Technical Report

Three-Data-Center Disaster Recovery Using NetApp SnapMirror for ONTAP 9.7

Cheryl George, NetApp

April 2020 | TR-4832

Abstract

This document describes a three-data-center disaster recovery configuration using NetApp® SnapMirror® technology for replication.

TABLE OF CONTENTS

LIST OF FIGURES

1 Executive Summary

In today's constantly connected global business environment, companies expect rapid recovery of critical application data with zero data loss. Especially financial organizations who have zero tolerance for data loss or application unavailability and must they adhere to General Data Protection Regulation (GDPR) and other regulatory mandates. Having a disaster recovery strategy in place enables an organization to maintain or quickly resume mission-critical functions following a disruption, which can be anything that puts an organization's operations at risk, from a cyberattack to power outage, equipment failure to natural disaster. Organizations devise effective disaster recovery plans by keeping the below requirements in mind:

- Data should be recoverable in the event of catastrophic failure at one or more data centers (disaster recovery).
- Data should be replicated and distributed in an optimal way, taking into consideration major business criteria such as cost of storage, protection level against site failures, and so on.
- What needs to be protected and for how long should be driven by:
  - Shrinking recovery point objective (RPO), to achieve zero data loss
  - Near-zero recovery time objective (RTO), for faster recovery of business-critical applications in case of disaster

Previously, the disaster recovery plan included dual-site configuration in which IT operations would transfer activity to another site in case the primary data center went down due to some event. However, the distance between these two sites decides the amount of data at risk of permanent loss since data transfer latency increases as the distance increases. Closer proximity between sites imposes the risk of both sites becoming unavailable during an event such as a natural calamity. The recovery time depends on various factors such as the speed of the link between the local and remote recovery nodes, the amount of data to be recovered, and the complexity of the recovery process. An effective data protection strategy is vital to prevent operations from stalling and lost productivity and revenue, which can eventually damage company's reputation.

In order to sustain business operations in the event of such vast disasters with minimal data loss, the current requirement is a three-data-center configuration to achieve zero data loss and geographic dispersion: two data centers located close to one another, and the third located outside the region defined by the other two. Not forgetting that the goal also is to maximize investments and get the most out of the IT infrastructure. Furthermore, the ability to reuse a secondary facility for business intelligence or development and testing can turn a backup and disaster recovery solution into a business accelerator.

This technical report covers implementation of the three-data-center disaster recovery protection by using NetApp SnapMirror technology for replication, in which one leg uses SnapMirror Synchronous (SM-S) technology and the other uses SnapMirror Asynchronous technology for protection across your data fabric powered by NetApp. This is the most optimal way of protecting critical business data across several geographic locations.

2 Introduction

NetApp SnapMirror software replicates 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 replicates to one or more NetApp storage systems and continually updates the secondary data, keeping your data current and available whenever you need it. It also delivers powerful data management capabilities for virtualization, protecting your critical data while providing the flexibility to move data between locations and storage tiers, including cloud service providers. The three-data-center configuration for disaster recovery uses a combination of technologies to enable zero data loss at greater distances. While SnapMirror Synchronous (SM-S) is a disaster recovery solution through zero data loss, you are bound by the requirement of round-trip time (RTT) being less than or equal to 10 milliseconds (distance of ~150 km) since it otherwise impacts application performance. In order to overcome this distance limitation, use SnapMirror Asynchronous replication, which allows you to protect your business even in the event of large-scale disaster (for instance, an earthquake), which would damage both primary and local sites.

3 Disaster Impact Radius

There are different degrees of failure, ranging anywhere from network to device to storage to the complete site itself. The worst-case scenario for a business disaster is obvious – some catastrophic event such as a natural calamity, fire, or man-made disaster which can physically destroy your complete site at one location. The other failures range from partial loss or corruption of data, security breaches, temporary service outage, or even loss of key personnel can constitute a disaster that impacts your day-to-day operations. It's important to plan your data center configuration and technology used so that your business-critical data can be recovered based on each of these failures, as shown in Figure 1.

Figure 1) Disaster impact radius.

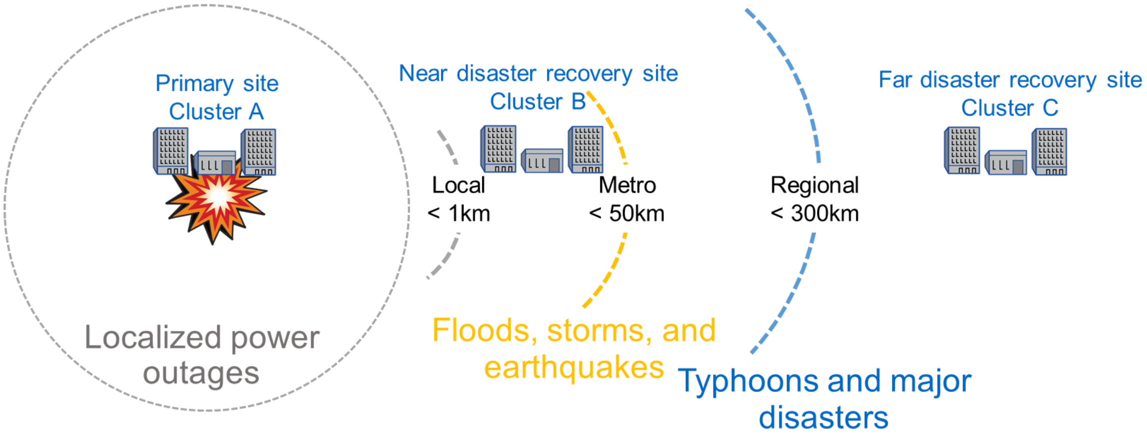
Primary site
Cluster A
Near disaster recovery site
Cluster B
Far disaster recovery site
Cluster C
Local
< 1km
Metro
< 50km
Regional
< 300km
Localized power outages
Floods, storms, and earthquakes
Typhoons and major disasters

4 Three-Data-Center Topology

The three-data-center configuration can be done in either a fan-out or cascade topology.

Fan-out topology consists of primary and near disaster recovery data centers within region sites (with replication network having <10ms round trip time), will synchronously replicate data between themselves to achieve zero RPO. And the primary asynchronously replicates data on a regular basis depending on the quantity of new data being generated to the remote or far disaster recovery out-of-region data center. If disaster strikes the primary site, the application can be started with zero data loss from the near disaster recovery data center, which would also take over asynchronous replication to the far disaster recovery data center, as shown in Figure 2.

Figure 2) Fan-out topology.

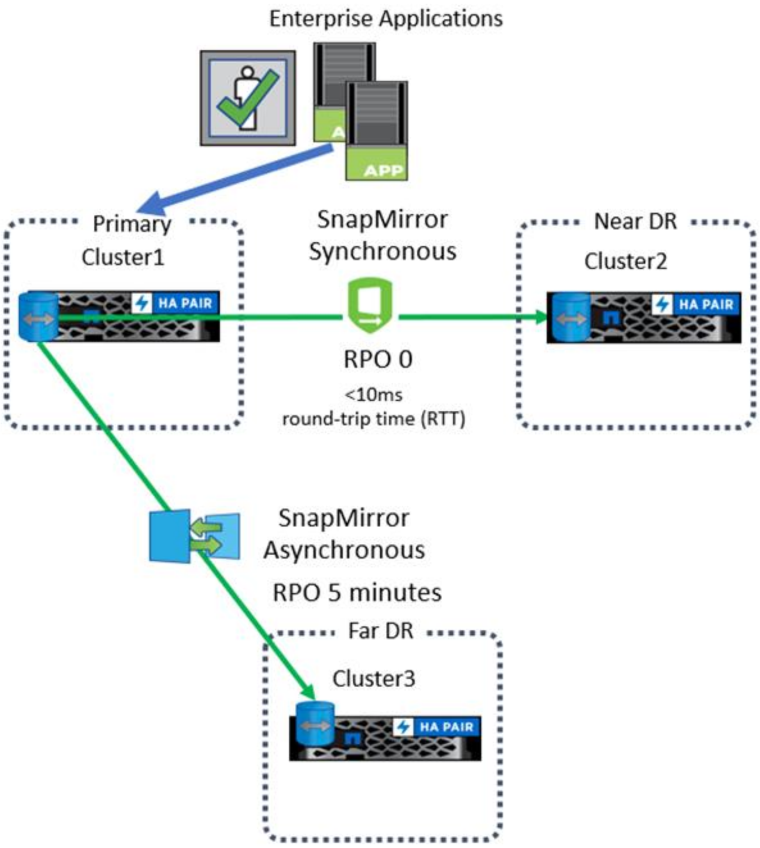

```
# Create the DP destination volume on Near DR site
cluster2::> volume create -vserver VServer -volume Volume_sync -aggregate aggr01 -size 20MB -type
DP

# Create the DP destination volume on Far DR site
cluster3::> volume create -vserver Vserver -volume Volume_async -aggregate aggr01 -size 20MB -
type DP

# Create the SnapMirror Synchronous relationship to Near DR site
cluster2::> snapmirror create -source-path Vserver:Volume -destination-path Vserver:Volume_sync -
type XDP -policy Sync

# Initialize the SnapMirror Synchronous relationship
cluster2::> snapmirror initialize -destination-path Vserver:Volume_sync

# Create the SnapMirror Asynchronous relationship to Far DR site
cluster3::> snapmirror create -source-path Vserver:Volume -destination-path Vserver:Volume_async
-type XDP -policy MirrorAllSnapshots -schedule 5min

# Initialize the SnapMirror Asynchronous relationship
Cluster3::> snapmirror initialize -destination-path Vserver:Volume_async
```

Another variation for this fan-out topology is the two-data-center colocation use case in small organizations. Here, synchronous replication occurs to another SVM within the same primary data center, while asynchronous replication continues to a far disaster recovery site. If the primary site fails, you can manually or script the failover to the far disaster recovery which requires an application restart and transaction logs to be applied to ensure RPO zero. While this can be relatively inexpensive, the three-data-center disaster recovery strategy ensures full redundancy of data.

Cascade topology is where the primary would synchronously replicate to the near disaster recovery site, and the near disaster recovery site would asynchronously replicate data to the far disaster recovery site. In case the intermediate near disaster recovery site goes down, you can always configure an asynchronous replication from the primary, directly to the far disaster recovery site over longer distance to ensure all the delta changes are being updated or replicated, as shown in Figure 3.

Figure 3) Cascade topology.

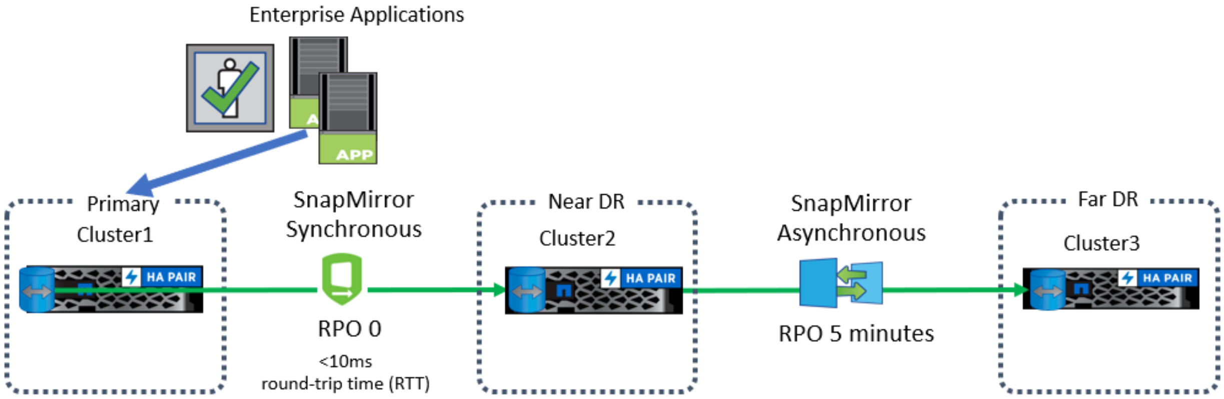

```
# Create the DP destination volume on Near DR site
cluster2::> volume create -vserver VServer -volume Volume_sync -aggregate aggr01 -size 20MB -type
DP

# Create the DP destination volume on Far DR site
cluster3::> volume create -vserver Vserver -volume Volume_async -aggregate aggr01 -size 20MB -
type DP

# Create the SnapMirror Synchronous relationship to Near DR site
cluster2::> snapmirror create -source-path Vserver:Volume -destination-path Vserver:Volume_sync -
type XDP -policy Sync

# Initialize the SnapMirror Synchronous relationship
cluster2::> snapmirror initialize -destination-path Vserver:Volume_sync

# Create the SnapMirror Asynchronous relationship to Far DR site
cluster3::> snapmirror create -source-path Vserver:Volume_sync -destination-path
Vserver:Volume_async -type XDP -policy MirrorAllSnapshots -schedule 5min

# Initialize the SnapMirror Asynchronous relationship
cluster3::> snapmirror initialize -destination-path Vserver:Volume_async
```

Note: You can create SnapMirror relationships only between clusters and storage virtual machine (SVM) that have NetApp ONTAP® peering setup.

Your choice of a failover site affects the capabilities of your disaster recovery methods, whether it entails restoring existing infrastructure, buying new infrastructure or moving to a production cloud. A disaster recovery plan is a necessity for business continuity. Going forward, this technical report covers the implementation of the three-data-center configuration by using the fan-out topology, implemented with simple CLI commands. ONTAP System Manager also can be used which further simplifies configuration and management.

5 Failover and Failback Preparation

Failover and failback strategies are integral part of an effective disaster recovery plan. Should a disaster occur, you need to be familiar with the order to bring your business back online and in production as quickly as possible. Make sure that the below are configured across all sites to enable the failover and failback procedures mentioned in the upcoming sections:

1. Mount any/all disaster recovery volumes to namespaces as needed.
   For more information about namespaces or how to mount volumes, see the Logical Storage Management Guide.
2. Create data LIFs as needed.
   For more information about creating data LIFs or other basic networking questions, see the Network Management Guide.
3. Configure the name resolution as needed.
4. Create/apply NAS/SAN configurations as needed:
   a. Create apply NFS export policies.
      For more information about configuring NFS, see the NFS Configuration Express Guide.
   b. Create CIFS server and CIFS shares.
      For more information about configuring CIFS and SMB, see the SMB/CIFS Configuration Express Guide.
   c. Create iGroups and LUN mappings.
      For more information about SAN configurations, see the SAN Configuration Guide.
5. Apply schedules and policies to the disaster recovery volumes as needed for client traffic/access.

6 Failover and Failback Operations

When you experience software or hardware failure, failover or switchover is the process of transferring mission-critical workloads from the production data center and recovering the system at a secondary disaster recovery location. What you need to be prepared for is that failover during an actual failover will be an unplanned event, which will temporarily halt I/O at the primary location, suspend mirroring activity that might be going on, and bring applications and I/O up from the remote location. The main goal of failover is to mitigate the negative impact of a disaster or service disruption on business services and customers.

After you recover your production site after a disaster and resolve any associated issues, you can transfer business operations back to the source location. Failback or switchback is normally a planned event, which is the process of resynchronizing the data back to the primary location, halting application I/O and application activity once again and cutting back over to the original location from the secondary or disaster recovery location.

7 Disaster Recovery Operations Between Primary and Near Disaster Recovery Sites

7.1 Failover from the Primary Site to the Near Disaster Recovery Site

In the event of a disaster at the primary site, failover to the near disaster recovery site (with zero data loss), as shown in Figure 4.

Figure 4) Fail over to the near disaster recovery site.

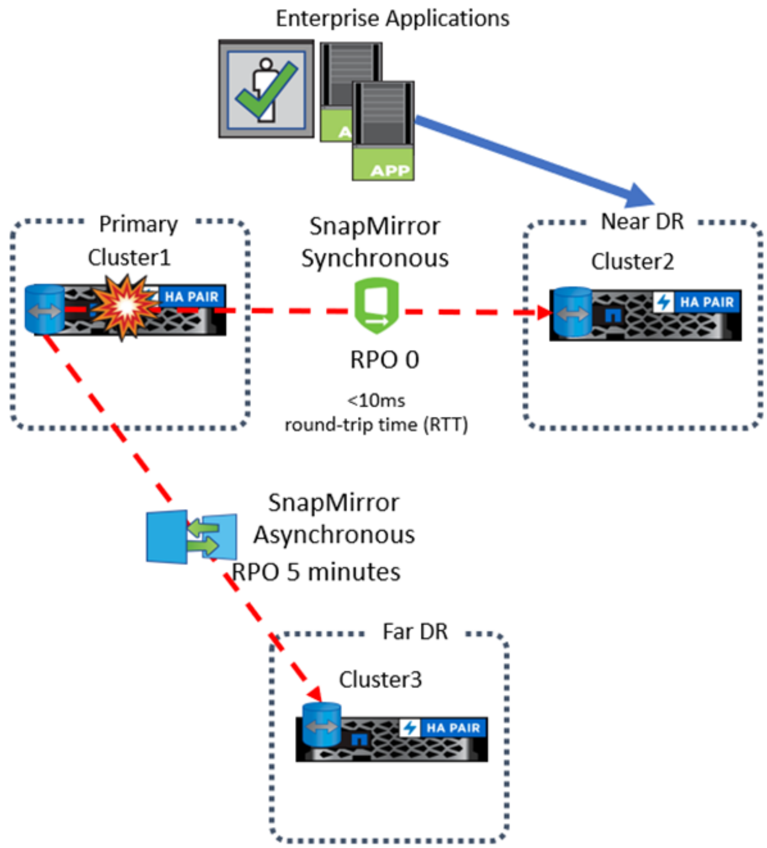

Unplanned failover to the near disaster recovery site consists of the following steps:

Note: Because the primary site is down, the volumes that were of type ‘dp’ on the destination cluster are now ‘rw’.

1. Mount LUNs from the near disaster recovery site on the application host.
2. Bring up the application at the near disaster recovery site with an RPO of zero.

After production is from the near disaster recovery site, complete the following steps:

1. Create a new SnapMirror Asynchronous relationship from the near disaster recovery site to the far disaster recovery site.

```
cluster3::> snapmirror create -source-path Vserver:Volume_sync -destination-path
Vserver:Volume_async -type xdp -policy MirrorAllSnapshots -schedule 5min
clusterc::> snapmirror resync -destination-path Vserver:Volume_async
```

Figure 5) Delta sync from the near disaster recovery site to the far disaster recovery site.

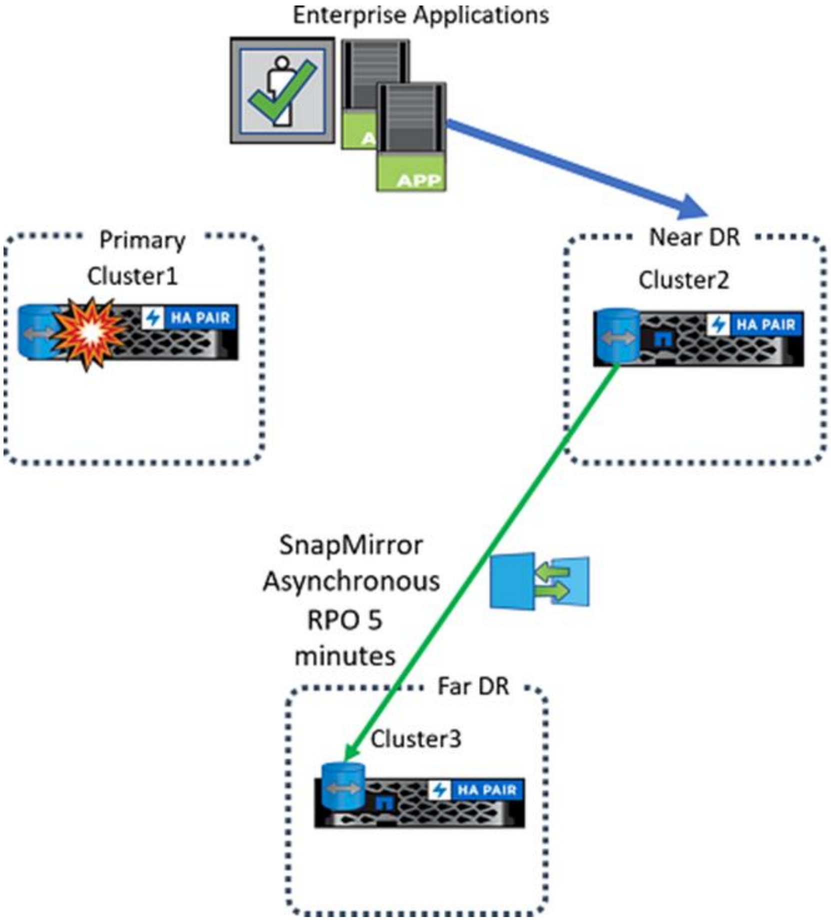

a. This process synchronizes the delta from the near disaster recovery site to the far disaster recovery site, to ensure the far disaster recovery site is up to date (RPO is five minutes).

b. Manually replay the transaction logs from the near disaster recovery site and perform an application-level recovery to ensure the delta is updated at the far disaster recovery site for an RPO of zero.

After the issue on the primary site is rectified, fail back to restore operations to the original source location, such as the primary site.

7.2 Failback from the Near Disaster Recovery Site to the Primary Site

To perform planned failback from the near disaster recovery site to the primary site, complete the following steps:

1. Stop all I/O to the volumes on the near disaster recovery site.
2. To synchronize the delta with the primary site to ensure zero data loss, reverse resync from the near disaster recovery site.

```
# Reverse resync from Near DR to Primary
cluster1::> snapmirror resync -destination-path Vserver:Volume# Check SnapMirror relationship
status is “Idle”
cluster2::> snapmirror show -destination-volume Volume
```

3. To restore production back to the primary site, mount the LUNs on the application host.

4. In addition, you are required to perform the following steps on the SnapMirror Asynchronous relationship that exists between the near disaster recovery and far disaster recovery sites:
    a. Quiesce the SnapMirror relationship to disable future transfers.
    b. Break this SnapMirror relationship.
    c. Release this relationship (with -relationship-info-only true).
    d. Delete the SnapMirror Asynchronous relationship.

```
# Quiesce the SnapMirror Asynchronous relationship
cluster3::> snapmirror quiesce -destination-path Vserver:Volume_async -source-path
Vserver:Volume_sync

# Break the SnapMirror relationship
cluster3::> snapmirror break -destination-path Vserver:Volume_async

# Release the SnapMirror relationship
cluster2::> snapmirror release -destination-path Vserver:Volume_async -relationship-info-only
true

# Delete the SnapMirror relationship
cluster2::> snapmirror delete -destination-path Vserver:Volume_async
```

5. To restore normal operational conditions, make sure the relationships are available or created as follows:
    a. SnapMirror Synchronous from the primary site to the near disaster recovery site.
    b. SnapMirror Asynchronous from the primary site to the far disaster recovery site.

Figure 6) Production operational again from the primary site.

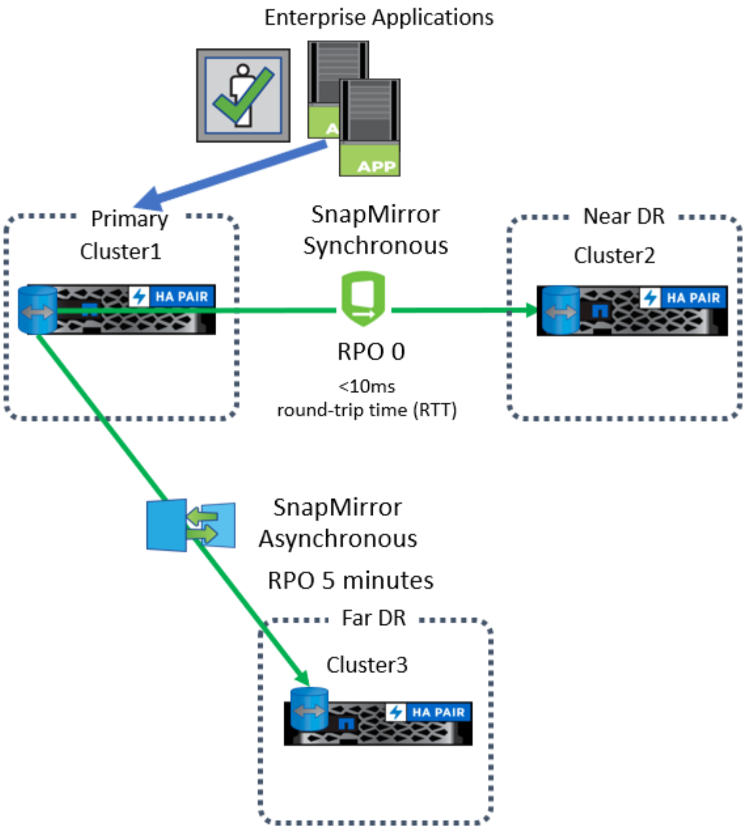

8 Disaster Recovery Operation Between Primary and Far Disaster Recovery Sites

8.1 Failover from the Primary Site to the Far Disaster Recovery Site

In the event of a disaster at the primary site, as illustrated in Figure 7, failover to the far disaster recovery site consists of the following steps:

Figure 7) Failover to the far disaster recovery site.

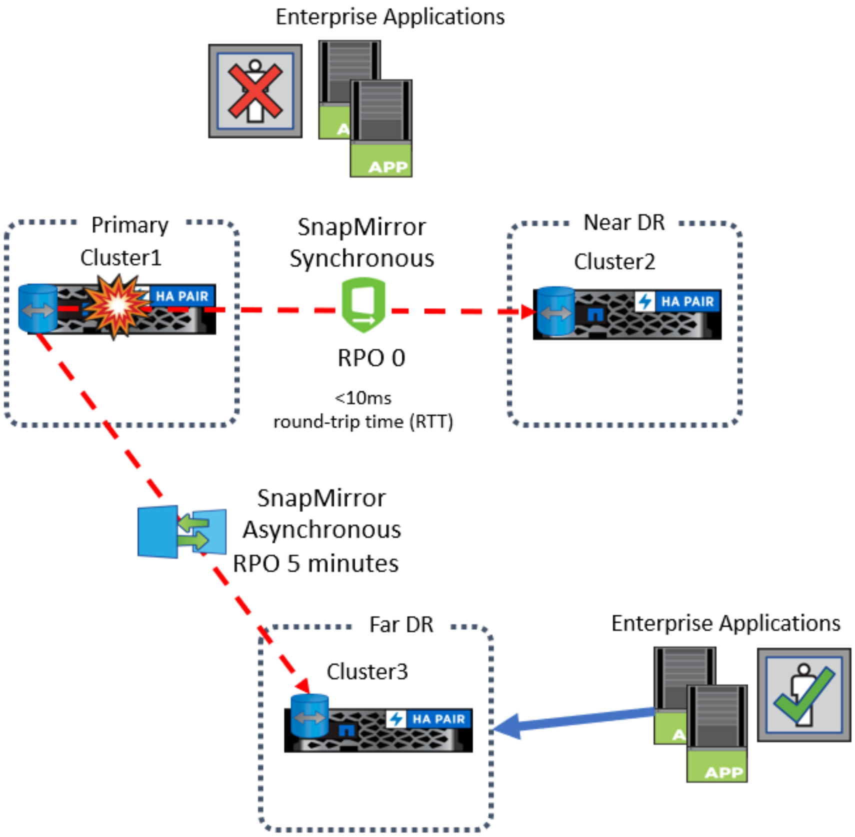

Note: Because the primary site is down, the volumes that were of type ‘dp’ on the destination cluster are now ‘rw’.

1. Mount the LUNs from the far disaster recovery site on the application host.
2. Bring up the application from the far disaster recovery site.
3. Manually replay the transaction logs from the near disaster recovery site and perform an application-level recovery at the far disaster recovery site for an RPO of zero.

   Note: If the disaster renders the primary and near disaster recovery sites inoperable, then your RPO will be five minutes at the far disaster recovery site since the transaction logs from the near disaster recovery site are not available.

After production is from the far disaster recovery site and the issue on the primary site is rectified, complete the following steps:

1. Create a new SnapMirror Asynchronous relationship from the far disaster recovery to the near disaster recovery (and the far disaster recovery to Primary) sites accordingly.

```
cluster2::> snapmirror create -source-path Vserver:Volume_async -destination-path
Vserver:Volume_sync -type xdp -policy MirrorAllSnapshots -schedule 5min
clusterb::> snapmirror resync -destination-path Vserver:Volume_sync

cluster1::> snapmirror create -source-path Vserver:Volume_async -destination-path
Vserver:Volume_sync -type xdp -policy MirrorAllSnapshots -schedule 5min
cluster1::> snapmirror resync -destination-path Vserver:Volume_sync
```

Figure 8) SnapMirror Asynchronous relationships from the far disaster recovery site to the near disaster recovery and primary sites.

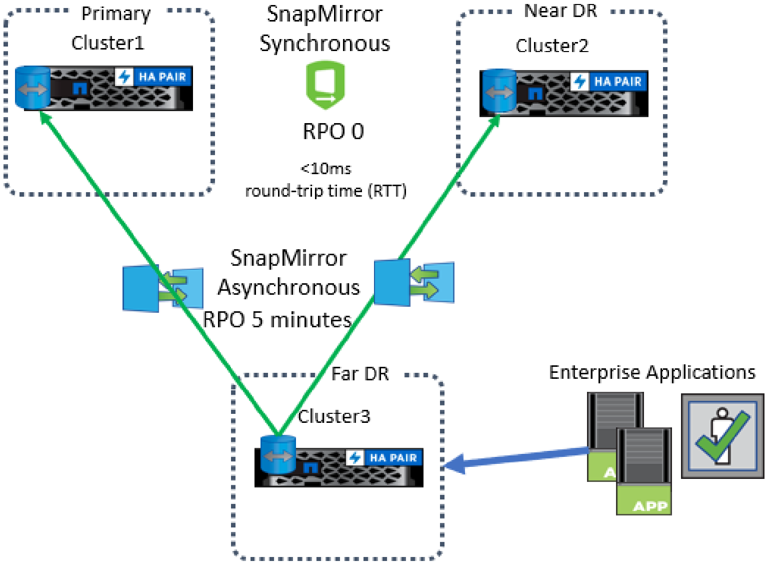

a. This process synchronizes the delta from the far disaster recovery to the near disaster recovery (and far disaster recovery to primary) sites, to ensure the near disaster recovery (and primary) site is up to date (with an RPO of five minutes).

b. Manually replay the transaction logs from the far disaster recovery site and perform an application-level recovery at the near disaster recovery (and primary) site for an RPO of zero.

Note: SnapMirror replication typically maintains two common Snapshot copies (by default) for a volume in order to be able to resynchronize. If you perform a disaster recovery drill within a short period of time, there is a chance that the common Snapshot copies will be overwritten and there exists no common Snapshot copy between the sites to allow resynchronization. In order to mitigate this issue during a short disaster recovery drill, you are required not set up the SnapMirror Asynchronous relationship from the far disaster recovery to the near disaster recovery sites just yet, only the asynchronous relationship from the far disaster recovery to DC. After failback to the primary (with the primary now in rw), only then can you create the SnapMirror Asynchronous relationship from the far disaster recovery to the near disaster recovery site.

2. You must fail back to restore operations at the original source location.

8.2 Failback from the Far Disaster Recovery Site to the Primary Site

To perform planned failback from the far disaster recovery site to the primary site, complete the following steps:

1. Stop all I/O to the volumes on the far disaster recovery site.
2. Unmount SAN LUNs from the far disaster recovery site on the application host.

3. Quiesce the existing SnapMirror Asynchronous relationship from the far disaster recovery to the near disaster recovery (and the far disaster recovery to the primary) sites to disable future transfers.
4. Break this SnapMirror relationship.
5. Release this relationship (with -relationship-info-only true).
6. Delete the SnapMirror Asynchronous relationship.

```
# Quiesce the SnapMirror Asynchronous relationship
cluster2::> snapmirror quiesce -destination-path Vserver:Volume -source-path Vserver:Volume_async

# Break the SnapMirror relationship
cluster2::> snapmirror break -destination-path Vserver:Volume

# Release the SnapMirror relationship
cluster3::> snapmirror release -destination-path Vserver:Volume -relationship-info-only true

# Delete the SnapMirror relationship
cluster3::> snapmirror delete -destination-path Vserver:Volume
```

Note: The cluster details (cluster 1 and cluster 3) in the above steps must be tweaked for the SnapMirror Asynchronous relationship between the far disaster recovery and the primary sites.

7. Mount the LUNs at the application host from the primary site.
8. Bring up the application from the primary site (currently with an RPO of five minutes).
9. Manually replay the transaction logs from the far disaster recovery site and perform an application-level recovery, to ensure zero data loss at the primary site.

   Production is now live from the primary site.
10. To restore normal operational conditions, make sure that the relationships are available or created as follows:
    a. SnapMirror Synchronous from the primary site to the near disaster recovery site.
    b. SnapMirror Asynchronous from the primary site to the far disaster recovery site.

Figure 9) Production operational again from the primary site.

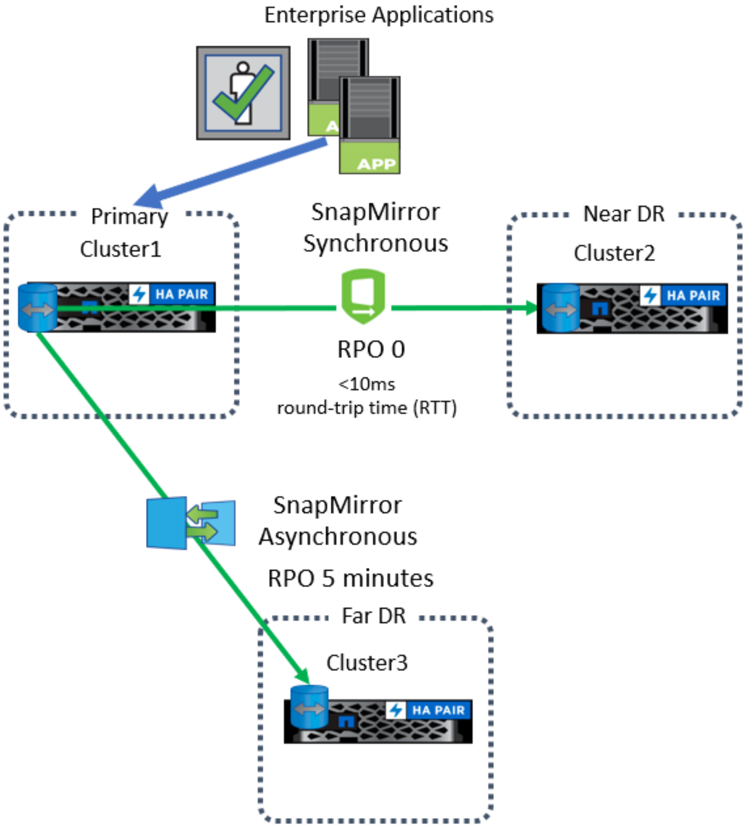

9 Best Practices for Disaster Recovery

Make sure there always exists a common Snapshot copy between the sites to allow them to resynchronize. NetApp recommends that you have a Snapshot copy schedule on the SnapMirror Asynchronous and that the synchronous relationships almost match or are more frequent.

- **Clearly write out your disaster recovery plan and define the scope.** If you don't have a written plan, you'll have to figure everything out in the middle of the emergency. And that practically guarantees that you'll make mistakes, spend more money than you need to, and stay offline longer than you would like. Also, the scope of your disaster recovery plan really determines which systems should be protected and identifies the expected results as well as any possible limitations.
- **Establish well-defined SLAs.** RTOs and RPOs should be established for applications, primarily based on the priorities of your organization during a disaster scenario. Increasing the frequency of the replication jobs considerably improves your RPO. Shorter RTOs should be assigned to the components of the highest priority, which should be recovered first.
- **Plan, review, and compliance** Some organizations operate with very sensitive and confidential data that require them to comply with regulations such as HIPAA or PCI DSS. If this is applicable to you, then you must verify whether your disaster recovery strategies for failover and failback operations meet the applicable security standards.
- **Delegate responsibilities efficiently.** Decide who is responsible for failover and failback operations. Designate and train members of the recovery team with specific responsibilities assigned. Train IT staff in failover and failback operations.
- **Check licensing.** Review your software documentation and determine whether there are any licensing limitations in your application stack that requires them to be addressed beforehand.

- **Test your plan.** While we can put together impeccable disaster recovery plans, every organization should have a schedule to test its disaster recovery plan periodically under realistic conditions. That means creating conditions where you attempt to bring your systems online after a power outage, for example. This also ensures that you can recover business-critical applications even if the catastrophes impact multiple locations.
- **Take advantage of automation and orchestration tools.** To simplify deployment, configuration, and management of data protection across multiple data centers.

10 Conclusion

In today's complex environments, organizations can face long-term effects of data loss, loss of customers, as well as the possibility of nonrecovery from disaster events. To avoid losing revenue, clients, and even production, be sure to design a comprehensive disaster recovery plan to protect your organization against any unpredicted disaster that might disrupt your production environment. Formulating and implementing the three-data-center disaster recovery infrastructure, you can enable application recovery for region-spanning disasters which can help organizations sustain operations and also mitigate the potential negative impact of future events.

Call to Action

Contact your NetApp sales office to engage Professional Services to assist with three-way disaster recovery planning, implementation, and automation.

For a hands-on lab experience, see the NetApp Lab on Demand at Early Adopter Lab for SnapMirror Sync 3-way DR v1.0.

Where to Find Additional Information

To learn more about the information that is described in this document, review the following documents and/or websites:

- Guidelines for SnapMirror Disaster Recovery Testing and Reverse SnapMirror in Clustered ONTAP
https://kb.netapp.com/app/answers/answer_view/a_id/1073855
- TR-4015: SnapMirror Configuration Best Practices for ONTAP 9.7
https://www.netapp.com/us/media/tr-4015.pdf
- TR-4733: SnapMirror Synchronous for ONTAP 9.7
https://www.netapp.com/us/media/tr-4733.pdf
- Early Adopter Lab on Demand
https://labondemand.netapp.com/lod/SM-S-DR
- Disaster Recovery Solution Videos
https://www.youtube.com/playlist?list=PLdXI3bZJEw7myCEHPWY6KKx56KAXL8war
- SM-S Lab on Demand
https://labondemand.netapp.com/sm-s
- ONTAP Documentation Center
https://docs.netapp.com/ontap-9/index.jsp
- ONTAP and ONTAP System Manager Documentation Resources page
https://www.netapp.com/us/documentation/ontap-and-oncommand-system-manager.aspx

Version History

Version	Date	Document Version History
Version 1.0	April 2020	Initial release.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

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.

Trademark Information

TR-4832-0420