



# **Automating SAP HANA System Copy and Clone Operations with SnapCenter**

NetApp solutions for SAP

NetApp  
August 19, 2025

# Table of Contents

Automating SAP HANA System Copy and Clone Operations with SnapCenter .....	1
TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter .....	1
Introduction .....	1
SAP system copy, refresh, and clone scenarios .....	2
Use cases for system refresh and cloning .....	3
Data refresh of QA, test, sandbox, or training systems .....	3
Address logical corruption .....	5
Disaster recovery testing .....	5
Supported infrastructure and scenarios .....	6
Overview of SAP system refresh workflow with SnapCenter .....	7
Overview of SAP system clone workflow with SnapCenter .....	9
Considerations for SAP HANA system refresh operations using storage snapshot backups .....	10
Tenant name(s) at target system .....	10
System refresh workflow with enabled SAP HANA encryption .....	12
Automation example scripts .....	15
Script sc-system-refresh.sh .....	16
Script sc-mount-volume.sh .....	17
Script to get SnapCenter environment variables .....	18
SAP HANA system refresh with SnapCenter .....	18
Prerequisites and limitations .....	20
Lab setup .....	20
Initial one-time preparation steps .....	21
Cloning from primary storage with tenant name equal to SID .....	22
Cloning from off-site backup storage .....	29
Cloning a SAP HANA system with multiple tenants .....	31
Clone delete operation .....	33
SAP HANA system refresh with clone split operation .....	36
SnapCenter workflow automation with PowerShell scripts .....	39
SAP system clone with SnapCenter .....	43
Prerequisites and limitations .....	44
Lab setup .....	45
Target host preparation .....	45
Cloning the HANA shared volume .....	48
Cloning additional SAP application services .....	52
Cloning the data volume and recovery of the HANA database .....	52
Where to find additional information and version history .....	56

# Automating SAP HANA System Copy and Clone Operations with SnapCenter

## TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter

In today's dynamic business environment, companies must provide ongoing innovation and react quickly to changing markets. Under these competitive circumstances, companies that implement greater flexibility in their work processes can adapt to market demands more effectively.

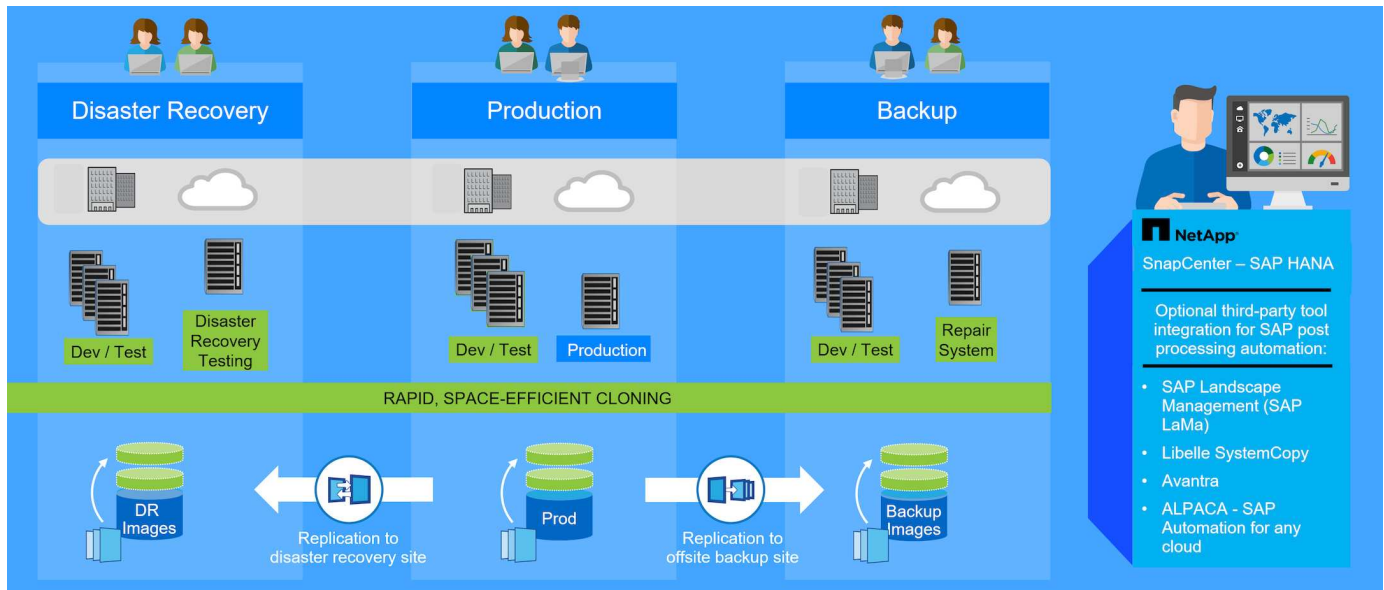
Author: Nils Bauer, NetApp

### Introduction

Changing market demands also affect a company's SAP environments such that they require regular integrations, changes, and updates. IT departments must implement these changes with fewer resources and over shorter time periods. Minimizing risk when deploying those changes requires thorough testing and training which require additional SAP systems with actual data from production.

Traditional SAP lifecycle-management approaches to provision these systems are primarily based on manual processes. These manual processes are often error-prone and time-consuming, delaying innovation and the response to business requirements.

NetApp solutions for optimizing SAP lifecycle management are integrated into SAP HANA database and lifecycle management tools, combining efficient application-integrated data protection with the flexible provisioning of SAP test systems, as is shown in the following figure. These solutions are available for SAP HANA running on-premises or running in the cloud on Azure NetApp Files (ANF) or Amazon FSx for NetApp ONTAP (FSx for ONTAP).



## Application-integrated Snapshot backup operations

The ability to create application-consistent Snapshot backups on the storage layer is the foundation for the system copy and system clone operations described in this document. Storage-based Snapshot backups are created by using the NetApp SnapCenter Plug-In for SAP HANA and interfaces provided by the SAP HANA database. SnapCenter registers Snapshot backups in the SAP HANA backup catalog so that the backups can be used for restore and recovery as well as for cloning operations.

## Off-site backup and/or disaster recovery data replication

Application-consistent Snapshot backups can be replicated on the storage layer to an off-site backup site or a disaster recovery site controlled by SnapCenter. Replication is based on changed and new blocks and is therefore space and bandwidth efficient.

## Use any Snapshot backup for SAP system copy or clone operations

NetApp technology and software integration allows you to use any Snapshot backup of a source system for an SAP system copy or clone operation. This Snapshot backup can be either selected from the same storage that is used for the SAP production systems, the storage that is used for off-site backups, or the storage at the disaster recovery site. This flexibility allows you to separate development and test systems from production if required and covers other scenarios, such as the testing of disaster recovery at the disaster recovery site.



Cloning from the off-site backup or disaster recovery storage is supported for on-premises NetApp systems and for Amazon FSx for NetApp ONTAP. With Azure NetApp Files clones can only be created at the source volume.

## Automation with integration

There are various scenarios and use cases for the provisioning of SAP test systems, and you might also have different requirements for the level of automation. NetApp software products for SAP integrate into database and lifecycle management products from SAP to support different scenarios and levels of automation.

NetApp SnapCenter with the plug-in for SAP HANA is used to provision the required storage volumes based on an application-consistent Snapshot backup and to execute all required host and database operations up to a started SAP HANA database. Depending on the use case, SAP system copy, system clone, system refresh, or additional manual steps such as SAP postprocessing might be required. More details are covered in the next section.

A fully automated, end-to-end provision of SAP test systems can be performed by using third-party tools and integration of NetApp features. More details are available at:

[TR-4953: NetApp SAP Landscape Management Integration using Ansible](#)

[TR-4929: Automating SAP system copy operations with Libelle SystemCopy \(netapp.com\)](#)

[Automating SAP system copy, refresh, and clone workflows with ALPACA and NetApp SnapCenter](#)

[Automating SAP system copy, refresh, and clone workflows with Avantra and NetApp SnapCenter](#)

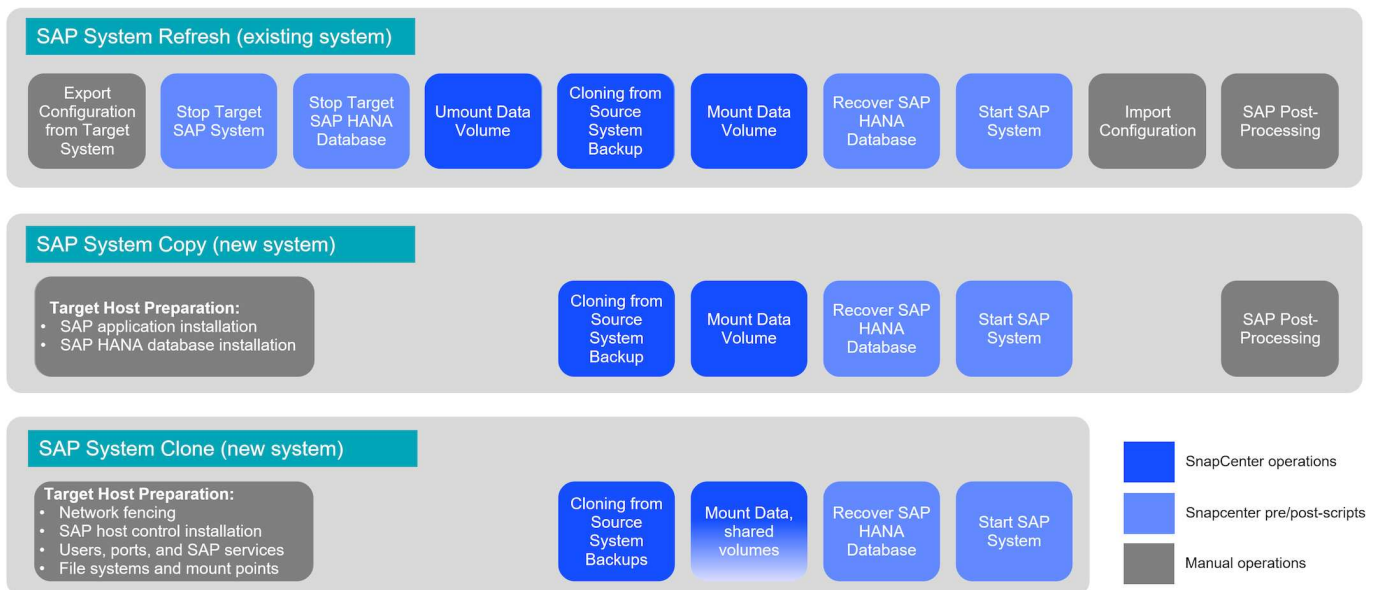
## SAP system copy, refresh, and clone scenarios

The term SAP system copy is often used as a synonym for three different processes: SAP system refresh, SAP system copy, or SAP system clone operations. It is important to

distinguish between the different operations because the workflows and use cases differ for each one.

- **SAP system refresh.** An SAP system refresh is a refresh of an existing target SAP system with data from a source SAP system. The target system is typically part of an SAP transport landscape, for example a quality assurance system, that is refreshed with data from the production system. The hostname, instance number, and SID are different for the source and target systems.
- **SAP system copy.** An SAP system copy is a setup of a new target SAP system with data from a source SAP system. The new target system could be, for example, an additional test system with data from the production system. The hostname, instance number, and SID are different for the source and target systems.
- **SAP system clone.** An SAP system clone is an identical clone of a source SAP system. SAP system clones are typically used to address logical corruption or to test disaster recovery scenarios. With a system clone operation, the hostname, instance number, and SID remain the same. It is therefore important to establish proper network fencing for the target system to make sure that there is no communication with the production environment.

The figure below illustrates the main steps that must be performed during a system refresh, system copy, or system clone operation. The blue boxes indicate steps that can be automated with SnapCenter, while the gray boxes indicate steps that must be performed outside of SnapCenter, either manually or by using third-party tools.



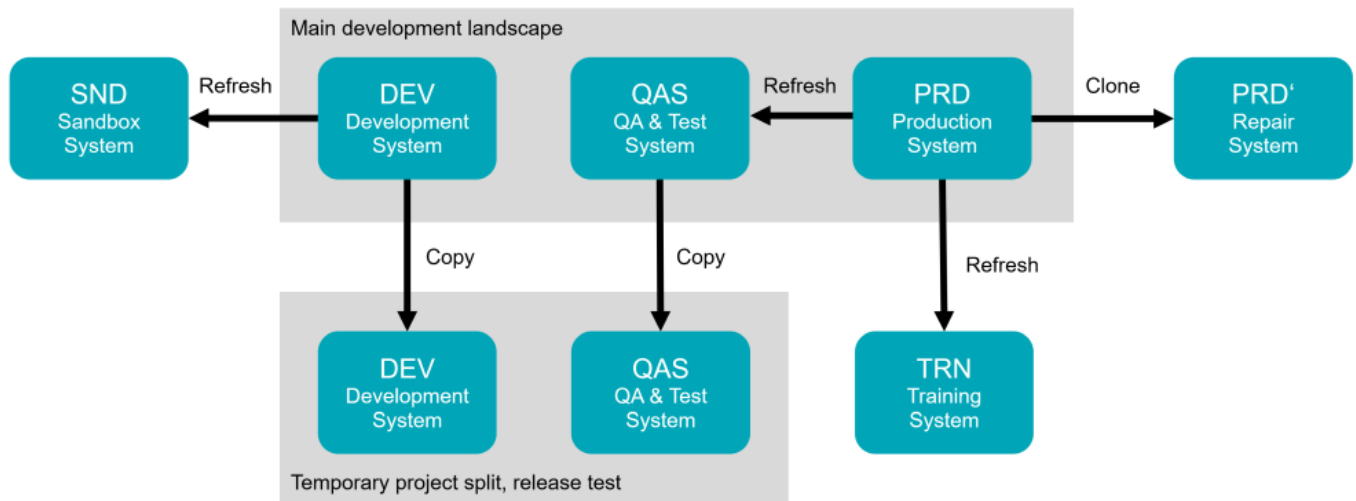
## Use cases for system refresh and cloning

NetApp solutions for optimizing SAP lifecycle management are integrated into SAP HANA database and lifecycle management tools, combining efficient application-integrated data protection with the flexible provisioning of SAP test systems.

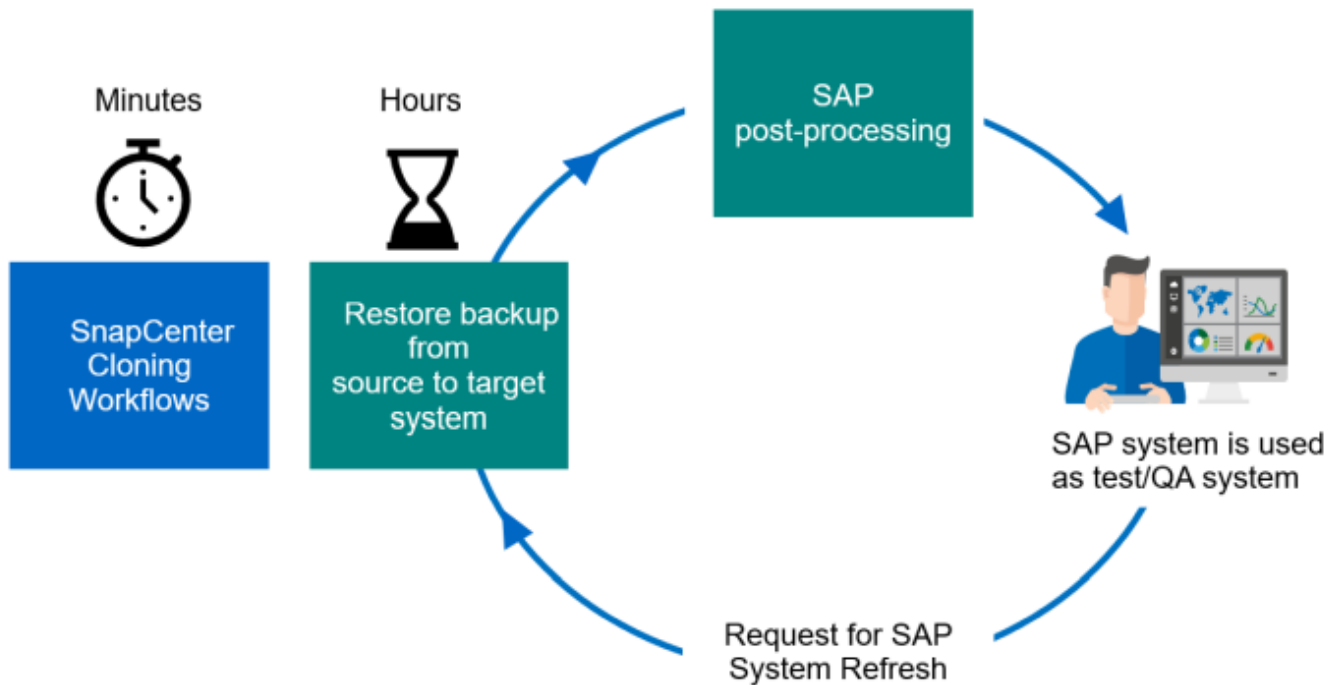
### Data refresh of QA, test, sandbox, or training systems

There are multiple scenarios in which data from a source system must be made available to a target system for testing or training purposes. These test and training systems must be updated with data from the source system on a regular basis to make sure that testing and training is performed with the current data set. These

system refresh operations consist of multiple tasks on the infrastructure, database, and application layers, and they can take multiple days depending on the level of automation.



SnapCenter cloning workflows can be used to accelerate and automate the required tasks at the infrastructure and database layers. Instead of restoring a backup from the source system to the target system, SnapCenter uses NetApp Snapshot copy and NetApp FlexClone technology, so that required tasks up to a started SAP HANA database can be performed in minutes instead of hours. The time needed for the cloning process is independent from the size of the database, therefore even very large systems can be created in a couple of minutes. The startup time just depends on the database size and the connectivity between the database server and the storage system.



The workflow for system-refresh operations is described in the section [“SAP HANA system refresh with SnapCenter.”](#)

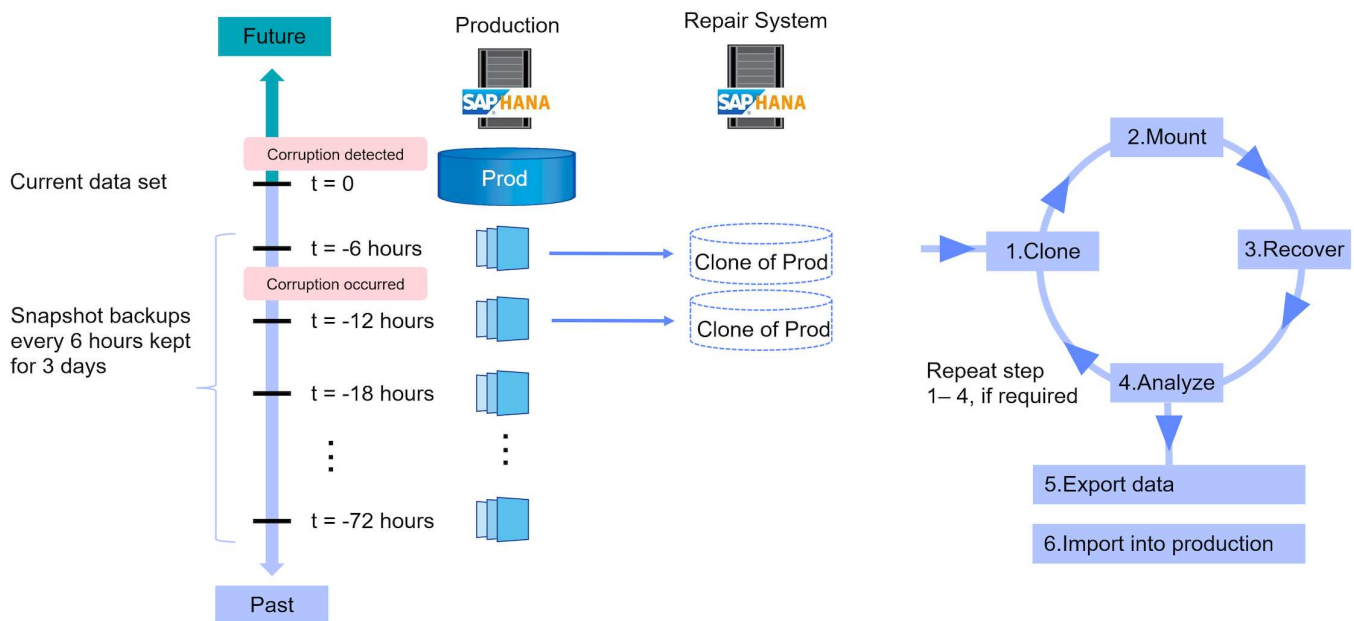
## Address logical corruption

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, minimal downtime and maximum data loss requirements can sometimes not be fulfilled.

The worst case is 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. 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 production system does not need to be stopped, and, in the best-case scenario, no data or only a small fraction of data is lost.

When setting up the repair system, flexibility and agility is crucial. When using 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.



The workflow of the repair system creation is described in the section [“SAP system clone with SnapCenter.”](#)

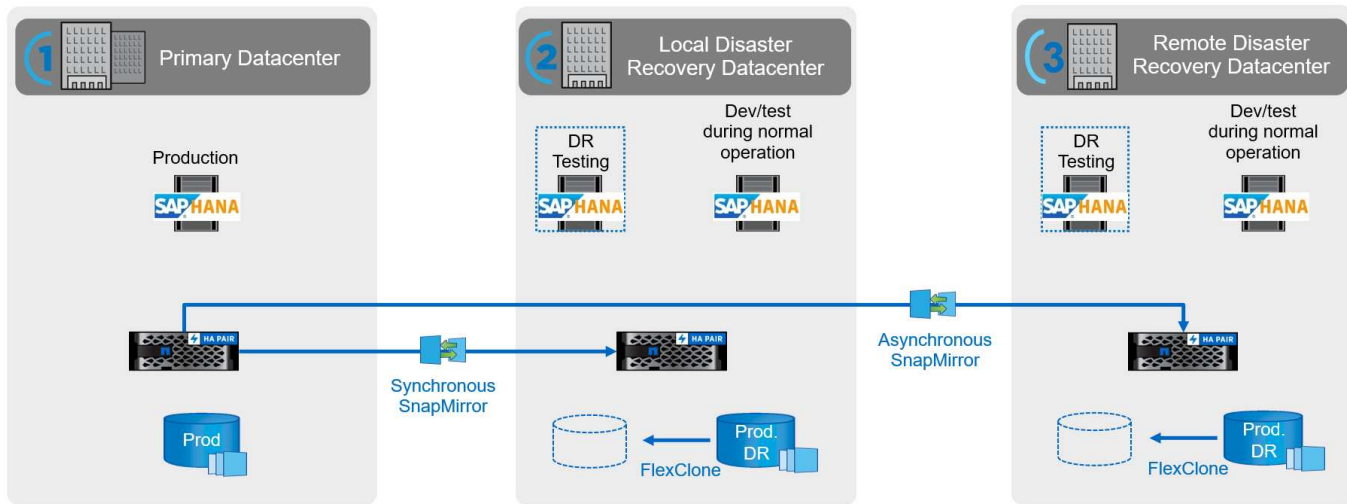
## Disaster recovery testing

An effective disaster recovery strategy needs testing the required workflow. Testing demonstrates whether the strategy works and whether the internal documentation is sufficient. It also allows administrators to train 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 performed without interrupting data replication.

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



A detailed step-by-step description can be found in the technical reports

[TR-4646: SAP HANA Disaster Recovery with Storage Replication \(netapp.com\)](#)

[TR-4891: SAP HANA disaster recovery with Azure NetApp Files](#)

## Supported infrastructure and scenarios

This document covers SAP system refresh and cloning scenarios for SAP HANA systems running on on-premises NetApp systems, on Amazon FSx for NetApp ONTAP systems and on Azure NetApp Files. However not all features and scenarios are available on every storage platform. The table below summarizes the supported configurations.

Within the document, we are using an SAP HANA landscape running on on-premises NetApp systems with NFS as the storage protocol. Most workflow steps are identical across the different platforms, and if there are differences, they are highlighted in this document.

	On-premises NetApp systems	AWS FSx for NetApp ONTAP	Azure NetApp Files
Storage protocol	NFS, Fibre Channel	NFS	NFS
Thin clone (FlexClone)	Yes	Yes	No, with the current ANF version, cloned volume is always split
Clone split operation	Yes	Yes	N/A
Cloning from primary	Yes	Yes	Yes
Cloning from off-site backup	Yes	Yes	No
Cloning at DR site	Yes	Yes	Yes, but not integrated into SnapCenter

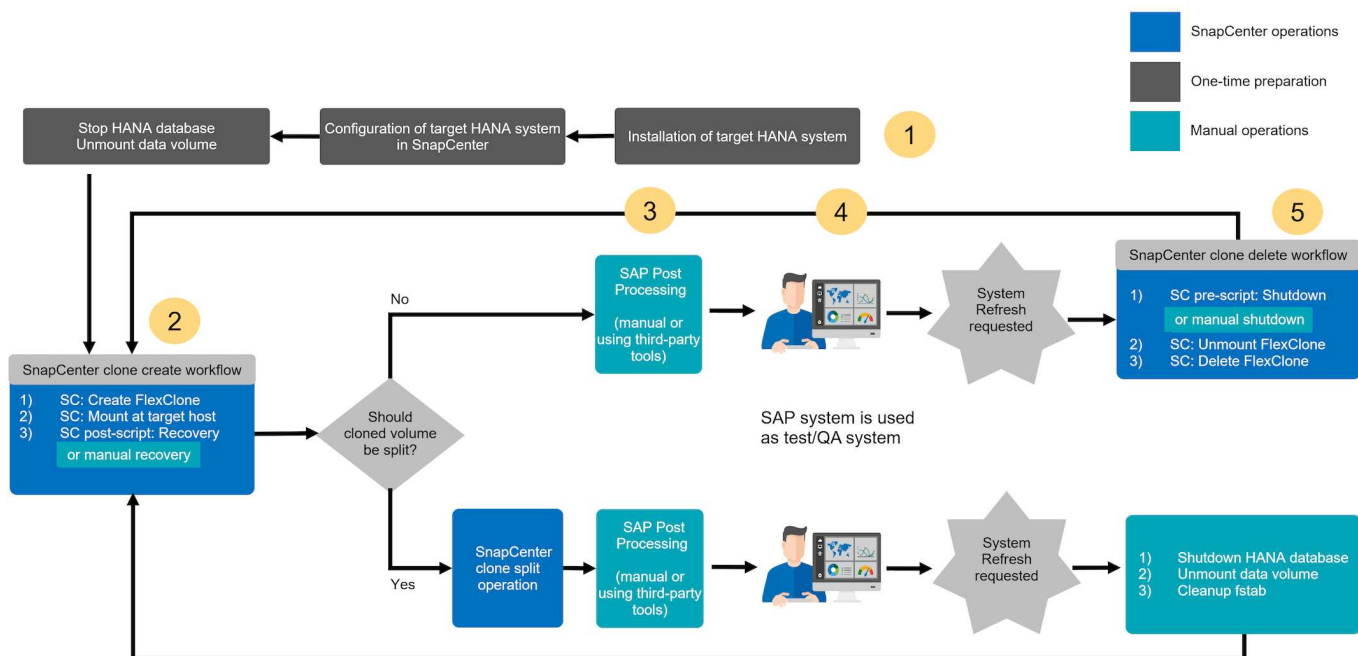
# Overview of SAP system refresh workflow with SnapCenter

SnapCenter provides workflows that allow you to manage clones of data sets from any existing Snapshot backup. This cloned data set, a FlexClone volume, can be used to rapidly provision a HANA data volume from a source system and attach it to a target system. It is therefore a perfect fit for executing system refresh operations for QA, test, sandbox, or training systems.

The SnapCenter cloning workflows handle all required operations on the storage layer and can be extended using scripts to execute host-specific and HANA database-specific operations. In this document, we use a script to perform HANA database recovery and shutdown operations. SnapCenter workflows with further automation using the script handle all required HANA database operations but do not cover any required SAP post-processing steps. SAP post processing must be performed manually or with third-party tools.

The SAP system refresh workflow with SnapCenter consists of five main steps as shown in the below figure.

1. A one-time, initial installation and preparation of the target system
  - a. The SnapCenter HANA plugin must be installed on the new target system and the HANA system must be configured in SnapCenter
  - b. The target system must be stopped, and the HANA data volume must be unmounted
2. The SnapCenter clone create workflow
  - a. SnapCenter creates a FlexClone volume of the selected Snapshot of the source system
  - b. SnapCenter mounts the FlexClone volume at the target system
  - c. Recovery of the target HANA database can be automated using the `sc-system-refresh` script as a post-script or can be executed manually
3. SAP post processing (manual or with a third-party tool)
4. The system can now be used as test/QA system.
5. When a new system refresh is requested, the SnapCenter clone delete workflow is used to remove the FlexClone volume
  - a. If the target HANA system has been protected in SnapCenter, the protection must be removed before the clone delete workflow is started.
  - b. The HANA system must be stopped manually or stopped automatically using the `sc-system-refresh` script as a SnapCenter pre-script
  - c. SnapCenter unmounts the HANA data volume
  - d. SnapCenter deletes the FlexClone volume
  - e. A refresh is restarted with step 2.



In most cases, target test/QA systems are used for at least a couple of weeks. Since the FlexClone volume is blocking the Snapshot of the source system volume, this Snapshot will require additional capacity based on the block change rate at the source system volume. For production source systems and an average change rate of 20% per day, the blocked Snapshot will reach 100% after 5 days. Therefore, NetApp recommends splitting the FlexClone volume either immediately or after a couple of days, if the clone is based on a production source system. The clone split operation does not block use of the cloned volume and can therefore be performed at any time while the HANA database is in use.



When splitting the FlexClone volume, SnapCenter deletes all backups that were created at the target system.



With SnapCenter and Azure NetApp Files, the clone split operation is not available, since Azure NetApp Files always splits the clone after creation.

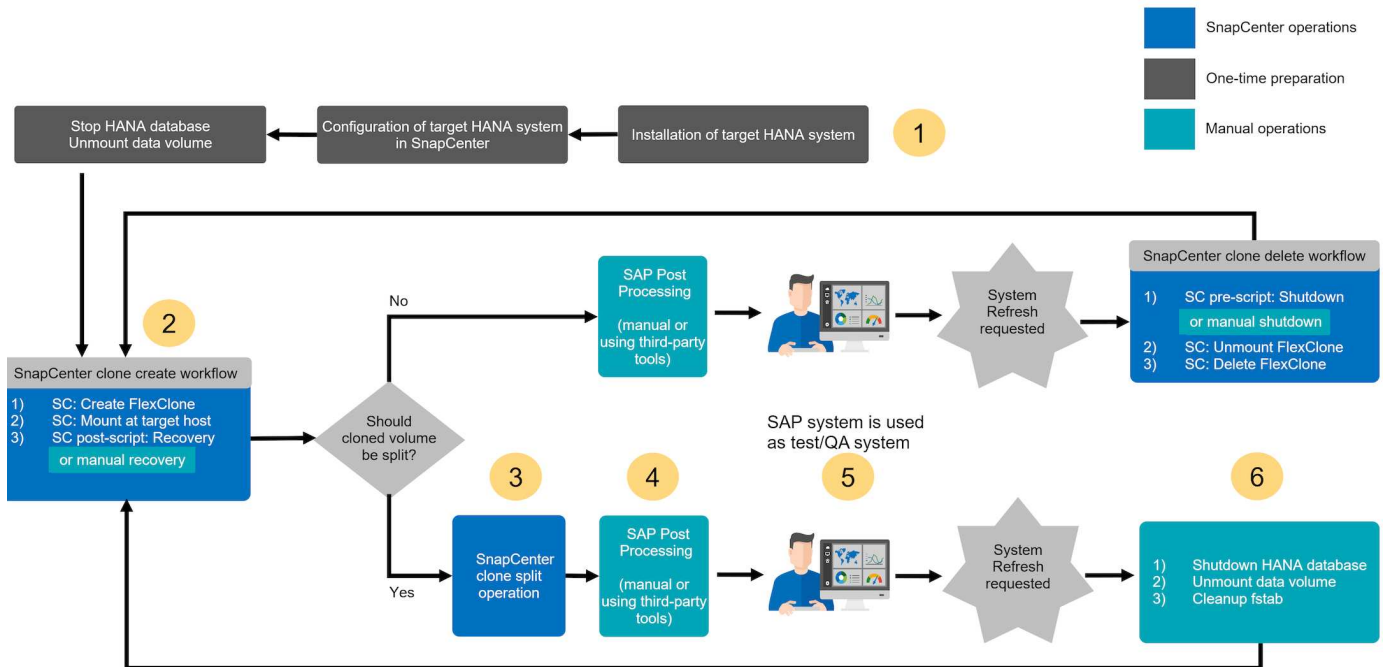
The refresh operation including the clone split consists of the following steps.

- A one-time, initial installation and preparation of the target system
  - The SnapCenter HANA plugin must be installed on the new target system and the HANA system must be configured in SnapCenter
  - The target system must be stopped, and the HANA data volume must be unmounted
- The SnapCenter clone create workflow
  - SnapCenter creates a FlexClone volume of the selected Snapshot of the source system
  - SnapCenter mounts the FlexClone volume at the target system
  - Recovery of the target HANA database can be automated using the `sc-system-refresh` script as a post-script or can be executed manually
- The FlexClone volume is split using the SnapCenter clone split workflow.
- SAP post processing (manual or with a third-party tool)
- The system can now be used as test/QA system.

6. When a new system refresh is requested, the cleanup is done with the following manual steps
  - a. If the target HANA system has been protected in SnapCenter, the protection must be removed.
  - b. The HANA system must be stopped manually
  - c. The HANA data volume must be unmounted and the fstab entry from SnapCenter must be removed (manual task)
  - d. A refresh is restarted with step 2.



The old data volume, which was split previously, must be deleted manually on the storage system.



The section [“SAP HANA system refresh with SnapCenter”](#) provides a detailed step-by-step description of both system-refresh workflows.

## Overview of SAP system clone workflow with SnapCenter

As discussed in the previous section, SnapCenter can manage clones of data sets from any existing Snapshot backup and can rapidly provision these data sets to any target system. The flexible and agile provisioning of production data to a repair system to address logical corruption is critical, since it is often necessary to reset the repair system and to choose a different production data set. FlexClone technology enables a rapid provisioning process, and provides significant capacity savings, since the repair system is typically only used for a short time.

The figure below summarizes the required steps for an SAP system clone operation using SnapCenter.

1. Prepare the target host.
2. SnapCenter clone create workflow for the SAP HANA shared volume.

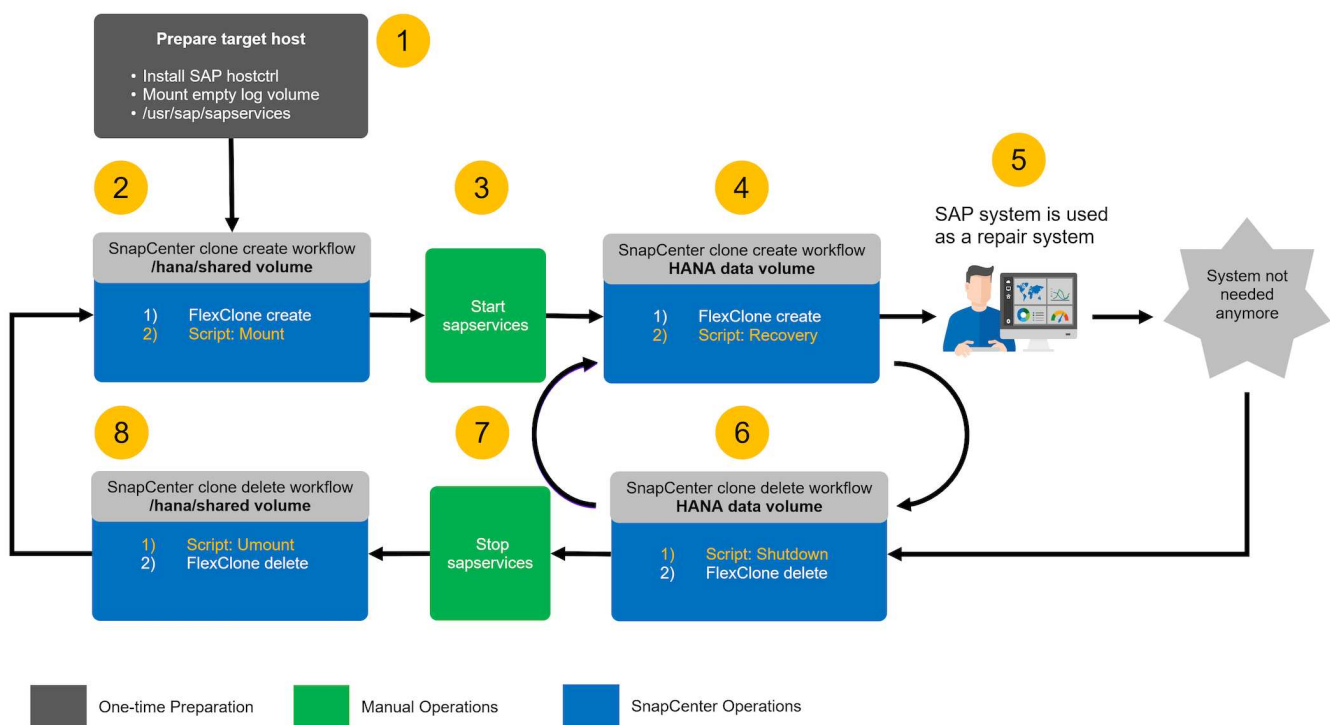
3. Start SAP HANA services.
4. SnapCenter clone create workflow for the SAP HANA data volume including database recovery.
5. The SAP HANA system can now be used as a repair system.

If the system is not needed anymore, the clean-up process is performed with the following steps.

6. SnapCenter clone delete workflow for the SAP HANA data volume including database shutdown (when using the automation script).
7. Stop SAP HANA services.
8. SnapCenter clone delete workflow for the SAP HANA shared volume.



If you must reset the system to a different Snapshot backup, then step 6 and step 4 are sufficient. A refresh of the SAP HANA shared volume is not required.



The section [“SAP system clone with SnapCenter”](#) provides a detailed step-by-step description of the system clone workflow.

## Considerations for SAP HANA system refresh operations using storage snapshot backups

NetApp solutions for optimizing SAP lifecycle management are integrated into SAP HANA database and lifecycle management tools, combining efficient application-integrated data protection with the flexible provisioning of SAP test systems.

### Tenant name(s) at target system

The steps required to perform an SAP HANA system refresh depend on the source system tenant

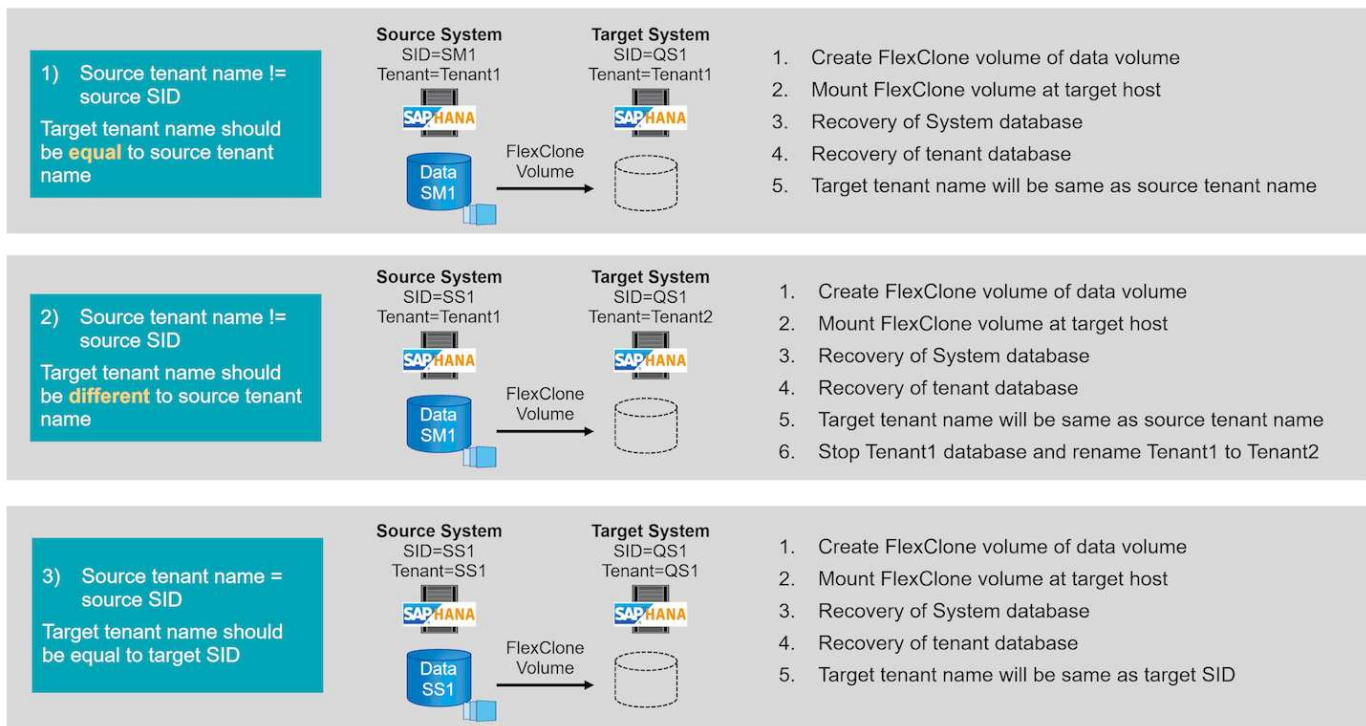
configuration and the required tenant name at the target system, as shown in the figure below.

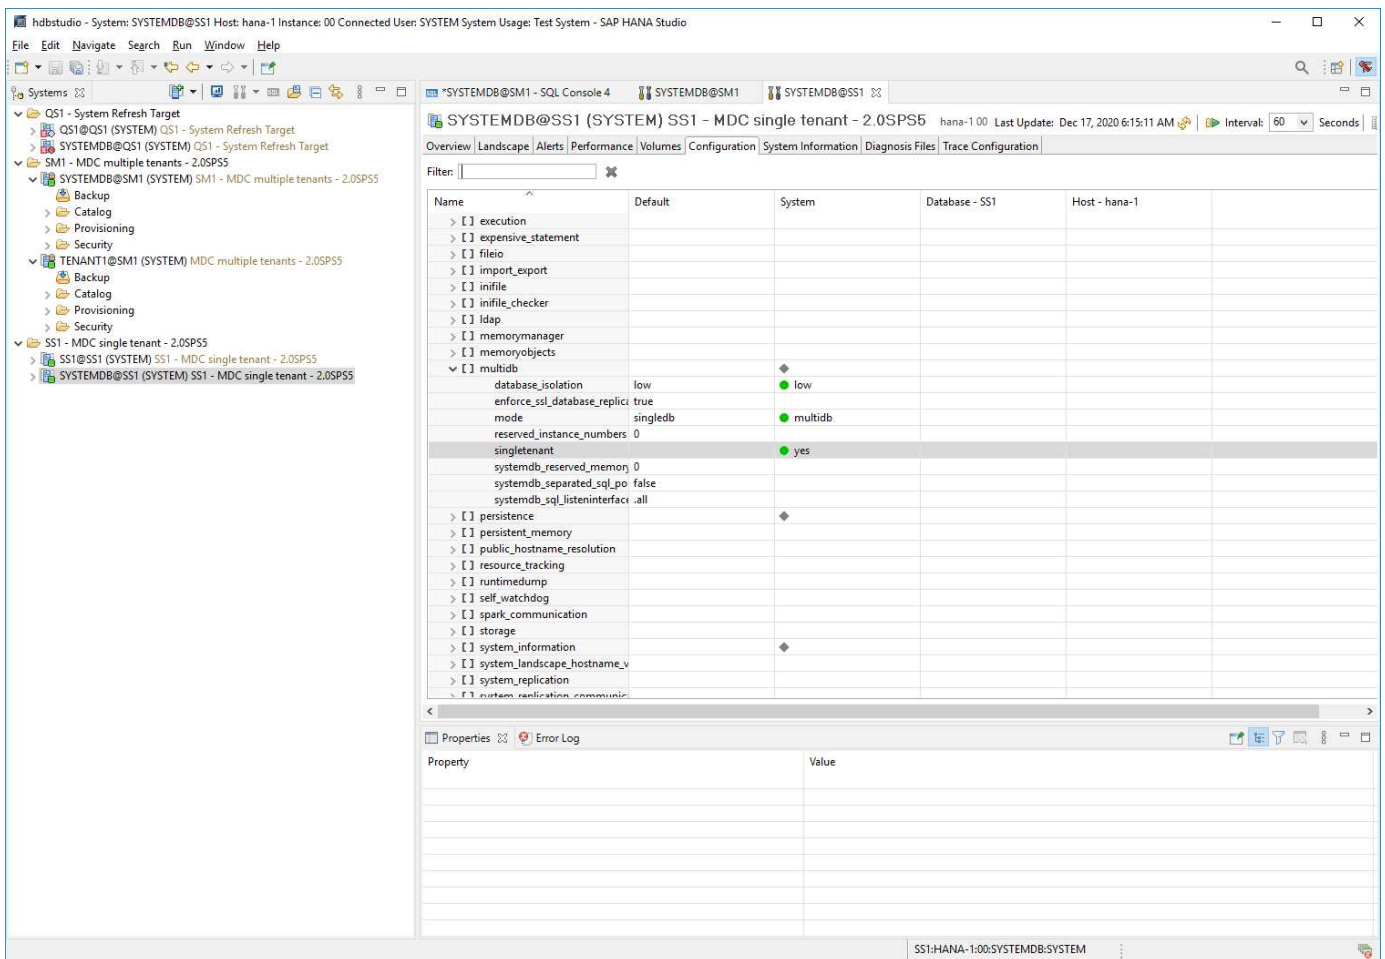
Since the tenant name is configured in the system database, the tenant name of the source system is also available at the target system after the recovery of the system database. Therefore, the tenant at the target system can only be recovered with the same name as the source tenant as shown in option 1. If the tenant name at the target system must be different, the tenant must first be recovered with the same name as the source tenant and then renamed to the required target tenant name. This is option 2.

An exception of this rule is an SAP HANA system with a single tenant, where the tenant name is identical to the system SID. This configuration is the default after an initial SAP HANA installation. This specific configuration is flagged by the SAP HANA database. In this case, tenant recovery at the target system can be executed with the tenant name of the target system, which must be also identical to the system SID of the target system. This workflow is shown in option 3.



As soon as any tenant create, rename, or drop operation is executed at the source system, this configuration flag is deleted by the SAP HANA database. Therefore, even if the configuration has been brought back to tenant = SID, the flag is no longer available and the exception regarding tenant recovery with workflow 3 is no longer possible. In this case, option 2 is the required workflow.





## System refresh workflow with enabled SAP HANA encryption

When SAP HANA persistence encryption is enabled, additional steps are required before you can recover the SAP HANA database at the target system.

At the source system you need to create a backup of the encryption root keys for the system database, as well as for all tenant databases. The backup files must be copied to the target system and the root keys must be imported from the backup before the recovery operation is executed.

See also [SAP HANA Administration Guide](#).

### Backup of root keys

A backup of the root keys is always required, if any changes to the root keys have been made. The backup command requires the dbid as a CLI parameter. The dbid's can be identified using the below SQL statement.

SYSTEMDB@SS1 (SYSTEM) hana-1 00

SQL Result

```

SELECT DATABASE_NAME,
CASE WHEN (DBID = " AND
DATABASE_NAME = 'SYSTEMDB')
THEN 1
WHEN (DBID = " AND
DATABASE_NAME <> 'SYSTEMDB')
THEN 3
ELSE TO_INT(DBID)
END DATABASE_ID
FROM (SELECT DISTINCT DATABASE_NAME, SUBSTR_AFTER (SUBPATH, '.') AS DBID FROM SYS_DATABASES.M_VOLUMES)

```

	DATABASE_NAME	DATABASE_ID
1	SYSTEMDB	1
2	SS1	3

The SQL statement and further documentation is available in the SAP HANA Admin Guide at [Back Up Root Keys | SAP Help Portal](#)

The following steps are illustrating the required operations for a HANA system with a single tenant SS1 and are executed at the source system.

1. Set backup password for system and tenant (SS1) databases (if not done yet).

```

hdbsql SYSTEMDB=> ALTER SYSTEM SET ENCRYPTION ROOT KEYS BACKUP PASSWORD
Netappl23;
0 rows affected (overall time 3658.128 msec; server time 3657.967 msec)
hdbsql SYSTEMDB=>
hdbsql SS1=> ALTER SYSTEM SET ENCRYPTION ROOT KEYS BACKUP PASSWORD
Netappl23;
0 rows affected (overall time 2424.236 msec; server time 2424.010 msec)
hdbsql SS1=>

```

2. Create backup of root keys for system and tenant (SS1) databases.

```

ssladm@hana-1:/usr/sap/SS1/home> /usr/sap/SS1/HDB00/exe/hdbnsutil
-backupRootKeys root-key-backup-SS1-SYSTEMDB.rkb --dbid=1 --type='ALL'
Exporting root key backup for database SYSTEMDB (DBID: 1) to
/usr/sap/SS1/home/root-key-backup-SS1-SYSTEMDB.rkb
done.
ssladm@hana-1:/usr/sap/SS1/home> /usr/sap/SS1/HDB00/exe/hdbnsutil
-backupRootKeys root-key-backup-SS1-SS1.rkb --dbid=3 --type='ALL'
Exporting root key backup for database SS1 (DBID: 3) to
/usr/sap/SS1/home/root-key-backup-SS1-SS1.rkb
done.

```

3. Validate root key backups (optional)

```

ssladm@hana-1:/usr/sap/SS1/home> ls -al root*
-rw-r----- 1 ssladm sapsys 1440 Apr 24 07:00 root-key-backup-SS1-SS1.rkb
-rw-r----- 1 ssladm sapsys 1440 Apr 24 06:54 root-key-backup-SS1-
SYSTEMDB.rkb
ssladm@hana-1:/usr/sap/SS1/home>

ssladm@hana-1:/usr/sap/SS1/home> /usr/sap/SS1/HDB00/exe/hdbnsutil
-validateRootKeysBackup root-key-backup-SS1-SS1.rkb
Please Enter the password:
Successfully validated SSFS backup file /usr/sap/SS1/home/root-key-backup-
SS1-SS1.rkb
done.

ssladm@hana-1:/usr/sap/SS1/home> /usr/sap/SS1/HDB00/exe/hdbnsutil
-validateRootKeysBackup root-key-backup-SS1-SYSTEMDB.rkb
Please Enter the password:
Successfully validated SSFS backup file /usr/sap/SS1/home/root-key-backup-
SS1-SYSTEMDB.rkb
done.

```

### Import of root keys at the target system

The import of the root keys is required initially for the first system refresh operation. If the root keys are not changed at the source system, no additional import is required.

The import command requires the dbid as a CLI parameter. The dbid's can be identified in the same way as described for the root key backup.

1. In our setup the root keys are copied from the source system to an NFS share

```

hana-1:~ # cp /usr/sap/SS1/home/root-key-backup-SS1-SS1.rkb /mnt/sapcc-
share/SAP-System-Refresh/
hana-1:~ # cp /usr/sap/SS1/home/root-key-backup-SS1-SYSTEMDB.rkb
/mnt/sapcc-share/SAP-System-Refresh/

```

2. The root keys can now be imported using hdbnsutil. The dbid for the system and tenant database must be provided with the command. The backup password is also required.

```

qsladm@hana-7:/usr/sap/QS1/HDB11> ./exe/hdbnsutil -recoverRootKeys
/mnt/sapcc-share/SAP-System-Refresh/root-key-backup-SS1-SYSTEMDB.rkb
--dbid=1 --type=ALL
Please Enter the password:
Importing root keys for DBID: 1 from /mnt/sapcc-share/SAP-System-
Refresh/root-key-backup-SS1-SYSTEMDB.rkb
Successfully imported root keys from /mnt/sapcc-share/SAP-System-
Refresh/root-key-backup-SS1-SYSTEMDB.rkb
done.

qsladm@hana-7:/usr/sap/QS1/HDB11> ./exe/hdbnsutil -recoverRootKeys
/mnt/sapcc-share/SAP-System-Refresh/root-key-backup-SS1-SS1.rkb --dbid=3
--type=ALL Please Enter the password:
Importing root keys for DBID: 3 from /mnt/sapcc-share/SAP-System-
Refresh/root-key-backup-SS1-SS1.rkb
Successfully imported root keys from /mnt/sapcc-share/SAP-System-
Refresh/root-key-backup-SS1-SS1.rkb
done.
qsladm@hana-7:/usr/sap/QS1/HDB11>

```

## Root key import, if dbid does not exist at target

As described in the chapter before, the dbid is required to import the root key for the system and all tenant databases. While the system database has always dbid=0, the tenant databases can have different dbid's.

SQL		Result
<pre> SELECT DATABASE_NAME, CASE WHEN (DBID = " AND DATABASE_NAME = 'SYSTEMDB') THEN 1 WHEN (DBID = " AND DATABASE_NAME &lt;&gt; 'SYSTEMDB') THEN 3 ELSE TO_INT(DBID) END DATABASE_ID FROM (SELECT DISTINCT DATABASE_NAME, SUBSTR_AFTER (SUBPATH,':') AS DBID FROM SYS_DATABASES.M_VOLUMES) </pre>		
DATABASE_NAME	DATABASE_ID	
1 TENANT1	4	
2 SYSTEMDB	1	
3 TENANT2	3	

The output above shows two tenants with dbid=3 and dbid=4. If the target system has not yet hosted a tenant with dbid=4, the import of the root key will fail. In that case you need to recover the system database first and then import the key for the tenant with dbid=4.

## Automation example scripts

In this document, two scripts are used to further automate SnapCenter clone create and clone delete operations.

- The script `sc-system-refresh.sh` is used for the system refresh and the system clone workflow to execute recovery and shutdown operations of the SAP HANA database.
- The script `sc-mount-volume.sh` is used for the system clone workflow to execute mount and unmount operations for the SAP HANA shared volume.



The example scripts are provided as is and are not supported by NetApp. You can request the scripts via email to [ng-sapcc@netapp.com](mailto:ng-sapcc@netapp.com).

## Script `sc-system-refresh.sh`

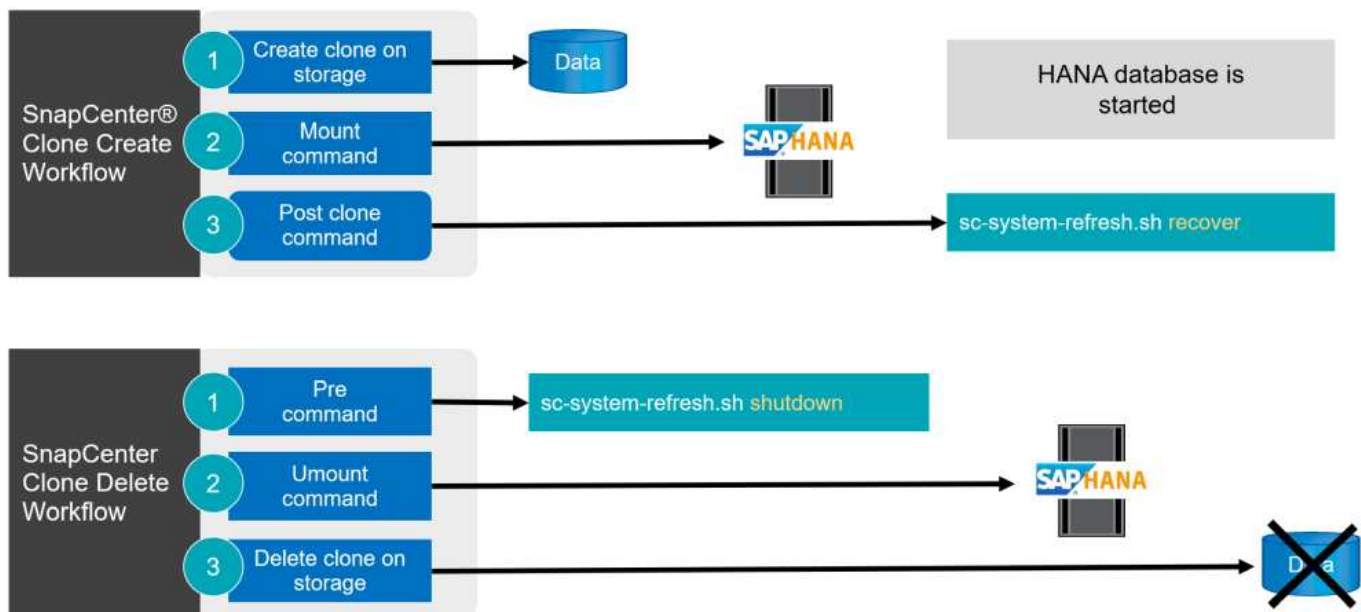
The example script `sc-system-refresh.sh` is used to execute recovery and shutdown operations. The script is called with specific command-line options within the SnapCenter workflows clone create and clone delete, as shown in the figure below.

The script is generic and reads all required parameters, like the SID from the target system. The script must be available at the target host of the system refresh operation. An hdb user store key must be configured for the user `<SID>adm` at the target system. The key must allow access to the SAP HANA system database and provide privileges for recovery operations. The key must have the name `<TARGET-SID>KEY`.

The script writes a log file `sc-system-refresh-SID.log`, to the same directory, where it gets executed.



The current version of the script supports single host systems MDC single tenant, or MDC multiple tenant configurations. It does not support SAP HANA multiple-host systems.



## Supported tenant recovery operations

As described in the section "SAP HANA system refresh operation workflows using storage snapshot" the possible tenant recovery operations at the target system depend on the tenant configuration of the source system. The script `sc-system-refresh.sh` supports all tenant recovery operations, which are possible dependent on the source system configuration, as shown in the table below.

If a different tenant name is required at the target system, the tenant must be renamed manually after the

recovery operation.

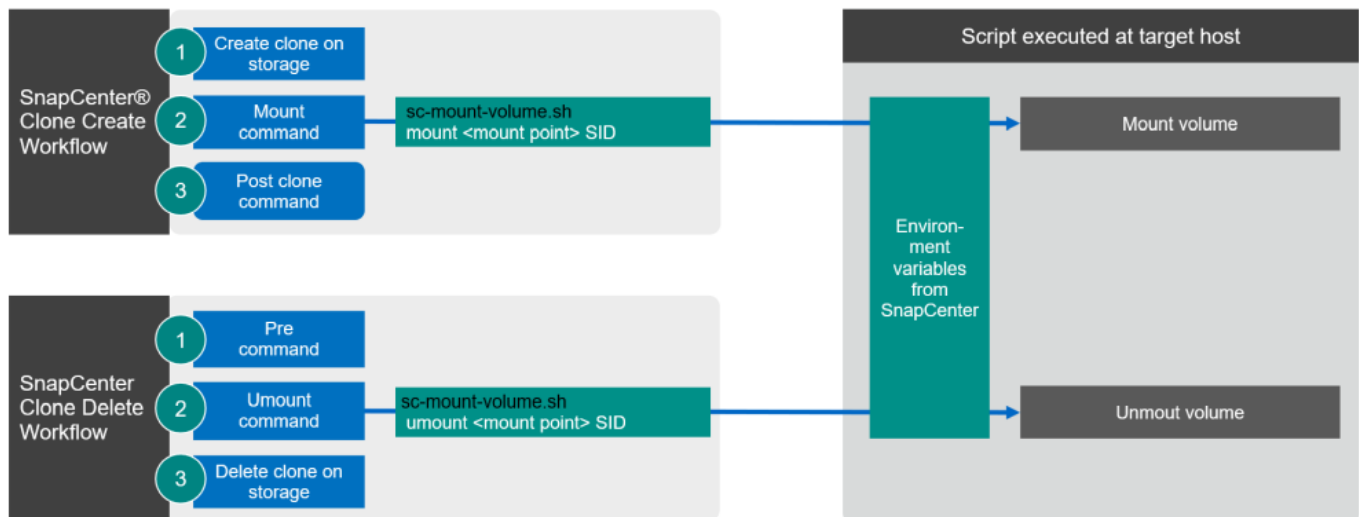
SAP HANA system	Tenant configuration at source system	Resulting tenant configuration at target system
MDC single tenant	Source tenant name equal to source SID	Target tenant name is equal to target SID
MDC single tenant	Source tenant name not equal to source SID	Target tenant name is equal to source tenant name
MDC multiple tenants	Any tenant names	All tenants are recovered and will have the same name as the source tenants.

## Script `sc-mount-volume.sh`

The example script `sc-mount-volume.sh` is used to execute mount and unmount for any volume. The script is used to mount the SAP HANA shared volume with the SAP HANA system clone operation. The script is called with specific command-line options within the SnapCenter workflows clone create and clone delete, as shown in the figure below.



The script supports SAP HANA systems using NFS as a storage protocol.



## SnapCenter environment variables

SnapCenter provides a set of environment variables that are available within the script that is executed at the target host. The script uses these variables to determine relevant configuration settings.

- The script variables `STORAGE`, `JUNCTION_PATH` are used for the mount operation.
- Derived from `CLONED_VOLUMES_MOUNT_PATH` environment variable.
- `CLONED_VOLUMES_MOUNT_PATH=${STORAGE}:/${JUNCTION_PATH}`
- For example:  
`CLONED_VOLUMES_MOUNT_PATH=192.168.175.117:/SS1_shared_Clone_05112206115489411`

## Script to get SnapCenter environment variables

If the automation scripts should not be used and the steps should be executed manually, you need to know the storage system junction path of the FlexClone volume. The junction path is not visible within SnapCenter, so you need to either look up the junction path directly at the storage system, or you could use a simple script that provides the SnapCenter environment variables at the target host. This script needs to be added as a mount operation script within the SnapCenter clone create operation.

```
ssladm@hana-1:/mnt/sapcc-share/SAP-System-Refresh> cat get-env.sh
#!/bin/bash
env > /tmp/env-from-sc.txt
ssladm@hana-1:/mnt/sapcc-share/SAP-System-Refresh>
```

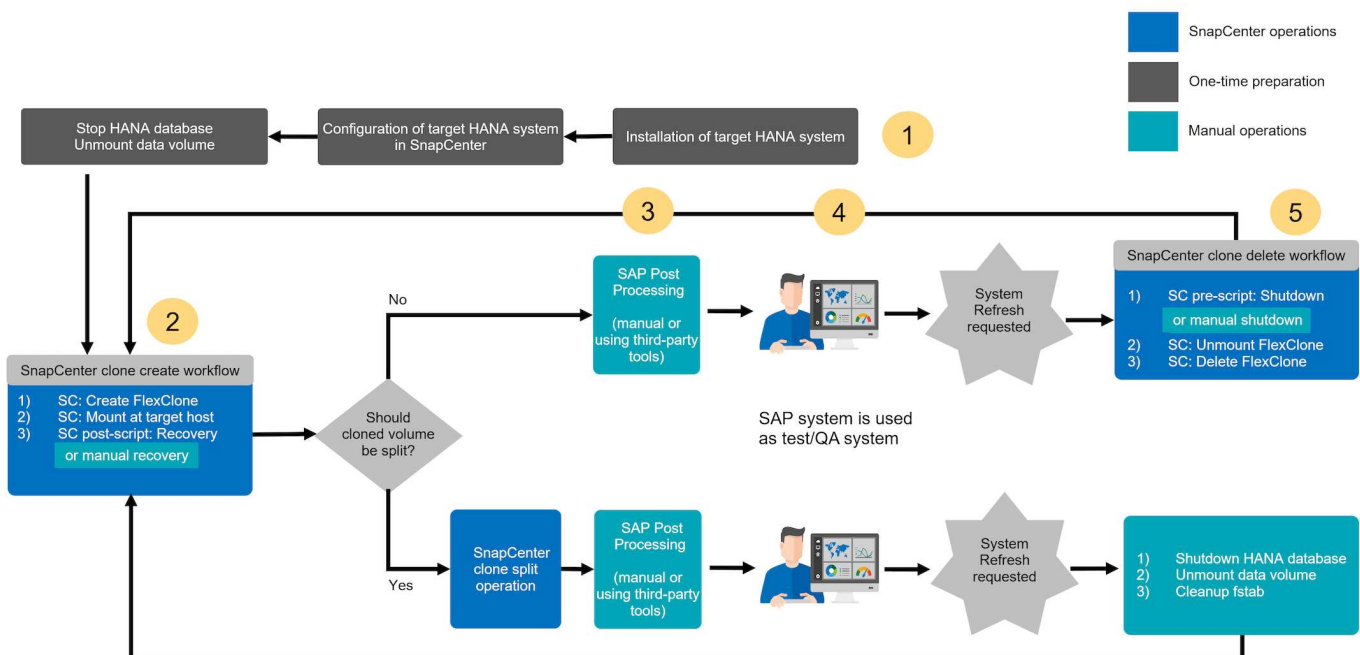
Within the `env-from-sc.txt` file, look for the variable `CLONED_VOLUMES_MOUNT_PATH` to get the storage system IP address and junction path of the FlexClone volume.

For example:

```
CLONED_VOLUMES_MOUNT_PATH=192.168.175.117:/SS1_data_mnt00001_Clone_05112206115489
411
```

## SAP HANA system refresh with SnapCenter

The following section provides a step-by-step description for the different system refresh operation options of an SAP HANA database.



Depending on the SAP HANA database configuration additional steps are executed or need to be prepared. The table below provides a summary.

Source system	Source system configuration	SnapCenter and SAP HANA operations
MDC single tenant SID = tenant name	Standard configuration	SnapCenter clone operation and optional recovery script execution.
	SAP HANA persistence encryption	Initially, or if root keys have been changed at the source system, root key backup(s) must be imported at the target system before recovery can be executed.
	SAP HANA system replication source	No additional steps required. If target system has no HSR configured it will stay a standalone system.
	SAP HANA multiple partitions	No additional steps required, but mount points for SAP HANA volume partitions must be available at the target system with same naming convention (only SID is different).
MDC multiple tenants or MDC single tenant with SID <> tenant name	Standard configuration	SnapCenter clone operation and optional recovery script execution. Script recovers all tenants. If tenants or tenant names does not exist at the target system names, required directories will be automatically created during the SAP HANA recovery operation. Tenant names will be same as source and need to be renamed after recovery, if required.
	SAP HANA persistence encryption	If a DBID of the source system does not exist before at the target system, the system database must be recovered first, before the root key backup of this tenant can be imported.
	HANA system replication source	No additional steps required. If target system has no HSR configured it will stay a standalone system.
	HANA multiple partitions	No additional steps required, but mount points for SAP HANA volume partitions must be available at the target system with same naming convention (only SID is different).

Within this section, the following scenarios are covered.

- SAP HANA system refresh without a clone split operation.
- Cloning from primary storage with tenant name equal to the SID
- Cloning from off-site backup storage
- Cloning from primary storage with multiple tenants
- Clone delete operation
- SAP HANA system refresh with a clone split operation
- Cloning from primary storage with tenant name equal to the SID
- Clone split operation

## Prerequisites and limitations

The workflows described in the following sections have a few prerequisites and limitations regarding the SAP HANA system architecture and the SnapCenter configuration.

- The described workflows are only valid for the SnapCenter 5.0 release or higher.
- The described workflows are valid for single host SAP HANA MDC systems with single or multiple tenants. SAP HANA multiple host systems are not covered.
- The SnapCenter SAP HANA plug-in must be deployed on the target host to enable SnapCenter auto discovery and the execution of automation scripts.
- The workflows are valid for SAP HANA systems using NFS or FCP on physical hosts, or for virtual hosts using in-guest NFS mounts.

## Lab setup

The figure below shows the lab setup that was used for the different system refresh operation options.

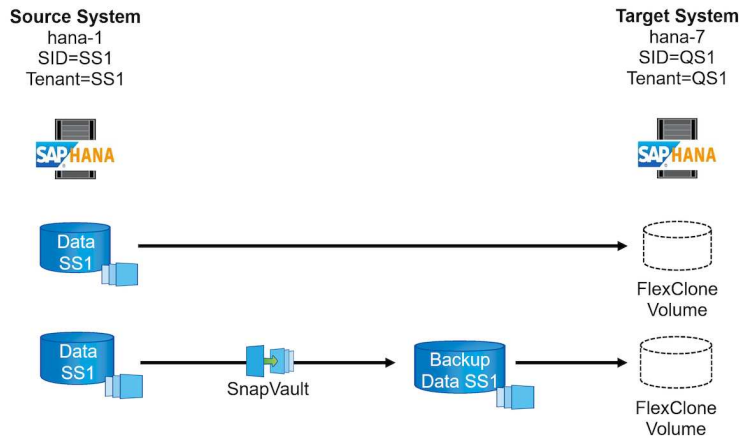
- Cloning from primary storage or off-site backup storage; tenant name is equal to the SID.
  - Source SAP HANA system: SS1 with Tenant SS1
  - Target SAP HANA system: QS1 with Tenant QS1
- Cloning from primary storage; multiple tenants.
  - Source SAP HANA system: SM1 with Tenant1 and Tenant2
  - Target SAP HANA system: QS1 with Tenant1 and Tenant2

The following software versions were used:


- SnapCenter 5.0
- SAP HANA systems: HANA 2.0 SPS7 rev.73
- SLES 15
- ONTAP 9.14P1

All SAP HANA systems must be configured based on the configuration guide [SAP HANA on NetApp AFF systems with NFS](#). SnapCenter and the SAP HANA resources were configured based on the best practice guide [SAP HANA Backup and Recovery with SnapCenter](#).

Cloning from primary or offsite backup storage, Tenant name = SID



Cloning from primary storage, Tenant name != SID



## Initial one-time preparation steps

As an initial step, the target SAP HANA system must be configured within SnapCenter.

1. Installation of SAP HANA target system
2. Configuration of SAP HANA system in SnapCenter  
as described in [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
  - a. Configuration of SAP HANA database user for SnapCenter backup operations  
This user must be identical at the source and the target system.
  - b. Configuration of hdbuserstore key for the <sid>adm with above backup user. If the automation script is used for recovery the key name must be <SID>KEY
  - c. Deployment of SnapCenter SAP HANA plug-in at target host. SAP HANA system is auto discovered by SnapCenter.
  - d. Configuration of SAP HANA resource protection (optional)

The first SAP system refresh operation after the initial installation is prepared with the following steps:

3. Shutdown target SAP HANA system
4. Unmount SAP HANA data volume.

You must add the scripts that should be executed at the target system to the SnapCenter allowed commands config file.

```
hana-7:/opt/NetApp/snapcenter/scc/etc # cat
/opt/NetApp/snapcenter/scc/etc/allowed_commands.config
command: mount
command: umount
command: /mnt/sapcc-share/SAP-System-Refresh/sc-system-refresh.sh
hana-7:/opt/NetApp/snapcenter/scc/etc #
```

## Cloning from primary storage with tenant name equal to SID

This section describes the SAP HANA system refresh workflow where the tenant name at the source and the target system is identical to the SID. The storage cloning is executed at the primary storage and the recovery is automated with the script `sc-system-refresh.sh`.



The workflow consists of the following steps:

1. If SAP HANA persistence encryption is enabled at the source system, the encryption root keys must be imported once. An import is also required if the keys have been changed at the source system. See chapter [“Considerations for SAP HANA system refresh operations using storage snapshot backups”](#)
2. If the target SAP HANA system has been protected in SnapCenter, the protection must be removed first.
3. SnapCenter clone create workflow.
  - a. Select Snapshot backup from the source SAP HANA system SS1.
  - b. Select target host and provide storage network interface of target host.
  - c. Provide SID of the target system, in our example QS1
  - d. Optionally, provide script for recovery as a post-clone operation.
4. SnapCenter cloning operation.
  - a. Creates FlexClone volume based on selected Snapshot backup of source SAP HANA system.
  - b. Exports FlexClone volume to target host storage network interface or igroup.
  - c. Executes mount operation of Mounts FlexClone volume at target host.
  - d. Executes post-clone operation recovery script, if configured before. Otherwise, recovery needs to be done manually when SnapCenter workflow is finished.
    - Recovery of system database.
    - Recovery of tenant database with tenant name = QS1.
5. Optionally, protect the target SAP HANA resource in SnapCenter.

The following screenshots show the required steps.

1. Select a Snapshot backup from the source system SS1 and click Clone.

The screenshot shows the NetApp SnapCenter interface. On the left, a sidebar lists systems: DT1, QS1, SM1, SS1, SS2, and SS2. The main area displays the 'SS1 Topology' with a 'Manage Copies' section showing 'Local copies' (14 Backups, 0 Clones) and 'Vault copies' (5 Backups, 0 Clones). A 'Summary Card' on the right shows '21 Backups', '19 Snapshot based backups', '2 File Based backups', '0 Clones', and '0 Snapshots Locked'. Below this is a 'Primary Backup(s)' section with a search bar and a table of backups.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-1_LocalSnap_Hourly_04-24-2024_07:00:01.4581		1	04/24/2024 7:01:07 AM
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_04-24-2024_05:00:02.7727		1	04/24/2024 5:01:08 AM
SnapCenter_hana-1_LocalSnap_Hourly_04-24-2024_03:00:01.5010		1	04/24/2024 3:01:07 AM
SnapCenter_hana-1_LocalSnap_Hourly_04-23-2024_23:00:01.5030		1	04/23/2024 11:01:06 PM
SnapCenter_hana-1_LocalSnap_Hourly_04-23-2024_19:00:01.4662		1	04/23/2024 7:01:06 PM
SnapCenter_hana-1_LocalSnap_Hourly_04-23-2024_15:00:01.4033		1	04/23/2024 3:01:06 PM
SnapCenter_hana-1_LocalSnap_Hourly_04-23-2024_11:00:02.4537		1	04/23/2024 11:01:08 AM
SnapCenter_LocalSnap_Hourly_04-23-2024_07:00:01.1635		1	04/23/2024 7:01:06 AM
SnapCenter_LocalSnapAndSnapVault_Daily_04-23-2024_05:00:02.1345		1	04/23/2024 5:01:08 AM
SnapCenter_LocalSnap_Hourly_04-23-2024_03:00:01.2634		1	04/23/2024 3:01:06 AM
SnapCenter_LocalSnap_Hourly_04-22-2024_23:00:01.2297		1	04/22/2024 11:01:06 PM

2. Select the host where the target system QS1 is installed. Enter QS1 as the target SID. The NFS export IP address must be the storage network interface of the target host.



The target SID which is entered controls how SnapCenter manages the cloned resource. If a resource with the target SID is already configured in SnapCenter and matches the plug-in host, SnapCenter just assigns the clone to this resource. If the SID is not configured on the target host, SnapCenter creates a new resource.



It is crucial that the target system resource and host has been configured in SnapCenter before you start the cloning workflow. Otherwise, the new resource created by SnapCenter will not support auto discovery and the described workflows won't work.

The screenshot shows the 'Clone From Backup' dialog box. It has a sidebar with steps: 1 Location, 2 Scripts, 3 Notification, and 4 Summary. The 'Location' step is active, showing a form to 'Select the host to create the clone'. The form has three fields: 'Plug-in host' (hana-7.sapcc.stl.netapp.com), 'Target Clone SID' (QS1), and 'NFS Export IP Address' (192.168.175.75). Each field has an information icon to its right.

In a Fibre Channel SAN setup, no export IP address is required, but you need to provide the used protocol in the next screen.



The screenshots show a different lab setup using a FibreChannel connectivity.

Clone From Backup

1 Location

2 Settings

3 Scripts

4 Notification

5 Summary

Select the host to create the clone

Plug-in host

cbc-demosrv02.muccbc.hq.netapp.com

Target Clone SID

H12

NFS Export IP Address

Clone From Backup

1 Location

2 Settings

3 Scripts

4 Notification

5 Summary

LUN Map Settings

Igroup protocol

FCP

With Azure NetApp Files and a manual QoS capacity pool, you need to provide the maximum throughput for the new volume. Make sure that the capacity pool has enough headroom, otherwise the cloning workflow will fail.



The screenshots show a different lab setup running in Microsoft Azure with Azure NetApp Files.

Clone From Backup

1 Location

2 Scripts

3 Notification

4 Summary

Select the host to create the clone

Plug-in host

vm-s01.1h05kdpkcgaujd4qsseqldygg.bx.i

Target Clone SID

S01

NFS Export IP Address

10.1.8.101

Capacity Pool Max. Throughput (MiB/s)

25

- Enter the optional post-clone scripts with the required command-line options. With our example we use a post clone script to execute the SAP HANA database recovery.

Clone From Backup
×

1 Location
2 Scripts
3 Notification
4 Summary

The following commands will run on the Plug-in Host: hana-7.sapcc.stl.netapp.com

Enter optional commands to run before performing a clone operation ⓘ

Pre clone command

Enter optional commands to run after performing a clone operation ⓘ

Post clone command




As discussed before, the usage of the recovery script is optional. The recovery can also be done manually after the SnapCenter cloning workflow is finished.



The script for the recovery operation recovers the SAP HANA database to the point in time of the Snapshot using the clear logs operation and does not execute any forward recovery. If a forward recovery to a specific point in time is required, the recovery must be performed manually. A manual forward recovery also requires that the log backups from the source system are available at the target host.

- The Job Details screen in SnapCenter shows the progress of the operation. The job details also show that the overall runtime including database recovery has been less than 3 minutes.



- The logfile of the `sc-system-refresh` script shows the different steps that were executed for the recovery operation. The script reads the list of tenants from the system database and executes a recovery of all existing tenants.

```
20240425112328###hana-7###sc-system-refresh.sh: Script version: 3.0
hana-7:/mnt/sapcc-share/SAP-System-Refresh # cat sap-system-refresh-
QS1.log
20240425112328###hana-7###sc-system-refresh.sh: *****
Starting script: recovery operation *****
20240425112328###hana-7###sc-system-refresh.sh: Recover system database.
```

```

20240425112328###hana-7###sc-system-refresh.sh:
/usr/sap/QS1/HDB11/exe/Python/bin/python
/usr/sap/QS1/HDB11/exe/python_support/recoverSys.py --command "RECOVER
DATA USING SNAPSHOT CLEAR LOG"
20240425112346###hana-7###sc-system-refresh.sh: Wait until SAP HANA
database is started ....
20240425112347###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425112357###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425112407###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425112417###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425112428###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425112438###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425112448###hana-7###sc-system-refresh.sh: Status: GREEN
20240425112448###hana-7###sc-system-refresh.sh: HANA system database
started.
20240425112448###hana-7###sc-system-refresh.sh: Checking connection to
system database.
20240425112448###hana-7###sc-system-refresh.sh:
/usr/sap/QS1/SYS/exe/hdb/hdbsql -U QS1KEY 'select * from sys.m_databases;'
DATABASE_NAME,DESCRIPTION,ACTIVE_STATUS,ACTIVE_STATUS_DETAILS,OS_USER,OS_G
ROUP,RESTART_MODE,FALLBACK_SNAPSHOT_CREATE_TIME
"SYSTEMDB","SystemDB-QS1-11","YES","","","","DEFAULT",?
"QS1","QS1-11","NO","ACTIVE","","","DEFAULT",?
2 rows selected (overall time 16.225 msec; server time 860 usec)
20240425112448###hana-7###sc-system-refresh.sh: Successfully connected to
system database.
20240425112449###hana-7###sc-system-refresh.sh: Tenant databases to
recover: QS1
20240425112449###hana-7###sc-system-refresh.sh: Found inactive
tenants(QS1) and starting recovery
20240425112449###hana-7###sc-system-refresh.sh: Recover tenant database
QS1.
20240425112449###hana-7###sc-system-refresh.sh:
/usr/sap/QS1/SYS/exe/hdb/hdbsql -U QS1KEY RECOVER DATA FOR QS1 USING
SNAPSHOT CLEAR LOG
0 rows affected (overall time 22.138599 sec; server time 22.136268 sec)
20240425112511###hana-7###sc-system-refresh.sh: Checking availability of
Indexserver for tenant QS1.
20240425112511###hana-7###sc-system-refresh.sh: Recovery of tenant
database QS1 succesfully finished.
20240425112511###hana-7###sc-system-refresh.sh: Status: GREEN
20240425112511###hana-7###sc-system-refresh.sh: *****
Finished script: recovery operation *****
hana-7:/mnt/sapcc-share/SAP-System-Refresh

```

6. When the SnapCenter job is finished, the clone is visible within the topology view of the source system.

The screenshot shows the NetApp SnapCenter interface for managing SAP HANA databases. The left sidebar displays a list of systems: DT1, Q51, SM1, SS1, SS2, and SS2. The main panel shows the 'Manage Copies' section for the 'SS1 Topology'. It displays a diagram with 'Local copies' (14 Backups, 1 Clone) and 'Vault copies' (5 Backups, 0 Clones). A 'Summary Card' on the right provides a quick overview of the backup status: 21 Backups, 19 Snapshot based backups, 2 File Based backups, 1 Clone, and 0 Snapshots Locked. Below the diagram, a table lists the 'Primary Clone(s)' with columns for Clone SID, Clone Host, Clone Name, Start Date, and End date. The table shows one clone with SID Q51, Host hana-7.sapcc.stl.netapp.com, and a Start Date of 04/24/2024 9:47:10 AM. The bottom status bar indicates 1 Completed, 2 Warnings, 0 Failed, 0 Canceled, 2 Running, and 0 Queued jobs.

7. The SAP HANA database is now running.
8. If you want to protect the target SAP HANA system, you need to run the auto discovery by clicking on the target system resource.

The 'Configure Database' dialog box is shown, with the following fields and values:

- Plug-in host: hana-7.sapcc.stl.netapp.com
- HDBSQL OS User: qs1adm
- HDB Secure User Store Key: QS1KEY

At the bottom right, there are 'Cancel' and 'OK' buttons.

When the auto discovery process is finished, the new cloned volume is listed in the storage footprint section.

**NetApp SnapCenter®**

SAP HANA

Search databases

System

DT1

Q51

SM1

SS1

SS2

SS2

Total 6

Resource - Details

Details for selected resource

Type	Multitenant Database Container
HANA System Name	Q51
SID	Q51
Tenant Databases	Q51
Plug-in Host	hana-7.sapcc.stl.netapp.com
HDB Secure User Store Key	Q51KEY
HDBSQL OS User	qs1adm
Log backup location	/usr/sap/Q51/HDB11/backup/log
Backup catalog location	/usr/sap/Q51/HDB11/backup/log
System Replication	None
plug-in name	SAP HANA
Last backup	None
Resource Groups	None
Policy	None
Discovery Type	Auto
Backup Name	SnapCenter_hana-1_LocalSnap_Hourly_06-25-2024_03.00.01.8458
Backup Name of Clone	SnapCenter_hana-1_LocalSnap_Hourly_06-25-2024_03.00.01.8458

Storage Footprint

SVM	Volume	Junction Path	LUN/Qtree
hana-primary	SS1_data_mnt00001_Clone_06252405324571927	/SS1_data_mnt00001_Clone_06252405324571927	

Activity The 5 most recent jobs are displayed

4 Completed 0 Warnings 1 Failed 0 Canceled 0 Running 0 Queued

By clicking on the resource again, data protection can be configured for the refreshed Q51 system.

**NetApp SnapCenter®**

SAP HANA

Search databases

System

DT1

Q51

SM1

SS1

SS2

SS2

Multitenant Database Container - Protect

Protect the resource by selecting protection policies, schedules, and notification settings.

Configure an SMTP Server to send email notifications for scheduled or on demand jobs by going to [Settings>Global Settings>Notification Server Settings](#).

1 Resource 2 Application Settings 3 Policies 4 Notification 5 Summary

Provide format for custom snapshot name

☐ Use custom name format for Snapshot copy

## Cloning from off-site backup storage

This section describes the SAP HANA system refresh workflow for which the tenant name at the source and the target system is identical to the SID. Storage cloning is executed at the off-site backup storage and further automated using the script `sc-system-refresh.sh`.

## Source System

hana-1  
SID=SS1  
Tenant=SS1



SnapVault



## Target System

hana-7  
SID=QS1  
Tenant=QS1



FlexClone  
Volume

The only difference in the SAP HANA system refresh workflow between primary and off-site backup storage cloning is the selection of the Snapshot backup in SnapCenter. For off-site backup storage cloning, the secondary backups must be selected first, followed by the selection of the Snapshot backup.

The screenshot shows the NetApp SnapCenter interface for SAP HANA. The left sidebar displays the system hierarchy: SAP HANA > System > QS1 > SM1 > SS1 > SS2 > SS2. The main panel shows the 'Manage Copies' section for 'SS1 Topology'. It displays 'Local copies' with 14 Backups and 0 Clones, and 'Vault copies' with 9 Backups and 0 Clones. A 'Summary Card' on the right shows 25 Backups, 23 Snapshot based backups, 2 File Based backups, and 0 Clones. Below this, the 'Secondary Vault Backup(s)' section shows a table of backup details.

Backup Name	Count	IF	End Date
SnapCenter_LocalSnapAndSnapVault_Daily_05-11-2022_05.00.02.9288	1		05/11/2022 5:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-10-2022_05.00.02.9444	1		05/10/2022 5:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-09-2022_05.00.02.9432	1		05/09/2022 5:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-08-2022_05.00.02.9894	1		05/08/2022 5:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-07-2022_05.00.02.9253	1		05/07/2022 5:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-06-2022_05.00.02.9333	1		05/06/2022 5:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-05-2022_05.00.03.8844	1		05/05/2022 5:01:02 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-04-2022_05.00.03.0342	1		05/04/2022 5:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_05-03-2022_05.00.02.9761	1		05/03/2022 5:01:01 AM

If there are multiple secondary storage locations for the selected backup, you need to choose the required destination volume.

Clone From Backup
×

1 Location

2 Scripts

3 Notification

4 Summary

Select the host to create the clone

Plug-in host

hana-7.sapcc.stl.netapp.com

ⓘ

Target Clone SID

QS1

ⓘ

NFS Export IP Address

192.168.175.75

ⓘ

Secondary storage location : Snap Vault / Snap Mirror

Source Volume

Destination Volume

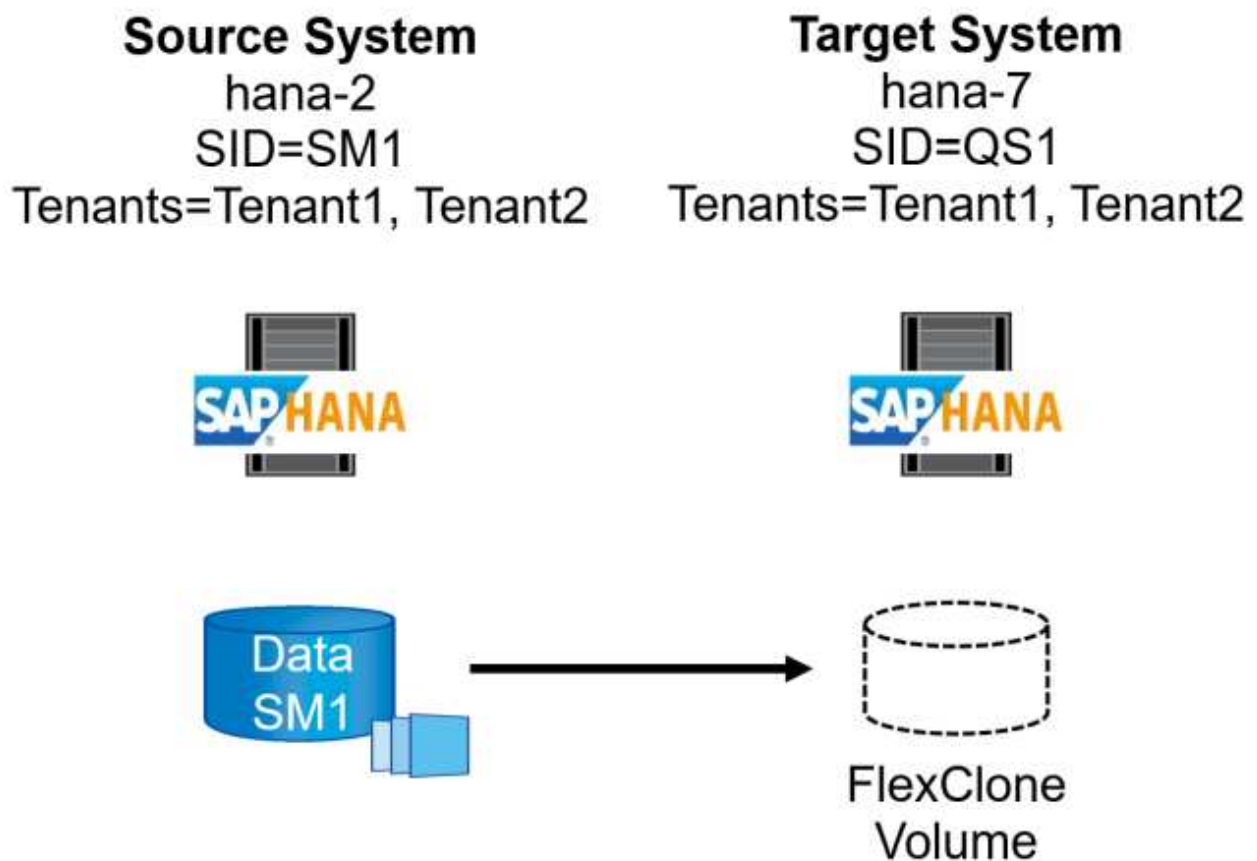
hana-primary.sapcc.stl.netapp.com:SS1\_data\_mnt00001

hana-backup.sapcc.stl.netapp.com:SS1\_data

All subsequent steps are identical to the workflow for cloning from primary storage.

## Cloning a SAP HANA system with multiple tenants

This section describes the SAP HANA system refresh workflow with multiple tenants. Storage cloning is executed at the primary storage and further automated using the script `sc-system-refresh.sh`.



The required steps in SnapCenter are identical to what has been described in the section “Cloning from

primary storage with tenant name equal to SID." The only difference is in the tenant recovery operation within the script `sc-system-refresh.sh`, where all tenants are recovered.

```
20240430070214###hana-7###sc-system-refresh.sh:
*****
*****
20240430070214###hana-7###sc-system-refresh.sh: Script version: 3.0
20240430070214###hana-7###sc-system-refresh.sh: *****
Starting script: recovery operation *****
20240430070214###hana-7###sc-system-refresh.sh: Recover system database.
20240430070214###hana-7###sc-system-refresh.sh:
/usr/sap/QS1/HDB11/exe/Python/bin/python
/usr/sap/QS1/HDB11/exe/python_support/recoverSys.py --command "RECOVER
DATA USING SNAPSHOT CLEAR LOG"
[140310725887808, 0.008] >> starting recoverSys (at Tue Apr 30 07:02:15
2024)
[140310725887808, 0.008] args: ()
[140310725887808, 0.008] keys: \{'command': 'RECOVER DATA USING SNAPSHOT
CLEAR LOG'}
using logfile /usr/sap/QS1/HDB11/hana-7/trace/backup.log
recoverSys started: =====2024-04-30 07:02:15 =====
testing master: hana-7
hana-7 is master
shutdown database, timeout is 120
stop system
stop system on: hana-7
stopping system: 2024-04-30 07:02:15
stopped system: 2024-04-30 07:02:15
creating file recoverInstance.sql
restart database
restart master nameserver: 2024-04-30 07:02:20
start system: hana-7
sapcontrol parameter: ['-function', 'Start']
sapcontrol returned successfully:
2024-04-30T07:02:32-04:00 P0023828 18f2eab9331 INFO RECOVERY RECOVER DATA
finished successfully
recoverSys finished successfully: 2024-04-30 07:02:33
[140310725887808, 17.548] 0
[140310725887808, 17.548] << ending recoverSys, rc = 0 (RC_TEST_OK), after
17.540 secs
20240430070233###hana-7###sc-system-refresh.sh: Wait until SAP HANA
database is started ....
20240430070233###hana-7###sc-system-refresh.sh: Status: GRAY
20240430070243###hana-7###sc-system-refresh.sh: Status: GRAY
20240430070253###hana-7###sc-system-refresh.sh: Status: GRAY
20240430070304###hana-7###sc-system-refresh.sh: Status: GRAY
```

```

20240430070314###hana-7###sc-system-refresh.sh: Status: GREEN
20240430070314###hana-7###sc-system-refresh.sh: HANA system database
started.
20240430070314###hana-7###sc-system-refresh.sh: Checking connection to
system database.
20240430070314###hana-7###sc-system-refresh.sh:
/usr/sap/QS1/SYS/exe/hdb/hdbsql -U QS1KEY 'select * from sys.m_databases;'
20240430070314###hana-7###sc-system-refresh.sh: Succesfully connected to
system database.
20240430070314###hana-7###sc-system-refresh.sh: Tenant databases to
recover: TENANT2
TENANT1
20240430070314###hana-7###sc-system-refresh.sh: Found inactive
tenants(TENANT2
TENANT1) and starting recovery
20240430070314###hana-7###sc-system-refresh.sh: Recover tenant database
TENANT2.
20240430070314###hana-7###sc-system-refresh.sh:
/usr/sap/QS1/SYS/exe/hdb/hdbsql -U QS1KEY RECOVER DATA FOR TENANT2 USING
SNAPSHOT CLEAR LOG
20240430070335###hana-7###sc-system-refresh.sh: Checking availability of
Indexserver for tenant TENANT2.
20240430070335###hana-7###sc-system-refresh.sh: Recovery of tenant
database TENANT2 succesfully finished.
20240430070335###hana-7###sc-system-refresh.sh: Status: GREEN
20240430070335###hana-7###sc-system-refresh.sh: Recover tenant database
TENANT1.
20240430070335###hana-7###sc-system-refresh.sh:
/usr/sap/QS1/SYS/exe/hdb/hdbsql -U QS1KEY RECOVER DATA FOR TENANT1 USING
SNAPSHOT CLEAR LOG
20240430070349###hana-7###sc-system-refresh.sh: Checking availability of
Indexserver for tenant TENANT1.
20240430070350###hana-7###sc-system-refresh.sh: Recovery of tenant
database TENANT1 succesfully finished.
20240430070350###hana-7###sc-system-refresh.sh: Status: GREEN
20240430070350###hana-7###sc-system-refresh.sh: *****
Finished script: recovery operation *****

```

## Clone delete operation

A new SAP HANA system refresh operation is started by cleaning up the target system using the SnapCenter clone delete operation.

If the target SAP HANA system has been protected in SnapCenter, the protection must be removed first. Within the topology view of the target system, click Remove Protection.

The clone delete workflow is now executed with the following steps.

1. Select the clone within the topology view of the source system and click Delete.

NetApp SnapCenter®

SAP HANA

SS1 Topology

Search databases

Backup to Object Store Remove Protection Back up Now Modify Maintenance Details Configure Database Refresh

Manage Copies

Local copies: 14 Backups, 1 Clone

Vault copies: 6 Backups, 0 Clones

Summary Card

- 22 Backups
- 20 Snapshot based backups
- 2 File-based backups w/
- 1 Clone
- 0 Snapshots Locked

Primary Clone(s)

search

Clone SID	Clone Host	Clone Name	Start Date	End date
Q51	hana-7.sapcc.stl.netapp.com	hana-1_sapcc_stl_netapp_com_hana_MDC_SS1_clone_102336_MDC_SS1_04-22-2024_09:54:34	04/25/2024 10:41:50 AM	04/25/2024 10:42:38 AM

Total 6 Total 1

Activity The 5 most recent jobs are displayed

Completed Versions Failed Canceled Running Queued

2. Enter the pre-clone and unmount scripts with the required command line options.

Delete Clone

Cloned volume will be deleted. SnapCenter backups and HANA backup catalog must be deleted manually.

Enter commands to execute before clone deletion

Pre clone delete :

```
/mnt/sapcc-share/SAP-System-Refresh/sc-system-refresh.sh  
shutdown
```

The selected clone(s) will be permanently deleted. If the selected clone contains other resource(s) it will also be deleted.

If the cloned databases are protected then the protection needs to be removed to delete the clone.

Do you want to proceed?

☐ Force Delete

Cancel OK

3. The job details screen in SnapCenter shows the progress of the operation.

Job Details

X

Deleting clone 'hana-1\_sapcc\_stl\_netapp\_com\_ha.....S1\_\_clone\_\_102336\_MDC\_SS1\_04-22-2024\_09.54.34'

▼ Deleting clone 'hana-1\_sapcc\_stl\_netapp\_com\_hana\_MDC\_SS1\_\_clone\_\_102336\_MDC\_SS1\_04-22-2024\_09.54.34'

▼ hana-7.sapcc.stl.netapp.com

▼ Delete Clone

▶ Validate Plugin Parameters

▼ Delete Pre Clone Commands

▶ Unmount Filesystem

▼ Delete Storage Clone

▼ Unregister Clone Metadata

▶ Filesystem Clone Metadata Cleanup

▶ Agent Finalize Workflow

Task Name: Unmount Filesystem Start Time: 04/25/2024 11:11:56 AM End Time: 04/25/2024 11:12:08 AM

View Logs

Cancel Job

Close

4. The log file of the `sc-system-refresh` script shows the shutdown and unmount operation steps.

```

20240425111042###hana-7###sc-system-refresh.sh:
*****
*****
20240425111042###hana-7###sc-system-refresh.sh: Script version: 3.0
20240425111042###hana-7###sc-system-refresh.sh: *****
Starting script: shutdown operation *****
20240425111042###hana-7###sc-system-refresh.sh: Stopping HANA database.
20240425111042###hana-7###sc-system-refresh.sh: sapcontrol -nr 11
-function StopSystem HDB
25.04.2024 11:10:42
StopSystem
OK
20240425111042###hana-7###sc-system-refresh.sh: Wait until SAP HANA
database is stopped ....
20240425111042###hana-7###sc-system-refresh.sh: Status: GREEN
20240425111052###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425111103###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425111113###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425111123###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425111133###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425111144###hana-7###sc-system-refresh.sh: Status: YELLOW
20240425111154###hana-7###sc-system-refresh.sh: Status: GRAY
20240425111154###hana-7###sc-system-refresh.sh: SAP HANA database is
stopped.
20240425111154###hana-7###sc-system-refresh.sh: *****
Finished script: shutdown operation *****

```

5. The SAP HANA refresh operation can now be started again using the SnapCenter clone create operation.

## SAP HANA system refresh with clone split operation

If the target system of the system refresh operation is planned to be used for a longer timeframe, it makes sense to split the FlexClone volume as part of the system refresh operation.

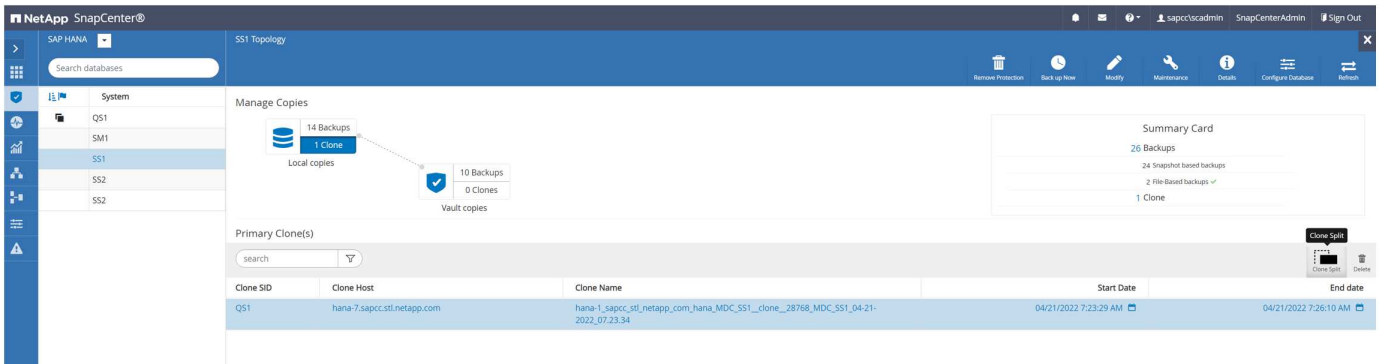


The clone split operation does not block the use of the cloned volume and can therefore be executed at any time while the SAP HANA database is in use.

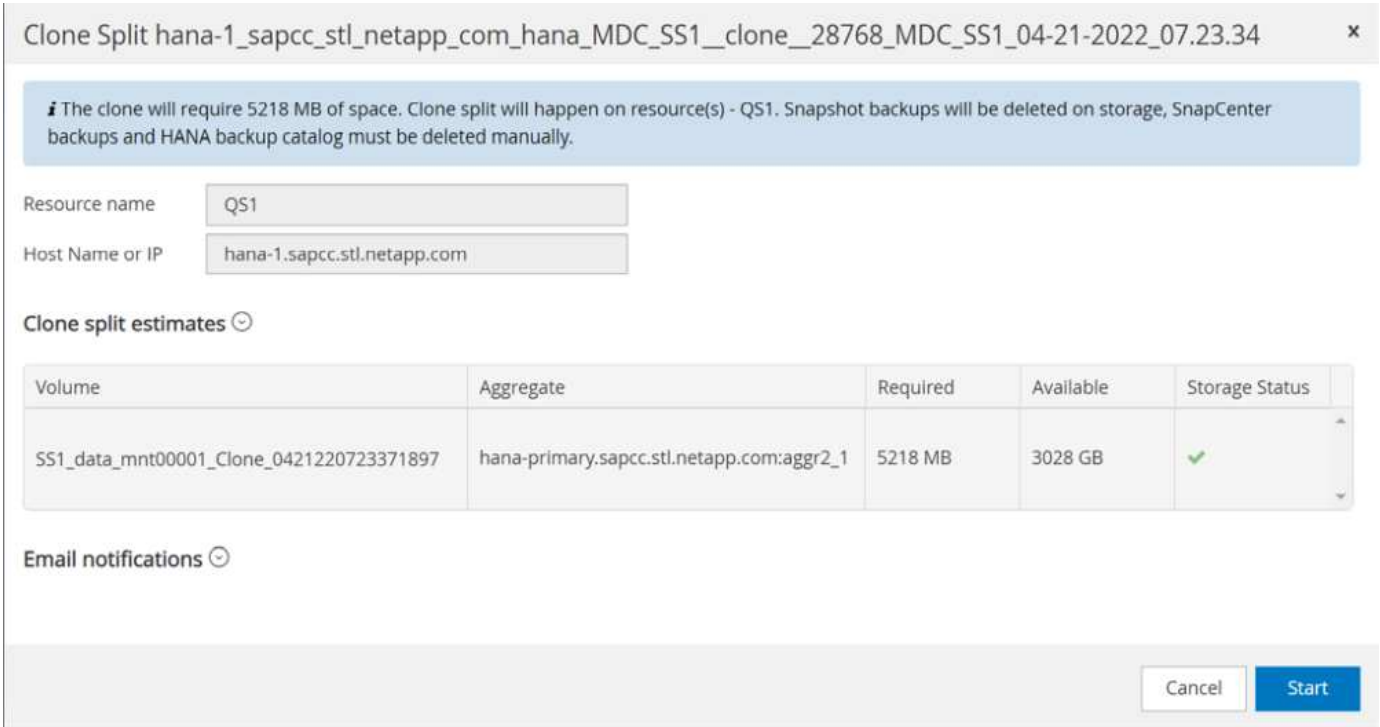


With Azure NetApp Files, the clone split operation is not available, since Azure NetApp Files always splits the clone after creation.

The clone split workflow in SnapCenter is initiated in the topology view of the source system by selecting the clone and clicking on clone split.



A preview is shown in the next screen, which provides information on the required capacity for the split volume.



The SnapCenter job log shows the progress of the clone split operation.

Job Details

Clone Split Start of Resource 'hana-1\_sapcc\_stl\_ne.....MDC\_SS1\_\_clone\_\_28768\_MDC\_SS1\_04-21-2022\_07.23.34'

▼ Clone Split Start of Resource 'hana-1\_sapcc\_stl\_netapp\_com\_hana\_MDC\_SS1\_\_clone\_\_28768\_MDC\_SS1\_04-21-2022\_07.23.34'

▼ SnapCenter.sapcc.stl.netapp.com

▶ Volume Clone Estimate

▶ Volume Clone Split Start

▶ Delete Backups of Clone

▶ Volume Clone Split Status

▶ Clone Split Status for volume SS1\_data\_mnt00001\_Clone\_0421220723371897 is 'In Progress'

▶ Clone Split Status for volume SS1\_data\_mnt00001\_Clone\_0421220723371897'Completed'

▶ Register Clone Split

▶ Data Collection

▶ Send EMS Messages

Task Name: Volume Clone Split Status Start Time: 04/21/2022 7:51:16 AM End Time:

View Logs

Cancel Job

Close

In the resource view in SnapCenter the target system QS1 is now not marked as a cloned resource anymore. When going back to the topology view of the source system, the clone is not visible anymore. The split volume is now independent from the Snapshot backup of the source system.

38

System	System ID (SID)	Tenant Databases	Replication	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
Q51	Q51	Q51	None	hana-7.sapcc.stl.netapp.com		LocalSnap	04/21/2022 7:30:50 AM	Backup succeeded
SM1	SM1	TENANT1	None	hana-2.sapcc.stl.netapp.com		LocalSnap	04/21/2022 4:01:01 AM	Backup succeeded
SS1	SS1	SS1	None	hana-1.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap LocalSnapAndSnapVault LocalSnap-OnDemand	04/21/2022 7:01:01 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Primary)	hana-3.sapcc.stl.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnap LocalSnapKeep2	04/21/2022 7:57:22 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Secondary)	hana-4.sapcc.stl.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapKeep2	04/11/2022 2:57:21 AM	Backup succeeded

**SS1 Topology**

Manage Copies

- Local copies: 14 Backups, 0 Clones
- Vault copies: 10 Backups, 0 Clones

Summary Card

- 26 Backups
- 24 Snapshot based backups
- 2 File Based backups w/
- 0 Clones

Primary Backup(s)

Backup Name	Count	End Date
SnapCenter_LocalSnap_Hourly_04-21-2022_07.00.02.7865	1	04/21/2022 7:01:01 AM
SnapCenter_LocalSnapAndSnapVault_Daily_04-21-2022_05.00.02.8215	1	04/21/2022 5:01:02 AM
SnapCenter_LocalSnap_Hourly_04-21-2022_03.00.01.7085	1	04/21/2022 3:01:00 AM
SnapCenter_LocalSnap_Hourly_04-20-2022_23.00.01.7142	1	04/20/2022 11:01:00 PM
SnapCenter_LocalSnap_Hourly_04-20-2022_19.00.01.9499	1	04/20/2022 7:01:00 PM

The refresh workflow after a clone split operation looks slightly different than the operation without clone split. After a clone split operation, there is no clone delete operation required, because the target data volume is not a FlexClone volume anymore.

The workflow consists of the following steps:

1. If the target SAP HANA system has been protected in SnapCenter, the protection must be removed first.
2. The SAP HANA database must be shut down, the data volume must be unmounted and the fstab entry created by SnapCenter must be removed. These steps need to be executed manually.
3. Now the SnapCenter clone create workflow can be executed as described in sections before.
4. After the refresh operation, the old target data volume still exists and it must be deleted manually with, for example, ONTAP System Manager.

## SnapCenter workflow automation with PowerShell scripts

In the previous sections, the different workflows were executed using the SnapCenter UI. All the workflows can also be executed with PowerShell scripts or REST API calls, allowing further automation. The following sections describe basic PowerShell script examples for the following workflows.

- Create clone
- Delete clone



The example scripts are provided as is and are not supported by NetApp.

All scripts must be executed in a PowerShell command window. Before the scripts can be run, a connection to the SnapCenter server must be established using the `Open-SmConnection` command.

## Create clone

The simple script below demonstrates how a SnapCenter clone create operation can be executed using PowerShell commands. The SnapCenter `New-SmClone` command is executed with the required command line option for the lab environment and the automation script discussed before.

```
$BackupName='SnapCenter_hana-1_LocalSnap_Hourly_06-25-2024_03.00.01.8458'
$JobInfo=New-SmClone -AppPluginCode hana -BackupName $BackupName
-Resources @\{"Host"="hana-1.sapcc.stl.netapp.com";"UID"="MDC\SS1"}
-CloneToInstance hana-7.sapcc.stl.netapp.com -postclonecreatecommands
'/mnt/sapcc-share/SAP-System-Refresh/sc-system-refresh.sh recover'
-NFSEExportIPs 192.168.175.75 -CloneUid 'MDC\QS1'
# Get JobID of clone create job
$Job=Get-SmJobSummaryReport | ?\{$_.JobType -eq "Clone" } | ?\{$_.JobName
-Match $BackupName} | ?\{$_.Status -eq "Running"}
$JobId=$Job.SmJobId
Get-SmJobSummaryReport -JobId $JobId
# Wait until job is finished
do \{ $Job=Get-SmJobSummaryReport -JobId $JobId; write-host $Job.Status;
sleep 20 } while ( $Job.Status -Match "Running" )
Write-Host " "
Get-SmJobSummaryReport -JobId $JobId
Write-Host "Clone create job has been finshed."
```

The screen output shows the execution of the clone create PowerShell script.

```

PS C:\Windows\system32> C:\NetApp\clone-create.ps1
SmJobId : 110382
JobCreatedDateTime :
JobStartDateTime : 6/26/2024 9:55:34 AM
JobEndDateTime :
JobDuration :
JobName : Clone from backup 'SnapCenter_hana-1_LocalSnap_Hourly_06-25-
2024_03.00.01.8458'
JobDescription :
Status : Running
IsScheduled : False
JobError :
JobType : Clone
PolicyName :
JobResultData :
Running
Running
Running
Running
Running
Running
Running
Running
Running
Running
Running
Running
Completed
SmJobId : 110382
JobCreatedDateTime :
JobStartDateTime : 6/26/2024 9:55:34 AM
JobEndDateTime : 6/26/2024 9:58:50 AM
JobDuration : 00:03:16.6889170
JobName : Clone from backup 'SnapCenter_hana-1_LocalSnap_Hourly_06-25-
2024_03.00.01.8458'
JobDescription :
Status : Completed
IsScheduled : False
JobError :
JobType : Clone
PolicyName :
JobResultData :
Clone create job has been finshed.

```

## Delete clone

The simple script below demonstrates how a SnapCenter clone delete operation can be executed using

PowerShell commands. The SnapCenter `Remove-SmClone` command is executed with the required command line option for the lab environment and the automation script discussed before.

```
$CloneInfo=Get-SmClone |?{$_.CloneName -Match "hana-  
1_sapcc_stl_netapp_com_hana_MDC_SS1" }  
$JobInfo=Remove-SmClone -CloneName $CloneInfo.CloneName -PluginCode hana  
-PreCloneDeleteCommands '/mnt/sapcc-share/SAP-System-Refresh/sc-system-  
refresh.sh shutdown QS1' -UnmountCommands '/mnt/sapcc-share/SAP-System-  
Refresh/sc-system-refresh.sh umount QS1' -Confirm: $False  
Get-SmJobSummaryReport -JobId $JobInfo.Id  
# Wait until job is finished  
do \{ $Job=Get-SmJobSummaryReport -JobId $JobInfo.Id; write-host  
$Job.Status; sleep 20 } while ( $Job.Status -Match "Running" )  
Write-Host " "  
Get-SmJobSummaryReport -JobId $JobInfo.Id  
Write-Host "Clone delete job has been finshed."  
PS C:\NetApp>
```

The screen output shows the execution of the clone –delete.ps1 PowerShell script.

```

PS C:\Windows\system32> C:\NetApp\clone-delete.ps1
SmJobId : 110386
JobCreatedDateTime :
JobStartDateTime : 6/26/2024 10:01:33 AM
JobEndDateTime :
JobDuration :
JobName : Deleting clone 'hana-
1_sapcc_stl_netapp_com_hana_MDC_SS1__clone__110382_MDC_SS1_04-22-
2024_09.54.34'
JobDescription :
Status : Running
IsScheduled : False
JobError :
JobType : DeleteClone
PolicyName :
JobResultData :
Running
Running
Running
Running
Completed
SmJobId : 110386
JobCreatedDateTime :
JobStartDateTime : 6/26/2024 10:01:33 AM
JobEndDateTime : 6/26/2024 10:02:38 AM
JobDuration : 00:01:05.5658860
JobName : Deleting clone 'hana-
1_sapcc_stl_netapp_com_hana_MDC_SS1__clone__110382_MDC_SS1_04-22-
2024_09.54.34'
JobDescription :
Status : Completed
IsScheduled : False
JobError :
JobType : DeleteClone
PolicyName :
JobResultData :
Clone delete job has been finshed.
PS C:\Windows\system32>

```

## SAP system clone with SnapCenter

This section provides a step-by-step description for the SAP system clone operation, which can be used to set up a repair system to address logical corruption.

The figure below summarizes the required steps for an SAP system clone operation using SnapCenter.


1. Prepare the target host.
2. SnapCenter clone create workflow for the SAP HANA shared volume.
3. Start SAP HANA services.
4. SnapCenter clone create workflow for the SAP HANA data volume including database recovery.
5. The SAP HANA system can now be used as a repair system.



If you must reset the system to a different Snapshot backup, then step 6 and step 4 are sufficient. The SAP HANA shared volume can continue to be mounted.

If the system is not needed anymore, the clean-up process is performed with the following steps.

6. SnapCenter clone delete workflow for the SAP HANA data volume including database shutdown.
7. Stop SAP HANA services.
8. SnapCenter clone delete workflow for the SAP HANA shared volume.



## Prerequisites and limitations

The workflows described in the following sections have a few prerequisites and limitations regarding the SAP HANA system architecture and the SnapCenter configuration.

- The described workflow is valid for single host SAP HANA MDC systems. Multiple host systems are not supported.
- The SnapCenter SAP HANA plug-in must be deployed on the target host to enable the execution of automation scripts.
- The workflow has been validated for NFS. The automation script `sc-mount-volume.sh`, which is used to mount the SAP HANA shared volume, does not support FCP. This step must be either done manually or by extending the script.

- The described workflow is only valid for the SnapCenter 5.0 release or higher.

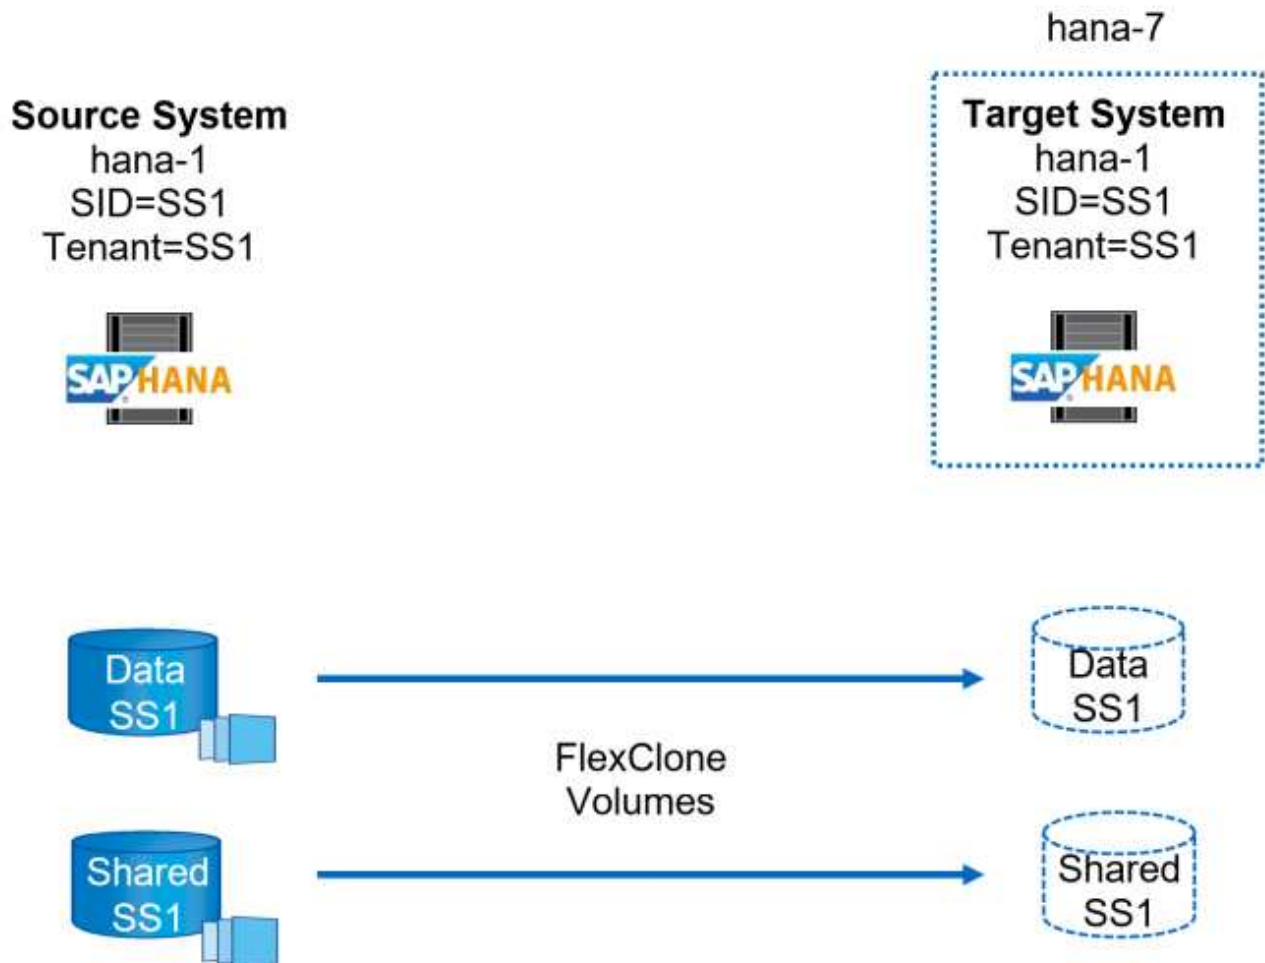
## Lab setup

The figure below shows the lab setup used for system clone operation.

The following software versions were used:

- SnapCenter 5.0
- SAP HANA systems: HANA 2.0 SPS6 rev.61
- SLES 15
- ONTAP 9.7P7

All SAP HANA systems must be configured based on the configuration guide [SAP HANA on NetApp AFF systems with NFS](#). SnapCenter and the SAP HANA resources were configured based on the best practice guide [SAP HANA Backup and Recovery with SnapCenter](#).



## Target host preparation

This section describes the preparation steps required at a server that is used as a system clone target.

During normal operation, the target host might be used for other purposes, for example, as an SAP HANA QA or test system. Therefore, most of the described steps must be executed when the system clone operation is

requested. 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 target host preparation also includes shutting down the SAP 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.



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.



In our lab setup, we changed the host name of the target system only internally from the target system perspective. Externally the host was still accessible with the hostname `hana-7`. When logged into the host, the host itself is `hana-1`.

### 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.

The SnapCenter SAP HANA plug-in must be deployed on the target host using the add host operation within SnapCenter.

### 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 SAP 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.

```
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/SS1/HDB00/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/SS1/HDB00/exe/sapstartsrv
pf=/usr/sap/SS1/SYS/profile/SS1_HDB00_hana-1 -D -u ssladm
limit.descriptors=1048576
```

### Prepare log and log backup volume

Because you do not need to clone the log volume from the source system and any recovery is performed with the clear log option, an empty log volume must be prepared at the target host.

Because the source system has been configured with a separate log backup volume, an empty log backup volume must be prepared and mounted to the same mount point as at the source system.

```
hana-1:/# cat /etc/fstab
192.168.175.117:/SS1_repair_log_mnt00001 /hana/log/SS1/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0
0
192.168.175.117:/SS1_repair_log_backup /mnt/log-backup nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0
0
```

Within the log volume hdb\*, you must create subdirectories in the same way as at the source system.

```
hana-1:/ # ls -al /hana/log/SS1/mnt00001/
total 16
drwxrwxrwx 5 root root 4096 Dec 1 06:15 .
drwxrwxrwx 1 root root 16 Nov 30 08:56 ..
drwxr-xr-- 2 ssladm sapsys 4096 Dec 1 06:14 hdb00001
drwxr-xr-- 2 ssladm sapsys 4096 Dec 1 06:15 hdb00002.00003
drwxr-xr-- 2 ssladm sapsys 4096 Dec 1 06:15 hdb00003.00003
```

Within the log backup volume, you must create subdirectories for the system and the tenant database.

```
hana-1:/ # ls -al /mnt/log-backup/
total 12
drwxr-xr-- 2 ssladm sapsys 4096 Dec 1 04:48 .
drwxr-xr-- 2 ssladm sapsys 4896 Dec 1 03:42 ..
drwxr-xr-- 2 ssladm sapsys 4096 Dec 1 06:15 DB_SS1
drwxr-xr-- 2 ssladm sapsys 4096 Dec 1 06:14 SYSTEMDB
```

## Prepare file system mounts

You must prepare mount points for the data and the shared volume.

With our example, the directories /hana/data/SS1/mnt00001, /hana/shared and usr/sap/SS1 must be created.


## Prepare script execution

You must add the scripts, that should be executed at the target system to the SnapCenter allowed commands config file.

```
hana-7:/opt/NetApp/snapcenter/scc/etc # cat
/opt/NetApp/snapcenter/scc/etc/allowed_commands.config
command: mount
command: umount
command: /mnt/sapcc-share/SAP-System-Refresh/sc-system-refresh.sh
command: /mnt/sapcc-share/SAP-System-Refresh/sc-mount-volume.sh
hana-7:/opt/NetApp/snapcenter/scc/etc #
```


## Cloning the HANA shared volume

1. Select a Snapshot backup from the source system SS1 shared volume and click Clone.



Backup Name	Count	End Date
SnapCenter_LocalSnap_Hourly_05-13-2022_05.04.01.8012	1	05/13/2022 5:04:12 AM
SnapCenter_LocalSnap_Hourly_05-13-2022_01.04.01.9799	1	05/13/2022 1:04:12 AM
SnapCenter_LocalSnap_Hourly_05-12-2022_21.04.01.8899	1	05/12/2022 9:04:12 PM

1. Select the host where the target repair system has been prepared. The NFS export IP address must be the storage network interface of the target host. As target SID keep the same SID as the source system. In our example SS1.



**Clone From Backup**

**1 Location**

Select the host to create the clone

Plug-In host: hana-7.sapcc.stl.netapp.com

Target Clone SID: SS1

NFS Export IP Address: 192.168.175.75

2 Scripts

3 Notification

4 Summary

3. Enter the mount script with the required command line options.



The SAP HANA system uses a single volume for /hana/shared as well as for /usr/sap/SS1, separated in subdirectories as recommended in the configuration guide [SAP HANA on NetApp AFF systems with NFS](#). The script `sc-mount-volume.sh` supports this configuration using a special command line option for the mount path. If the mount path command line option is equal to `usr-sap-and-shared`, the script mounts the subdirectories `shared` and `usr-sap` in the volume accordingly.

Clone From Backup

1 Location

2 Scripts

3 Notification

4 Summary

Enter optional commands to run before performing a clone operation

Pre clone command

Enter optional commands to mount a file system to a host

Mount command

Enter optional commands to run after performing a clone operation


Post clone command

Configure an SMTP Server to send email notifications for Clone jobs by going to [Settings>Global Settings>Notification Server Settings.](#)

Previous

Next

4. The Job Details screen in SnapCenter shows the progress of the operation.



5. The logfile of the `sc-mount-volume.sh` script shows the different steps executed for the mount operation.

```

20201201041441###hana-1###sc-mount-volume.sh: Adding entry in /etc/fstab.
20201201041441###hana-1###sc-mount-volume.sh:
192.168.175.117://SS1_shared_Clone_05132205140448713/usr-sap /usr/sap/SS1
nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0
0
20201201041441###hana-1###sc-mount-volume.sh: Mounting volume: mount
/usr/sap/SS1.
20201201041441###hana-1###sc-mount-volume.sh:
192.168.175.117://SS1_shared_Clone_05132205140448713/shared /hana/shared
nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0
0
20201201041441###hana-1###sc-mount-volume.sh: Mounting volume: mount
/hana/shared.
20201201041441###hana-1###sc-mount-volume.sh: usr-sap-and-shared mounted
successfully.
20201201041441###hana-1###sc-mount-volume.sh: Change ownership to ssladm.

```


6. When the SnapCenter workflow is finished, the /usr/sap/SS1 and the /hana/shared filesystems are mounted at the target host.

```

hana-1:~ # df
Filesystem 1K-blocks Used Available Use% Mounted on
192.168.175.117://SS1_repair_log_mnt00001 262144000 320 262143680 1%
/hana/log/SS1/mnt00001
192.168.175.100:/sapcc_share 1020055552 53485568 966569984 6% /mnt/sapcc-
share
192.168.175.117://SS1_repair_log_backup 104857600 256 104857344 1%
/mnt/log-backup
192.168.175.117://SS1_shared_Clone_05132205140448713/usr-sap 262144064
10084608 252059456 4% /usr/sap/SS1
192.168.175.117://SS1_shared_Clone_05132205140448713/shared 262144064
10084608 252059456 4% /hana/shared

```

7. Within SnapCenter, a new resource for the cloned volume is visible.



The screenshot shows the NetApp SnapCenter web interface. The top navigation bar includes the NetApp logo, 'SnapCenter', and user information. A sidebar on the left contains navigation links: Dashboard, Resources, Monitor, Reports, Hosts, Storage Systems, Settings, and Alerts. The main content area displays a table of resources under the 'Resources' tab. The table has columns for Name, Associated System ID (SID), Plug-in Host, Resource Groups, Policies, Last backup, and Overall Status. Two resources are listed: 'SS1-Shared-Volume' and 'SS1-Shared-Volume'.

Name	Associated System ID (SID)	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
SS1-Shared-Volume	SS1	hana-1.sapcc.stl.netapp.com		LocalSnap LocalSnap-OnDemand	05/13/2022 5:04:12 AM	Backup succeeded
SS1-Shared-Volume	SS1	hana-7.sapcc.stl.netapp.com				Not protected

8. Now that the /hana/shared volume is available, the SAP HANA services can be started.

```
hana-1:/mnt/sapcc-share/SAP-System-Refresh # systemctl start sapinit
```

9. SAP Host Agent and sapstartsrv processes are now started.

```
hana-1:/mnt/sapcc-share/SAP-System-Refresh # ps -ef |grep sap
root 12377 1 0 04:34 ? 00:00:00 /usr/sap/hostctrl/exe/saphostexec
pf=/usr/sap/hostctrl/exe/host_profile
sapadm 12403 1 0 04:34 ? 00:00:00 /usr/lib/systemd/systemd --user
sapadm 12404 12403 0 04:34 ? 00:00:00 (sd-pam)
sapadm 12434 1 1 04:34 ? 00:00:00 /usr/sap/hostctrl/exe/sapstartsrv
pf=/usr/sap/hostctrl/exe/host_profile -D
root 12485 12377 0 04:34 ? 00:00:00 /usr/sap/hostctrl/exe/saphostexec
pf=/usr/sap/hostctrl/exe/host_profile
root 12486 12485 0 04:34 ? 00:00:00 /usr/sap/hostctrl/exe/saposcol -l -w60
pf=/usr/sap/hostctrl/exe/host_profile
ssladm 12504 1 0 04:34 ? 00:00:00 /usr/sap/SS1/HDB00/exe/sapstartsrv
pf=/usr/sap/SS1/SYS/profile/SS1_HDB00_hana-1 -D -u ssladm
root 12582 12486 0 04:34 ? 00:00:00 /usr/sap/hostctrl/exe/saposcol -l -w60
pf=/usr/sap/hostctrl/exe/host_profile
root 12585 7613 0 04:34 pts/0 00:00:00 grep --color=auto sap
hana-1:/mnt/sapcc-share/SAP-System-Refresh #
```

## Cloning additional SAP application services

Additional SAP application services are cloned in the same way as the SAP HANA shared volume as described in the section “Cloning the SAP HANA shared volume.” Of course, the required storage volume(s) of the SAP application servers must be protected with SnapCenter as well.

You must add the required services entries to /usr/sap/sapservices, and the ports, users, and the file system mount points (for example, /usr/sap/SID) must be prepared.

## Cloning the data volume and recovery of the HANA database

1. Select an SAP HANA Snapshot backup from the source system SS1.

The screenshot shows the NetApp SnapCenter web interface. On the left, a sidebar lists systems: SAP HANA, Q51, SM1, SS1 (selected), SS2, and SS2. The main area displays the 'SS1 Topology' with a 'Manage Copies' section showing 15 Backups (0 Clones) for Local copies and 11 Backups (0 Clones) for Vault copies. A 'Summary Card' on the right shows 28 Backups, 26 Snapshot based backups, 2 File-Based backups, and 0 Clones. Below this, a table lists 'Primary Backup(s)' with columns for Backup Name, Count, and End Date. The table contains four entries, all with a count of 1.

Backup Name	Count	End Date
SnapCenter_LocalSnapAndSnapVault_Daily_05-13-2022_05.00.03.0030	1	05/13/2022 5:01:01 AM
SnapCenter_LocalSnap_Hourly_05-13-2022_03.00.01.8016	1	05/13/2022 3:01:00 AM
SnapCenter_LocalSnap_Hourly_05-12-2022_23.00.01.8743	1	05/12/2022 11:01:00 PM
SnapCenter_LocalSnap_Hourly_05-12-2022_19.00.01.9803	1	05/12/2022 7:01:00 PM

2. Select the host where the target repair system has been prepared. The NFS export IP address must be the storage network interface of the target host. As target SID keep the same SID as the source system. In our example SS1

Clone From Backup

1 Location

2 Scripts

3 Notification

4 Summary

Select the host to create the clone

Plug-in host

hana-7.sapcc.stl.netapp.com

i

Target Clone SID

SS1

i

NFS Export IP Address

192.168.175.75

i

3. Enter the post-clone scripts with the required command line options.



The script for the recovery operation recovers the SAP HANA database to the point in time of the Snapshot operation and does not execute any forward recovery. If a forward recovery to a specific point in time is required, the recovery must be performed manually. A manual forward recovery also requires that the log backups from the source system are available at the target host.

Clone From Backup

1 Location

2 Scripts

3 Notification

4 Summary

The following commands will run on the Plug-in Host: hana-7.sapcc.stl.netapp.com

Enter optional commands to run before performing a clone operation i

Pre clone command

Enter optional commands to run after performing a clone operation i

Post clone command

/mnt/sapcc-share/SAP-System-Refresh/sc-system-refresh.sh  
recover

The job details screen in SnapCenter shows the progress of the operation.

Job Details

Clone from backup 'SnapCenter\_LocalSnap\_Hourly\_05-13-2022\_03.00.01.8016'

✓

▼ Clone from backup 'SnapCenter\_LocalSnap\_Hourly\_05-13-2022\_03.00.01.8016'

✓

▼ hana-7.sapcc.stl.netapp.com

✓

▼ Clone

✓

▶ Application Pre Clone

✓

▶ Storage Clone

✓

▶ Application Post Clone

✓

▶ Register Clone Metadata

✓

▶ Application Clean-Up

✓

▶ Data Collection

✓

▶ Agent Finalize Workflow

Task Name: Clone Start Time: 05/13/2022 5:24:36 AM End Time: 05/13/2022 5:25:05 AM

View Logs

Cancel Job

Close

The logfile of the `sc-system-refresh` script shows the different steps that are executed for the mount and the recovery operation.

54

```

20201201052124###hana-1###sc-system-refresh.sh: Recover system database.
20201201052124###hana-1###sc-system-refresh.sh:
/usr/sap/SS1/HDB00/exe/Python/bin/python
/usr/sap/SS1/HDB00/exe/python_support/recoverSys.py --command "RECOVER
DATA USING SNAPSHOT CLEAR LOG"
20201201052156###hana-1###sc-system-refresh.sh: Wait until SAP HANA
database is started ....
20201201052156###hana-1###sc-system-refresh.sh: Status: GRAY
20201201052206###hana-1###sc-system-refresh.sh: Status: GREEN
20201201052206###hana-1###sc-system-refresh.sh: SAP HANA database is
started.
20201201052206###hana-1###sc-system-refresh.sh: Source system has a single
tenant and tenant name is identical to source SID: SS1
20201201052206###hana-1###sc-system-refresh.sh: Target tenant will have
the same name as target SID: SS1.
20201201052206###hana-1###sc-system-refresh.sh: Recover tenant database
SS1.
20201201052206###hana-1###sc-system-refresh.sh:
/usr/sap/SS1/SYS/exe/hdb/hdbsql -U SS1KEY RECOVER DATA FOR SS1 USING
SNAPSHOT CLEAR LOG
0 rows affected (overall time 34.773885 sec; server time 34.772398 sec)
20201201052241###hana-1###sc-system-refresh.sh: Checking availability of
Indexserver for tenant SS1.
20201201052241###hana-1###sc-system-refresh.sh: Recovery of tenant
database SS1 succesfully finished.
20201201052241###hana-1###sc-system-refresh.sh: Status: GREEN
After the recovery operation, the HANA database is running and the data
volume is mounted at the target host.
hana-1:/mnt/log-backup # df
Filesystem 1K-blocks Used Available Use% Mounted on
192.168.175.117:/SS1_repair_log_mnt00001 262144000 760320 261383680 1%
/hana/log/SS1/mnt00001
192.168.175.100:/sapcc_share 1020055552 53486592 966568960 6% /mnt/sapcc-
share
192.168.175.117:/SS1_repair_log_backup 104857600 512 104857088 1%
/mnt/log-backup
192.168.175.117:/SS1_shared_Clone_05132205140448713/usr-sap 262144064
10090496 252053568 4% /usr/sap/SS1
192.168.175.117:/SS1_shared_Clone_05132205140448713/shared 262144064
10090496 252053568 4% /hana/shared
192.168.175.117:/SS1_data_mnt00001_Clone_0421220520054605 262144064
3732864 258411200 2% /hana/data/SS1/mnt00001

```

The SAP HANA system is now available and can be used, for example, as a repair system.

## Where to find additional information and version history

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

- [SAP Business Application and SAP HANA Database Solutions \(netapp.com\)](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4436: SAP HANA on NetApp All Flash FAS Systems with Fibre Channel Protocol](#)
- [TR-4435: SAP HANA on NetApp All Flash FAS Systems with NFS](#)
- [TR-4926: SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter](#)
- [TR-4953: NetApp SAP Landscape Management Integration using Ansible](#)
- [TR-4929: Automating SAP system copy operations with Libelle SystemCopy \(netapp.com\)](#)
- [Automating SAP system copy; refresh; and clone workflows with ALPACA and NetApp SnapCenter](#)
- [Automating SAP system copy; refresh; and clone workflows with Avantra and NetApp SnapCenter](#)

Version	Date	Document Version History
Version 1.0	February 2018	Initial release.
Version 2.0	February 2021	Complete rewrite covering SnapCenter 4.3 and improved automation scripts. New workflow description for system refresh and system clone operations.
Version 3.0	May 2022	Adapted to changed workflow with SnapCenter 4.6 P1
Version 4.0	July 2024	Document covers NetApp systems on-premises, FSx for ONTAP and Azure NetApp Files New SnapCenter 5.0 operations mount and unmount during clone create and delete workflows Added specific steps for Fibre Channel SAN Added specific steps for Azure NetApp Files Adapted and simplified <code>sc-system-refresh</code> script Included required steps for enabled SAP HANA volume encryption

## Copyright information

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

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

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

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

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

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

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

## Trademark information

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