Active Optimized (A/O). ONTAP Mediator is required as part of the deployment, primarily to perform failover (planned or unplanned) in the event of a storage outage on the primary.

SnapMirror active sync uses ALUA, a mechanism that allows an application host multipathing software with paths advertised with priorities and access availability for the application host communication with the storage array. ALUA marks active optimized paths to the controllers owning the LUN and others as active non-optimized paths, used only if the primary path fails.

### **Clustered applications**

Clustered applications including VMware Metro Storage Cluster, Oracle RAC, and Windows Failover Clustering with SQL require simultaneous access so the VMs can be failed over to other site without any performance overhead. SnapMirror active sync symmetric active/active serves IO locally with bidirectional replication to meet the requirements of clustered applications.

### **Disaster scenario**

Synchronously replicate multiple volumes for an application between sites at geographically dispersed locations. You can automatically failover to the secondary copy in case of disruption of the primary, thus enabling business continuity for tier one applications. When the site hosting the primary cluster experiences a disaster, the host multipathing software marks all paths through the cluster as down and uses paths from the secondary cluster. The result is a non-disruptive failover enabled by ONTAP Mediator to the mirror copy.

### **Windows failover**

SnapMirror active sync provides flexibility with easy-to-use application-level granularity and automatic failover. SnapMirror active sync uses proven SnapMirror Synchronous replication over IP network to replicate data at high speeds over LAN or WAN, to achieve high data availability and fast data replication for your business-critical applications such as Oracle, Microsoft SQL Server, and so on, in both virtual and physical environments.


SnapMirror active sync enables mission-critical business services to continue operating even through a complete site failure, with TAF to the secondary copy. No manual intervention or no additional scripting are required to trigger this failover.

## **Deployment strategy and best practices for SnapMirror active sync**

It is important that your data protection strategy clearly identifies the workloads threats need to be protected for business continuity. The most critical step in your data protection strategy is to have clarity in your enterprise application data layout so that you can decide how you are distributing the volumes and protecting business continuity. Because failover occurs at the consistency group level on a per-application basis, make sure to add the necessary data volumes to the consistency group.

### **SVM configuration**

The diagram captures a recommended storage VM (SVM) configuration for SnapMirror active sync.



- For data volumes:
  - Random read workloads are isolated from sequential writes; therefore, depending on the database size, the data and log files are typically placed on separate volumes.
    - For large critical databases, the single data file is on FlexVol 1 and its corresponding log file is on FlexVol 2.
    - For better consolidation, small-to-medium-size noncritical databases are grouped such that all the data files are on FlexVol 1 and their corresponding log files are on FlexVol 2. However, you will lose application-level granularity through this grouping.
  - Another variant is to have all the files within the same FlexVol 3, with data files in LUN1 and its log files in LUN 2.
- If your environment is virtualized, you would have all the VMs for various enterprise applications shared in a datastore. Typically, the VMs and application binaries are asynchronously replicated using SnapMirror.

## Plan

### Prerequisites

When planning your SnapMirror active sync deployment, ensure you have met the various hardware, software, and system configuration requirements.

## Hardware

- Only two-node HA clusters are supported
- Both clusters must be either AFF (including AFF C-Series) or All-Flash SAN Array (ASA; including C-Series). Mixing is not supported.

## Software

- ONTAP 9.9.1 or later
- ONTAP Mediator 1.2 or later
- A Linux server or virtual machine for the ONTAP Mediator running one of the following:

ONTAP Mediator version	Supported Linux versions
1.7	<ul style="list-style-type: none"><li>• Red Hat Enterprise Linux: 8.5, 8.6, 8.7, 8.8, 8.9, 9.0, 9.1, 9.2, and 9.3</li><li>• Rocky Linux 8 and 9</li></ul>
1.6	<ul style="list-style-type: none"><li>• Red Hat Enterprise Linux: 8.4, 8.5, 8.6, 8.7, 8.8, 9.0, 9.1, 9.2</li><li>• Rocky Linux 8 and 9</li></ul>
1.5	<ul style="list-style-type: none"><li>• Red Hat Enterprise Linux: 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5</li><li>• CentOS: 7.6, 7.7, 7.8, 7.9</li></ul>
1.4	<ul style="list-style-type: none"><li>• Red Hat Enterprise Linux: 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5</li><li>• CentOS: 7.6, 7.7, 7.8, 7.9</li></ul>
1.3	<ul style="list-style-type: none"><li>• Red Hat Enterprise Linux: 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3</li><li>• CentOS: 7.6, 7.7, 7.8, 7.9</li></ul>
1.2	<ul style="list-style-type: none"><li>• Red Hat Enterprise Linux: 7.6, 7.7, 7.8, 8.1</li><li>• CentOS: 7.6, 7.7, 7.8</li></ul>

## Licensing

- SnapMirror synchronous license must be applied on both clusters.
- SnapMirror license must be applied on both clusters.



If your ONTAP storage systems were purchased before June 2019, see [NetApp ONTAP Master License Keys](#) to get the required SnapMirror synchronous license.

## Networking environment

- Inter-cluster latency round trip time (RTT) must be less than 10 milliseconds.
- Beginning with ONTAP 9.14.1, [SCSI-3 persistent reservations](#) are supported with SnapMirror active sync.

## Supported protocols

- Only SAN protocols are supported (not NFS/SMB).
- Only Fibre Channel and iSCSI protocols are supported.
- The default IPspace is required by SnapMirror active sync for cluster peer relationships. Custom IPspace is not supported.

## NTFS Security Style

NTFS security style is **not** supported on SnapMirror active sync volumes.

## ONTAP Mediator

- The ONTAP Mediator must be provisioned externally and attached to ONTAP for transparent application failover.
- To be fully functional and to enable automatic unplanned failover, the external ONTAP mediator should be provisioned and configured with ONTAP clusters.
- The ONTAP Mediator must be installed in a third failure domain, separate from the two ONTAP clusters.
- When installing the ONTAP Mediator, you should replace the self-signed certificate with a valid certificate signed by a mainstream reliable CA.
- For more information about the ONTAP Mediator, see [Prepare to install the ONTAP Mediator service](#).

## Read-write destination volumes

- SnapMirror active sync relationships are not supported on read-write destination volumes. Before you can use a read-write volume, you must convert it to a DP volume by creating a volume-level SnapMirror relationship and then deleting the relationship. For details, see [Convert an existing SnapMirror relationships to SnapMirror active sync](#).

## Further information

- [Hardware Universe](#)
- [ONTAP Mediator overview](#)

## SnapMirror active sync interoperability

SnapMirror active sync is compatible with numerous operating systems, application hosts, and other features in ONTAP.

### Application hosts

SnapMirror active sync support applications hosts including Hyper-V, Red Hat Enterprise Linux (RHEL), VMware, VMware vSphere Metro Storage Cluster (vMSC), Windows Server, and, beginning with ONTAP 9.14.1, Windows Server Failover Cluster.

### Operating systems

SnapMirror active sync is supported with numerous operating systems, including:

- AIX (beginning ONTAP 9.11.1)

- HP-UX (beginning ONTAP 9.10.1)
- Solaris 11.4 (beginning ONTAP 9.10.1)

## AIX

Beginning with ONTAP 9.11.1, AIX is supported with SnapMirror active sync. With an AIX configuration, the primary cluster is the "active" cluster.

In an AIX configuration, failovers are disruptive. With each failover, you will need to perform a re-scan on the host for I/O operations to resume.

To configure for AIX host with SnapMirror active syn, refer to the Knowledge Base article [How to configure an AIX host for SnapMirror active sync](#).

## HP-UX

Beginning in ONTAP 9.10.1, SnapMirror active sync for HP-UX is supported.

### Automatic unplanned failover with HP-UX

An automatic unplanned failover (AUFO) event on the isolated master cluster may be caused by dual event failure when the connection between the primary and the secondary cluster is lost and the connection between the primary cluster and the mediator is also lost. This is considered a rare event, unlike other AUFO events.

- In this scenario, it might take more than 120 seconds for I/O to resume on the HP-UX host. Depending on the applications that are running, this might not lead to any I/O disruption or error messages.
- To remediate, you must restart applications on the HP-UX host that have a disruption tolerance of less than 120 seconds.

## Solaris

Beginning with ONTAP 9.10.1, SnapMirror active sync supports Solaris 11.4.

To ensure the Solaris client applications are non-disruptive when an unplanned site failover switchover occurs in an SnapMirror active sync environment, modify the default Solaris OS settings. To configure Solaris with the recommended settings, see the Knowledge Base article [Solaris Host support recommended settings in SnapMirror active sync](#).

## ONTAP interoperability


SnapMirror active sync integrates with components of ONTAP to extends its data protection capabilities.

### FabricPool

SnapMirror active sync supports source and destination volumes on FabricPool aggregates with tiering policies of None, Snapshot or Auto. SnapMirror active sync does not support FabricPool aggregates using a tiering policy of All.

### Fan-out configurations

In [fan-out configurations](#), your source volume can be mirrored to a SnapMirror active sync destination endpoint and to one or more SnapMirror asynchronous relationships.



SnapMirror active sync supports [fan-out configurations](#) with the `MirrorAllSnapshots` policy and, beginning in ONTAP 9.11.1, the `MirrorAndVault` policy. Fan-out configurations are not supported in SnapMirror active sync with the `XDPDefault` policy.

Beginning with ONTAP 9.15.1, SnapMirror active sync supports automatic reconfiguration in the fan-out leg after a failover event. If the failover from the primary to the secondary site has succeeded, the tertiary site is automatically reconfigured to treat the secondary site as the source. Reconfiguration is triggered by either a planned or unplanned failover. Reconfiguration also occurs upon failback to the primary site.

For information about managing your fan-out configuration in earlier releases of ONTAP, see [resume protection in the fan-out configuration](#).

#### NDMP restore

Beginning in ONTAP 9.13.1, you can use [NDMP to copy and restore data](#) with SnapMirror active sync. Using NDMP allows you to move data onto the SnapMirror active sync source to complete a restore without pausing protection. This is particularly useful in fan-out configurations.

#### SnapCenter

SnapMirror active sync is supported with SnapCenter beginning with [SnapCenter 5.0](#). SnapCenter enables the creation of Snapshots that can be used to protect and recover applications and virtual machines, enabling always available storage solutions with application-level granularity.

#### SnapRestore

SnapMirror active sync supports partial and single file SnapRestore.

#### Single file SnapRestore

Beginning with ONTAP 9.11.1, [single-file SnapRestore](#) is supported for SnapMirror active sync volumes. You can restore a single file from a Snapshot copy replicated from the SnapMirror active sync source to the destination. Because volumes can contain one or more LUNs, this feature helps you implement a less disruptive restore operation, granularly restoring a single LUN without disrupting the other LUNs. Single File SnapRestore has two options: in-place and out-of-place.

#### Partial file SnapRestore

Beginning in ONTAP 9.12.1, [partial LUN restore](#) is supported for SnapMirror active sync volumes. You can restore a data from application-created Snapshot copies that have been replicated between the SnapMirror active sync source (volume) and the destination (Snapshot copy) volumes. Partial LUN or file restore may be necessary if you need to restore a database on a host that stores multiple databases on the same LUN. Using this functionality requires you to know the starting byte offset of the data and byte count.

## Large LUNs and large volumes

Support for large LUNs and large volumes (greater than 100 TB) depends on the version of ONTAP you are using and your platform.

### ONTAP 9.12.1P2 and later

- For ONTAP 9.12.1 P2 and later, SnapMirror active sync supports Large LUNs and large volumes greater than 100TB on ASA and AFF (including C-Series).



For ONTAP Releases 9.12.1P2 and later, you must ensure that both the primary and secondary clusters are either All-Flash SAN Arrays (ASA) or All Flash Array (AFF), and that they both have ONTAP 9.12.1 P2 or later installed. If the secondary cluster is running a version earlier than ONTAP 9.12.1P2 or if the array type is not the same as primary cluster, the synchronous relationship can go out of sync if the primary volume grows larger than 100 TB.

### ONTAP 9.9.1 - 9.12.1P1

- For ONTAP releases between ONTAP 9.9.1 and 9.12.1 P1 (inclusive), Large LUNs and large volumes greater than 100TB are supported only on All-Flash SAN Arrays.



For ONTAP releases between ONTAP 9.9.1 and 9.12.1 P2, you must ensure that both the primary and secondary clusters are All-Flash SAN Arrays, and that they both have ONTAP 9.9.1 or later installed. If the secondary cluster is running a version earlier than ONTAP 9.9.1 or if it is not an All-Flash SAN Array, the synchronous relationship can go out of sync if the primary volume grows larger than 100 TB.

## Object limits for SnapMirror active sync

When preparing to use SnapMirror active sync, be aware of the following object limits.

### Consistency groups in a cluster

Consistency group limits for a cluster with SnapMirror active sync are calculated based on relationships and depend on the version of ONTAP used. Limits are platform-independent.

ONTAP version	Maximum number of relationships
ONTAP 9.11.1 and later	50
ONTAP 9.10.1	20
ONTAP 9.9.1	5

### Volumes per consistency group

The maximum number of volumes per consistency group with SnapMirror active sync is platform independent.

ONTAP version	Maximum number of volumes supported in a consistency group relationship
ONTAP 9.15.1 and later	80



ONTAP version	Maximum number of volumes supported in a consistency group relationship
ONTAP 9.10.1-9.14.1	16
ONTAP 9.9.1	12

## Volumes

Volume limits in SnapMirror active sync are calculated based on the number of endpoints, not the number of relationships. A consistency group with 12 volumes contributes 12 endpoints on both the primary and secondary cluster. Both SnapMirror active sync and SnapMirror Synchronous relationships contribute to the total number of endpoints.

The maximum endpoints per platform are included in the following table.

S. No	Platform	Endpoints per HA for SnapMirror active sync			Overall sync and SnapMirror active sync endpoints per HA		
		ONTAP 9.11.1 and later	ONTAP 9.10.1	ONTAP 9.9.1	ONTAP 9.11.1 and later	ONTAP 9.10.1	ONTAP 9.9.1
1	AFF	400	200	60	400	200	80
2	ASA	400	200	60	80	200	400

## SAN object limits

SAN object limits are included in the following table. The limits apply regardless of the platform.

Object in an SnapMirror active sync relationship	Count
LUNs per volume	256
LUN maps per node	<ul style="list-style-type: none"> <li>• 4096 (ONTAP 9.10 and later)</li> <li>• 2048 (ONTAP 9.9.1 and earlier)</li> </ul>
LUN maps per cluster	<ul style="list-style-type: none"> <li>• 8192 (ONTAP 9.10 and later)</li> <li>• 4096 (ONTAP 9.9.1 and earlier)</li> </ul>
LIFs per SVM (with at least one volume in an SnapMirror active sync relationship)	256
Inter-cluster LIFs per node	4
Inter-cluster LIFs per cluster	8

## Related information

- [Hardware Universe](#)
- [Consistency group limits](#)

# Configure

## Configure the ONTAP Mediator and clusters for SnapMirror active sync

SnapMirror active sync utilizes peered clusters to ensure your data is available in the event of a failover scenario. The ONTAP Mediator is a key resource ensuring business continuity, monitoring the health of each cluster. To configure SnapMirror active sync, you must first install the ONTAP Mediator and ensure you primary and secondary clusters are configured properly.

Once you have installed the ONTAP Mediator and configured your clusters, you must [Initialize the ONTAP mediator for SnapMirror active sync](#) the ONTAP Mediator for use with SnapMirror active sync. You must then [create, initialize, and map the consistency group for SnapMirror active sync](#).

### ONTAP Mediator

The ONTAP Mediator establishes a quorum for the ONTAP clusters in an SnapMirror active sync relationship. It coordinates automated failover when a failure is detected, determining which cluster acts as the primary and ensuring data is served to and from the correct destination.

### Prerequisites for the ONTAP Mediator

- The ONTAP Mediator includes its own set of prerequisites. You must meet these prerequisites before installing the mediator.

For more information, see [Prepare to install the ONTAP Mediator service](#).

- By default, the ONTAP Mediator provides service through TCP port 31784. You should make sure that port 31784 is open and available between the ONTAP clusters and the mediator.

### Install the ONTAP Mediator and confirm cluster configuration

Proceed through each of the following steps. For each step, you should confirm that the specific configuration has been performed. Use the link included after each step to get more information as needed.

#### Steps

1. Install the ONTAP Mediator service before you ensure that your source and destination clusters are configured properly.

[Prepare to install or upgrade the ONTAP Mediator service](#)

2. Confirm that a cluster peering relationship exists between the clusters.



The default IPspace is required by SnapMirror active sync for cluster peer relationships. A custom IPspace is not supported.

[Configure peer relationships](#)

3. Confirm that the Storage VMs are created on each cluster.

[Creating an SVM](#)

4. Confirm that a peer relationship exists between the Storage VMs on each cluster.

### [Creating an SVM peering relationship](#)

5. Confirm that the volumes exist for your LUNs.

### [Creating a volume](#)

6. Confirm that at least one SAN LIF is created on each node in the cluster.

### [Considerations for LIFs in a cluster SAN environment](#)

### [Creating a LIF](#)

7. Confirm that the necessary LUNs are created and mapped to an igroup, which is used to map LUNs to the initiator on the application host.

### [Create LUNs and map igroups](#)

8. Rescan the application host to discover any new LUNs.

## **Initialize the ONTAP mediator for SnapMirror active sync**

Once you have installed the ONTAP Mediator and confirmed your cluster configuration, you must initialize the ONTAP Mediator for cluster monitoring. You can initialize the ONTAP Mediator using System Manager or the ONTAP CLI.

## System Manager

With System Manager, you can configure the ONTAP Mediator server for automated failover. You can also replace the self-signed SSL and CA with the third party validated SSL Certificate and CA if you have not already done so.



From ONTAP 9.8 through 9.14.1, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

### Steps

1. Navigate to **Protection > Overview > Mediator > Configure**.
2. Select **Add**, and enter the following ONTAP Mediator server information:
  - IPv4 address
  - Username
  - Password
  - Certificate

### CLI

You can initialize the ONTAP Mediator from either the primary or secondary cluster using the ONTAP CLI. When you issue the `mediator add` command on one cluster, the ONTAP Mediator is automatically added on the other cluster.

### Steps

1. Initialize Mediator on one of the clusters:

```
snapmirror mediator add -mediator-address IP_Address -peer-cluster  
cluster_name -username user_name
```

### Example

```
cluster1::> snapmirror mediator add -mediator-address 192.168.10.1  
-peer-cluster cluster2 -username mediatoradmin  
Notice: Enter the mediator password.  
  
Enter the password: *****  
Enter the password again: *****
```

2. Check the status of the Mediator configuration:

```
snapmirror mediator show
```

Mediator Address	Peer Cluster	Connection Status	Quorum Status
192.168.10.1	cluster-2	connected	true

Quorum Status indicates whether the SnapMirror consistency group relationships are synchronized

with the mediator; a status of `true` indicates successful synchronization.

## Protect with SnapMirror active sync

SnapMirror active sync offers asymmetric protection and, beginning with ONTAP 9.15.1, symmetric active/active protection.

### Configure asymmetric protection

Configuring asymmetric protection using SnapMirror active sync involves selecting LUNs on the ONTAP source cluster and adding them to a consistency group.

#### Before you begin

- You must have a SnapMirror synchronous license.
- You must be a cluster or storage VM administrator.
- All constituent volumes in a consistency group must be in a single storage VM (SVM).
  - LUNs can reside on different volumes.
- The source and destination cluster cannot be the same.
- You cannot establish SnapMirror active sync consistency group relationships across ASA clusters and non-ASA clusters.
- The default IPspace is required by SnapMirror active sync for cluster peer relationships. Custom IPspace is not supported.
- The name of the consistency group must be unique.
- The volumes on the secondary (destination) cluster must be type DP.
- The primary and the secondary SVMs must be in a peered relationship.

#### Steps

You can configure a consistency group using the ONTAP CLI or System Manager.

Beginning in ONTAP 9.10.1, ONTAP offers a consistency group endpoint and menu in System Manager, offering additional management utilities. If you are using ONTAP 9.10.1 or later, see [Configure a consistency group](#) then [configure protection](#) to create an SnapMirror active sync relationship.



From ONTAP 9.8 through 9.14.1, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

## System Manager

1. On the primary cluster, navigate to **Protection > Overview > Protect for Business Continuity > Protect LUNs**.
2. Select the LUNs you want to protect and add them to a protection group.
3. Select the destination cluster and SVM.
4. **Initialize relationship** is selected by default. Click **Save** to begin protection.
5. Go to **Dashboard > Performance** to verify IOPS activity for the LUNs.
6. On the destination cluster, use System Manager to verify that the protection for business continuity relationship is in sync: **Protection > Relationships**.

## CLI

1. Create a consistency group relationship from the destination cluster.

```
destination::> snapmirror create -source-path source-path -destination-path destination-path -cg-item-mappings volume-paths -policy policy-name
```

You can map up to 12 constituent volumes using the `cg-item-mappings` parameter on the `snapmirror create` command.

The following example creates two consistency groups: `cg_src_` on the source with `vol1` and `vol2` and a mirrored destination consistency group, `cg_dst`.

```
destination::> snapmirror create -source-path vs1_src:/cg/cg_src -destination-path vs1_dst:/cg/cg_dst -cg-item-mappings vol_src1:@vol_dst1,vol_src2:@vol_dst2 -policy AutomatedFailOver
```

2. From the destination cluster, initialize the consistency group.

```
destination::> snapmirror initialize -destination-path destination-consistency-group
```

3. Confirm that the initialization operation completed successfully. The status should be `InSync`.

```
snapmirror show
```

4. On each cluster, create an igroup so you can map LUNs to the initiator on the application host.

```
lun igroup create -igroup name -protocol fc|iscsi -ostype os -initiator initiator_name
```

5. On each cluster, map LUNs to the igroup:

```
lun map -path path_name -igroup igroup_name
```

6. Verify the LUN mapping completed successfully with the `lun map` command. Then, you can discover the new LUNs on the application host.

## Configure symmetric active/active protection

You can establish symmetric protection using System Manager or the ONTAP CLI. In both interfaces, there are different steps for [uniform and non-uniform configurations](#).

## Before you begin

- Both clusters must be running ONTAP 9.15.1 or later.
- Symmetric active/active configurations require the `AutomatedFailoverDuplex` protection policy. Alternately, you can [create a custom SnapMirror policy](#) provided the `-type` is `automated-failover-duplex`.

## Example 1. Steps

### System Manager

#### Steps for a uniform configuration

1. On the primary site, [create a consistency group using new LUNs](#).
  - a. When creating the consistency group, specify host initiators to create igroups.
  - b. Select the checkbox to **Enable SnapMirror** then choose the `AutomatedFailoverDuplex` policy.
  - c. In the dialog box that appears, select the **Replicate initiator groups** checkbox to replicate igroups. In **Edit proximal settings**, set proximal SVMs for your hosts.
  - d. Select **Save**.

#### Steps for a non-uniform configuration

1. On the primary site, [create a consistency group using new LUNs](#).
  - a. When creating the consistency group, specify host initiators to create igroups.
  - b. Select the checkbox to **Enable SnapMirror** then choose the `AutomatedFailoverDuplex` policy.
  - c. Select **Save** to create the LUNs, consistency group, igroup, SnapMirror relationship, and igroup mapping.
2. On the secondary site, create an igroup and map the LUNs.
  - a. Navigate to **Hosts > SAN Initiator Groups**.
  - b. Select **+Add** to create a new igroup.
  - c. Provide a **Name**, select the **Host Operating System**, then choose **Initiator Group Members**.
  - d. Select **Save**.
3. Map the new igroup to the destination LUNs.
  - a. Navigate to **Storage > LUNs**.
  - b. Select all the LUNs to map to the igroup.
  - c. Select **More** then **Map to Initiator Groups**.

### CLI

#### Steps for a uniform configuration

1. Create a new SnapMirror relationship grouping all the volumes in the application. Ensure you designate the `AutomatedFailOverDuplex` policy to establish bidirectional sync replication.

```
snapmirror create -source-path source_path -destination-path  
destination_path -cg-item-mappings source_volume:@destination_volume  
-policy AutomatedFailOverDuplex
```

2. Confirm the operation has succeeded by waiting for the `Mirrored` State to show as `SnapMirrored` and the `Relationship` Status as `Insync`.

```
snapmirror show -destination-path destination_path
```

3. On your host, configure host connectivity with access to each cluster according to your needs.



4. Establish the igroup configuration. Set the preferred paths for initiators on the local cluster. Specify the option to replicate the configuration to the peer cluster inverse affinity.

```
SiteA::> igroup create -vserver svm_name -igroup igroup_name -replication
-peer peer_svm_name -initiators host -proximal-vserver local
```

```
SiteA::> igroup add -vserver svm_name -igroup igroup_name -initiators host
-proximal-vserver peer_svm
```

5. From the host, discover the paths and verify the hosts have an active/optimized path to the storage LUN from the preferred cluster.
6. Deploy the application and distribute the VM workloads across clusters to achieve the required load balancing.

### Steps for a non-uniform configuration

1. Create a new SnapMirror relationship grouping all the volumes in the application. Ensure you designate the `AutomatedFailOverDuplex` policy to establish bidirectional sync replication.

```
snapmirror create -source-path source_path -destination-path
destination_path -cg-item-mappings source_volume:@destination_volume
-policy AutomatedFailOverDuplex
```

2. Confirm the operation has succeeded by waiting for the Mirrored State to show as SnapMirrored and the Relationship Status as Insync.

```
snapmirror show -destination-path destination_path
```

3. On your host, configure host connectivity with access to each cluster according to your needs.
4. Establish the igroup configurations on both the source and destination clusters.

```
# primary site
SiteA::> igroup create -vserver svm_name -igroup igroup_name -initiators
host_1_name
```

```
# secondary site
SiteB::> igroup create -vserver svm_name -igroup igroup_name -initiators
host_2_name
```

5. From the host, discover the paths and verify the hosts have an active/optimized path to the storage LUN from the preferred cluster.
6. Deploy the application and distribute the VM workloads across clusters to achieve the required load balancing.

## Convert an existing SnapMirror relationship to SnapMirror active sync relationship

If you've configured SnapMirror protection, you can convert the relationship to SnapMirror active sync. Beginning with ONTAP 9.15.1, you can convert the relationship to use symmetric active/active protection.

## Convert an existing SnapMirror relationship to an asymmetric SnapMirror active sync relationship

If you have an existing SnapMirror synchronous relationship between a source and destination cluster, you can convert it to an asymmetric SnapMirror active sync relationship. This allows you to associate the mirrored volumes with a consistency group, ensuring zero RPO across a multi-volume workload. Additionally, you can retain existing SnapMirror snapshots if you need to revert to a point in time prior to establishing the SnapMirror active sync relationship.

### About this task

- You must be a cluster and SVM administrator on the primary and secondary clusters.
- You cannot convert zero RPO to zero RTO sync by changing the SnapMirror policy.
- You must ensure the LUNs are unmapped before issuing the `snapmirror create` command.

If existing LUNs on the secondary volume are mapped and the `AutomatedFailover` policy is configured, the `snapmirror create` command triggers an error.

### Before you begin

- A zero RPO SnapMirror synchronous relationship must exist between the primary and secondary cluster.
- All LUNs on the destination volume must be unmapped before the zero RTO SnapMirror relationship can be created.
- SnapMirror active sync only supports SAN protocols (not NFS/CIFS). Ensure no constituent of the consistency group is mounted for NAS access.

### Steps

1. From the secondary cluster, perform a SnapMirror update on the existing relationship:

```
SiteB::>snapmirror update -destination-path vs1_dst:vol1
```

2. Verify that the SnapMirror update completed successfully:

```
SiteB::>snapmirror show
```

3. Pause each of the zero RPO synchronous relationships:

```
SiteB::>snapmirror quiesce -destination-path vs1_dst:vol1
```

```
SiteB::>snapmirror quiesce -destination-path vs1_dst:vol2
```

4. Delete each of the zero RPO synchronous relationships:

```
SiteB::>snapmirror delete -destination-path vs1_dst:vol1
```

```
SiteB::>snapmirror delete -destination-path vs1_dst:vol2
```

5. Release the source SnapMirror relationship but retain the common Snapshot copies:

```
SiteA::>snapmirror release -relationship-info-only true -destination-path vs1_dst:vol1
```

```
SiteA::>snapmirror release -relationship-info-only true -destination-path vs1_dst:vol2
```

6. Create a zero RTO Snapmirror synchronous relationship:

```
SiteB::> snapmirror create -source-path vs1_src:/cg/cg_src -destination-path
vs1_dst:/cg/cg_dst -cg-item-mappings vol1:@vol1,vol2:@vol2 -policy
AutomatedFailover
```

7. Resynchronize the consistency group:

```
SiteB::> snapmirror resync -destination-path vs1_dst:/cg/cg_dst
```

8. Rescan host LUN I/O paths to restore all paths to the LUNs.

### Convert an existing SnapMirror relationship to symmetric active/active

Beginning with ONTAP 9.15.1, you can convert an existing SnapMirror relationship to a SnapMirror active sync symmetric active/active relationship.

#### Before you begin

- You must be running ONTAP 9.15.1 or later.
- A zero RPO SnapMirror synchronuous relationship must exist between the primary and secondary cluster.
- All LUNs on the destination volume must be unmapped before the zero RTO SnapMirror relationship can be created.
- SnapMirror active sync only supports SAN protocols (not NFS/CIFS). Ensure no constituent of the consistency group is mounted for NAS access.

#### Steps

1. From the secondary cluster, perform a SnapMirror update on the existing relationship:

```
SiteB::>snapmirror update -destination-path vs1_dst:vol1
```

2. Verify that the SnapMirror update completed successfully:

```
SiteB::>snapmirror show
```

3. Pause each of the zero RPO synchronous relationships:

```
SiteB::>snapmirror quiesce -destination-path vs1_dst:vol1
```

```
SiteB::>snapmirror quiesce -destination-path vs1_dst:vol2
```

4. Delete each of the zero RPO synchronous relationships:

```
SiteB::>snapmirror delete -destination-path vs1_dst:vol1
```

```
SiteB::>snapmirror delete -destination-path vs1_dst:vol2
```

5. Release the source SnapMirror relationship but retain the common Snapshot copies:

```
SiteA::>snapmirror release -relationship-info-only true -destination-path
vs1_dst:vol1
```

```
SiteA::>snapmirror release -relationship-info-only true -destination-path
```

```
vs1_dst:vol2
```

6. Create a zero RTO Snapmirror synchronous relationship with the AutomatedFailoverDuplex policy:

```
SiteB::> snapmirror create -source-path vs1_src:/cg/cg_src -destination-path  
vs1_dst:/cg/cg_dst -cg-item-mappings vol1:@vol1,vol2:@vol2 -policy  
AutomatedFailoverDuplex
```

7. If the existing hosts are local the primary cluster, add the host to the secondary cluster and establish connectivity with respective access to each cluster.
8. On the secondary site, delete the LUN maps on the igroups associated with remote hosts.



Ensure the igroup does not contain maps for non-replicated LUNs.

```
SiteB::> lun mapping delete -vserver svm_name -igroup igroup -path <>
```

9. On the primary site, modify the initiator configuration for existing hosts to set the proximal path for initiators on the local cluster.

```
SiteA::> igroup initiator add-proximal-vserver -vserver svm_name -initiator  
host -proximal-vserver server
```

10. Add a new igroup and initiator for the new hosts and set the host proximity for host affinity to its local site. Enable igroup replication to replicate the configuration and invert the host locality on the remote cluster.

```
SiteA::> igroup modify -vserver vsA -igroup ig1 -replication-peer vsB  
SiteA::> igroup initiator add-proximal-vserver -vserver vsA -initiator host2  
-proximal-vserver vsB
```

11. Discover the paths on the hosts and verify the hosts have an Active/Optimized path to the storage LUN from the preferred cluster
12. Deploy the application and distribute the VM workloads across clusters.
13. Resynchronize the consistency group:

```
SiteB::> snapmirror resync -destination-path vs1_dst:/cg/cg_dst
```

14. Rescan host LUN I/O paths to restore all paths to the LUNs.

## Convert SnapMirror active sync relationship type

Beginning with ONTAP 9.15.1, you can convert between types of SnapMirror active sync protection: from asymmetric to symmetric active/active and vice versa.

### Convert to a symmetric active/active relationship

You can convert a SnapMirror active sync relationship with asynchronous protection to use symmetric active/active.

#### Before you begin

- Both clusters must be running ONTAP 9.15.1 or later.
- Symmetric active/active configurations require the AutomatedFailoverDuplex protection policy. Alternately, you can [create a custom SnapMirror policy](#) provided the `-type` is `automated-failover-duplex`.

## System Manager

### Steps for a uniform configuration

1. Remove the destination igroup:
  - a. On the destination cluster, navigate to **Hosts > SAN Initiator Groups**.
  - b. Select the igroup with the SnapMirror relationship, then **Delete**.
  - c. In the dialog box, select the **Unmap the associated LUNs** box then **Delete**.
2. Edit the SnapMirror active sync relationship.
  - a. Navigate to **Protection > Relationships**.
  - b. Select the kabob menu next to the relationship you want to modify then **Edit**.
  - c. Modify the **Protection Policy** to AutomatedFailoverDuplex.
  - d. Selecting `AutoMatedFailoverDuplex` prompts a dialog box to modify host proximity settings. For the initiators, select the appropriate option for **Initiator proximal to** then **Save**.
  - e. Select **Save**.
3. In the **Protection** menu, confirm the operation succeeded when the relationship displays as `InSync`.

### Steps for a non-uniform configuration

1. Remove the destination igroup:
  - a. On the secondary site, navigate to **Hosts > SAN Initiator Groups**.
  - b. Select the igroup with the SnapMirror relationship, then **Delete**.
  - c. In the dialog box, select the **Unmap the associated LUNs** box then **Delete**.
2. Create a new igroup:
  - a. In the **SAN Initiator Groups** menu on the destination site, select **Add**.
  - b. Provide a **Name**, select the **Host Operating System**, then choose **Initiator Group Members**.
  - c. Select **Save**.
3. Map the new igroup to the destination LUNs.
  - a. Navigate to **Storage > LUNs**.
  - b. Select all the LUNs to map to the igroup.
  - c. Select **More** then **Map to Initiator Groups**.
4. Edit the SnapMirror active sync relationship.
  - a. Navigate to **Protection > Relationships**.
  - b. Select the kabob menu next to the relationship you want to modify then **Edit**.
  - c. Modify the **Protection Policy** to AutomatedFailoverDuplex.
  - d. Selecting `AutoMatedFailoverDuplex` initiates the option to modify host proximity settings. For the initiators, select the appropriate option for **Initiator proximal to** then **Save**.
  - e. Select **Save**.
5. In the **Protection** menu, confirm the operation succeeded when the relationship displays as `InSync`.

## CLI

### Steps for a uniform configuration

1. Modify the SnapMirror policy from AutomatedFailover to AutomatedFailoverDuplex:

```
snapmirror modify -destination-path destination_path -policy  
AutomatedFailoverDuplex
```

2. Modifying the policy triggers a resync. Wait for the resync to complete and confirm the relationship is Insync:

```
snapmirror show -destination-path destination_path
```

3. If the existing hosts are local the primary cluster, add the host to the second cluster and establish connectivity with respective access to each cluster.
4. On the secondary site, delete the LUN maps on the igroups associated with remote hosts.



Ensure the igroup does not contain maps for non-replicated LUNs.

```
SiteB::> lun mapping delete -vserver svm_name -igroup igroup -path <>
```

5. On the primary site, modify the initiator configuration for existing hosts to set the proximal path for initiators on the local cluster.

```
SiteA::> igroup initiator add-proximal-vserver -vserver svm_name -initiator  
host -proximal-vserver server
```

6. Add a new igroup and initiator for the new hosts and set the host proximity for host affinity to its local site. Enable igroup replication to replicate the configuration and invert the host locality on the remote cluster.

```
SiteA::> igroup modify -vserver vsA -igroup ig1 -replication-peer vsB  
SiteA::> igroup initiator add-proximal-vserver -vserver vsA -initiator  
host2 -proximal-vserver vsB
```

7. Discover the paths on the hosts and verify the hosts have an Active/Optimized path to the storage LUN from the preferred cluster
8. Deploy the application and distribute the VM workloads across clusters.

### Steps for a non-uniform configuration

1. Modify the SnapMirror policy from AutomatedFailover to AutomatedFailoverDuplex:

```
snapmirror modify -destination-path destination_path -policy  
AutomatedFailoverDuplex
```

2. Modifying the policy triggers a resync. Wait for the resync to complete and confirm the relationship is Insync:

```
snapmirror show -destination-path destination_path
```

3. If the existing hosts are local to the primary cluster, add the host to the second cluster and establish connectivity with respective access to each cluster.

4. On the secondary site, delete the LUN maps on the igroups associated with remote hosts.



Ensure the igroup does not contain maps for non-replicated LUNs.

```
SiteB::> lun mapping delete -vserver svm_name -igroup igroup -path <>
```

5. On the primary site, modify the initiator configuration for existing hosts to set the proximal path for initiators on the local cluster.

```
SiteA::> igroup initiator add-proximal-vserver -vserver Svm_name -initiator host -proximal-vserver server
```

6. On the secondary site, add a new igroup and initiator for the new hosts and set the host proximity for host affinity to its local site. Map the LUNs to the group.

```
SiteB::> igroup create -vserver svm_name -igroup igroup_name
SiteB::> igroup add -vserver svm_name -igroup igroup_name -initiator host_name
SiteB::> lun mapping create -igroup igroup_name -path path_name
```

7. Discover the paths on the hosts and verify the hosts have an Active/Optimized path to the storage LUN from the preferred cluster
8. Deploy the application and distribute the VM workloads across clusters.

## Convert from symmetric active/active to an asymmetric relationship

If you've configured symmetric active/active protection, you can convert the relationship to asymmetric protection using the ONTAP CLI.

### Steps

1. Move all the VM workloads to the host local to the source cluster.
2. Remove the igroup configuration for the hosts that are not managing the VM instances and modify the igroup configuration to terminate igroup replication.

code

3. On the secondary site, unmap the LUNs.

```
SiteB::> lun mapping delete -vserver svm_name -igroup igroup_name -path <>
```

4. On the secondary site, delete the symmetric active/active relationship.

```
SiteB::> snapmirror delete -destination-path destination_path
```

5. On the primary site, release the symmetric active/active relationship.

```
SiteA::> snapmirror release -destination-path destination_path -relationship -info-only true
```

6. From the secondary site, create a relationship to the same set of volumes with the AutomatedFailover policy the resynchronize the relationship.



```
SiteB::> snapmirror create -source-path source_path -destination-path
destination_path -cg-item-mappings source:@destination -policy
AutomatedFailover
SiteB::> snapmirror resync -destination-path vs1:/cg/cg1_dst
```



The consistency group on the secondary site needs [to be deleted](#) before recreating the relationship. The destination volumes [must be converted to type DP](#).

7. Confirm the relationship Mirror State is Snapmirrored the Relationship Status is Insync.

```
snapmirror show -destination-path destination_path
```

8. Re-discover the paths from the host.

## Manage SnapMirror active sync and protect data

### Create a common Snapshot copy

In addition to the regularly scheduled Snapshot copy operations, you can manually create a common [Snapshot copy](#) between the volumes in the primary SnapMirror consistency group and the volumes in the secondary SnapMirror consistency group.

#### About this task

The scheduled Snapshot creation interval is 12 hours.

#### Before you begin

- The SnapMirror group relationship must be in sync.

#### Steps

1. Create a common Snapshot copy:

```
destination::>snapmirror update -destination-path vs1_dst:/cg/cg_dst
```

2. Monitor the progress of the update:

```
destination::>snapmirror show -fields -newest-snapshot
```

### Perform a planned failover of clusters in a SnapMirror active sync relationship

In a planned failover of ONTAP clusters in a SnapMirror active sync relationship, you switch the roles of the primary and secondary clusters, so that the secondary cluster takes over from the primary cluster. During a failover, what is normally the secondary cluster processes input and output requests locally without disrupting client operations.

You may want to perform a planned failover to test the health of your disaster recovery configuration or to perform maintenance on the primary cluster.

#### About this task

A planned failover is initiated by the administrator of the secondary cluster. The operation requires switching the primary and secondary roles so that the secondary cluster takes over from the primary. The new primary cluster can then begin processing input and output requests locally without disrupting client operations.

### Before you begin

- The SnapMirror active sync relationship must be in sync.
- You cannot initiate a planned failover when a nondisruptive operation is in process. Nondisruptive operations include volume moves, aggregate relocations, and storage failovers.
- The ONTAP Mediator must be configured, connected, and in quorum.

### Steps

You can perform a planned failover using the ONTAP CLI or System Manager.

#### System Manager



From ONTAP 9.8 through 9.14.1, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

1. In System Manager, select **Protection > Overview > Relationships**.
2. Identify the SnapMirror active sync relationship you want to failover. Next to its name, select the ... next to the relationship's name, then select **Failover**.
3. To monitor the status of the failover, use the `snapmirror failover show` in the ONTAP CLI.

#### CLI

1. From the destination cluster, initiate the failover operation:

```
destination::>snapmirror failover start -destination-path  
vs1_dst:/cg/cg_dst
```

2. Monitor the progress of the failover:

```
destination::>snapmirror failover show
```

3. When the failover operation is complete, you can monitor the SnapMirror synchronous protection relationship status from the destination:

```
destination::>snapmirror show
```

## Recover from automatic unplanned failover operations

An automatic unplanned failover (AUFO) operation occurs when the primary cluster is down or isolated. The ONTAP Mediator detects when a failover occurs and, and executes an automatic unplanned failover to the the secondary cluster. The secondary cluster is converted to the primary and begins serving clients. This operation is performed only with assistance from the ONTAP Mediator.



After the automatic unplanned failover, it is important to rescan the host LUN I/O paths so that there is no loss of I/O paths.

## Reestablish the protection relationship after an unplanned failover


You can reestablish the protection relationship using System Manager or the ONTAP CLI.

### System Manager



#### Steps

From ONTAP 9.8 through 9.14.1, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

1. Navigate to **Protection > Relationships** and wait for the relationship state to show “InSync.”
2. To resume operations on the original source cluster, click  and select **Failover**.

### CLI

You can monitor the status of the automatic unplanned failover using the `snapmirror failover show` command.

For example:

```
ClusterB::> snapmirror failover show -instance
Start Time: 9/23/2020 22:03:29
    Source Path: vs1:/cg/scg3
    Destination Path: vs3:/cg/dcg3
    Failover Status: completed
    Error Reason:
        End Time: 9/23/2020 22:03:30
Primary Data Cluster: cluster-2
Last Progress Update: -
    Failover Type: unplanned
Error Reason codes: -
```

Refer to the [EMS reference](#) to learn about event messages and about corrective actions.

## Resume protection in a fan-out configuration after failover

Beginning with ONTAP 9.15.1, SnapMirror active sync supports automatic reconfiguration in the fan-out leg after a failover event. For more information, see [fan-out configurations](#).

If you're using ONTAP 9.14.1 or earlier and you experience a failover on the secondary cluster in the SnapMirror active sync relationship, the SnapMirror asynchronous destination becomes unhealthy. You must manually restore protection by deleting and recreating the relationship with the SnapMirror asynchronous endpoint.

### Steps

1. Verify the failover has completed successfully:  
`snapmirror failover show`
2. On the Snapmirror asynchronous endpoint, delete the fan-out endpoint:  
`snapmirror delete -destination-path destination_path`
3. On the third site, create a SnapMirror asynchronous relationships between the new SnapMirror active sync primary volume and the async fan-out destination volume:  
`snapmirror create -source-path source_path -destination-path destination_path -policy MirrorAllSnapshots -schedule schedule`
4. Resynchronize the relationship:  
`snapmirror resync -destination-path destination_path`
5. Verify the relationship status and health:  
`snapmirror show`

## Monitor SnapMirror active sync operations

You can monitor the following SnapMirror active sync operations to ensure the health of your SnapMirror active sync configuration:

- ONTAP Mediator
- Planned failover operations
- Automatic unplanned failover operations
- SnapMirror active sync availability



Beginning in ONTAP 9.15.1, System Manager displays the status of your SnapMirror active sync relationship from either cluster. You can also monitor the ONTAP Mediator's status from either cluster in System Manager.

### ONTAP Mediator

During normal operations, the ONTAP Mediator state should be connected. If it's in any other state, this might indicate an error condition. You can review the [Event Management System \(EMS\) messages](#) to determine the error and appropriate corrective actions.

### Planned failover operations

You can monitor status and progress of a planned failover operation using the `snapmirror failover show` command. For example:

```
ClusterB::> snapmirror failover start -destination-path vs1:/cg/dcg1
```

Once the failover operation is complete, you can monitor the SnapMirror protection status from the new destination cluster. For example:

```
ClusterA::> snapmirror show
```

Refer to the [EMS reference](#) to learn about event messages and corrective actions.

### Automatic unplanned failover operations

During an unplanned automatic failover, you can monitor the status of the operation using the `snapmirror failover show` command.

```
ClusterB::> snapmirror failover show -instance
Start Time: 9/23/2020 22:03:29
      Source Path: vs1:/cg/scg3
      Destination Path: vs3:/cg/dcg3
      Failover Status: completed
      Error Reason:
            End Time: 9/23/2020 22:03:30
Primary Data Cluster: cluster-2
Last Progress Update: -
      Failover Type: unplanned
Error Reason codes: -
```

Refer to the [EMS reference](#) to learn about event messages and about corrective actions.

### SnapMirror active sync availability

You can check the availability of the SnapMirror active sync relationship using a series of commands, either on the primary cluster, the secondary cluster, or both.

Commands you use include the `snapmirror mediator show` command on both the primary and secondary cluster to check the connection and quorum status, the `snapmirror show` command, and the `volume show` command. For example:

```

SMBC_A::*> snapmirror mediator show
Mediator Address Peer Cluster      Connection Status Quorum Status
-----
10.236.172.86   SMBC_B           connected         true

SMBC_B::*> snapmirror mediator show
Mediator Address Peer Cluster      Connection Status Quorum Status
-----
10.236.172.86   SMBC_A           connected         true

SMBC_B::*> snapmirror show -expand

Progress
Source          Destination Mirror Relationship Total
Last
Path            Type Path          State Status Progress Healthy
Updated
-----
-----
vs0:/cg/cg1 XDP vs1:/cg/cg1_dp Snapmirrored Insync - true -
vs0:vol1 XDP vs1:vol1_dp Snapmirrored Insync - true -
2 entries were displayed.

SMBC_A::*> volume show -fields is-smbc-master,smbc-consensus,is-smbc-
failover-capable -volume vol1
vserver volume is-smbc-master is-smbc-failover-capable smbc-consensus
-----
vs0 vol1 true false Consensus

SMBC_B::*> volume show -fields is-smbc-master,smbc-consensus,is-smbc-
failover-capable -volume vol1_dp
vserver volume is-smbc-master is-smbc-failover-capable smbc-consensus
-----
vs1 vol1_dp false true No-consensus

```

## Add or remove volumes to a consistency group

As your application workload requirements change, you may need to add or remove volumes from a consistency group to ensure business continuity. The process of adding and removing volumes in an active SnapMirror active sync relationship depends on the version of ONTAP you are using.

In most instances, this is a disruptive process requiring you to break the SnapMirror relationship, modify the consistency group, then resume protection. Beginning with ONTAP 9.13.1, adding volumes to a consistency group with an active SnapMirror relationship is a non-disruptive operation.

**About this task**

- In ONTAP 9.9.1, you can add or remove volumes to a consistency group using the ONTAP CLI.
- Beginning with ONTAP 9.10.1, it is recommended that you manage [consistency groups](#) through System Manager or with the ONTAP REST API.

If you want to change the composition of the consistency group by adding or removing a volume, you must first delete the original relationship and then create the consistency group again with the new composition.

- Beginning in ONTAP 9.13.1, you can non-disruptively add volumes to a consistency group with an active SnapMirror relationship from the source or destination.

Removing volumes is a disruptive operation. You must break the SnapMirror relationship before proceeding removing volumes.

## ONTAP 9.9.1-9.13.0

### Before you begin

- You cannot begin to modify the consistency group while it is in the InSync state.
- The destination volume should be of type DP.
- The new volume you add to expand the consistency group must have a pair of common Snapshot copies between the source and destination volumes.

### Steps

The examples shown in two volume mappings: `vol_src1 ↔ vol_dst1` and `vol_src2 ↔ vol_dst2`, in a consistency group relationship between the end points `vs1_src:/cg/cg_src` and `vs1_dst:/cg/cg_dst`.

1. On the source and destination clusters, verify there is a common Snapshot between the source and destination clusters with the command `snapshot show -vserver svm_name -volume volume_name -snapshot snapmirror`

```
source::>snapshot show -vserver vs1_src -volume vol_src3 -snapshot snapmirror*
```

```
destination::>snapshot show -vserver vs1_dst -volume vol_dst3 -snapshot snapmirror*
```

2. If no common Snapshot copy exists, create and initialize a FlexVol SnapMirror relationship:

```
destination::>snapmirror initialize -source-path vs1_src:vol_src3 -destination-path vs1_dst:vol_dst3
```

3. Delete the consistency group relationship:

```
destination::>snapmirror delete -destination-path vs1_dst:vol_dst3
```

4. Release the source SnapMirror relationship and retain the common Snapshot copies:

```
source::>snapmirror release -relationship-info-only true -destination-path vs1_dst:vol_dst3
```

5. Unmap the LUNs and delete the existing consistency group relationship:

```
destination::>lun mapping delete -vserver vs1_dst -path <lun_path> -igroup <igroup_name>
```



The destination LUNs are unmapped, while the LUNs on the primary copy continue to serve the host I/O.

```
destination::>snapmirror delete -destination-path vs1_dst:/cg/cg_dst
```

```
source::>snapmirror release -destination-path vs1_dst:/cg/cg_dst -relationship-info-only true
```

6. If you are using ONTAP 9.10.1 through 9.13.0, delete and recreate the consistency group on



the source with the correct composition. Follow the steps in [Delete a consistency group](#) and then [Configure a single consistency group](#). In ONTAP 9.10.1 and later, you must perform the delete and create operations in System Manager or with the ONTAP REST API; there is no CLI procedure.

**If you are using ONTAP 9.9.1, skip to the next step.**

7. Create the new consistency group on the destination with the new composition:

```
destination::>snapmirror create -source-path vs1_src:/cg/cg_src
-destination-path vs1_dst:/cg/cg_dst -cg-item-mappings vol_src1:@vol_dst1,
vol_src2:@vol_dst2, vol_src3:@vol_dst3
```

8. Resynchronize the zero RTO consistency group relationship to ensure it is in sync:

```
destination::>snapmirror resync -destination-path vs1_dst:/cg/cg_dst
```

9. Remap the LUNs that you unmapped in Step 5:

```
destination::> lun map -vserver vs1_dst -path lun_path -igroup igroup_name
```

10. Rescan host LUN I/O paths to restore all paths to the LUNs.

#### ONTAP 9.13.1 and later


Beginning in ONTAP 9.13.1, you can non-disruptively add volumes to a consistency group with an active SnapMirror active sync relationship. SnapMirror active sync supports adding volumes from both the source or destination.



From ONTAP 9.8 through 9.14.1, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

For details on adding volumes from the source consistency group, see [Modify a consistency group](#).

#### Add a volume from the destination cluster

1. On the destination cluster, select **Protection > Relationships**.
2. Find the SnapMirror configuration you want to add volumes to. Select  then **Expand**.
3. Select the volume relationships whose volumes are to be added to consistency group
4. Select **Expand**.

## Upgrade and revert ONTAP with SnapMirror active sync

SnapMirror active sync is supported beginning with ONTAP 9.9.1. Upgrading and reverting your ONTAP cluster has implications on your SnapMirror active sync relationships depending on the ONTAP version to which you are upgrading or reverting.

### Upgrade ONTAP with SnapMirror active sync

To use SnapMirror active sync, all nodes on the source and destination clusters must be running ONTAP 9.9.1 or later.

When upgrading ONTAP with active SnapMirror active sync relationships, you should use [automated](#)

[nondisruptive upgrade \(ANDU\)](#). Using ANDU ensures your SnapMirror active sync relationships are in sync and healthy during the upgrade process.

There are no configuration steps to prepare SnapMirror active sync deployments for ONTAP upgrades. However, it is recommended that before and after the upgrade, you should check that:

- SnapMirror active sync relationships are in sync.
- There are no errors related to SnapMirror in the event log.
- The Mediator is online and healthy from both clusters.
- All hosts can see all paths properly to protect LUNs.



When you upgrade clusters from ONTAP 9.9.1 or 9.9.1 to ONTAP 9.10.1 and later, ONTAP creates new [consistency groups](#) on both source and destination clusters for SnapMirror active sync relationships that can be configured using System Manager.



The `snapmirror quiesce` and `snapmirror resume` commands are not supported with SnapMirror active sync.

## Revert to ONTAP 9.9.1 from ONTAP 9.10.1

To revert relationships from 9.10.1 to 9.9.1, SnapMirror active sync relationships must be deleted, followed by the 9.10.1 consistency group instance. Consistency groups with an active SnapMirror active sync relationship cannot be deleted. Any FlexVol volumes that were upgraded to 9.10.1 previously associated with another Smart Container or Enterprise App in 9.9.1 or earlier will no longer be associated on revert. Deleting consistency groups does not delete the constituent volumes or volume granular snapshots. Refer to [Delete a consistency group](#) for more information on this task in ONTAP 9.10.1 and later.

## Revert from ONTAP 9.9.1



SnapMirror active sync is not supported with mixed ONTAP clusters that include releases earlier than ONTAP 9.9.1.

When you revert from ONTAP 9.9.1 to an earlier release of ONTAP, you must be aware of the following:

- If the cluster hosts an SnapMirror active sync destination, reverting to ONTAP 9.8 or earlier is not allowed until the relationship is broken and deleted.
- If the cluster hosts an SnapMirror active sync source, reverting to ONTAP 9.8 or earlier is not allowed until the relationship is released.
- All user-created custom SnapMirror active sync policies must be deleted before reverting to ONTAP 9.8 or earlier.

To meet these requirements, see [Remove a SnapMirror active sync configuration](#).

## Steps

1. Confirm your readiness to revert, entering the following command from one of the clusters in the SnapMirror active sync relationship:

```
cluster::> system node revert-to -version 9.7 -check-only
```

The following sample output shows a cluster that is not ready to revert with instructions for clean up.

```

cluster::> system node revert-to -version 9.7 -check-only
Error: command failed: The revert check phase failed. The following
issues must be resolved before revert can be completed. Bring the data
LIFs down on running vservers. Command to list the running vservers:
vserver show -admin-state running Command to list the data LIFs that are
up: network interface show -role data -status-admin up Command to bring
all data LIFs down: network interface modify {-role data} -status-admin
down
Disable snapshot policies.
    Command to list snapshot policies: "snapshot policy show".
    Command to disable snapshot policies: "snapshot policy modify
-vserver
* -enabled false"

Break off the initialized online data-protection (DP) volumes and
delete
Uninitialized online data-protection (DP) volumes present on the
local
node.
    Command to list all online data-protection volumes on the local
node:
volume show -type DP -state online -node <local-node-name>
    Before breaking off the initialized online data-protection volumes,
quiesce and abort transfers on associated SnapMirror relationships
and
wait for the Relationship Status to be Quiesced.
    Command to quiesce a SnapMirror relationship: snapmirror quiesce
    Command to abort transfers on a SnapMirror relationship: snapmirror
abort
    Command to see if the Relationship Status of a SnapMirror
relationship
is Quiesced: snapmirror show
    Command to break off a data-protection volume: snapmirror break
    Command to break off a data-protection volume which is the
destination
of a SnapMirror relationship with a policy of type "vault":
snapmirror
break -delete-snapshots
    Uninitialized data-protection volumes are reported by the
"snapmirror
break" command when applied on a DP volume.
    Command to delete volume: volume delete

Delete current version snapshots in advanced privilege level.
    Command to list snapshots: "snapshot show -fs-version 9.9.1"
    Command to delete snapshots: "snapshot prepare-for-revert -node

```

```
<nodename>"
```

```
Delete all user-created policies of the type active-strict-sync-mirror
```

```
and active-sync-mirror.
```

```
The command to see all active-strict-sync-mirror and active-sync-mirror
```

```
type policies is:
```

```
snapmirror policy show -type  
active-strict-sync-mirror,active-sync-mirror
```

```
The command to delete a policy is :
```

```
snapmirror policy delete -vserver <SVM-name> -policy <policy-name>
```

2. Once you've satisfied the requirements of the revert check, see [Revert ONTAP](#).

## Remove a SnapMirror active sync configuration

If you no longer require zero RTO SnapMirror synchronous protection, you can delete your SnapMirror active sync relationship.

### Remove an asymmetric configuration

- Before you delete the SnapMirror active sync relationship, all LUNs in the destination cluster must be unmapped.
- After the LUNs are unmapped and the host is rescanned, the SCSI target notifies the hosts that the LUN inventory has changed. The existing LUNs on the zero RTO secondary volumes change to reflect a new identity after the zero RTO relationship is deleted. Hosts discover the secondary volume LUNs as new LUNs that have no relationship to the source volume LUNs.
- The secondary volumes remain DP volumes after the relationship is deleted. You can issue the `snapmirror break` command to convert them to read/write.
- Deleting the relationship is not allowed in the failed-over state when the relationship is not reversed.

### Steps

1. From the secondary cluster, remove the SnapMirror active sync consistency group relationship between the source endpoint and destination endpoint:

```
destination::>snapmirror delete -destination-path vs1_dst:/cg/cg_dst
```

2. From the primary cluster, release the consistency group relationship and the Snapshot copies created for the relationship:

```
source::>snapmirror release -destination-path vs1_dst:/cg/cg_dst
```

3. Perform a host rescan to update the LUN inventory.
4. Beginning with ONTAP 9.10.1, deleting the SnapMirror relationship does not delete the consistency group. If you want to delete the consistency group, you must use System Manager or the ONTAP REST API. See [Delete a consistency group](#) for more information.

## Remove a symmetric active/active configuration

You can remove a symmetric configuration using System Manager or the ONTAP CLI. In both interfaces, there are different steps for [uniform](#) and [non-uniform configurations](#).

## System Manager

### Steps for a uniform configuration

1. On the primary site, remove the remote hosts from the igroup and terminate replication.
  - a. Navigate to **Hosts > SAN Initiator Groups**.
  - b. Select the igroup you want to modify then **Edit**.
  - c. Remove the remote initiator and terminate igroup replication. Select **Save**.
2. On the secondary site, delete the replicated relationship by unmapping the LUNs.
  - a. Navigate to **Hosts > SAN Initiator Groups**.
  - b. Select the igroup with the SnapMirror relationship, then **Delete**.
  - c. In the dialog box, select the **Unmap the associated LUNs** box then **Delete**.
  - d. Navigate to **Protection > Relationships**.
  - e. Select the SnapMirror active sync relationship then **Release** to delete the relationships.

### Steps for a non-uniform configuration

1. On the primary site, remove the remote hosts from the igroup and terminate replication.
  - a. Navigate to **Hosts > SAN Initiator Groups**.
  - b. Select the igroup you want to modify then **Edit**.
  - c. Remove the remote initiator and terminate igroup replication. Select **Save**.
2. On the secondary site, remove the SnapMirror active sync relationship.
  - a. Navigate to **Protection > Relationships**.
  - b. Select the SnapMirror active sync relationship then **Release** to delete the relationships.

## CLI

### Steps for a uniform configuration

1. Move all the VM workloads to the host local to source cluster of SnapMirror active sync.
2. On the source cluster, remove the initiators from the igroup and modify the igroup configuration to terminate igroup replication.

```
SiteA::> igroup remove -vserver svm_name -igroup igroup_name -initiator host2
SiteA::> igroup modify -vserver svm_name -igroup igroup_name -replication -peer "-"
```

3. On the secondary site, delete the LUN mapping and remove the igroup configuration:

```
SiteB::> lun mapping delete -vserver svm_name -igroup igroup_name -path <>
SiteB::> igroup delete -vserver svm_name -igroup igroup_name
```

4. On the secondary site, delete the SnapMirror active sync relationship.

```
SiteB::> snapmirror delete -destination-path destination_path
```

5. On the primary site, release the SnapMirror active sync relationship from primary site.

```
SiteA::> snapmirror release -destination-path destination_path
```

6. Rediscover the paths to verify that only the local path is available to the host.

#### Steps for a non-uniform configuration

1. Move all the VM workloads to the host local to source cluster of SnapMirror active sync.
2. On the source cluster, remove the initiators from the igroup.

```
SiteA::> igroup remove -vserver svm_name -igroup igroup_name -initiator
host2
```

3. On the secondary site, delete the LUN mapping and remove the igroup configuration:

```
SiteB::> lun mapping delete -vserver svm_name -igroup igroup_name -path <>
SiteB::> igroup delete -vserver svm_name -igroup igroup_name
```

4. On the secondary site, delete the SnapMirror active sync relationship.

```
SiteB::> snapmirror delete -destination-path destination_path
```

5. On the primary site, release the SnapMirror active sync relationship from primary site.

```
SiteA::> snapmirror release -destination-path destination_path
```

6. Rediscover the paths to verify that only the local path is available to the host.

## Remove ONTAP Mediator

If you want to remove an existing ONTAP Mediator configuration from your ONTAP clusters, you can do so by using the `snapmirror mediator remove` command.

#### Steps

1. Remove ONTAP Mediator:

```
snapmirror mediator remove -mediator-address 12.345.678.90 -peer-cluster
cluster_xyz
```

## Troubleshoot

### SnapMirror delete operation fails in takeover state

#### Issue:

When ONTAP 9.9.1 is installed on a cluster, executing the `snapmirror delete` command fails when an SnapMirror active sync consistency group relationship is in takeover state.

**Example:**

```
C2_cluster::> snapmirror delete vs1:/cg/dd  
  
Error: command failed: RPC: Couldn't make connection
```

**Solution**

When the nodes in an SnapMirror active sync relationship are in takeover state, perform the SnapMirror delete and release operation with the "-force" option set to true.

**Example:**

```
C2_cluster::> snapmirror delete vs1:/cg/dd -force true  
  
Warning: The relationship between source "vs0:/cg/ss" and destination  
        "vs1:/cg/dd" will be deleted, however the items of the  
destination  
        Consistency Group might not be made writable, deletable, or  
modifiable  
        after the operation. Manual recovery might be required.  
Do you want to continue? {y|n}: y  
Operation succeeded: snapmirror delete for the relationship with  
destination "vs1:/cg/dd".
```

## Failure creating a SnapMirror relationship and initializing consistency group

**Issue:**

Creation of SnapMirror relationship and consistency group initialization fails.

**Solution:**

Ensure that you have not exceeded the limit of consistency groups per cluster. Consistency group limits in SnapMirror active sync are platform independent and differ based on the version of ONTAP. See [Object limits](#) for guidance specific to your ONTAP version.

**Error:**

If the consistency group is stuck initializing, check the status of your consistency group initializations with the ONTAP REST API, System Manager or the command `sn show -expand`.



From ONTAP 9.8 through 9.14.1, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

**Solution:**


If consistency groups fail to initialize, remove the SnapMirror active sync relationship, delete the consistency group, then recreate the relationship and initialize it. This workflow differs depending on the version of ONTAP you are using.

If you are using ONTAP 9.9.1

If you are using ONTAP 9.10.1 or later



1. [Remove the SnapMirror active sync configuration](#)
2. [Create a consistency group relationship then Initialize the consistency group relationship](#)

1. Under **Protection > Relationships**, find the SnapMirror active sync relationship on the consistency group. Select , then **Delete** to remove the SnapMirror active sync relationship.
2. [Delete the consistency group](#)
3. [Configure the consistency group](#)

## Planned failover unsuccessful

### Issue:

After executing the `snapmirror failover start` command, the output for the `snapmirror failover show` command displays a message indicates that a nondisruptive operation is in progress.

### Example:

```
Cluster1::> snapmirror failover show
Source Destination Error
Path Path Type Status start-time end-time Reason
-----
vs1:/cg/cg vs0:/cg/cg planned failed 10/1/2020 10/1/2020 SnapMirror
Failover cannot start because a volume move is running. Retry the command
once volume move has finished.
08:35:04
08:35:04
```

### Cause:

A planned failover cannot begin when a nondisruptive operation is in progress, including volume move, aggregate relocation, and storage failover.

### Solution:

Wait for the nondisruptive operation to complete and try the failover operation again.

## ONTAP Mediator not reachable or Mediator quorum status is false

### Issue:

After executing the `snapmirror failover start` command, the output for the `snapmirror failover show` command displays a message indicating that the ONTAP Mediator is not configured.

See [Configure the ONTAP Mediator and clusters for SnapMirror active sync](#).

**Example:**

```
Cluster1::> snapmirror failover show
Source Destination Error
Path Path Type Status start-time end-time Reason
-----
vs0:/cg/cg vs1:/cg/cg planned failed 10/1/2020 10/1/2020 SnapMirror
failover cannot start because the source-side precheck failed. reason:
Mediator not configured.
05:50:42 05:50:43
```

**Cause:**

Mediator is not configured or there are network connectivity issues.

**Solution:**

If the ONTAP Mediator is not configured, you must configure the ONTAP Mediator before you can establish a SnapMirror active sync relationship. Fix any network connectivity issues. Make sure Mediator is connected and quorum status is true on both the source and destination site using the `snapmirror mediator show` command. For more information, see [Configure the ONTAP Mediator](#).

**Example:**

```
cluster::> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
-----
10.234.10.143 cluster2 connected true
```

**Automatic unplanned failover not triggered on Site B****Issue:**

A failure on Site A does not trigger an unplanned failover on Site B.

**Possible cause #1:**

The ONTAP Mediator is not configured. To determine if this is the cause, issue the `snapmirror mediator show` command on the Site B cluster.

**Example:**

```
Cluster2::*> snapmirror mediator show
This table is currently empty.
```

This example indicates that ONTAP Mediator is not configured on Site B.

**Solution:**

Ensure that ONTAP Mediator is configured on both clusters, that the status is connected, and quorum is set to True.

**Possible cause #2:**

SnapMirror consistency group is out of sync. To determine if this is the cause, view the event log to view if the consistency group was in sync during the time at which the Site A failure occurred.

**Example:**

```
cluster::*> event log show -event *out.of.sync*

Time                Node                Severity            Event
-----
-----
10/1/2020 23:26:12  sti42-vsimsim-ucs511w ERROR              sms.status.out.of.sync:
Source volume "vs0:zrto_cg_556844_511u_RW1" and destination volume
"vs1:zrto_cg_556881_511w_DP1" with relationship UUID "55ab7942-03e5-11eb-
ba5a-005056a7dc14" is in "out-of-sync" status due to the following reason:
"Transfer failed."
```

**Solution:**

Complete the following steps to perform a forced failover on Site B.

1. Unmap all LUNs belonging to the consistency group from Site B.
2. Delete the SnapMirror consistency group relationship using the `force` option.
3. Enter the `snapmirror break` command on the consistency group constituent volumes to convert volumes from DP to R/W, to enable I/O from Site B.
4. Boot up the Site A nodes to create a zero RTO relationship from Site B to Site A.
5. Release the consistency group with `relationship-info-only` on Site A to retain common Snapshot copy and unmap the LUNs belonging to the consistency group.
6. Convert volumes on Site A from R/W to DP by setting up a volume level relationship using either the Sync policy or Async policy.
7. Issue the `snapmirror resync` to synchronize the relationships.
8. Delete the SnapMirror relationships with the Sync policy on Site A.
9. Release the SnapMirror relationships with Sync policy using `relationship-info-only true` on Site B.
10. Create a consistency group relationship from Site B to Site A.
11. Perform a consistency group resync from Site A, and then verify that the consistency group is in sync.
12. Rescan host LUN I/O paths to restore all paths to the LUNs.

**Link between Site B and mediator down and Site A down**

To check on the connection of the ONTAP Mediator, use the `snapmirror mediator show` command. If the connection status is unreachable and Site B is unable to reach Site A, you will have an output similar to the one below. Follow the steps in the solution to restore connection

## Example:

```
cluster::*> snapmirror mediator show
Mediator Address Peer Cluster      Connection Status Quorum Status
-----
10.237.86.17      C1_cluster      unreachable      true
SnapMirror consistency group relationship status is out of sync.

C2_cluster::*> snapmirror show -expand
Source          Destination Mirror Relationship Total
Last
Path           Type Path           State Status           Progress Healthy
Updated
-----
-----
vs0:/cg/src_cg_1 XDP vs1:/cg/dst_cg_1 Snapmirrored OutOfSync - false -
vs0:zrto_cg_655724_188a_RW1 XDP vs1:zrto_cg_655755_188c_DP1 Snapmirrored
OutOfSync - false -
vs0:zrto_cg_655733_188a_RW2 XDP vs1:zrto_cg_655762_188c_DP2 Snapmirrored
OutOfSync - false -
vs0:zrto_cg_655739_188b_RW1 XDP vs1:zrto_cg_655768_188d_DP1 Snapmirrored
OutOfSync - false -
vs0:zrto_cg_655748_188b_RW2 XDP vs1:zrto_cg_655776_188d_DP2 Snapmirrored
OutOfSync - false -
5 entries were displayed.

Site B cluster is unable to reach Site A.
C2_cluster::*> cluster peer show
Peer Cluster Name          Cluster Serial Number Availability
Authentication
-----
-----
C1_cluster                  1-80-000011          Unavailable          ok
```

## Solution

Force a failover to enable I/O from Site B and then establish a zero RTO relationship from Site B to Site A. Complete the following steps to perform a forced failover on Site B.

1. Unmap all LUNs belonging to the consistency group from Site B.
2. Delete the SnapMirror consistency group relationship using the force option.
3. Enter the SnapMirror break command (`snapmirror break -destination_path svm:_volume_`) on the consistency group constituent volumes to convert volumes from DP to RW, to enable I/O from Site B.

You must issue the SnapMirror break command for each relationship in the consistency group. For example, if there are three volumes in the consistency group, you will issue the command for each volume.

4. Boot up the Site A nodes to create a zero RTO relationship from Site B to Site A.

5. Release the consistency group with relationship-info-only on Site A to retain common Snapshot copy and unmap the LUNs belonging to the consistency group.
6. Convert volumes on Site A from RW to DP by setting up a volume level relationship using either Sync policy or Async policy.
7. Issue the `snapmirror resync` command to synchronize the relationships.
8. Delete the SnapMirror relationships with Sync policy on Site A.
9. Release the SnapMirror relationships with Sync policy using `relationship-info-only true` on Site B.
10. Create a consistency group relationship between Site B to Site A.
11. From the source cluster, resynchronize the consistency group. Verify the consistency group state is in sync.
12. Rescan the host LUN I/O paths to restore all paths to the LUNs.

## Link between Site A and mediator down and Site B down

When using SnapMirror active sync, you may lose connectivity between the ONTAP Mediator or your peered clusters. You can diagnose the issue by checking the connection, availability, and consensus status of the different parts of the SnapMirror active sync relationship then forcefully resuming connection.

**Table 1. Determining the cause**

What to check	CLI command	Indicator
Mediator from Site A	<code>snapmirror mediator show</code>	The connection status displays as unreachable
Site B connectivity	<code>cluster peer show</code>	Availability displays as unavailable
Consensus status of the SnapMirror active sync volume	<code>volume show volume_name -fields smbc-consensus</code>	The <code>sm-bc consensus</code> field displays <code>Awaiting-consensus</code>

For additional information about diagnosing and resolving this issue, refer to the Knowledge Base article [Link between Site A and Mediator down and Site B down when using SnapMirror active sync](#).

## SnapMirror delete operation fails when fence is set on destination volume

### Issue:

SnapMirror delete operation fails when any of the destination volumes have redirection fence set.

### Solution

Performing the following operations to retry the redirection and remove the fence from the destination volume.

- SnapMirror resync
- SnapMirror update

## Volume move operation stuck when primary is down

### Issue:

A volume move operation is stuck indefinitely in cutover deferred state when the primary site is down in a SnapMirror active sync relationship.

When the primary site is down, the secondary site performs an automatic unplanned failover (AUFO). When a volume move operation is in progress when the AUFO is triggered the volume move becomes stuck.

**Solution:**

Abort the volume move instance that is stuck and restart the volume move operation.

## **SnapMirror release fails when unable to delete Snapshot copy**

**Issue:**

The SnapMirror release operation fails when the Snapshot copy cannot be deleted.

**Solution:**

The Snapshot copy contains a transient tag. Use the `snapshot delete` command with the `-ignore-owners` option to remove the transient Snapshot copy.

```
snapshot delete -volume <volume_name> -snapshot <snapshot_name> -ignore-owners true -force true
```

Retry the `snapmirror release` command.

## **Volume move reference Snapshot copy shows as the newest**

**Issue:**

After performing a volume move operation on a consistency group volume, the volume move reference Snapshot copy might incorrectly display as the newest for the SnapMirror relationship.

You can view the newest Snapshot copy with the following command:

```
snapmirror show -fields newest-snapshot status -expand
```

**Solution:**

Manually perform a `snapmirror resync` or wait for the next automatic resync operation after the volume move operation completes.

## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

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