



Technical Report

# SAP HANA System Replication

## Backup and Recovery with SnapCenter

Nils Bauer, NetApp  
October 2018 | TR-4719

### **Abstract**

This document describes how NetApp® SnapCenter® technology and the SAP HANA plug-in can be used for backup and recovery in an SAP HANA System Replication environment.

## TABLE OF CONTENTS

|          |   |           |
|----------|---|-----------|
| <b>1</b> | <b>SAP HANA System Replication Overview .....</b>                           | <b>4</b>  |
| 1.1      | High Availability with an RPO of Zero and a Minimal RTO .....               | 4         |
| 1.2      | Disaster Recovery over a Large Distance .....                               | 4         |
| <b>2</b> | <b>Storage Snapshot Backups and SAP System Replication .....</b>            | <b>5</b>  |
| <b>3</b> | <b>SnapCenter Configuration Options for SAP System Replication .....</b>    | <b>6</b>  |
| 3.1      | SnapCenter Configuration with Separate Resources .....                      | 7         |
| 3.2      | SnapCenter Configuration with a Single Resource .....                       | 8         |
| 3.3      | Summary .....   | 10        |
| <b>4</b> | <b>SnapCenter Backup Operations .....</b>                                   | <b>11</b> |
| 4.1      | Backup Operation with Separate SnapCenter Resources .....                   | 11        |
| 4.2      | Backup Operation with a Single SnapCenter Resource .....                    | 13        |
| <b>5</b> | <b>Restore and Recovery .....</b>   | <b>17</b> |
| 5.1      | Overview of Restore and Recovery Operations .....                           | 17        |
| 5.2      | Restore and Recovery with Separate SnapCenter Resources Configuration ..... | 17        |
| 5.3      | Restore and Recovery with a Single SnapCenter Resource Configuration .....  | 18        |
| 5.4      | Restore and Recovery from a Backup Created at the Other Host .....          | 23        |
|          | <b>Where to Find Additional Information .....</b>                           | <b>26</b> |
|          | <b>Version History .....</b>  | <b>26</b> |

## LIST OF FIGURES

|  |    |
|--|----|
| Figure 1) SAP System Replication as a high-availability solution.....              | 4  |
| Figure 2) SAP System Replication as a disaster recovery solution.....              | 5  |
| Figure 3) Backup operation with SAP System Replication.....                        | 5  |
| Figure 4) Restore operation with SAP System Replication.....                       | 6  |
| Figure 5) SnapCenter configuration options for SAP System Replication.....         | 6  |
| Figure 6) Backup operation with host 1 as the primary host.....                    | 7  |
| Figure 7) Backup operation with host 2 as the primary host.....                    | 8  |
| Figure 8) Data and log backup housekeeping.....                                    | 8  |
| Figure 9) Backup operation with host 1 as the primary host.....                    | 9  |
| Figure 10) Backup operation with host 2 as the primary host.....                   | 10 |
| Figure 11) Identification of the backup host.....                                  | 10 |
| Figure 12) Summary of configuration options.....                                   | 11 |
| Figure 13) Lab setup: SnapCenter with separate resources.....                      | 12 |
| Figure 14) Maintenance mode activation.....  | 12 |
| Figure 15) Lab setup for SnapCenter with a single resource.....                    | 13 |
| Figure 16) SnapCenter resource configuration.....                                  | 14 |
| Figure 17) SnapCenter resource configuration: storage footprint.....               | 14 |
| Figure 18) Resource protection configuration.....                                  | 15 |
| Figure 19) Backup job log with host 1 as the primary host.....                     | 15 |
| Figure 20) Backup job log with host 2 as the primary host.....                     | 16 |
| Figure 21) SAP HANA backup catalog.....  | 16 |
| Figure 22) Overview of restore and recovery operations.....                        | 17 |
| Figure 23) Restore operations with a single SnapCenter resource configuration..... | 18 |
| Figure 24) SnapCenter restore of the valid backup only.....                        | 19 |
| Figure 25) SAP HANA Studio before the restore operation.....                       | 19 |
| Figure 26) Restore and recovery with SAP HANA Studio.....                          | 20 |
| Figure 27) File-level restore with SnapCenter.....                                 | 20 |
| Figure 28) Backup and log backup selection.....                                    | 21 |
| Figure 29) Start of secondary host and resynchronization.....                      | 21 |
| Figure 30) SnapCenter restore of valid backup and crash image.....                 | 22 |
| Figure 31) Complete resource restore operation.....                                | 22 |
| Figure 32) Start of secondary host and resynchronization.....                      | 23 |
| Figure 33) Restore and recovery from a backup created at the other host.....       | 23 |
| Figure 34) SAP HANA Studio before restore operation.....                           | 24 |
| Figure 35) SnapCenter clone workflow.....  | 25 |
| Figure 36) Junction path information.....  | 25 |

# 1 SAP HANA System Replication Overview

SAP HANA System Replication is commonly used as a high-availability or disaster recovery solution for SAP HANA databases. SAP HANA System Replication provides different operating modes that you can use depending on the use case or availability requirements.

There are two primary use cases that can also be combined:

- High availability with a recovery point objective (RPO) of zero and a minimal recovery time objective (RTO).
- Disaster recovery over a large distance. The secondary SAP HANA host can also be used for development or testing during normal operation.

## 1.1 High Availability with an RPO of Zero and a Minimal RTO

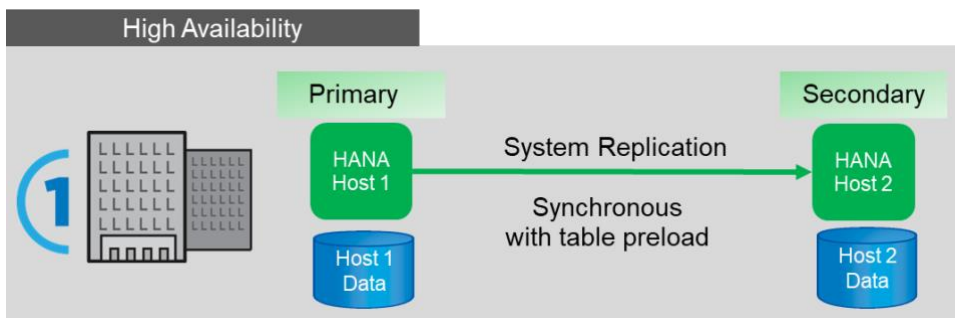
System Replication is configured with synchronous replication using tables preloaded into memory at the secondary SAP HANA host. This high-availability solution can be used to address hardware or software failures and also to reduce planned downtime during SAP HANA software upgrades (near-zero downtime).

Failover operations are often automated by using third-party cluster software or with a one-click workflow with SAP Landscape Management software.

From a backup requirement perspective, you must be able to create backups independent of which SAP HANA host is primary or secondary. A shared backup infrastructure is used to restore any backup, regardless of which host the backup has been created on.

The rest of this document focuses on backup operations with SAP System Replication configured as a high-availability solution.

Figure 1) SAP System Replication as a high-availability solution.



## 1.2 Disaster Recovery over a Large Distance

System Replication can be configured with asynchronous replication with no table preloaded into memory at the secondary host. This solution is used to address data center failures, and failover operations are typically performed manually.

Regarding backup requirements, you must be able to create backups during normal operation in data center 1 and during disaster recovery in data center 2. A separate backup infrastructure is available in data centers 1 and 2, and backup operations are activated as a part of disaster failover. The backup infrastructure is typically not shared, and a restore operation of a backup that was created at the other data center is not possible.

Figure 2) SAP System Replication as a disaster recovery solution.



## 2 Storage Snapshot Backups and SAP System Replication

Backup operations are always performed at the primary SAP HANA host. The required SQL commands for the backup operation cannot be performed at the secondary SAP HANA host.

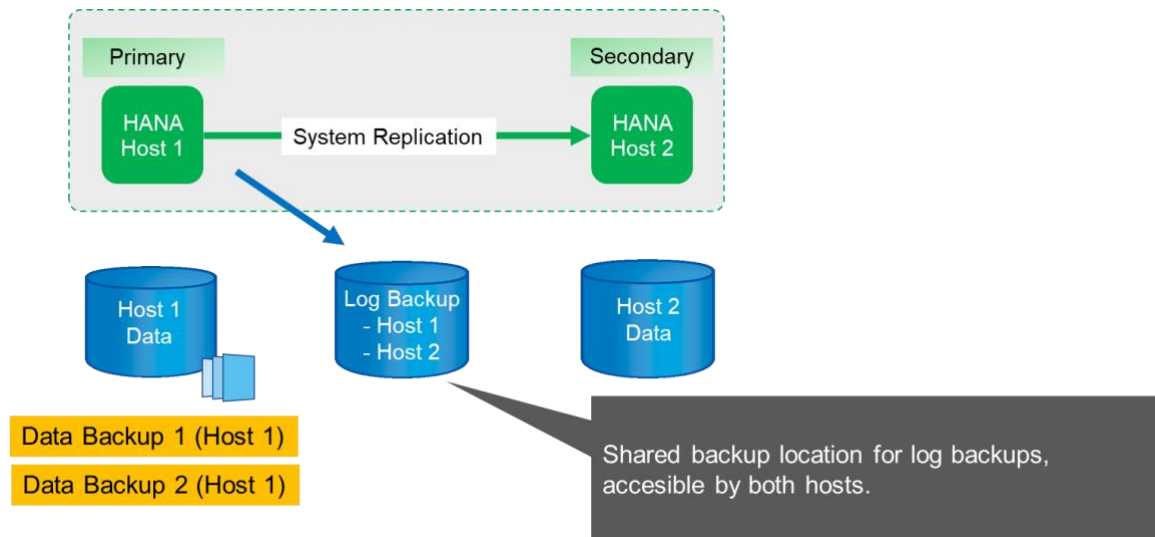
For SAP HANA backup operations, the primary and secondary SAP HANA hosts are a single entity. They share the same SAP HANA backup catalog, and they use backups for restore and recovery, regardless of whether the backup was created at the primary or secondary SAP HANA host.

The ability to use any backup for restore and to do forward recovery using log backups from both hosts requires a shared log backup location that is accessible from both hosts. NetApp recommends that you use a shared storage volume. However, you should also separate the log backup destination into subdirectories within the shared volume.

Each SAP HANA host has its own storage volume. When you use a storage-based NetApp® Snapshot™ copy to perform a backup, a database-consistent Snapshot copy is created on the primary SAP HANA host's storage volume.

Figure 3 shows an overview of a backup operation with SAP System Replication in which two Snapshot backups have been created at the primary host.

Figure 3) Backup operation with SAP System Replication.

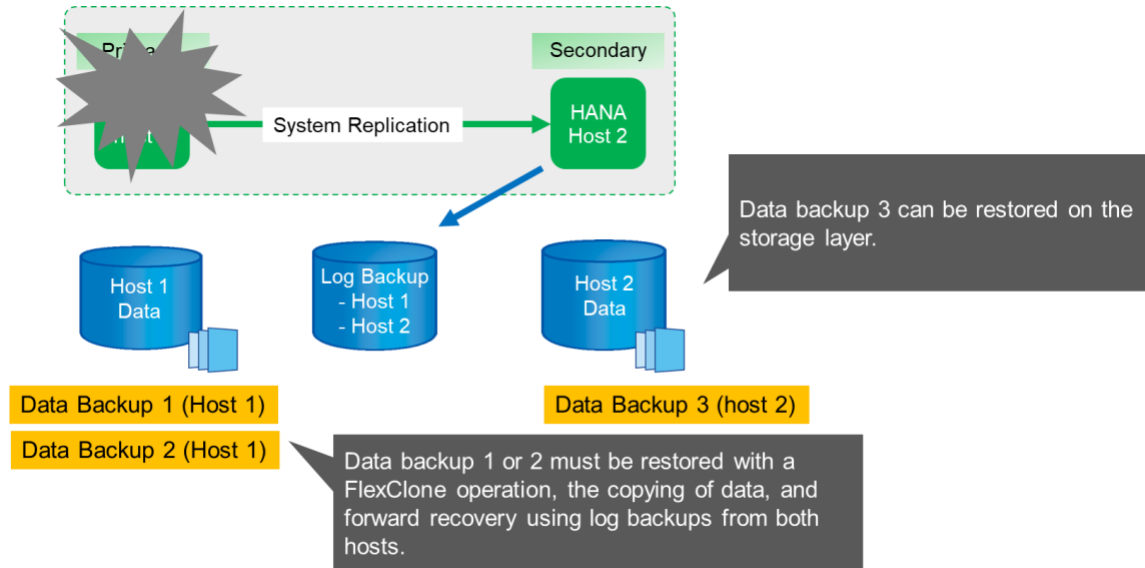


When a failover is performed for host 2, the backups are executed at host 2 and Snapshot copies are created at the storage volume of host 2.

The backup created at host 2 can be restored directly at the storage layer. If you must use a backup created at host 1, then the backup must be copied from the host 1 storage volume to the host 2 storage volume. Forward recovery uses the log backups from both hosts.

Figure 4 shows an overview of the two different restore operations.

Figure 4) Restore operation with SAP System Replication.

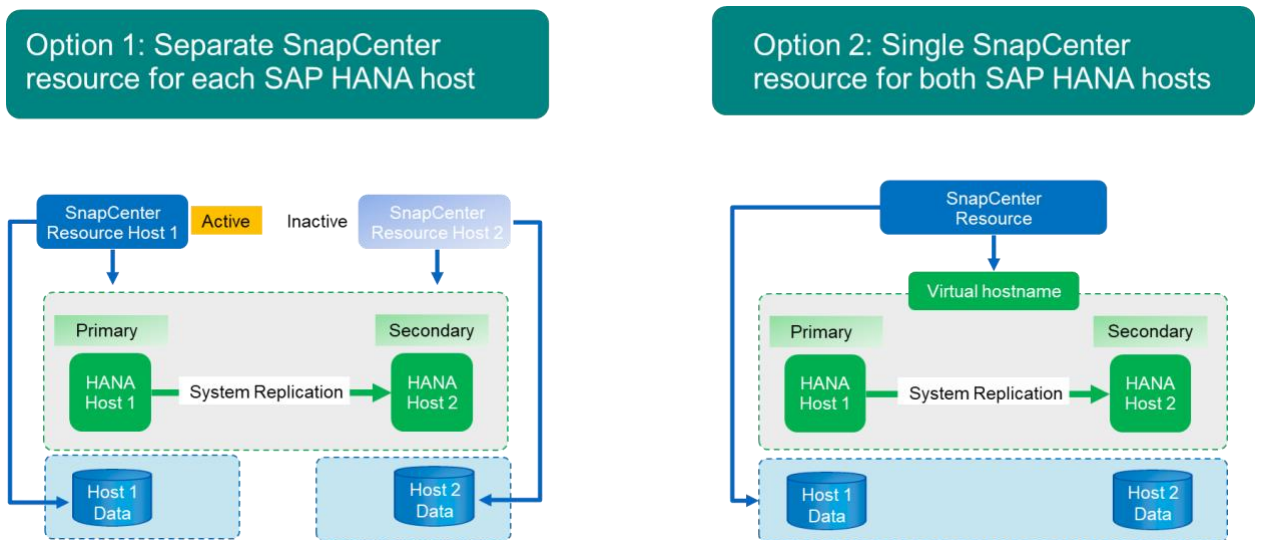


### 3 SnapCenter Configuration Options for SAP System Replication

There are two options for configuring data protection with NetApp SnapCenter® software in an SAP HANA System Replication environment, as Figure 5 shows:

- A separate SnapCenter resource for each SAP HANA host
- A single SnapCenter resource for both SAP HANA hosts

Figure 5) SnapCenter configuration options for SAP System Replication.



With separate resources for each SAP HANA host, SnapCenter is configured in the same way as with SAP HANA hosts without SAP System Replication. Each host is configured using its physical IP address (host name) and its individual data volume on the storage layer. Scheduled backup operations are activated and deactivated in SnapCenter, depending on which host is primary or secondary.

With a single-resource configuration for both SAP HANA hosts, the SnapCenter resource is configured using the virtual IP address of the SAP HANA System Replication hosts. Both data volumes of the SAP HANA hosts are included in the SnapCenter resource.

The two options are discussed in more detail in the following sections.

### 3.1 SnapCenter Configuration with Separate Resources

Figure 6 shows SnapCenter configuration with two separate resources. Each SAP HANA host of the SAP HANA System Replication environment is configured with its individual physical IP address (host name) and data volume on the storage layer.

SnapCenter policies and schedules are also configured for each resource. Because the SAP HANA backup operation can be performed only at the primary host, the secondary host must be put into maintenance mode in SnapCenter. You can activate and deactivate maintenance mode in the topology view of each resource.

When a backup is created at host 1, the primary host at this point in time, a Snapshot copy is created on the data volume of host 1. The backup is registered in the SAP HANA backup catalog and in the SnapCenter resource for host 1. For the inactive resource (host 2, the secondary host), no backup is created.

Figure 6) Backup operation with host 1 as the primary host.

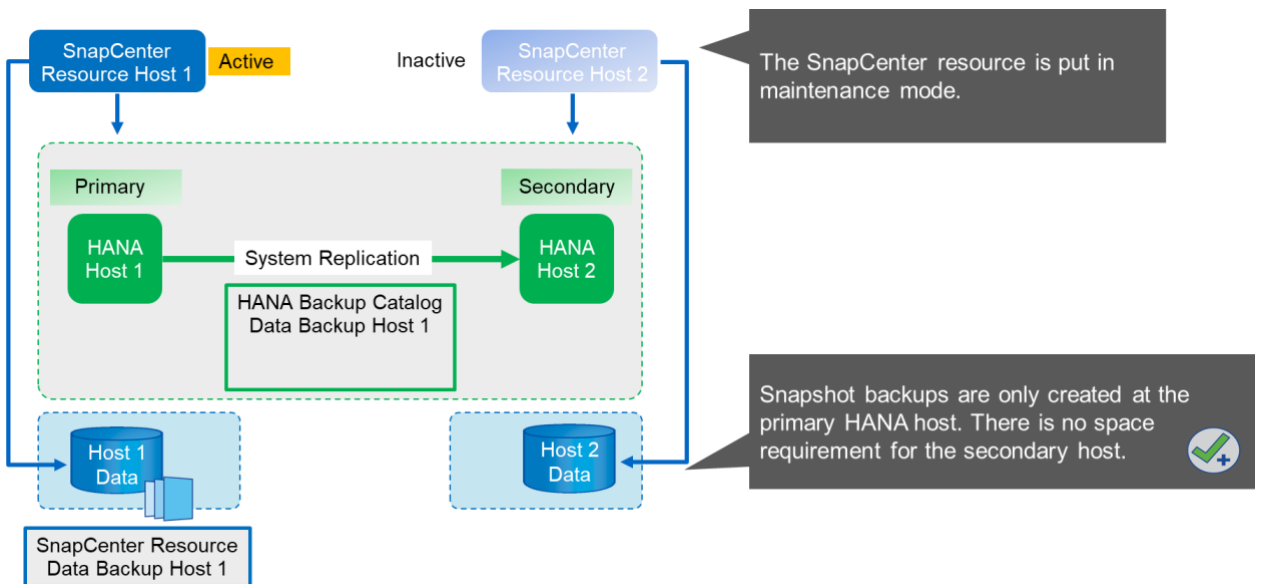
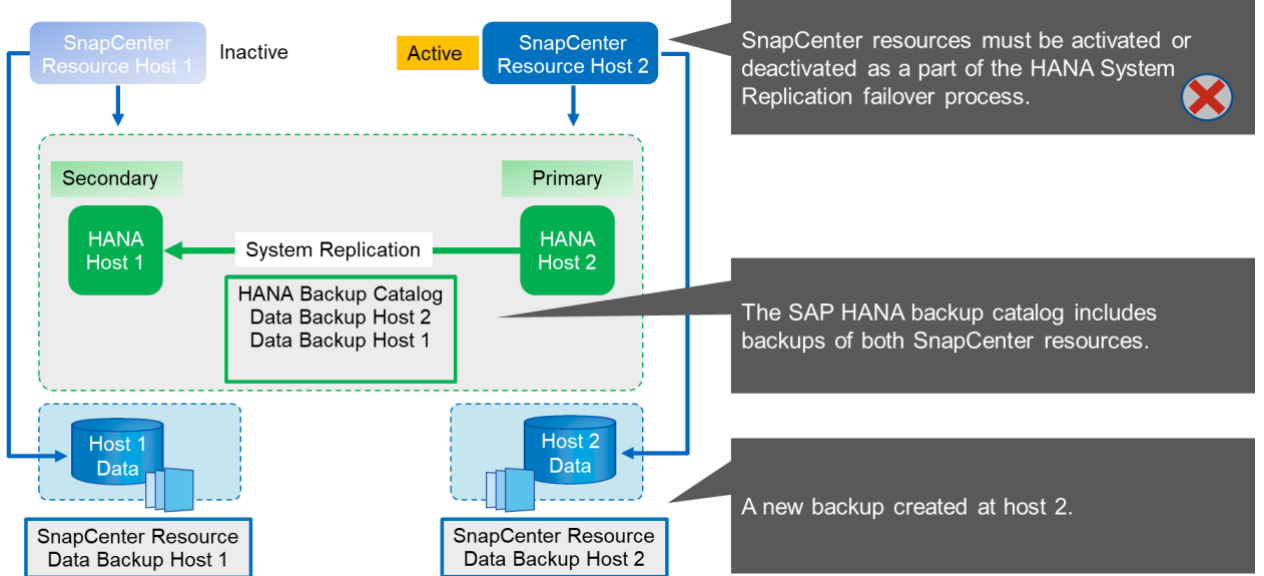


Figure 7 shows the backup operation after a failover to host 2 and a replication from host 2 to host 1. As part of the SAP HANA System Replication failover workflow, you must put the SnapCenter resource for host 1 in maintenance mode, and you must put the resource of host 2 in production mode. Backups are now executed at host 2, and Snapshot copies are created at the data volume of host 2. The SAP HANA backup catalog now includes a backup that has been created at host 1, and another backup is created at host 2. The SnapCenter resource for host 2 includes only the backup created at host 2.

As discussed in the section “Storage Snapshot Backups and SAP System Replication,” the restore operation with storage-based Snapshot backups is different, depending on which backup needs to be restored. It is important to identify where the backup was created to determine whether the restore can be performed at the local storage volume or must be performed from the other host’s storage volume. With two separate SnapCenter resources, this identification is performed on the SnapCenter resource level.

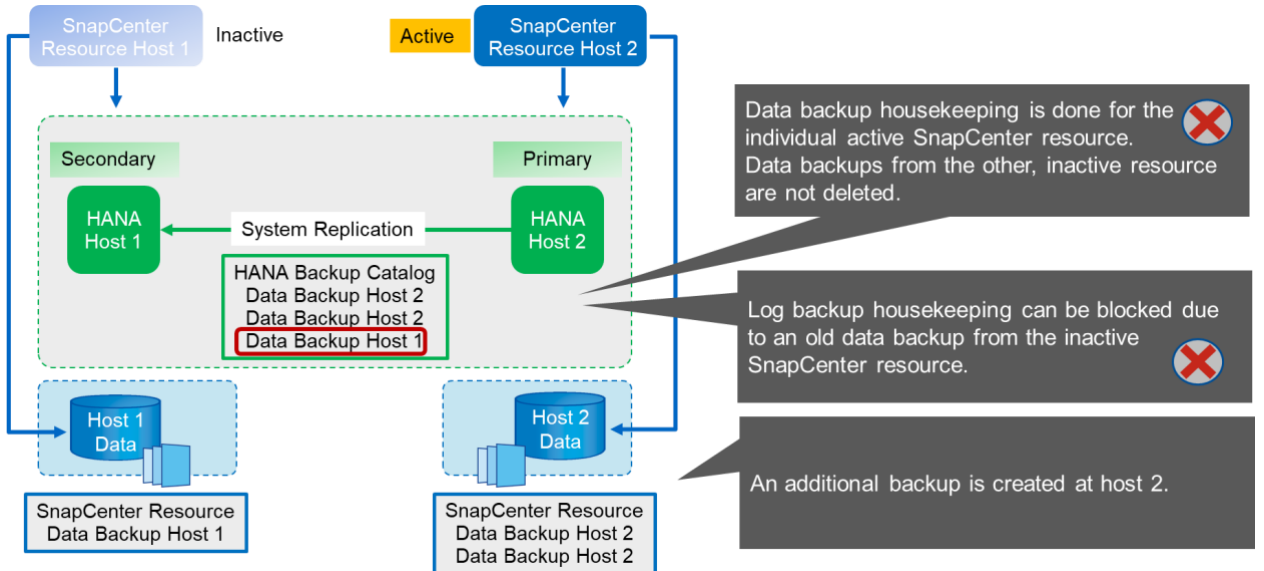
Figure 7) Backup operation with host 2 as the primary host.



The housekeeping of data backups, based on the retention policy defined in SnapCenter, is done only for the backups of the active SnapCenter resource. As Figure 8 shows, the backup that was created at host 1 is not deleted as long as host 2 is the active resource. Therefore, the backup from host 1 becomes the oldest backup, and all log backups that are required for forward recovery of this backup are not deleted.

To clean up Snapshot copies on the data volume of host 1 and to clean up log backups, you must delete the data backup of host 1 manually in SnapCenter and the SAP HANA backup catalog.

Figure 8) Data and log backup housekeeping.



### 3.2 SnapCenter Configuration with a Single Resource

Figure 9 shows a SnapCenter configuration with a single resource. The SnapCenter resource is configured with the virtual IP address (host name) of the SAP System Replication environment. With this approach, SnapCenter always communicates with the primary host, regardless of whether host 1 or host 2 is primary. The data volumes of both SAP HANA hosts are included in the SnapCenter resource.

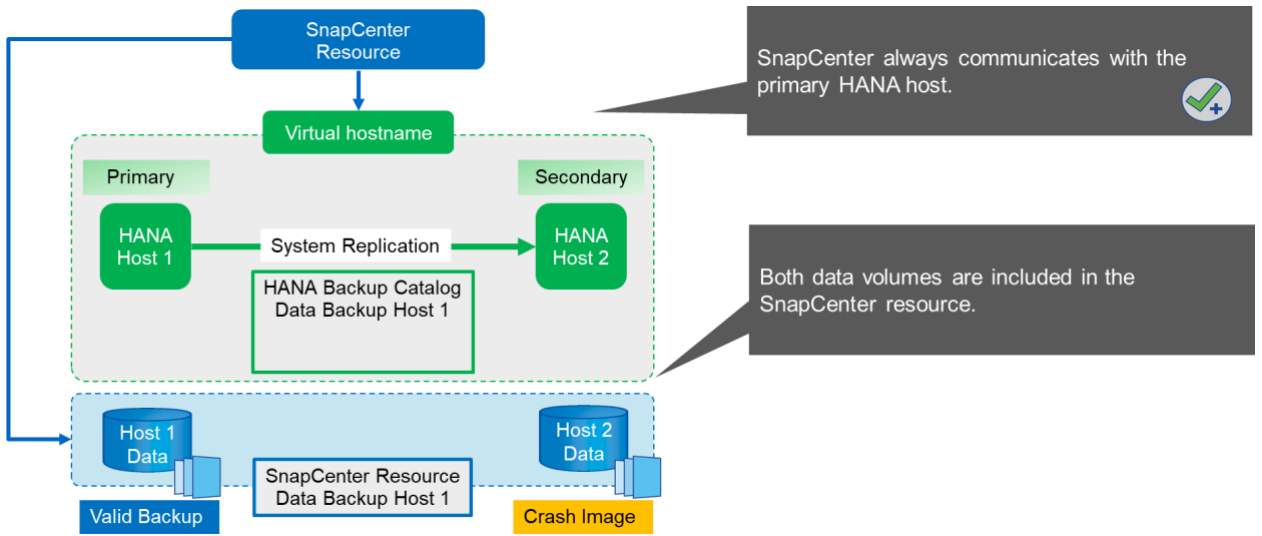


**Note:** We assume that the virtual IP address is always bound to the primary SAP HANA host. The failover of the virtual IP address is performed outside SnapCenter as part of the SAP System Replication failover workflow.

When a backup is executed with host 1 as the primary host, a database-consistent Snapshot backup is created at the data volume of host 1. Because the data volume of host 2 is part of the SnapCenter resource, another Snapshot copy is created for this volume. This Snapshot copy is not database consistent; rather, it is just a crash image of the secondary host.

The SAP HANA backup catalog and the SnapCenter resource includes the backup created at host 1.

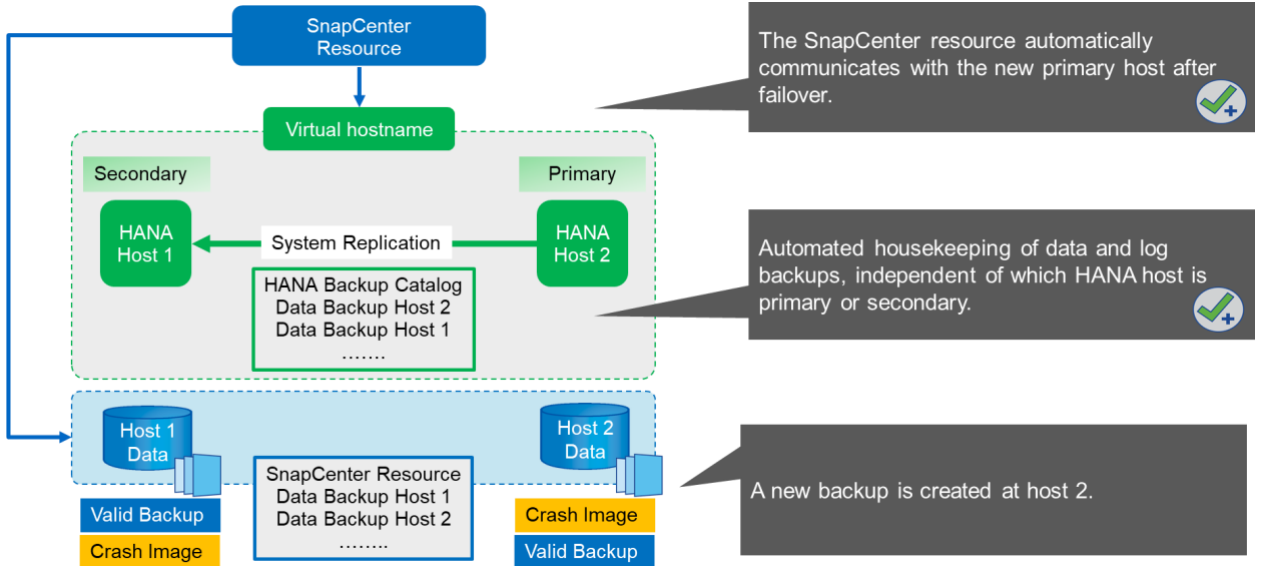
**Figure 9) Backup operation with host 1 as the primary host.**



**Error! Reference source not found.** shows the backup operation after failover to host 2 and replication from host 2 to host 1. SnapCenter automatically communicates with host 2 by using the virtual IP address configured in the SnapCenter resource. Backups are now created at host 2. Two Snapshot copies are created by SnapCenter: a database-consistent backup at the data volume at host 2 and a crash image Snapshot copy at the data volume at host 1. The SAP HANA backup catalog and the SnapCenter resource now include the backup created at host 1 and the backup created at host 2.

Housekeeping of data and log backups is based on the defined SnapCenter retention policy, and backups are deleted regardless of which host is primary or secondary.

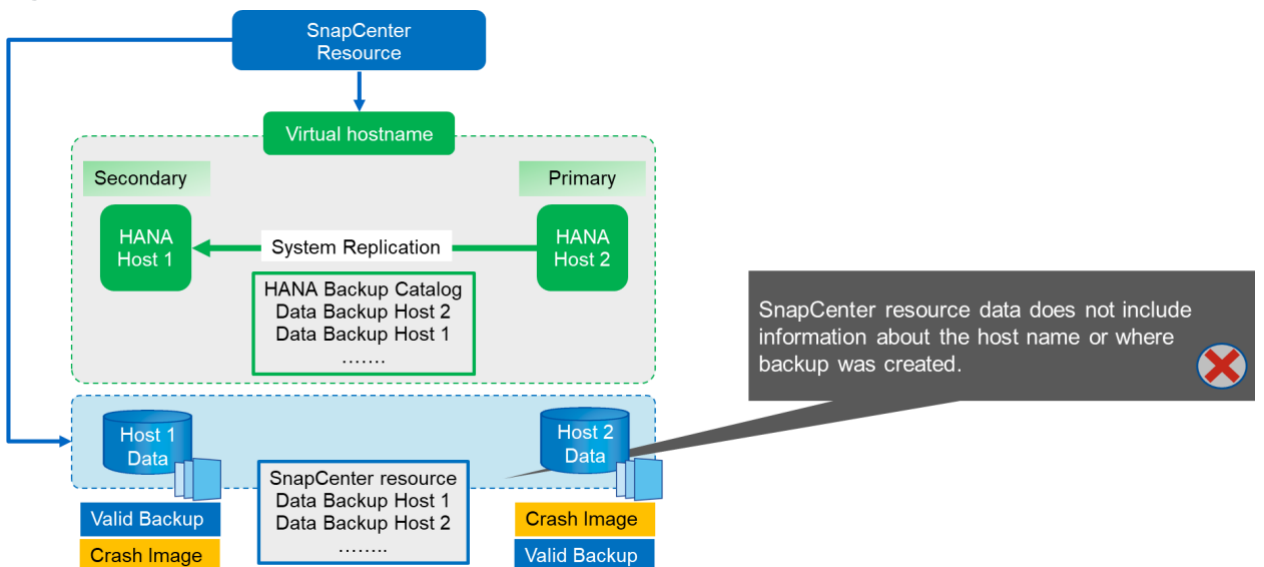
Figure 10) Backup operation with host 2 as the primary host.



As discussed in the section “Storage Snapshot Backups and SAP System Replication,” a restore operation with storage-based Snapshot backups is different, depending on which backup must be restored. It is important to identify which host the backup was created at to determine if the restore can be performed at the local storage volume, or if the restore must be performed at the other host’s storage volume.

With a single-resource SnapCenter configuration, SnapCenter is not aware of where the backup was created. Therefore, NetApp recommends that you add a prebackup script to the SnapCenter backup workflow to identify which host is currently the primary SAP HANA host.

Figure 11) Identification of the backup host.



### 3.3 Summary

Figure 12 summarizes the pros and cons of the two configuration options.

Figure 12) Summary of configuration options.

|   | Separate SnapCenter Resources                              | Single SnapCenter Resource                                       |
|---|--|--|
| Backup operation after failover               | ❌ Manual activation or deactivation of SnapCenter resource | ✅ Automatic, using virtual hostname                              |
| Backup housekeeping                           | ❌ Manual housekeeping required for inactive resource       | ✅ Automatic, using single resource                               |
| Backup capacity requirements                  | ✅ Backups are only created at active SnapCenter resource   | ❌ Backups are always created for data volumes of both HANA hosts |
| Identification, which host created the backup | ✅ Visible, using separate SnapCenter resources             | ❌ Prebackup script to identify which host is active              |

## Separate SnapCenter Resources

A SnapCenter configuration with separate resources requires additional steps and manual operations to make sure that backups are created at the primary SAP HANA host. These operations are also needed to clean up data and log backups after an SAP HANA failover.

On the other hand, backups are created only at the active resource (the primary SAP HANA host), and no additional storage capacity is required for the inactive resource. NetApp recommends this setup if there is a preferred primary HANA host and you avoid operations in failover mode.

## Single SnapCenter Resource

A SnapCenter configuration with a single resource simplifies backup operations because SnapCenter always communicates with the correct HANA host (the primary host, through its virtual IP address). Additionally, housekeeping of data and log backups works out of the box.

As a downside, additional capacity is required for crash-image Snapshot copies at the secondary host. To identify where a backup has been created, you must add a prebackup script to the SnapCenter backup workflow.

NetApp recommends this setup if backup operation simplicity is important and there is no preferred primary HANA host.

## Valid for Both Configuration Options

A shared log backup volume is required to enable forward recovery with a backup that has been created at the other SAP HANA host. It is also required if a failover has occurred between backup creation and the actual point in time that you want to recover to.

Restoring from a backup that was created at the other HANA host requires manual steps and cannot be performed with SnapCenter.

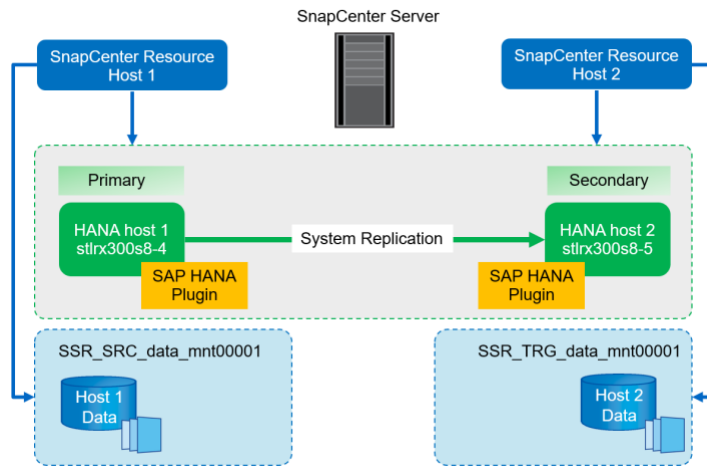
# 4 SnapCenter Backup Operations

## 4.1 Backup Operation with Separate SnapCenter Resources

### SnapCenter Configuration

Figure 13 shows the lab setup and an overview of the required NetApp SnapCenter configuration.

Figure 13) Lab setup: SnapCenter with separate resources.



To configure SnapCenter with separate resources, the SAP HANA plug-in must be deployed on each SAP HANA host.

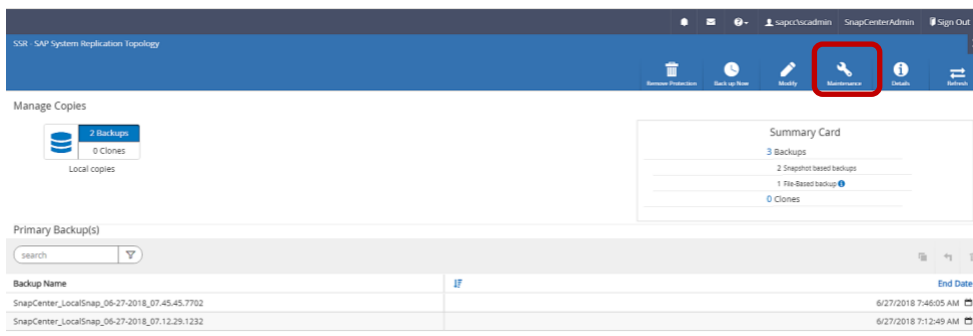
**Note:** SnapCenter uses the security identifier (SID), the tenant name, and the `hdbsql` client host as a unique resource identifier. Since the SID and the tenant name are identical for the system replication resources, the `hdb` client host must be different. Therefore, a central communication host cannot be used.

The configuration of both resources is identical to a non-system replication setup and is described in the technical report [SAP HANA Backup and Recovery with SnapCenter](#).

## SnapCenter Backup Operation

As discussed in the section “SnapCenter Configuration with Separate Resources,” the SnapCenter resource of the secondary SAP HANA host must be put into maintenance mode so that backup operations are executed only for the primary SAP HANA host. Figure 14 shows the topology view of one of the resources. Maintenance mode can be activated or deactivated using the Maintenance button in the upper right of the screen.

Figure 14) Maintenance mode activation.



Backup operations are performed as usual. If an SAP HANA System Replication failover occurs, the old primary resource must be put into maintenance mode, and the old secondary resource must be put into production mode.

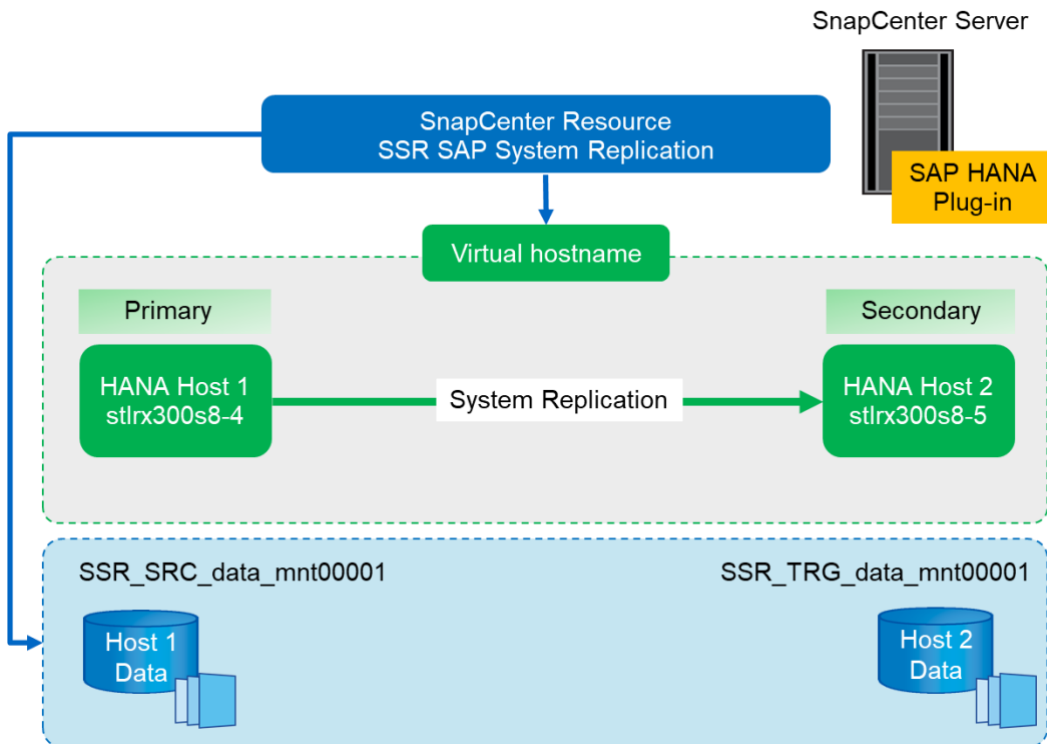
After failover, SnapCenter does not delete the backups and SAP HANA catalog entries of the inactive resource. You must delete these backups manually in SnapCenter and the SAP HANA backup catalog to make sure that log backup housekeeping is not blocked by the old backup of the inactive resource.

## 4.2 Backup Operation with a Single SnapCenter Resource

### SnapCenter Configuration

Figure 15 shows the lab setup and an overview of the required SnapCenter configuration.

Figure 15) Lab setup for SnapCenter with a single resource.



To perform backup operations regardless of which SAP HANA host is primary and even whether one host is down, the SnapCenter SAP HANA plug-in must be deployed on a central `hdbsql` communication host. In our lab setup, we used the SnapCenter server as a central communication host, and we deployed the SAP HANA plug-in on the SnapCenter server.

A user was created in the HANA database to perform backup operations. A user store key was configured at the SnapCenter server on which the SAP HANA plug-in was installed. The user store key includes the virtual IP address of the SAP HANA System Replication hosts (`ssr-vip`).

```
hdbuserstore.exe -u SYSTEM set SSRKEY ssr-vip:31013 SNAPCENTER Netapp123
```

You can find more information about SAP HANA plug-in deployment options and user store configuration in the technical report [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

In SnapCenter, the resource is configured as shown in Figure 16 using the user store key, configured before, and the SnapCenter server as the `hdbsql` communication host.

**Figure 16) SnapCenter resource configuration.**

The screenshot shows the 'Add SAP HANA Database' configuration window. The 'Provide Resource Details' step is active. The 'Resource Type' is set to 'Multitenant Database Container (MDC) - Single Tenant'. The fields are filled with the following values:

|                            |                              |
|----------------------------|------------------------------|
| HANA System Name           | SSR - SAP System Replication |
| SID                        | SSR                          |
| Tenant Database            | SSR                          |
| HDBSQL Client Host         | SC30-V2.sapcc.stl.netapp.com |
| HDB Secure User Store Keys | SSRKEY                       |
| HDBSQL OS User             | SYSTEM                       |

The data volumes of both SAP HANA hosts are included in the storage footprint configuration as Figure 17 shows.

**Figure 17) SnapCenter resource configuration: storage footprint.**

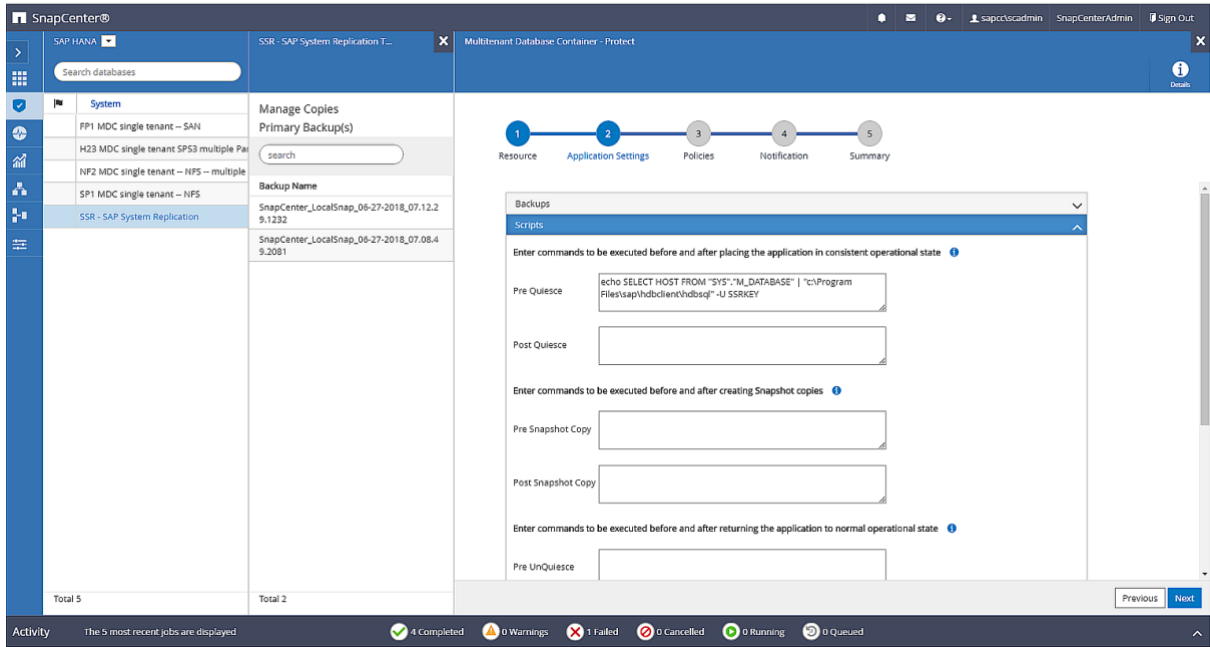
The screenshot shows the 'Add SAP HANA Database' configuration window. The 'Provide Storage Footprint Details' step is active. The 'Storage Systems for storage footprint' is set to 'hana'. A 'Modify hana' dialog box is open, showing the selection of volumes and LUNs or Qtrees:

| Volume Name           | LUNs or Qtrees                    |
|-----------------------|-----------------------------------|
| SSR_TRG_data_mnt00001 | Default is 'None' or type to find |
| SSR_SRC_data_mnt00001 | Default is 'None' or type to find |

As discussed in the section “SnapCenter Configuration with a Single Resource,” SnapCenter is not aware of where the backup was created. NetApp therefore recommends that you add a prebackup script in the SnapCenter backup workflow to identify which host is currently the primary SAP HANA host. You can perform this identification using a SQL statement that is added to the backup workflow, as Figure 18 shows.

```
Select host from "SYS".M_DATABASE
```

Figure 18) Resource protection configuration.



## SnapCenter Backup Operation

Backup operations are now executed as usual. Housekeeping of data and log backups is performed independent of which SAP HANA host is primary or secondary.

The backup job logs include the output of the SQL statement, which allows you to identify the SAP HANA host where the backup was created, as Figure 19 and Figure 20 show.

Figure 19) Backup job log with host 1 as the primary host.

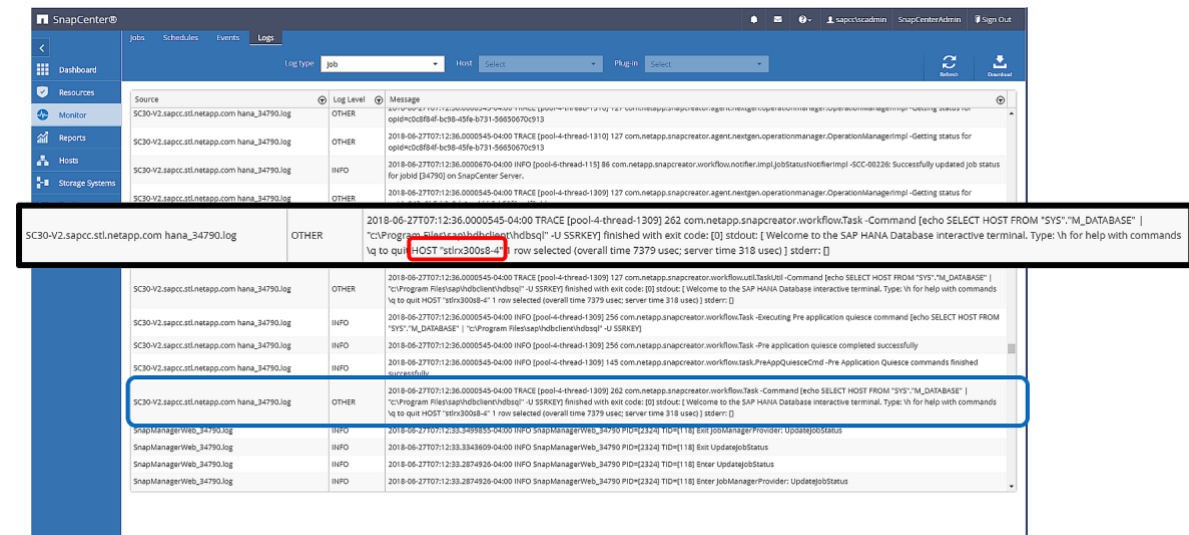


Figure 20) Backup job log with host 2 as the primary host.

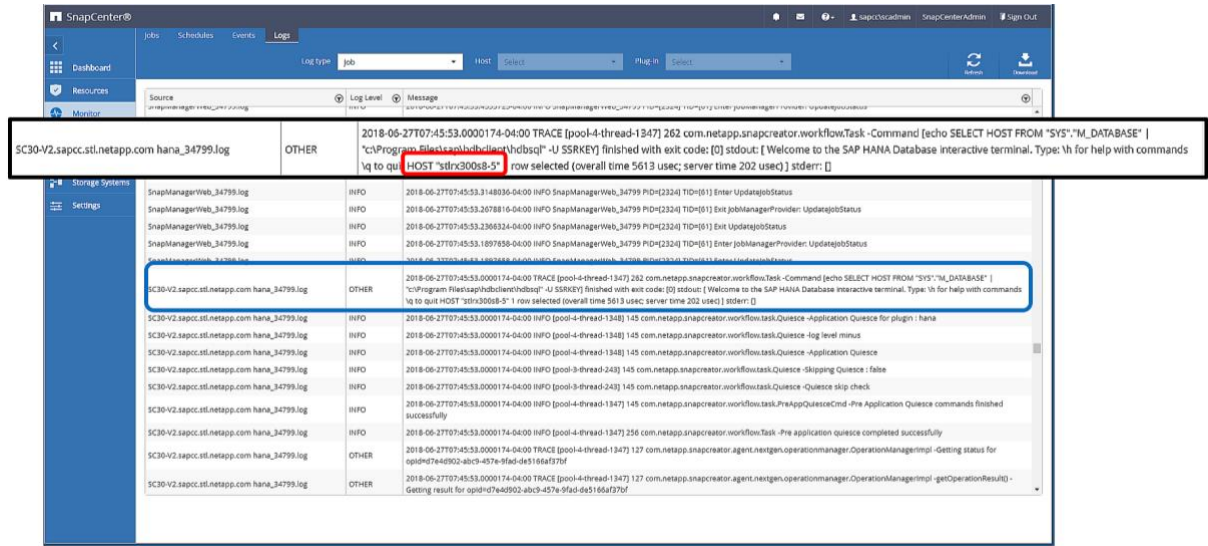
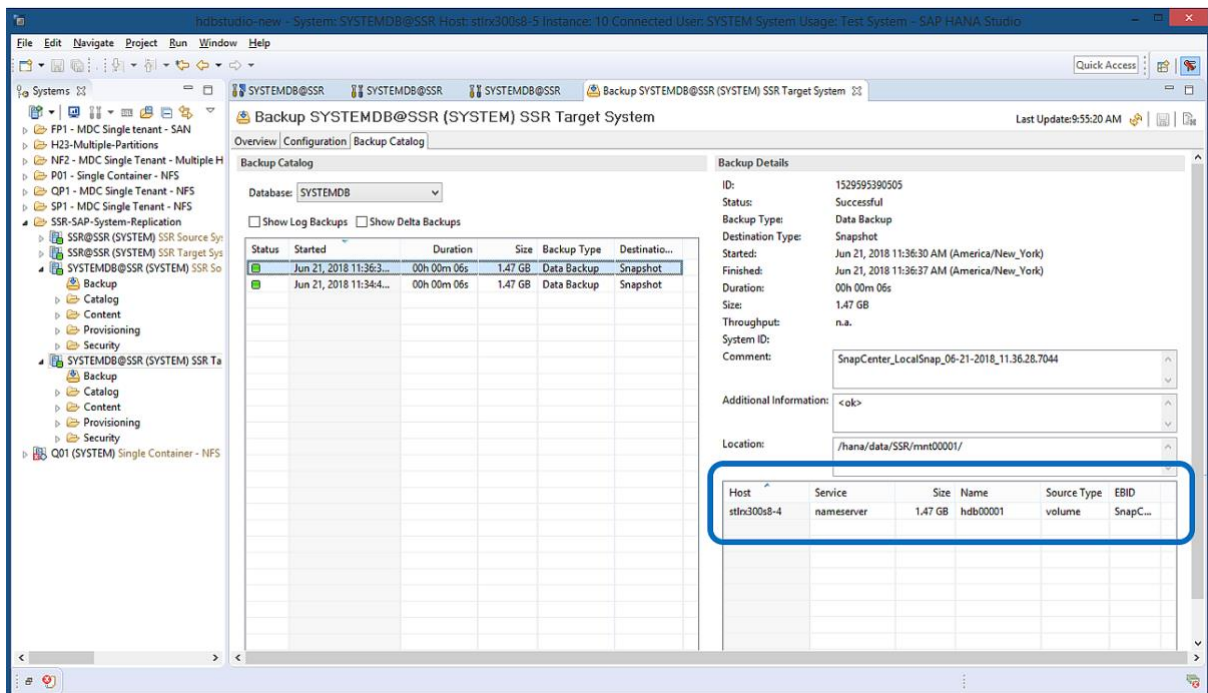


Figure 21 shows the SAP HANA backup catalog in SAP HANA Studio. When the SAP HANA database is online, the SAP HANA host where the backup was created is visible in SAP HANA Studio.

**Note:** The SAP HANA backup catalog on the file system, which is used during a restore and recovery operation, does not include the host name where the backup was created. The only way to identify the host when the database is down is to combine the backup catalog entries with the backup.log file of both SAP HANA hosts.

Figure 21) SAP HANA backup catalog.





## 5 Restore and Recovery

### 5.1 Overview of Restore and Recovery Operations

Independent of the NetApp SnapCenter configuration (either separate or single resource), there are two different restore and recovery operations.

- Restore from a backup that has been created at the current primary host. See the sections “Restore and Recovery with Separate SnapCenter Resources Configuration” and “Restore and Recovery with a Single SnapCenter Resource Configuration.”
- Restore from a backup that has been created at the other host, which is currently not the primary host. See the section “Restore and Recovery from a Backup Created at the Other Host.”

As discussed before, you must be able to identify where the selected backup was created to define the required restore operation. If the SAP HANA database is still online, you can use SAP HANA Studio to identify the host at which the backup was created. If the database is offline, the information is available at the following locations:


- With the host-specific SnapCenter resources (separate resource configuration)
- In the backup job log (single-resource configuration)

Figure 22 illustrates the different restore operations depending on the selected backup.

If a restore operation must be performed after timestamp T3 and host 1 is the primary, you can restore the backup created at T1 or T3 by using SnapCenter. These Snapshot backups are available at the storage volume attached to host 1.

If you need to restore using the backup created at host 2 (T2), which is a Snapshot copy at the storage volume of host 2, the backup needs to be made available to host 1. You can make this backup available by creating a NetApp FlexClone® copy from the backup, mounting the FlexClone copy to host 1, and copying the data to the original location.

Figure 22) Overview of restore and recovery operations.



| Restore Operation With |   |
|------------------------|---|
| Backup T1              | SnapCenter  |
| Backup T2              | Create FlexClone from „Backup host 2“, mount and copy |
| Backup T3              | SnapCenter  |

### 5.2 Restore and Recovery with Separate SnapCenter Resources Configuration

Restore and recovery operations for a SnapCenter configuration with separate resources is identical to a non-System Replication setup. For further details, see the technical report [SAP HANA Backup and Recovery with SnapCenter](#).

A restore operation from a backup that was created at the other host is described in the section “Restore and Recovery from a Backup Created at the Other Host.”

### 5.3 Restore and Recovery with a Single SnapCenter Resource Configuration

With a single SnapCenter resource configuration, Snapshot copies are created at both storage volumes of both SAP HANA System Replication hosts. Only the Snapshot backup that is created at the storage volume of the primary SAP HANA host is valid to use for forward recovery. The Snapshot copy created at the storage volume of the secondary SAP HANA host is a crash image that cannot be used for forward recovery.

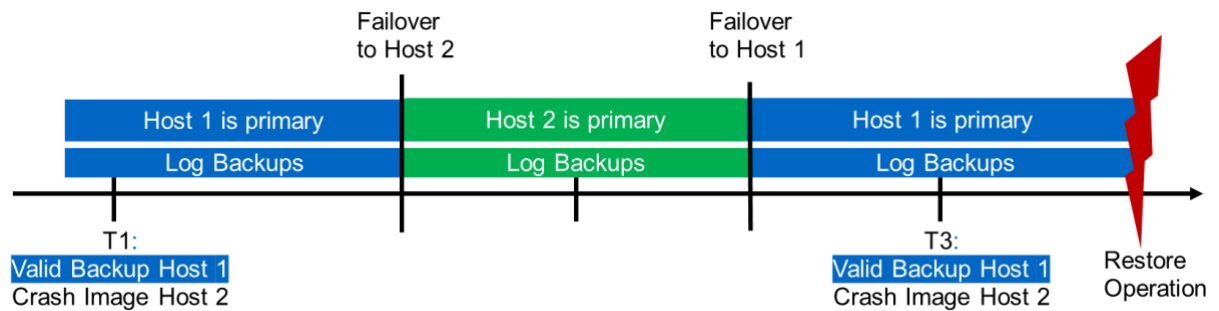
A restore operation with SnapCenter can be performed in two different ways:

- Restore only the valid backup
- Restore the complete resource, including the valid backup and the crash image

The following sections discuss the two different restore operations in more detail.

A restore operation from a backup that was created at the other host is described in the section “Restore and Recovery from a Backup Created at the Other Host.”

Figure 23) Restore operations with a single SnapCenter resource configuration.



#### SnapCenter Restore of the Valid Backup Only

Figure 24 shows an overview of the restore and recovery scenario described in this section.

A backup has been created at T1 at host 1. A failover has been performed to host 2. After a certain point in time, another failover back to host 1 was performed. At the current point in time, host 1 is the primary host.

1. A failure occurred and you must restore to the backup created at T1 at host 1.
2. The secondary host (host 2) is shut down, but no restore operation is executed.
3. The storage volume of host 1 is restored to the backup created at T1.
4. A forward recovery is performed with logs from host 1 and host 2.
5. Host 2 is started, and a system replication resynchronization of host 2 is automatically started.

Figure 24) SnapCenter restore of the valid backup only.

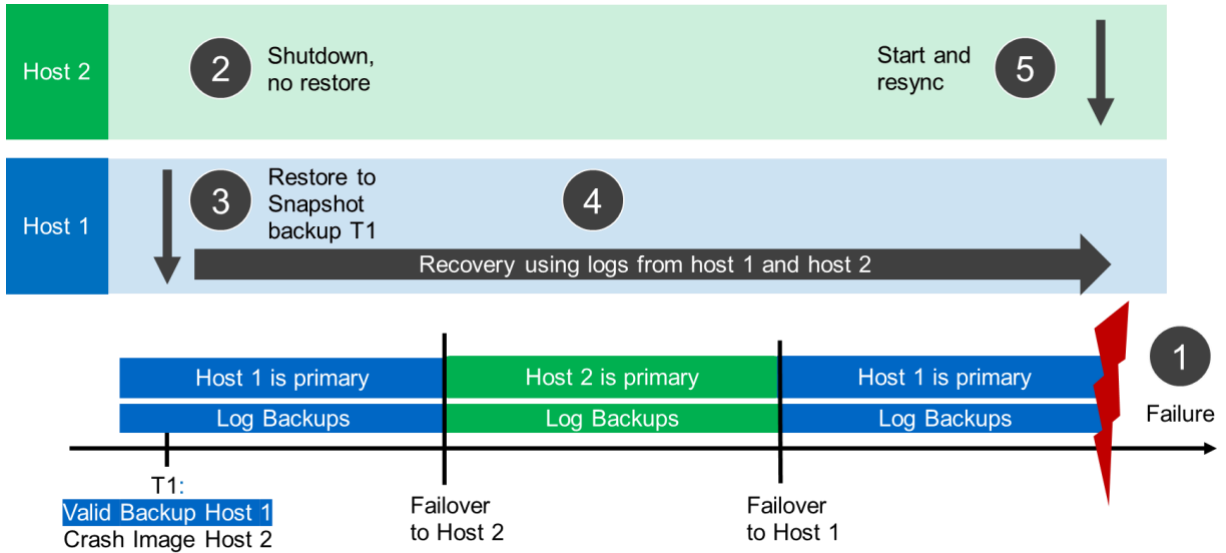
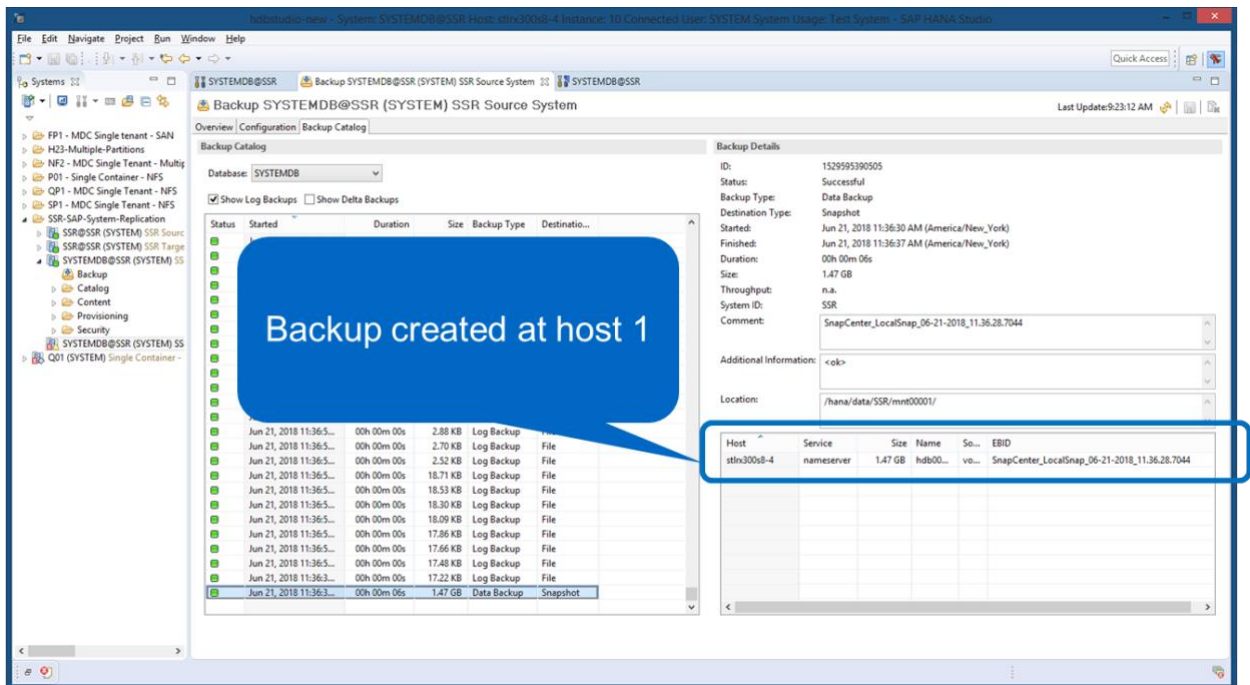


Figure 25 shows the SAP HANA backup catalog in SAP HANA Studio. The highlighted backup shows the backup created at T1 at host 1.

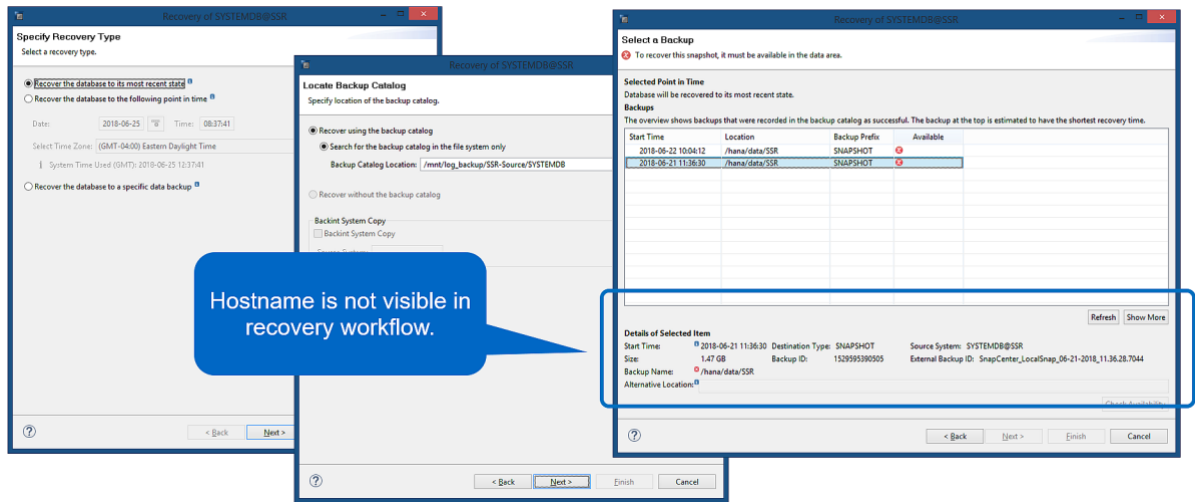
Figure 25) SAP HANA Studio before the restore operation.



A restore and recovery operation is started in SAP HANA Studio. As Figure 26 shows, the name of the host where the backup was created is not visible in the restore and recovery workflow.

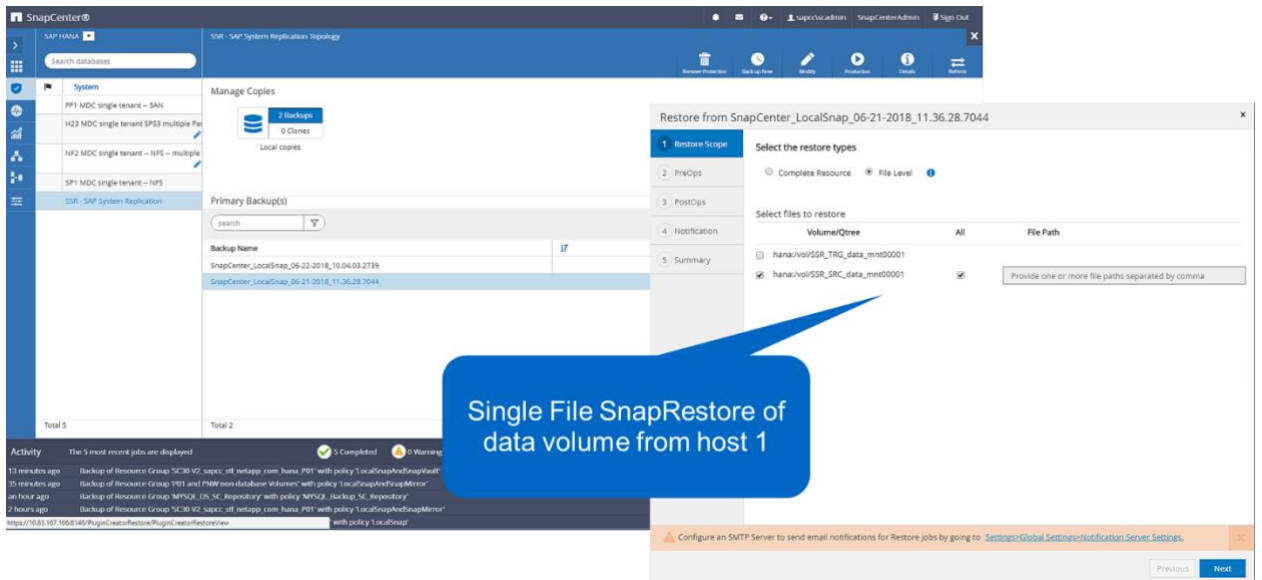
**Note:** In our test scenario, we were able to identify the correct backup (the backup created at host 1) in SAP HANA Studio when the database was still online. If the database is not available, you must check the SnapCenter backup job log to identify the right backup.

Figure 26 Restore and recovery with SAP HANA Studio.



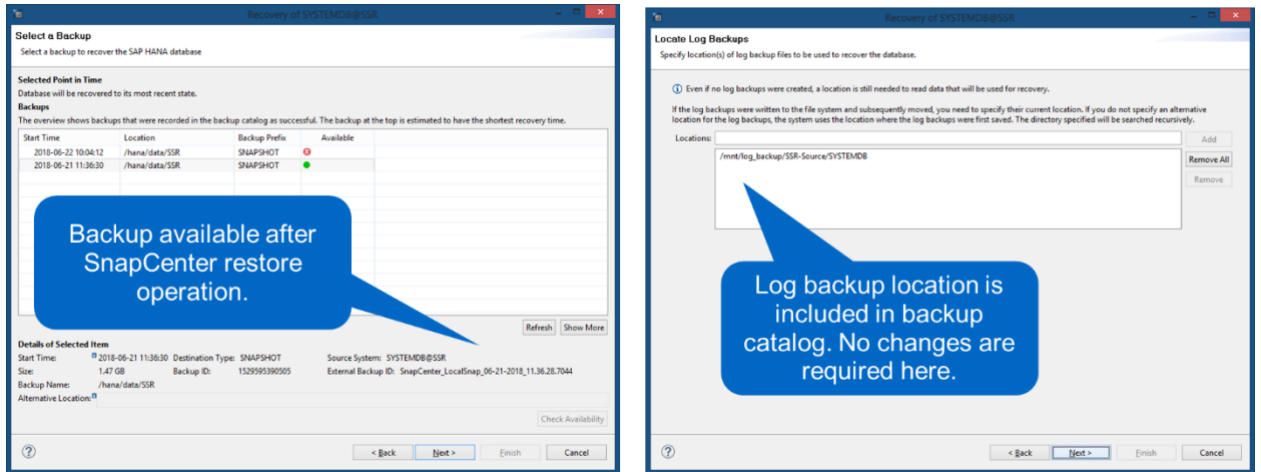
In SnapCenter, the backup is selected and a file-level restore operation is performed. On the file-level restore screen, only the host 1 volume is selected so that only the valid backup is restored.

Figure 27) File-level restore with SnapCenter.



After the restore operation, the backup is highlighted in green in SAP HANA Studio. You don't have to enter an additional log backup location, because the file path of log backups of host 1 and host 2 are included in the backup catalog.

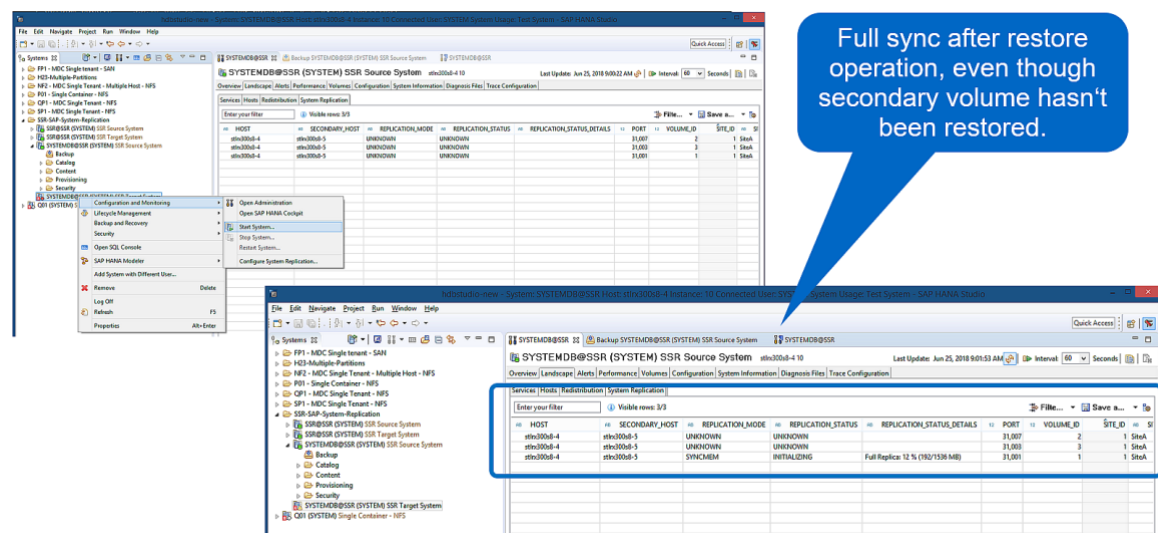
Figure 28) Backup and log backup selection.



After forward recovery has finished, the secondary host (host 2) is started and SAP HANA System Replication resynchronization is started.

**Note:** Even though the secondary host is up-to-date (no restore operation was performed for host 2), SAP HANA executes a full replication of all data. This behavior is standard after a restore and recovery operation with SAP HANA System Replication.

Figure 29) Start of secondary host and resynchronization.



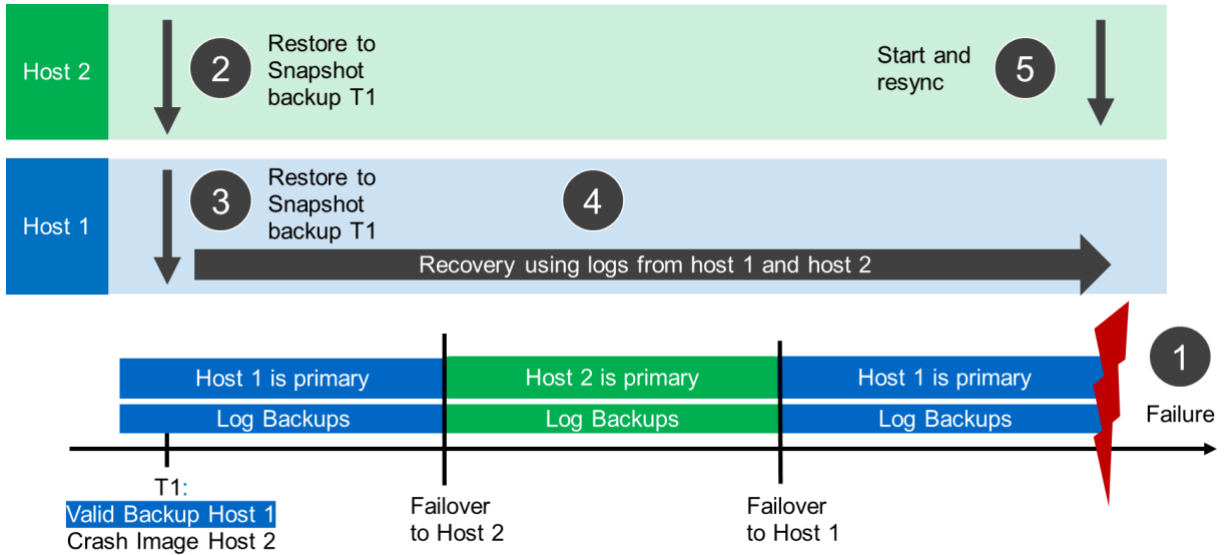
## SnapCenter Restore of Valid Backup and Crash Image

Figure 30 shows an overview of the restore and recovery scenario described in this section.

A backup has been created at T1 at host 1. A failover has been performed to host 2. After a certain point in time, another failover back to host 1 has been performed. At the current point in time, host 1 is the primary host.

1. A failure occurred and you must restore to the backup created at T1 at host 1.
2. The secondary host (host 2) is shut down, and the T1 crash image is restored.
3. The storage volume of host 1 is restored to the backup created at T1.
4. A forward recovery is performed with logs from host 1 and host 2.
5. Host 2 is started, and a system replication resynchronization of host 2 is automatically started.

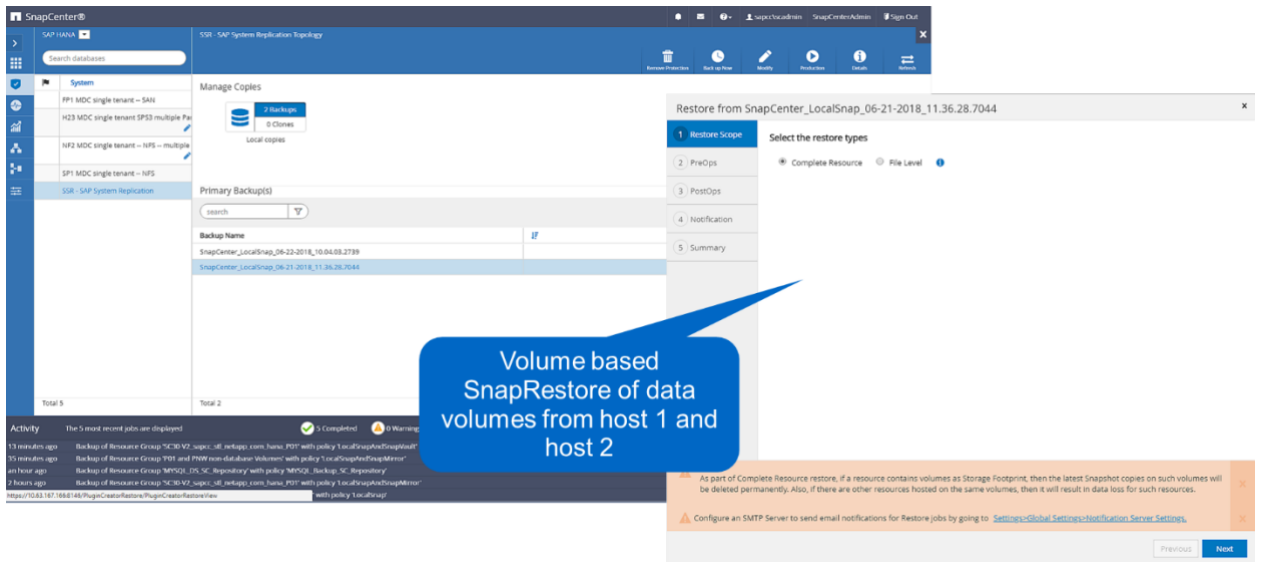
Figure 30) SnapCenter restore of valid backup and crash image.



The restore and recovery operation with SAP HANA Studio is identical to the steps described in the section “SnapCenter Restore of the Valid Backup Only.”

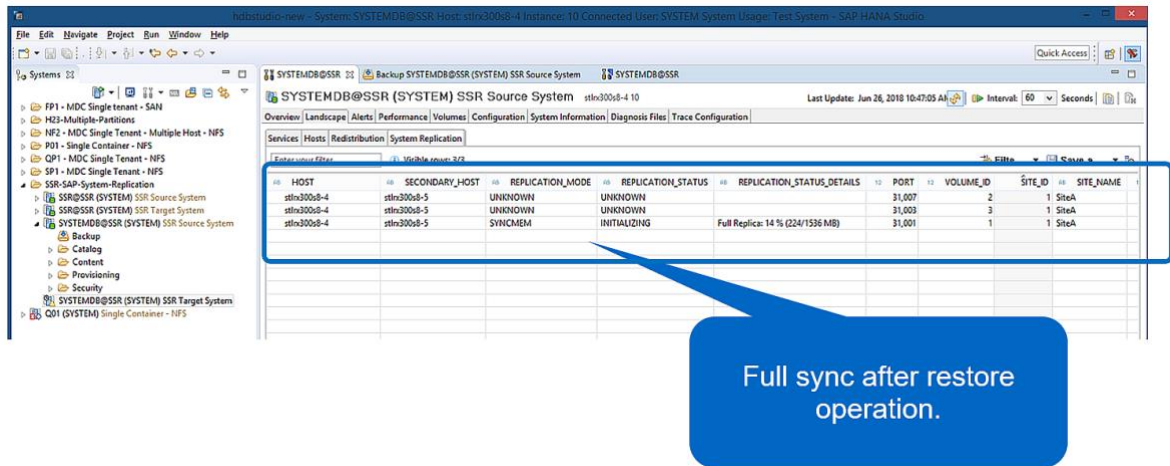
To perform the restore operation, select Complete Resource in SnapCenter. The volumes of both hosts are restored.

Figure 31) Complete resource restore operation.



After forward recovery has been completed, the secondary host (host 2) is started and SAP HANA System Replication resynchronization is started. Full replication of all data is executed.

Figure 32) Start of secondary host and resynchronization.



### 5.4 Restore and Recovery from a Backup Created at the Other Host

A restore operation from a backup that has been created at the other SAP HANA host is a valid scenario for both SnapCenter configuration options.

Figure 33 shows an overview of the restore and recovery scenario described in this section.

A backup has been created at T1 at host 1. A failover has been performed to host 2. At the current point in time, host 2 is the primary host.

1. A failure occurred and you must restore to the backup created at T1 at host 1.
2. The primary host (host 1) is shut down.
3. The backup data T1 of host 1 is restored to host 2.
4. A forward recovery is performed using logs from host 1 and host 2.
5. Host 1 is started, and a system replication resynchronization of host 1 is automatically started.

Figure 33) Restore and recovery from a backup created at the other host.

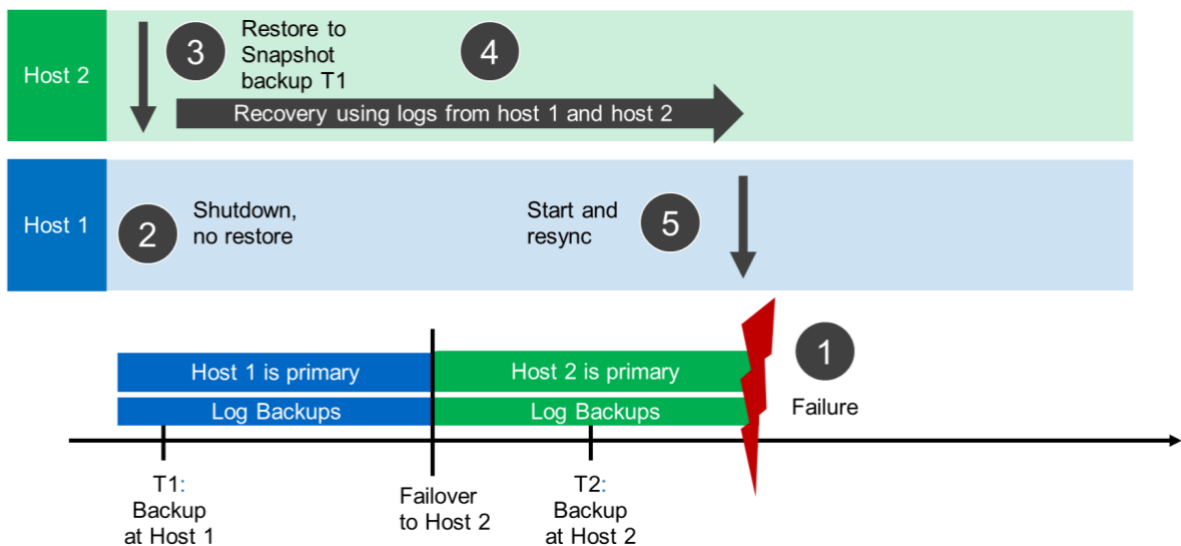
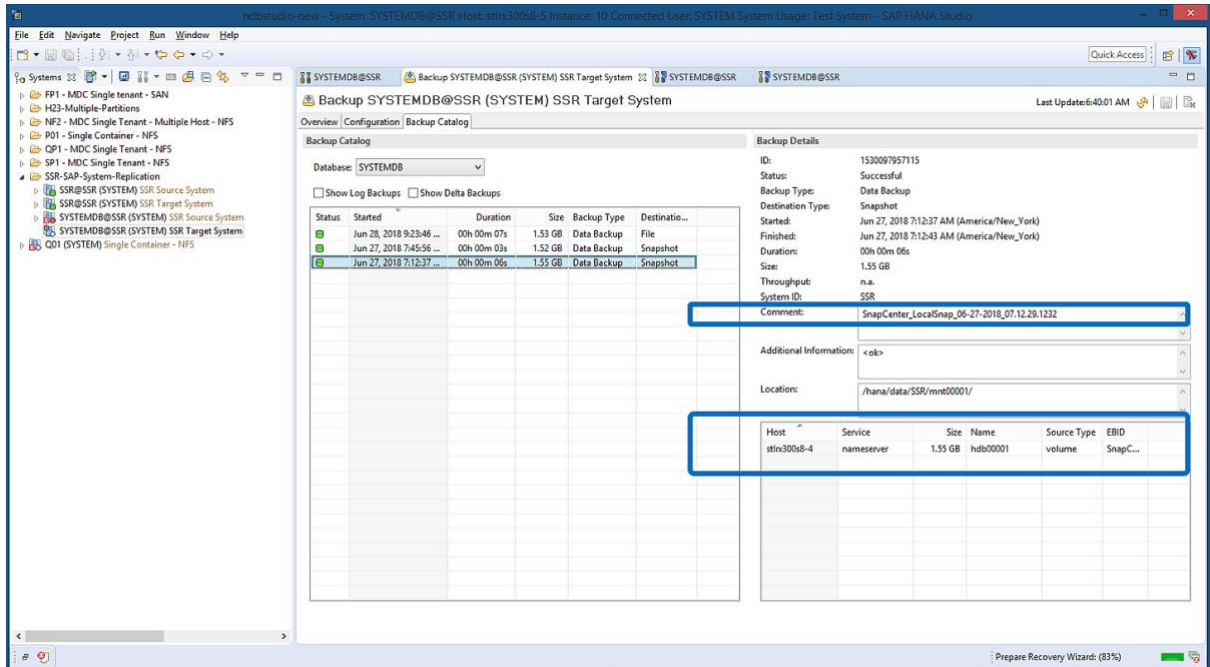


Figure 34 shows the SAP HANA backup catalog and highlights the backup, created at host 1, that was used for the restore and recovery operation.

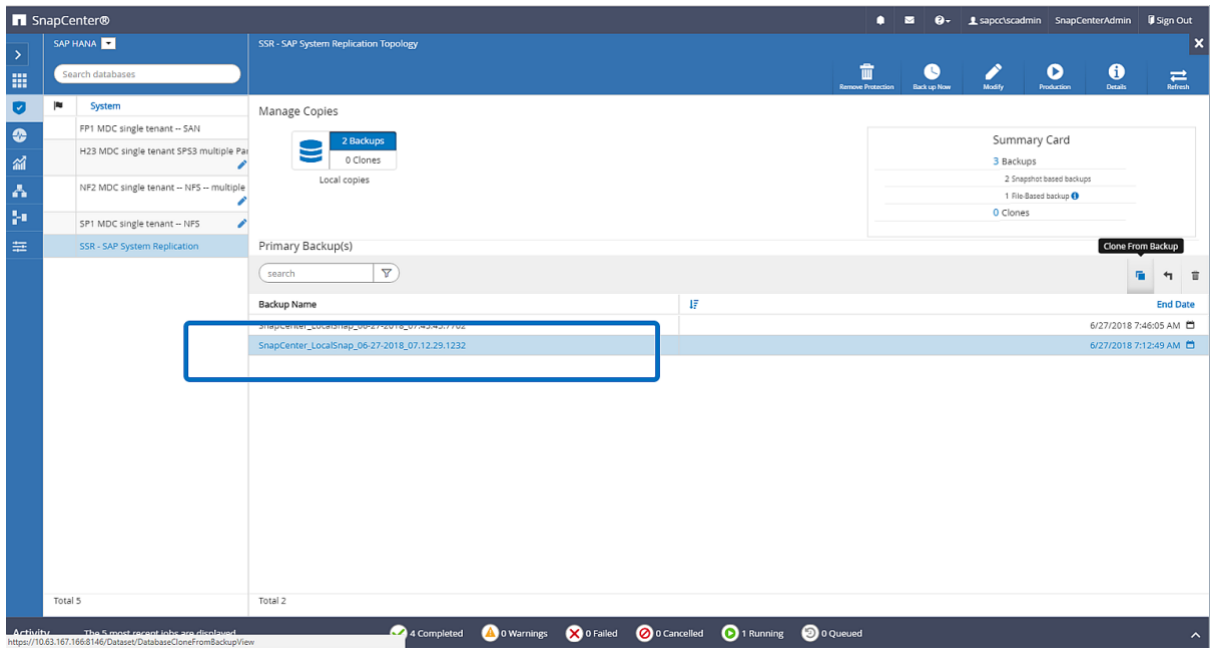
Figure 34) SAP HANA Studio before restore operation.



The restore operation involves the following steps:

1. Create a clone from the backup created at host 1.
2. Mount the cloned volume at host 2.
3. Copy the data from the cloned volume to the original location.

In SnapCenter, the backup is selected and the clone operation is started.



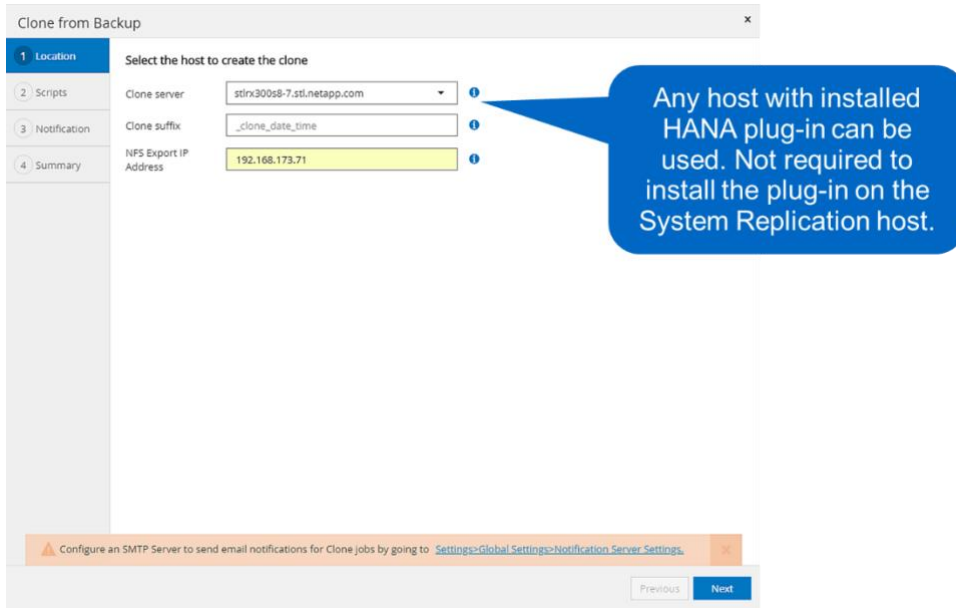
You must provide the clone server and the NFS export IP address.

**Note:** In a SnapCenter single-resource configuration, the SAP HANA plug-in is not installed at the database host. To execute the SnapCenter clone workflow, any host with an installed HANA plug-in can be used as a clone server.



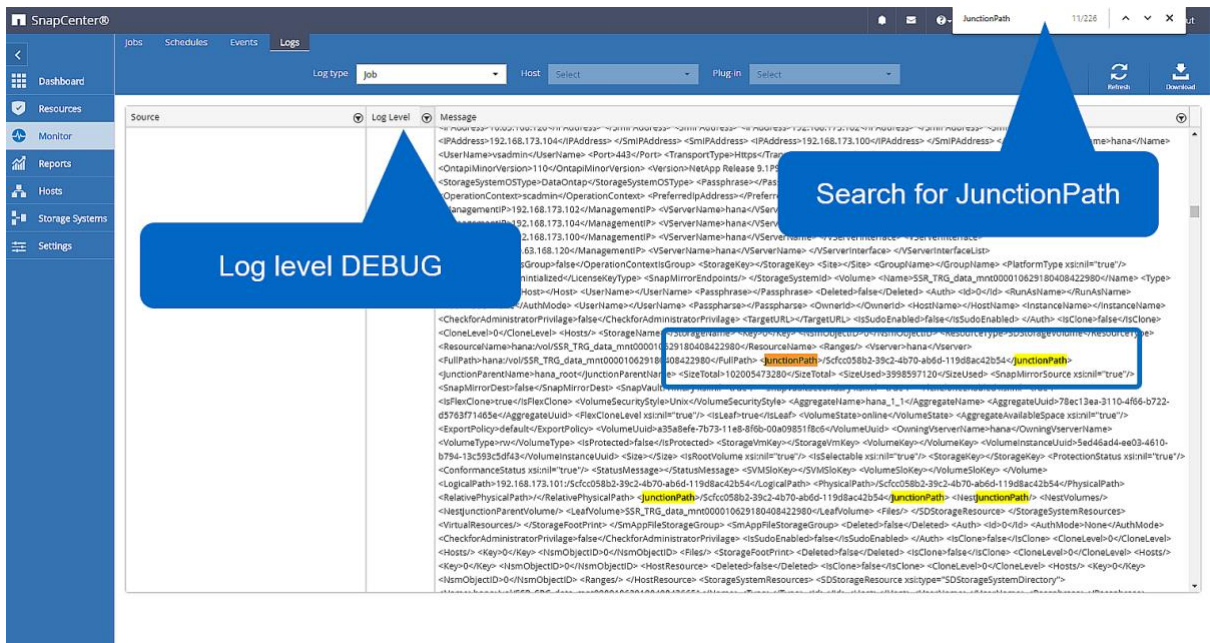
In a SnapCenter configuration with separate resources, the HANA database host is selected as a clone server, and a mount script is used to mount the clone to the target host.

Figure 35) SnapCenter clone workflow.



To determine the junction path that is required to mount the cloned volume, check the job log of the cloning job, as Figure 36 shows.

Figure 36) Junction path information.



The cloned volume can now be mounted.

```
stlr300s8-5:/mnt/tmp # mount 192.168.173.101:/Sc373da37-00ff-4694-b1e1-8153dbd46caf /mnt/tmp
```

The cloned volume contains the data of the HANA database.

```
stlr300s8-5:/mnt/tmp/# ls -al
drwxr-x--x 2 ssradm sapsys 4096 Jun 27 11:12 hdb00001
drwx----- 2 ssradm sapsys 4096 Jun 21 09:38 hdb00002.00003
```

```
drwx----- 2 ssradm sapsys 4096 Jun 27 11:12 hdb00003.00003
-rw-r--r-- 1 ssradm sapsys   22 Jun 27 11:12 nameserver.lck
```

The data is copied to the original location.

```
stlrx300s8-5:/mnt/tmp # cp -Rp hdb00001 /hana/data/SSR/mnt00001/
stlrx300s8-5:/mnt/tmp # cp -Rp hdb00002.00003/ /hana/data/SSR/mnt00001/
stlrx300s8-5:/mnt/tmp # cp -Rp hdb00003.00003/ /hana/data/SSR/mnt00001/
```

The recovery with SAP HANA Studio is performed as described in the section “SnapCenter Restore of the Valid Backup Only.”

## Where to Find Additional Information

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

- SAP HANA Backup and Recovery with SnapCenter  
<https://www.netapp.com/us/media/tr-4614.pdf>
- Automating SAP System Copies Using the SnapCenter 4.0 SAP HANA Plug-In  
<https://www.netapp.com/us/media/tr-4667.pdf>
- SAP HANA Disaster Recovery with Asynchronous Storage Replication  
<https://www.netapp.com/us/media/tr-4646.pdf>
- SAP HANA on VMware vSphere with NetApp FAS and All Flash FAS Systems  
<https://www.netapp.com/us/media/tr-4338.pdf>

## Version History

| Version     | Date    | Document Version History |
|-------------|---------|--------------------------|
| Version 1.0 | October | Initial version          |

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 © 2018 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

### **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.