



# Introduction

## ONTAP 9

NetApp  
March 06, 2025

# Table of Contents

- Introduction..... 1
  - SnapMirror active sync overview in ONTAP ..... 1
    - Benefits ..... 1
    - Key concepts ..... 1
  - SnapMirror active sync architecture ..... 4
    - ONTAP Mediator ..... 5
    - Operations ..... 6
    - Symmetric active/active ..... 6
  - Use cases for SnapMirror active sync ..... 7
  - Deployment strategy and best practices for SnapMirror active sync ..... 8
    - SVM configuration ..... 8

# Introduction

## SnapMirror active sync overview in ONTAP

SnapMirror active sync (also referred to 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. There is no manual intervention or custom scripting required to trigger a failover with SnapMirror active sync.

Available beginning with ONTAP 9.9.1, SnapMirror active sync is supported on AFF clusters, All-Flash SAN Array (ASA) clusters, and 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](#). Symmetric active/active enable read and write I/O operations from both copies of a protected LUN with bidirectional synchronous replication so that both LUN copies can serve I/O operations locally. Prior to ONTAP 9.15.1, SnapMirror active sync only supports asymmetric active/active configurations, in which data on the secondary site is proxied to a LUN. Beginning with ONTAP 9.16.1 GA, symmetric active/active is also supported in four-node ONTAP clusters.



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 sites.
- 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 uses 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 define the failover units in a SnapMirror active sync relationship); 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 with 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 site 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 the 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

- **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 distances.
- **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.

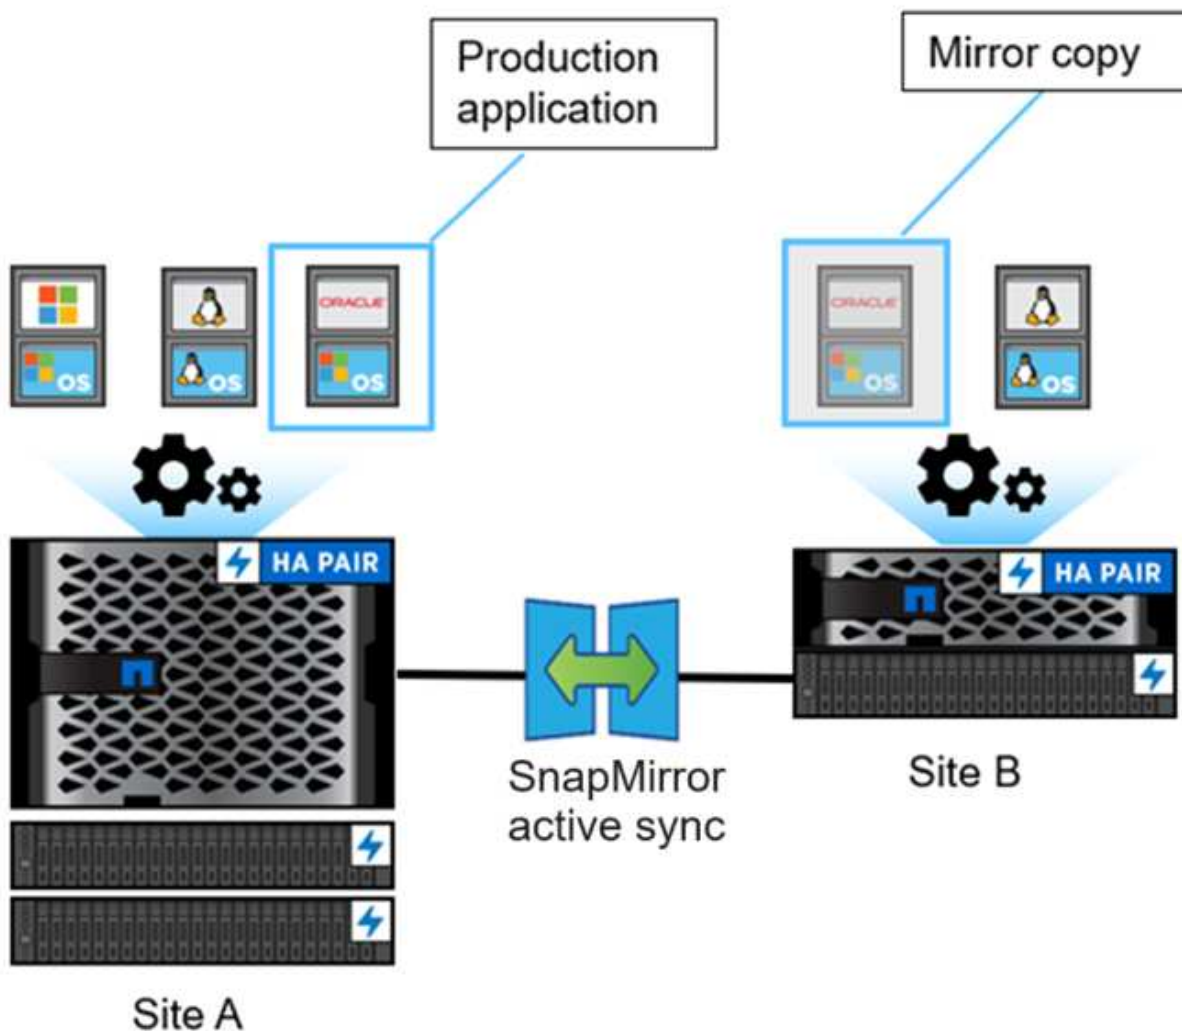
## Zero RTO

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 consistency group
- Secondary ONTAP cluster hosting the mirror consistency group
- 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 uses 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 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)



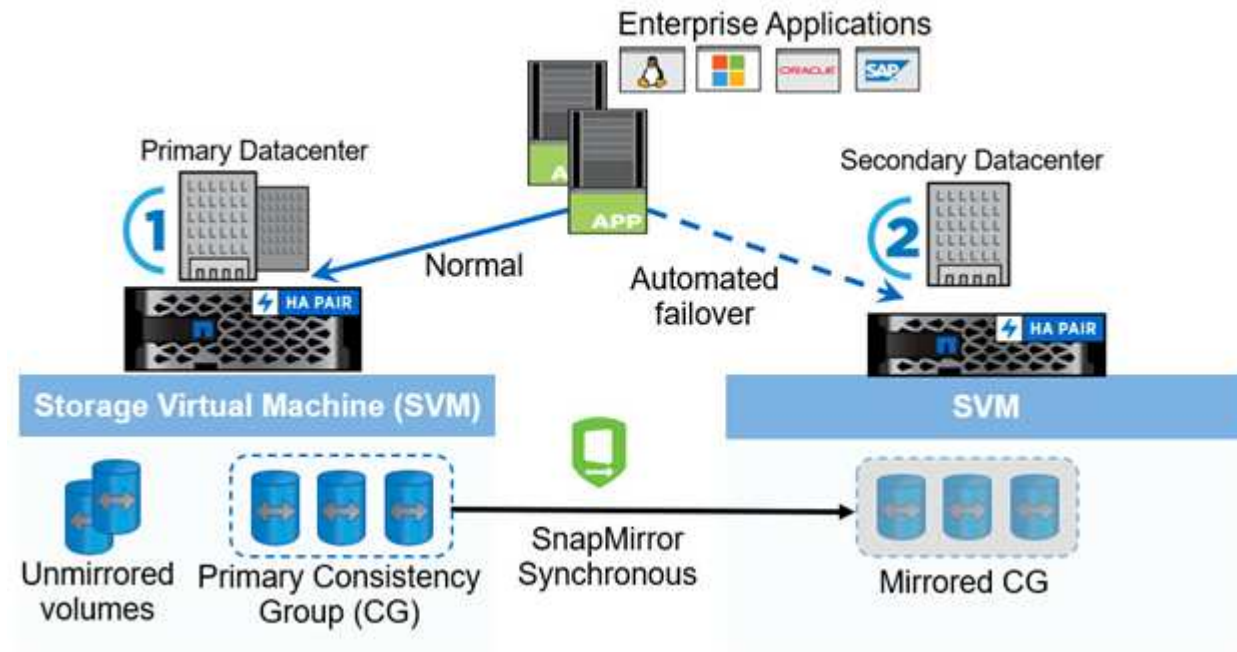
ONTAP Mediator 1.7 can manage ten cluster pairs for the purpose of business continuity.



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.

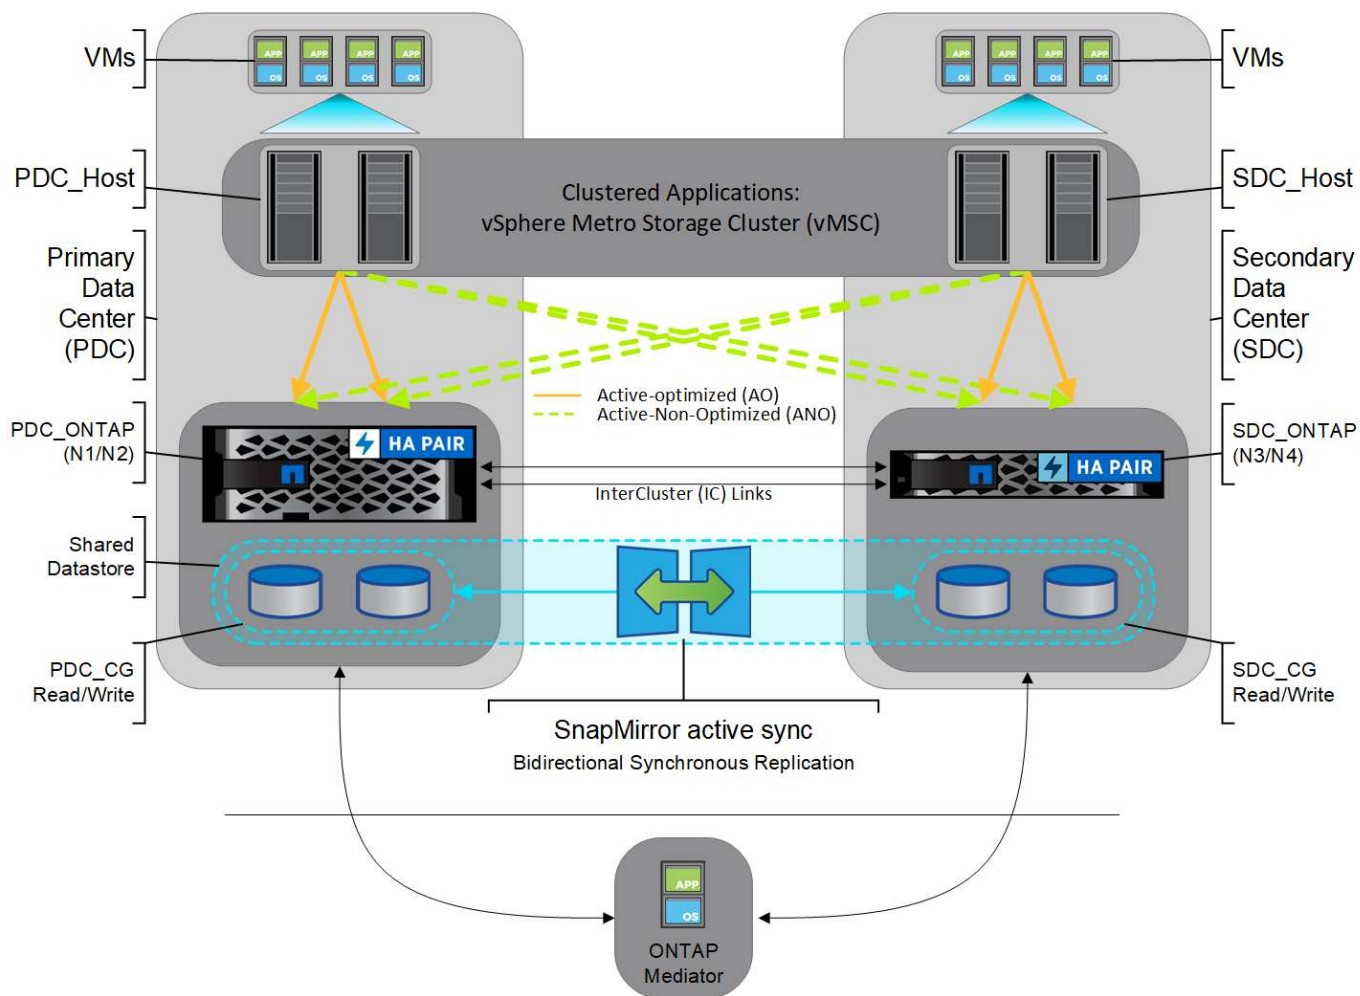
## Symmetric active/active

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. Beginning with ONTAP 9.16.1, symmetric active/active is supported on clusters with up to four nodes.





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 objective (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 transparent application failover (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. Beginning with ONTAP 9.16.1, symmetric active/active is supported in a configuration in four-node clusters, expanding from the two-node cluster limit in ONTAP 9.15.1.

### **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.

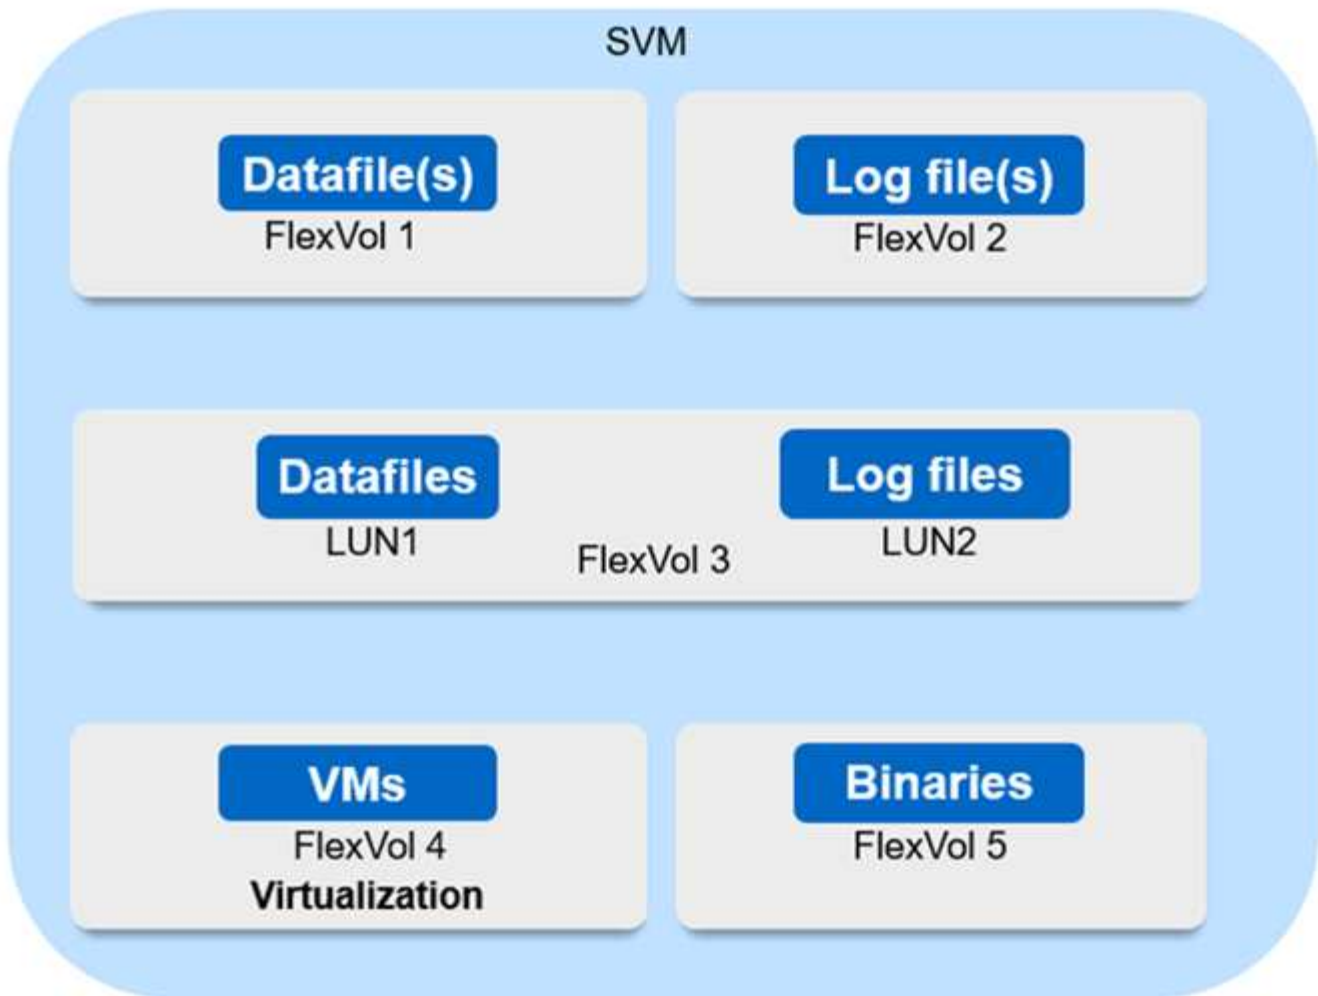
SnapMirror active sync enables mission-critical business services to continue operating even through a complete site failure, with TAF to the secondary copy. No manual intervention or no additional scripting are required to trigger this failover.

## **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.

## Copyright information

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

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

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

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

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

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

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