



Technical Report

# SAP HANA Disaster Recovery with Storage Replication

Nils Bauer, NetApp  
July 2020 | TR-4646

## Abstract

This document is an overview of the options for disaster recovery protection for SAP HANA. It includes detailed setup information and a use case description of a three-site disaster recovery solution based on synchronous and asynchronous NetApp® SnapMirror® storage replication. The described solution uses NetApp SnapCenter® with the SAP HANA plug-in to manage database consistency.

## TABLE OF CONTENTS

|   |           |
|---|-----------|
| <b>SAP HANA Disaster Recovery with Storage Replication .....</b>      | <b>1</b>  |
| Nils Bauer, NetApp July 2020   TR-4646 .....                          | 1         |
| <b>1 Data Protection Overview .....</b>                               | <b>6</b>  |
| 1.1 Business Application Requirements.....                            | 6         |
| 1.2 High Availability.....  | 6         |
| 1.3 Logical Corruption.....   | 7         |
| 1.4 Backups .....   | 7         |
| 1.5 Synchronous or Asynchronous Data Replication .....                | 8         |
| 1.6 HANA System Replication with or Without Data Preload .....        | 8         |
| <b>2 Disaster Recovery Solution Comparison.....</b>                   | <b>8</b>  |
| 2.1 Backup and Recovery.....  | 9         |
| 2.2 SAP HANA Disaster Recovery Using Storage Replication .....        | 11        |
| 2.3 SAP HANA System Replication .....                                 | 13        |
| 2.4 Summary of Disaster Recovery Solutions.....                       | 15        |
| <b>3 SAP HANA Disaster Recovery with SnapMirror Replication .....</b> | <b>16</b> |
| 3.1 Lab Setup.....  | 17        |
| 3.2 SnapCenter Support for Synchronous SnapMirror.....                | 18        |
| 3.3 Configuration Steps for Synchronous SnapMirror Replication .....  | 19        |
| 3.4 Configuration Steps for Asynchronous SnapMirror Replication ..... | 21        |
| <b>4 Overview of Disaster Recovery Testing .....</b>                  | <b>27</b> |
| <b>5 Synchronous SnapMirror Disaster Recovery Testing .....</b>       | <b>28</b> |
| 5.1 Prepare the Target Host .....                                     | 29        |
| 5.2 Create Snapshot Backups at the Source Storage System.....         | 31        |
| 5.3 Create FlexClone Volumes at the Disaster Recovery Storage.....    | 31        |
| 5.4 Mount FlexClone Volumes at the Target Host .....                  | 33        |
| 5.5 Start the HANA Database .....                                     | 33        |
| 5.6 Cleanup Operations .....  | 34        |
| <b>6 Asynchronous SnapMirror Disaster Recovery Testing.....</b>       | <b>34</b> |
| 6.1 Prepare the Target Host .....                                     | 35        |
| 6.2 Create FlexClone Volumes at the Disaster Recovery Storage.....    | 37        |
| 6.3 Mount the FlexClone Volumes at the Target Host .....              | 39        |
| 6.4 Check Consistency of Latest Log Backups.....                      | 40        |
| 6.5 Recover the HANA Database .....                                   | 41        |

|           |  |           |
|-----------|--|-----------|
| <b>7</b>  | <b>Overview of Disaster Recovery Failover</b>                  | <b>52</b> |
| <b>8</b>  | <b>Synchronous SnapMirror Disaster Recovery Failover</b>       | <b>52</b> |
| 8.1       | Prepare Target Host  | 53        |
| 8.2       | Break SnapMirror Relationships                                 | 54        |
| 8.3       | Mount Volumes at Target Host                                   | 58        |
| 8.4       | Start the HANA Database  | 58        |
| <b>9</b>  | <b>Asynchronous SnapMirror Disaster Recovery Failover</b>      | <b>58</b> |
| 9.1       | Prepare Target Host  | 59        |
| 9.2       | Break SnapMirror Relationships                                 | 61        |
| 9.3       | Restore Data Volume to Latest SnapCenter Backup                | 65        |
| 9.4       | Mount FlexClone Volumes at Target Host                         | 66        |
| 9.5       | Check Consistency of Latest Log Backups                        | 67        |
| 9.6       | Recovery of the HANA Database                                  | 67        |
| <b>10</b> | <b>Different Steps Required in a Fibre Channel Environment</b> | <b>67</b> |
| 10.1      | Mapping LUNs to Disaster Recovery Server                       | 67        |
| 10.2      | Mount File Systems   | 68        |
| 10.3      | Cleanup Operation  | 69        |
|           | <b>Where to Find Additional Information</b>                    | <b>70</b> |
|           | <b>Version History</b>   | <b>70</b> |

## LIST OF TABLES

|          |   |    |
|----------|---|----|
| Table 1) | Disaster recovery solution comparison                     | 15 |
| Table 2) | Comparison of asynchronous storage replication approaches | 17 |
| Table 3) | Replication relationships                                 | 18 |
| Table 4) | Disaster recovery testing – required steps                | 28 |
| Table 5) | Volume names of FlexClone volumes                         | 30 |
| Table 6) | Volume names of FlexClone copies                          | 36 |
| Table 7) | Disaster recovery testing – required steps                | 52 |
| Table 8) | Volume names of FlexClone volumes                         | 54 |
| Table 9) | Volume names of FlexClone volumes                         | 60 |

## LIST OF FIGURES

|           |                            |   |
|-----------|----------------------------|---|
| Figure 1) | Data protection overview   | 6 |
| Figure 2) | Disaster recovery overview | 9 |

|  |    |
|--|----|
| Figure 3) NetApp storage-based backups .....   | 11 |
| Figure 4) Disaster recovery with synchronous and asynchronous SnapMirror.....                    | 12 |
| Figure 5) NetApp MetroCluster combined with asynchronous SnapMirror.....                         | 13 |
| Figure 6) SAP HANA System Replication. ....  | 14 |
| Figure 7) SAP HANA disaster recovery with SnapMirror replication. ....                           | 16 |
| Figure 8) Lab setup. ....  | 17 |
| Figure 9) SnapMirror update warning message. ....  | 19 |
| Figure 10) Synchronous SnapMirror - Create Protection Relationship.....                          | 20 |
| Figure 11) Synchronous SnapMirror volume relationships.....                                      | 20 |
| Figure 12) SnapCenterSync Protection Policies.....   | 21 |
| Figure 13) Combined backup and disaster recovery replication.....                                | 22 |
| Figure 14) Replication of data volume only. ....   | 23 |
| Figure 15) Asynchronous SnapMirror - Create Protection Relationship – Configuration Details..... | 23 |
| Figure 16) Asynchronous SnapMirror - Create Protection Relationship – Schedule. ....             | 24 |
| Figure 17) Asynchronous SnapMirror Volume Relationships. ....                                    | 24 |
| Figure 18) SnapCenter topology view for HANA database.....                                       | 26 |
| Figure 19) SnapCenter topology view for HANA shared volume.....                                  | 26 |
| Figure 20) SnapCenter policies. ....   | 27 |
| Figure 21) Asynchronous SnapMirror policy.....   | 27 |
| Figure 22) Synchronous SnapMirror disaster recovery testing.....                                 | 29 |
| Figure 23) Create FlexClone volume using crash-consistent Snapshot backup.....                   | 32 |
| Figure 24) List of FlexClone volumes. ....   | 32 |
| Figure 25) Junction paths configuration. ....  | 32 |
| Figure 26) Asynchronous SnapMirror disaster recovery testing. ....                               | 35 |
| Figure 27) Create a FlexClone volume based on an application-consistent SnapCenter backup. ....  | 37 |
| Figure 28) Create a FlexClone copy of the log backup volume. ....                                | 38 |
| Figure 29) List of FlexClone volumes. ....   | 38 |
| Figure 30) Junction path configuration. ....   | 39 |
| Figure 31) Recovery to a specific data backup: Recovery Type.....                                | 42 |
| Figure 32) Recovery with log backups: Recovery Type. ....  | 42 |
| Figure 33) Recovery to a specific data backup: Specify backup catalog location.....              | 43 |
| Figure 34) Recovery with log backups: Backup catalog location.....                               | 43 |
| Figure 35) Recovery to a specific data backup: Select Destination Type.....                      | 43 |
| Figure 36) Recovery with log backups: Backup selection.....                                      | 44 |
| Figure 37) Recovery with log backups: Log backup location.....                                   | 44 |
| Figure 38) Recovery to a specific data backup: Other Settings.....                               | 45 |
| Figure 39) Recovery with log backups: Other Settings.....  | 45 |
| Figure 40) Recovery to a specific data backup: Review Recovery Settings.....                     | 46 |
| Figure 41) Recovery with log backups: Review Recovery Settings.....                              | 46 |
| Figure 42) Recovery to a specific data backup: Recovery Execution Summary.....                   | 47 |

|  |    |
|--|----|
| Figure 43) Recovery of Tenant Database in DRS. ....                                | 47 |
| Figure 44) Recovery to a specific data backup: Recovery Type.....                  | 47 |
| Figure 45) Recovery with log backups: Specify Recovery Type.....                   | 48 |
| Figure 46) Recovery to a specific data backup:Specify backup catalog location..... | 48 |
| Figure 47) Recovery with log backups: Specify backup catalog location. ....        | 48 |
| Figure 48) Recovery to a specific data backup: Select Destination Type.....        | 49 |
| Figure 49) Recovery with log backups: Select a Backup.....                         | 49 |
| Figure 50) Recovery to a specific data backup: Other Settings.....                 | 50 |
| Figure 51) Recovery to a specific data backup: Review Recovery Settings.....       | 50 |
| Figure 52) Recovery Execution Summary.....   | 51 |
| Figure 53) HANA Studio: System recovered.....                                      | 51 |
| Figure 54) Synchronous SnapMirror disaster recovery failover.....                  | 53 |
| Figure 55) SnapMirror quiesce operation—step 1.....                                | 55 |
| Figure 56) SnapMirror quiesce operation—step 2.....                                | 55 |
| Figure 57) All SnapMirror target volumes are quiesced.....                         | 56 |
| Figure 58) SnapMirror break operation—step 1.....                                  | 56 |
| Figure 59) SnapMirror break operation—step 2.....                                  | 57 |
| Figure 60) All SnapMirror target volumes are broken off.....                       | 57 |
| Figure 61) Junction path configuration.....  | 58 |
| Figure 62) Synchronous SnapMirror disaster recovery failover.....                  | 59 |
| Figure 63) SnapMirror quiesce operation—step 1.....                                | 61 |
| Figure 64) SnapMirror quiesce operation—step 2.....                                | 62 |
| Figure 65) All SnapMirror target volumes are quiesced.....                         | 62 |
| Figure 66) SnapMirror break operation—step 1.....                                  | 63 |
| Figure 67) SnapMirror break operation—step 2.....                                  | 63 |
| Figure 68) All SnapMirror target volumes are broken off.....                       | 64 |
| Figure 69) Junction path configuration.....  | 64 |

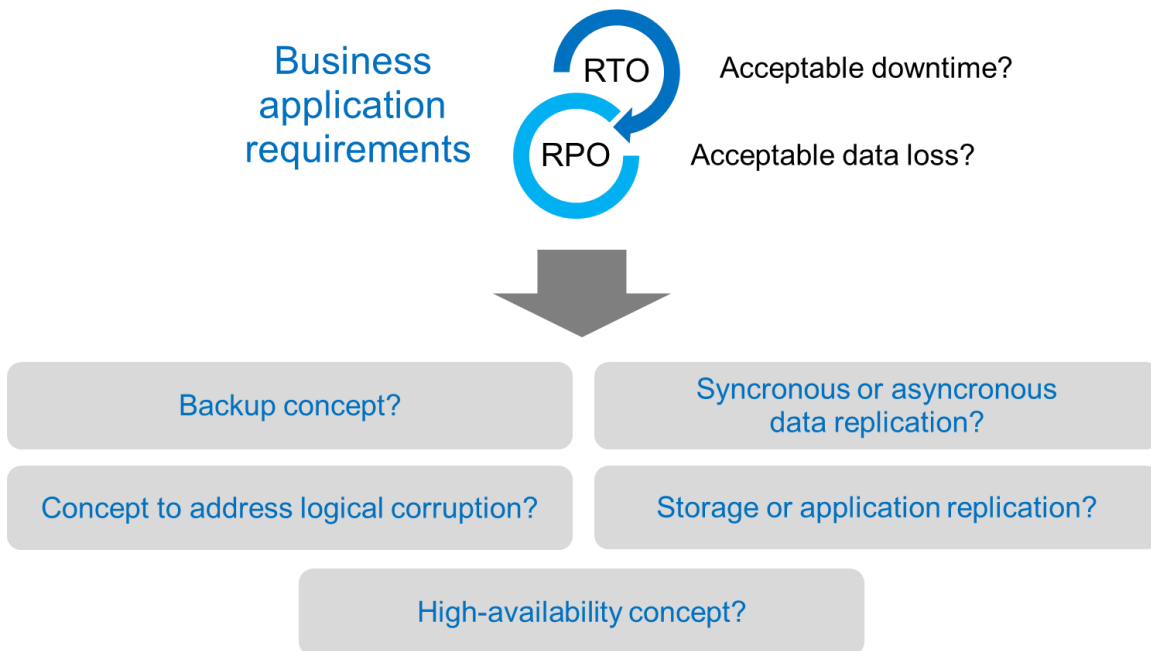
# 1 Data Protection Overview

Studies have shown that business application downtime has a significant negative impact on the business of enterprises. In addition to the financial impact, downtime can also damage the company's reputation, staff morale, and customer loyalty. Surprisingly, not all companies have a comprehensive disaster recovery policy.

Running SAP HANA on NetApp® storage gives customers access to additional features that extend and improve the built-in data protection and disaster recovery capabilities of SAP HANA. This overview section explains these options to help customers select options that support their business needs.

To develop a comprehensive disaster recovery policy, customers must understand the business application requirements and technical capabilities they need for data protection and disaster recovery (Figure 1).

Figure 1) Data protection overview.



## 1.1 Business Application Requirements

There are two key indicators for business applications:

- The recovery point objective (RPO), or the maximum tolerable data loss
- The recovery time objective (RTO), or the maximum tolerable business application downtime

These requirements are defined by the kind of application used and the nature of your business data. The RPO and the RTO might differ if you are protecting against hardware failures at a single site. They might also differ if you are preparing for catastrophic disasters such as the loss of a complete data center. It's important to evaluate the business requirements that define the RPO and RTO, because these requirements have a significant impact on the technical options that are available.

## 1.2 High Availability

The infrastructure hardware for SAP HANA, such as server, network, and storage, must have redundant components to make sure that there is no single point of failure.

Failures on the network side are typically addressed with redundant network paths to different network components. Storage systems usually offer failover capabilities to another storage controller. Therefore, failures in these redundant systems should not cause any application downtime.

To provide high availability on the server and application side, standby SAP HANA hosts can be configured for built-in high availability with an SAP HANA multiple-host system. If a server or an SAP HANA service fails, the SAP HANA service fails over to the standby host, which causes application downtime.

If application downtime is not acceptable in the case of server or application failure, you can also use SAP HANA system replication as a high-availability solution that enables failover in a very short time frame. SAP system replication is discussed in detail in section 2.3, SAP HANA System Replication. SAP HANA customers use HANA system replication not only to address high availability for unplanned failures, but also to minimize downtime for planned operations, such as HANA software upgrades.

### 1.3 Logical Corruption

Logical corruption can be caused by software errors, human errors, or sabotage. Unfortunately, logical corruption often cannot be addressed with standard high-availability and disaster recovery solutions. As a result, depending on the layer, application, file system, or storage where the logical corruption occurred, RTO and RPO requirements can sometimes not be fulfilled.

The worst case is a logical corruption in an SAP application. SAP applications often operate in a landscape in which different applications communicate with each other and exchange data. Therefore, restoring and recovering an SAP system in which a logical corruption has occurred is not the recommended approach. Restoring the system to a point in time before the corruption occurred results in data loss, so the RPO becomes larger than zero. Also, the SAP landscape would no longer be in sync and would require additional postprocessing.

Instead of restoring the SAP system, the better approach is to try to fix the logical error within the system, by analyzing the problem in a separate repair system. Root cause analysis requires the involvement of the business process and application owner. For this scenario, you create a repair system (a clone of the production system) based on data stored before the logical corruption occurred. Within the repair system, the required data can be exported and imported to the production system. With this approach, the productive system does not need to be stopped, and, in the best-case scenario, no data or only a small fraction of data is lost.

### 1.4 Backups

Backups are created to enable restore and recovery from different point-in-time datasets. Typically, these backups are kept for a couple of days to a few weeks.

Depending on the kind of corruption, restore and recovery can be performed with or without data loss. If the RPO must be zero, even when the primary and backup storage is lost, backup must be combined with synchronous data replication.

The RTO for restore and recovery is defined by the required restore time, the recovery time (including database start), and the loading of data into memory. For large databases and traditional backup approaches, the RTO can easily be several hours, which might not be acceptable. To achieve very low RTO values, a backup must be combined with a hot-standby solution, which includes preloading data into memory.

In contrast, a backup solution must address logical corruption, because data replication solutions cannot cover all kinds of logical corruption. For details, see section 2.1, Backup and Recovery.

## 1.5 Synchronous or Asynchronous Data Replication

The RPO primarily determines which data replication method you should use. If the RPO must be zero, even when the primary and backup storage is lost, the data must be replicated synchronously. However, there are technical limitations for synchronous replication, such as the distance between the two data centers. In most cases, synchronous replication is not appropriate for distances greater than 100km. Indeed, synchronous replication over a large distance places significant demands on the network infrastructure between the two data centers and therefore can be very expensive.

If a larger RPO is acceptable, asynchronous replication can be used over large distances. The RPO in this case is defined by the replication frequency.

## 1.6 HANA System Replication with or Without Data Preload

The startup time for an SAP HANA database is much longer than that of traditional databases because a large amount of data must be loaded into memory before the database can provide the expected performance. Therefore, a significant part of the RTO is the time needed to start the database. With any storage-based replication, the SAP HANA database must be started in case of failover to the disaster recovery site.

SAP HANA system replication offers an operation mode in which the data is preloaded and continuously updated at the disaster recovery server. This mode enables very low RTO values, but it also requires a dedicated server that is only used to receive the replication data from the source system.

## 2 Disaster Recovery Solution Comparison

A comprehensive disaster recovery solution must enable customers to recover from a complete failure of the primary site. Therefore, data must be transferred to a secondary site, and a complete infrastructure is necessary to run the required production SAP HANA systems in case of site failure. Depending on the availability requirements of the application and the kind of disaster you want to be protected from, a two-site or three-site disaster recovery solution must be considered.

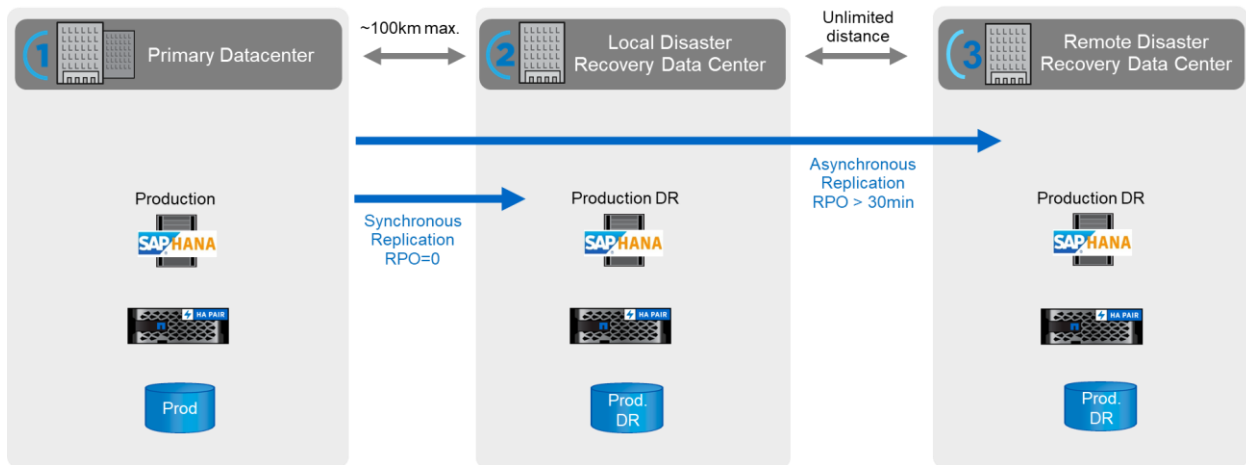
Figure 2 shows a typical three-site configuration, where the secondary data center is close to the primary. This allows you to replicate the data synchronously and to achieve an RPO of zero.

In addition, the data is also replicated asynchronously to a third, remote data center to be protected from disasters, which would influence operations in the primary as well as in the local disaster recovery data center. The minimum achievable RPO depends on the data replication frequency, which is limited by the available bandwidth between the primary and the remote disaster recovery data center. A typical minimal RPO is in the range of 30 minutes to multiple hours. Instead of running a third remote disaster recovery data center, a cloud provider IaaS (AWS, Azure, Google Cloud) could be used to provide the required resources.

This document discusses the different implementation options of a three-site disaster recovery solution. All of the use cases and scenarios described are also valid for two-site disaster recovery solutions, either with just a local or just a remote disaster recovery data center.



Figure 2) Disaster recovery overview.



In addition to the RPO and RTO, there are additional infrastructure and business metrics that can help you identify the best implementation for your needs. Additional metrics include:

- Resource usage at the local and remote disaster recovery site during standard operations. Are the servers available for different workloads, or are they allocated explicitly for the disaster recovery setup?
  - Servers at the disaster recovery site are available for dev/test during standard operations, and the database data is not preloaded into memory.
  - Servers at the disaster recovery site are exclusively allocated for disaster recovery, and the database data is preloaded into memory.
  - Costs for dedicated disaster recovery servers.
- Distance between the sites:
  - Physical limitations for synchronous replication because of increasing latency.
  - Availability and costs for the network connectivity between the sites.
- Impact on the required bandwidth to synchronize the data between the sites:
  - Bandwidth requirements increase for lower RPO values.
- Could the data at the second site be used as the basis for dev or test systems?

These options are compared and discussed in detail in the following sections.

## 2.1 Backup and Recovery

SAP HANA supports these methods for database backups:

- File-based backup to a file system, typically an NFS share
- Backups using the SAP HANA BACKINT API and certified third-party backup tools
- Storage-based NetApp® Snapshot™ backups

To choose the method that's best for them, customers must understand the infrastructure and performance impact as well as the additional required features of the selected HANA backup method. The following subsections provide a few examples.

### File-Based Backups

With file-based backups or backups made using the BACKINT API, the SAP HANA database server reads the data from the primary storage. The database server then either writes the data to an NFS share

or streams the data to a backup server using the third-party backup tool. Both approaches have a significant impact on the performance of the SAP HANA database in the following ways:

- Additional CPU load at the SAP HANA database server
- Additional I/O load at the primary storage
- The load on the backup network

In addition, the backup run time, specifically for larger databases, can be significant, resulting in lower operation speed during backup. The restore and recovery process can also be a challenge because of the long run time.

## Storage-Based Snapshot Backups

NetApp storage-based Snapshot backups address the challenges discussed earlier. Independently of the size of the SAP HANA database, a Snapshot backup is executed within a few seconds instead of hours. The complete run time of the backup operation depends mainly on the required HANA savepoint operation, which needs to be done before the storage Snapshot operation. Customer data shows that the average backup run time is around 1 minute. The Snapshot backup is executed at the storage layer, and there is no impact on the performance of the SAP HANA database. Also, the restore process occurs in a matter of seconds, which has a significant impact on the RTO if a restore operation is required.

NetApp SnapCenter® with the SAP HANA plug-in can facilitate an automated and fully integrated HANA backup based on Snapshot technology, including the automation of SAP HANA block integrity checks.

SnapCenter also handles the scheduling and housekeeping of backups on the storage and within the SAP HANA backup catalog based on flexible, configurable retention policies. In addition, nondatabase files can be secured with SnapCenter.

For more information, see the [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#) best practices guide.

## Addressing Logical Corruption

As discussed in section 1, Data Protection Overview, a logical corruption in a production SAP system can typically not be addressed by a point-in-time recovery of the SAP HANA database. A point-in-time recovery would result in data loss and in an inconsistent SAP landscape if multiple SAP systems are exchanging data with each other. Rather, NetApp recommends fixing the logical corruption by setting up a repair system, exporting the required data, and importing that data back to the production system.

When setting up the repair system, flexibility and speed are crucial. With NetApp storage-based Snapshot backups, multiple consistent database images are available to create a clone of the production system by using NetApp FlexClone® technology. FlexClone volumes can be created in a matter of seconds rather than multiple hours if a redirected restore from a file-based backup is used to set up the repair system. Section 4, Overview of Disaster Recovery Testing, describes the process for setting up a production clone for disaster recovery testing purposes. The same workflow can be used to set up a repair system.

Figure 3) NetApp storage-based backups.

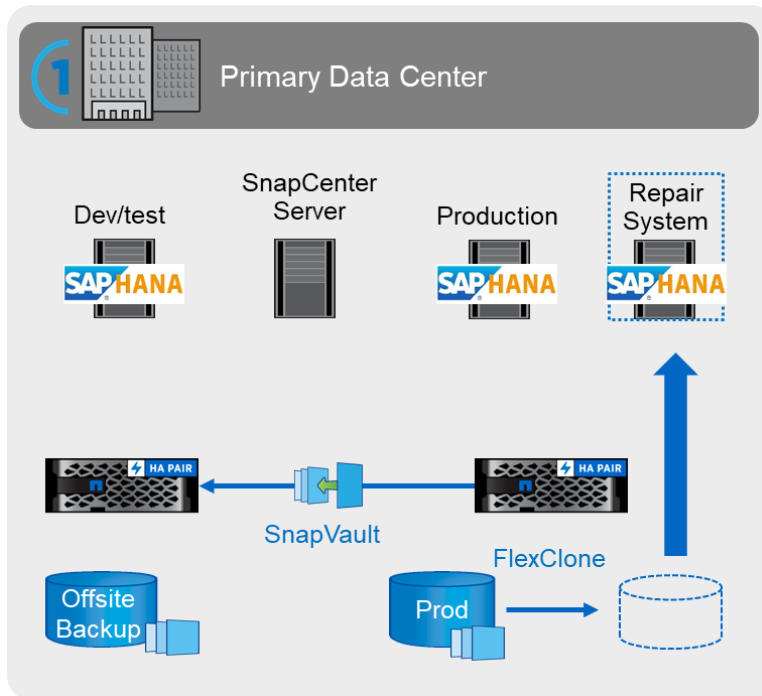


Figure 3 is an example of a backup solution using NetApp storage-based Snapshot backups to secure the system on the primary site. Backups can be automatically transferred by using NetApp SnapVault® backup software to a dedicated off-site backup storage system. This process is controlled by SnapCenter. All the backups available at the primary or off-site backup storage site can be used to create FlexClone volumes and to set up a repair system to address logical corruption.

The following sections combine this backup approach with different options for disaster recovery replication.

## 2.2 SAP HANA Disaster Recovery Using Storage Replication

NetApp supports three different disaster recovery data replication methods for SAP HANA.

- **Asynchronous replication with SnapMirror.** NetApp SnapMirror® replication technology is built into ONTAP® for disaster recovery solutions. SnapMirror is configured through a data protection relationship between two storage volumes on a primary and a secondary storage system. SnapMirror updates the secondary volume by using efficient block delta replications. Update schedules can be either defined on the storage system itself or triggered externally. For full information, see [TR-4015: SnapMirror Configuration and Best Practices Guide for ONTAP 9](#).
- **Synchronous replication with SnapMirror.** Synchronous SnapMirror (SM-S) was introduced with ONTAP 9.5. SM-S is configured through a data protection relationship between two storage volumes on a primary and a secondary storage system. SM-S is targeted at relatively short distances to provide exact replicas with an RPO of zero. For full information, see [TR-4733: SnapMirror Synchronous for ONTAP 9.7](#).
- **Synchronous replication with NetApp MetroCluster.** NetApp MetroCluster™ configurations are used to provide high availability, zero data loss, and nondisruptive operations both within and beyond the data center. MetroCluster is a free feature of ONTAP software that synchronously mirrors data and configuration between two ONTAP clusters in separate locations or failure domains. MetroCluster provides continuously available storage for applications by automatically handling two objectives:
  - Zero RPO by synchronously mirroring data written to the cluster

- Near-zero RTO by mirroring configuration and automating access to data at the second site MetroCluster provides simplicity with automatic mirroring of data and configuration between the two independent clusters located in the two sites. As storage is provisioned in one cluster, it is automatically mirrored to the second cluster at the second site. For full information, see [TR-4689: NetApp MetroCluster IP](#) and [TR-4375: NetApp MetroCluster FC](#)

## Synchronous SnapMirror Combined with Asynchronous SnapMirror

Figure 4 shows a three-site disaster recovery solution, using synchronous SnapMirror replication to the local disaster recovery data center and asynchronous SnapMirror to replicate the data to the remote disaster recovery data center.

Data replication using synchronous SnapMirror provides an RPO of zero. The distance between the primary and the local disaster recovery data center is limited to around 100km.

Protection against failures of both the primary and local disaster recover sites is done by replicating the data to a third remote disaster recovery data center using asynchronous SnapMirror. The RPO depends on the frequency of replication updates and how fast they can be transferred. In theory, the distance is unlimited, but the limit depends on the amount of data that must be transferred and the connection that is available between the data centers. Typical RPO values are in the range of 30 minutes to multiple hours.

The RTO for both replication methods depends mainly on the time needed to start the HANA database at the disaster recovery site and to load the data into memory. With the assumption that the data is read with a throughput of 1000MB/s, loading 1TB of data would take approximately 18 minutes.

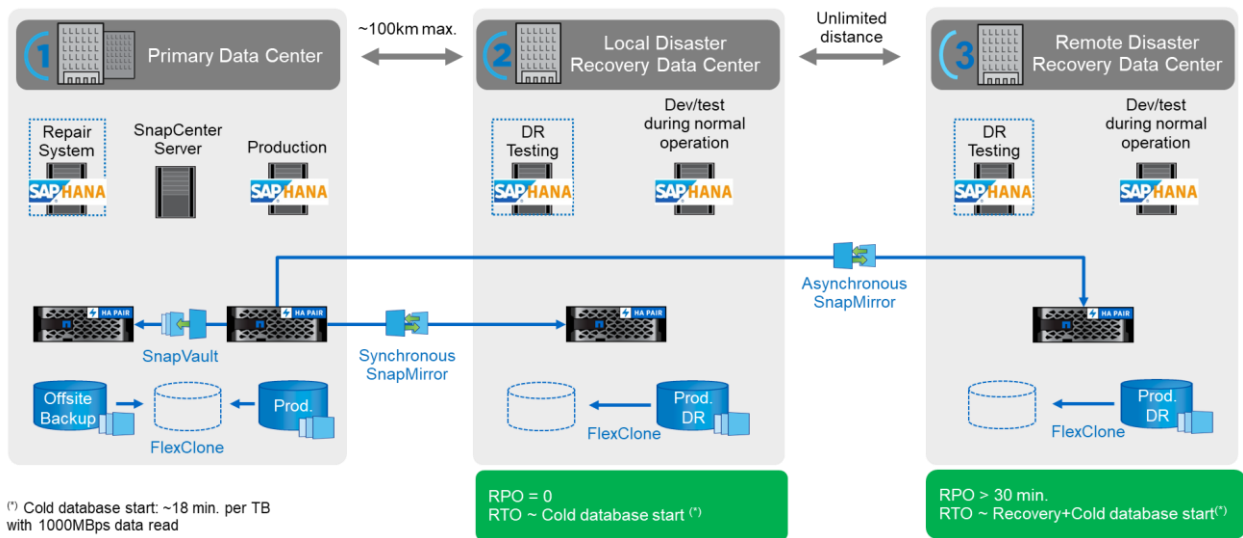
The servers at the disaster recovery sites can be used as dev/test systems during normal operation. In case of a disaster, the dev/test systems need to be shut down and started as DR production servers.

Both replication methods allow you to test a DR workflow without influencing the RPO and RTO. FlexClone volumes are created on the storage and are attached to the DR testing servers.

Leveraging cloud provider IaaS offerings from AWS, Azure, or Google is another option for data replication to a third location. Instead of running your own remote data center resources, the data could be replicated to Cloud Volumes ONTAP by using asynchronous SnapMirror. Cloud Volumes ONTAP is available at all three main cloud providers, AWS, Azure, and Google.

The following sections cover the DR configuration, the DR workflow testing, and the DR failover workflow steps in detail.

Figure 4) Disaster recovery with synchronous and asynchronous SnapMirror.



## NetApp MetroCluster Combined with Asynchronous SnapMirror

Figure 5 shows a three-site disaster recovery solution, using NetApp MetroCluster for synchronous replication to the local disaster recovery data center and asynchronous SnapMirror to replicate the data to the remote disaster recovery data center.

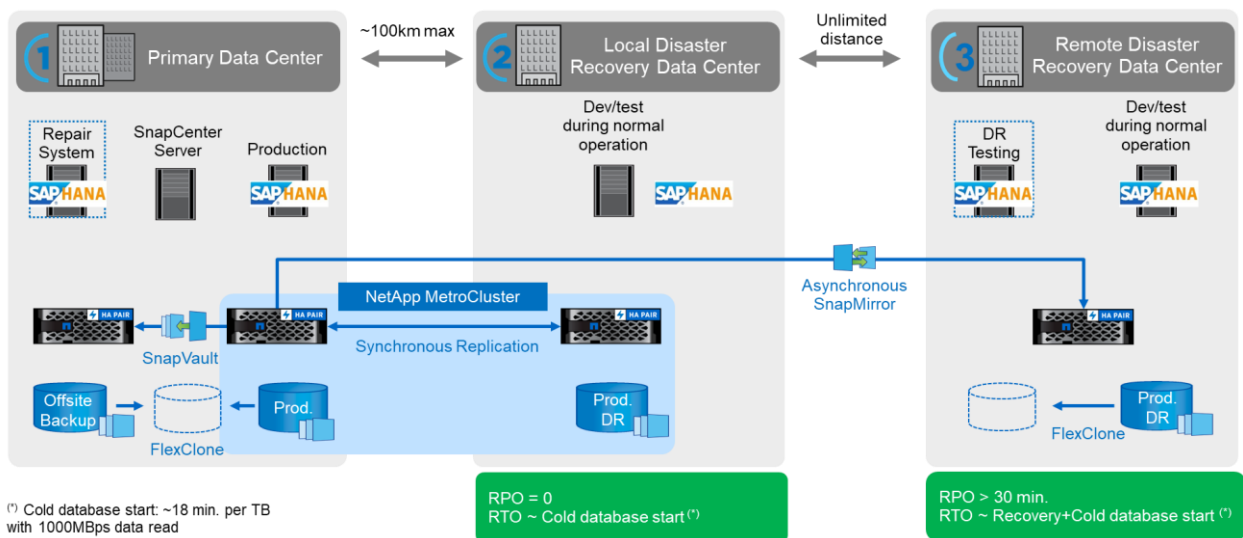
Synchronous storage replication based on NetApp MetroCluster provides RPO=0. This disaster recovery solution does not require any additional configuration at the SAP HANA level.

NetApp MetroCluster is supported up to 700km for MetroCluster IP. However, the maximum distance is determined by the maximum acceptable latency for the application; in most cases it is in the range of 100km.

The SAP HANA data and log volumes and the nondatabase data are synchronously replicated to the disaster recovery site, as shown in Figure 5. During normal operation, the disaster recovery servers can run development or test systems. In the event of a disaster, the dev/test systems must be shut down, and MetroCluster failover must be initiated at the storage layer to make the mirrored plexes available to the disaster recovery server. A third-party software solution, such as [ProLion ClusterLion automatic switchover for NetApp MetroCluster](#), can automate the MetroCluster failover process.

After mounting the data at the disaster recovery server, you must run a normal SAP HANA database start, including crash recovery. The RTO for this cold standby approach depends on the size of the database and the read throughput during the load of the row and column store. With the assumption that the data is read with a throughput of 1000MB/s, loading 1TB of data takes approximately 18 minutes.

Figure 5) NetApp MetroCluster combined with asynchronous SnapMirror.



All storage Snapshot copies stored at the primary site are also available at the secondary site. So even after a disaster failover, multiple replication images are available to address logical corruption.

### 2.3 SAP HANA System Replication

SAP HANA System Replication works at the application layer. The solution is based on an additional SAP HANA system at the disaster recovery site that receives the changes from the primary system. This secondary system must be identical to the primary system.

SAP HANA System Replication can be operated in one of two modes:

- With data preload into memory, with a dedicated server at the disaster recovery site:
  - The server is used exclusively as an SAP HANA System Replication secondary host.

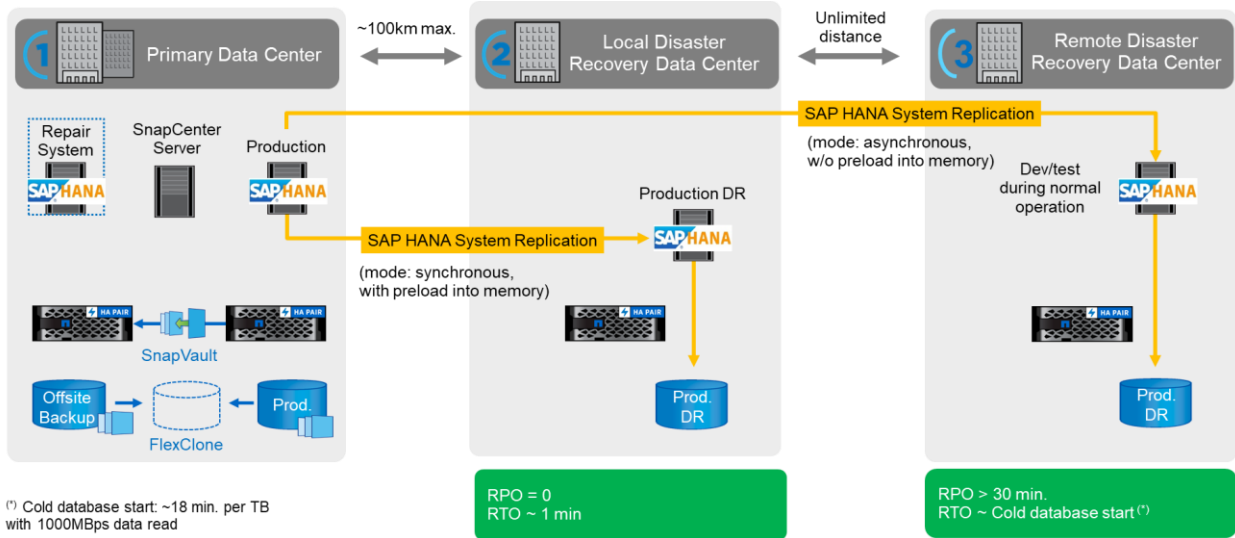
- Very low RTO values can be achieved, because the data is already loaded into memory and no database start is required in case of a failover.
- Without data preload into memory, with a shared server at the disaster recovery site:
  - The server is shared as an SAP HANA System Replication secondary and as a dev/test system.
  - RTO depends mainly on the time required to start the database and load the data into memory.

For a full description of all configuration options and replication scenarios, see the [SAP HANA Administration Guide](#).

Figure 6 shows the setup of a three-site disaster recovery solution with SAP HANA System Replication. Synchronous replication with data preload into memory is used for the local disaster recovery data center. Asynchronous replication without data preload is configured for the remote disaster recovery data center.

**Note:** SAP HANA System Replication can also be combined with storage replication. In this case, HANA System Replication would be used for data replication to the local DR data center and asynchronous SnapMirror would be used for replication to the remote DR data center.

Figure 6) SAP HANA System Replication.



## SAP HANA System Replication with Data Preload into Memory

Very low RTO values with SAP HANA can be achieved only with SAP HANA System Replication with data preload into memory. Operating SAP HANA System Replication with a dedicated secondary server at the disaster recovery site allows an RTO value of approximately 1 minute or less. The replicated data is received and preloaded into memory at the secondary system. Because of this low failover time, SAP HANA System Replication is also often used for near-zero-downtime maintenance operations, like HANA software upgrades.

Typically, SAP HANA System Replication is configured to replicate synchronously when data preload is chosen. The maximum supported distance for synchronous replication is in the range of 100km.

SAP HANA System Replication does not include replication of nondatabase files, so any system changes outside of the database, like SAP application server, require an additional replication method. Therefore, SAP HANA System Replication is often combined with storage-based replication for nondatabase data.

## SAP System Replication Without Data Preload into Memory

For less stringent RTO requirements, you can use SAP HANA System Replication without data preload. In this operational mode, the data at the disaster recovery site is not loaded into memory. The server at



the DR site is still used to process SAP HANA System Replication running all the required SAP HANA processes. However, most of the server's memory is available to run other workloads, such as SAP HANA dev/test systems.

In the event of a disaster, the dev/test system must be shut down, failover must be initiated, and the data must be loaded into memory. The RTO of this cold standby approach depends on the size of the database and the read throughput during the load of the row and column store. With the assumption that the data is read with a throughput of 1000MB/s, loading 1TB of data should take approximately 18 minutes.

## 2.4 Summary of Disaster Recovery Solutions

Table 1 compares the disaster recovery solutions discussed in this section and highlights the most important indicators.

The key findings are:

- If a very low RTO is required, SAP HANA System Replication with preload into memory is the only option.
  - Storage replication is also needed to replicate nondatabase data.
- For medium RTO requirements, storage replication can also be used to:
  - Combine database and nondatabase data replication.
  - Cover additional use cases such as disaster recovery testing and dev/test refresh.
- All disaster recovery solutions must be combined with a backup solution that addresses logical corruption.

Table 1) Disaster recovery solution comparison.

|   | Storage Replication                                  |   | SAP HANA System Replication                          |  |
|---|--|---|--|--|
|   | NetApp SnapMirror                                    | NetApp MetroCluster                               | With Data Preload                                    | Without Data Preload                                 |
| RTO   | Low to medium, depending on database startup time    | Low to medium, depending on database startup time | Very low   | Low to medium, depending on database startup time    |
| RPO   | RPO=0 for synchronous<br>RPO> 30min for asynchronous | RPO=0   | RPO=0 for synchronous<br>RPO> 30min for asynchronous | RPO=0 for synchronous<br>RPO> 30min for asynchronous |
| Servers at DR site can be used for dev/test         | Yes  | Yes   | No   | Yes  |
| Replication of nondatabase data                     | Yes  | Yes   | No   | No   |
| DR data can be used for refresh of dev/test systems | Yes, for asynchronous<br>No, for synchronous         | No  | No   | No   |
| DR testing without affecting RTO and RPO            | Yes  | Yes   | No   | No   |
| DR configuration effort                             | For each storage volume used by the databases        | All storage volumes are automatically replicated  | For each database                                    | For each database                                    |

### 3 SAP HANA Disaster Recovery with SnapMirror Replication

Figure 7 shows the volume replication relationships for a three-site disaster recovery configuration with synchronous and asynchronous SnapMirror replication. In addition to the SnapMirror replication, the HANA data and shared volume are also replicated with SnapVault to offsite backup storage.

**Note:** The following description and the lab setup focus on the HANA database. SAP application servers would have additional storage volumes, which would be protected and replicated in the same way as the HANA shared volume.

With synchronous SnapMirror replication, all three HANA volumes, the data, the log, and the shared volume must be replicated to the local disaster recovery data center. Replication of application-consistent Snapshot copies is not required. In case of a disaster failover, the replication relationship must be broken, the volumes must be mounted to the DR production server, and the HANA database must be started and will execute a crash recovery. Section 8, Synchronous SnapMirror Disaster Recovery Failover, describes the required steps.

With asynchronous SnapMirror replication, the HANA data and the HANA shared volume need to be replicated. With HANA data volume replication, typical RPO values are in the range of multiple hours. If lower RPO values are required, the HANA log backups must be replicated in addition for forward recovery. To enable HANA savepoint or forward recovery, application-consistent Snapshot data backups must be included in the HANA data volume at the disaster recovery site. This is accomplished with SnapCenter backups created at the primary site, which are then replicated to the DR site. In case of a disaster failover, the replication relationship must be broken, the volumes must be mounted to the DR production server, and the HANA database must be recovered, either to the last HANA savepoint or with forward recovery using the replicated log backups. Section 9, Asynchronous SnapMirror Disaster Recovery Failover, describes the required steps.

**Note:** The HANA shared volume includes the /hana/shared as well as the /usr/sap file system of the HANA system. The installation was done as described in [TR-4435: SAP HANA on NetApp AFF Systems with NFS](#).

Figure 7) SAP HANA disaster recovery with SnapMirror replication.

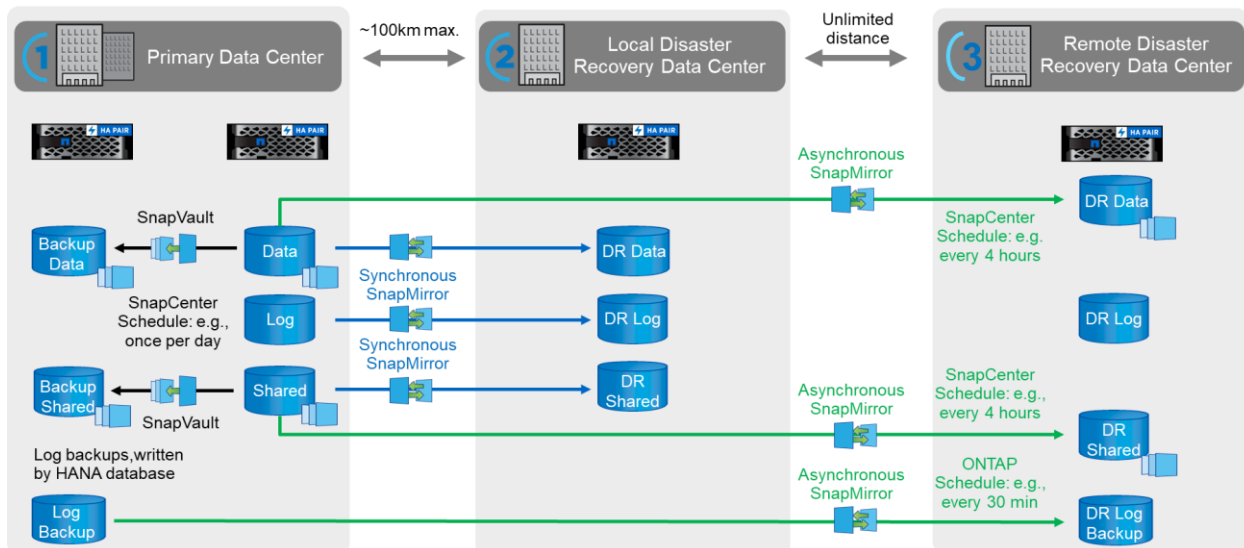


Table 2 compares the different storage replication approaches.



Table 2) Comparison of asynchronous storage replication approaches.

| Replication   | RPO   | RTO   |
|---|---|---|
| Synchronous SnapMirror  | <ul style="list-style-type: none"> <li>RPO = 0</li> </ul>   | <ul style="list-style-type: none"> <li>Break mirror and mount volumes</li> <li>HANA crash recovery</li> <li>Database start (~ 18 min per TB)</li> </ul>                 |
| Asynchronous SnapMirror with data and log backup volume replication | <ul style="list-style-type: none"> <li>Depends on the log backup replication frequency and the log backup interval.</li> <li>Typical RPO &gt; 30 min</li> </ul> | <ul style="list-style-type: none"> <li>Break mirror and mount volumes</li> <li>HANA recovery with forward recovery</li> <li>Database start (~ 18 min per TB)</li> </ul> |
| Asynchronous SnapMirror with data volume replication only           | <ul style="list-style-type: none"> <li>Depends on the data volume replication frequency</li> <li>Typical RPO=multiple hours</li> </ul>                          | <ul style="list-style-type: none"> <li>Break mirror and mount volumes</li> <li>HANA recovery to last savepoint</li> <li>Database start (~ 18 min per TB)</li> </ul>     |

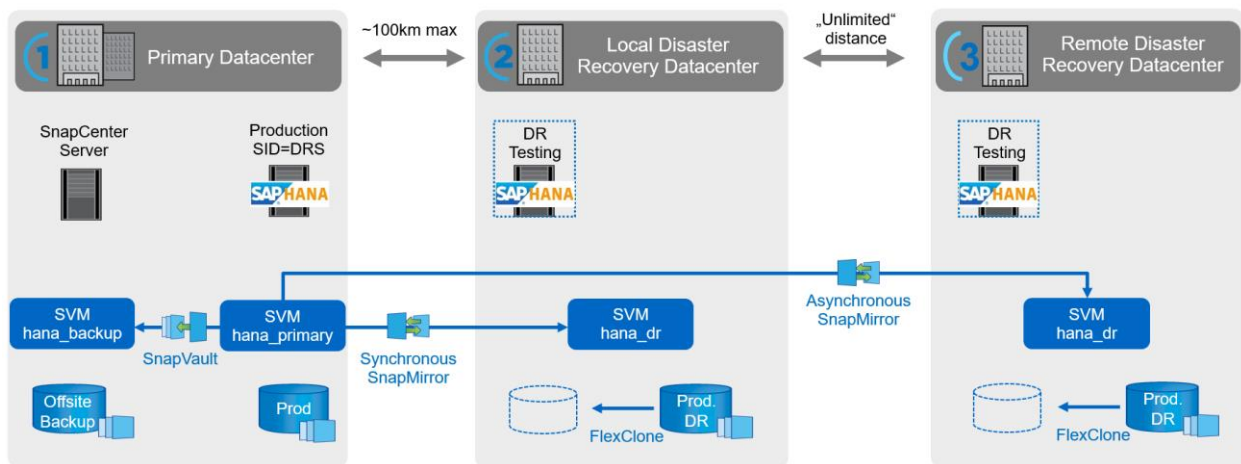
**Note:** Database start and data load into memory takes around 18 minutes per TB, with a data read throughput of 1000MB/s.

### 3.1 Lab Setup

Figure 8 shows the schematic lab setup with storage virtual machines (SVM) used for primary, backup, and disaster recovery data. SnapCenter is used to create Snapshot data backups of the HANA system with the System Identifier (SID) DRS. These backups are replicated to the offsite backup as well as to the remote disaster recovery site. For full information about configuring SnapCenter and the HANA plug-in for data protection, see [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

**Note:** For the lab setup, NFS is used to connect the storage volumes to the HANA hosts. The differences for a HANA host using Fibre Channel (FC) are highlighted in the subsections of section 5, 6, 8, and 9.

Figure 8) Lab setup.



The following software versions were used in the lab setup:

- SAP HANA 2.0 SPS4 (MDC single tenant)
- SUSE Linux SLES for SAP 15 SP1
- NetApp ONTAP 9.7

- SnapCenter 4.3.1P1

Table 3 shows the replication relationships, which have been configured in the lab setup. The table lists the same replication relationships that are shown in Figure 7.

**Table 3) Replication relationships.**

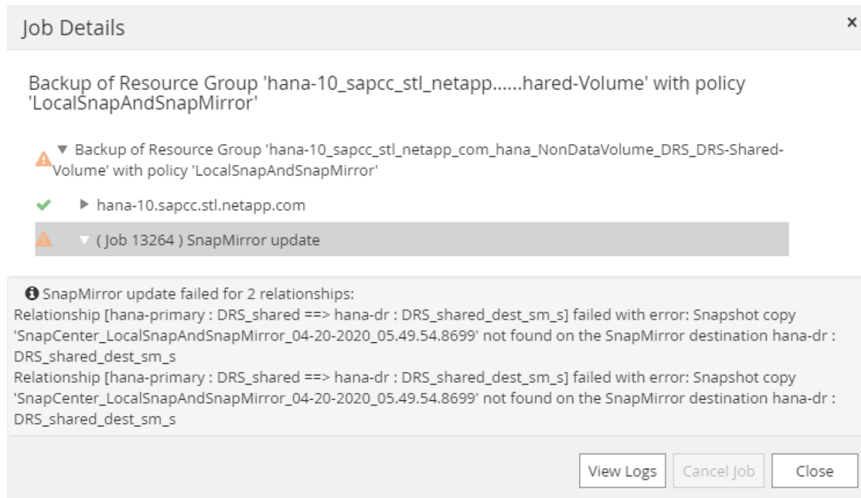
| Source Volume     | Source SVM   | Target Volume               | Target SVM  | Replication             | Schedule         | Schedule Managed By |
|-------------------|--------------|-----------------------------|-------------|-------------------------|------------------|---------------------|
| DRS_data_mnt00001 | hana_primary | DRS_data_mnt00001_dest      | hana_backup | SnapVault               | Once per day     | SnapCenter          |
| DRS_shared        | hana_primary | DRS_shared_dest             | hana_backup | SnapVault               | Once per day     | SnapCenter          |
| DRS_data_mnt00001 | hana_primary | DRS_data_mnt00001_dest_sm_s | hana_dr     | Synchronous SnapMirror  | NA               | NA                  |
| DRS_log_mnt00001  | hana_primary | DRS_log_mnt00001_dest_sm_s  | hana_dr     | Synchronous SnapMirror  | NA               | NA                  |
| DRS_shared        | hana_primary | DRS_shared_dest_sm_s        | hana_dr     | Synchronous SnapMirror  | NA               | NA                  |
| DRS_data_mnt00001 | hana_primary | DRS_data_mnt00001_dest      | hana_dr     | Asynchronous SnapMirror | Every 4 hours    | SnapCenter          |
| DRS_shared        | hana_primary | DRS_shared_dest             | hana_dr     | Asynchronous SnapMirror | Every 4 hours    | SnapCenter          |
| DRS_log_backup    | hana_backup  | DRS_log_backup_dest         | hana_dr     | Asynchronous SnapMirror | Every 30 minutes | ONTAP               |

### 3.2 SnapCenter Support for Synchronous SnapMirror

The current SnapCenter release does not have an integration of synchronous SnapMirror. Synchronous SnapMirror can be used together with SnapCenter, but you need to be aware of a few shortcomings. Synchronous SnapMirror will be supported in future SnapCenter releases.

- In the SnapCenter topology view of a HANA resource, a synchronous SnapMirror relationship is not visible.
- Snapshot backups that are created with SnapCenter are not replicated to the synchronous SnapMirror target.
  - A new protection policy (SnapCenterSync) for synchronous SnapMirror that enables Snapshot replication was introduced in ONTAP 9.7. SnapCenter does not yet support the required APIs, so SnapCenter Snapshot backups are not replicated.
- When synchronous SnapMirror is combined with asynchronous SnapMirror, SnapCenter checks the availability of the Snapshot backups at the synchronous SnapMirror target with each SnapMirror update job.
  - Because the backups are not replicated to the synchronous SnapMirror target, SnapCenter waits and tries the operation again until the timeout of 120 minutes. Therefore, the SnapMirror update job runs for about 2 hours and delays the visibility of the backup at the asynchronous SnapMirror target.
  - After the timeout, SnapCenter displays a warning message that the backup at the synchronous SnapMirror target was not found.

Figure 9) SnapMirror update warning message.



### 3.3 Configuration Steps for Synchronous SnapMirror Replication

Synchronous SnapMirror replication is configured on the storage system. Figure 10 shows the protection relationship dialog box from ONTAP System Manager, which is used to configure the replication. The dialog box includes four main configuration options.

- Replication type:
  - Synchronous
- Synchronization mode:
  - Sync or StrictSync
  - With StrictSync, the application using the source volume gets an I/O error if the replication to the target is not possible. With StrictSync, an RPO of zero can be guaranteed. However, if the secondary site is not available, the result will be application downtime.
- Source volume and target volume suffix
- Protection policy:
  - Can be Sync or SnapCenterSync for synchronous replication.
  - To enable Snapshot replication to the synchronous SnapMirror target, SnapCenterSync must be selected. Using this policy, Snapshot copies that are created with the ONTAP CLI with the label `app_consistent` are replicated to the target. As discussed in section 3.2, SnapCenter Support for Synchronous SnapMirror, SnapCenter backups cannot currently be replicated. Snapshot replication based on ONTAP CLI Snapshot copies, is used for disaster recovery testing, as described in section 5, Synchronous SnapMirror Disaster Recovery Testing.

Figure 10) Synchronous SnapMirror - Create Protection Relationship.

**Create Protection Relationship**

Data protection refers to backing up data and being able to recover it. Depending on your data protection and backup needs, you can select an appropriate method to protect your data against accidental, malicious, or disaster-induced data loss.  
[Tell me more about different types of data protection relationships.](#)

**Relationship Type**

Replication: Synchronous [Help me Choose](#)

Synchronization Mode: Sync

**Source Volume**

Cluster: a700-marco

SVM: hana-primary

Volume: DRS\_data\_mnt00001 [Browse...](#)

**Destination Volume**

SVM: hana-dr

Volume Name Suffix: \_dest\_sm\_s

**Configuration Details**

Sync Policy: SnapCenterSync [Browse...](#)

Initialize Relationship

SnapLock for SnapVault SnapLock for SnapVault is applicable only for Vault protection relationships.

[Create](#) [Cancel](#)

Figure 11 shows the required volume relationships. The SAP HANA data, log, and shared volume must be replicated with synchronous SnapMirror.

Figure 11) Synchronous SnapMirror volume relationships.

**ONTAP System Manager**

Switch to the new experience

Type: All Search all Objects

**Volume Relationships**

| Source Storage | Source Volume     | Destination Volume          | Destination SVM | Is Healthy | Object Type | Relationship | Transfer Status | Relationship Type | Lag Time | Policy Name    | Policy Type |
|----------------|-------------------|-----------------------------|-----------------|------------|-------------|--------------|-----------------|-------------------|----------|----------------|-------------|
| hana-primary   | DRS_data_mnt00001 | DRS_data_mnt00001_dest_sm_s | hana-dr         | Yes        | Volume      | Snapmirrored | InSync          | Sync              | None     | SnapCenterSync | Sync        |
| hana-primary   | DRS_log_mnt00001  | DRS_log_mnt00001_dest_sm_s  | hana-dr         | Yes        | Volume      | Snapmirrored | InSync          | Sync              | None     | SnapCenterSync | Sync        |
| hana-primary   | DRS_shared        | DRS_shared_dest_sm_s        | hana-dr         | Yes        | Volume      | Snapmirrored | InSync          | Sync              | None     | SnapCenterSync | Sync        |

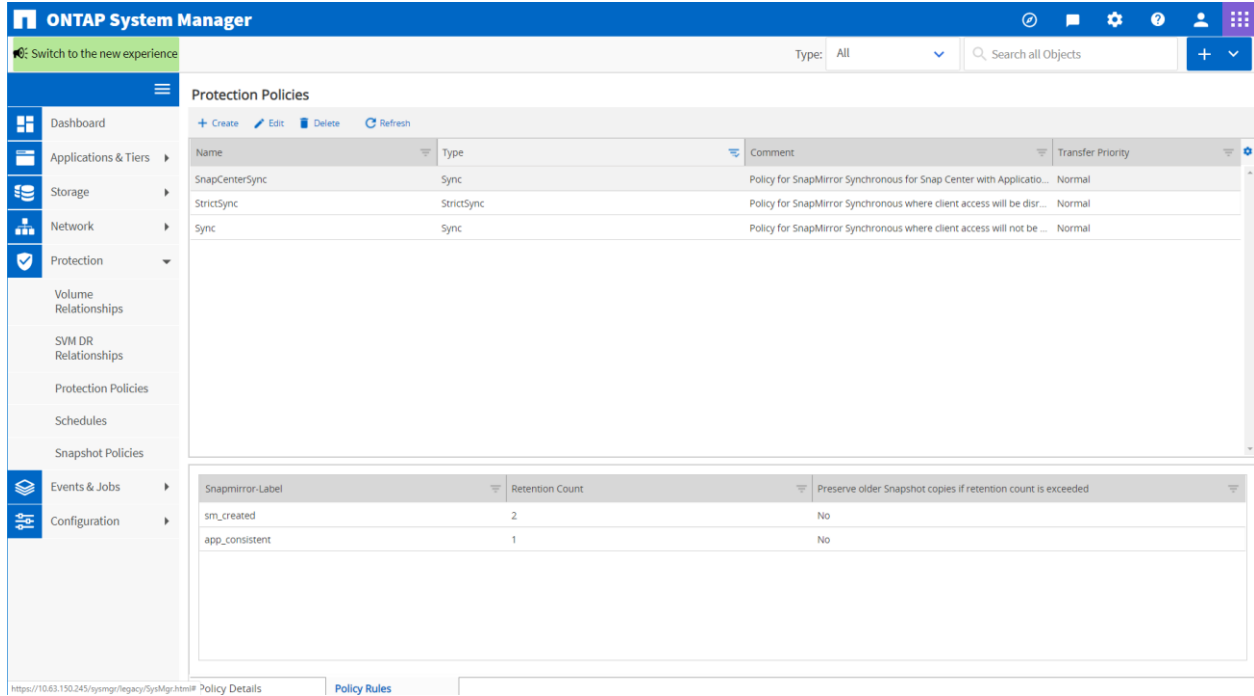
**Details**

|   |  |  |
|---|--|--|
| Source Location: hana-primary:DRS_data...       | Is Healthy: <span style="color: green;">●</span> Yes | Transfer Status: InSync  |
| Destination Location: hana-dr:DRS_data_mnt00... | Relationship State: Snapmirrored                     | Current Transfer Type: None  |
| Source Cluster: a700-marco                      | Network Compression Ratio: Not Applicable            | Current Transfer Error: None   |
| Destination Cluster: a700-marco                 |  | Current Transfer Progress: None  |
| Transfer Schedule: None                         |  | Last Transfer Error: None  |
| Data Transfer Rate: Unlimited                   |  | Last Transfer Type: Resync   |
| Lag Time: None                                  |  | Latest Snapshot Timestamp: 04/20/2020 04:05:00   |
|   |  | Latest Snapshot Copy: snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901726.2020-04-20_040500 |

[Details](#) [Policy Details](#) [Snapshot Copies](#)

Figure 12 shows the policy rules of the SnapCenterSync protection policy. The SnapMirror label `app_consistent` enables the replication of Snapshot copies to the target volume. Snapshot replication is required for disaster recovery testing, as described in section 5, Synchronous SnapMirror Disaster Recovery Testing.

Figure 12) SnapCenterSync Protection Policies.



### 3.4 Configuration Steps for Asynchronous SnapMirror Replication

SAP HANA disaster recovery with asynchronous SnapMirror replication requires database-consistent Snapshot backups at the production site, which are then replicated to the disaster recover site. The database-consistent Snapshot backups are created with SnapCenter and the SAP HANA plug-in.

First, the asynchronous SnapMirror protection relationship must be configured on the storage layer. The initial complete data transfer is also part of this configuration.

SnapCenter is then used to initiate a SnapMirror update replication during the normal Snapshot backup workflow. The following tasks are performed during the backup workflow with SnapCenter:

1. Trigger an SAP HANA database snapshot, which includes a backup savepoint to get a consistent image on the persistence layer.
2. Create a storage Snapshot copy at the production site.
3. Register a backup in the SAP HANA backup catalog.
4. Initiate the replication update to the disaster recovery site by using SnapMirror.
5. Perform retention management and housekeeping for data and log backups.

Because the replication to the disaster recovery site is part of the overall backup process, the replication frequency of the SAP HANA data volume depends on the backup frequency.

### Data Volume Replication Combined with Log Backup Volume Replication

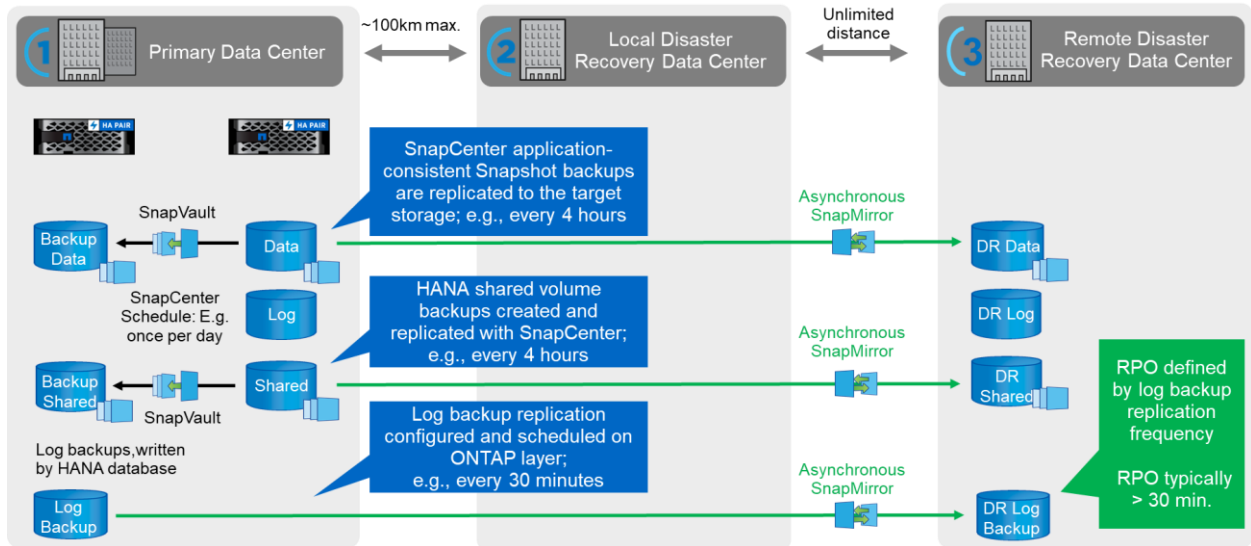
Figure 13 shows a disaster recovery solution based on the replication of the SAP HANA database data volume plus the replication of the log backup volume.

As discussed, data volume replication is tied to the backup workflow in SnapCenter. A typical backup frequency is “every 4 hours.” In our example the data volume and the shared volume are backed up and replicated by SnapCenter every 4 hours.

The RPO is defined by the replication frequency of the log backup volume. With the standard SAP HANA log backup interval of 15 minutes and a log backup replication interval of, for example, 30 minutes, the minimal possible RPO is >30 minutes.

In the event of a disaster failover, the RTO is defined by the time required to recover the database without applying logs, plus the time needed for infrastructure preparations. Section 9, Asynchronous SnapMirror Disaster Recovery Failover, describes the required steps.

Figure 13) Combined backup and disaster recovery replication.



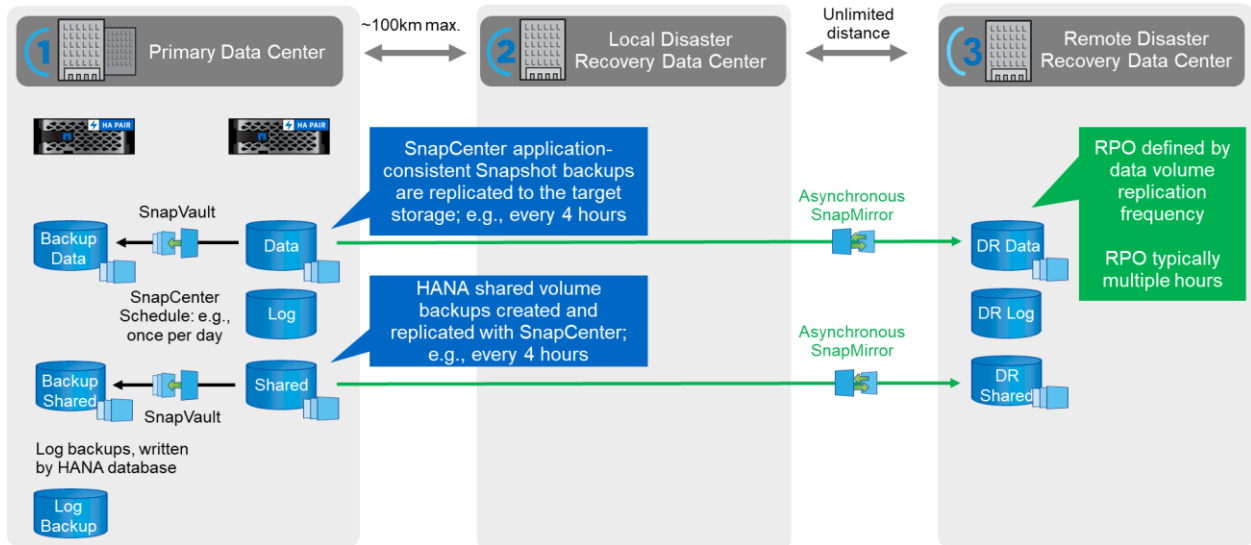
## Replication of Data Volume Only

Figure 14 shows a disaster recovery solution that is based on the replication of the SAP HANA database data volume. The log backup volume is not replicated.

The RPO is defined by the replication frequency of the data volume. Replication to the DR site is part of the normal backup workflow. Therefore, a higher replication frequency can be achieved only by adopting a higher backup frequency. In general, NetApp does not recommend having a backup interval of less than 1 hour. Based on this recommendation, the lowest achievable RPO is 1 to 2 hours.

In the event of a disaster failover, the RTO is defined by the time required to recover the database without applying logs plus the time needed for infrastructure preparations. Section 9, Asynchronous SnapMirror Disaster Recovery Failover, describes the required steps.

Figure 14) Replication of data volume only.

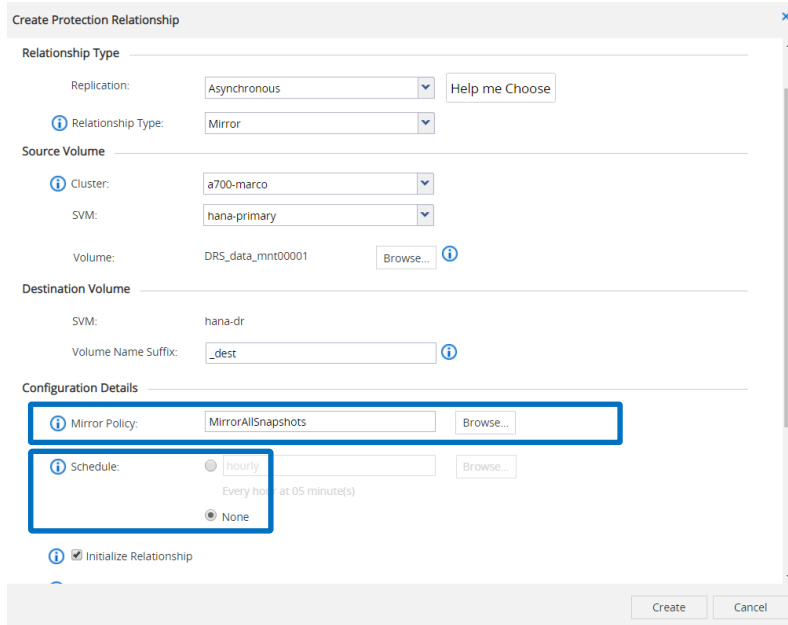


## Configuration of ONTAP Replication Relationships

Figure 15 shows the ONTAP System Manager dialog box for the configuration of the HANA data and log volumes. For both volumes, the selected Protection Policy is MirrorAllSnapshots, so that every backup that is created by SnapCenter is replicated to the DR site.

Because SnapCenter triggers the SnapMirror update operation as part of each backup workflow, schedule None needs to be configured in ONTAP.

Figure 15) Asynchronous SnapMirror - Create Protection Relationship – Configuration Details.



For the log backup volume, a schedule is configured on the storage layer. because SnapCenter is not executing any backup operation for this volume. In our example, a new schedule, Every-30-Minutes, has been created.

Figure 16) Asynchronous SnapMirror - Create Protection Relationship – Schedule.

Figure 17 shows all configured asynchronous SnapMirror relationships. As discussed, log backup volume replication is not required if a data-volume-only replication concept has been chosen.

Figure 17) Asynchronous SnapMirror Volume Relationships.

| Source Storage | Source Volume     | Destination Volume     | Destination... | Is ... | Object ... | Relationship ... | Relationship Type | Policy Name                  | Policy Type        |                     |
|----------------|-------------------|------------------------|----------------|--------|------------|------------------|-------------------|------------------------------|--------------------|---------------------|
| hana-backup    | DRS_log_backup    | DRS_log_backup_dest    | hana-dr        | Yes    | Volume     | Snapmirrored     | Idle              | Asynchronous Version-Flex... | MirrorAllSnapshots | Asynchronous Mirror |
| hana-primary   | DRS_data_mnt00001 | DRS_data_mnt00001_dest | hana-dr        | Yes    | Volume     | Snapmirrored     | Idle              | Asynchronous Version-Flex... | MirrorAllSnapshots | Asynchronous Mirror |
| hana-primary   | DRS_shared        | DRS_shared_dest        | hana-dr        | Yes    | Volume     | Snapmirrored     | Idle              | Asynchronous Version-Flex... | MirrorAllSnapshots | Asynchronous Mirror |

|                       |                           |                            |                |                            |  |
|-----------------------|---------------------------|----------------------------|----------------|----------------------------|--|
| Source Location:      | hana-backup:DRS_log_ba... | Is Healthy:                | Yes            | Transfer Status:           | Idle   |
| Destination Location: | hana-dr:DRS_log_backup... | Relationship State:        | Snapmirrored   | Current Transfer Type:     | None   |
| Source Cluster:       | a700-marco                | Network Compression Ratio: | Not Applicable | Current Transfer Error:    | None   |
| Destination Cluster:  | a700-marco                |                            |                | Current Transfer Progress: | None   |
| Transfer Schedule:    | Every-30-Minutes          |                            |                | Last Transfer Error:       | None   |
| Data Transfer Rate:   | Unlimited                 |                            |                | Last Transfer Type:        | Update   |
| Lag Time:             | None                      |                            |                | Last Snapshot Timestamp:   | 04/23/2020 06:30:00  |
|                       |                           |                            |                | Latest Snapshot Copy:      | snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901723.2020-04-23_063000 |



## SnapCenter Configuration

This section describes the SnapCenter configuration for disaster recovery. Basic operations like adding storage systems and configuring policies are described in [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

Two resources need to be configured in SnapCenter:

- HANA database resource for the backup operations of the HANA database
- Non-data-volume resource for the backup operations of the /hana/shared volume

To configure the SnapCenter resources, follow these steps:

1. Add storage systems (if they are not already configured).  
For our lab setup, the storage SVMs hana-primary, hana-backup, and hana-dr need to be added.
2. Add policies (if they are not already configured).  
LocalSnapAndSnapMirror, LocalSnapAndSnapVault, BlockIntegrityCheck.
3. Add SAP HANA system DRS.  
Execute auto discovery by deploying the SnapCenter HANA plug-in on the database host.
4. Configure resource protection for HANA system DRS.
  - a. Select policies.
  - b. Define schedules.
5. Add non-data-volume resource for the DRS-shared volume.
6. Configure resource protection for the DRS-shared volume.
  - a. Select policies.
  - b. Define schedules.

Figure 18 shows the SnapCenter topology view for the HANA database resource and Figure 19 shows the SnapCenter topology view for the HANA shared volume.

**Note:** As discussed in section 3.2, in the current SnapCenter release, synchronous SnapMirror relationships are not displayed in the topology view.

Figure 18) SnapCenter topology view for HANA database.

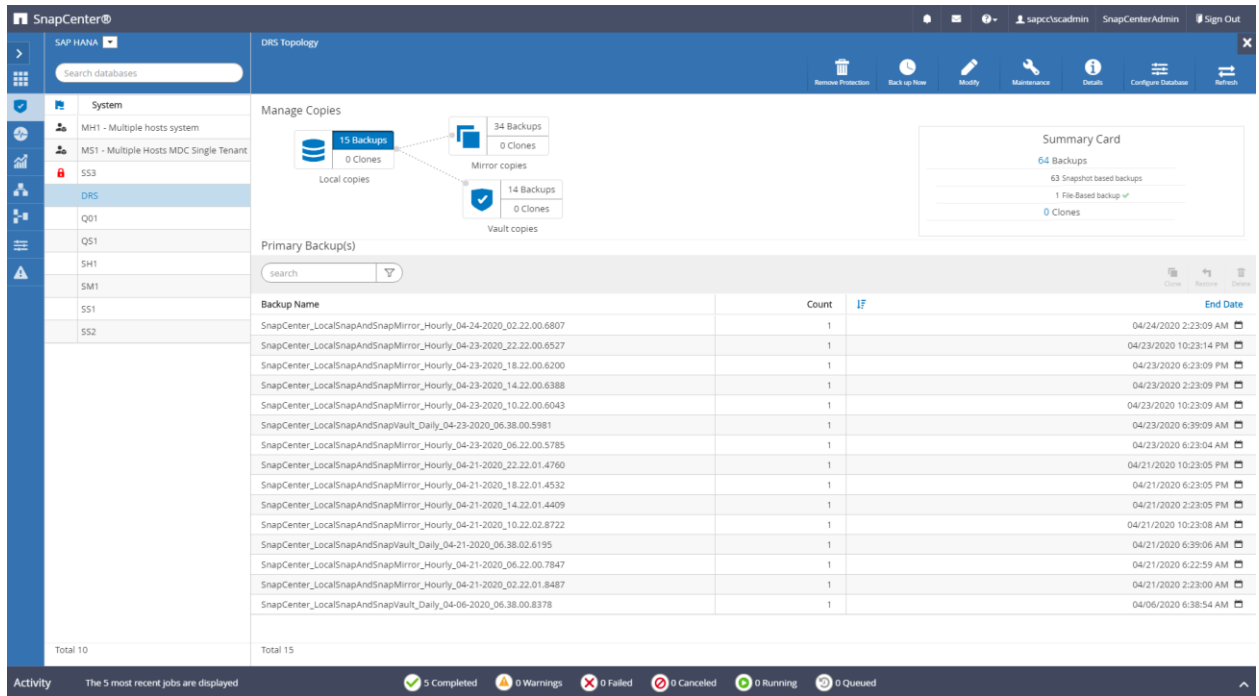
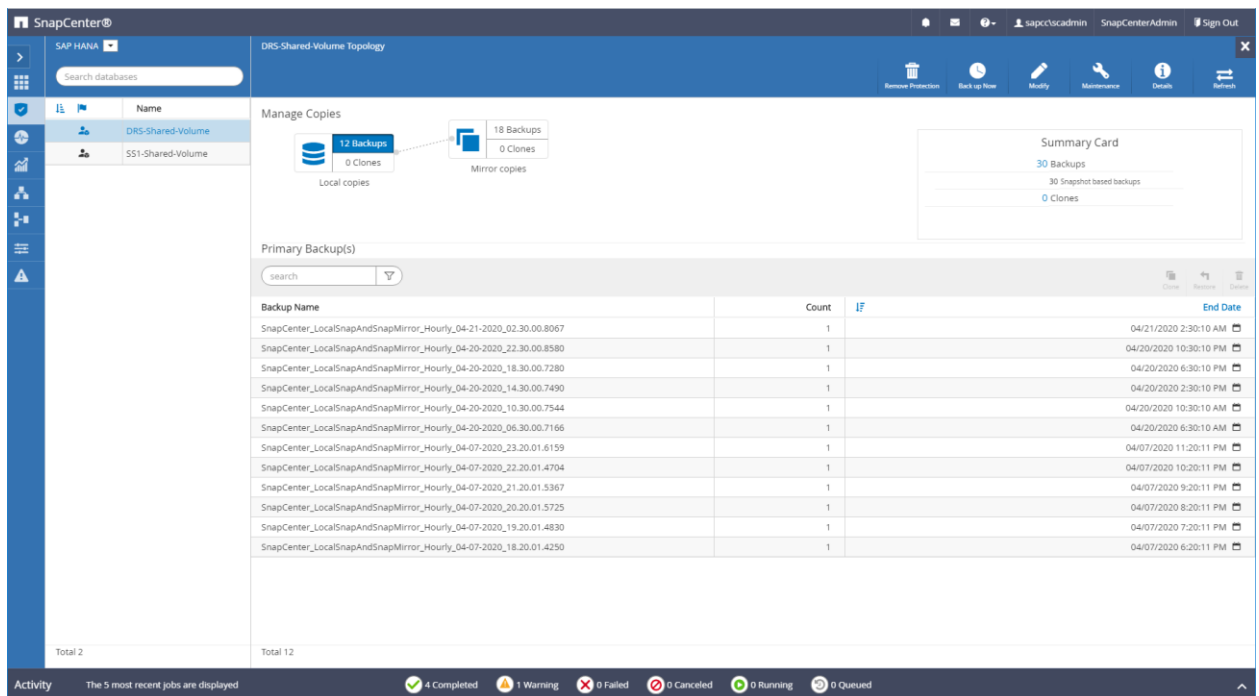


Figure 19) SnapCenter topology view for HANA shared volume.



The details screen of the HANA resource shows the configured SnapCenter policies (Figure 20).

Figure 20) SnapCenter policies.

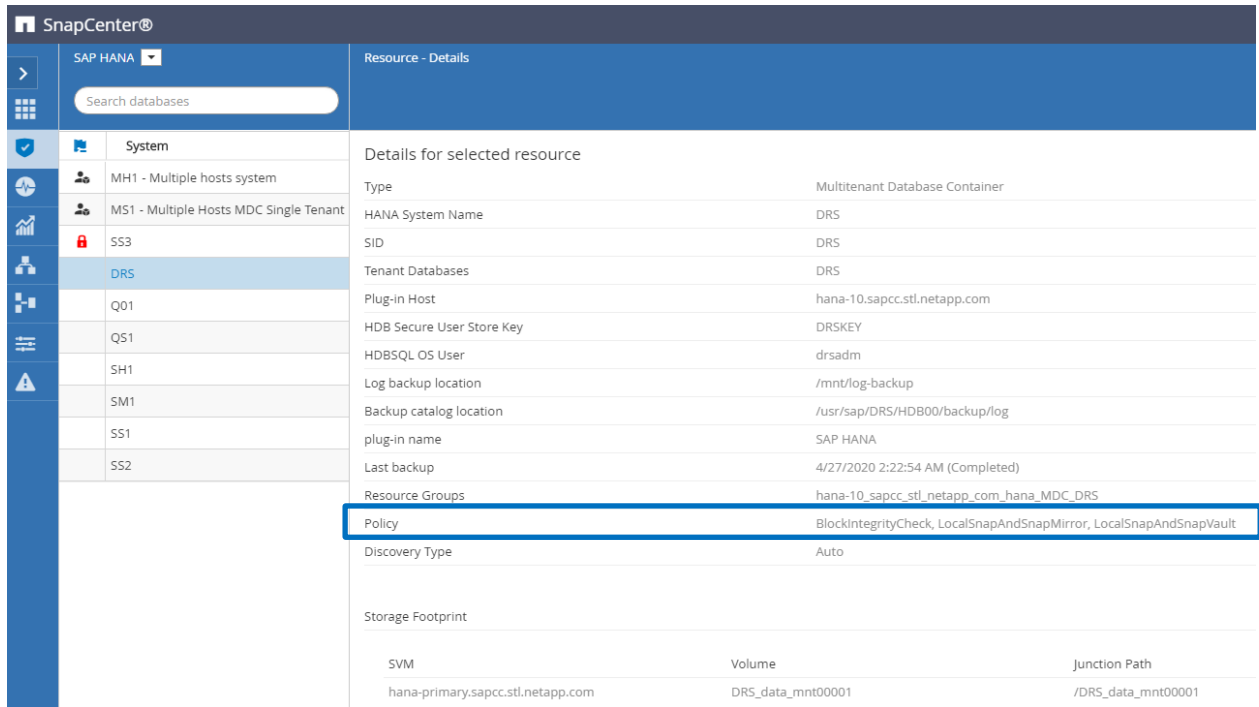
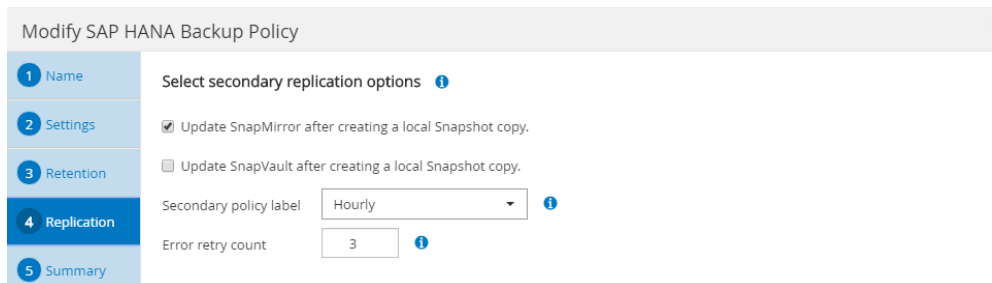


Figure 21 shows the SnapCenter policy for asynchronous SnapMirror. The selected secondary label is not relevant, because MirrorAllSnapshots is configured with the ONTAP protection relationship.

Figure 21) Asynchronous SnapMirror policy.



## 4 Overview of Disaster Recovery Testing

An effective disaster recovery strategy requires testing the required workflow. Testing demonstrates whether the strategy works, and whether the internal documentation is sufficient, and it also allows administrators to train on the required procedures.

Storage replication with SnapMirror makes it possible to execute disaster recovery testing without putting RTO and RPO at risk. Disaster recovery testing can be done without interrupting data replication.

Disaster recovery testing for both asynchronous and synchronous SnapMirror uses Snapshot backups and FlexClone volumes at the disaster recovery target.

Table 4 is a high-level overview of the required steps. The following sections describe the disaster recovery failover test workflow in detail.

Table 4) Disaster recovery testing – required steps.

|   | Synchronous SnapMirror  | Asynchronous SnapMirror  |
|---|---|--|
| Prepare target host for testing   | <ol style="list-style-type: none"> <li>1. Install SAP host agent.</li> <li>2. Configure user, ports, SAP services.</li> <li>3. Create mount points.</li> <li>4. Prepare <code>/etc/fstab</code>.</li> </ol>                 | <ol style="list-style-type: none"> <li>1. Install SAP host agent.</li> <li>2. Configure user, ports, SAP services.</li> <li>3. Create mount points.</li> <li>4. Mount new, empty log volume, and create subdirectories identical to source system.</li> <li>5. Prepare <code>/etc/fstab</code>.</li> </ol> |
| Provide Snapshot backup at the source storage system  | <p>ONTAP CLI:</p> <ul style="list-style-type: none"> <li>• Manual, on-demand Snapshot creation for data (crash consistent), log, and shared volume. ONTAP CLI must be used to set the required SnapMirror label.</li> </ul> | <p>SnapCenter:</p> <ul style="list-style-type: none"> <li>• Application-consistent Snapshot backups of the data volume are available.</li> <li>• Snapshot backups of the shared volume are available.</li> </ul>   |
| Create FlexClone volumes at the disaster recovery storage   | <p>ONTAP System Manager or CLI:</p> <ol style="list-style-type: none"> <li>1. Create FlexClone volumes for data, log, and shared volume, based on Snapshot created before.</li> </ol>                                       | <p>ONTAP System Manager or CLI:</p> <ol style="list-style-type: none"> <li>1. Create FlexClone volumes for data and shared volume, based on SnapCenter Snapshot backups.</li> <li>2. Create FlexClone volume of the log backup volume (if log backup replication is part of the DR concept).</li> </ol>    |
| Mount FlexClone volumes at target host<br><br><b>Note:</b> This step requires additional LUN discovery operations in an FC SAN environment. | <ol style="list-style-type: none"> <li>1. Mount data, log, and shared volume.</li> </ol>  | <ol style="list-style-type: none"> <li>1. Mount data and shared volume.</li> <li>2. Mount log backup volume (if log backup replication is part of the DR concept).</li> </ol>  |
| Start or recover the HANA database  | <ol style="list-style-type: none"> <li>1. Start SAP services.</li> <li>2. Start HANA database.</li> </ol> <p>Crash recovery is executed.</p>  | <ol style="list-style-type: none"> <li>1. Start SAP services.</li> </ol> <p>HANA database is recovered to the last backup (or recovered with forward recovery using log backups if log backup replication is part of the DR concept).</p>  |

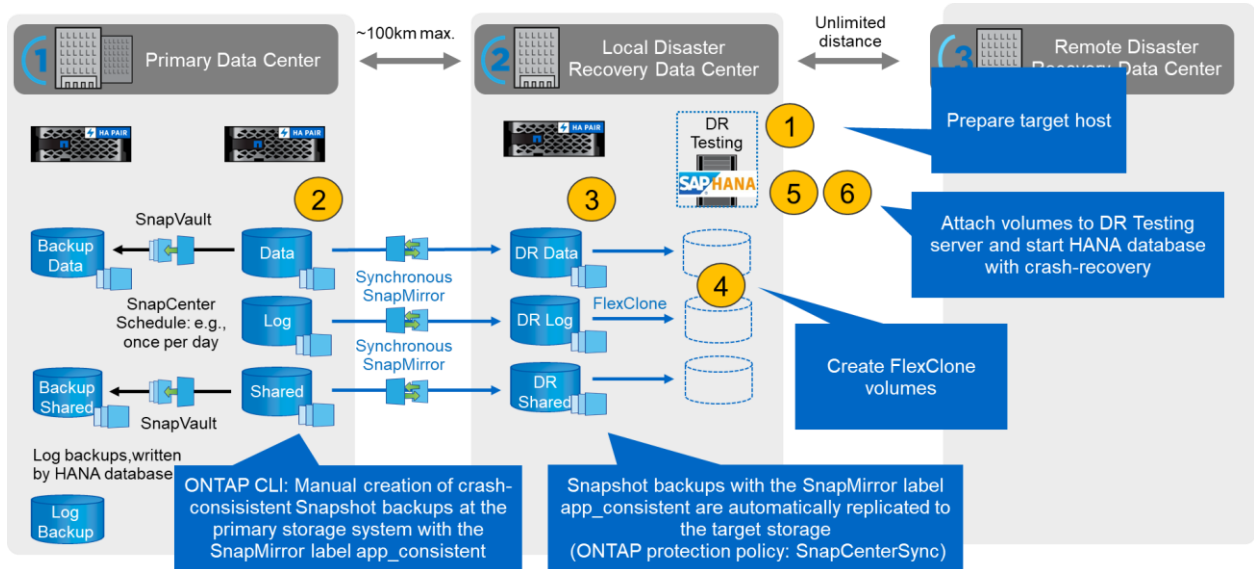
## 5 Synchronous SnapMirror Disaster Recovery Testing

To configure synchronous SnapMirror disaster recovery testing, follow these steps, as shown in Figure 22:

1. Prepare the target host.
2. Create Snapshot backups at the source storage system.
3. Snapshot backups are replicated to the target storage.
4. Create FlexClone volumes at the disaster recovery storage.
5. Mount FlexClone volumes at the target host.
6. Start the HANA database.

The following subsections describe these steps in detail.

Figure 22) Synchronous SnapMirror disaster recovery testing.



## 5.1 Prepare the Target Host

This section describes the preparation steps required at the server, which is used for the disaster recovery failover testing.

During normal operation, the target host is typically used for other purposes, for example, as a HANA QA or test system. Therefore, most of the described steps must be executed when disaster failover testing is executed. On the other hand, the relevant configuration files, like `/etc/fstab` and `/usr/sap/sapservices` can be prepared and then put in production by simply copying the configuration file. The disaster recovery testing procedure ensures that the relevant prepared configuration files are configured correctly.

The target host preparation also includes shutting down the HANA QA or test system.

### Target Server Host Name and IP Address

The host name of the target server must be identical to the host name of the source system. The IP address can be different.

**Note:** Proper fencing of the target server must be established so that it cannot communicate with other systems. If proper fencing is not in place, then the cloned production system might exchange data with other production systems, resulting in logically corrupted data.

### Install Required Software

The SAP host agent software must be installed at the target server. For full information, see the [SAP Host Agent](#) at the SAP help portal.

**Note:** If the host is used as a HANA QA or test system, the SAP host agent software is already installed.

### Configure Users, Ports, and SAP Services

The required users and groups for the SAP HANA database must be available at the target server. Typically, central user management is used; therefore, no configuration steps are necessary at the target server. The required ports for the HANA database must be configured at the target hosts. The

configuration can be copied from the source system by copying the `/etc/services` file to the target server.

The required SAP services entries must be available at the target host. The configuration can be copied from the source system by copying the `/usr/sap/sapservices` file to the target server. The following output shows the required entries for the SAP HANA database used in the lab setup.

```
hana-10:~ # cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/DRS/HDB00/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/DRS/HDB00/exe/sapstartsrv pf=/usr/sap/DRS/SYS/profile/DRS_HDB00_hana-10
-D -u drsadm
limit.descriptors=1048576
```

### Prepare Log Backup Volume

Because the source system is configured with a separate volume for the HANA log backups, a log backup volume must also be available at the target host.

A volume for the log backups must be configured and mounted at the target host.

### Prepare File System Mounts

Table 5 shows the naming conventions used in the lab setup. The volume names of the FlexClone copies at the disaster recovery storage are included in `/etc/fstab`. These volume names are used in the FlexClone copy creation step in the next section.

Table 5) Volume names of FlexClone volumes.

| HANA DRS Volumes | FlexClone Volume at Disaster Recovery Storage                                     | Mount Point at Target Host    |
|------------------|---|-------------------------------|
| Data volume      | DRS_data_mnt00001_dest_sm_s_clone   | /hana/data/DRS/mnt00001       |
| Log volume       | DRS_log_mnt00001_dest_sm_s_clone  | /hana/log/DRS/mnt00001        |
| Shared volume    | DRS_hana_shared_dest_sm_s_clone/shared<br>DRS_hana_shared_dest_sm_s_clone/usr-sap | /hana/shared<br>/usr/sap//DRS |

The mount points in Table 5 must be created at the target host.

Here are the required `/etc/fstab` entries.

```
hana-10:/mnt/log-backup # cat /etc/fstab
192.168.175.116:/DR_testing_log_backup /mnt/log-backup nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0

192.168.175.116:/DRS_data_mnt00001_dest_sm_s_clone /hana/data/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_log_mnt00001_dest_sm_s_clone /hana/log/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_shared_dest_sm_s_clone/shared /hana/shared nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_shared_dest_sm_s_clone/usr-sap /usr/sap/DRS nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
```

**Note:** In a Fibre Channel setup, the `fstab` entries depend on the multipath configuration. Because the SCSI IDs are different with each operation, the `fstab` file is not static and must be adapted after the LUN discovery process.

## 5.2 Create Snapshot Backups at the Source Storage System

The crash-consistent Snapshot backups must be created using the ONTAP CLI to include the SnapMirror label `app_consistent`, which is required so that the Snapshot backups are also replicated to the target storage system.

```
a700-marco::> snap create -vserver hana-primary -volume DRS_data_mnt00001 -snapshot crash-consistent -snapmirror-label app_consistent

a700-marco::> snap create -vserver hana-primary -volume DRS_log_mnt00001 -snapshot crash-consistent -snapmirror-label app_consistent

a700-marco::> snap create -vserver hana-primary -volume DRS_shared -snapshot crash-consistent -snapmirror-label app_consistent
```

The following command output shows the Snapshot backups at the synchronous SnapMirror targets.

Data volume:

```
a700-marco::> snap show -vserver hana-dr -volume DRS_data_mnt00001_dest_sm_s
---Blocks---
Vserver Volume Snapshot Size Total% Used%
-----
hana-dr DRS_data_mnt00001_dest_sm_s
        snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901726.2020-05-11_002402
                                                180KB 0% 0%
        snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901726.2020-05-11_002501
                                                100.5MB 1% 1%
        crash-consistent
                                                184KB 0% 0%
3 entries were displayed.
```

Log volume:

```
a700-marco::> snap show -vserver hana-dr -volume DRS_log_mnt00001_dest_sm_s
---Blocks---
Vserver Volume Snapshot Size Total% Used%
-----
hana-dr DRS_log_mnt00001_dest_sm_s
        snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901728.2020-05-10_230500
                                                400KB 0% 0%
        snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901728.2020-05-11_000500
                                                388KB 0% 0%
        crash-consistent
                                                220KB 0% 0%
3 entries were displayed.
```

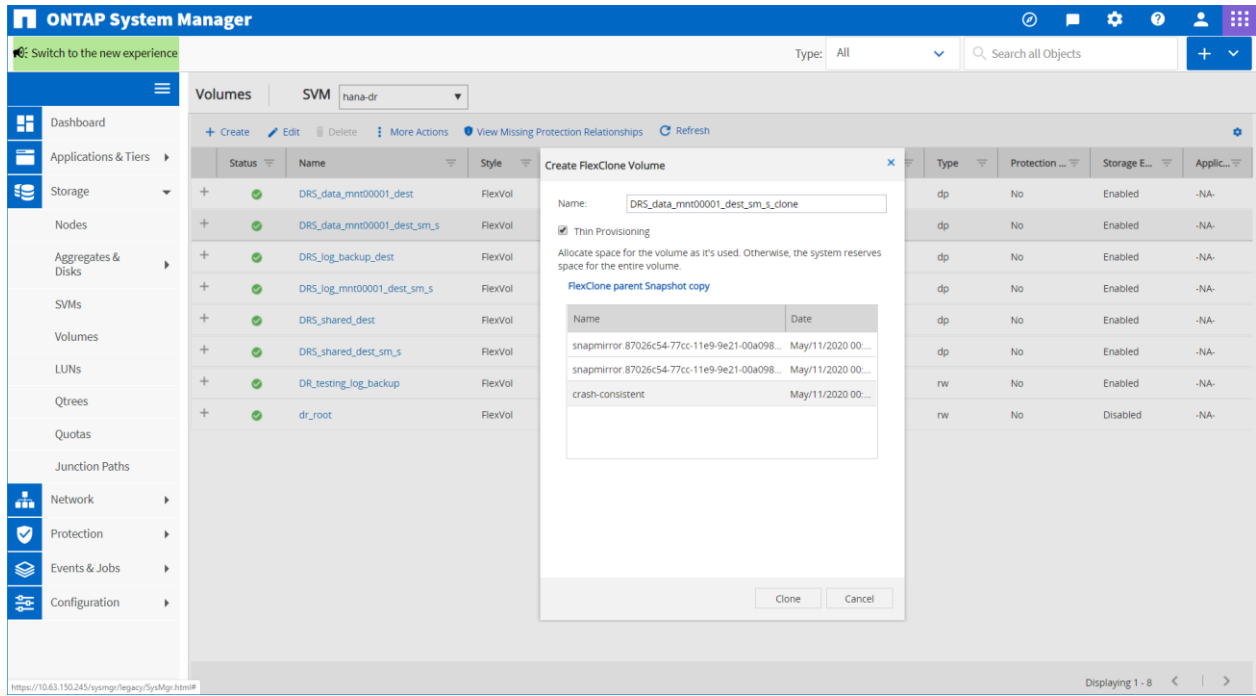
Shared volume:

```
a700-marco::> snap show -vserver hana-dr -volume DRS_shared_dest_sm_s
---Blocks---
Vserver Volume Snapshot Size Total% Used%
-----
hana-dr DRS_shared_dest_sm_s
        snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901727.2020-05-10_230500
                                                392KB 0% 0%
        snapmirror.87026c54-77cc-11e9-9e21-00a098d994db_2157901727.2020-05-11_000500
                                                428KB 0% 0%
        crash-consistent
                                                268KB 0% 0%
3 entries were displayed.
```

## 5.3 Create FlexClone Volumes at the Disaster Recovery Storage

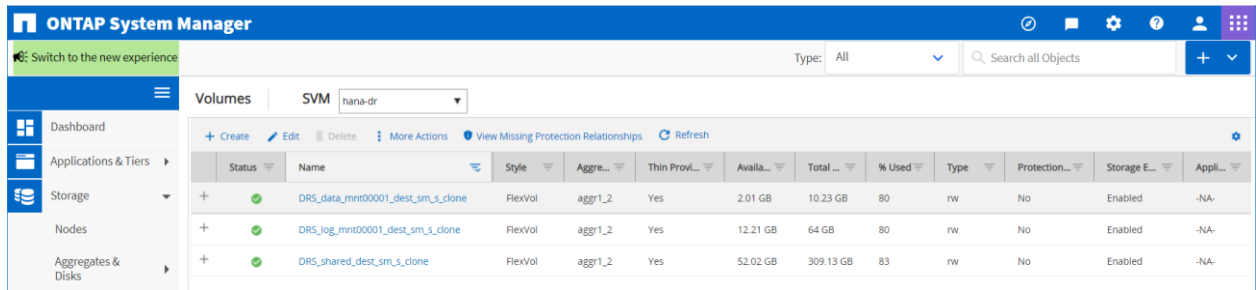
The FlexClone volumes can now be created at the disaster recovery storage by using ONTAP System Manager, as shown in Figure 23.

Figure 23) Create FlexClone volume using crash-consistent Snapshot backup.



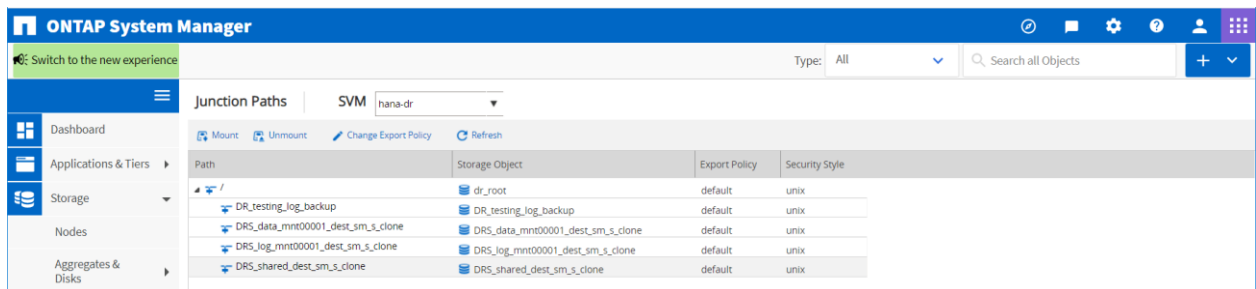
The same operation must be done for the log and the shared volume. Figure 24 shows the list of FlexClone volumes.

Figure 24) List of FlexClone volumes.



All three FlexClone volumes must be mounted to the namespace, as shown in Figure 25.

Figure 25) Junction paths configuration.





**Note:** In a Fibre Channel setup, the LUNs in the FlexClone volumes must be mapped to the initiator group of the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

## 5.4 Mount FlexClone Volumes at the Target Host

The FlexClone volumes can now be mounted at the target host.

```
hana-10:/mnt/log-backup # mount -a
```

The following output shows the required file systems.

```
hana-10:/mnt/log-backup # df
Filesystem                                1K-blocks      Used  Available  Use%  Mounted
on
192.168.175.116:/DR_testing_log_backup    104857600        256  104857344    1%
/mnt/log-backup
192.168.175.116:/DRS_data_mnt00001_dest_sm_s_clone  10725184   8612864    2112320   81%
/hana/data/DRS/mnt00001
192.168.175.116:/DRS_log_mnt00001_dest_sm_s_clone  67106816  54305408   12801408   81%
/hana/log/DRS/mnt00001
192.168.175.116:/DRS_shared_dest_sm_s_clone/shared  324150976 269605952   54545024   84%
/hana/shared
192.168.175.116:/DRS_shared_dest_sm_s_clone/usr-sap  324150976 269605952   54545024   84%
/usr/sap/DRS
hana-10:/mnt/log-backup #
```

**Note:** In a Fibre Channel setup, additional steps are required before the LUNs can be mounted at the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

## 5.5 Start the HANA Database

Start the required SAP services.

```
hana-10:/mnt/log-backup # systemctl start sapinit
```

The following output shows the required processes.

```
hana-10:/mnt/log-backup # ps -ef | grep sap
drsadm  17468      1  3 02:25 ?                00:00:00 /usr/sap/DRS/HDB00/exe/sapstartsrv
pf=/usr/sap/DRS/SYS/profile/DRS_HDB00_hana-10 -D -u drsadm
root    17485    5714  0 02:25 pts/0          00:00:00 grep --color=auto sap
root    24223      1  0 Mar25 ?                00:04:14 /usr/sap/hostctrl/exe/saphostexec
pf=/usr/sap/hostctrl/exe/host_profile
sapadm  24227      1  0 Mar25 ?                00:50:25 /usr/sap/hostctrl/exe/sapstartsrv
pf=/usr/sap/hostctrl/exe/host_profile -D
root    24443      1  0 Mar25 ?                01:01:43 /usr/sap/hostctrl/exe/saposcol -l -w60
pf=/usr/sap/hostctrl/exe/host_profile
hana-10:/mnt/log-backup #
```

The HANA database can now be started.

During the startup, the HANA database executes a crash recovery.

```
hana-10:/mnt/log-backup # su - drsadm
drsadm@hana-10:/usr/sap/DRS/HDB00> sapcontrol -nr 00 -function StartSystem HDB

11.05.2020 02:27:02
StartSystem
OK
```

Use the `sapcontrol` command to check the Start procedure.

```
drsadm@hana-10:/usr/sap/DRS/HDB00> sapcontrol -nr 00 -function GetSystemInstanceList

11.05.2020 02:27:11
```

```
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
hana-10, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GRAY
11.05.2020 02:27:56
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
hana-10, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
drsadm@hana-10:/usr/sap/DRS/HDB00>
```

The HANA database is now up and running. Additional customer-dependent failover tests can be executed as required.

## 5.6 Cleanup Operations

After finishing the testing process, follow these steps to clean up the environment.

1. Stop the HANA database.
2. Stop the SAP services.
3. Unmount the FlexClone volumes from the target host.
4. Delete the FlexClone volumes.
5. Delete the Snapshot backups.

**Note:** When the Snapshot backups are deleted at the source storage system, there can be delay until they are also deleted at the synchronous SnapMirror target. As soon as a new automated SnapMirror Snapshot is created (for resync operations), the Snapshot delete operation also happens at the target storage. If it is required to delete the Snapshot immediately, it can be deleted directly at the target storage.

## 6 Asynchronous SnapMirror Disaster Recovery Testing

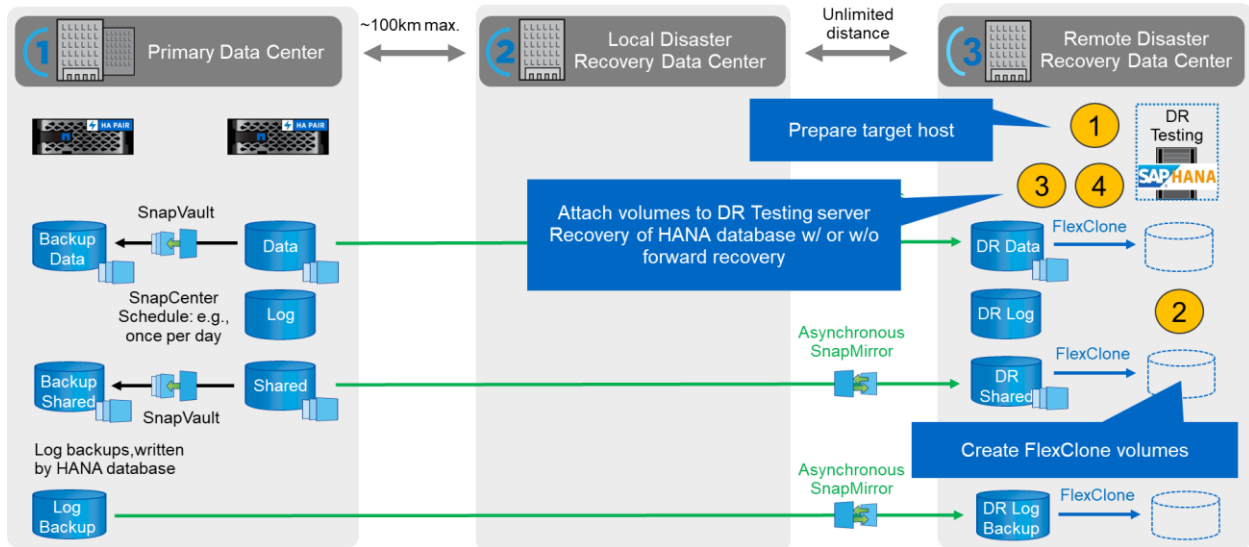
Depending on whether the log backup replication is part of the disaster recovery setup, the steps for disaster recovery are slightly different. This section describes the disaster recovery testing for data-backup-only replication. It also documents the differences if the log backups are also replicated.

To configure asynchronous SnapMirror disaster recovery testing, follow these steps.

1. Prepare the target host.
2. Create FlexClone volumes at the disaster recovery storage.
3. Mount the FlexClone volumes at the target host.
4. Recover the HANA database.
  - Data volume recovery only.
  - Forward recovery using replicated log backups.

The following subsections describe these steps in detail, as shown in Figure 26.

Figure 26) Asynchronous SnapMirror disaster recovery testing.



## 6.1 Prepare the Target Host

This section describes the preparation steps required at the server, which is used for the disaster recovery failover testing.

During normal operation, the target host is typically used for other purposes, for example, as a HANA QA or test system. Therefore, most of the described steps must be executed when disaster failover testing is executed. On the other hand, the relevant configuration files, like `/etc/fstab` and `/usr/sap/sapservices`, can be prepared and then put in production by simply copying the configuration file. The disaster recovery testing procedure ensures that the relevant prepared configuration files are configured correctly.

The target host preparation also includes shutting down the HANA QA or test system.

### Target Server Host Name and IP Address

The host name of the target server must be identical to the host name of the source system. The IP address can be different.

**Note:** Proper fencing of the target server must be established so that it cannot communicate with other systems. If proper fencing is not in place, then the cloned production system might exchange data with other production systems, resulting in logically corrupted data.

### Install Required Software

The SAP host agent software must be installed at the target server. For full information, see the [SAP Host Agent](#) at the SAP help portal.

**Note:** If the host is used as a HANA QA or test system, the SAP host agent software is already installed.

### Configure Users, Ports, and SAP Services

The required users and groups for the SAP HANA database must be available at the target server. Typically, central user management is used; therefore, no configuration steps are necessary at the target server. The required ports for the HANA database must be configured at the target hosts. The configuration can be copied from the source system by copying the `/etc/services` file to the target server.

The required SAP services entries must be available at the target host. The configuration can be copied from the source system by copying the `/usr/sap/sapservices` file to the target server. The following output shows the required entries for the SAP HANA database used in the lab setup.

```
hana-10:~ # cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/DRS/HDB00/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/DRS/HDB00/exe/sapstartsrv pf=/usr/sap/DRS/SYS/profile/DRS_HDB00_hana-10
-D -u drsadm
limit.descriptors=1048576
```

### Prepare HANA Log Volume

Because the HANA log volume is not part of the SnapMirror replication, an empty log volume must exist at the target host. The log volume must include the same subdirectories as the source HANA system.

```
hana-10:/hana/log/DRS/mnt00001 # ls -al
total 84
drwxr-x--- 5 drsadm sapsys 4096 Apr  2 06:51 .
drwxr-x--- 1 drsadm sapsys  16 Apr  2 06:45 ..
drwxr-x--- 2 drsadm sapsys 61440 Apr 24 07:11 hdb00001
drwxr-xr-- 2 drsadm sapsys 12288 Apr 24 02:59 hdb00002.00003
drwxr-xr-- 2 drsadm sapsys  4096 Apr 24 08:31 hdb00003.00003
hana-10:/hana/log/DRS/mnt00001 #
```

### Prepare Log Backup Volume

Because the source system is configured with a separate volume for the HANA log backups, a log backup volume must also be available at the target host. A volume for the log backups must be configured and mounted at the target host.

**Note:** If log backup volume replication is part of the disaster recovery setup, a FlexClone volume is mounted at the target host and it is not necessary to prepare an additional log backup volume.

### Prepare File System Mounts

Table 6 shows the naming conventions used in the lab setup. The volume names of the FlexClone volumes at the disaster recovery storage are included in `/etc/fstab`. These volume names are used in the FlexClone volume creation step in the next section.

**Table 6) Volume names of FlexClone copies.**

| HANA DRS Volumes  | FlexClone Volume at Disaster Recovery Storage                           | Mount Point at Target Host    |
|-------------------|---|-------------------------------|
| Data volume       | DRS_data_mnt00001_dest_clone  | /hana/data/DRS/mnt00001       |
| Shared volume     | DRS_hana_shared_dest_clone/shared<br>DRS_hana_shared_dest_clone/usr-sap | /hana/shared<br>/usr/sap//DRS |
| Log backup volume | DRS_log_backup_dest_clone   | /mnt/log-backup               |

**Note:** A FlexClone copy of the log backup volume is required only if log backups are also replicated to the DR site.

**Note:** The mount points in Table 6 must be created at the target host.

Here are the required `/etc/fstab` entries.

```
hana-10:~ # cat /etc/fstab
192.168.175.116:/DRS_log_mnt00001 /hana/log/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0

192.168.175.116:/DRS_data_mnt00001_dest_clone /hana/data/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
```

```

192.168.175.116:/DRS_shared_dest_clone/shared /hana/shared nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_shared_dest_clone/usr-sap /usr/sap/DRS nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_log_backup_dest_clone /mnt/log-backup nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
hana-10:~ #

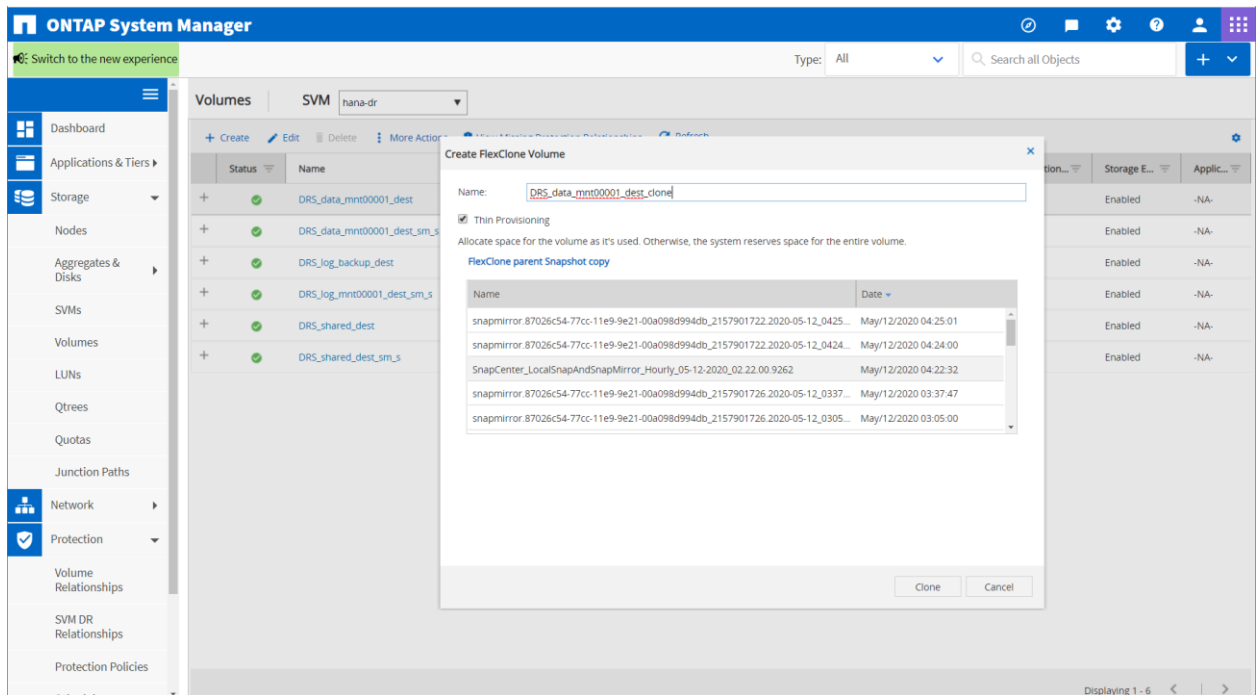
```

**Note:** In a Fibre Channel setup, the `fstab` entries depend on the multipath configuration. Because the SCSI IDs are different with each operation, the `fstab` file is not static and must be adapted after the LUN discovery process.

## 6.2 Create FlexClone Volumes at the Disaster Recovery Storage

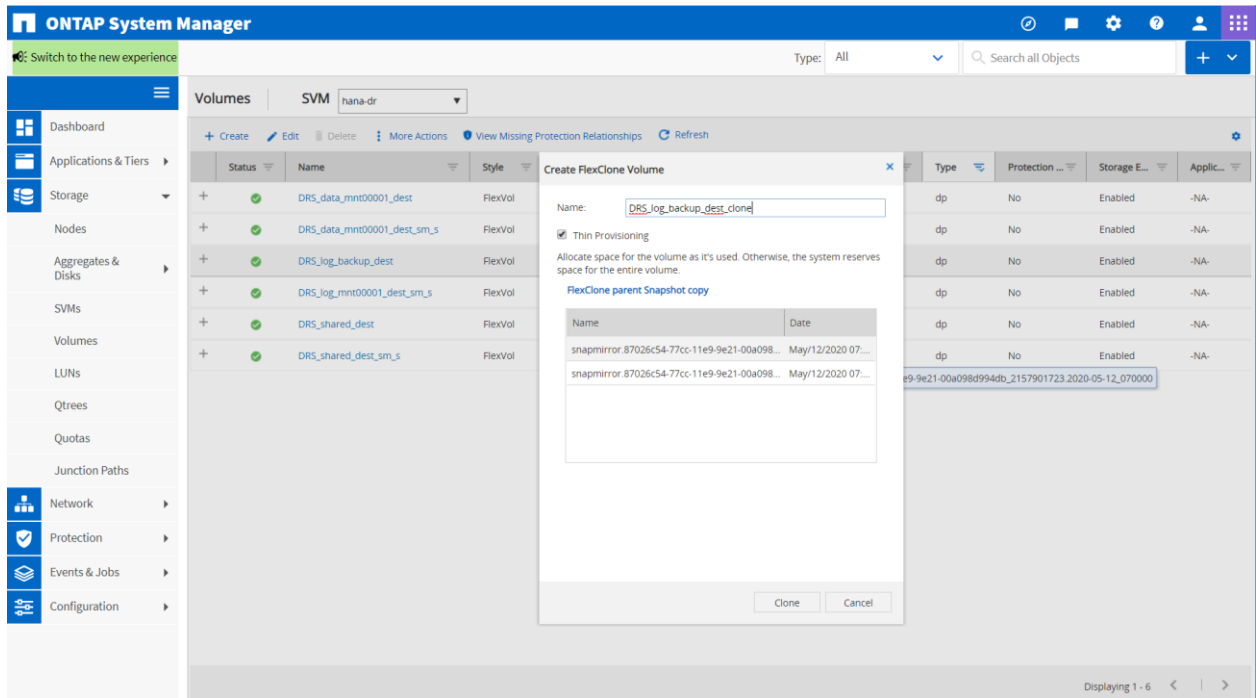
Depending on the disaster recovery setup—with or without log backup replication—two or three FlexClone volumes must be created. In both cases, a FlexClone volume of the data and the shared volume must be created. A FlexClone volume of the log backup volume must be created if the log backup data is also replicated. Figure 27 through Figure 30 show the required steps using ONTAP System Manager.

Figure 27) Create a FlexClone volume based on an application-consistent SnapCenter backup.



One of the SnapCenter backups is selected as a volume source for the FlexClone volume of the HANA data volume.

Figure 28) Create a FlexClone copy of the log backup volume.



As a source for the log backup FlexClone volume, one of the SnapMirror Snapshot copies is selected.

Figure 29) List of FlexClone volumes.

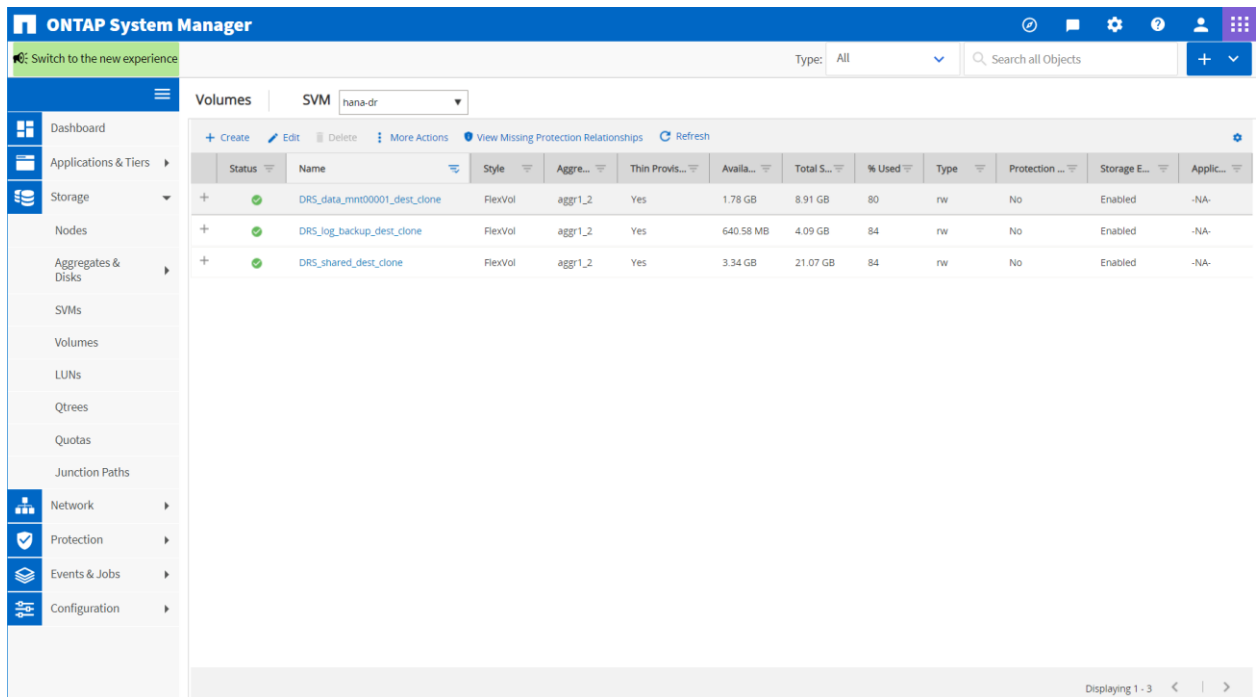
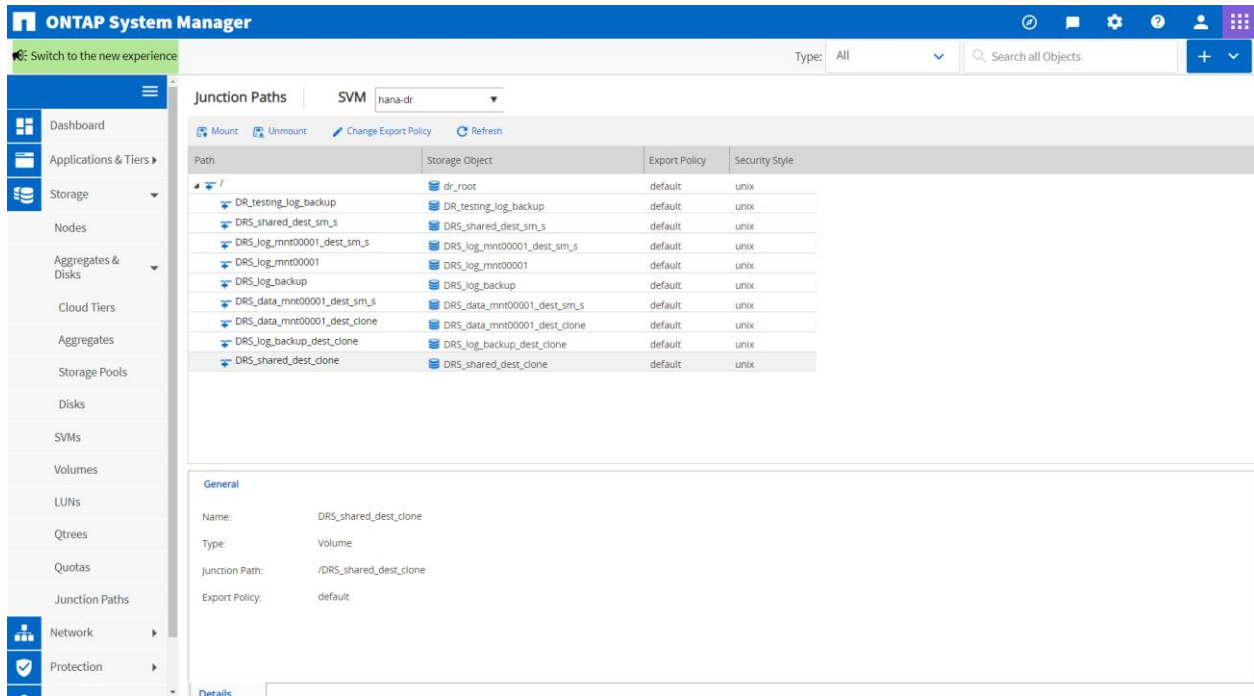


Figure 30) Junction path configuration.



All volumes must be mounted to the namespace.

**Note:** In a Fibre Channel setup, the LUNs in the FlexClone volumes must be mapped to the initiator group of the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

### 6.3 Mount the FlexClone Volumes at the Target Host

The FlexClone volumes can now be mounted at the target host.

```
hana-10:~ # mount -a
```

The following output shows the required file systems.

```
hana-10:~ # df
Filesystem                                1K-blocks      Used  Available Use% Mounted on
192.168.175.116:/DRS_log_mnt00001        104857600         512  104857088   1% /hana/log/DRS/mnt00001
192.168.175.116:/DRS_data_mnt00001_dest_clone  9343040    7479488   1863552   81% /hana/data/DRS/mnt00001
192.168.175.116:/DRS_shared_dest_clone/shared  22094464  18594048   3500416   85% /hana/shared
192.168.175.116:/DRS_shared_dest_clone/usr-sap  22094464  18594048   3500416   85% /usr/sap/DRS
192.168.175.116:/DR_testing_log_backup    104857600         256  104857344   1% /mnt/log-backup
```

**Note:** If log backups are also replicated, the log backup FlexClone volume would be mounted instead of the above log backup volume.

```
192.168.175.116:/DRS_log_backup_dest_clone  4288448    3632448   656000   85% /mnt/log-backup
hana-10:~ #
```

**Note:** In a Fibre Channel setup, additional steps are required before the LUNs can be mounted at the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

## 6.4 Check Consistency of Latest Log Backups

Because the log backup volume replication is performed independently of the log backup process executed by the SAP HANA database, there might be open, inconsistent log backup files at the disaster recovery site. Only the latest log backup files might be inconsistent, and those files should be checked before a forward recovery is performed at the disaster recovery site.

To check the consistency of the latest log backups, follow these steps.

1. List the latest log backups in the log backup directory. The following output shows the latest log backups for the system and the tenant database in the lab environment.

```
drsdm@hana-10:/usr/sap/DRS/home> ls -altr /mnt/log-backup/SYSTEMDB
.....
-rw-r----- 1 drsdm sapsys 1396736 May 12 08:38
log_backup_1_0_1619642560_1619664192.1589287111110
-rw-r----- 1 drsdm sapsys 1232896 May 12 08:53
log_backup_1_0_1619664192_1619683264.1589288011113
-rw-r----- 1 drsdm sapsys 1236992 May 12 09:08
log_backup_1_0_1619683264_1619702400.1589288911115
-rw-r----- 1 drsdm sapsys 561152 May 12 09:08 log_backup_0_0_0_0.1589288911144
-rw-r----- 1 drsdm sapsys 561152 May 12 09:22 log_backup_0_0_0_0.1589289757361
-rw-r----- 1 drsdm sapsys 561152 May 12 09:23 log_backup_0_0_0_0.1589289787462
-rw-r----- 1 drsdm sapsys 561152 May 12 09:23 log_backup_0_0_0_0.1589289788236
-rw-r----- 1 drsdm sapsys 1945600 May 12 09:23
log_backup_1_0_1619702400_1619732608.1589289811118
-rw-r----- 1 drsdm sapsys 561152 May 12 09:23 log_backup_0_0_0_0.1589289811148
```

```
drsdm@hana-10:/usr/sap/DRS/home> ls -altr /mnt/log-backup/DB_DRS/
-rw-r----- 1 drsdm sapsys 118784 May 12 07:48
log_backup_3_0_1016494848_1016496512.1589284100889
-rw-r----- 1 drsdm sapsys 65536 May 12 07:48 log_backup_2_0_14169280_14170112.1589284124059
-rw-r----- 1 drsdm sapsys 61440 May 12 08:03
log_backup_3_0_1016496512_1016497280.1589285000892
-rw-r----- 1 drsdm sapsys 49152 May 12 08:18
log_backup_3_0_1016497280_1016497856.1589285900896
-rw-r----- 1 drsdm sapsys 45056 May 12 08:33
log_backup_3_0_1016497856_1016498368.1589286800899
-rw-r----- 1 drsdm sapsys 118784 May 12 08:48
log_backup_3_0_1016498368_1016500032.1589287700902
-rw-r----- 1 drsdm sapsys 65536 May 12 08:48 log_backup_2_0_14170112_14170944.1589287724082
-rw-r----- 1 drsdm sapsys 90112 May 12 09:03
log_backup_3_0_1016500032_1016501248.1589288600905
-rw-r----- 1 drsdm sapsys 602112 May 12 09:03 log_backup_0_0_0_0.1589288600925
-rw-r----- 1 drsdm sapsys 49152 May 12 09:18
log_backup_3_0_1016501248_1016501824.1589289500908
-rw-r----- 1 drsdm sapsys 606208 May 12 09:18 log_backup_0_0_0_0.1589289500928
-rw-r----- 1 drsdm sapsys 606208 May 12 09:22 log_backup_0_0_0_0.1589289757363
-rw-r----- 1 drsdm sapsys 606208 May 12 09:23 log_backup_0_0_0_0.1589289787926
-rw-r----- 1 drsdm sapsys 606208 May 12 09:23 log_backup_0_0_0_0.1589289788730
drwxr-xr-- 2 drsdm sapsys 176128 May 12 09:23 .
drsdm@hana-10:/usr/sap/DRS/home>
```

2. Check the latest log backups for these services by running the `hdbbackupcheck` command.

```
drsdm@hana-10:/usr/sap/DRS/home> hdbbackupcheck /mnt/log-
backup/SYSTEMDB/log_backup_0_0_0_0.1589289811148
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivercache'
Backup '/mnt/log-backup/SYSTEMDB/log_backup_0_0_0_0.1589289811148' successfully checked.
drsdm@hana-10:/usr/sap/DRS/home> hdbbackupcheck /mnt/log-
backup/SYSTEMDB/log_backup_1_0_1619702400_1619732608.1589289811118
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivercache'
```



```

Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivescache'
Backup '/mnt/log-backup/SYSTEMDB/log_backup_1_0_1619702400_1619732608.1589289811118' successfully
checked.
drsadm@hana-10:/usr/sap/DRS/home>

```

```

drsadm@hana-10:/usr/sap/DRS/home> hdbbackupcheck /mnt/log-
backup/DB_DRS/log_backup_0_0_0_0.1589289788730
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivescache'
Backup '/mnt/log-backup/DB_DRS/log_backup_0_0_0_0.1589289788730' successfully checked.

drsadm@hana-10:/usr/sap/DRS/home> hdbbackupcheck /mnt/log-
backup/DB_DRS/log_backup_3_0_1016501248_1016501824.1589289500908
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivescache'
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivescache'
Backup '/mnt/log-backup/DB_DRS/log_backup_3_0_1016501248_1016501824.1589289500908' successfully
checked.

drsadm@hana-10:/usr/sap/DRS/home> hdbbackupcheck /mnt/log-
backup/DB_DRS/log_backup_2_0_14170112_14170944.1589287724082
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivescache'
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivescache'
Backup '/mnt/log-backup/DB_DRS/log_backup_2_0_14170112_14170944.1589287724082' successfully
checked.
drsadm@hana-10:/usr/sap/DRS/home>

```

If the output of the `hdbbackupcheck` tool shows that the latest log backups are consistent, you can perform a recovery that includes them.

If the `hdbbackupcheck` tool reports an error for the latest log backups, the latest set of log backups must be removed or renamed in the log backup.

## 6.5 Recover the HANA Database

Start the required SAP services.

```
hana-10:~ # systemctl start sapinit
```

The following output shows the required processes.

```

hana-10:~ # ps -ef | grep sap
root      15919      1  0 06:36 ?          00:00:00 /usr/sap/hostctrl/exe/saphostexec
pf=/usr/sap/hostctrl/exe/host_profile
sapadm    15941      1  0 06:36 ?          00:00:00 /usr/lib/systemd/systemd --user
sapadm    15943 15941  0 06:36 ?          00:00:00 (sd-pam)
sapadm    15963      1  3 06:36 ?          00:00:00 /usr/sap/hostctrl/exe/sapstartsrv
pf=/usr/sap/hostctrl/exe/host_profile -D
drsadm    16027      1  1 06:36 ?          00:00:00 /usr/sap/DRS/HDB00/exe/sapstartsrv
pf=/usr/sap/DRS/SYS/profile/DRS_HDB00_hana-10 -D -u drsadm
root      16161      1  0 06:36 ?          00:00:00 /usr/sap/hostctrl/exe/saposcol -l -w60
pf=/usr/sap/hostctrl/exe/host_profile
root      16163    5714  0 06:36 pts/0    00:00:00 grep --color=auto sap
hana-10:~ #

```

The following subsections describe the recovery process using HANA Studio for the system database and the tenant database. The screenshots show recovery without and with forward recovery using replicated log backups.

If HANA Studio is used for recovery, the file

`/hana/shared/DRS/global/hdb/metadata/net_publicname/hana-10` must be adapted to

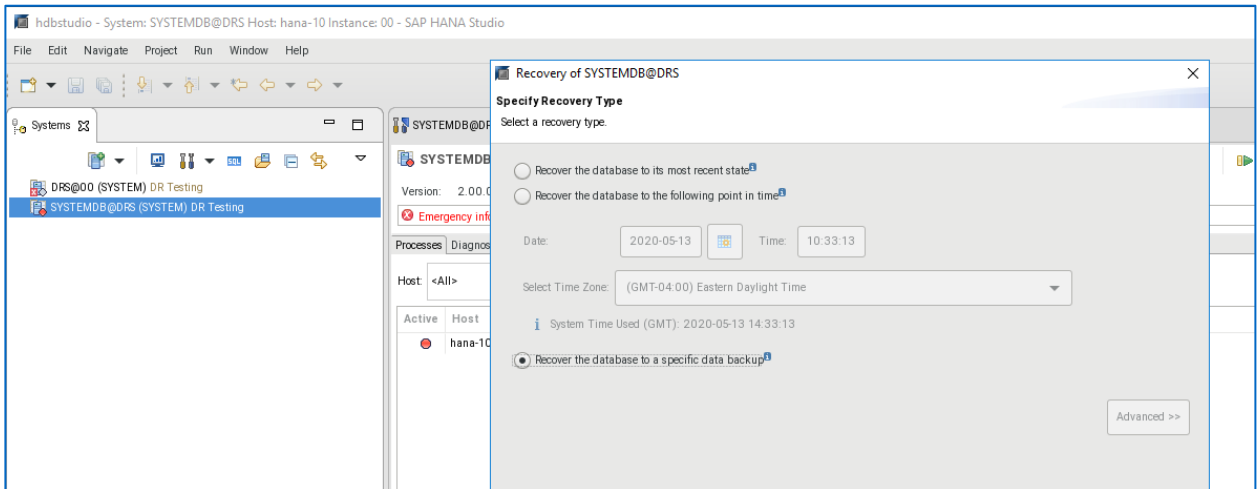
include the correct IP address of the disaster recovery testing server. Otherwise, HANA Studio will try to access the original HANA system.

```
drsdadm@hana-10: /> cat /hana/shared/DRS/global/hdb/metadata/net_publicname/hana-10
10.63.167.175
```

## Recover the System Database

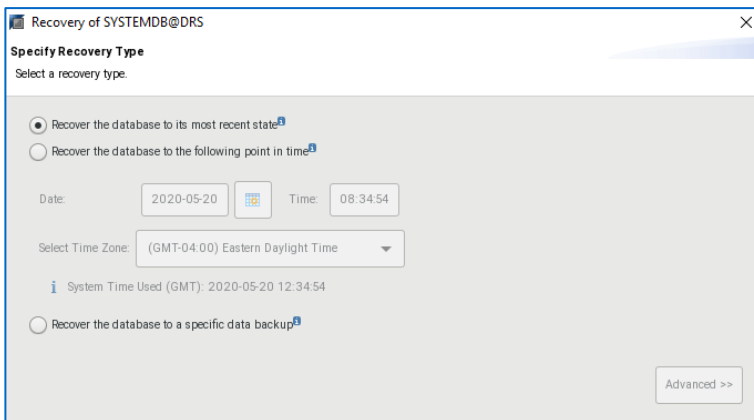
In HANA Studio, select Recover System Database and then select Recover to a Specific Data Backup.

Figure 31) Recovery to a specific data backup: Recovery Type.



**Note:** For forward recovery using the replicated log backups, select the recovery type Recover to Most Recent State.

Figure 32) Recovery with log backups: Recovery Type.



Select Recover Without Backup Catalog.

**Figure 33) Recovery to a specific data backup: Specify backup catalog location.**

The screenshot shows a dialog box titled "Recovery of SYSTEMDB@DRS" with a close button (X) in the top right corner. The main heading is "Specify Backup Location". Below the heading, there is a sub-heading "Specify Backup Location" and a paragraph: "Choose whether you want to select a backup from a backup catalog or enter the name and the path of a backup in the next step." There are two radio button options: "Recover using the backup catalog" (which is unselected) and "Recover without the backup catalog" (which is selected). Under the first option, there is a sub-option "Search for the backup catalog in the file system only" and a text input field for "Backup Catalog Location" containing the path "/usr/sap/DRS/HDB00/backup/log/SYSTEMDB". Under the second option, there is a sub-option "Backint System Copy" with an unchecked checkbox and a text input field for "Source System".

**Note:** For forward recovery using the replicated log backups, the backup catalog location must be specified.

**Figure 34) Recovery with log backups: Backup catalog location.**

The screenshot shows a dialog box titled "Recovery of SYSTEMDB@DRS" with a close button (X) in the top right corner. The main heading is "Locate Backup Catalog". Below the heading, there is a sub-heading "Locate Backup Catalog" and a paragraph: "Specify location of the backup catalog." There are two radio button options: "Recover using the backup catalog" (which is selected) and "Recover without the backup catalog" (which is unselected). Under the first option, there is a sub-option "Search for the backup catalog in the file system only" and a text input field for "Backup Catalog Location" containing the path "/mnt/log-backup/SYSTEMDB". Under the second option, there is a sub-option "Backint System Copy" with an unchecked checkbox and a text input field for "Source System".

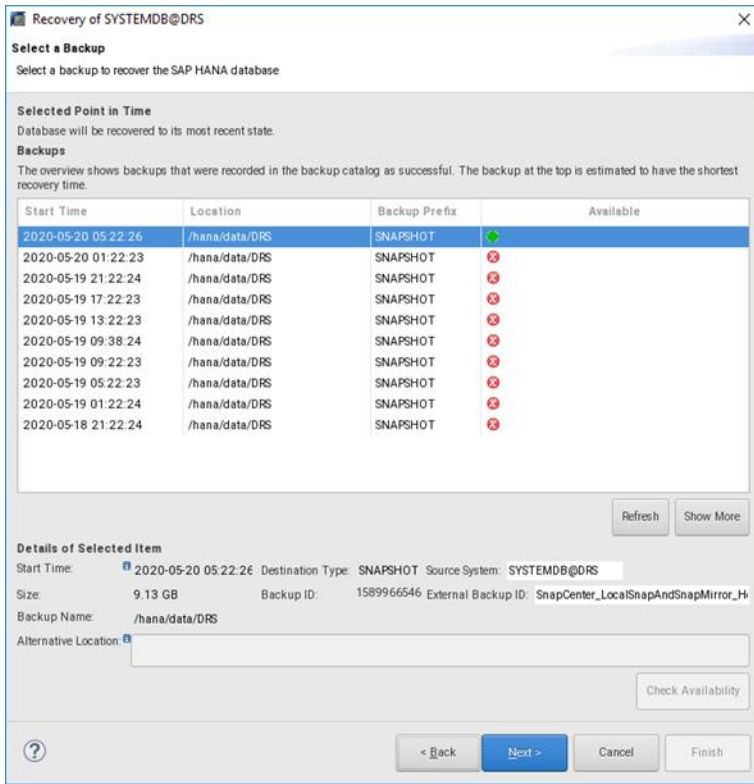
Select Snapshot as the destination type.

**Figure 35) Recovery to a specific data backup: Select Destination Type.**

The screenshot shows a dialog box titled "Recovery of SYSTEMDB@DRS" with a close button (X) in the top right corner. The main heading is "Specify the Backup to Recover". Below the heading, there is a sub-heading "Specify the Backup to Recover" and a paragraph: "Specify the backup to be recovered." There is a "Destination Type" dropdown menu set to "Snapshot". Below this, there is a section "Locate the Data Backup" with a paragraph: "Specify the destination of the data backup that you want to use to recover the database." There are two text input fields: "Location" containing the path "/hana/data/DRS" and "Backup Prefix" which is empty.

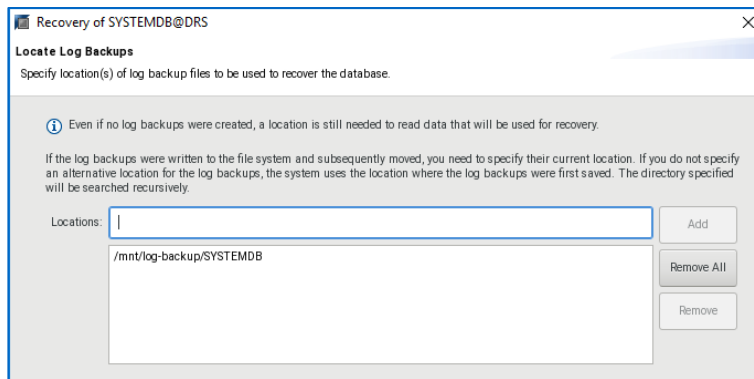
**Note:** For forward recovery using the replicated log backups, HANA Studio shows the availability of the Snapshot copy with a green icon.

Figure 36) Recovery with log backups: Backup selection.



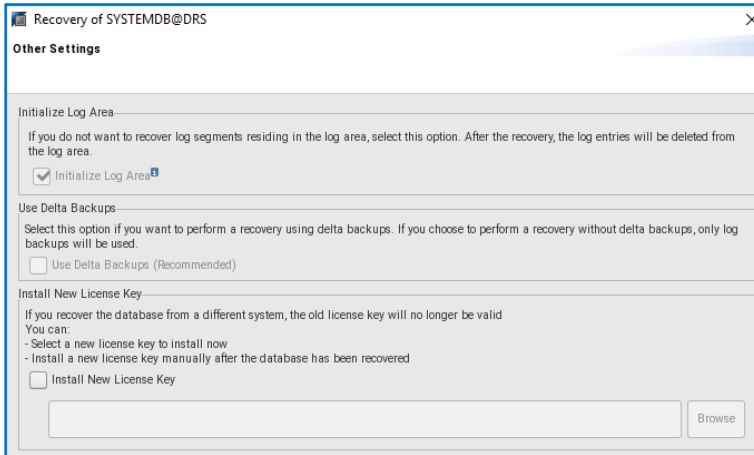
Provide a log backup destination for recovery.

Figure 37) Recovery with log backups: Log backup location.



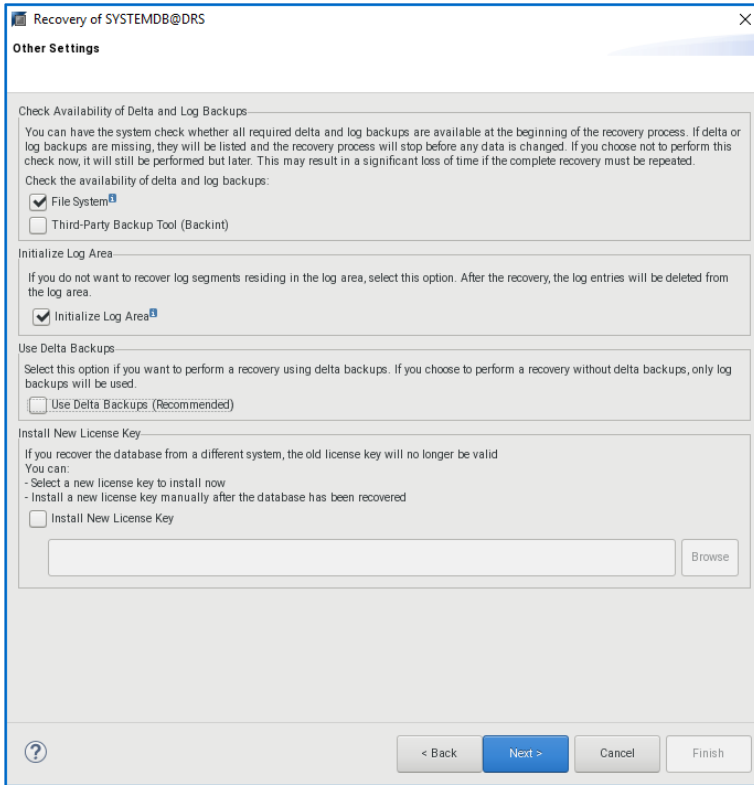
For recovery to a specific backup, Initialize Log Area is preselected.

Figure 38) Recovery to a specific data backup: Other Settings.



**Note:** For forward recovery using the replicated log backups, Initialize Log Area must be selected.

Figure 39) Recovery with log backups: Other Settings.



Figures 40, 41, and 42 summarize the different recovery processes.

Figure 40) Recovery to a specific data backup: Review Recovery Settings.

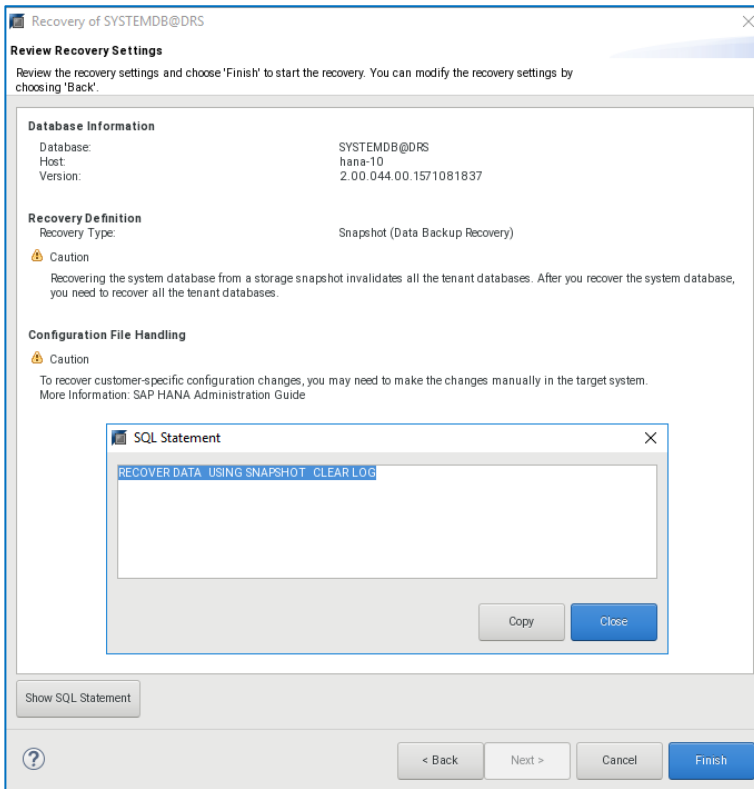


Figure 41) Recovery with log backups: Review Recovery Settings.

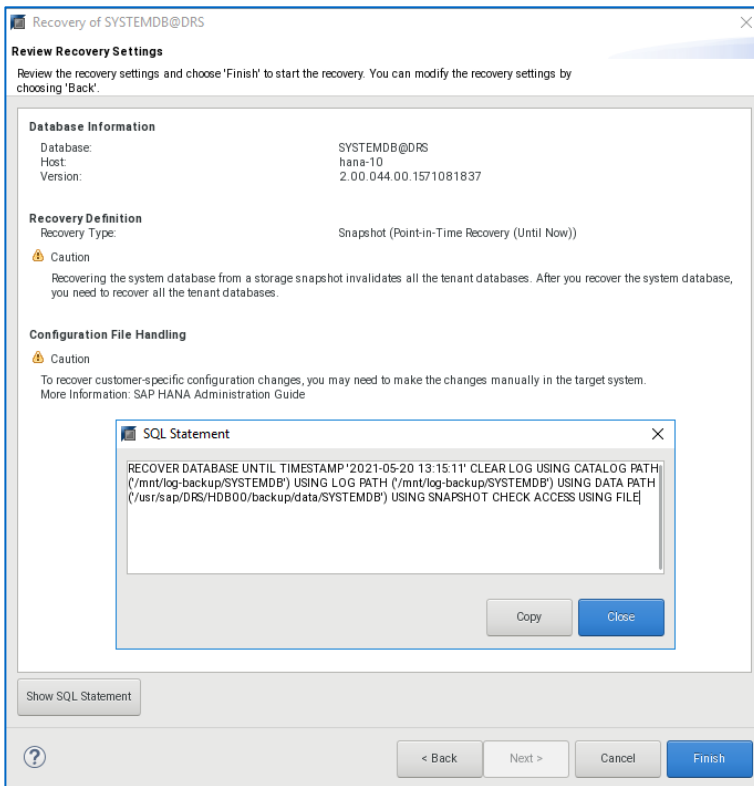
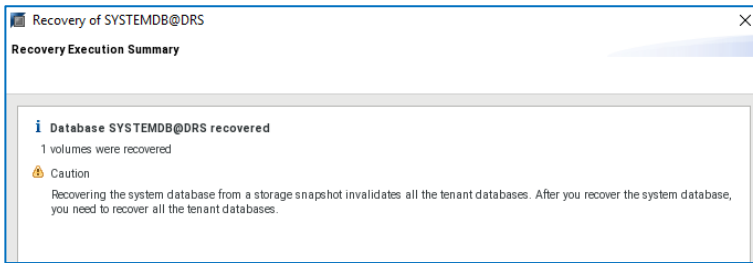


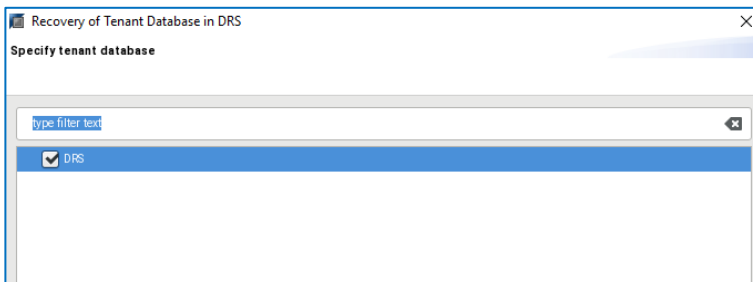
Figure 42) Recovery to a specific data backup: Recovery Execution Summary.



## Recover a Tenant Database

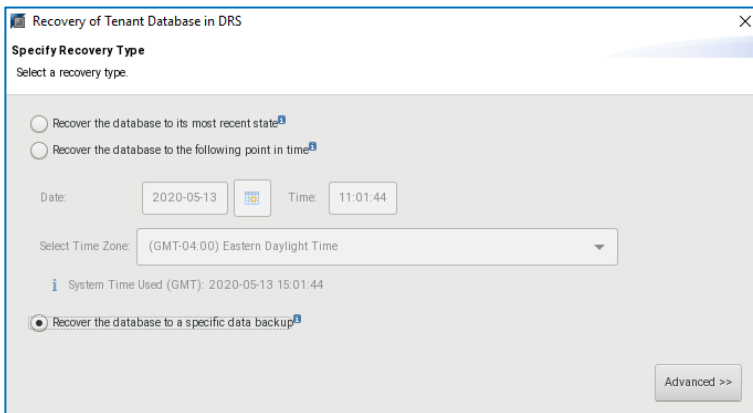
In HANA Studio, select Recover Tenant Database and then select a tenant.

Figure 43) Recovery of Tenant Database in DRS.



Select the recovery type Recover the Database to a Specific Data Backup.

Figure 44) Recovery to a specific data backup: Recovery Type.



**Note:** For forward recovery using the replicated log backups, select the recovery type Recover the Database to the Most Recent State.

Figure 45) Recovery with log backups: Specify Recovery Type.

The screenshot shows a dialog box titled "Recovery of Tenant Database in DRS" with a close button (X) in the top right corner. The main heading is "Specify Recovery Type" and the instruction is "Select a recovery type." There are three radio button options: "Recover the database to its most recent state" (which is selected), "Recover the database to the following point in time", and "Recover the database to a specific data backup". The "Recover the database to the following point in time" option is expanded to show a "Date" field with "2020-05-25", a "Time" field with "05:21:03", and a "Select Time Zone" dropdown menu set to "(GMT-04:00) Eastern Daylight Time". Below this is a small information icon and the text "System Time Used (GMT): 2020-05-25 09:21:03". The "Recover the database to a specific data backup" option is also expanded to show a "Source System" text field. At the bottom right, there is an "Advanced >>" button.

Select Recover Without the Backup Catalog.

Figure 46) Recovery to DRS a specific data backup: Specify backup catalog location.

The screenshot shows a dialog box titled "Recovery of Tenant Database in DRS" with a close button (X) in the top right corner. The main heading is "Specify Backup Location" and the instruction is "Choose whether you want to select a backup from a backup catalog or enter the name and the path of a backup in the next step." There are two radio button options: "Recover using the backup catalog" and "Recover without the backup catalog" (which is selected). Under "Recover using the backup catalog", there is a sub-option "Search for the backup catalog in the file system only" and a "Backup Catalog Location" text field containing "/usr/sap/DRS/HDB00/backup/log/DB\_DRS". Under "Recover without the backup catalog", there is a sub-section "Backint System Copy" with a "Backint System Copy" checkbox (unchecked) and a "Source System" text field.

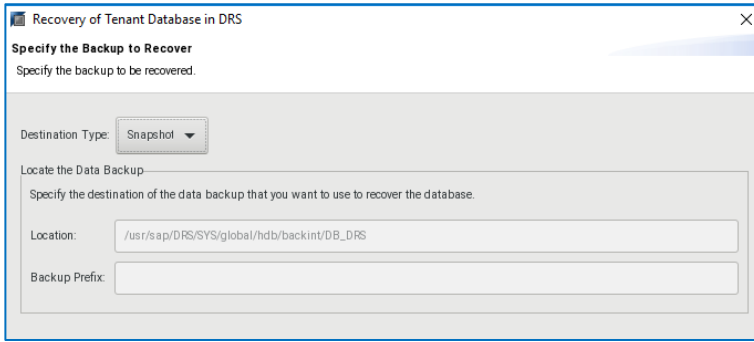
**Note:** For forward recovery using the replicated log backups, the backup catalog location must be specified.

Figure 47) Recovery with log backups: Specify backup catalog location.

The screenshot shows a dialog box titled "Recovery of Tenant Database in DRS" with a close button (X) in the top right corner. The main heading is "Locate Backup Catalog" and the instruction is "Specify location of the backup catalog." There are two radio button options: "Recover using the backup catalog" and "Recover without the backup catalog" (which is selected). Under "Recover using the backup catalog", there is a sub-option "Search for the backup catalog in the file system only" (which is selected) and a "Backup Catalog Location" text field containing "/mnt/log-backup/DB\_DRS". Under "Recover without the backup catalog", there is a sub-section "Backint System Copy" with a "Backint System Copy" checkbox (unchecked) and a "Source System" text field.

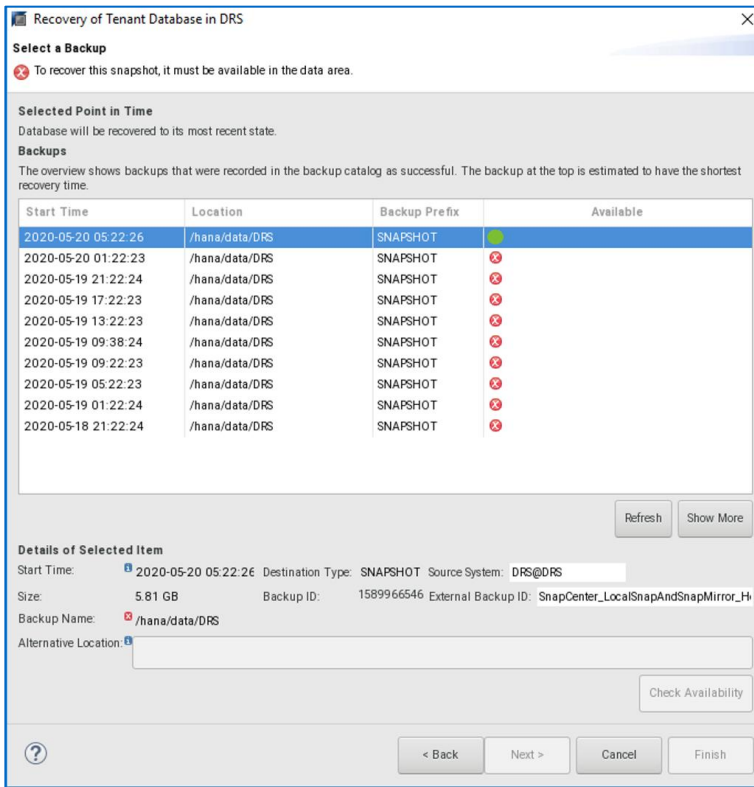


Figure 48) Recovery to a specific data backup: Select Destination Type.



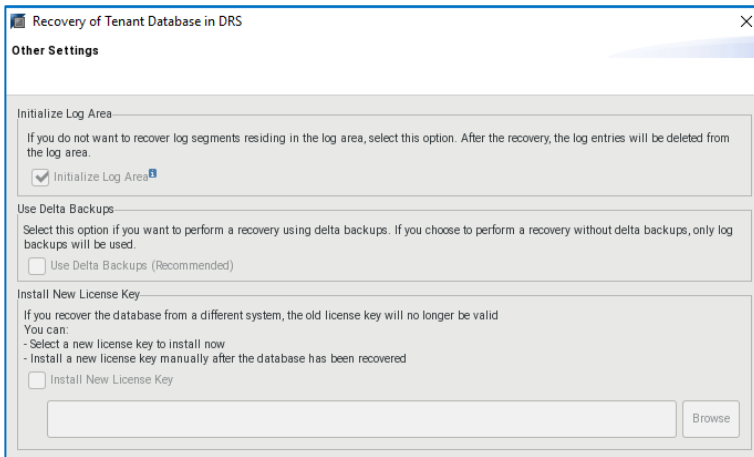
**Note:** For forward recovery using the replicated log backups, HANA Studio shows the availability of the Snapshot copy with a green icon.

Figure 49) Recovery with log backups: Select a Backup.



For recovery to a specific backup. Initialize Log Area is preselected.

Figure 50) Recovery to a specific data backup: Other Settings.



**Note:** For forward recovery using the replicated log backups, the Initialize Log Area option must be selected.

Figure 51) Recovery to a specific data backup: Review Recovery Settings.

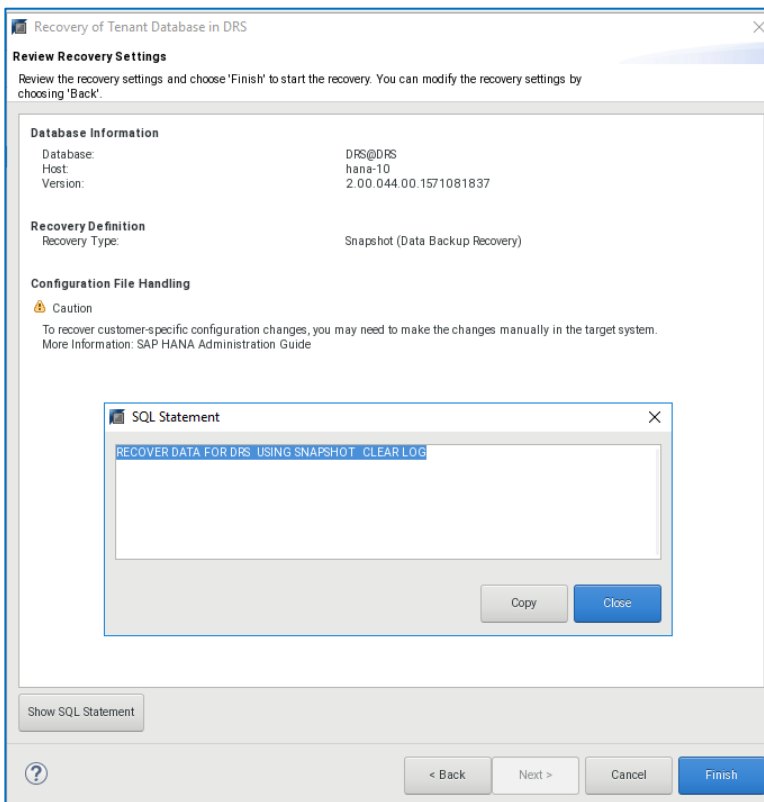


Figure 52) Recovery Execution Summary.

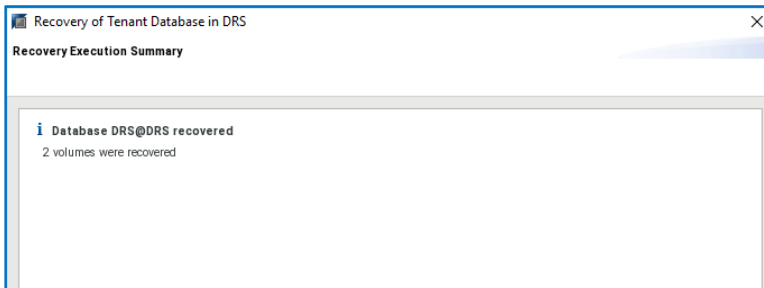
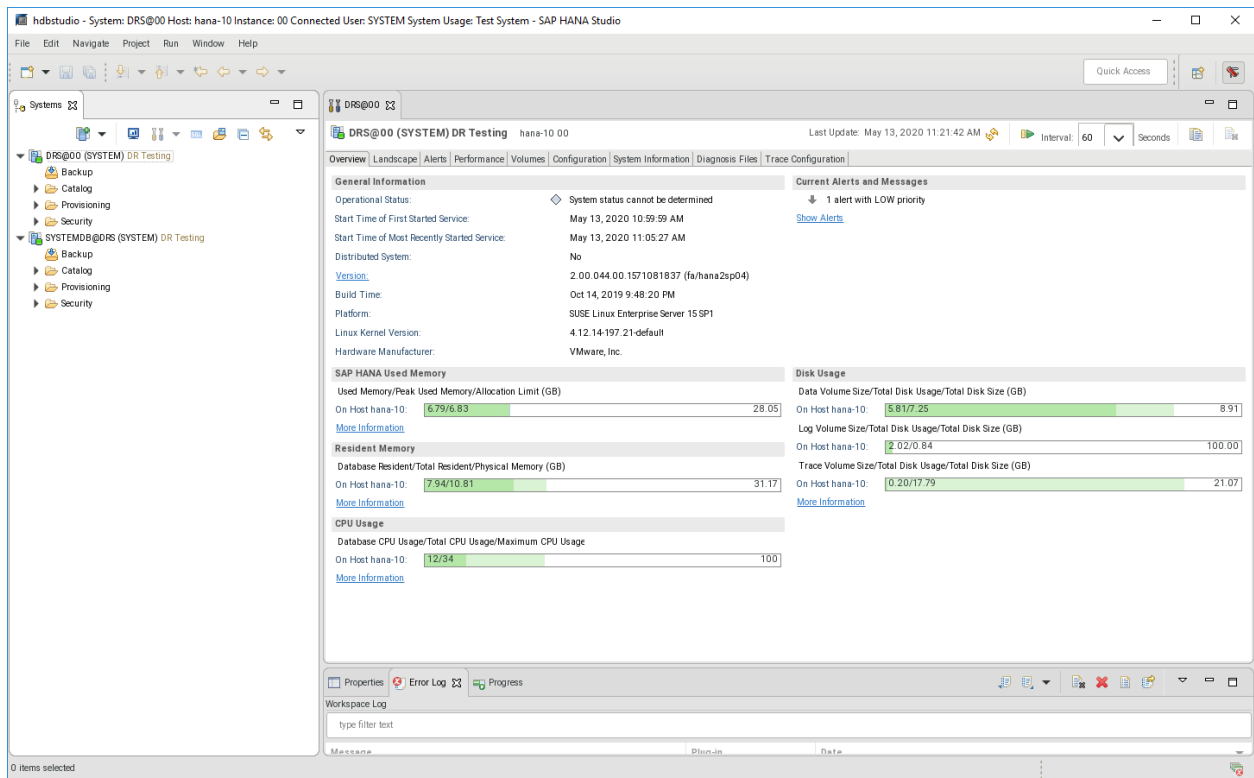


Figure 53) HANA Studio: System recovered.



## Recover Using the Command Line

Instead of using HANA Studio, recovery can also be done by using the command line, or it can be automated by using scripts.

### System Database Recovery

- Without log backups.

```
recoverSys.py --command "RECOVER DATA USING SNAPSHOT CLEAR LOG"
```

- With log backups.

```
recoverSys.py --command "RECOVER DATABASE UNTIL TIMESTAMP 'timestamp' CLEAR LOG USING CATALOG PATH ('/mnt/log-backup/SYSTEM-DB') USING LOG PATH ('/mnt/log-backup/SYSTEM-DB') USING SNAPSHOT"
```

### Tenant Database Recovery

- Without log backups.

```
Within hdbsql: RECOVER DATA FOR DRS USING SNAPSHOT CLEAR LOG
```

- With log backups.

```
Within hdbsql: RECOVER DATABASE UNTIL TIMESTAMP 'timestamp' CLEAR LOG USING CATALOG PATH ('/mnt/log-backup/DRS-DB') USING LOG PATH ('/mnt/log-backup/DRS-DB') USING SNAPSHOT
```

## 7 Overview of Disaster Recovery Failover

Most of the required steps for disaster failover are identical to those described in section 4, Overview of Disaster Recovery Testing.

Table 4 is a high-level overview of the required steps. The following sections describe the disaster recovery failover workflow in detail.

Table 7) Disaster recovery testing – required steps.

|   | Synchronous SnapMirror  | Asynchronous SnapMirror   |
|---|---|---|
| Prepare target host   | <ol style="list-style-type: none"> <li>1. Install SAP host agent.</li> <li>2. Configure user, ports, SAP services.</li> <li>3. Create mount points.</li> <li>4. Prepare <code>/etc/fstab</code>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Install SAP host agent.</li> <li>2. Configure user, ports, and SAP services.</li> <li>3. Create mount points.</li> <li>4. Mount a new, empty log volume and create subdirectories identical to the source system.</li> <li>5. Prepare <code>/etc/fstab</code>.</li> </ol> |
| Break the SnapMirror relation at the target storage   | ONTAP System Manager: <ol style="list-style-type: none"> <li>1. SnapMirror quiesce, SnapMirror break for data, log, and shared volume.</li> </ol>   | ONTAP System Manager: <ol style="list-style-type: none"> <li>1. SnapMirror quiesce, SnapMirror break for data, and shared volume.</li> </ol> Same for log backup volume (if log backup replication is part of the DR concept).  |
| Mount volumes at target host<br><b>Note:</b> This step requires additional LUN discovery operations in an FC SAN environment. | <ol style="list-style-type: none"> <li>1. Mount data, log, and shared volume.</li> </ol>  | <ol style="list-style-type: none"> <li>1. Mount data and shared volume.</li> <li>2. Mount log backup volume (if log backup replication is part of the DR concept).</li> </ol>   |
| Start or recover the HANA database  | <ol style="list-style-type: none"> <li>1. Start SAP services.</li> <li>2. Start HANA database. Crash recovery is executed.</li> </ol>   | <ol style="list-style-type: none"> <li>1. Start SAP services.</li> </ol> The HANA database is recovered to the last backup (or recovered with forward recovery using log backups, if log backup replication is part of the DR concept).   |

## 8 Synchronous SnapMirror Disaster Recovery Failover

To configure synchronous SnapMirror disaster recovery failover, follow these steps.

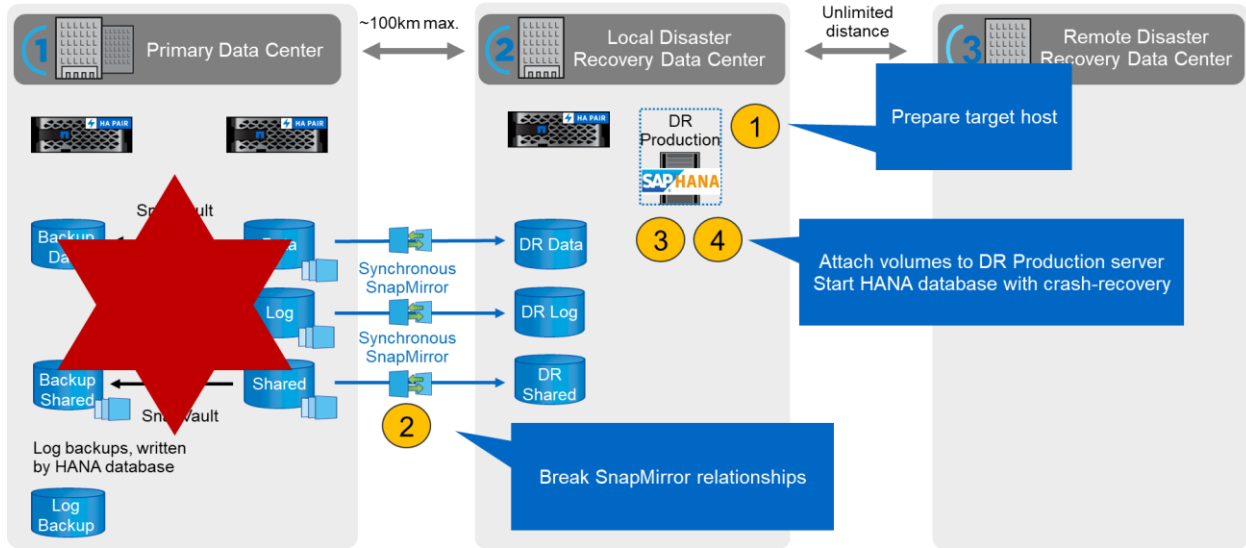
1. Prepare the target host.
2. Break SnapMirror relationships.

3. Mount volumes at the target host.
4. Start the HANA database.

The following sections describe these steps in detail.

**Note:** Manual steps are described for the different operations. All the steps could also be automated by using scripts or other automation tools.

Figure 54) Synchronous SnapMirror disaster recovery failover.



## 8.1 Prepare Target Host

This section describes the preparation steps required at the server, which is used for the disaster recovery failover.

During normal operation, the target host is typically used for other purposes—for example, as a HANA QA or test system. Therefore, most of the described steps must be executed when a disaster failover happens. On the other hand, the relevant configuration files, like `/etc/fstab` and `/usr/sap/sapservices`, can be prepared and then put in production simply by copying the configuration file. The disaster recovery testing procedure, as described in section 5, Synchronous SnapMirror Disaster Recovery Testing, ensures that the relevant prepared configuration files are configured correctly.

The target host preparation also includes shutting down the HANA QA or test system.

### Target Server Host Name and IP Address

The host name of the target server must be identical to the host name of the source system. The IP address can be different. As an alternative, a virtual IP address concept can be used as well.

### Install Required Software

The SAP host agent software must be installed at the target server. For full information, see the [SAP Host Agent](#) at the SAP help portal.

### Configure Users, Ports, and SAP Services

The required users and groups for the SAP HANA database must be available at the target server. Typically, central user management is used; therefore, no configuration steps are necessary at the target server. The required ports for the HANA database must be configured at the target hosts. The

configuration can be copied from the source system by copying the `/etc/services` file to the target server.

The required SAP services entries must be available at the target host. The configuration can be copied from the source system by copying the `/usr/sap/sapservices` file to the target server. The following output shows the required entries for the SAP HANA database used in the lab setup.

```
hana-10:~ # cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/DRS/HDB00/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/DRS/HDB00/exe/sapstartsrv pf=/usr/sap/DRS/SYS/profile/DRS_HDB00_hana-10
-D -u drsadm
limit.descriptors=1048576
```

### Prepare Log Backup Volume

Because the source system is configured with a separate volume for the HANA log backups, the log backup volume must also be available at the target host.

A volume for the log backups must be configured and mounted at the target host.

### Prepare File System Mounts

Table 8 shows the naming conventions that were used in the lab setup. The volume names of the FlexClone volumes at the disaster recovery storage were included in `/etc/fstab`. These volume names were used in the FlexClone volume creation step in the next section.

Table 8) Volume names of FlexClone volumes.

| HANA DRS Volumes | Volume at Disaster Recovery Storage                         | Mount Point at Target Host    |
|------------------|---|-------------------------------|
| Data volume      | DRS_data_mnt00001_dest_sm_s                                 | /hana/data/DRS/mnt00001       |
| Log volume       | DRS_log_mnt00001_dest_sm_s                                  | /hana/log/DRS/mnt00001        |
| Shared volume    | DRS_hana_shared_sm_s/shared<br>DRS_hana_shared_sm_s/usr-sap | /hana/shared<br>/usr/sap//DRS |

The mount points in Table 8 must be created at the target host.

The required `/etc/fstab` entries are the following.

```
hana-10:~ # cat /etc/fstab
192.168.175.116:/DRS_log_backup /mnt/log-backup nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_data_mnt00001_dest_sm_s /hana/data/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_log_mnt00001_dest_sm_s /hana/log/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_shared_dest_sm_s/shared /hana/shared nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_shared_dest_sm_s/usr-sap /usr/sap/DRS nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
```

**Note:** In a Fibre Channel setup, the `fstab` entries depend on the multipath configuration. Because the SCSI IDs are different with each operation, the `fstab` file is not static and needs to be adapted after the LUN discovery process.

## 8.2 Break SnapMirror Relationships

The synchronous SnapMirror relationships for the data, the log, and the shared volume must be quiesced and broken off. In addition, the volumes must be mounted to the namespace.

Figure 55 through Figure 61 show the required steps using ONTAP System Manager for the quiesce and break operations.

Figure 55) SnapMirror quiesce operation—step 1.

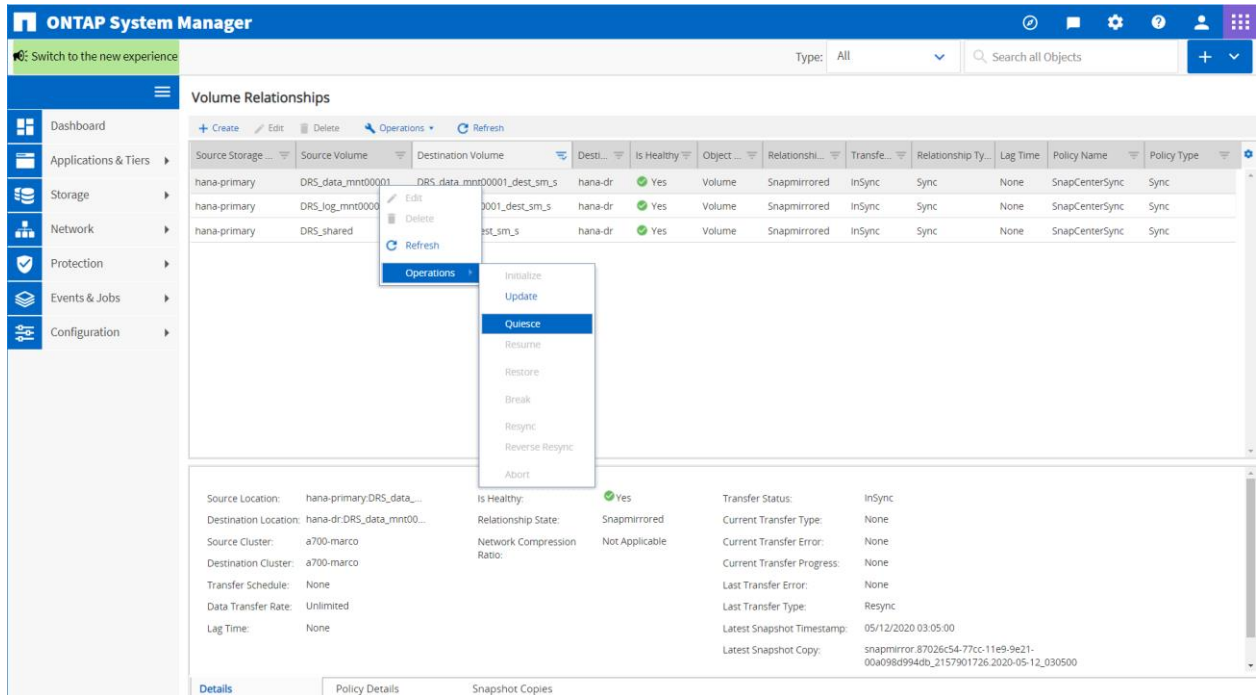


Figure 56) SnapMirror quiesce operation—step 2.

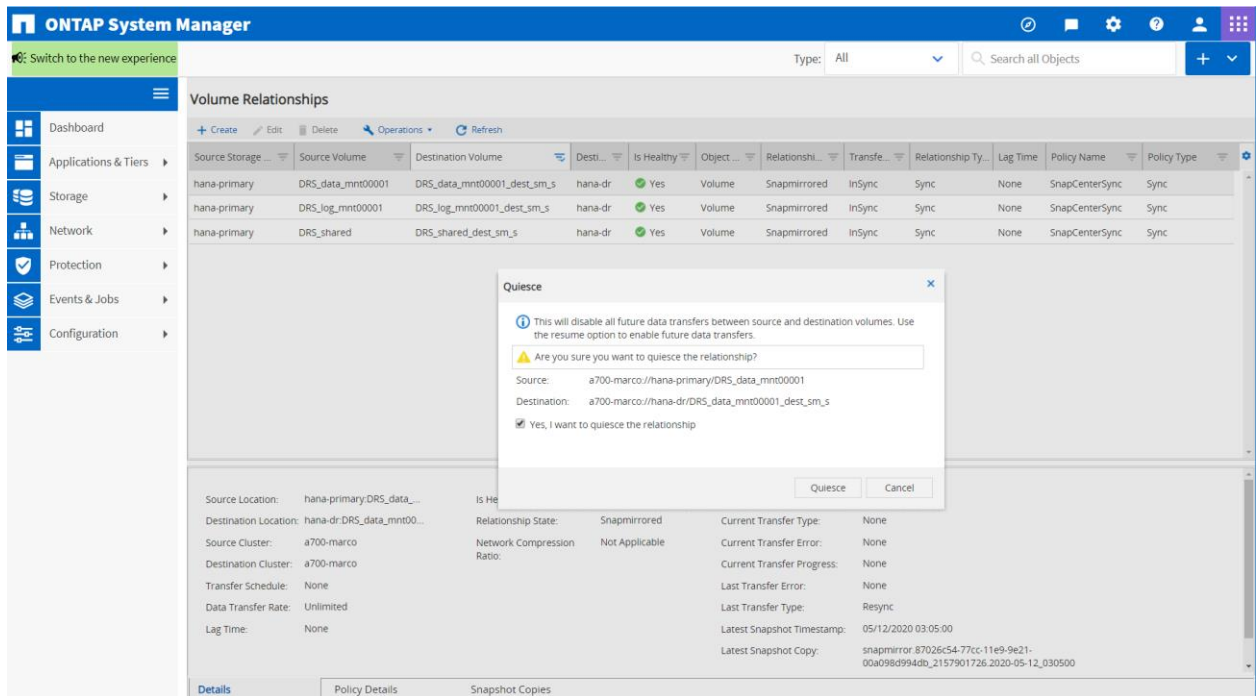


Figure 57) All SnapMirror target volumes are quiesced.

The screenshot shows the ONTAP System Manager interface. The 'Volume Relationships' section contains a table with the following data:

| Source Storage | Source Volume     | Destination Volume          | Desti... | Is Healthy | Object ... | Relations... | Transfe... | Relationship Ty... | Lag Time | Policy Name    | Policy Type |
|----------------|-------------------|-----------------------------|----------|------------|------------|--------------|------------|--------------------|----------|----------------|-------------|
| hana-primary   | DRS_data_mnt00001 | DRS_data_mnt00001_dest_sm_s | hana-dr  | No         | Volume     | Snapmirrored | Quiesced   | Sync               | 1 min(s) | SnapCenterSync | Sync        |
| hana-primary   | DRS_log_mnt00001  | DRS_log_mnt00001_dest_sm_s  | hana-dr  | No         | Volume     | Snapmirrored | Quiesced   | Sync               | 1 min(s) | SnapCenterSync | Sync        |
| hana-primary   | DRS_shared        | DRS_shared_dest_sm_s        | hana-dr  | No         | Volume     | Snapmirrored | Quiesced   | Sync               | 1 min(s) | SnapCenterSync | Sync        |

Below the table, the 'Details' section for a selected relationship shows:

- Source Location: hana-primary:DRS\_shared
- Destination Location: hana-dr:DRS\_shared\_des...
- Source Cluster: a700-marco
- Destination Cluster: a700-marco
- Transfer Schedule: None
- Data Transfer Rate: Unlimited
- Lag Time: 1 min(s)
- Is Healthy: No
- Relationship State: Snapmirrored
- Network Compression Ratio: Not Applicable
- Transfer Status: Quiesced
- Current Transfer Type: None
- Current Transfer Error: None
- Current Transfer Progress: None
- Last Transfer Error: None
- Last Transfer Type: Resync
- Latest Snapshot Timestamp: 05/12/2020 03:38:13
- Latest Snapshot Copy: snapmirror\_8702654-77cc-11e9-9e21-00a098d994db\_2157901727.2020-05-12\_033813

Figure 58) SnapMirror break operation—step 1.

The screenshot shows the ONTAP System Manager interface with the 'Operations' menu open over the 'Volume Relationships' table. The 'Break' option is selected. The table data is as follows:

| Source Storage | Source Volume     | Destination Volume          | Desti... | Is Healthy | Object ... | Relations... | Transfe... | Relationship Ty... | Lag Time | Policy Name    | Policy Type |
|----------------|-------------------|-----------------------------|----------|------------|------------|--------------|------------|--------------------|----------|----------------|-------------|
| hana-primary   | DRS_data_mnt00001 | DRS_data_mnt00001_dest_sm_s | hana-dr  | No         | Volume     | Snapmirrored | Quiesced   | Sync               | 1 min(s) | SnapCenterSync | Sync        |
| hana-primary   | DRS_log_mnt00001  | DRS_log_mnt00001_dest_sm_s  | hana-dr  | No         | Volume     | Snapmirrored | Quiesced   | Sync               | 1 min(s) | SnapCenterSync | Sync        |
| hana-primary   | DRS_shared        | DRS_shared_dest_sm_s        | hana-dr  | No         | Volume     | Snapmirrored | Quiesced   | Sync               | 1 min(s) | SnapCenterSync | Sync        |

The 'Operations' menu includes the following options: Initialize, Update, Quiesce, Resume, Restore, Break, Resync, Reverse Resync, and Abort.

Below the table, the 'Details' section for a selected relationship shows:

- Source Location: hana-primary:DRS\_data...
- Destination Location: hana-dr:DRS\_data\_mnt00...
- Source Cluster: a700-marco
- Destination Cluster: a700-marco
- Transfer Schedule: None
- Data Transfer Rate: Unlimited
- Lag Time: 1 min(s)
- Is Healthy: No
- Relationship State: Snapmirrored
- Network Compression Ratio: Not Applicable
- Transfer Status: Quiesced
- Current Transfer Type: None
- Current Transfer Error: None
- Current Transfer Progress: None
- Last Transfer Error: None
- Last Transfer Type: Resync
- Latest Snapshot Timestamp: 05/12/2020 03:37:47
- Latest Snapshot Copy: snapmirror\_8702654-77cc-11e9-9e21-00a098d994db\_2157901726.2020-05-12\_033747



Figure 59) SnapMirror break operation—step 2.

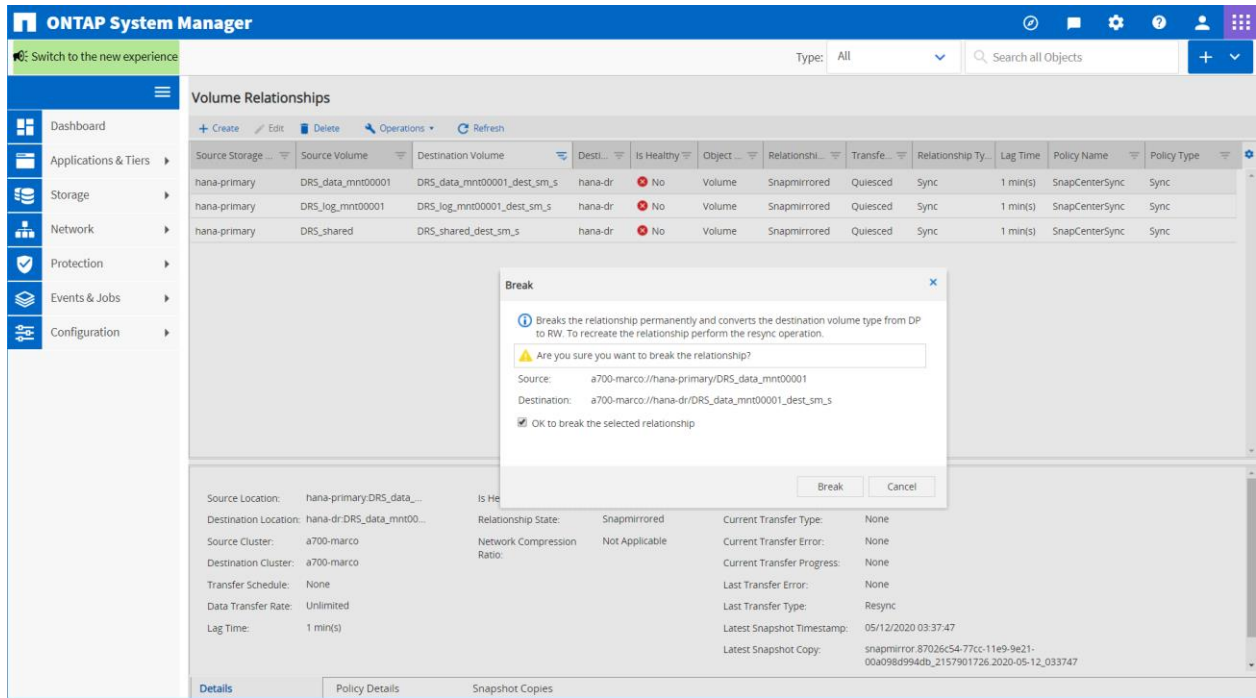


Figure 60) All SnapMirror target volumes are broken off.

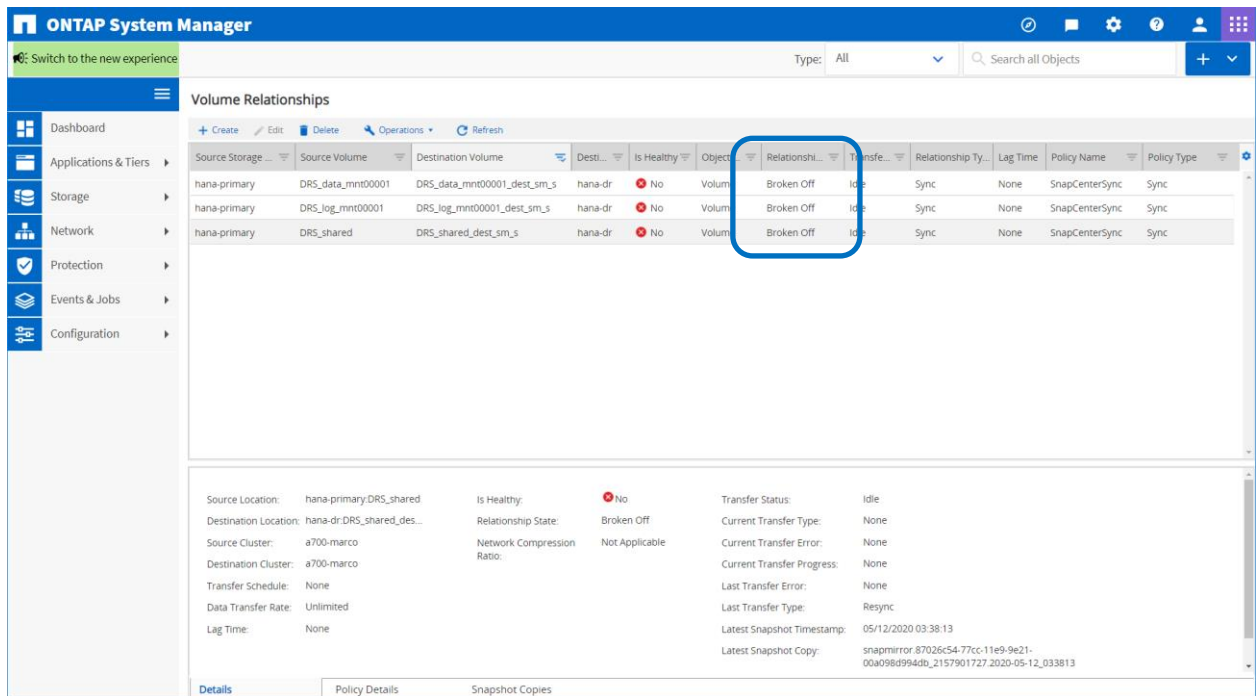
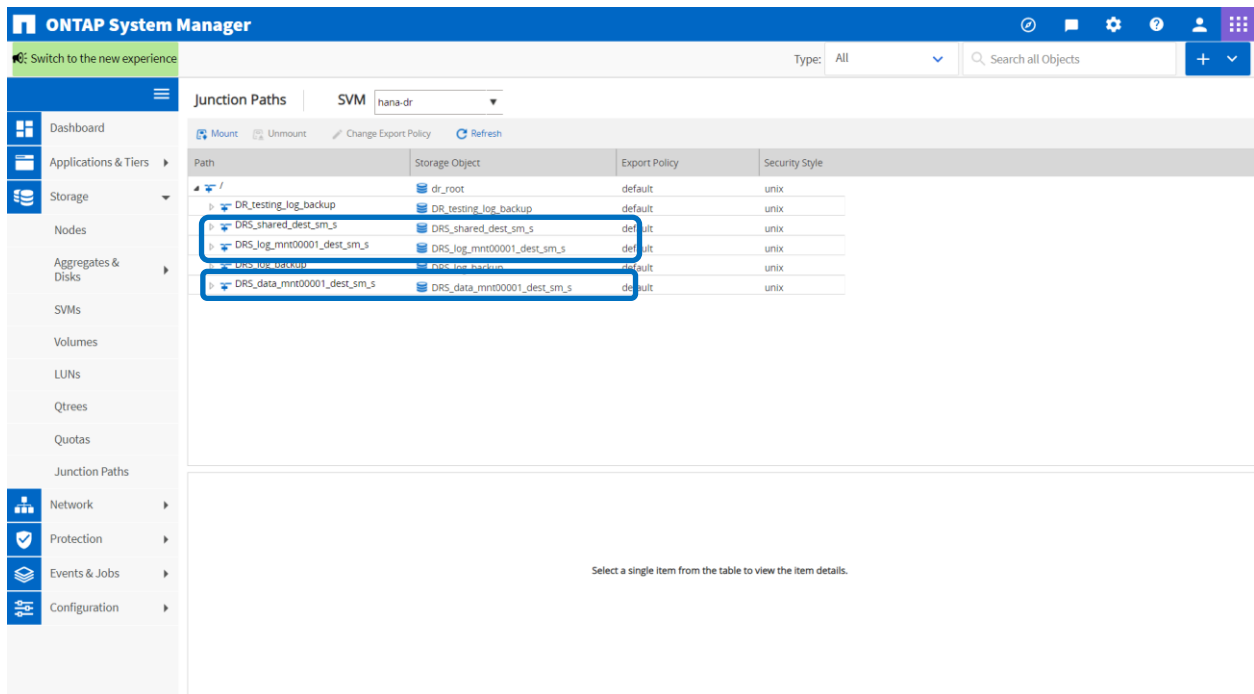


Figure 61 shows the required junction path configuration.

Figure 61) Junction path configuration.



**Note:** In a Fibre Channel setup, the LUNs in the FlexClone volumes must be mapped to the initiator group of the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

### 8.3 Mount Volumes at Target Host

The volumes can now be mounted at the target host.

```
hana-10:~ # mount -a
```

The following output shows the required file systems.

```
hana-10:~ # df
Filesystem                1K-blocks      Used  Available Use% Mounted on
192.168.175.116:/DRS_log_backup
192.168.175.116:/DRS_data_mnt00001_dest_sm_s
192.168.175.116:/DRS_log_mnt00001_dest_sm_s
192.168.175.116:/DRS_shared_dest_sm_s/shared
192.168.175.116:/DRS_shared_dest_sm_s/usr-sap
104857600          256    104857344    1% /mnt/log-backup
10725184          8694080    2031104    82% /hana/data/DRS/mnt00001
67106816          54338816    12768000    81% /hana/log/DRS/mnt00001
324150976          269775616    54375360    84% /hana/shared
324150976          269775616    54375360    84% /usr/sap/DRS
hana-10:~ #
```

**Note:** In a Fibre Channel setup, additional steps are required before the LUNs can be mounted at the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

### 8.4 Start the HANA Database

The steps to start the HANA database are described in section 5.5, Start the HANA Database.

## 9 Asynchronous SnapMirror Disaster Recovery Failover

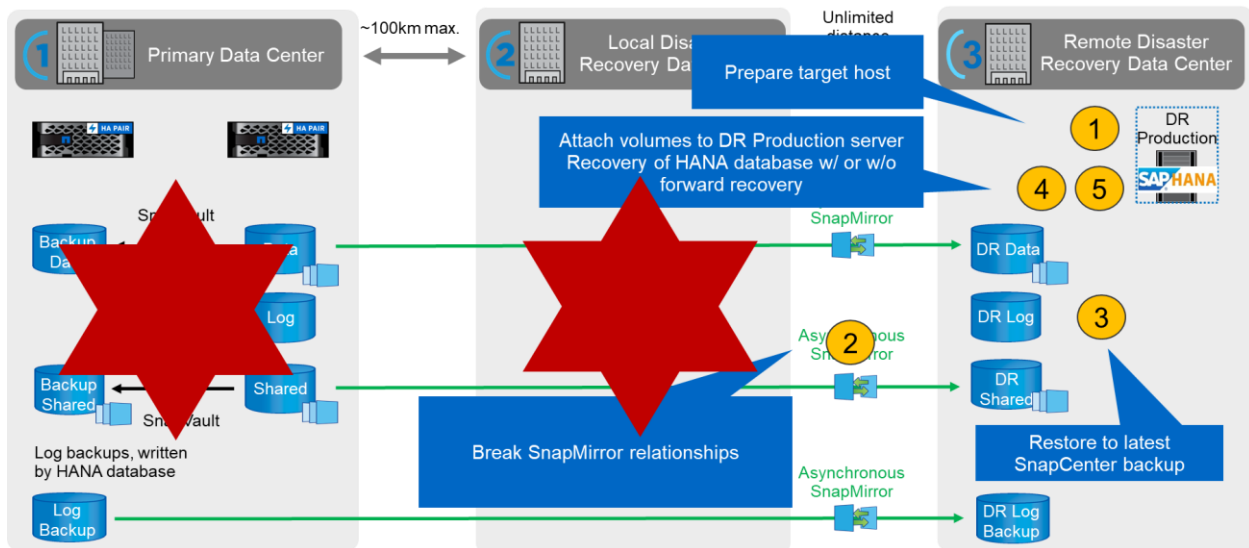
Follow these steps to execute the disaster recovery failover:

1. Prepare the target host.
2. Break SnapMirror relationships.
3. Restore to the latest SnapCenter Snapshot backup.
4. Mount volumes at the target host.
5. Recover the HANA database.

The following chapters describe these steps in detail.

**Note:** Manual steps are described for the different operations. All the steps could also be automated by using scripts or other automation tools.

Figure 62) Synchronous SnapMirror disaster recovery failover.



## 9.1 Prepare Target Host

This section describes the preparation steps required at the server, which is used for the disaster recovery failover testing.

During normal operation, the target host is typically used for other purposes—for example, as a HANA QA or test system. Therefore, most of the described steps must be executed when a disaster failover testing is executed. On the other hand, the relevant configuration files, such as `/etc/fstab` and `/usr/sap/sapservices`, can be prepared and then put in production simply by copying the configuration file. The disaster recovery testing procedure ensures that the relevant prepared configuration files are configured correctly.

The target host preparation also includes shutting down the HANA QA or test system.

### Target Server Host Name and IP Address

The host name of the target server must be identical to the host name of the source system. The IP address can be different. As an alternative a virtual IP address concept can be used as well.

### Install Required Software

The SAP host agent software must be installed at the target server. For full information, see the [SAP Host Agent](#) at the SAP help portal.

**Note:** If the host is used as a HANA QA or test system, the SAP host agent software is already installed.

## Configure Users, Ports, and SAP Services

The required users and groups for the SAP HANA database must be available at the target server. Typically, central user management is used; therefore, no configuration steps are necessary at the target server. The required ports for the HANA database must be configured at the target hosts. The configuration can be copied from the source system by copying the `/etc/services` file to the target server.

The required SAP services entries must be available at the target host. The configuration can be copied from the source system by copying the `/usr/sap/sapservices` file to the target server. The following output shows the required entries for the SAP HANA database used in the lab setup.

```
hana-10:~ # cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/DRS/HDB00/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/DRS/HDB00/exe/sapstartsrv pf=/usr/sap/DRS/SYS/profile/DRS_HDB00_hana-10
-D -u drsadm
limit.descriptors=1048576
```

## Prepare the HANA Log Volume

Because the HANA log volume is not part of the SnapMirror replication, an empty log volume must exist at the target host. The log volume must include the same subdirectories as the source HANA system.

```
hana-10:/hana/log/DRS/mnt00001 # ls -al
total 84
drwxr-x--- 5 drsadm sapsys 4096 Apr  2 06:51 .
drwxr-x--- 1 drsadm sapsys  16 Apr  2 06:45 ..
drwxr-x--- 2 drsadm sapsys 61440 Apr 24 07:11 hdb00001
drwxr-xr-- 2 drsadm sapsys 12288 Apr 24 02:59 hdb00002.00003
drwxr-xr-- 2 drsadm sapsys  4096 Apr 24 08:31 hdb00003.00003
hana-10:/hana/log/DRS/mnt00001 #
```

## Prepare Log Backup Volume

Because the source system is configured with a separate volume for the HANA log backups, a log backup volume must also be available at the target host. A volume for the log backups must be configured and mounted at the target host.

**Note:** If log backup volume replication is part of the disaster recovery setup, a FlexClone volume will be mounted at the target host and there is no need to prepare an additional log backup volume.

## Prepare File System Mounts

Table 9 shows the naming conventions that were used in the lab setup. The volume names of the FlexClone volumes at the disaster recovery storage were included in `/etc/fstab`. These volume names were used in the FlexClone copy creation step in the next section.

**Table 9) Volume names of FlexClone volumes.**

| HANA DRS Volumes  | Volume at Disaster Recovery Storage                         | Mount Point at Target Host    |
|-------------------|---|-------------------------------|
| Data volume       | DRS_data_mnt00001_dest                                      | /hana/data/DRS/mnt00001       |
| Shared volume     | DRS_hana_shared_dest/shared<br>DRS_hana_shared_dest/usr-sap | /hana/shared<br>/usr/sap//DRS |
| Log backup volume | DRS_log_backup_dest   | /mnt/log-backup               |

**Note:** A FlexClone of the log backup volume is required only if log backups are also replicated to the DR site.

**Note:** The mount points in Table 9 must be created at the target host.

The required `/etc/fstab` entries are the following.

```
hana-10:/etc # cat /etc/fstab

192.168.175.116:/DRS_log_mnt00001 /hana/log/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_data_mnt00001_dest /hana/data/DRS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_shared_dest/shared /hana/shared nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_shared_dest/usr-sap /usr/sap/DRS nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
192.168.175.116:/DRS_log_backup_dest /mnt/log-backup nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0

hana-10:/etc #hana-10:~ #
```

**Note:** In a Fibre Channel setup, the `fstab` entries depend on the multipath configuration. Because the SCSI IDs are different with each operation, the `fstab` file will not be static and must be adapted after the LUN discovery process.

## 9.2 Break SnapMirror Relationships

The asynchronous SnapMirror relationships for the data and the shared volume must be quiesced and broken off. If log backup replication is part of the disaster recovery setup, the replication relationship of the log backup volume must be quiesced and broken off as well.

In addition, the volumes must be mounted to the namespace.

Figure 63 through Figure 69 show the steps using ONTAP System Manager for the quiesce and break operations.

Figure 63) SnapMirror quiesce operation—step 1.

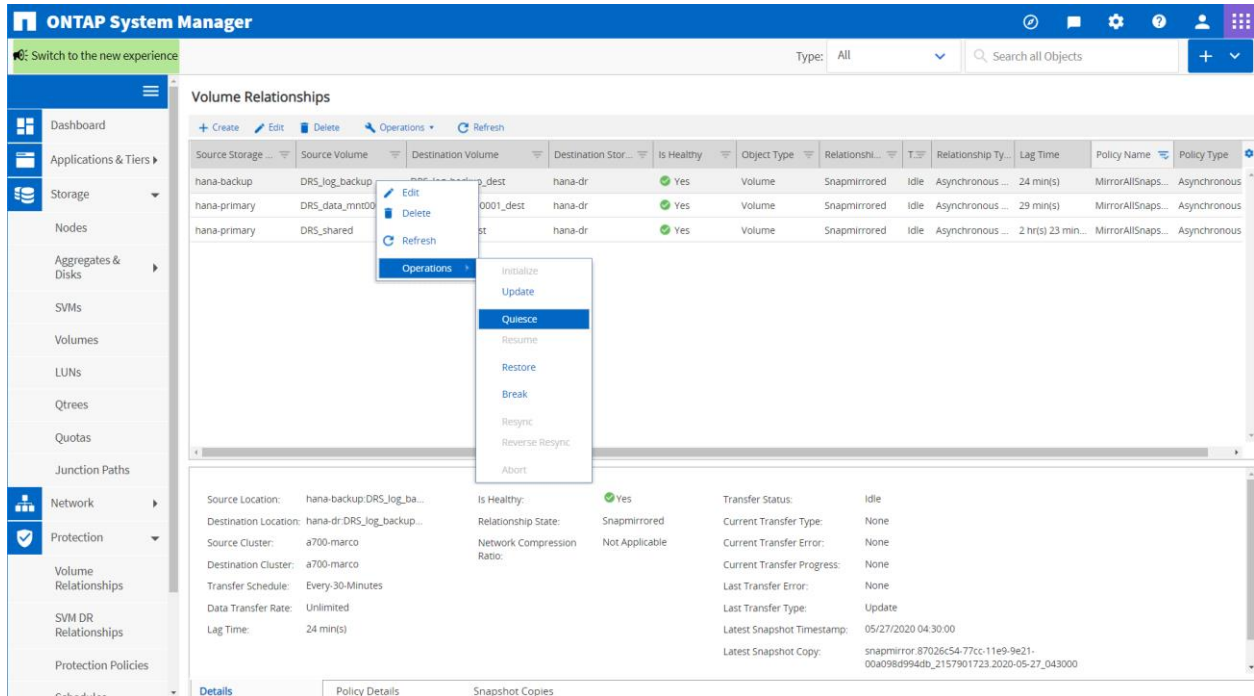
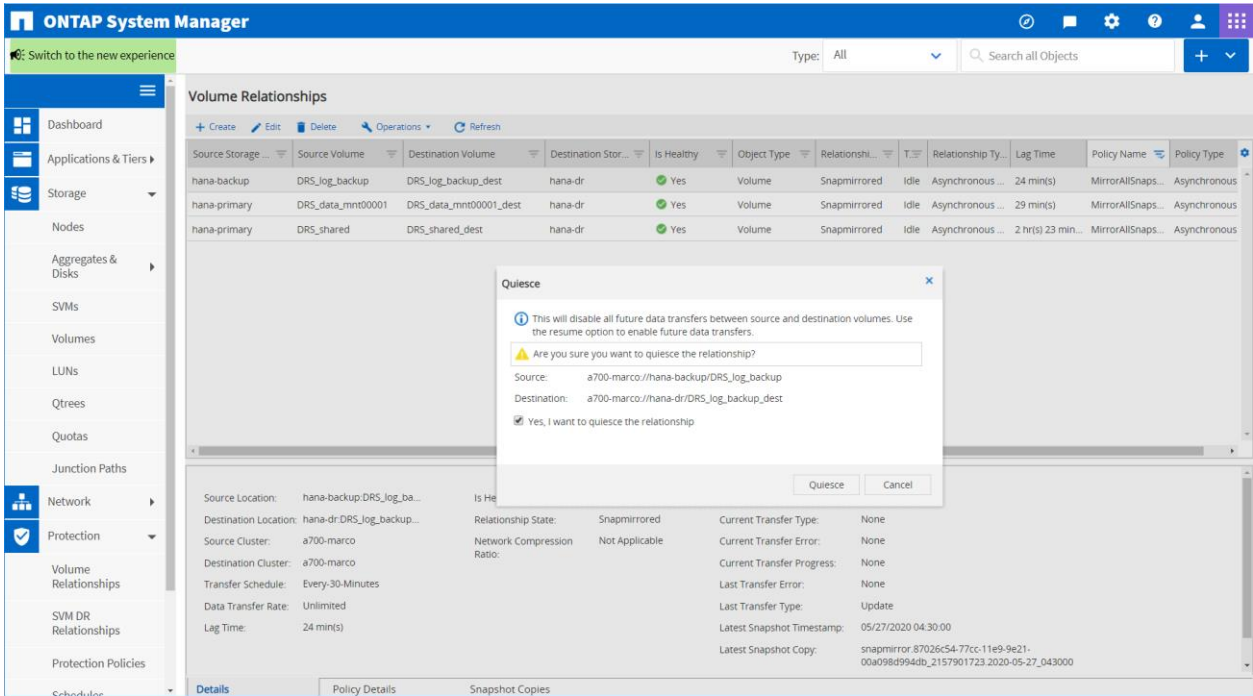


Figure 64) SnapMirror quiesce operation—step 2.



As a source for the log backup FlexClone volume, one of the SnapMirror Snapshot copies is selected.

Figure 65) All SnapMirror target volumes are quiesced.

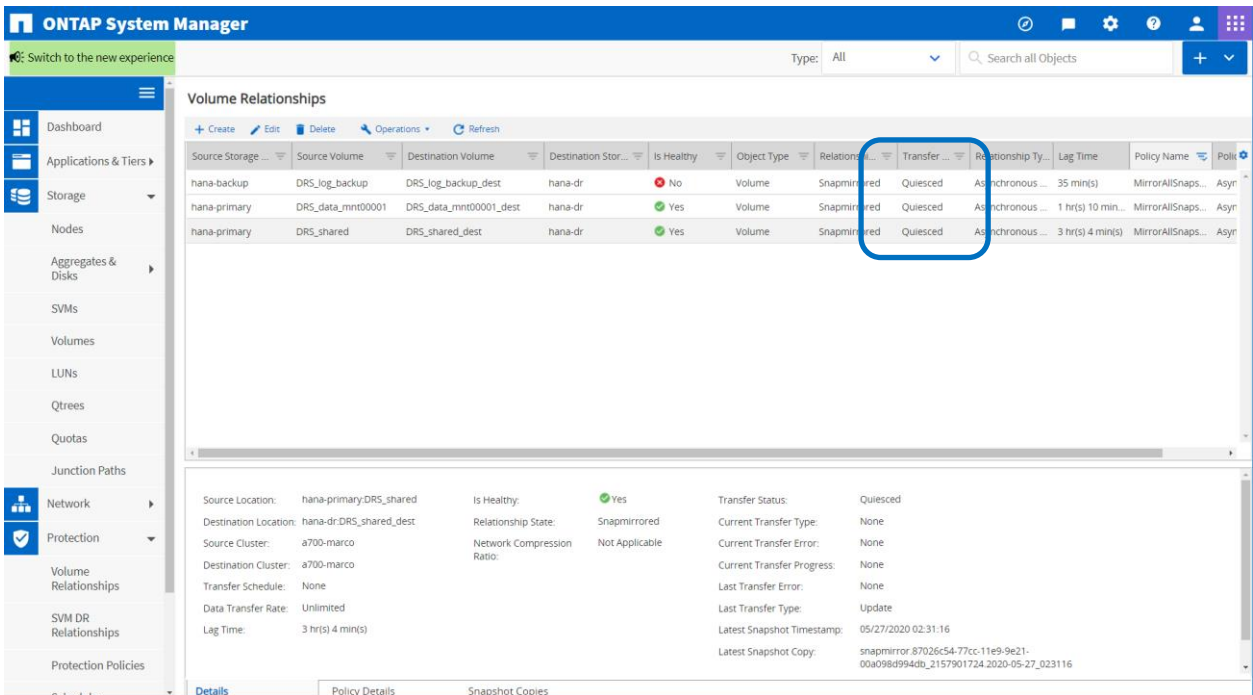




Figure 66) SnapMirror break operation—step 1.

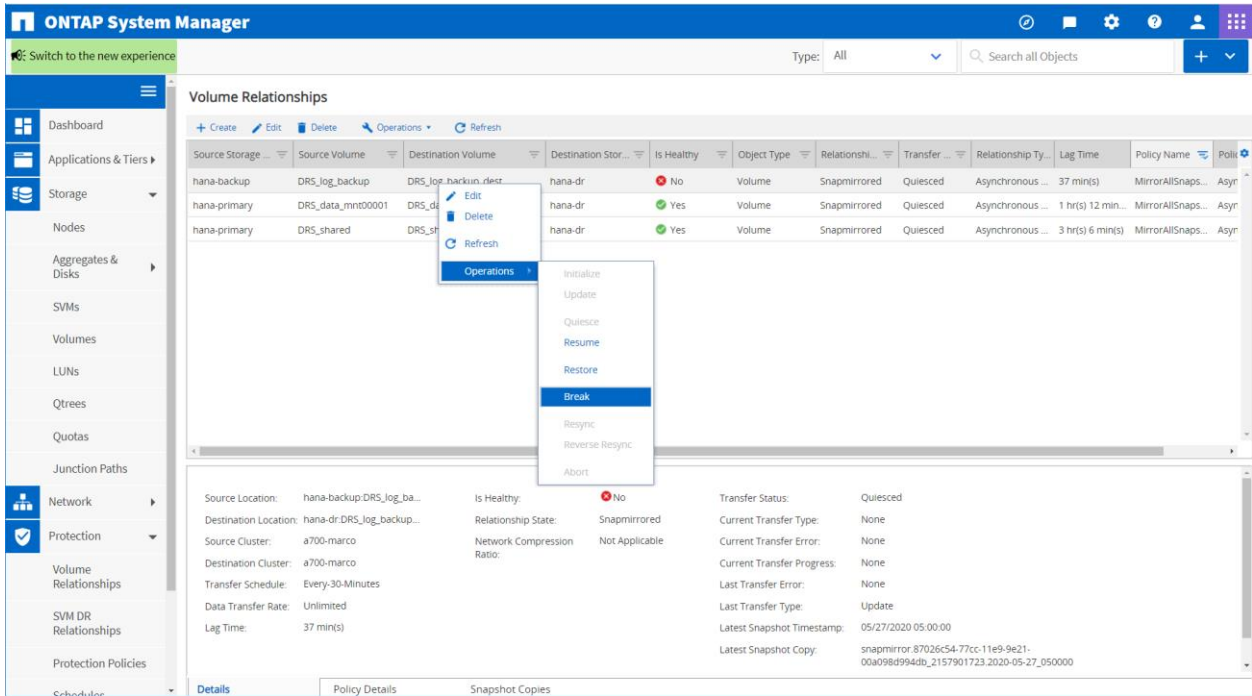


Figure 67) SnapMirror break operation—step 2.

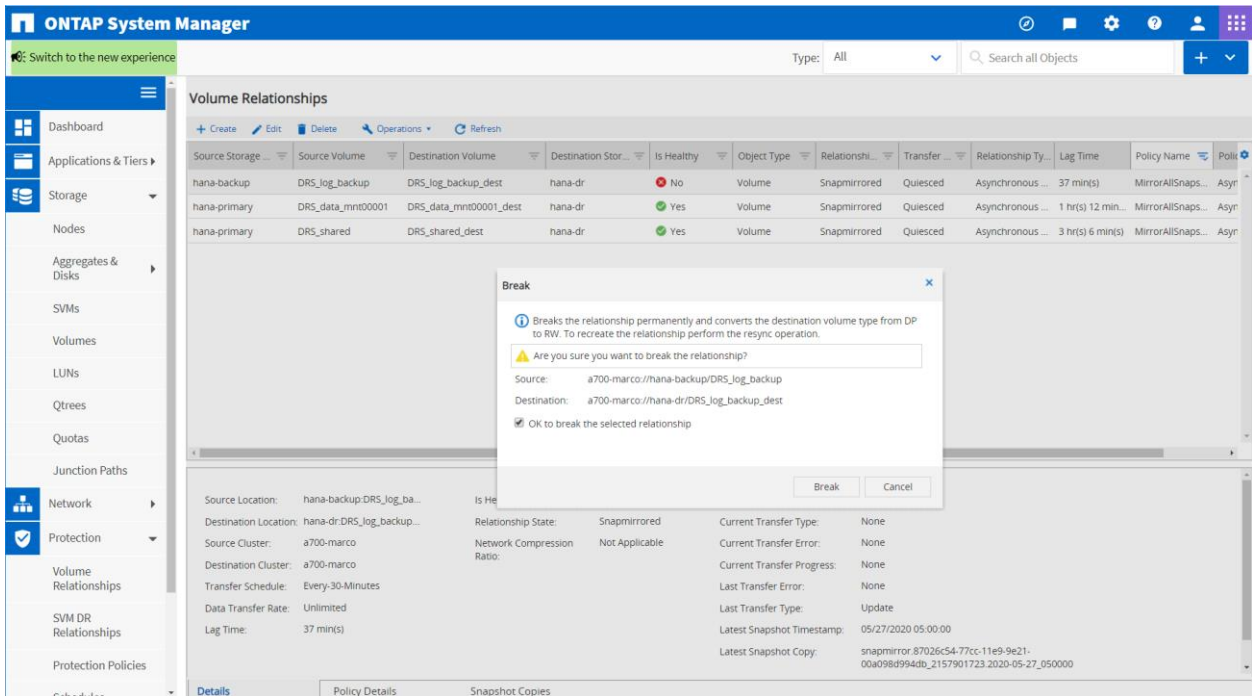
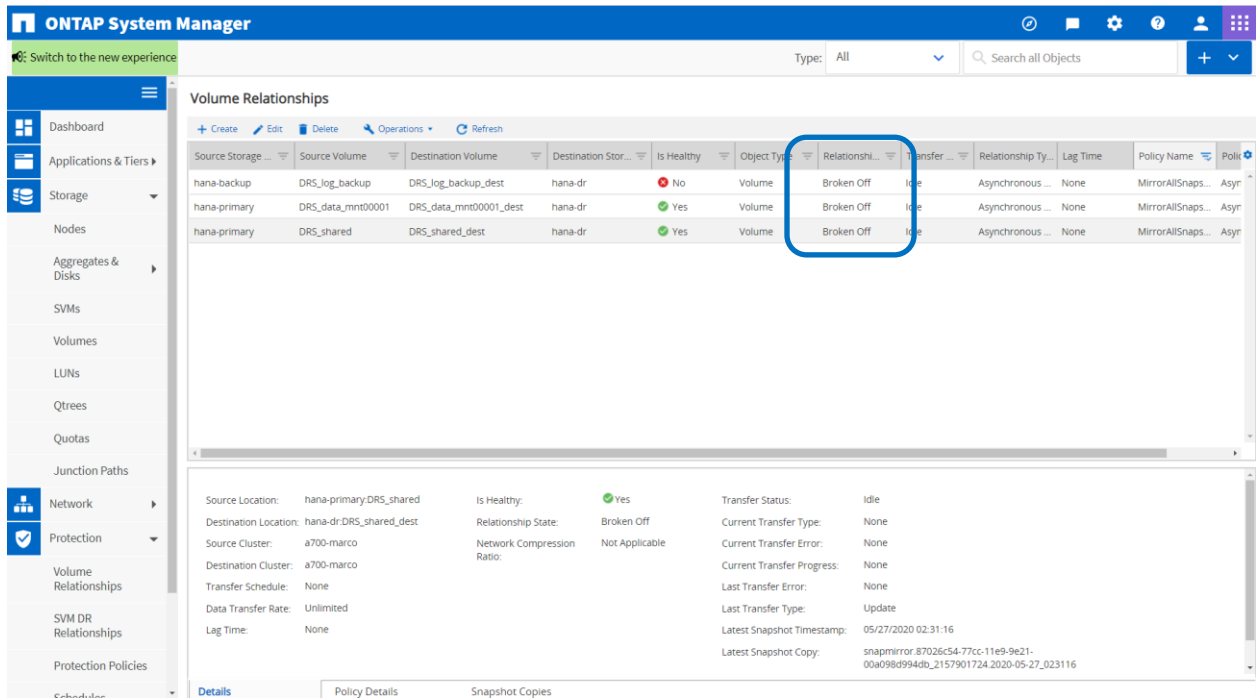
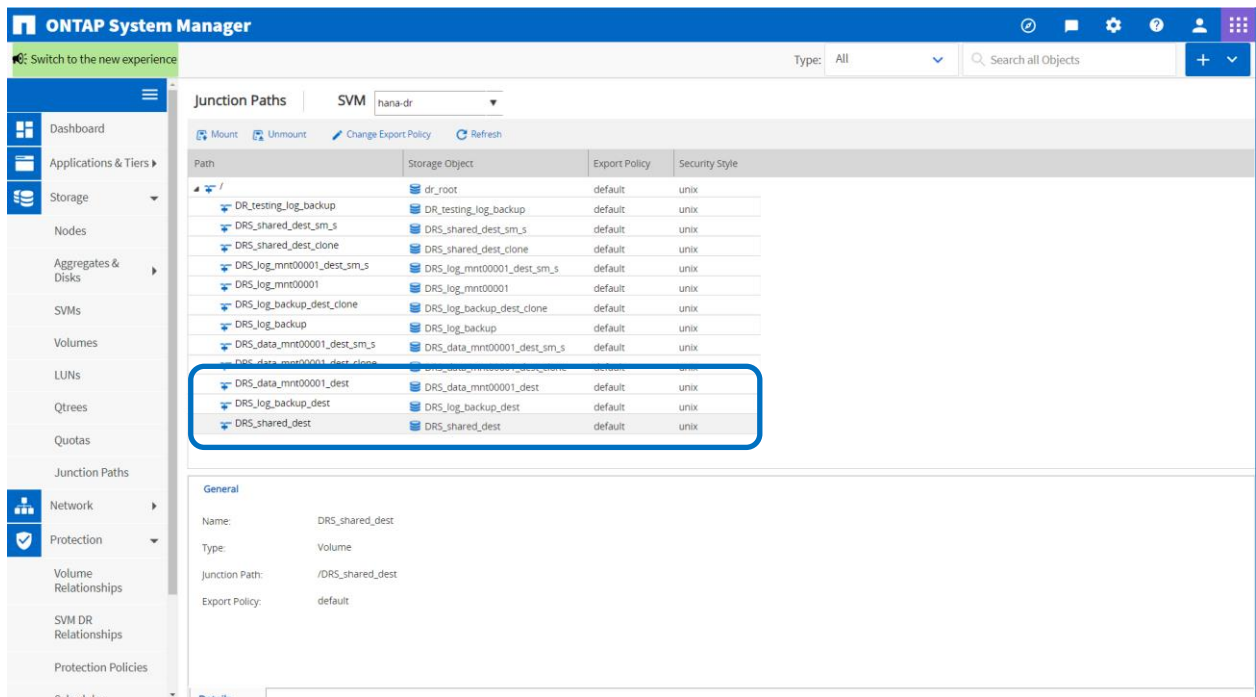


Figure 68) All SnapMirror target volumes are broken off.



All volumes must be mounted to the namespace.

Figure 69) Junction path configuration.



**Note:** In a Fibre Channel setup, the LUNs in the FlexClone volumes must be mapped to the initiator group of the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

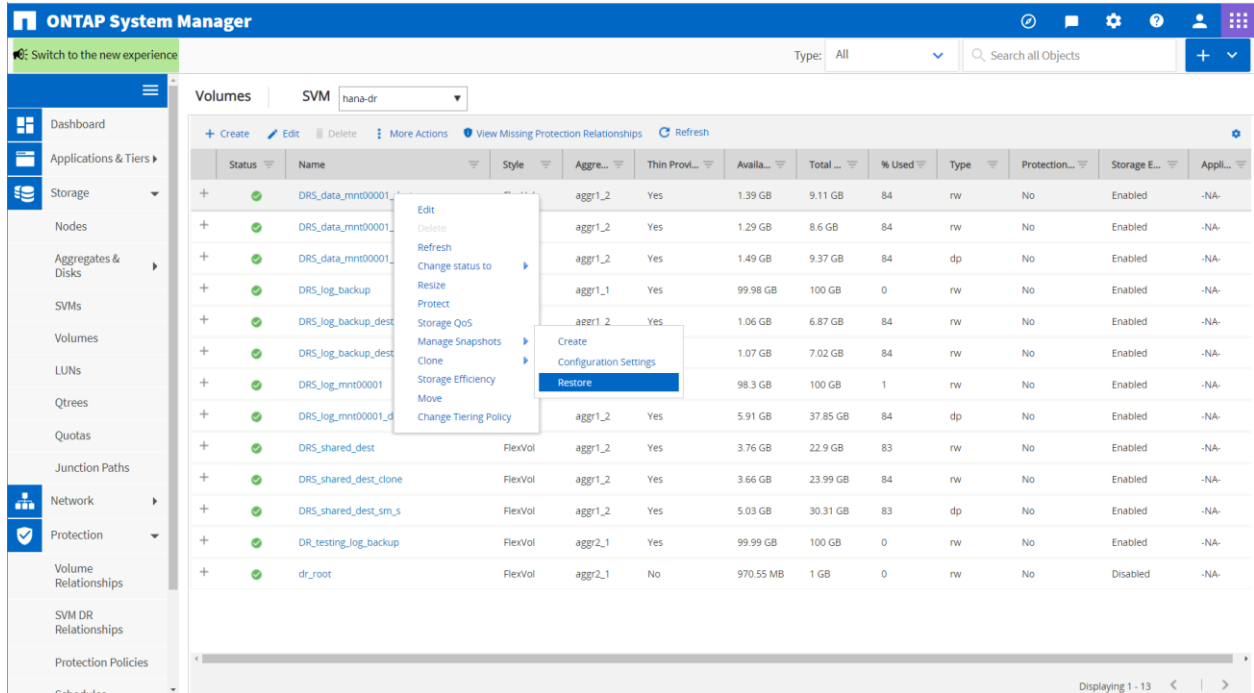


### 9.3 Restore Data Volume to Latest SnapCenter Backup

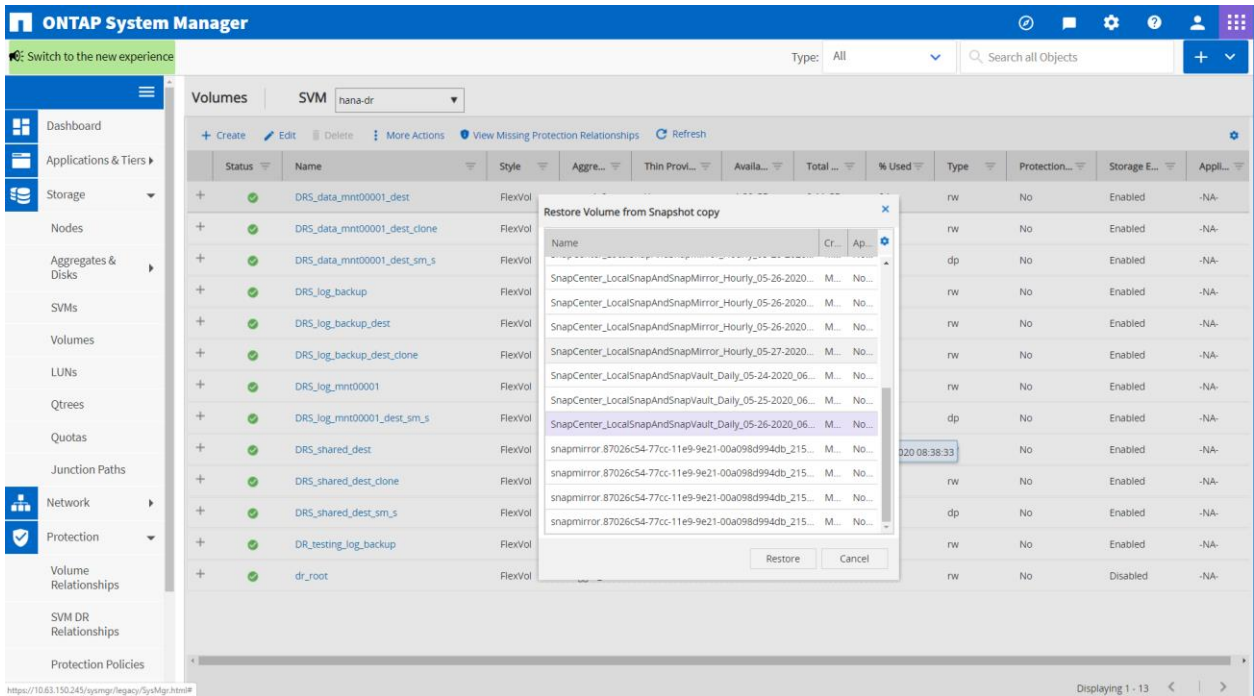
The current active file system in the database data volume is not consistent from the SAP HANA database point of view because it is based on a SnapMirror Snapshot copy. This Snapshot copy was created after SnapCenter issued the `unquiesce` command for the SAP HANA database. Therefore, the SAP HANA database data volume must be restored to the latest backup that was created with SnapCenter to get a consistent database image.

To restore the SAP HANA database data volume, follow these steps.

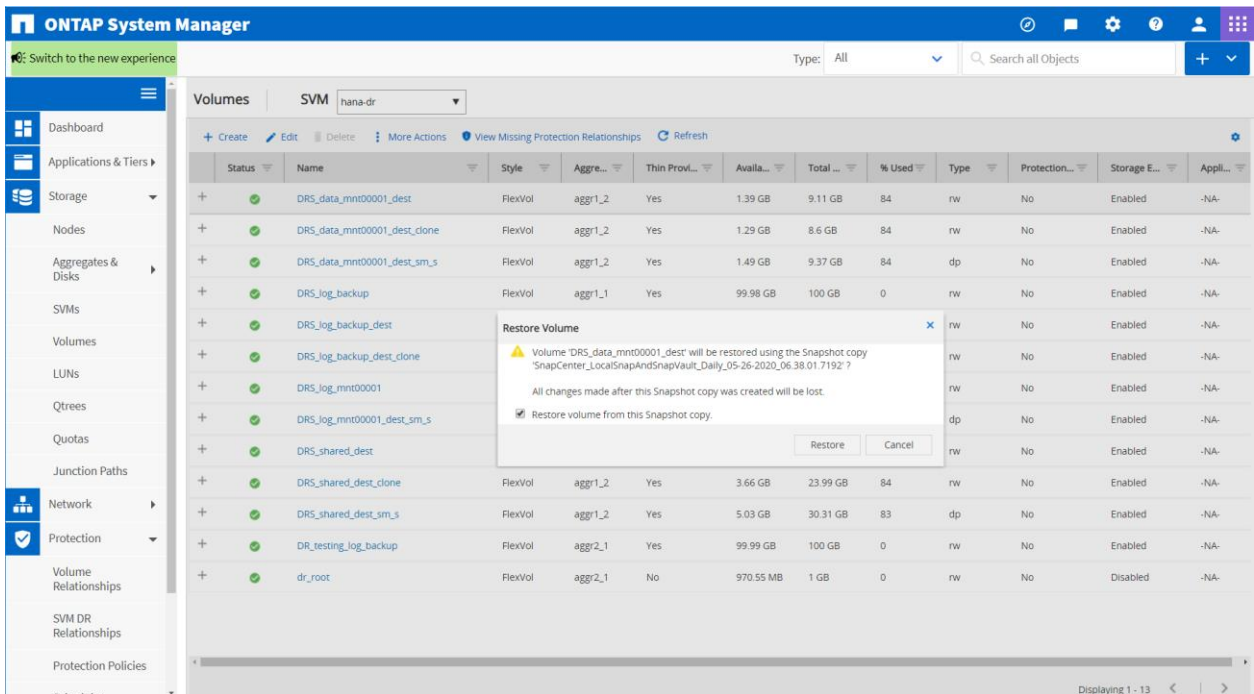
1. In the disaster recovery SVM, select Volume. Select the data volume and then select Manage Snapshots and Restore.



2. Select a SnapCenter Snapshot backup.



### 3. Execute the restore operation.



## 9.4 Mount FlexClone Volumes at Target Host

The FlexClone volumes can now be mounted at the target host.

```
hana-10:~ # mount -a
```

The output shows the required file systems.

```

hana-10:/etc # df
Filesystem                                1K-blocks      Used  Available Use% Mounted on
192.168.175.116:/DRS_log_mnt00001        104857600     1778240  103079360   2% /hana/log/DRS/mnt00001
192.168.175.116:/DRS_data_mnt00001_dest   9553472      8096192   1457280    85% /hana/data/DRS/mnt00001
192.168.175.116:/DRS_shared_dest/shared   24013312     20064512   3948800    84% /hana/shared
192.168.175.116:/DRS_shared_dest/usr-sap  24013312     20064512   3948800    84% /usr/sap/DRS
192.168.175.116:/DRS_log_backup_dest      7200512      6084992   1115520    85% /mnt/log-backup

```

**Note:** This output includes a replicated log backup volume..

**Note:** In Fibre Channel setup, additional steps are required before the LUNs can be mounted at the target host. See section 10, Different Steps Required in a Fibre Channel Environment.

## 9.5 Check Consistency of Latest Log Backups

The consistency check is done in the same way as described in the disaster recovery testing section 6.4, Check Consistency of Latest Log Backups.

## 9.6 Recovery of the HANA Database

The recovery of the SAP HANA database is done in the same way as described in the disaster recovery testing section 6.5, Recover the HANA Database.

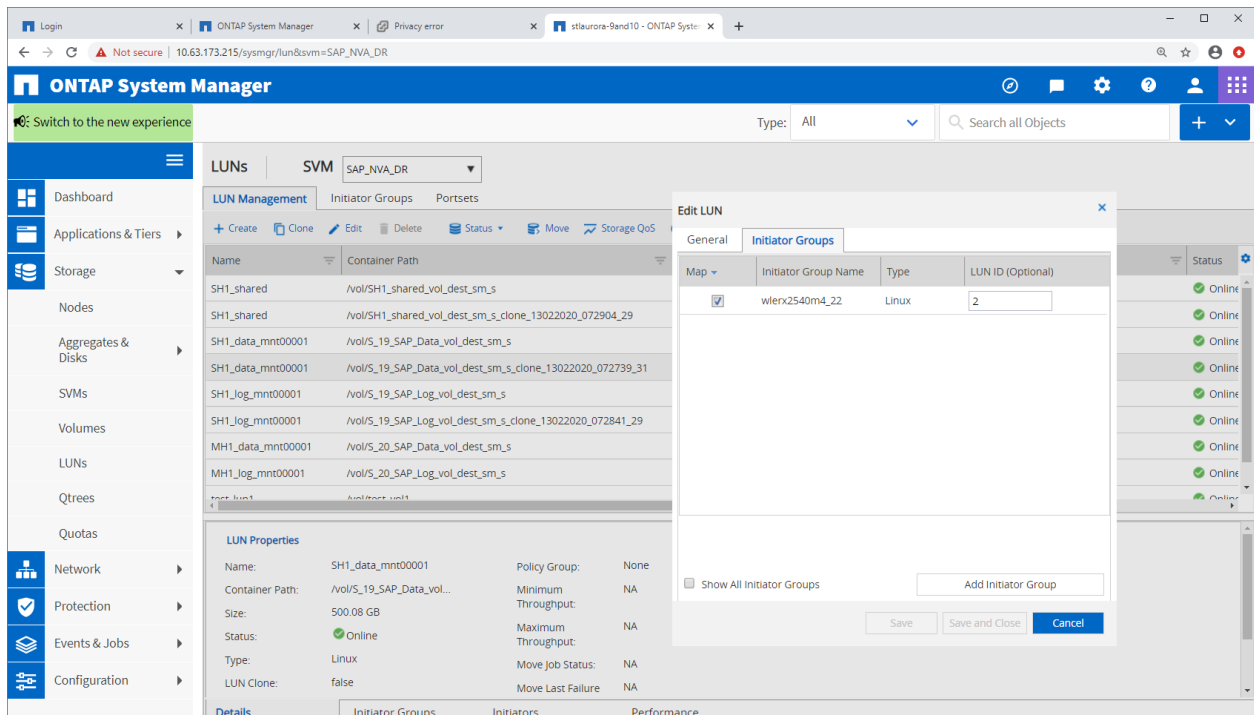
# 10 Different Steps Required in a Fibre Channel Environment

Most of the steps required for disaster recovery testing and disaster recovery failover are identical with those for NFS and Fibre Channel. This section describes the steps that are different.

## 10.1 Mapping LUNs to Disaster Recovery Server

After FlexClone volumes have been created or SnapMirror relationships have been broken off, the LUNs must be mapped to the initiator group of the disaster recovery server.

Figure 70) Mapping LUN to target host.



## 10.2 Mount File Systems

After the LUNs have been mapped, the discover process at the target host can be started.

```
wlerx2540m4-19:/usr/sap # rescan-scsi-bus.sh -a
```

The device files of the different LUNs can be determined with the NetApp `sanlun` utility. Depending on the multipathing configuration, there will be multiple device files. The following output shows the output for the data volume LUN. The same command must be executed for the log and the shared volume LUNs.

```
wlerx2540m4-19:/usr/sap # sanlun lun show | grep SAP_NVA_DR | grep data
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddp
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddm
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddj
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddg
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddd
host11 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sdda
host11 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sdcx
host11 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sdcu
host11 FCP 500.1g cDOT
```

One of the device files can now be selected to determine the UUID of the LUN. The same step must be done for the log and shared LUNs.

```
wlrx2540m4-19:/usr/sap # /lib/udev/scsi_id -g -u -d /dev/sddp  
3600a098038304132793f4f7050757369
```

This UUIDs are now used to mount the file systems.

```
wlrx2540m4-19:~ # cat /etc/fstab  
/dev/system/swap swap defaults 0 0  
/dev/system/root / btrfs defaults 0 0  
/dev/system/root /.snapshots btrfs subvol=@/.snapshots 0 0  
/dev/system/root /var btrfs subvol=@/var 0 0  
/dev/system/root /usr/local btrfs subvol=@/usr/local 0 0  
/dev/system/root /tmp btrfs subvol=@/tmp 0 0  
/dev/system/root /srv btrfs subvol=@/srv 0 0  
/dev/system/root /root btrfs subvol=@/root 0 0  
/dev/system/root /opt btrfs subvol=@/opt 0 0  
/dev/system/root /home btrfs subvol=@/home 0 0  
/dev/system/root /boot/grub2/x86_64-efi btrfs subvol=@/boot/grub2/x86_64-efi 0 0  
/dev/system/root /boot/grub2/i386-pc btrfs subvol=@/boot/grub2/i386-pc 0 0  
UUID=AD43-17E1 /boot/efi vfat defaults 0 0  
  
/dev/mapper/3600a098038304132793f4f705075736a /hana/log/SH1/mnt00001 xfs  
relatime,inode64,nobarrier,noauto 0 0  
/dev/mapper/3600a098038304132793f4f7050757369 /hana/data/SH1/mnt00001 xfs  
relatime,inode64,noauto 0 0  
/dev/mapper/3600a098038304133535d4f7466746a2f /hana/shared xfs defaults,noauto 0 0
```

```
wlrx2540m4-19:/usr/sap # df -h | grep 3600  
/dev/mapper/3600a098038304132793f4f7050757369 500G 7.1G 493G 2% /hana/data/SH1/mnt00001  
/dev/mapper/3600a098038304132793f4f705075736a 500G 5.7G 495G 2% /hana/log/SH1/mnt00001  
/dev/mapper/3600a098038304133535d4f7466746a2f 500G 67G 434G 14% /hana/shared
```

### 10.3 Cleanup Operation

Unmount data, log, and shared file systems.

```
wlrx2540m4-22:~ # umount /hana/data/SH1/mnt00001 /hana/log/SH1/mnt00001 /hana/shared
```

Check in the output of `multipath -ll` for the WWID and the multipath devices of the LUN to be deleted.

```
wlrx2540m4-22:~ # multipath -ll  
3600a098038304132793f4f7050757369 dm-5 NETAPP,LUN C-Mode  
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw  
`-+- policy='service-time 0' prio=50 status=active  
|- 11:0:7:3 sdbi 67:192 active ready running  
|- 11:0:0:3 sde 8:64 active ready running  
|- 11:0:2:3 sdt 65:48 active ready running  
|- 11:0:6:3 sdbe 67:128 active ready running  
|- 12:0:7:3 sddq 71:128 active ready running  
|- 12:0:1:3 sdbx 68:176 active ready running  
|- 12:0:2:3 sdcx 68:240 active ready running  
`- 12:0:6:3 sddm 71:64 active ready running
```

Remove the WWID by using `multipath -f <WWID>`.

```
wlrx2540m4-19:~ # multipath -f /dev/mapper/3600a098038304132793f4f7050757369  
wlrx2540m4-19:~ # multipath -f /dev/mapper/3600a098038304132793f4f705075736a  
wlrx2540m4-19:~ # multipath -f /dev/mapper/3600a098038304133535d4f7466746a2f  
wlrx2540m4-19:~ #
```

To remove the multipath device from *all* paths to the LUN, run  
echo 1 > /sys/bus/scsi/devices/\${H:B:T:L}/delete  
(where H = host:B = bus:T = target:L = lun).

```
wlerx2540m4-19:~ # echo 1 > /sys/bus/scsi/devices/11:0:7:3/delete
```

**Note:** This command must be executed for all device files listed in the `multipath -ll` output.

See also [HOWTO: Add, Resize and Remove LUN without restarting SLES or OES Linux](#).

## Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- TR-4436: SAP HANA on NetApp All Flash FAS Systems with Fibre Channel Protocol  
<http://www.netapp.com/us/media/tr-4436.pdf>
- TR-4435: SAP HANA on NetApp All Flash FAS Systems with NFS  
<http://www.netapp.com/us/media/tr-4435.pdf>
- TR-4614: SAP HANA Backup and Recovery with SnapCenter  
<https://www.netapp.com/us/media/tr-4614.pdf>
- TR-4667: Automating SAP System Copies Using the SnapCenter  
<https://www.netapp.com/us/media/tr-4667.pdf>
- TR-4719: SAP HANA System Replication, Backup and Recovery with SnapCenter  
<https://www.netapp.com/us/media/tr-4719.pdf>
- NetApp SnapCenter Software Resource page:  
<http://mysupport.netapp.com/snapcenter/resources>
- SAP Software Solutions product page:  
<http://www.netapp.com/us/solutions/applications/sap/index.aspx>

## Version History

| Version     | Date          | Document Version History  |
|-------------|---------------|---|
| Version 1.0 | November 2017 | Initial release.  |
| Version 1.1 | April 2018    | Update to cover SnapCenter 4.0  |
| Version 2.0 | July 2020     | Complete rewrite: <ul style="list-style-type: none"><li>• Three-site disaster recovery solution with synchronous and asynchronous SnapMirror replication</li><li>• DR testing and failover workflows for synchronous and asynchronous replication</li><li>• New software versions, SnapCenter 4.3, HANA 2.0 SPS4, SLES 15SP1, ONTAP 9.7</li></ul> |

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

### **Copyright Information**

Copyright © 2017–2020 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

### **Trademark Information**

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

TR-4646-0720