



# **SAP Business Application and SAP HANA Database Solutions**

NetApp Solutions SAP

NetApp  
August 01, 2025

# Table of Contents

- SAP Business Application and SAP HANA Database Solutions ..... 1
- Best Practices ..... 2
  - SAP HANA on NetApp AFF Systems with FCP Configuration Guide ..... 2
    - SAP HANA on NetApp AFF Systems with Fibre Channel Protocol ..... 2
    - SAP HANA using VMware vSphere ..... 3
    - Architecture ..... 4
    - Storage sizing ..... 8
    - Infrastructure setup and configuration ..... 13
    - Where to find additional information ..... 49
    - Update history ..... 50
  - SAP HANA on NetApp AFF Systems with NFS Configuration Guide ..... 50
    - SAP HANA on NetApp AFF Systems with NFS - Configuration Guide ..... 51
    - Architecture ..... 53
    - Storage sizing ..... 58
    - Infrastructure setup and configuration ..... 63
    - Where to find additional information ..... 90
    - Update history ..... 91
  - SAP HANA on NetApp ASA Systems with FCP Configuration Guide ..... 92
    - SAP HANA on NetApp ASA Systems with Fibre Channel Protocol ..... 92
    - SAP HANA using VMware vSphere ..... 93
    - Architecture ..... 94
    - Storage sizing ..... 97
    - Infrastructure setup and configuration ..... 102
    - Where to find additional information ..... 137
    - Update history ..... 137
  - SAP HANA on NetApp FAS Systems with NFS Configuration Guide ..... 138
    - SAP HANA on NetApp FAS systems with NFS Configuration guide ..... 138
    - Architecture ..... 140
    - Storage sizing ..... 145
    - Infrastructure setup and configuration ..... 150
    - Where to find additional information ..... 180
    - Update history ..... 181
  - SAP HANA on FAS Systems with FCP Configuration Guide ..... 182
    - SAP HANA on NetApp FAS Systems with Fibre Channel Protocol Configuration Guide ..... 182
    - Architecture ..... 184
    - Storage sizing ..... 188
    - Infrastructure setup and configuration ..... 193
    - Where to find additional information ..... 232
    - Update history ..... 233
- Backup, Restore and Disaster Recovery ..... 234
  - SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter ..... 234
    - TR-4926: SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter... 234
    - Backup and recovery using Amazon FSx for ONTAP ..... 234

SnapCenter architecture	238
SnapCenter configuration	243
SnapCenter backup operations	260
Backup of non-data volumes	270
Restore and recover	277
Backup replication with SnapVault	285
Where to find additional information	299
SAP HANA backup and recovery with SnapCenter	300
TR-4614: SAP HANA backup and recovery with SnapCenter	300
SnapCenter architecture	305
SnapCenter SAP HANA backup solution	305
SnapCenter concepts and best practices	309
Lab setup used for this report	327
SnapCenter configuration	328
Initial SnapCenter configuration	330
SnapCenter resource-specific configuration for SAP HANA database backups	343
SnapCenter resource-specific configuration for non-data volume backups	362
Database backups	367
Block integrity check	376
Restore and recovery	380
Advanced configuration and tuning	433
Where to find additional information and version history	441
SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster	443
SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster	443
Overview SAP HANA high availability	444
Example configuration overview	447
HANA system provisioning and installation	448
SnapMirror active sync configuration	456
SnapCenter configuration	462
SnapCenter backup operations	467
SnapCenter restore and recovery	470
SAP System refresh operation	472
SnapCenter non-data volumes	472
Failover scenarios	474
Additional information and version history	478
SAP HANA data protection with SnapCenter with VMware VMFS and NetApp ASA systems	478
SAP HANA data protection with SnapCenter with VMware VMFS and NetApp ASA systems	479
Lab setup used for this document	479
HANA system provisioning and installation	480
HANA configuration	487
SnapCenter configuration	488
Backup operations	495
Restore and recovery operations	498

SAP System Refresh .....	502
Additional information and version history .....	511
BlueXP Backup and Recovery for SAP HANA - Cloud Object storage as backup destination .....	512
BlueXP Backup and Recovery for SAP HANA - Cloud Object storage as backup destination .....	512
Configuring BlueXP Backup and Recovery for SAP HANA .....	514
Restoring SAP HANA BlueXP Backup .....	531
Additional Information and Version History .....	533
SAP HANA System Replication Backup and Recovery with SnapCenter .....	533
TR-4719: SAP HANA System Replication - Backup and Recovery with SnapCenter .....	534
Storage Snapshot backups and SAP System Replication .....	535
SnapCenter configuration options for SAP System Replication .....	536
SnapCenter 4.6 configuration using a resource group .....	538
SnapCenter configuration with a single resource .....	549
Restore and recovery from a backup created at the other host .....	562
Where to find additional information .....	566
Version history .....	566
SAP HANA Disaster Recovery with Azure NetApp Files .....	566
TR-4891: SAP HANA disaster recovery with Azure NetApp Files .....	567
Disaster recovery solution comparison .....	569
ANF Cross-Region Replication with SAP HANA .....	573
Disaster recovery testing .....	585
Disaster recovery failover .....	598
Update history .....	609
TR-4646: SAP HANA Disaster Recovery with Storage Replication .....	609
TR-4711: SAP HANA Backup and Recovery Using NetApp Storage Systems and Commvault Software .....	610
SnapCenter Integration for SAP ASE Database .....	610
Introduction .....	610
Additional information and version history .....	623
SnapCenter Integration for IBM DB2 Database .....	624
Introduction .....	624
Example configuration overview .....	625
Demo Environment .....	625
Additional information and version history .....	632
SnapCenter Integration for SAP MaxDB Database .....	633
Introduction .....	633
Example configuration overview .....	633
Demo Environment .....	634
Software versions .....	634
MaxDB Volume Design .....	634
Steps to Protect Database M02 .....	635
Prerequisites on Host .....	635
Prerequisites for the Database – Create Backup Templates, Enable Logbackup .....	635
Deploy SnapCenter Agent to Host sap-Inx25 .....	636
Create SnapCenter Resource Configuration for Database M02 .....	637
Sequence to Recover System M02 .....	644

Recover Instance M02 . . . . .	644
Additional information and version history . . . . .	650
Lifecycle Management . . . . .	651
NetApp SAP Landscape Management Integration using Ansible . . . . .	651
TR-4953: NetApp SAP Landscape Management Integration using Ansible . . . . .	651
SAP system clone, copy, and refresh scenarios . . . . .	651
Use cases for system refresh, copy, and cloning . . . . .	652
NetApp SAP LaMa integration using Ansible . . . . .	655
Example implementation . . . . .	655
SAP LaMa provisioning workflow - clone system . . . . .	660
SAP LaMa deprovisioning workflow - system destroy . . . . .	668
SAP LaMa provisioning workflow - copy system . . . . .	671
SAP LaMa provisioning workflow - system refresh . . . . .	675
Provider script configuration and Ansible playbooks . . . . .	677
Conclusion . . . . .	690
Automating SAP HANA System Copy and Clone Operations with SnapCenter . . . . .	691
TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter . . . . .	691
SAP system copy, refresh, and clone scenarios . . . . .	693
Use cases for system refresh and cloning . . . . .	694
Supported infrastructure and scenarios . . . . .	697
Overview of SAP system refresh workflow with SnapCenter . . . . .	697
Overview of SAP system clone workflow with SnapCenter . . . . .	700
Considerations for SAP HANA system refresh operations using storage snapshot backups . . . . .	701
Automation example scripts . . . . .	706
SAP HANA system refresh with SnapCenter . . . . .	709
SAP system clone with SnapCenter . . . . .	734
Where to find additional information and version history . . . . .	747
Automating SAP system copy operations with Libelle SystemCopy . . . . .	747
TR-4929: Automating SAP system copy operations with Libelle SystemCopy . . . . .	747
SAP HANA system refresh with LSC and SnapCenter . . . . .	750
SAP HANA system refresh with LSC, AzAcSnap, and Azure NetApp Files . . . . .	762
Where to find additional information and version history . . . . .	774
Solution Briefs . . . . .	776
SB-3978: Lifecycle Management for SAP HANA . . . . .	776
SB-3965: Backup and Recovery for SAP HANA . . . . .	776
The challenge . . . . .	776
The Solution . . . . .	776
SB-3968: Disaster Recovery for SAP HANA . . . . .	776
The challenge . . . . .	776
The Solution . . . . .	777
SB-4292: SAP automation with Ansible . . . . .	777
Solution overview . . . . .	777
Conclusion . . . . .	780
Where to find additional information . . . . .	780
Version history . . . . .	781

SB-4293: Automating SAP system copy, refresh, and clone workflows with ALPACA and NetApp	
SnapCenter .....	781
Solution overview .....	781
Conclusion .....	785
Where to find additional information .....	785
Version history .....	786
SB-4294: Automating SAP system copy, refresh, and clone workflows with Avantra and NetApp	
SnapCenter .....	786
Solution overview .....	786
Conclusion .....	789
Where to find additional information .....	790
Version history .....	790
Legal notices .....	791
Copyright .....	791
Trademarks .....	791
Patents .....	791
Privacy policy .....	791

# **SAP Business Application and SAP HANA Database Solutions**

# Best Practices

## SAP HANA on NetApp AFF Systems with FCP Configuration Guide

### SAP HANA on NetApp AFF Systems with Fibre Channel Protocol

The NetApp AFF product family is certified for use with SAP HANA in TDI projects. This guide provides best practices for SAP HANA on this platform for FCP.

Marco Schoen, NetApp

#### Introduction

The NetApp AFF A-Series and AFF C-Series product families have been certified for use with SAP HANA in tailored data center integration (TDI) projects.

This certification is valid for the following models:

- AFF A150, AFF A20, AFF A250, AFF A30, AFF A400, AFF A50, AFF A70, AFF A800, AFF A90, AFF A900, AFF A1K
- AFF C250, AFF C30, AFF C400, AFF C60, AFF C800, AFF C80
- ASA A250, ASA A400, ASA A800, ASA A900
- ASA C250, ASA C400, ASA C800



NetApp AFF and ASA C-Series requires NetApp ONTAP 9.13.1 or later

For a complete list of NetApp certified storage solutions for SAP HANA, see the [Certified and supported SAP HANA hardware directory](#).

This document describes AFF configurations that use the Fibre Channel Protocol (FCP).



The configuration described in this paper is necessary to achieve the required SAP HANA KPIs and the best performance for SAP HANA. Changing any settings or using features not listed herein might cause performance degradation or unexpected behavior and should only be done if advised by NetApp support.

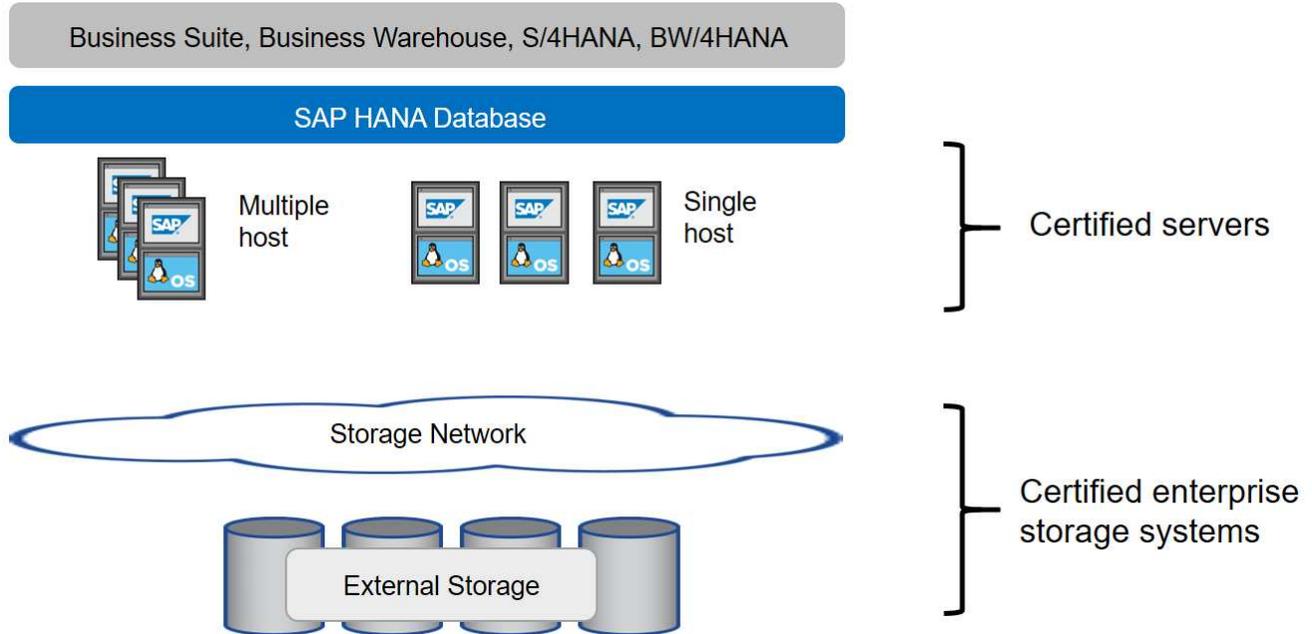
The configuration guides for AFF systems using NFS and NetApp FAS systems can be found using the following links:

- [SAP HANA on NetApp FAS Systems with FCP](#)
- [SAP HANA on NetApp ASA Systems with FCP](#)
- [SAP HANA on NetApp FAS Systems with NFS](#)
- [SAP HANA on NetApp AFF Systems with NFS](#)

In an SAP HANA multiple-host environment, the standard SAP HANA storage connector is used to provide fencing in the event of an SAP HANA host failover. Always refer to the relevant SAP notes for operating system configuration guidelines and HANA specific Linux kernel dependencies. For more information, see [SAP Note](#)

## SAP HANA tailored data center integration

NetApp AFF storage systems are certified in the SAP HANA TDI program using both NFS (NAS) and FC (SAN) protocols. They can be deployed in any of the current SAP HANA scenarios, such as SAP Business Suite on HANA, S/4HANA, BW/4HANA, or SAP Business Warehouse on HANA in either single-host or multiple-host configurations. Any server that is certified for use with SAP HANA can be combined with NetApp certified storage solutions. The following figure shows an architecture overview.



For more information regarding the prerequisites and recommendations for productive SAP HANA systems, see the following resource:

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

## SAP HANA using VMware vSphere

There are several options to connect storage to virtual machines (VMs). The preferred one is to connect the storage volumes with NFS directly out of the guest operating system. This option is described in [SAP HANA on NetApp AFF Systems with NFS](#).

Raw device mappings (RDM), FCP datastores, or VVOL datastores with FCP are supported as well. For both datastore options, only one SAP HANA data or log volume must be stored within the datastore for productive use cases.

For more information about using vSphere with SAP HANA, see the following links:

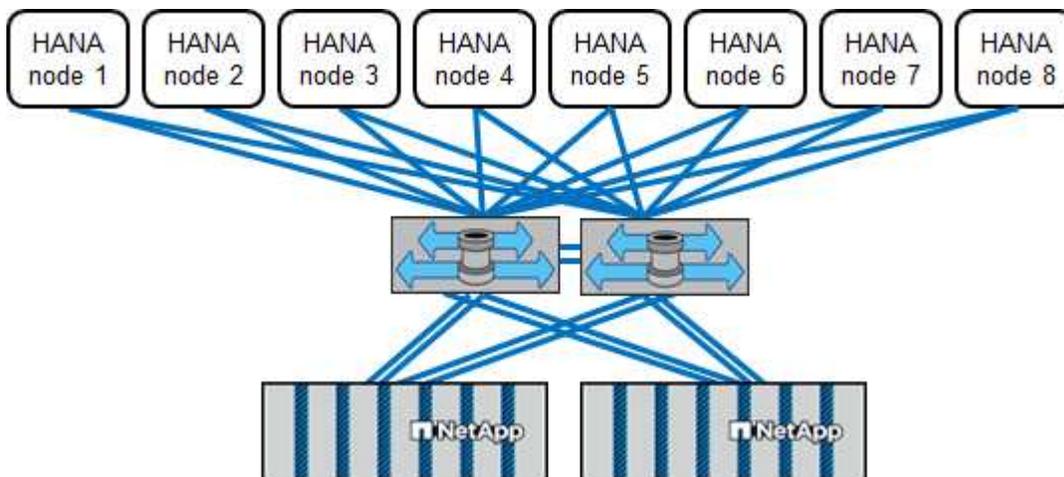
- [SAP HANA on VMware vSphere - Virtualization - Community Wiki](#)
- [SAP HANA on VMware vSphere Best Practices Guide](#)
- [2161991 - VMware vSphere configuration guidelines - SAP ONE Support Launchpad \(Login required\)](#)

## Architecture

SAP HANA hosts are connected to storage controllers using a redundant FCP infrastructure and multipath software. A redundant FCP switch infrastructure is required to provide fault-tolerant SAP HANA host-to-storage connectivity in case of switch or host bus adapter (HBA) failure. Appropriate zoning must be configured at the switch to allow all HANA hosts to reach the required LUNs on the storage controllers.

Different models of the AFF system product family can be mixed and matched at the storage layer to allow for growth and differing performance and capacity needs. The maximum number of SAP HANA hosts that can be attached to the storage system is defined by the SAP HANA performance requirements and the model of NetApp controller used. The number of required disk shelves is only determined by the capacity and performance requirements of the SAP HANA systems.

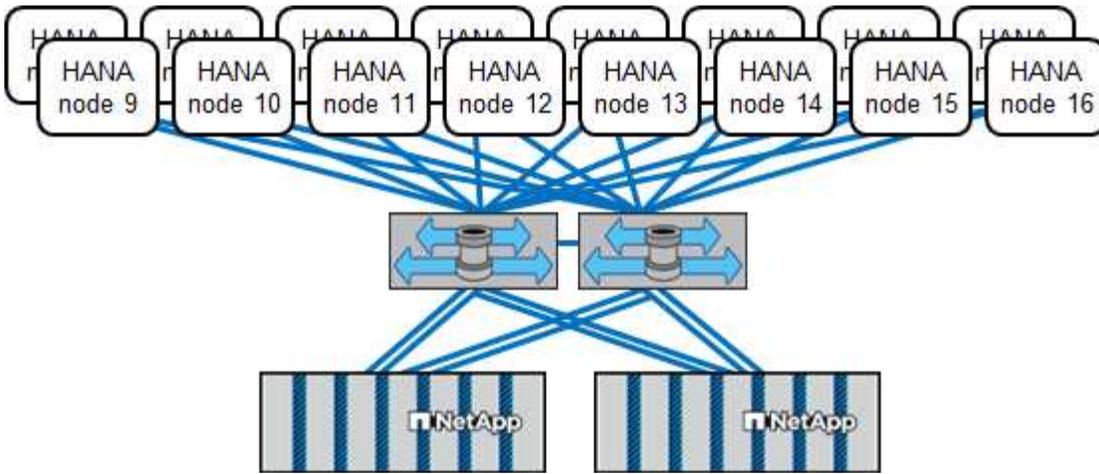
The following figure shows an example configuration with eight SAP HANA hosts attached to a storage HA pair.



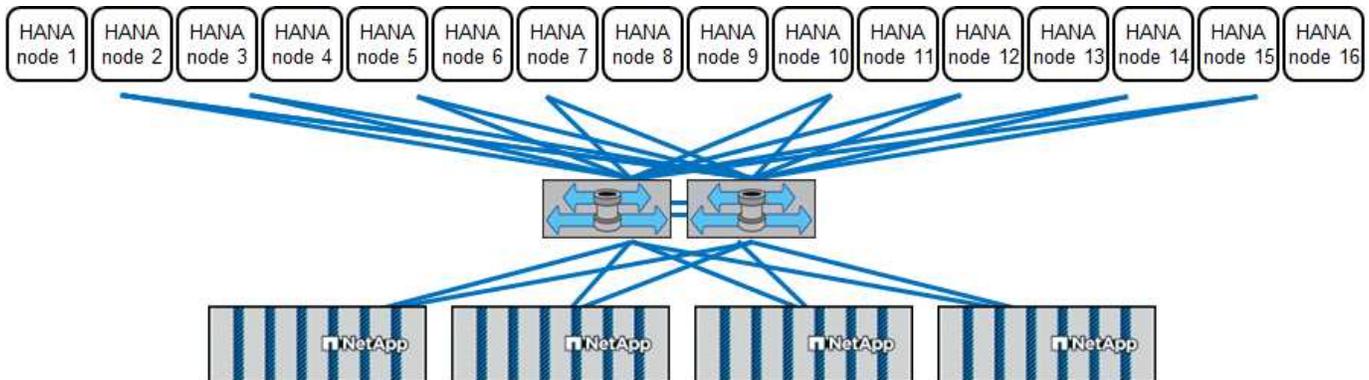
This architecture can be scaled in two dimensions:

- By attaching additional SAP HANA hosts and storage capacity to the existing storage, if the storage controllers provide enough performance to meet the current SAP HANA KPIs
- By adding more storage systems with additional storage capacity for the additional SAP HANA hosts

The following figure shows a configuration example in which more SAP HANA hosts are attached to the storage controllers. In this example, more disk shelves are necessary to meet the capacity and performance requirements of the 16 SAP HANA hosts. Depending on the total throughput requirements, you must add additional FC connections to the storage controllers.



Independent of the deployed AFF system, the SAP HANA landscape can also be scaled by adding any certified storage controllers to meet the desired node density, as shown in the following figure.



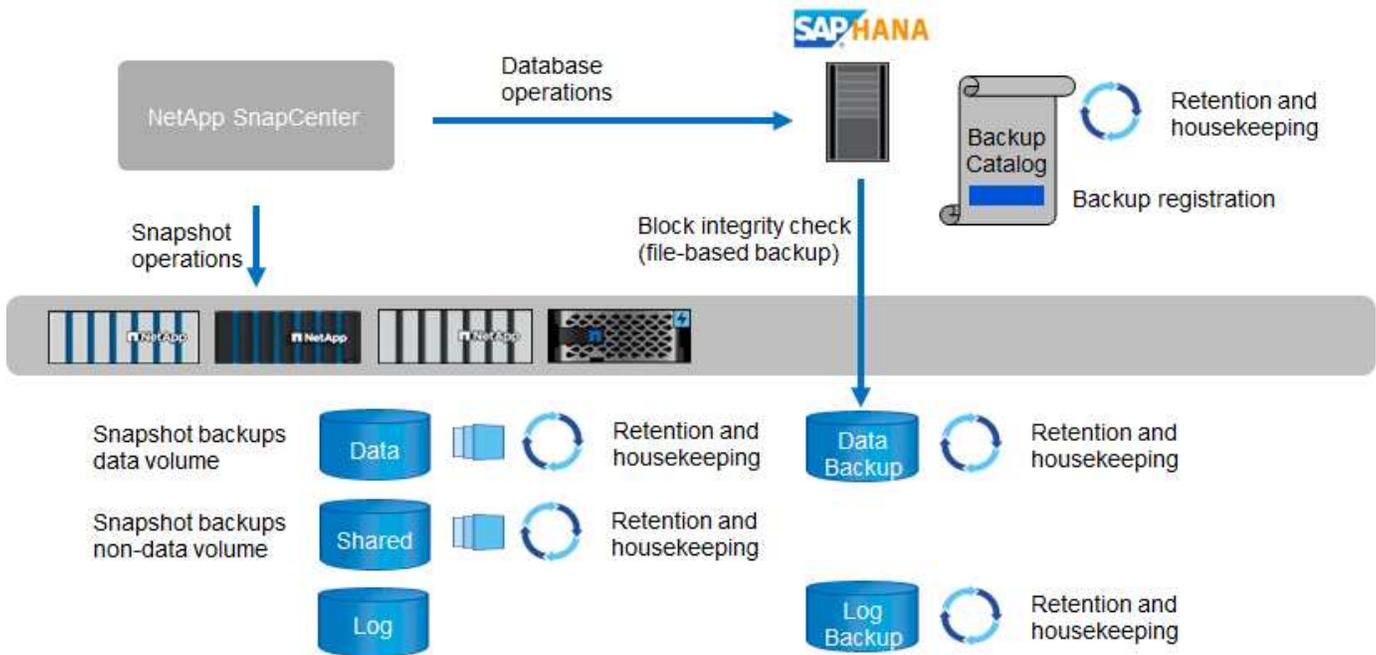
### SAP HANA backup

The ONTAP software present on all NetApp storage controllers provides a built-in mechanism to back up SAP HANA databases while in operation with no effect on performance. Storage-based NetApp Snapshot backups are a fully supported and integrated backup solution available for SAP HANA single containers and for SAP HANA MDC systems with a single tenant or multiple tenants.

Storage-based Snapshot backups are implemented by using the NetApp SnapCenter plug-in for SAP HANA. This allows users to create consistent storage-based Snapshot backups by using the interfaces provided natively by SAP HANA databases. SnapCenter registers each of the Snapshot backups into the SAP HANA backup catalog. Therefore, backups taken by SnapCenter are visible within SAP HANA Studio or Cockpit where they can be selected directly for restore and recovery operations.

NetApp SnapMirror technology allows for Snapshot copies that were created on one storage system to be replicated to a secondary backup storage system that is controlled by SnapCenter. Different backup retention policies can then be defined for each of the backup sets on the primary storage and also for the backup sets on the secondary storage systems. The SnapCenter Plug-in for SAP HANA automatically manages the retention of Snapshot copy-based data backups and log backups, including the housekeeping of the backup catalog. The SnapCenter Plug-in for SAP HANA also allows for the execution of a block integrity check of the SAP HANA database by executing a file-based backup.

The database logs can be backed up directly to the secondary storage by using an NFS mount, as shown in the following figure.



Storage-based Snapshot backups provide significant advantages compared to conventional file-based backups. These advantages include, but are not limited to the following:

- Faster backup (a few minutes)
- Reduced RTO due to a much faster restore time on the storage layer (a few minutes) as well as more frequent backups
- No performance degradation of the SAP HANA database host, network, or storage during backup and recovery operations
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes

For detailed information about the SAP HANA backup and recovery solution, see [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

### SAP HANA disaster recovery

SAP HANA disaster recovery can be done either on the database layer by using SAP HANA system replication or on the storage layer by using storage replication technologies. The following section provides an overview of disaster recovery solutions based on storage replication.

For detailed information about the SAP HANA disaster recovery solutions, see [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#).

#### Storage replication based on SnapMirror

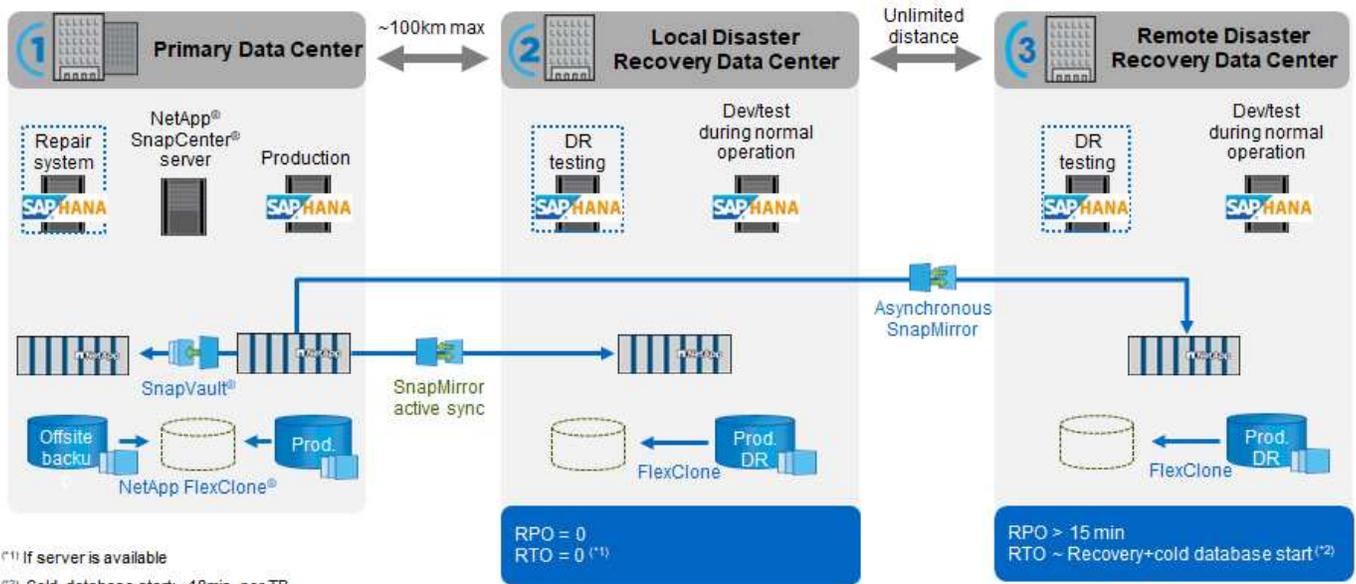
The following figure shows a three-site disaster recovery solution using synchronous SnapMirror active sync to the local DR datacenter and asynchronous SnapMirror to replicate the data to the remote DR datacenter. SnapMirror active sync enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy (RPO=0 and RTO=0). There is no manual intervention or custom scripting required to trigger a failover with SnapMirror active sync. Beginning with ONTAP 9.15.1, SnapMirror active sync supports a symmetric active/active capability. Symmetric active/active enable read and write I/O operations from both copies of a protected LUN with bidirectional synchronous replication so that both LUN copies can serve I/O operations locally.

More details can be found at [SnapMirror active sync overview in ONTAP](#).

The RTO for the asynchronous SnapMirror replication primarily depends on the time needed to start the HANA database at the DR site and load the data into memory. With the assumption that the data is read with a throughput of 1000MBps, loading 1TB of data would take approximately 18 minutes.

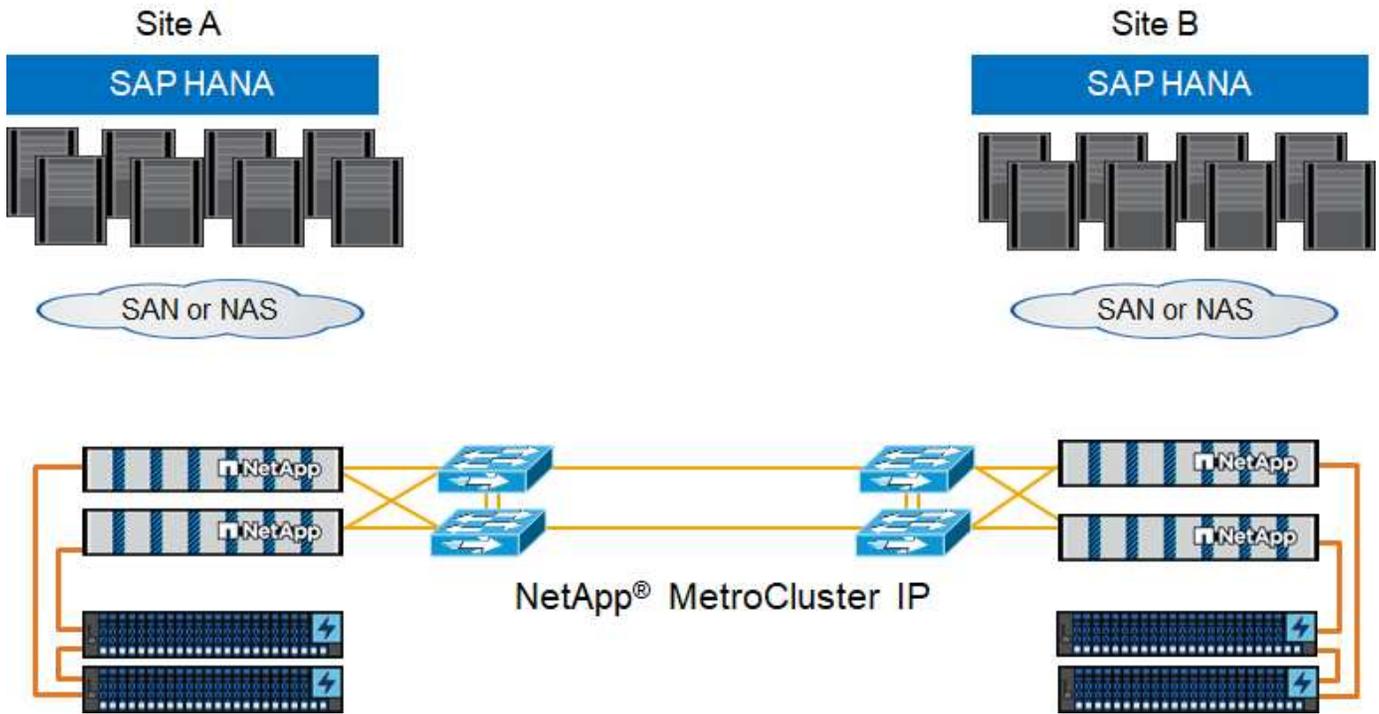
The servers at the DR sites can be used as dev/test systems during normal operation. In the case of a disaster, the dev/test systems would need to be shut down and started as DR production servers.

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



### Storage replication based on NetApp MetroCluster

The following figure shows a high-level overview of the solution. The storage cluster at each site provides local high availability and is used for the production workload. The data of each site is synchronously replicated to the other location and is available in case of disaster failover.



## Storage sizing

The following section provides an overview of performance and capacity considerations required for sizing a storage system for SAP HANA.



Contact your NetApp or NetApp partner sales representative to support the storage sizing process and to assist you with creating a properly sized storage environment.

## Performance considerations

SAP has defined a static set of storage key performance indicators (KPIs). These KPIs are valid for all production SAP HANA environments independent of the memory size of the database hosts and the applications that use the SAP HANA database. These KPIs are valid for single-host, multiple-host, Business Suite on HANA, Business Warehouse on HANA, S/4HANA, and BW/4HANA environments. Therefore, the current performance sizing approach depends on only the number of active SAP HANA hosts that are attached to the storage system.



Storage performance KPIs are only mandated for production SAP HANA systems, but you can implement them in for all HANA system.

SAP delivers a performance test tool which must be used to validate the storage systems performance for active SAP HANA hosts attached to the storage.

NetApp tested and predefined the maximum number of SAP HANA hosts that can be attached to a specific storage model, while still fulfilling the required storage KPIs from SAP for production-based SAP HANA systems.

The maximum number of SAP HANA hosts that can be run on a disk shelf and the minimum number of SSDs required per SAP HANA host were determined by running the SAP performance test tool. This test does not consider the actual storage capacity requirements of the hosts. You must also calculate the capacity requirements to determine the actual storage configuration needed.

## SAS disk shelf

With the 12Gb SAS disk shelf (DS224C), the performance sizing is performed by using fixed disk- shelf configurations:

- Half-loaded disk shelves with 12 SSDs
- Fully loaded disk shelves with 24 SSDs

Both configurations use advanced drive partitioning (ADPv2). A half-loaded disk shelf supports up to 9 SAP HANA hosts; a fully loaded shelf supports up to 14 hosts in a single disk shelf. The SAP HANA hosts must be equally distributed between both storage controllers.



The DS224C disk shelf must be connected by using 12Gb SAS to support the number of SAP HANA hosts.

The 6Gb SAS disk shelf (DS2246) supports a maximum of 4 SAP HANA hosts. The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers. The following figure summarizes the supported number of SAP HANA hosts per disk shelf.

	<b>6Gb SAS shelves (DS2246) Fully loaded with 24 SSDs</b>	<b>12Gb SAS shelves (DS224C) Half-loaded with 12 SSDs and ADPv2</b>	<b>12Gb SAS shelves (DS224C) Fully loaded with 24 SSDs and ADPv2</b>
Maximum number of SAP HANA hosts per disk shelf	4	9	14



This calculation is independent of the storage controller used. Adding more disk shelves does not increase the maximum number of SAP HANA hosts that a storage controller can support.

## NS224 NVMe shelf

One NVMe SSDs (data) supports up to 2/5 SAP HANA hosts depending on the specific NVMe disk being used.

The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers. The same applies to the internal NVMe disks of AFF and ASA systems.



Adding more disk shelves does not increase the maximum number of SAP HANA hosts that a storage controller can support.

## Mixed workloads

SAP HANA and other application workloads running on the same storage controller or in the same storage aggregate are supported. However, it is a NetApp best practice to separate SAP HANA workloads from all other application workloads.

You might decide to deploy SAP HANA workloads and other application workloads on either the same storage controller or the same aggregate. If so, you must make sure that adequate performance is available for SAP HANA within the mixed workload environment. NetApp also recommends that you use quality of service (QoS) parameters to regulate the effect these other applications could have on SAP HANA applications and to guarantee throughput for SAP HANA applications.

The SAP HCMT test tool must be used to check if additional SAP HANA hosts can be run on an existing

storage controller that is already in use for other workloads. SAP application servers can be safely placed on the same storage controller and/or aggregate as the SAP HANA databases.

## Capacity considerations

A detailed description of the capacity requirements for SAP HANA is in the [SAP Note 1900823](#) white paper.



The capacity sizing of the overall SAP landscape with multiple SAP HANA systems must be determined by using SAP HANA storage sizing tools from NetApp. Contact NetApp or your NetApp partner sales representative to validate the storage sizing process for a properly sized storage environment.

## Configuration of performance test tool

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used. These parameters must also be set for the performance test tool from SAP when the storage performance is being tested with the SAP test tool.

NetApp conducted performance tests to define the optimal values. The following table lists the parameters that must be set within the configuration file of the SAP test tool.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For more information about the configuration of SAP test tool, see [SAP note 1943937](#) for HWCCT (SAP HANA 1.0) and [SAP note 2493172](#) for HCMT/HCOT (SAP HANA 2.0).

The following example shows how variables can be set for the HCMT/HCOT execution plan.

```
...
{
    "Comment": "Log Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "LogAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Data Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "DataAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
```

```

    "Comment": "Log Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "LogAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "DataAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
  },
  {
    "Comment": "Log Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
    "Name": "LogAsyncWriteSubmitBlocks",
    "Value": "all",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
    "Name": "DataAsyncWriteSubmitBlocks",
    "Value": "all",
    "Request": "false"
  },
  {
    "Comment": "Log Volume: Maximum number of parallel I/O requests
per completion queue",
    "Name": "LogExtMaxParallelIoRequests",
    "Value": "128",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Maximum number of parallel I/O requests
per completion queue",
    "Name": "DataExtMaxParallelIoRequests",
    "Value": "128",
    "Request": "false"
  }, ...

```

These variables must be used for the test configuration. This is usually the case with the predefined execution plans SAP delivers with the HCMT/HCOT tool. The following example for a 4k log write test is from an

execution plan.

```
...
  {
    "ID": "D664D001-933D-41DE-A904F304AEB67906",
    "Note": "File System Write Test",
    "ExecutionVariants": [
      {
        "ScaleOut": {
          "Port": "${RemotePort}",
          "Hosts": "${Hosts}",
          "ConcurrentExecution": "${FSConcurrentExecution}"
        },
        "RepeatCount": "${TestRepeatCount}",
        "Description": "4K Block, Log Volume 5GB, Overwrite",
        "Hint": "Log",
        "InputVector": {
          "BlockSize": 4096,
          "DirectoryName": "${LogVolume}",
          "FileOverwrite": true,
          "FileSize": 5368709120,
          "RandomAccess": false,
          "RandomData": true,
          "AsyncReadSubmit": "${LogAsyncReadSubmit}",
          "AsyncWriteSubmitActive":
"${LogAsyncWriteSubmitActive}",
          "AsyncWriteSubmitBlocks":
"${LogAsyncWriteSubmitBlocks}",
          "ExtMaxParallelIoRequests":
"${LogExtMaxParallelIoRequests}",
          "ExtMaxSubmitBatchSize": "${LogExtMaxSubmitBatchSize}",
          "ExtMinSubmitBatchSize": "${LogExtMinSubmitBatchSize}",
          "ExtNumCompletionQueues":
"${LogExtNumCompletionQueues}",
          "ExtNumSubmitQueues": "${LogExtNumSubmitQueues}",
          "ExtSizeKernelIoQueue": "${ExtSizeKernelIoQueue}"
        }
      },
    ]
  }
...

```

### Storage sizing process overview

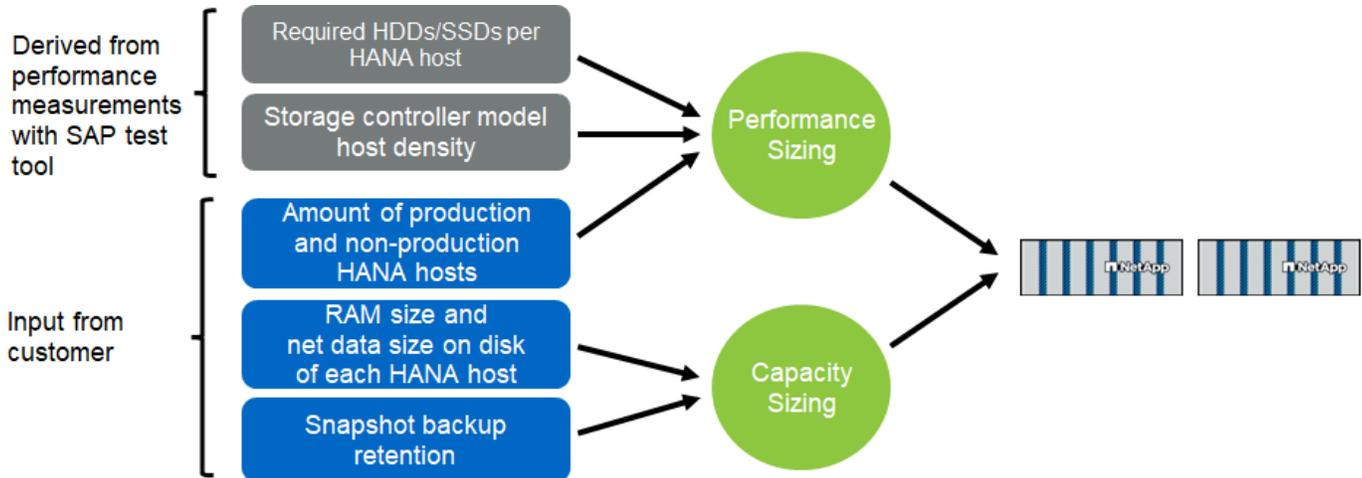
The number of disks per HANA host and the SAP HANA host density for each storage model were determined using the SAP HANA test tool.

The sizing process requires details such as the number of production and nonproduction SAP HANA hosts, the

RAM size of each host, and the backup retention of the storage-based Snapshot copies. The number of SAP HANA hosts determines the storage controller and the number of disks required.

The size of the RAM, net data size on the disk of each SAP HANA host, and the Snapshot copy backup retention period are used as inputs during capacity sizing.

The following figure summarizes the sizing process.



## Infrastructure setup and configuration

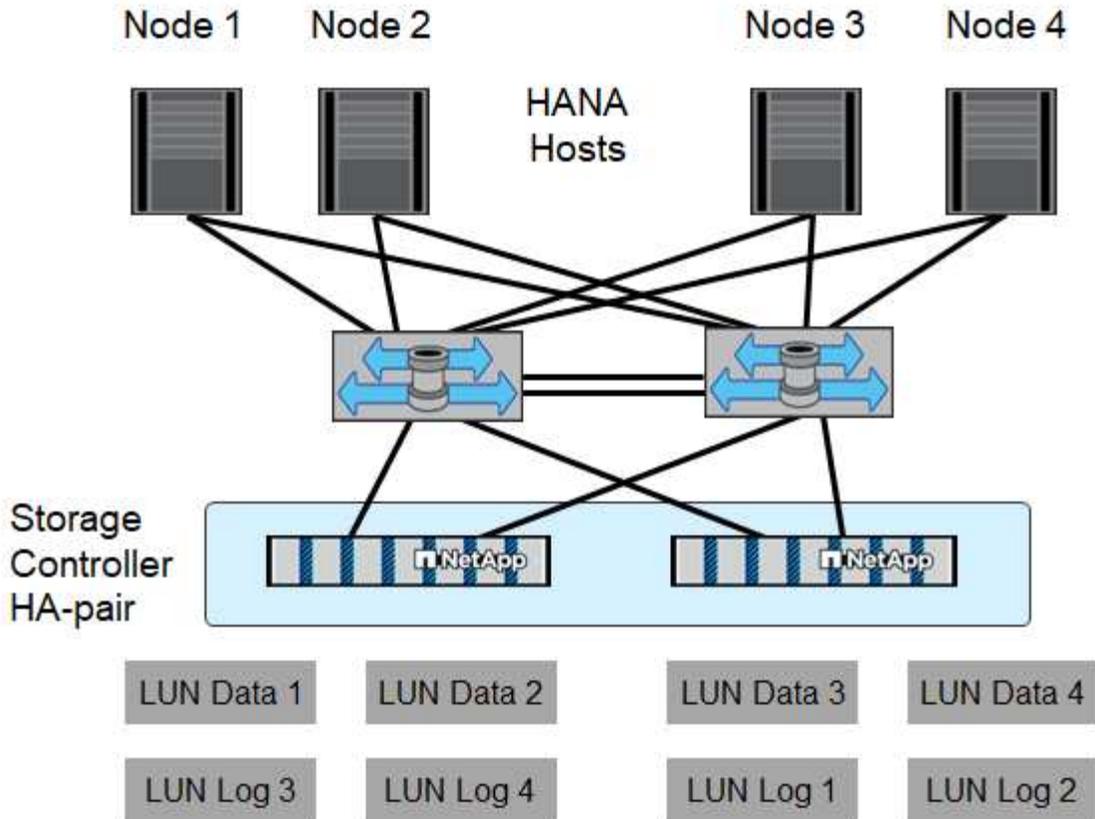
The following sections provide SAP HANA infrastructure setup and configuration guidelines and describes all the steps needed to set up an SAP HANA system. Within these sections, the following example configurations are used:

- HANA system with SID=FC5
  - SAP HANA single and multiple host using Linux logical volume manager (LVM)
  - SAP HANA single host using SAP HANA multiple partitions

## SAN fabric setup

Each SAP HANA server must have a redundant FCP SAN connection with a minimum of 8Gbps bandwidth. For each SAP HANA host attached to a storage controller, at least 8Gbps bandwidth must be configured at the storage controller.

The following figure shows an example with four SAP HANA hosts attached to two storage controllers. Each SAP HANA host has two FCP ports connected to the redundant fabric. At the storage layer, four FCP ports are configured to provide the required throughput for each SAP HANA host.



In addition to the zoning on the switch layer, you must map each LUN on the storage system to the hosts that connect to this LUN. Keep the zoning on the switch simple; that is, define one zone set in which all host HBAs can see all controller HBAs.

### Time synchronization

You must synchronize the time between the storage controllers and the SAP HANA database hosts. To do so, set the same time server for all storage controllers and all SAP HANA hosts.

### Storage controller setup

This section describes the configuration of the NetApp storage system. You must complete the primary installation and setup according to the corresponding Data ONTAP setup and configuration guides.

### Storage efficiency

Inline deduplication, cross-volume inline deduplication, inline compression, and inline compaction are supported with SAP HANA in an SSD configuration.

### NetApp FlexGroup Volumes

The usage of NetApp FlexGroup Volumes is not supported for SAP HANA. Due to the architecture of SAP HANA the usage of FlexGroup Volumes does not provide any benefit and may result in performance issues.

## NetApp Volume and Aggregate Encryption

The use of NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) are supported with SAP HANA.

### Quality of Service

QoS can be used to limit the storage throughput for specific SAP HANA systems or non-SAP applications on a shared controller.

### Production and Dev/Test

One use case would be to limit the throughput of development and test systems so that they cannot influence production systems in a mixed setup.

During the sizing process, you should determine the performance requirements of a nonproduction system. Development and test systems can be sized with lower performance values, typically in the range of 20% to 50% of a production-system KPI as defined by SAP.

Large write I/O has the biggest performance effect on the storage system. Therefore, the QoS throughput limit should be set to a percentage of the corresponding write SAP HANA storage performance KPI values in the data and log volumes.

### Shared Environments

Another use case is to limit the throughput of heavy write workloads, especially to avoid that these workloads have an impact on other latency sensitive write workloads.

In such environments it is best practice to apply a non-shared throughput ceiling QoS group-policy to each LUN within each Storage Virtual Machine (SVM) to restrict the max throughput of each individual storage object to the given value. This reduces the possibility that a single workload can negatively influence other workloads.

To do so, a group-policy needs to be created using the CLI of the ONTAP cluster for each SVM:

```
qos policy-group create -policy-group <policy-name> -vserver <vserver
name> -max-throughput 1000MB/s -is-shared false
```

and applied to each LUN within the SVM. Below is an example to apply the policy group to all existing LUNs within an SVM:

```
lun modify -vserver <vserver name> -path * -qos-policy-group <policy-
name>
```

This needs to be done for every SVM. The name of the QoS police group for each SVM needs to be different. For new LUNs, the policy can be applied directly:

```
lun create -vserver <vserver_name> -path /vol/<volume_name>/<lun_name>
-size <size> -ostype <e.g. linux> -qos-policy-group <policy-name>
```

It is recommended to use 1000MB/s as maximum throughput for a given LUN. If an application requires more throughput, multiple LUNs with LUN striping shall be used to provide the needed bandwidth. This guide

provides an example for SAP HANA based on Linux LVM in section [Host Setup](#).



The limit applies also to reads. Therefore use enough LUNs to fulfil the required SLAs for SAP HANA database startup time and for backups.

### NetApp FabricPool

NetApp FabricPool technology must not be used for active primary file systems in SAP HANA systems. This includes the file systems for the data and log area as well as the `/hana/shared` file system. Doing so results in unpredictable performance, especially during the startup of an SAP HANA system.

You can use the Snapshot-Only tiering policy along with FabricPool at a backup target such as SnapVault or SnapMirror destination.



Using FabricPool for tiering Snapshot copies at primary storage or using FabricPool at a backup target changes the required time for the restore and recovery of a database or other tasks such as creating system clones or repair systems. Take this into consideration for planning your overall lifecycle-management strategy, and check to make sure that your SLAs are still being met while using this function.

FabricPool is a good option for moving log backups to another storage tier. Moving backups affects the time needed to recover an SAP HANA database. Therefore, the option `tiering-minimum-cooling-days` should be set to a value that places log backups, which are routinely needed for recovery, on the local fast storage tier.

### Configure storage

The following overview summarizes the required storage configuration steps. Each step is covered in more detail in the subsequent sections. In this section, we assume that the storage hardware is set up and that the ONTAP software is already installed. Also, the connection of the storage FCP ports to the SAN fabric must already be in place.

1. Check the correct disk shelf configuration, as described in [Disk shelf connection](#).
2. Create and configure the required aggregates, as described in [Aggregate configuration](#).
3. Create a storage virtual machine (SVM), as described in [Storage virtual machine configuration](#).
4. Create logical interfaces (LIFs), as described in [Logical interface configuration](#).
5. Create initiator groups (igroups) with worldwide names (WWNs) of HANA servers as described in the section [Initiator groups](#).
6. Create and configure volumes and LUNs within the aggregates as described in the section [Single Host Setup](#) for single hosts or in section [Multiple Host Setup](#)

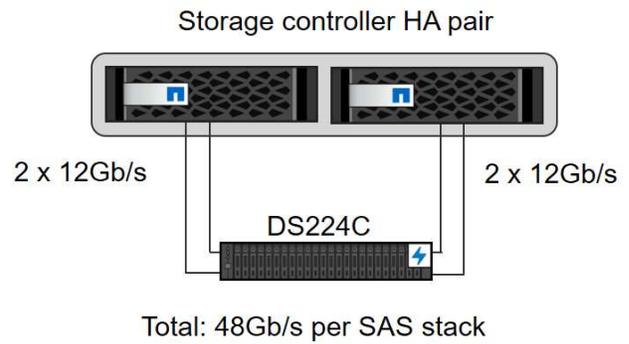
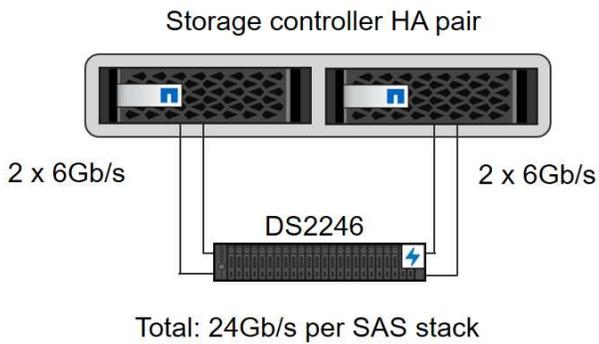
### Disk shelf connections

#### SAS-based disk shelves

A maximum of one disk shelf can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally between both controllers of the HA pair. ADPv2 is used with ONTAP 9 and the DS224C disk shelves.

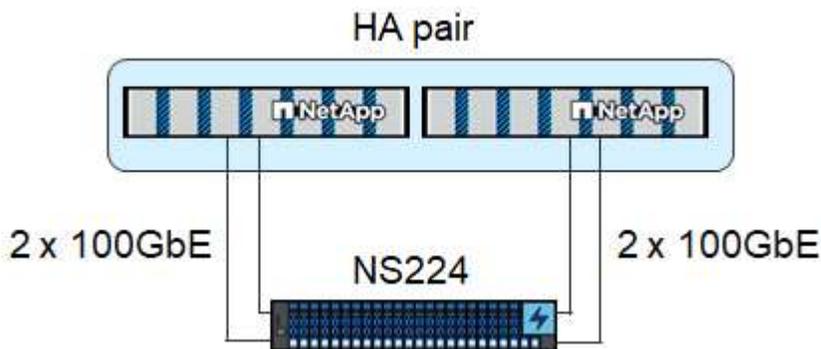


With the DS224C disk shelf, quad-path SAS cabling can also be used but is not required.



### NVMe-based disk shelves

Each NS224 NVMe disk shelf is connected with two 100GbE ports per controller, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair. ADPv2 is also used for the NS224 disk shelf.



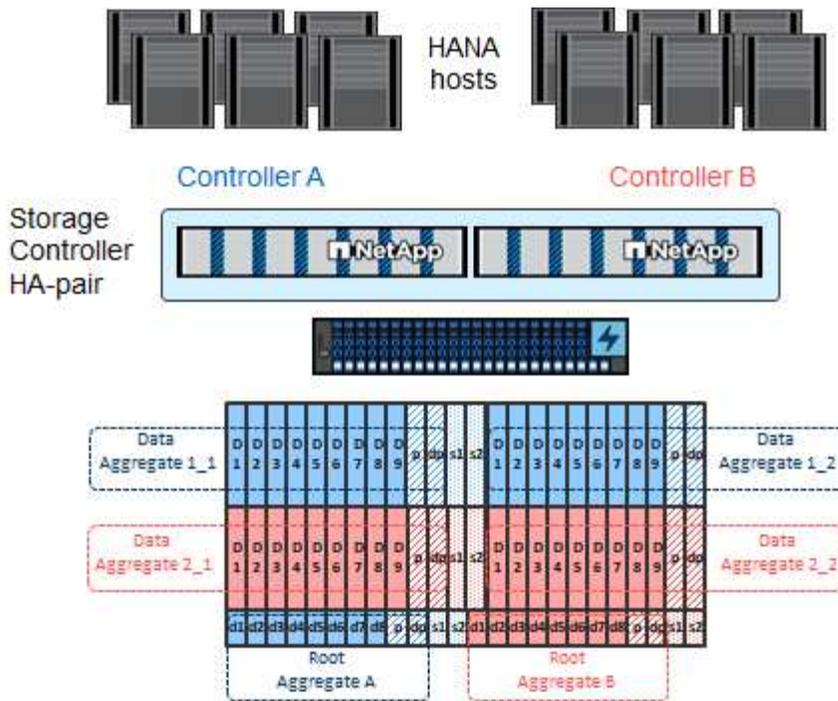
### Aggregate configuration

In general, you must configure two aggregates per controller, independent of which disk shelf or disk technology (SSD or HDD) is used. This step is necessary so that you can use all available controller resources.



ASA systems launched after August 2024 do not require this step as it is automatically done

The following figure shows a configuration of 12 SAP HANA hosts running on a 12Gb SAS shelf configured with ADPv2. Six SAP HANA hosts are attached to each storage controller. Four separate aggregates, two at each storage controller, are configured. Each aggregate is configured with 11 disks with nine data and two parity disk partitions. For each controller, two spare partitions are available.



### Storage virtual machine configuration

Multiple SAP landscapes with SAP HANA databases can use a single SVM. An SVM can also be assigned to each SAP landscape, if necessary, in case they are managed by different teams within a company.

If there is a QoS profile automatically created and assigned while creating a new SVM, remove this automatically created profile from the SVM to ensure the required performance for SAP HANA:

```
vserver modify -vserver <svm-name> -qos-policy-group none
```

### Logical interface configuration

Within the storage cluster configuration, one network interface (LIF) must be created and assigned to a dedicated FCP port. If, for example, four FCP ports are required for performance reasons, four LIFs must be created. The following figure shows a screenshot of the eight LIFs that were configured on the SVM.

NetApp ONTAP System Manager | a400-sapcc

Search actions, objects, and pages

- Dashboard
- Insights
- Storage
- Network**
  - Overview
  - Ethernet ports
  - FC ports
  - Events & jobs
  - Protection
  - Hosts
  - Cluster

### IPspaces

+ Add

Cluster	Broadcast domains Cluster
Default	Storage VMs BlueXPDR_SVM1 ,C30-HANA ,TCP-NVME ,abhi-a400 , hana-A400 ,infra-svm ,svm-dietmare-misc ,test_rdma Broadcast domains Default ,NFS ,NFS2 ,rdma ,vlan-data ,vlan-log

### Broadcast domains

Learn more

+ Add

Cluster	9000 MTU	IPspace: Cluster a400-sapcc-01 e3a e3b a400-sapcc-02 e3a e3b
Default	1500 MTU	IPspace: Default a400-sapcc-01 e0M a400-sapcc-02 e0M
NFS	9000 MTU	IPspace: Default a400-sapcc-01 a0a a400-sapcc-02 a0a
NFS2	9000 MTU	IPspace: Default

### Network interfaces

Subnets

+ Add

Name	Status	Storage VM	IPspace	Address	Current node	Current port	Portset	Protocols	Throughput (M)
lif_hana_345	✔	hana-A400		20:0b:d0:39:ea:2ef9:41	a400-sapcc-01	1a	FC	0	
lif_hana_965	✔	hana-A400		20:0c:d0:39:ea:2ef9:41	a400-sapcc-01	1b	FC	0	
lif_hana_205	✔	hana-A400		20:0d:d0:39:ea:2ef9:41	a400-sapcc-01	1c	FC	0	
lif_hana_314	✔	hana-A400		20:0e:d0:39:ea:2ef9:41	a400-sapcc-01	1d	FC	0	
lif_hana_908	✔	hana-A400		20:0f:d0:39:ea:2ef9:41	a400-sapcc-02	1a	FC	0	
lif_hana_726	✔	hana-A400		20:10:d0:39:ea:2ef9:41	a400-sapcc-02	1b	FC	0	
lif_hana_521	✔	hana-A400		20:11:d0:39:ea:2ef9:41	a400-sapcc-02	1c	FC	0	
lif_hana_946	✔	hana-A400		20:12:d0:39:ea:2ef9:41	a400-sapcc-02	1d	FC	0	

During the SVM creation with ONTAP System Manager, you can select all of the required physical FCP ports, and one LIF per physical port is created automatically.

NetApp ONTAP System Manager | a400-sapcc

Search actions, objects, and pages

Dashboard

Insights

Storage

- Overview
- Volumes
- LUNs
- NVMe namespaces
- Consistency groups
- Shares
- Qtrees
- Quotas
- Storage VMs
- Tiers

Network

Events & jobs

Protection

Hosts

Cluster

### Add storage VM

Storage VM name: hana

Access protocol: SMB/CIFS, NFS, iSCSI, **FC**, NVMe

Enable FC

Configure FC ports

Nodes	1a	1b	1c	1d
a400-sapcc-01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
a400-sapcc-02	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Storage VM administration

Enable maximum capacity limit  
The maximum capacity that all volumes in this storage VM can allocate. [Learn More](#)

Manage administrator account

User name: vsadmin

Password: .....

Confirm password: .....

Add a network interface for storage VM management.

Node: a400-sapcc-01

IP address: 10.10.10.10

Subnet mask: 255.255.255.0

Save Cancel

## Initiator groups

An igroup can be configured for each server or for a group of servers that require access to a LUN. The igroup configuration requires the worldwide port names (WWPNs) of the servers.

Using the `sanlun` tool, run the following command to obtain the WWPNs of each SAP HANA host:

```
stlrx300s8-6:~ # sanlun fcp show adapter
/sbin/udevadm
/sbin/udevadm

host0 ..... WWPN:2100000e1e163700
host1 ..... WWPN:2100000e1e163701
```



The `sanlun` tool is part of the NetApp Host Utilities and must be installed on each SAP HANA host. More details can be found in section [Host setup](#).

The initiator groups can be created using the CLI of the ONTAP Cluster.

```
lun igroup create -igroup <igroup name> -protocol fcp -ostype linux
-initiator <list of initiators> -vserver <SVM name>
```

## Single host

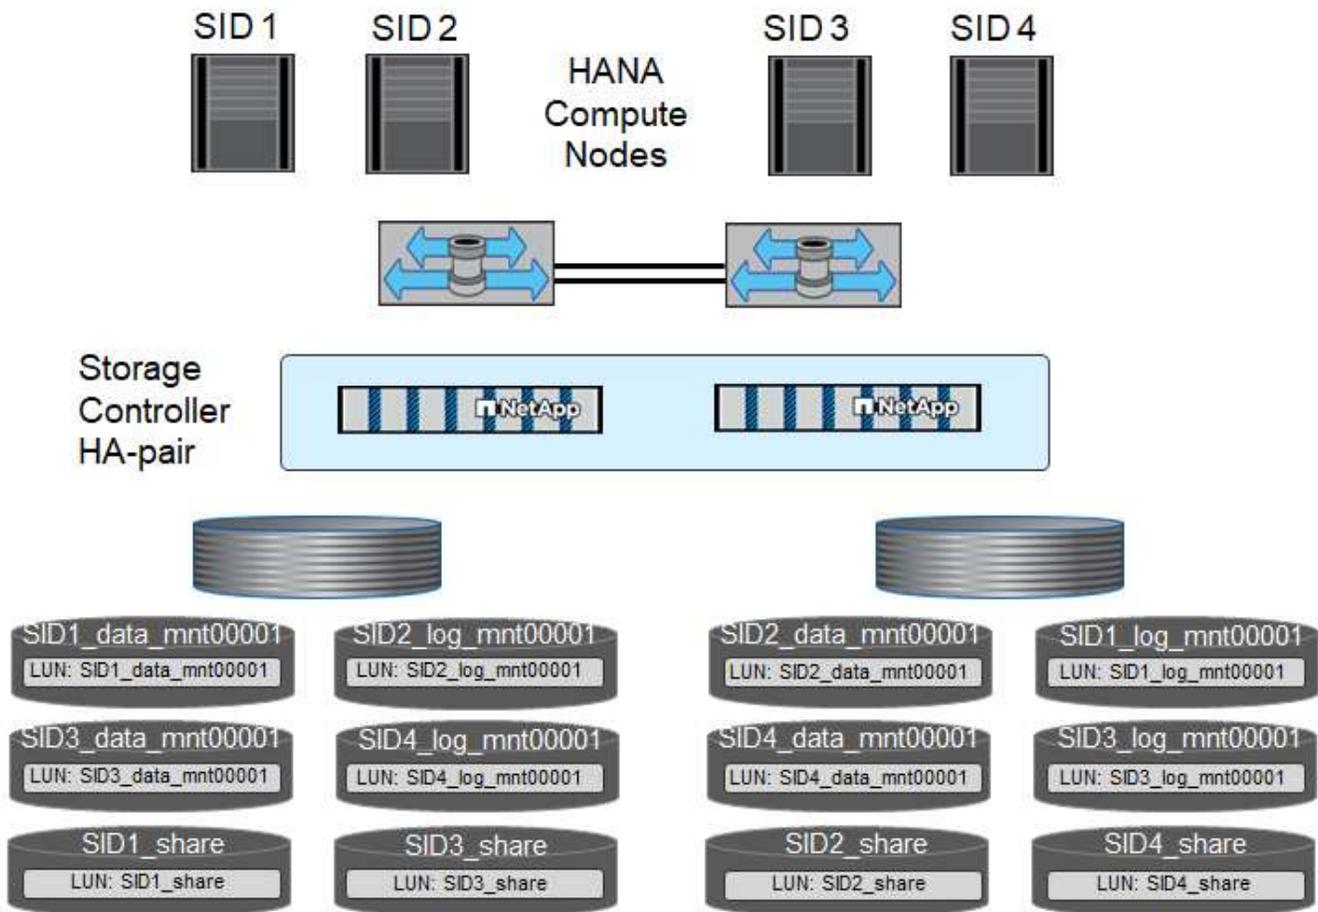
This section describes the configuration of the NetApp storage system specific to SAP HANA single-host systems

### Volume and LUN configuration for SAP HANA single-host systems

The following figure shows the volume configuration of four single-host SAP HANA systems. The data and log volumes of each SAP HANA system are distributed to different storage controllers. For example, volume `SID1_data_mnt00001` is configured on controller A, and volume `SID1_log_mnt00001` is configured on controller B. Within each volume, a single LUN is configured.



If only one storage controller of a HA pair is used for the SAP HANA systems, data volumes and log volumes can also be stored on the same storage controller.



For each SAP HANA host, a data volume, a log volume, and a volume for `/hana/shared` are configured. The following table shows an example configuration with four SAP HANA single-host systems.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data, log, and shared volumes for system SID1	Data volume: SID1_data_mnt00001	Shared volume: SID1_shared	–	Log volume: SID1_log_mnt00001
Data, log, and shared volumes for system SID2	–	Log volume: SID2_log_mnt00001	Data volume: SID2_data_mnt00001	Shared volume: SID2_shared
Data, log, and shared volumes for system SID3	Shared volume: SID3_shared	Data volume: SID3_data_mnt00001	Log volume: SID3_log_mnt00001	–
Data, log, and shared volumes for system SID4	Log volume: SID4_log_mnt00001	–	Shared volume: SID4_shared	Data volume: SID4_data_mnt00001

The following table shows an example of the mount point configuration for a single-host system.

LUN	Mount point at SAP HANA host	Note
SID1_data_mnt00001	<code>/hana/data/SID1/mnt00001</code>	Mounted using <code>/etc/fstab</code> entry

LUN	Mount point at SAP HANA host	Note
SID1_log_mnt00001	/hana/log/SID1/mnt00001	Mounted using /etc/fstab entry
SID1_shared	/hana/shared/SID1	Mounted using /etc/fstab entry



With the described configuration, the `/usr/sap/SID1` directory in which the default home directory of user `SID1adm` is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `SID1_shared` volume for the `/usr/sap/SID1` directory so that all file systems are on the central storage.

## Volume and LUN configuration for SAP HANA single-host systems using Linux LVM

The Linux LVM can be used to increase performance and to address LUN size limitations. The different LUNs of an LVM volume group should be stored within a different aggregate and at a different controller. The following table shows an example for two LUNs per volume group.



It is not necessary to use LVM with multiple LUNs to fulfill the SAP HANA KPIs, but it is recommended.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data, log, and shared volumes for LVM based system	Data volume: SID1_data_mnt00001	Shared volume: SID1_shared Log2 volume: SID1_log2_mnt00001	Data2 volume: SID1_data2_mnt00001	Log volume: SID1_log_mnt00001

## Volume options

The volume options listed in the following table must be verified and set on all volumes used for SAP HANA.

Action	ONTAP 9
Disable automatic Snapshot copies	<code>vol modify -vserver &lt;vserver-name&gt; -volume &lt;volname&gt; -snapshot-policy none</code>
Disable visibility of Snapshot directory	<code>vol modify -vserver &lt;vserver-name&gt; -volume &lt;volname&gt; -snapdir-access false</code>

## Creating LUNs and mapping LUNs to initiator groups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a SAP HANA single host system with SID FC5 using LVM and two LUNs per LVM volume group:

1. Create all necessary volumes.

```
vol create -volume FC5_data_mnt00001 -aggregate aggr1_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00001 -aggregate aggr1_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00001 -aggregate aggr1_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00001 -aggregate aggr1_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_shared -aggregate aggr1_1 -size 512g -state
online -policy default -snapshot-policy none -junction-path /FC5_shared
-encrypt false -space-guarantee none
```

## 2. Create all LUNs.

```
lun create -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
```

## 3. Create the initiator group for all ports belonging to sythe hosts of FC5.

```
lun igroup create -igroup HANA-FC5 -protocol fcp -ostype linux
-initiator 10000090fadcc5fa,10000090fadcc5fb -vserver hana
```

## 4. Map all LUNs to created initiator group.

```
lun map -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -igroup HANA-FC5
```

## Multiple hosts

This section describes the configuration of the NetApp storage system specific to SAP HANA multiple-hosts systems

### Volume and LUN configuration for SAP HANA multiple-host systems

The following figure shows the volume configuration of a 4+1 multiple-host SAP HANA system. The data volumes and log volumes of each SAP HANA host are distributed to different storage controllers. For example, the volume `SID_data_mnt00001` is configured on controller A and the volume `SID_log_mnt00001` is configured on controller B. One LUN is configured within each volume.

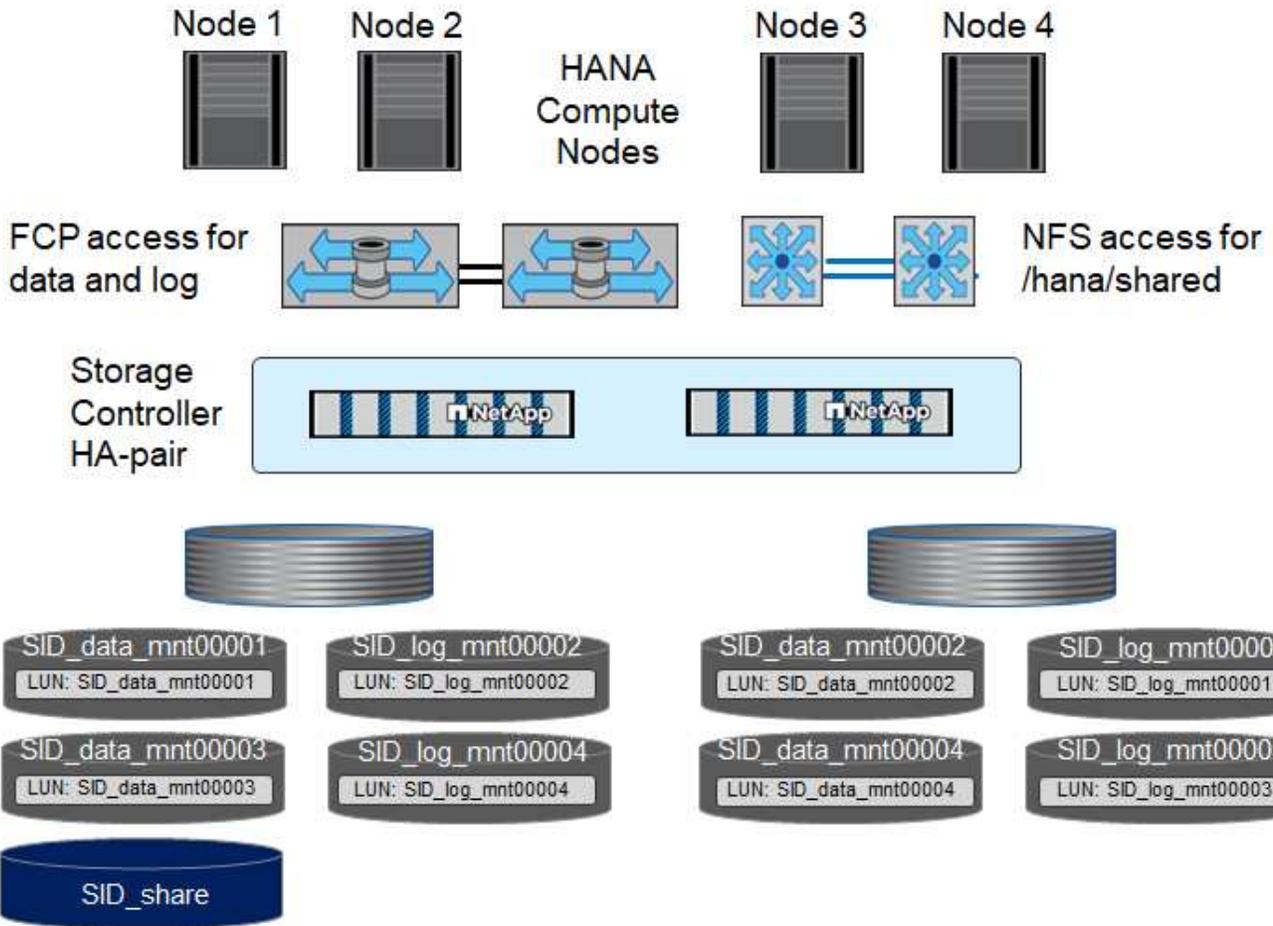
The `/hana/shared` volume must be accessible by all HANA hosts and is therefore exported by using NFS. Even though there are no specific performance KPIs for the `/hana/shared` file system, NetApp recommends using a 10Gb Ethernet connection.



If only one storage controller of an HA pair is used for the SAP HANA system, data and log volumes can also be stored on the same storage controller.



NetApp ASA systems do not support NFS as a protocol. NetApp recommends using an additional AFF or FAS system for the `/hana/shared` file system.



For each SAP HANA host, a data volume and a log volume are created. The `/hana/shared` volume is used by all hosts of the SAP HANA system. The following table shows an example configuration for a 4+1 multiple-host SAP HANA system.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	–	Log volume: SID_log_mnt00001	–
Data and log volumes for node 2	Log volume: SID_log_mnt00002	–	Data volume: SID_data_mnt00002	–
Data and log volumes for node 3	–	Data volume: SID_data_mnt00003	–	Log volume: SID_log_mnt00003
Data and log volumes for node 4	–	Log volume: SID_log_mnt00004	–	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared	–	–	–

The following table shows the configuration and the mount points of a multiple-host system with four active SAP HANA hosts.

LUN or volume	Mount point at SAP HANA host	Note
LUN: SID_data_mnt00001	/hana/data/SID/mnt00001	Mounted using storage connector
LUN: SID_log_mnt00001	/hana/log/SID/mnt00001	Mounted using storage connector
LUN: SID_data_mnt00002	/hana/data/SID/mnt00002	Mounted using storage connector
LUN: SID_log_mnt00002	/hana/log/SID/mnt00002	Mounted using storage connector
LUN: SID_data_mnt00003	/hana/data/SID/mnt00003	Mounted using storage connector
LUN: SID_log_mnt00003	/hana/log/SID/mnt00003	Mounted using storage connector
LUN: SID_data_mnt00004	/hana/data/SID/mnt00004	Mounted using storage connector
LUN: SID_log_mnt00004	/hana/log/SID/mnt00004	Mounted using storage connector
Volume: SID_shared	/hana/shared	Mounted at all hosts using NFS and /etc/fstab entry



With the described configuration, the `/usr/sap/SID` directory in which the default home directory of user SIDadm is stored, is on the local disk for each HANA host. In a disaster recovery setup with disk-based replication, NetApp recommends creating four additional subdirectories in the `SID_shared` volume for the `/usr/sap/SID` file system so that each database host has all its file systems on the central storage.

### Volume and LUN configuration for SAP HANA multiple-host systems using Linux LVM

The Linux LVM can be used to increase performance and to address LUN size limitations. The different LUNs of an LVM volume group should be stored within a different aggregate and at a different controller.



It is not necessary to use LVM to combine several LUN to fulfill the SAP HANA KPIs, but it is recommended

The following table shows an example for two LUNs per volume group for a 2+1 SAP HANA multiple host system.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	Log2 volume: SID_log2_mnt00001	Log volume: SID_log_mnt00001	Data2 volume: SID_data2_mnt00001
Data and log volumes for node 2	Log2 volume: SID_log2_mnt00002	Data volume: SID_data_mnt00002	Data2 volume: SID_data2_mnt00002	Log volume: SID_log_mnt00002
Shared volume for all hosts	Shared volume: SID_shared	–	–	–

### Volume options

The volume options listed in the following table must be verified and set on all SVMs.

Action	
Disable automatic Snapshot copies	vol modify -vserver <vserver-name> -volume <volname> -snapshot-policy none
Disable visibility of Snapshot directory	vol modify -vserver <vserver-name> -volume <volname> -snapdir-access false

## Creating LUNs, volumes, and mapping LUNs to initiator groups

You can use NetApp ONTAP System Manager to create storage volumes and LUNs and the map them to the igroups of the servers and the ONTAP CLI. This guide describes the usage of the CLI.

### Creating LUNs, volumes, and mapping LUNs to initiator groups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a 2+1 SAP HANA multiple host system with SID FC5 using LVM and two LUNs per LVM volume group:

1. Create all necessary volumes.

```
vol create -volume FC5_data_mnt00001 -aggregate aggr1_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00002 -aggregate aggr2_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00001 -aggregate aggr1_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data_mnt00002 -aggregate aggr2_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00001 -aggregate aggr1_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00002 -aggregate aggr2_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00001 -aggregate aggr1_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00002 -aggregate aggr2_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_shared -aggregate aggr1_1 -size 512g -state
online -policy default -snapshot-policy none -junction-path /FC5_shared
-encrypt false -space-guarantee none
```

## 2. Create all LUNs.

```
lun create -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data_mnt00002/FC5_data_mnt00002 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00002/FC5_data2_mnt00002 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log_mnt00002/FC5_log_mnt00002 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00002/FC5_log2_mnt00002 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
```

## 3. Create the initiator group for all servers belonging to system FC5.

```
lun igroup create -igroup HANA-FC5 -protocol fcp -ostype linux
-initiator
10000090fadcc5fa,10000090fadcc5fb,10000090fadcc5c1,10000090fadcc5c2,1000
0090fadcc5c3,10000090fadcc5c4 -vserver hana
```

## 4. Map all LUNs to created initiator group.

```
lun map -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_data_mnt00002/FC5_data_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00002/FC5_data2_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt00002/FC5_log_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00002/FC5_log2_mnt00002 -igroup HANA-FC5
```

## SAP HANA storage connector API

A storage connector is required only in multiple-host environments that have failover capabilities. In multiple-host setups, SAP HANA provides high-availability functionality so that an SAP HANA database host can fail over to a standby host.

In this case, the LUNs of the failed host are accessed and used by the standby host. The storage connector is used to make sure that a storage partition can be actively accessed by only one database host at a time.

In SAP HANA multiple-host configurations with NetApp storage, the standard storage connector delivered by SAP is used. The “SAP HANA Fibre Channel Storage Connector Admin Guide” can be found as an attachment to [SAP note 1900823](#).

## Host setup

Before setting up the host, NetApp SAN host utilities must be downloaded from the [NetApp Support](#) site and installed on the HANA servers. The host utility documentation includes information about additional software that must be installed depending on the FCP HBA used.

The documentation also contains information on multipath configurations that are specific to the Linux version used. This document covers the required configuration steps for SLES 12 SP1 or higher and RHEL 7. 2 or later, as described in the [Linux Host Utilities 7.1 Installation and Setup Guide](#).

## Configure multipathing



Steps 1 through 6 must be executed on all worker and standby hosts in an SAP HANA multiple-host configuration.

To configure multipathing, complete the following steps:

1. Run the Linux `rescan-scsi-bus.sh -a` command on each server to discover new LUNs.
2. Run the `sanlun lun show` command and verify that all required LUNs are visible. The following example shows the `sanlun lun show` command output for a 2+1 multiple-host HANA system with two data LUNs and two log LUNs. The output shows the LUNs and the corresponding device files, such as LUN

FC5\_data\_mnt00001 and the device file /dev/sdag Each LUN has eight FC paths from the host to the storage controllers.

```

sapcc-hana-tst:~ # sanlun lun show
controller(7mode/E-Series)/                               device
host                lun
vserver(cDOT/FlashRay)    lun-pathname                filename
adapter  protocol  size  product
-----
svm1                FC5_log2_mnt00002                /dev/sdbb
host21      FCP      500g    cDOT
svm1                FC5_log_mnt00002                /dev/sdba
host21      FCP      500g    cDOT
svm1                FC5_log2_mnt00001                /dev/sdaz
host21      FCP      500g    cDOT
svm1                FC5_log_mnt00001                /dev/sday
host21      FCP      500g    cDOT
svm1                FC5_data2_mnt00002                /dev/sdax
host21      FCP      1t      cDOT
svm1                FC5_data_mnt00002                /dev/sdaw
host21      FCP      1t      cDOT
svm1                FC5_data2_mnt00001                /dev/sdav
host21      FCP      1t      cDOT
svm1                FC5_data_mnt00001                /dev/sdau
host21      FCP      1t      cDOT
svm1                FC5_log2_mnt00002                /dev/sdat
host21      FCP      500g    cDOT
svm1                FC5_log_mnt00002                /dev/sdas
host21      FCP      500g    cDOT
svm1                FC5_log2_mnt00001                /dev/sdar
host21      FCP      500g    cDOT
svm1                FC5_log_mnt00001                /dev/sdaq
host21      FCP      500g    cDOT
svm1                FC5_data2_mnt00002                /dev/sdap
host21      FCP      1t      cDOT
svm1                FC5_data_mnt00002                /dev/sdao
host21      FCP      1t      cDOT
svm1                FC5_data2_mnt00001                /dev/sdan
host21      FCP      1t      cDOT
svm1                FC5_data_mnt00001                /dev/sdam
host21      FCP      1t      cDOT
svm1                FC5_log2_mnt00002                /dev/sdal
host20      FCP      500g    cDOT
svm1                FC5_log_mnt00002                /dev/sdak
host20      FCP      500g    cDOT

```

```

svm1          FC5_log2_mnt00001          /dev/sdaj
host20       FCP           500g          cDOT
svm1          FC5_log_mnt00001          /dev/sdai
host20       FCP           500g          cDOT
svm1          FC5_data2_mnt00002        /dev/sdah
host20       FCP           1t             cDOT
svm1          FC5_data_mnt00002        /dev/sdag
host20       FCP           1t             cDOT
svm1          FC5_data2_mnt00001        /dev/sdaf
host20       FCP           1t             cDOT
svm1          FC5_data_mnt00001        /dev/sdae
host20       FCP           1t             cDOT
svm1          FC5_log2_mnt00002        /dev/sdad
host20       FCP           500g          cDOT
svm1          FC5_log_mnt00002        /dev/sdac
host20       FCP           500g          cDOT
svm1          FC5_log2_mnt00001        /dev/sdab
host20       FCP           500g          cDOT
svm1          FC5_log_mnt00001        /dev/sdaa
host20       FCP           500g          cDOT
svm1          FC5_data2_mnt00002        /dev/sdz
host20       FCP           1t             cDOT
svm1          FC5_data_mnt00002        /dev/sdy
host20       FCP           1t             cDOT
svm1          FC5_data2_mnt00001        /dev/sdx
host20       FCP           1t             cDOT
svm1          FC5_data_mnt00001        /dev/sdw
host20       FCP           1t             cDOT

```

3. Run the `multipath -r` and `multipath -ll` command to get the worldwide identifiers (WWIDs) for the device file names.



In this example, there are eight LUNs.

```

sapcc-hana-tst:~ # multipath -r
sapcc-hana-tst:~ # multipath -ll
3600a098038314e63492b59326b4b786d dm-7 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:2 sdaf 65:240 active ready running
  |- 20:0:5:2 sdx 65:112 active ready running
  |- 21:0:4:2 sdav 66:240 active ready running
  `-- 21:0:6:2 sdan 66:112 active ready running
3600a098038314e63492b59326b4b786e dm-9 NETAPP,LUN C-Mode

```

```

size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:4 sdah 66:16 active ready running
  |- 20:0:5:4 sdz 65:144 active ready running
  |- 21:0:4:4 sdax 67:16 active ready running
  `- 21:0:6:4 sdap 66:144 active ready running
3600a098038314e63492b59326b4b786f dm-11 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:6 sdaj 66:48 active ready running
  |- 20:0:5:6 sdab 65:176 active ready running
  |- 21:0:4:6 sdaz 67:48 active ready running
  `- 21:0:6:6 sdar 66:176 active ready running
3600a098038314e63492b59326b4b7870 dm-13 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:8 sdal 66:80 active ready running
  |- 20:0:5:8 sdad 65:208 active ready running
  |- 21:0:4:8 sdbb 67:80 active ready running
  `- 21:0:6:8 sdat 66:208 active ready running
3600a098038314e63532459326d495a64 dm-6 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:1 sdae 65:224 active ready running
  |- 20:0:5:1 sdw 65:96 active ready running
  |- 21:0:4:1 sdau 66:224 active ready running
  `- 21:0:6:1 sdam 66:96 active ready running
3600a098038314e63532459326d495a65 dm-8 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:3 sdag 66:0 active ready running
  |- 20:0:5:3 sdy 65:128 active ready running
  |- 21:0:4:3 sdaw 67:0 active ready running
  `- 21:0:6:3 sdao 66:128 active ready running
3600a098038314e63532459326d495a66 dm-10 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:5 sdai 66:32 active ready running
  |- 20:0:5:5 sdaa 65:160 active ready running
  |- 21:0:4:5 sday 67:32 active ready running

```

```
`- 21:0:6:5 sdaq 66:160 active ready running
3600a098038314e63532459326d495a67 dm-12 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:7 sdak 66:64 active ready running
  |- 20:0:5:7 sdac 65:192 active ready running
  |- 21:0:4:7 sdba 67:64 active ready running
  `- 21:0:6:7 sdas 66:192 active ready running
```

4. Edit the `/etc/multipath.conf` file and add the WWIDs and alias names.



The example output shows the content of the `/etc/multipath.conf` file, which includes alias names for the four LUNs of a 2+1 multiple-host system. If there is no `multipath.conf` file available, you can create one by running the following command: `multipath -T > /etc/multipath.conf`.

```

sapcc-hana-tst:/ # cat /etc/multipath.conf
multipaths {
    multipath {
        wwid      3600a098038314e63492b59326b4b786d
        alias     svm1-FC5_data2_mnt00001
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786e
        alias     svm1-FC5_data2_mnt00002
    }
    multipath {
        wwid      3600a098038314e63532459326d495a64
        alias     svm1-FC5_data_mnt00001
    }
    multipath {
        wwid      3600a098038314e63532459326d495a65
        alias     svm1-FC5_data_mnt00002
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786f
        alias     svm1-FC5_log2_mnt00001
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b7870
        alias     svm1-FC5_log2_mnt00002
    }
    multipath {
        wwid      3600a098038314e63532459326d495a66
        alias     svm1-FC5_log_mnt00001
    }
    multipath {
        wwid      3600a098038314e63532459326d495a67
        alias     svm1-FC5_log_mnt00002
    }
}

```

5. Run the `multipath -r` command to reload the device map.
6. Verify the configuration by running the `multipath -ll` command to list all the LUNs, alias names, and active and standby paths.



The following example output shows the output of a 2+1 multiple-host HANA system with two data and two log LUNs.

```

sapcc-hana-tst:~ # multipath -ll
hsvm1-FC5_data2_mnt00001 (3600a098038314e63492b59326b4b786d) dm-7
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:2 sdaf 65:240 active ready running
  |- 20:0:5:2 sdx 65:112 active ready running
  |- 21:0:4:2 sdav 66:240 active ready running
  `-- 21:0:6:2 sdan 66:112 active ready running
svm1-FC5_data2_mnt00002 (3600a098038314e63492b59326b4b786e) dm-9
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:4 sdah 66:16 active ready running
  |- 20:0:5:4 sdz 65:144 active ready running
  |- 21:0:4:4 sdax 67:16 active ready running
  `-- 21:0:6:4 sdap 66:144 active ready running
svm1-FC5_data_mnt00001 (3600a098038314e63532459326d495a64) dm-6
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:1 sdae 65:224 active ready running
  |- 20:0:5:1 sdw 65:96 active ready running
  |- 21:0:4:1 sdau 66:224 active ready running
  `-- 21:0:6:1 sdam 66:96 active ready running
svm1-FC5_data_mnt00002 (3600a098038314e63532459326d495a65) dm-8
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:3 sdag 66:0 active ready running
  |- 20:0:5:3 sdy 65:128 active ready running
  |- 21:0:4:3 sdaw 67:0 active ready running
  `-- 21:0:6:3 sdao 66:128 active ready running
svm1-FC5_log2_mnt00001 (3600a098038314e63492b59326b4b786f) dm-11
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:6 sdaj 66:48 active ready running
  |- 20:0:5:6 sdab 65:176 active ready running
  |- 21:0:4:6 sdaz 67:48 active ready running
  `-- 21:0:6:6 sdar 66:176 active ready running

```

```

svm1-FC5_log2_mnt00002 (3600a098038314e63492b59326b4b7870) dm-13
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:8 sdal 66:80 active ready running
  |- 20:0:5:8 sdad 65:208 active ready running
  |- 21:0:4:8 sdbb 67:80 active ready running
  `-- 21:0:6:8 sdat 66:208 active ready running
svm1-FC5_log_mnt00001 (3600a098038314e63532459326d495a66) dm-10
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:5 sdai 66:32 active ready running
  |- 20:0:5:5 sdaa 65:160 active ready running
  |- 21:0:4:5 sday 67:32 active ready running
  `-- 21:0:6:5 sdaq 66:160 active ready running
svm1-FC5_log_mnt00002 (3600a098038314e63532459326d495a67) dm-12
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:7 sdak 66:64 active ready running
  |- 20:0:5:7 sdac 65:192 active ready running
  |- 21:0:4:7 sdba 67:64 active ready running
  `-- 21:0:6:7 sdas 66:192 active ready running

```

## Single host setup

This chapter describes the setup of an SAP HANA single host using LINUX LVM.

### LUN configuration for SAP HANA single-host systems

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

Logical volume/LUN	Mount point at SAP HANA host	Note
LV: FC5_data_mnt0000-vol	/hana/data/FC51/mnt00001	Mounted using /etc/fstab entry
LV: FC5_log_mnt00001-vol	/hana/log/FC5/mnt00001	Mounted using /etc/fstab entry
LUN: FC5_shared	/hana/shared/FC5	Mounted using /etc/fstab entry



With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user `FC5adm` is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `FC5_shared` volume for the `/usr/sap/FC5` directory so that all file systems are on the central storage.

## Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```
pvcreate /dev/mapper/hana-FC5_data_mnt00001
pvcreate /dev/mapper/hana-FC5_data2_mnt00001
pvcreate /dev/mapper/hana-FC5_log_mnt00001
pvcreate /dev/mapper/hana-FC5_log2_mnt00001
```

2. Create the volume groups for each data and log partition.

```
vgcreate FC5_data_mnt00001 /dev/mapper/hana-FC5_data_mnt00001
/dev/mapper/hana-FC5_data2_mnt00001
vgcreate FC5_log_mnt00001 /dev/mapper/hana-FC5_log_mnt00001
/dev/mapper/hana-FC5_log2_mnt00001
```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00001
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00001
```

4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```
modprobe dm_mod
pvscan
vgscan
lvscan
```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```
vgchange -a y
```

## Create file systems

Create the XFS file system on all data and log logical volumes and the hana shared LUN.

```
mkfs.xfs FC5_data_mnt00001-vol
mkfs.xfs FC5_log_mnt00001-vol
mkfs.xfs /dev/mapper/svm1-FC5_shared
```

## Create mount points

Create the required mount point directories, and set the permissions on the database host:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

## Mount file systems

To mount file systems during system boot using the `/etc/fstab` configuration file, add the required file systems to the `/etc/fstab` configuration file:

```
# cat /etc/fstab
/dev/mapper/hana-FC5_shared /hana/shared xfs defaults 0 0
/dev/mapper/FC5_log_mnt00001-vol /hana/log/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/FC5_data_mnt00001-vol /hana/data/FC5/mnt00001 xfs
relatime,inode64 0 0
```



The XFS file systems for the data and log LUNs must be mounted with the `relatime` and `inode64` mount options.

To mount the file systems, run the `mount -a` command at the host.

## Multiple hosts setup

This chapter describes the setup of a 2+1 SAP HANA multiple host system as example.

### LUN configuration for SAP HANA multiple-hosts systems

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

Logical volume (LV) or volume	Mount point at SAP HANA host	Note
LV: FC5_data_mnt00001-vol	/hana/data/FC5/mnt00001	Mounted using storage connector
LV: FC5_log_mnt00001-vol	/hana/log/FC5/mnt00001	Mounted using storage connector
LV: FC5_data_mnt00002-vol	/hana/data/FC5/mnt00002	Mounted using storage connector
LV: FC5_log_mnt00002-vol	/hana/log/FC5/mnt00002	Mounted using storage connector
Volume: FC5_shared	/hana/shared	Mounted at all hosts using NFS and /etc/fstab entry



With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user FC5adm is stored, is on the local disk for each HANA host. In a disaster recovery setup with disk-based replication, NetApp recommends creating four additional subdirectories in the `FC5_shared` volume for the `/usr/sap/FC5` file system so that each database host has all its file systems on the central storage.

## Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```
pvccreate /dev/mapper/hana-FC5_data_mnt00001
pvccreate /dev/mapper/hana-FC5_data2_mnt00001
pvccreate /dev/mapper/hana-FC5_data_mnt00002
pvccreate /dev/mapper/hana-FC5_data2_mnt00002
pvccreate /dev/mapper/hana-FC5_log_mnt00001
pvccreate /dev/mapper/hana-FC5_log2_mnt00001
pvccreate /dev/mapper/hana-FC5_log_mnt00002
pvccreate /dev/mapper/hana-FC5_log2_mnt00002
```

2. Create the volume groups for each data and log partition.

```
vgcreate FC5_data_mnt00001 /dev/mapper/hana-FC5_data_mnt00001
/dev/mapper/hana-FC5_data2_mnt00001
vgcreate FC5_data_mnt00002 /dev/mapper/hana-FC5_data_mnt00002
/dev/mapper/hana-FC5_data2_mnt00002
vgcreate FC5_log_mnt00001 /dev/mapper/hana-FC5_log_mnt00001
/dev/mapper/hana-FC5_log2_mnt00001
vgcreate FC5_log_mnt00002 /dev/mapper/hana-FC5_log_mnt00002
/dev/mapper/hana-FC5_log2_mnt00002
```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00001
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00001
```

#### 4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```
modprobe dm_mod
pvscan
vgscan
lvscan
```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```
vgchange -a y
```

#### Create file systems

Create the XFS file system on all data and log logical volumes.

```
mkfs.xfs FC5_data_mnt00001-vol
mkfs.xfs FC5_data_mnt00002-vol
mkfs.xfs FC5_log_mnt00001-vol
mkfs.xfs FC5_log_mnt00002-vol
```

#### Create mount points

Create the required mount point directories, and set the permissions on all worker and standby hosts:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00002
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00002
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

## Mount file systems

To mount the `/hana/shared` file systems during system boot using the `/etc/fstab` configuration file, add the `/hana/shared` file system to the `/etc/fstab` configuration file of each host.

```
sapcc-hana-tst:/ # cat /etc/fstab
<storage-ip>:/hana_shared /hana/shared nfs rw,vers=3,hard,timeo=600,
intr,noatime,nolock 0 0
```



All the data and log file systems are mounted through the SAP HANA storage connector.

To mount the file systems, run the `mount -a` command at each host.

## I/O Stack configuration for SAP HANA

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used.

NetApp conducted performance tests to define the ideal values. The following table lists the optimal values as inferred from the performance tests.

Parameter	Value
<code>max_parallel_io_requests</code>	128
<code>async_read_submit</code>	on
<code>async_write_submit_active</code>	on
<code>async_write_submit_blocks</code>	all

For SAP HANA 1.0 up to SPS12, these parameters can be set during the installation of the SAP HANA database, as described in SAP Note [2267798 – Configuration of the SAP HANA Database during Installation Using hdbparam](#).

Alternatively, the parameters can be set after the SAP HANA database installation by using the `hdbparam` framework.

```
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.max_parallel_io_requests=128
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_write_submit_active=on
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_read_submit=on
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_write_submit_blocks=all
```

Starting with SAP HANA 2.0, `hdbparam` is deprecated, and the parameters are moved to the `global.ini` file. The parameters can be set by using SQL commands or SAP HANA Studio. For more details, refer to SAP

note [2399079: Elimination of hdbparam in HANA 2](#). The parameters can be also set within the `global.ini` file.

```
SS3adm@stlrx300s8-6: /usr/sap/SS3/SYS/global/hdb/custom/config> cat
global.ini
...
[fileio]
async_read_submit = on
async_write_submit_active = on
max_parallel_io_requests = 128
async_write_submit_blocks = all
...
```

For SAP HANA 2.0 SPS5 and later, use the `setParameter.py` script to set the correct parameters.

```
fc5adm@sapcc-hana-tst-03:/usr/sap/FC5/HDB00/exe/python_support>
python setParameter.py
-set=SYSTEM/global.ini/fileio/max_parallel_io_requests=128
python setParameter.py -set=SYSTEM/global.ini/fileio/async_read_submit=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_active=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_blocks=all
```

## SAP HANA software installation

This section describes the preparation necessary to install SAP HANA on single-host and multiple-host systems.

### Installation on single-host system

SAP HANA software installation does not require any additional preparation for a single-host system.

### Installation on multiple-host system

Before beginning the installation, create a `global.ini` file to enable use of the SAP storage connector during the installation process. The SAP storage connector mounts the required file systems at the worker hosts during the installation process. The `global.ini` file must be available in a file system that is accessible from all hosts, such as the `/hana/shared` file system.

Before installing SAP HANA software on a multiple-host system, the following steps must be completed:

1. Add the following mount options for the data LUNs and the log LUNs to the `global.ini` file:
  - `relatime` and `inode64` for the data and log file system
2. Add the WWIDs of the data and log partitions. The WWIDs must match the alias names configured in the `/etc/multipath.conf` file.

The following output shows an example of a 2+1 multiple-host setup using LVM with SID=FC5.

```
sapcc-hana-tst-03:/hana/shared # cat global.ini
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/FC5
basepath_logvolumes = /hana/log/FC5
[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*_prtype = 5
partition_*_data__mountOptions = -o relatime,inode64
partition_*_log__mountOptions = -o relatime,inode64
partition_1_data__lvmname = FC5_data_mnt00001-vol
partition_1_log__lvmname = FC5_log_mnt00001-vol
partition_2_data__lvmname = FC5_data_mnt00002-vol
partition_2_log__lvmname = FC5_log_mnt00002-vol
sapcc-hana-tst-03:/hana/shared #
```

Using the SAP hdblcm installation tool, start the installation by running the following command at one of the worker hosts. Use the `addhosts` option to add the second worker (`sapcc-hana-tst-06`) and the standby host (`sapcc-hana-tst-07`).



The directory where the prepared `global.ini` file is stored is included with the `storage_cfg` CLI option (`--storage_cfg=/hana/shared`).



Depending on the OS version being used, it might be necessary to install Python 2.7 before installing the SAP HANA database.

```
./hdblcm --action=install --addhosts=sapcc-hana-tst-06:role=worker:storage_partition=2,sapcc-hana-tst-07:role=standby
--storage_cfg=/hana/shared/

AP HANA Lifecycle Management - SAP HANA Database 2.00.073.00.1695288802
*****

Scanning software locations...
Detected components:
    SAP HANA AFL (incl.PAL,BFL,OFL) (2.00.073.0000.1695321500) in
/mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
    SAP HANA Database (2.00.073.00.1695288802) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
```

SAP HANA Database Client (2.18.24.1695756995) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/HDB\_CLIENT\_LINUX\_X86\_64/SAP\_HANA\_CLIENT/client

SAP HANA Studio (2.3.75.000000) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/HDB\_STUDIO\_LINUX\_X86\_64/studio

SAP HANA Local Secure Store (2.11.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/HANA\_LSS\_24\_LINUX\_X86\_64/packages

SAP HANA XS Advanced Runtime (1.1.3.230717145654) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_RT\_10\_LINUX\_X86\_64/packages

SAP HANA EML AFL (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/HDB\_EML\_AFL\_10\_LINUX\_X86\_64/packages

SAP HANA EPM-MDS (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/SAP\_HANA\_EPM-MDS\_10/packages

Automated Predictive Library (4.203.2321.0.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/PAAPL4\_H20\_LINUX\_X86\_64/apl-4.203.2321.0-hana2sp03-linux\_x64/installer/packages

GUI for HALM for XSA (including product installer) Version 1 (1.015.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACALMPIUI15\_0.zip

XSAC FILEPROCESSOR 1.0 (1.000.102) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACFILEPROC00\_102.zip

SAP HANA tools for accessing catalog content, data preview, SQL console, etc. (2.015.230503) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSAC\_HRTT\_20/XSACHRTT15\_230503.zip

Develop and run portal services for customer applications on XSA (2.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACPORTALSERV07\_0.zip

The SAP Web IDE for HANA 2.0 (4.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSAC\_SAP\_WEB\_IDE\_20/XSACSAPWEBIDE07\_0.zip

XS JOB SCHEDULER 1.0 (1.007.22) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACSERVICES07\_22.zip

SAPUI5 FESV6 XSA 1 - SAPUI5 1.71 (1.071.52) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV671\_52.zip

SAPUI5 FESV9 XSA 1 - SAPUI5 1.108 (1.108.5) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV9108\_5.zip

SAPUI5 SERVICE BROKER XSA 1 - SAPUI5 Service Broker 1.0 (1.000.4) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

```
73/DATA_UNITS/XSA_CONTENT_10/XSACUI5SB00_4.zip
XSA Cockpit 1 (1.001.37) in /mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACXSACOCKPIT01_37.zip
```

SAP HANA Database version '2.00.073.00.1695288802' will be installed.

Select additional components for installation:

Index	Components	Description
1	all	All components
2	server	No additional components
3	client	Install SAP HANA Database Client version 2.18.24.1695756995
4	lss	Install SAP HANA Local Secure Store version 2.11.0
5	studio	Install SAP HANA Studio version 2.3.75.000000
6	xs	Install SAP HANA XS Advanced Runtime version 1.1.3.230717145654
7	afl	Install SAP HANA AFL (incl.PAL,BFL,OFL) version 2.00.073.0000.1695321500
8	eml	Install SAP HANA EML AFL version 2.00.073.0000.1695321500
9	epmmds	Install SAP HANA EPM-MDS version 2.00.073.0000.1695321500
10	sap_afl_sdk_apl	Install Automated Predictive Library version 4.203.2321.0.0

Enter comma-separated list of the selected indices [3,4]: 2,3

3. Verify that the installation tool installed all selected components at all worker and standby hosts.

### Adding additional data volume partitions for SAP HANA single-host systems

Starting with SAP HANA 2.0 SPS4, additional data volume partitions can be configured. This feature allows you to configure two or more LUNs for the data volume of an SAP HANA tenant database and to scale beyond the size and performance limits of a single LUN.



It is not necessary to use multiple partitions to fulfill the SAP HANA KPIs. A single LUN with a single partition fulfills the required KPIs.



Using two or more individual LUNs for the data volume is only available for SAP HANA single-host systems. The SAP storage connector required for SAP HANA multiple-host systems does only support one device for the data volume.

Adding additional data volume partitions can be done at any time but might require a restart of the SAP HANA database.

### Enabling additional data volume partitions

To enable additional data volume partitions, complete the following steps:

1. Add the following entry within the `global.ini` file.

```
[customizable_functionalities]
persistence_datavolume_partition_multipath = true
```

2. Restart the database to enable the feature. Adding the parameter through the SAP HANA Studio to the `global.ini` file by using the Systemdb configuration prevents the restart of the database.

### Volume and LUN configuration

The layout of volumes and LUNs is like the layout of a single host with one data volume partition, but with an additional data volume and LUN stored on a different aggregate as the log volume and the other data volume. The following table shows an example configuration of an SAP HANA single-host systems with two data volume partitions.

Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data volume: SID_data_mnt00001	Shared volume: SID_shared	Data volume: SID_data2_mnt00001	Log volume: SID_log_mnt00001

The following table shows an example of the mount point configuration for a single-host system with two data volume partitions.

LUN	Mount point at HANA host	Note
SID_data_mnt00001	/hana/data/SID/mnt00001	Mounted using /etc/fstab entry
SID_data2_mnt00001	/hana/data2/SID/mnt00001	Mounted using /etc/fstab entry
SID_log_mnt00001	/hana/log/SID/mnt00001	Mounted using /etc/fstab entry
SID_shared	/hana/shared/SID	Mounted using /etc/fstab entry

Create the new data LUNs using either ONTAP System Manager or the ONTAP CLI.

### Host configuration

To configure a host, complete the following steps:

1. Configure multipathing for the additional LUNs, as described in chapter [Host Setup](#).
2. Create the XFS file system on each additional LUN belonging to the HANA system:

```
stlrx300s8-6:/ # mkfs.xfs /dev/mapper/hana-FC5_data2_mnt00001
```

3. Add the additional file system/s to the `/etc/fstab` configuration file.



The XFS file systems for the data and log LUN must be mounted with the `relatime` and `inode64` mount options.

```
stlrx300s8-6:/ # cat /etc/fstab
/dev/mapper/hana-FC5_shared /hana/shared xfs defaults 0 0
/dev/mapper/hana-FC5_log_mnt00001 /hana/log/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/hana-FC5_data_mnt00001 /hana/data/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/hana-FC5_data2_mnt00001 /hana/data2/FC5/mnt00001 xfs
relatime,inode64 0 0
```

4. Create mount points and set permissions on the database host.

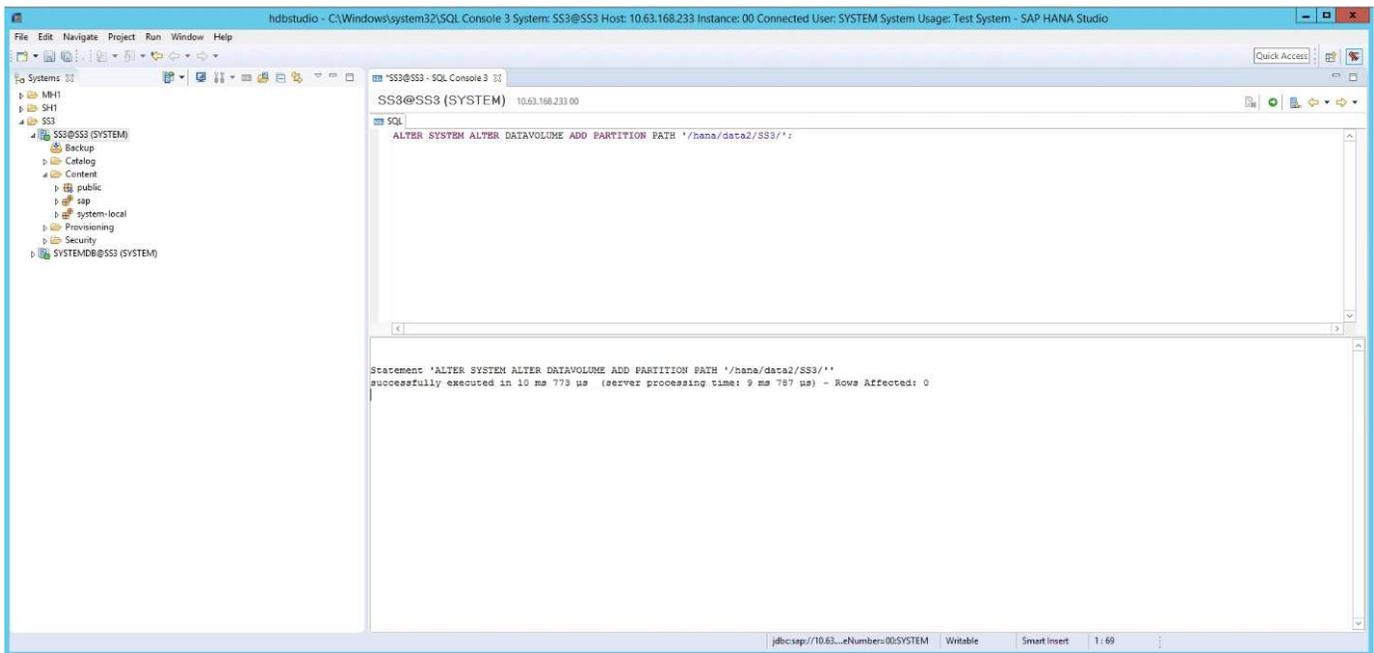
```
stlrx300s8-6:/ # mkdir -p /hana/data2/FC5/mnt00001
stlrx300s8-6:/ # chmod -R 777 /hana/data2/FC5
```

5. Mount the file systems, run the `mount -a` command.

#### Adding an additional datavolume partition

To add an additional datavolume partition to your tenant database, execute the following SQL statement against the tenant database. Each additional LUN can have a different path:

```
ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION PATH '/hana/data2/SID/';
```



## Where to find additional information

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

- [SAP HANA Software Solutions](#)
- [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- [NetApp Documentation Centers](#)

<https://www.netapp.com/support-and-training/documentation/>

- [SAP Certified Enterprise Storage Hardware for SAP HANA](#)

<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/>

- [SAP HANA Storage Requirements](#)

<https://www.sap.com/documents/2024/03/146274d3-ae7e-0010-bca6-c68f7e60039b.html>

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

<https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html>

- [SAP HANA on VMware vSphere Wiki](#)

[https://help.sap.com/docs/SUPPORT\\_CONTENT/virtualization/3362185751.html](https://help.sap.com/docs/SUPPORT_CONTENT/virtualization/3362185751.html)

- [SAP HANA on VMware vSphere Best Practices Guide](#)

[https://www.vmware.com/docs/sap\\_hana\\_on\\_vmware\\_vsphere\\_best\\_practices\\_guide-white-paper](https://www.vmware.com/docs/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper)

## Update history

The following technical changes have been made to this solution since its original publication.

Date	Update summary
October 2015	Initial Version
March 2016	Updated capacity sizing
February 2017	New NetApp storage systems and disk shelves New features of ONTAP 9 New OS releases (SLES12 SP1 and RHEL 7.2) New SAP HANA release
July 2017	Minor updates
September 2018	New NetApp storage systems New OS releases (SLES12 SP3 and RHEL 7.4) Additional minor updates SAP HANA 2.0 SPS3
November 2019	New NetApp storage systems and NVMe shelf New OS releases (SLES12 SP4, SLES 15, and RHEL 7.6) Additional minor updates
April 2020	New AFF ASA series storage systems Introduced multiple data partition feature available since SAP HANA 2.0 SPS4
June 2020	Additional information about optional functionalities Minor updates
February 2021	Linux LVM support New NetApp storage systems New OS releases (SLES15SP2, RHEL 8)
April 2021	VMware vSphere-specific information added
September 2022	New OS-Releases
August 2023	New Storage Systems (AFF C-Series)
May 2024	New Storage Systems (AFF A-Series)
September 2024	New Storage Systems (ASA A-Series)
November 2024	New Storage Systems
February 2025	New Storage Systems
July 2025	Minor updates

## SAP HANA on NetApp AFF Systems with NFS Configuration Guide

## SAP HANA on NetApp AFF Systems with NFS - Configuration Guide

The NetApp AFF A-Series and AFF C-Series product families have been certified for use with SAP HANA in tailored data center integration (TDI) projects. This guide provides best practices for SAP HANA on this platform for NFS.

Marco Schoen, NetApp

This certification is valid for the following models:

- AFF A150, AFF A20, AFF A250, AFF A30, AFF A400, AFF A50, AFF A70, AFF A800, AFF A900, AFF A90, AFF A1K
- AFF C250, AFF C30, AFF C400, AFF C60, AFF C800, AFF C80



NetApp AFF C-Series requires NetApp ONTAP 9.13.1 or later

A complete list of NetApp certified storage solutions for SAP HANA can be found at the [Certified and supported SAP HANA hardware directory](#).

This document describes the ONTAP configuration requirements for the NFS protocol version 3 (NFSv3) or the NFS protocol version 4 (NFSv4.1).



Only NFS versions 3 or 4.1 are supported. NFS versions 1, 2, 4.0, and 4.2 aren't supported.



The configuration described in this paper is necessary to achieve the required SAP HANA KPIs and the best performance for SAP HANA. Changing any settings or using features not listed herein might cause performance degradation or unexpected behavior and should only be done if advised by NetApp support.

The configuration guides for NetApp AFF systems using FCP and for FAS systems using NFS or FCP can be found at the following links:

- [SAP HANA on NetApp FAS Systems with FCP](#)
- [SAP HANA on NetApp FAS Systems with NFS](#)
- [SAP HANA on NetApp AFF Systems with FCP](#)
- [SAP HANA on NetApp ASA Systems with FCP](#)

The following table shows the supported combinations for NFS versions, NFS locking, and the required isolation implementations, depending on the SAP HANA database configuration.

For SAP HANA single-host systems or multiple hosts that do not use Host Auto-Failover, NFSv3 and NFSv4 are supported.

For SAP HANA multiple host systems with Host Auto-Failover, NetApp only supports NFSv4, while using NFSv4 locking as an alternative to a server-specific STONITH (SAP HANA HA/DR provider) implementation.

SAP HANA	NFS version	NFS locking	SAP HANA HA/DR provider
SAP HANA single host, multiple hosts without Host Auto-Failover	NFSv3	Off	n/a
	NFSv4	On	n/a
SAP HANA multiple hosts using Host Auto-Failover	NFSv3	Off	Server-specific STONITH implementation mandatory
	NFSv4	On	Not required



A server-specific STONITH implementation is not part of this guide. Contact your server vendor for such an implementation.

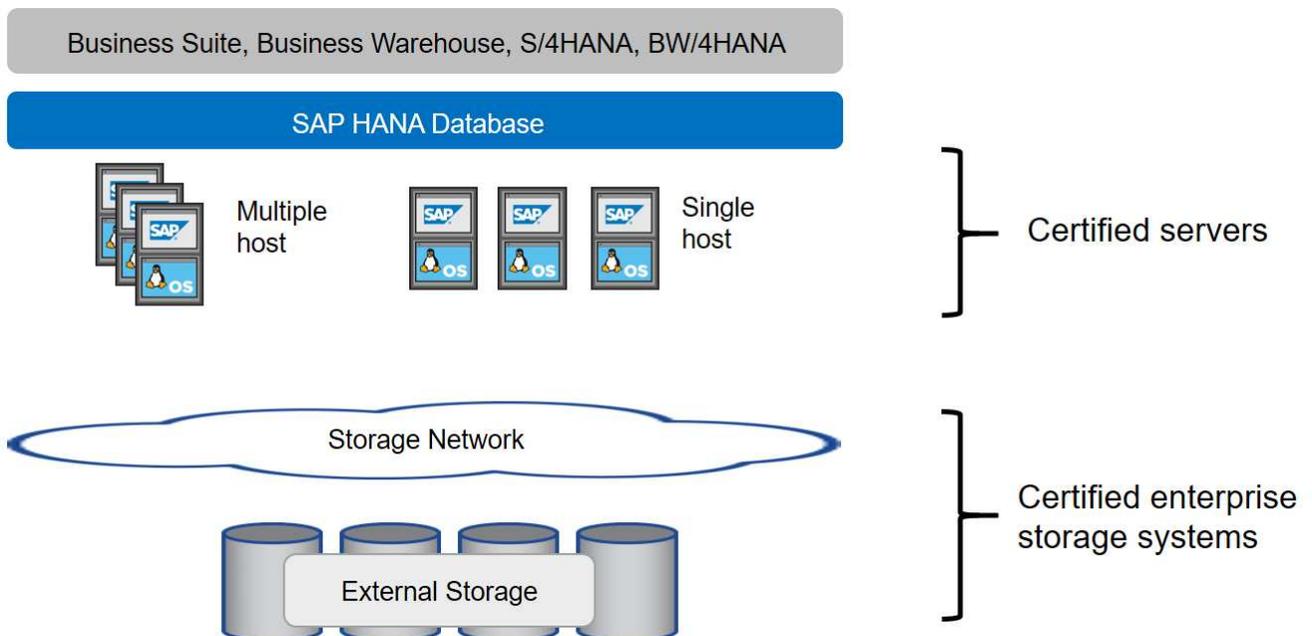
This document covers configuration recommendations for SAP HANA running on physical servers and on virtual servers that use VMware vSphere.



See the relevant SAP notes for operating system configuration guidelines and HANA-specific Linux kernel dependencies. For more information, see SAP note 2235581: SAP HANA Supported Operating Systems.

### SAP HANA tailored data center integration

NetApp AFF storage controllers are certified in the SAP HANA TDI program using both NFS (NAS) and FC (SAN) protocols. They can be deployed in any of the current SAP HANA scenarios, such as SAP Business Suite on HANA, S/4HANA, BW/4HANA, or SAP Business Warehouse on HANA in either single-host or multiple-host configurations. Any server that is certified for use with SAP HANA can be combined with NetApp certified storage solutions. See the following figure for an architecture overview of SAP HANA TDI.



For more information regarding the prerequisites and recommendations for producti SAP HANA systems, see the following resource:

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

## **SAP HANA using VMware vSphere**

There are several options for connecting storage to virtual machines (VMs). The preferred option is to connect the storage volumes with NFS directly out of the guest operating system. Using this option, the configuration of hosts and storage does not differ between physical hosts and VMs.

NFS datastores and VVOL datastores with NFS are supported as well. For both options, only one SAP HANA data or log volume must be stored within the datastore for production use cases.

This document describes the recommended setup with direct NFS mounts from the guest OS.

For more information about using vSphere with SAP HANA, see the following links:

- [SAP HANA on VMware vSphere - Virtualization - Community Wiki](#)
- [SAP HANA on VMware vSphere Best Practices Guide](#)
- [2161991 - VMware vSphere configuration guidelines - SAP ONE Support Launchpad \(Login required\)](#)

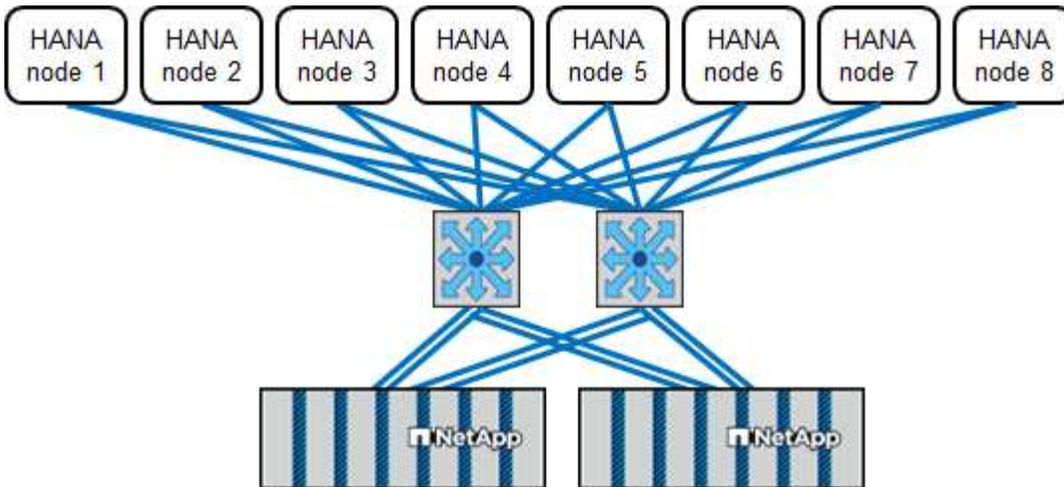
## **Architecture**

SAP HANA hosts are connected to storage controllers by using a redundant 10GbE or faster network infrastructure. Data communication between SAP HANA hosts and storage controllers is based on the NFS protocol. A redundant switching infrastructure is required to provide fault-tolerant SAP HANA host-to-storage connectivity in case of switch or network interface card (NIC) failure.

The switches might aggregate individual port performance with port channels in order to appear as a single logical entity at the host level.

Different models of the AFF system product family can be mixed and matched at the storage layer to allow for growth and differing performance and capacity needs. The maximum number of SAP HANA hosts that can be attached to the storage system is defined by the SAP HANA performance requirements and the model of NetApp controller used. The number of required disk shelves is only determined by the capacity and performance requirements of the SAP HANA systems.

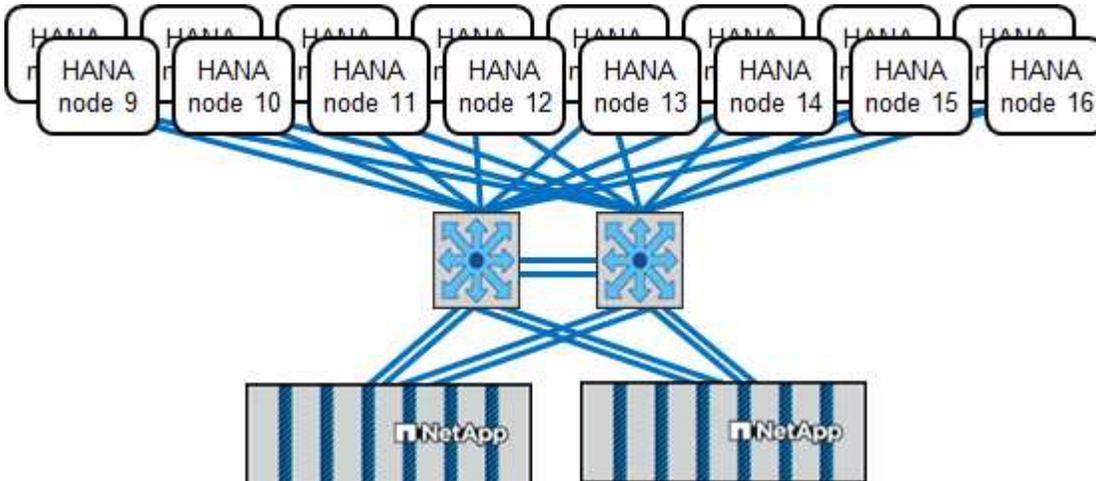
The following figure shows an example configuration with eight SAP HANA hosts attached to a storage high availability (HA) pair.



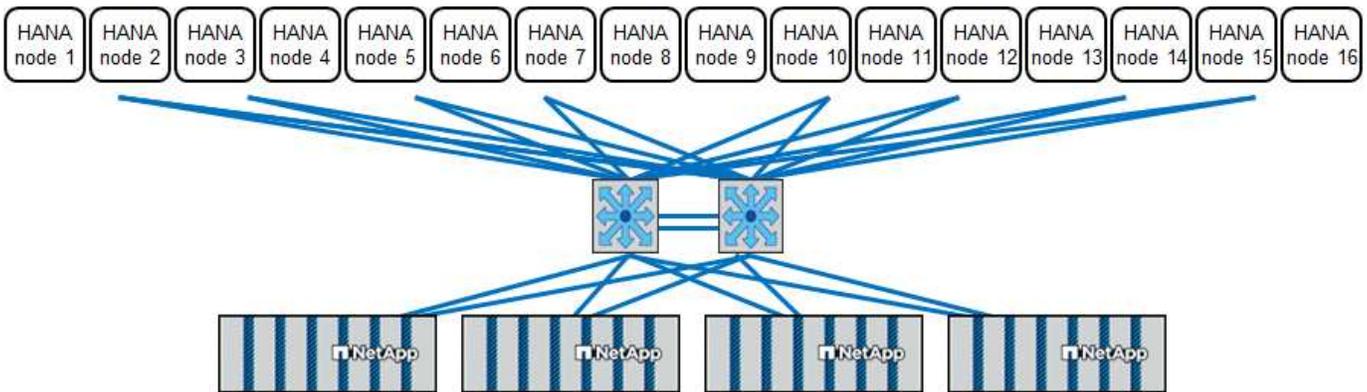
The architecture can be scaled in two dimensions:

- By attaching additional SAP HANA hosts and storage capacity to the existing storage, if the storage controllers provide enough performance to meet the current SAP HANA key performance indicators (KPIs).
- By adding more storage systems with additional storage capacity for the additional SAP HANA hosts

The following figure shows an example configuration in which more SAP HANA hosts are attached to the storage controllers. In this example, more disk shelves are necessary to fulfill the capacity and performance requirements of the 16 SAP HANA hosts. Depending on the total throughput requirements, you must add additional 10GbE or faster connections to the storage controllers.



Independent of the deployed AFF system, the SAP HANA landscape can also be scaled by adding any of the certified storage controllers to meet the desired node density, as shown in the following figure.



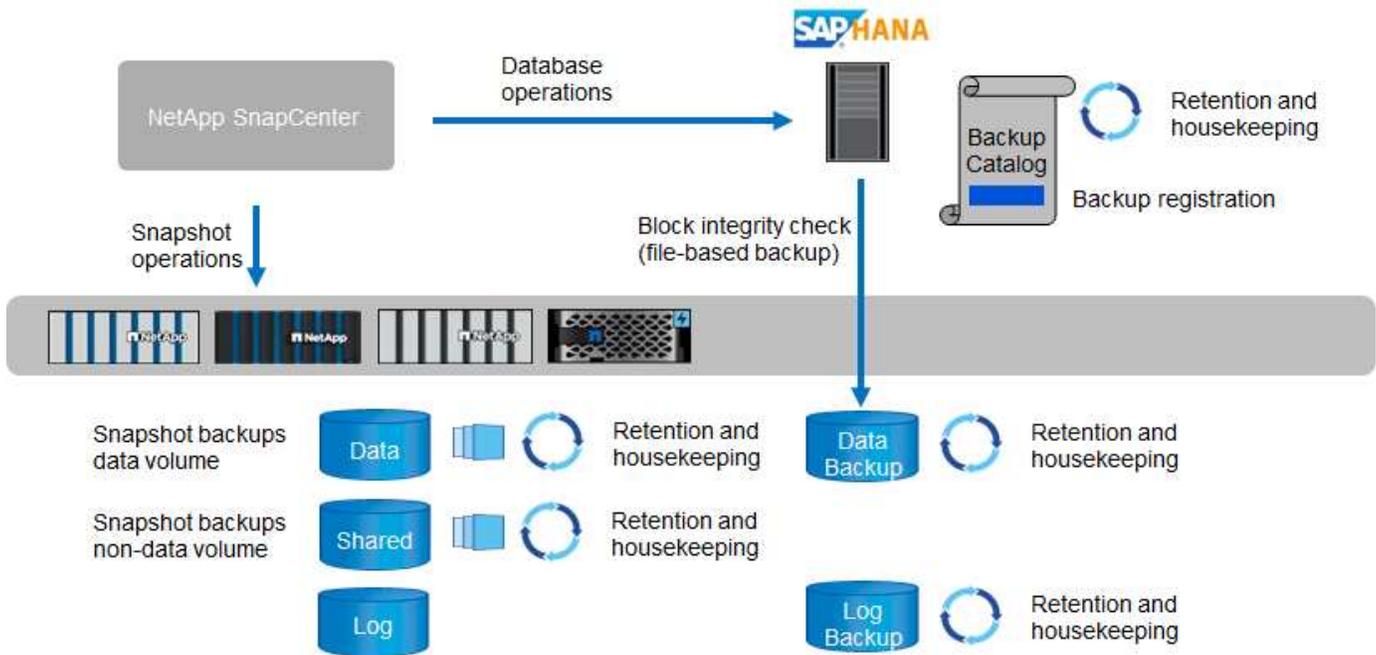
## SAP HANA backup

The ONTAP software present on all NetApp storage controllers provides a built-in mechanism to back up SAP HANA databases while in operation with no effect on performance. Storage-based NetApp Snapshot backups are a fully supported and integrated backup solution available for SAP HANA single containers and for SAP HANA Multitenant Database Containers (MDC) systems with a single tenant or multiple tenants.

Storage-based Snapshot backups are implemented by using the NetApp SnapCenter plug-in for SAP HANA. This allows users to create consistent storage-based Snapshot backups by using the interfaces provided natively by SAP HANA databases. SnapCenter registers each of the Snapshot backups into the SAP HANA backup catalog. Therefore, the backups taken by SnapCenter are visible within SAP HANA Studio and Cockpit where they can be selected directly for restore and recovery operations.

NetApp SnapMirror technology enables Snapshot copies that were created on one storage system to be replicated to a secondary backup storage system that is controlled by SnapCenter. Different backup retention policies can then be defined for each of the backup sets on the primary storage and for the backup sets on the secondary storage systems. The SnapCenter Plug-in for SAP HANA automatically manages the retention of Snapshot copy-based data backups and log backups, including the housekeeping of the backup catalog. The SnapCenter Plug-in for SAP HANA also allows the execution of a block integrity check of the SAP HANA database by executing a file-based backup.

The database logs can be backed up directly to the secondary storage by using an NFS mount, as shown in the following figure.



Storage-based Snapshot backups provide significant advantages compared to conventional file-based backups. These advantages include, but are not limited to, the following:

- Faster backup (a few minutes)
- Reduced recovery time objective (RTO) due to a much faster restore time on the storage layer (a few minutes) as well as more frequent backups
- No performance degradation of the SAP HANA database host, network, or storage during backup and recovery operations
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes



For detailed information about the SAP HANA backup and recovery solution see [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

## SAP HANA disaster recovery

SAP HANA disaster recovery (DR) can be done either on the database layer by using SAP HANA system replication or on the storage layer by using storage replication technologies. The following section provides an overview of disaster recovery solutions based on storage replication.

For detailed information about SAP HANA disaster recovery solutions, see [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#).

### Storage replication based on SnapMirror

The following figure shows a three-site disaster recovery solution using synchronous SnapMirror replication to the local DR datacenter and asynchronous SnapMirror to replicate the data to the remote DR datacenter.

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

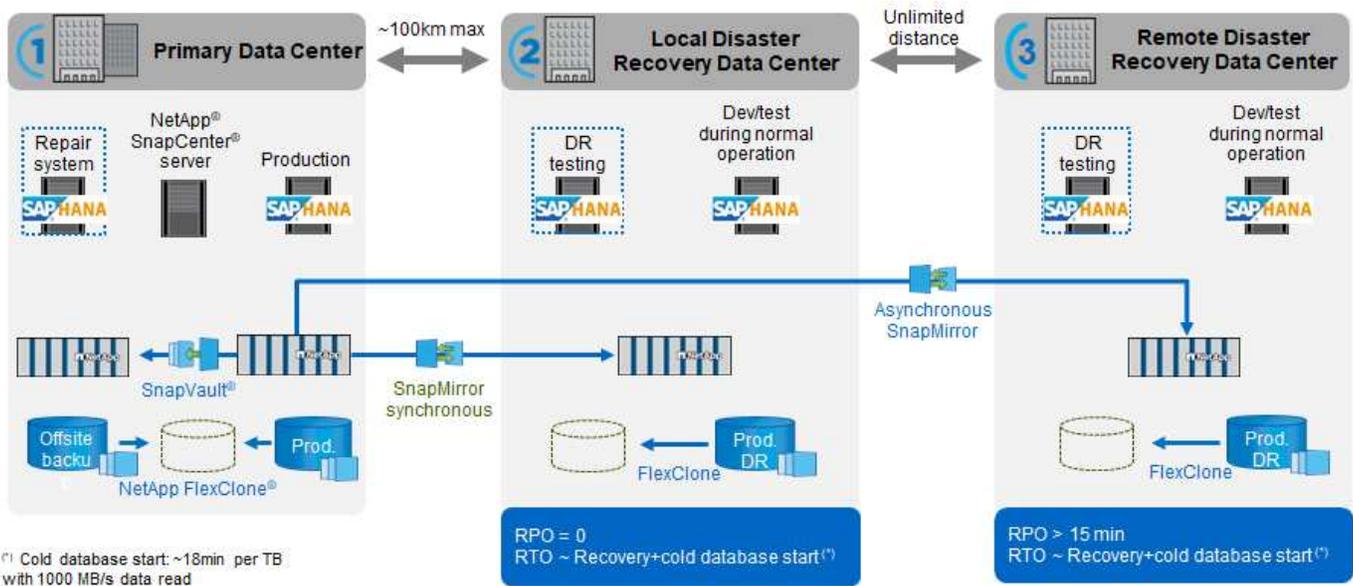
Protection against failures of both the primary and the local DR site is performed by replicating the data to a third remote DR datacenter using asynchronous SnapMirror. The RPO depends on the frequency of replication

updates and how fast they can be transferred. In theory, the distance is unlimited, but the limit depends on the amount of data that must be transferred and the connection that is available between the data centers. Typical RPO values are in the range of 30 minutes to multiple hours.

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

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

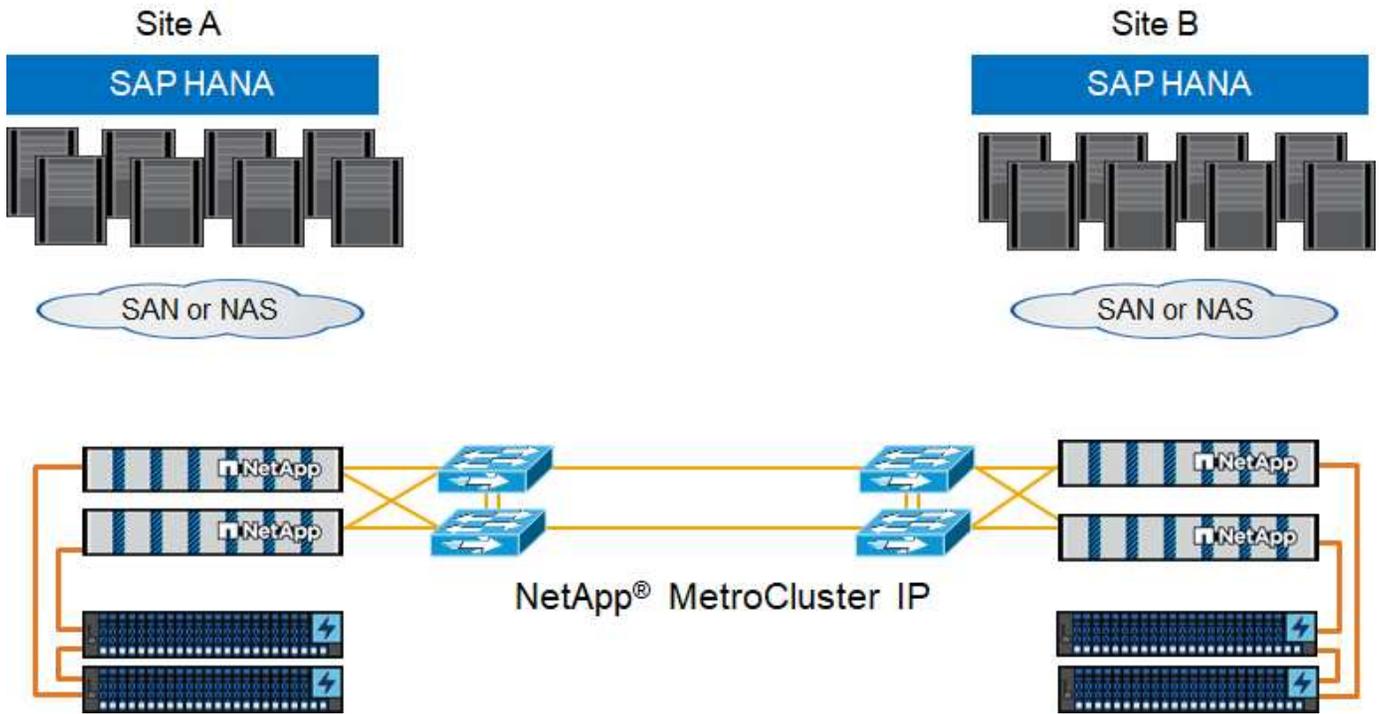
Both replication methods allow to you execute DR workflow testing without influencing the RPO and RTO. FlexClone volumes are created on the storage and are attached to the DR testing servers.



Synchronous replication offers StrictSync mode. If the write to secondary storage is not completed for any reason, the application I/O fails, thereby ensuring that the primary and secondary storage systems are identical. Application I/O to the primary resumes only after the SnapMirror relationship returns to the InSync status. If the primary storage fails, application I/O can be resumed on the secondary storage after failover with no loss of data. In StrictSync mode, the RPO is always zero.

### Storage replication based on MetroCluster

The following figure shows a high-level overview of the solution. The storage cluster at each site provides local high availability and is used for the production workload. The data of each site is synchronously replicated to the other location and is available in case of disaster failover.



## Storage sizing

The following section provides an overview of the required performance and capacity considerations needed for sizing a storage system for SAP HANA.



Contact NetApp or your NetApp partner sales representative to assist you in creating a properly sized storage environment.

## Performance considerations

SAP has defined a static set of storage KPIs. These KPIs are valid for all production SAP HANA environments independent of the memory size of the database hosts and the applications that use the SAP HANA database. These KPIs are valid for single-host, multiple-host, Business Suite on HANA, Business Warehouse on HANA, S/4HANA, and BW/4HANA environments. Therefore, the current performance sizing approach depends on only the number of active SAP HANA hosts that are attached to the storage system.



Storage performance KPIs are only mandated for production SAP HANA systems, but you can implement them in for all HANA system.

SAP delivers a performance test tool that must be used to validate the storage system's performance for active SAP HANA hosts attached to the storage.

NetApp tested and predefined the maximum number of SAP HANA hosts that can be attached to a specific storage model while still fulfilling the required storage KPIs from SAP for production-based SAP HANA systems.

The maximum number of SAP HANA hosts that can be run on a disk shelf and the minimum number of SSDs required per SAP HANA host were determined by running the SAP performance test tool. This test does not consider the actual storage capacity requirements of the hosts. You must also calculate the capacity requirements to determine the actual storage configuration needed.

## SAS disk shelf

With the 12Gb serial-attached SCSI (SAS) disk shelf (DS224C), performance sizing is performed by using the following fixed disk-shelf configurations:

- Half-loaded disk shelves with 12 SSDs
- Fully loaded disk shelves with 24 SSDs



Both configurations use Advanced Disk Partitioning (ADPv2). A half-loaded disk shelf supports up to nine SAP HANA hosts, whereas a fully loaded shelf supports up to 14 hosts in a single disk shelf. The SAP HANA hosts must be equally distributed between both storage controllers. The same applies to the internal disks of an AFF A700s system. The DS224C disk shelf must be connected using 12Gb SAS to support the number of SAP HANA hosts.

The 6Gb SAS disk shelf (DS2246) supports a maximum of four SAP HANA hosts. The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers.

The following table summarizes the supported number of SAP HANA hosts per disk shelf.

	<b>6Gb SAS shelves (DS2246) Fully loaded with 24 SSDs</b>	<b>12Gb SAS shelves (DS224C) Half loaded with 12 SSDs and ADPv2</b>	<b>12Gb SAS shelves (DS224C) Fully loaded with 24 SSDs and ADPv2</b>
Maximum number of SAP HANA hosts per disk shelf	4	9	14



This calculation is independent of the storage controller used. Adding more disk shelves do not increase the maximum amount of SAP HANA hosts a storage controller can support.

## NS224 NVMe shelf

One NVMe SSDs (data) supports up to 2/5 SAP HANA hosts depending on the used NVMe disks. The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers.

The same applies to the internal NVMe disks of AFF systems.



Adding more disk shelves does not increase the maximum amount of SAP HANA hosts a storage controller can support.

## Mixed workloads

SAP HANA and other application workloads running on the same storage controller or in the same storage aggregate are supported. However, it is a NetApp best practice to separate SAP HANA workloads from all other application workloads.

You might decide to deploy SAP HANA workloads and other application workloads on either the same storage controller or the same aggregate. If so, you must make sure that adequate performance is available for SAP HANA within the mixed workload environment. NetApp also recommends that you use quality of service (QoS) parameters to regulate the effect these other applications could have on SAP HANA applications and to guarantee throughput for SAP HANA applications.

The SAP performance test tool must be used to check if additional SAP HANA hosts can be run on an existing storage controller that is already in use for other workloads. SAP application servers can be safely placed on

the same storage controller and/or aggregate as the SAP HANA databases.

### Capacity considerations

A detailed description of the capacity requirements for SAP HANA is in the [SAP Note 1900823](#) white paper.



The capacity sizing of the overall SAP landscape with multiple SAP HANA systems must be determined by using SAP HANA storage sizing tools from NetApp. Contact NetApp or your NetApp partner sales representative to validate the storage sizing process for a properly sized storage environment.

### Configuring the performance test tool

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used. These parameters must also be set for the performance test tool from SAP when storage performance is being tested with the SAP performance test tool.

NetApp conducted performance tests to define the optimal values. The following table lists the parameters that must be set within the configuration file of the SAP performance test tool.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For more information about the configuration of the different SAP test tools, see [SAP note 1943937](#) for HWCCT (SAP HANA 1.0) and [SAP note 2493172](#) for HCMT/HCOT (SAP HANA 2.0).

The following example shows how variables can be set for the HCMT/HCOT execution plan.

```
...{
    "Comment": "Log Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "LogAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Data Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "DataAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Log Volume: Controls whether write requests can be
submitted asynchronously",
```

```

    "Name": "LogAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "DataAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
  },
  {
    "Comment": "Log Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
    "Name": "LogAsyncWriteSubmitBlocks",
    "Value": "all",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
    "Name": "DataAsyncWriteSubmitBlocks",
    "Value": "all",
    "Request": "false"
  },
  {
    "Comment": "Log Volume: Maximum number of parallel I/O requests
per completion queue",
    "Name": "LogExtMaxParallelIoRequests",
    "Value": "128",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Maximum number of parallel I/O requests
per completion queue",
    "Name": "DataExtMaxParallelIoRequests",
    "Value": "128",
    "Request": "false"
  },
  }, ...

```

These variables must be used for the test configuration. This is usually the case with the predefined execution plans SAP delivers with the HCMT/HCOT tool. The following example for a 4k log write test is from an execution plan.

```

...
{
  "ID": "D664D001-933D-41DE-A904F304AEB67906",
  "Note": "File System Write Test",
  "ExecutionVariants": [
    {
      "ScaleOut": {
        "Port": "${RemotePort}",
        "Hosts": "${Hosts}",
        "ConcurrentExecution": "${FSConcurrentExecution}"
      },
      "RepeatCount": "${TestRepeatCount}",
      "Description": "4K Block, Log Volume 5GB, Overwrite",
      "Hint": "Log",
      "InputVector": {
        "BlockSize": 4096,
        "DirectoryName": "${LogVolume}",
        "FileOverwrite": true,
        "FileSize": 5368709120,
        "RandomAccess": false,
        "RandomData": true,
        "AsyncReadSubmit": "${LogAsyncReadSubmit}",
        "AsyncWriteSubmitActive":
"${LogAsyncWriteSubmitActive}",
        "AsyncWriteSubmitBlocks":
"${LogAsyncWriteSubmitBlocks}",
        "ExtMaxParallelIoRequests":
"${LogExtMaxParallelIoRequests}",
        "ExtMaxSubmitBatchSize": "${LogExtMaxSubmitBatchSize}",
        "ExtMinSubmitBatchSize": "${LogExtMinSubmitBatchSize}",
        "ExtNumCompletionQueues":
"${LogExtNumCompletionQueues}",
        "ExtNumSubmitQueues": "${LogExtNumSubmitQueues}",
        "ExtSizeKernelIoQueue": "${ExtSizeKernelIoQueue}"
      }
    }, ...
  ]
}

```

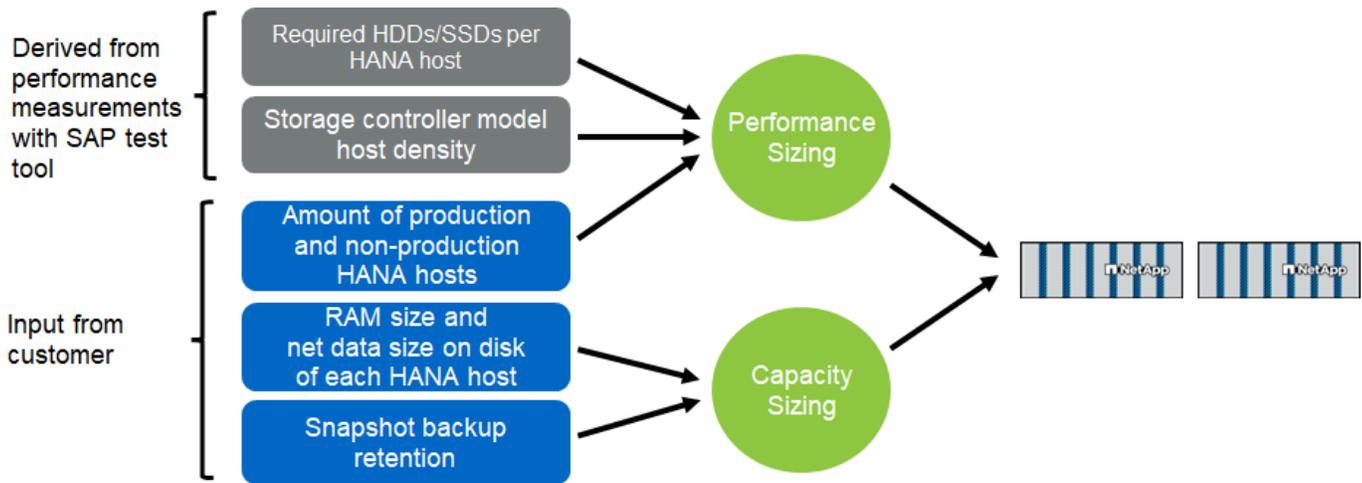
### Storage sizing process overview

The number of disks per HANA host and the SAP HANA host density for each storage model were determined with performance test tool.

The sizing process requires details such as the number of production and nonproduction SAP HANA hosts, the RAM size of each host, and backup retention of the storage-based Snapshot copies. The number of SAP HANA hosts determines the storage controller and the number of disks required.

The size of the RAM, net data size on the disk of each SAP HANA host, and the Snapshot copy backup retention period are used as inputs during capacity sizing.

The following figure summarizes the sizing process.



## Infrastructure setup and configuration

### Network setup

This section describes the dedicated storage network setup for SAP HANA hosts.

Use the following guidelines when configuring the network:

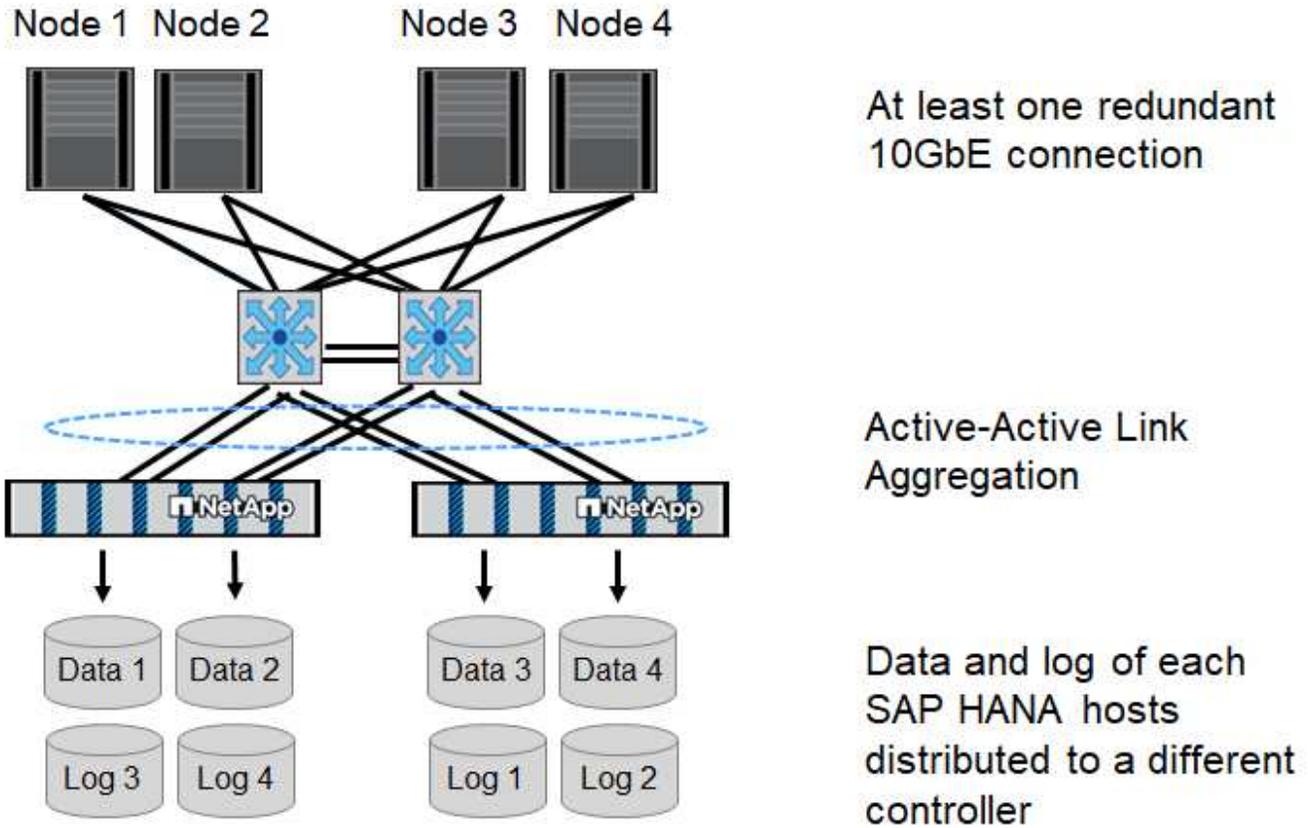
- A dedicated storage network must be used to connect the SAP HANA hosts to the storage controllers with a 10GbE or faster network.
- Use the same connection speed for storage controllers and SAP HANA hosts. If this is not possible, ensure that the network components between the storage controllers and the SAP HANA hosts are able to handle different speeds. For example, you must provide enough buffer space to allow speed negotiation at the NFS level between storage and hosts. Network components are usually switches, but other components within blade chassis, such as the back plane, must be considered as well.
- Disable flow control on all physical ports used for storage traffic on the storage network switch and host layer.
- Each SAP HANA host must have a redundant network connection with a minimum of 10Gb of bandwidth.
- Jumbo frames with a maximum transmission unit (MTU) size of 9,000 must be enabled on all network components between the SAP HANA hosts and the storage controllers.
- In a VMware setup, dedicated VMXNET3 network adapters must be assigned to each running virtual machine. Check the relevant papers mentioned in “Introduction” for further requirements.
- To avoid interference between each other, use separate network/IO paths for the log and data area.

The following figure shows an example with four SAP HANA hosts attached to a storage controller HA pair using a 10GbE network. Each SAP HANA host has an active-active connection to the redundant fabric.

At the storage layer, four active connections are configured to provide 10Gb throughput for each SAP HANA host.

At the storage layer, a broadcast domain with an MTU size of 9000 is configured, and all required physical interfaces are added to this broadcast domain. This approach automatically assigns these physical interfaces

to the same failover group. All logical interfaces (LIFs) that are assigned to these physical interfaces are added to this failover group.



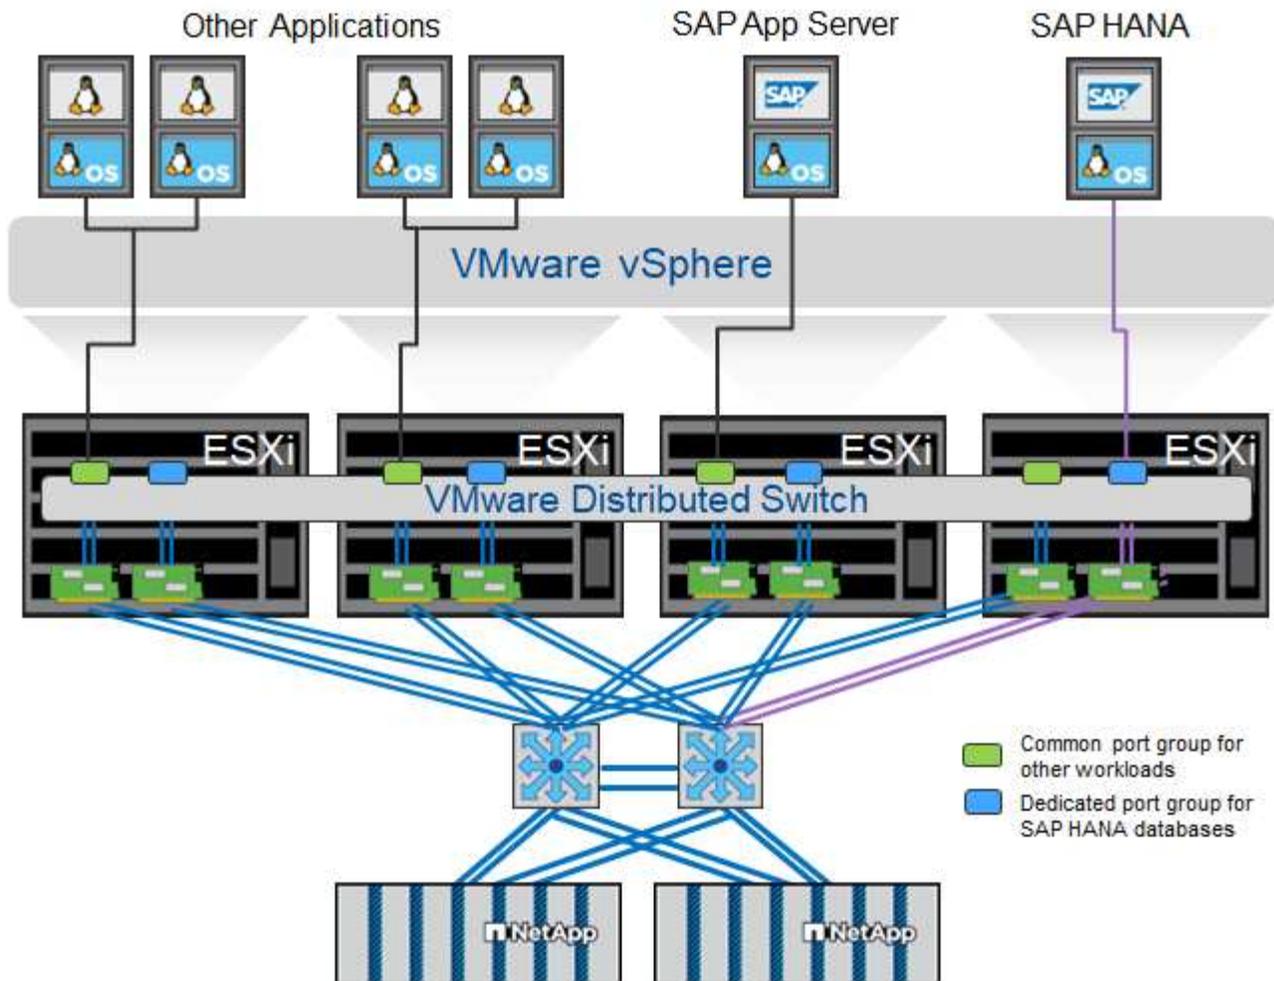
In general, it is recommended to use HA interface groups on the servers (bonds) and the storage systems (for example, Link Aggregation Control Protocol [LACP] and ifgroups). With HA interface groups, verify that the load is equally distributed between all interfaces within the group. The load distribution depends on the functionality of the network switch infrastructure.



Depending on the number of SAP HANA hosts and the connection speed used, different numbers of active physical ports are needed. For details, see the section "[LIF configuration](#)".

### VMware-specific network setup

Proper network design and configuration are crucial because all data for SAP HANA instances, including performance-critical data and log volumes for the database, is provided through NFS in this solution. A dedicated storage network is used to separate the NFS traffic from communication and user access traffic between SAP HANA nodes. Each SAP HANA node requires a redundant dedicated network connection with a minimum of 10Gb of bandwidth. Higher bandwidth is also supported. This network must extend end to end from the storage layer through network switching and computing up to the guest operating system hosted on VMware vSphere. In addition to the physical switching infrastructure, a VMware distributed switch (vDS) is used to provide adequate performance and manageability of network traffic at the hypervisor layer.



As shown in the preceding figure, each SAP HANA node uses a dedicated port group on the VMware distributed switch. This port group allows for enhanced quality of service (QoS) and dedicated assignment of physical network interface cards (NICs) on the ESX hosts. To use dedicated physical NICs while preserving HA capabilities in the event of NIC failure, the dedicated physical NIC is configured as an active uplink. Additional NICs are configured as standby uplinks in the teaming and failover settings of the SAP HANA port group. In addition, jumbo frames (MTU 9,000) must be enabled end to end on physical and virtual switches. In addition, turn off flow control on all ethernet ports used for storage traffic on servers, switches, and storage systems. The following figure shows an example of such a configuration.



LRO (large receive offload) must be turned off for interfaces used for NFS traffic. For all other network configuration guidelines, see the respective VMware best practices guides for SAP HANA.

t003-HANA-HV1 - Edit Settings

- General
- Advanced
- Security
- Traffic shaping
- VLAN
- Teaming and failover**
- Monitoring
- Traffic filtering and marking
- Miscellaneous

Load balancing:

Network failure detection:

Notify switches:

Failback:

Failover order

↑ ↓

**Active uplinks**

-  dvUplink2

**Standby uplinks**

-  dvUplink1

**Unused uplinks**

### Time synchronization

You must synchronize the time between the storage controllers and the SAP HANA database hosts. To do so, set the same time server for all storage controllers and all SAP HANA hosts.

### Storage controller setup

This section describes the configuration of the NetApp storage system. You must complete the primary installation and setup according to the corresponding ONTAP setup and configuration guides.

### Storage efficiency

Inline deduplication, cross-volume inline deduplication, inline compression, and inline compaction are supported with SAP HANA in an SSD configuration.

### NetApp FlexGroup Volumes

The usage of NetApp FlexGroup Volumes is not supported for SAP HANA. Due to the architecture of SAP HANA the usage of FlexGroup Volumes does not provide any benefit and may result in performance issues.

### NetApp Volume and Aggregate Encryption

The use of NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) are supported with SAP HANA.

### Quality of Service

QoS can be used to limit the storage throughput for specific SAP HANA systems or other applications on a shared-use controller. One use case would be to limit the throughput of development and test systems so that they cannot influence production systems in a mixed setup.

During the sizing process, you should determine the performance requirements of a nonproduction system. Development and test systems can be sized with lower performance values, typically in the range of 20% to 50% of a production- system KPI as defined by SAP.

Starting with ONTAP 9, QoS is configured on the storage volume level and uses maximum values for throughput (MBps) and the amount of I/O (IOPS).

Large write I/O has the biggest performance effect on the storage system. Therefore, the QoS throughput limit should be set to a percentage of the corresponding write SAP HANA storage performance KPI values in the data and log volumes.

### NetApp FabricPool

NetApp FabricPool technology must not be used for active primary file systems in SAP HANA systems. This includes the file systems for the data and log area as well as the `/hana/shared` file system. Doing so results in unpredictable performance, especially during the startup of an SAP HANA system.

Using the “snapshot-only” tiering policy is possible as well as using FabricPool in general at a backup target such as a NetApp SnapVault or SnapMirror destination.



Using FabricPool for tiering Snapshot copies at primary storage or using FabricPool at a backup target changes the required time for the restore and recovery of a database or other tasks such as creating system clones or repair systems. Take this into consideration for planning your overall lifecycle-management strategy and check to make sure that your SLAs are still being met while using this function.

FabricPool is a good option for moving log backups to another storage tier. Moving backups affects the time needed to recover an SAP HANA database. Therefore, the option “tiering-minimum-cooling-days” should be set to a value that places log backups, which are routinely needed for recovery, on the local fast storage tier.

### Storage configuration

The following overview summarizes the required storage configuration steps. Each step is covered in detail in the subsequent sections. In this section, we assume that the storage hardware is set up and that the ONTAP software is already installed. Also, the connections between the storage ports (10GbE or faster) and the network must already be in place.

1. Check the correct disk shelf configuration as described in "[Disk shelf connection](#)."
2. Create and configure the required aggregates as described in "[Aggregate configuration](#)."
3. Create a storage virtual machine (SVM) as described in "[SVM configuration](#)."
4. Create LIFs as described in "[LIF configuration](#)."
5. Create volumes within the aggregates as described in "[Volume configuration for SAP HANA single-host systems](#)" and "[Volume configuration for SAP HANA multiple-host systems](#)."
6. Set the required volume options as described in "[Volume options](#)."
7. Set the required options for NFSv3 as described in "[NFS configuration for NFSv3](#)" or for NFSv4 as described in "[NFS configuration for NFSv4](#)."
8. Mount the volumes to namespace and set export policies as described in "[Mount volumes to namespace and set export policies](#)."

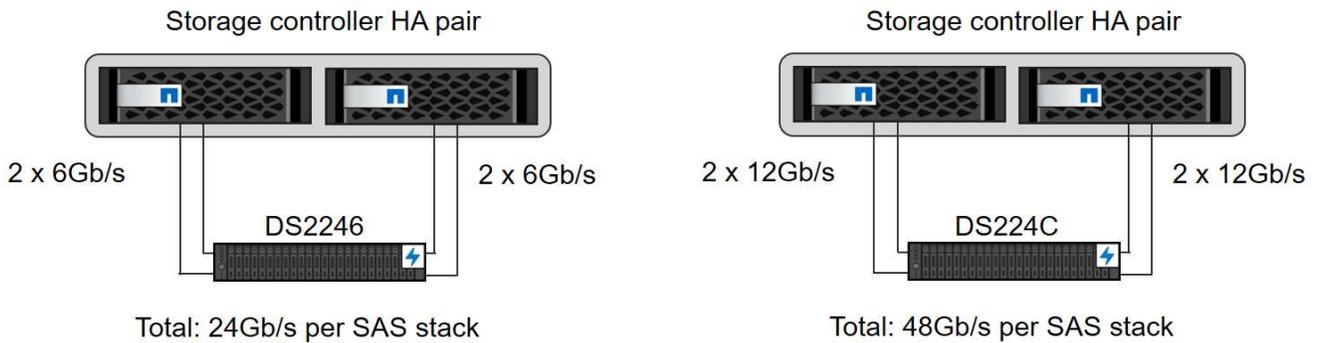
## Disk shelf connection

### SAS disk shelves

A maximum of one disk shelf can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair. ADPv2 is used with ONTAP 9 and the DS224C disk shelves.

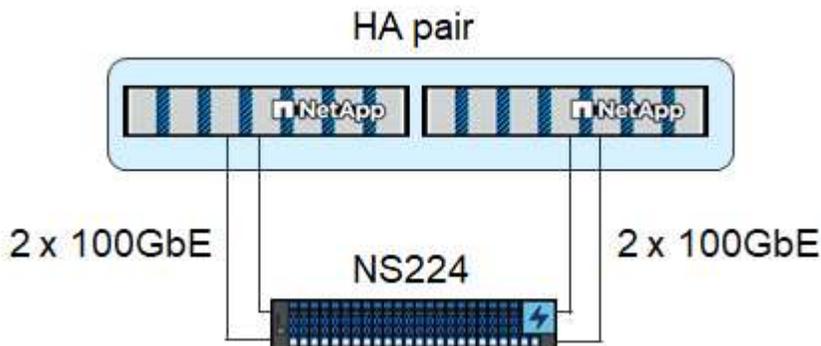


With the DS224C disk shelf, quad-path SAS cabling can also be used but is not required.



### NVMe (100GbE) disk shelves

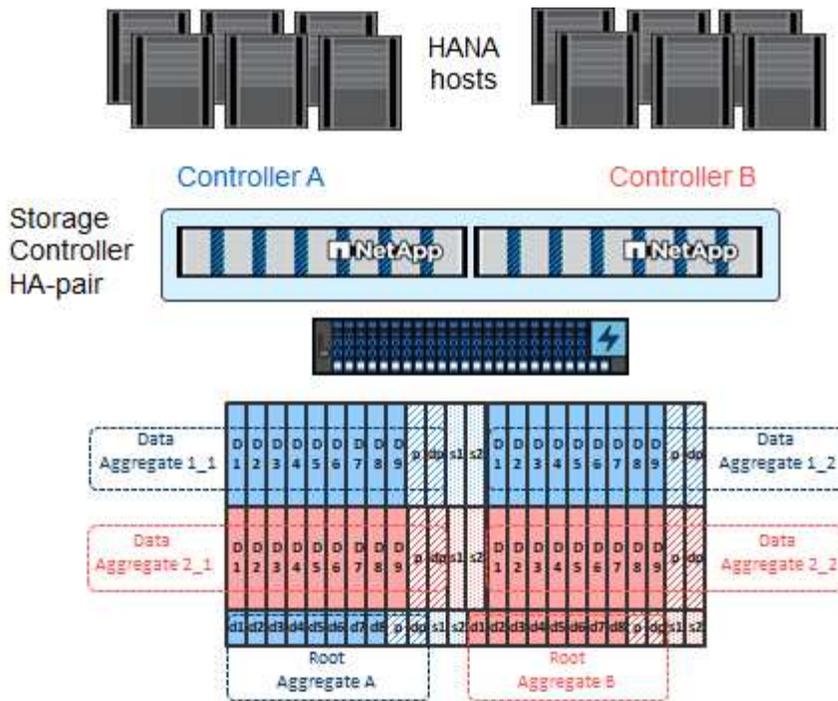
Each NS224 NVMe disk shelf is connected with two 100GbE ports per controller, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair. ADPv2, as described in the aggregate configuration chapter, is also used for the NS224 disk shelf.



### Aggregate configuration

In general, you must configure two aggregates per controller, independent of the disk shelf or drive technology (SAS SSDs or NVMe SSDs) that is used.

The following image shows a configuration of 12 SAP HANA hosts running on a 12Gb SAS shelf configured with ADPv2. Six SAP HANA hosts are attached to each storage controller. Four separate aggregates, two at each storage controller, are configured. Each aggregate is configured with 11 disks with nine data and two parity disk partitions. For each controller, two spare partitions are available.



### SVM configuration

Multiple SAP landscapes with SAP HANA databases can use a single SVM. An SVM can also be assigned to each SAP landscape, if necessary, in case they are managed by different teams within a company.

If there is a QoS profile automatically created and assigned while creating a new SVM, remove this automatically created profile from the SVM to enable the required performance for SAP HANA:

```
vserver modify -vserver <svm-name> -qos-policy-group none
```

### LIF configuration

For SAP HANA production systems, you must use different LIFs to mount the data volume and the log volume from the SAP HANA host. Therefore at least two LIFs are required.

The data and log volume mounts of different SAP HANA hosts can share a physical storage network port by either using the same LIFs or by using individual LIFs for each mount.

The maximum amount of data and log volume mounts per physical interface are shown in the following table.

Ethernet port speed	10GbE	25GbE	40GbE	100GeE
Maximum number of log or data volume mounts per physical port	3	8	12	30



Sharing one LIF between different SAP HANA hosts might require a remount of data or log volumes to a different LIF. This change avoids performance penalties if a volume is moved to a different storage controller.

Development and test systems can use more data and volume mounts or LIFs on a physical network interface.

For production, development, and test systems, the `/hana/shared` file system can use the same LIF as the data or log volume.

### Volume configuration for SAP HANA single-host systems

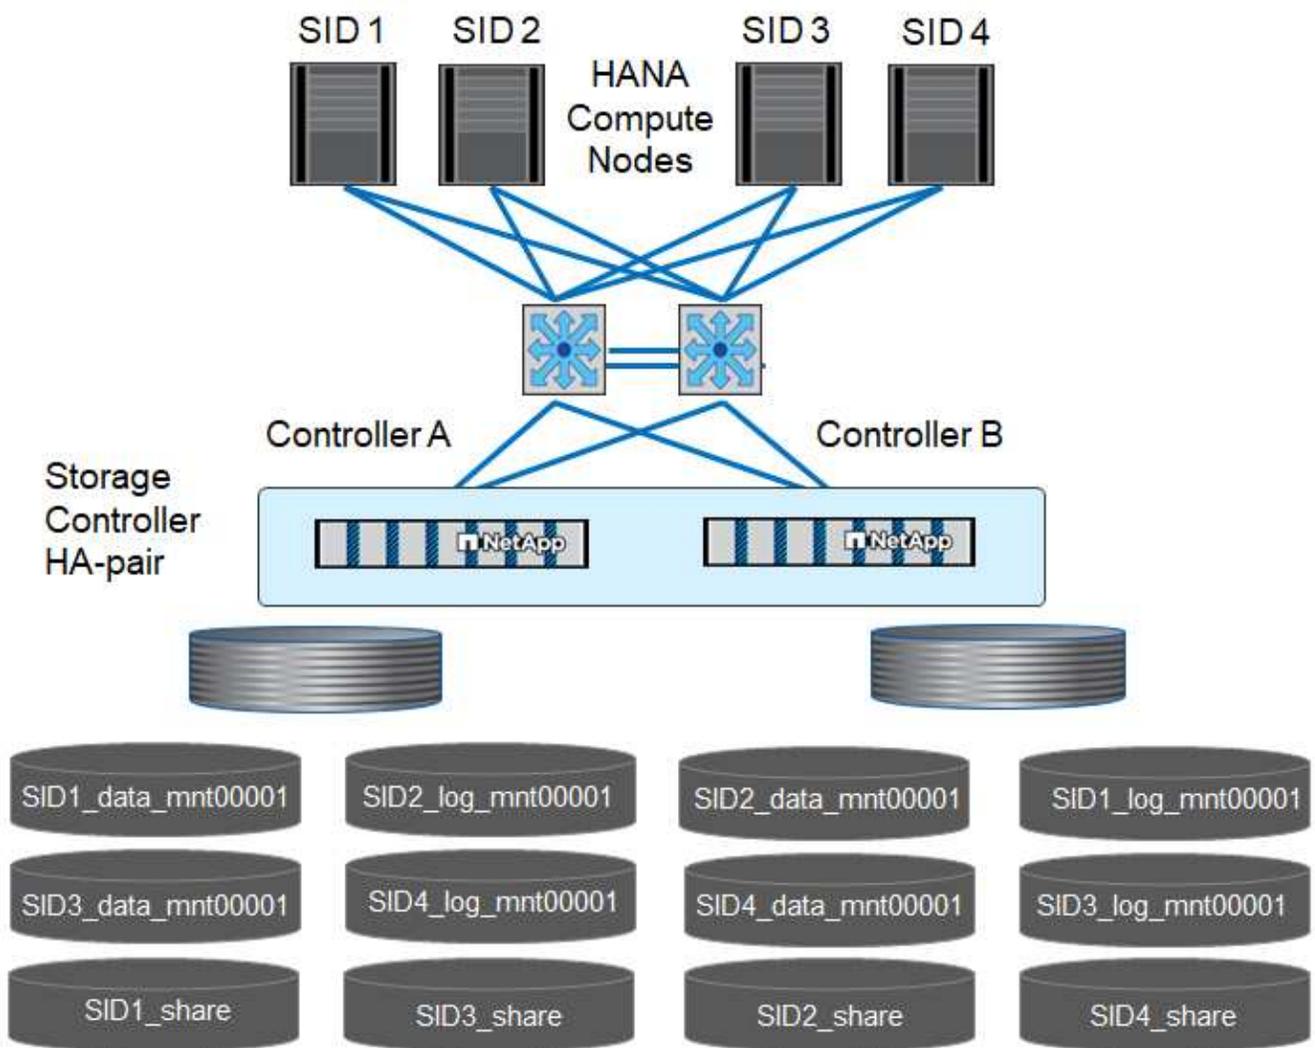
The following figure shows the volume configuration of four single-host SAP HANA systems. The data and log volumes of each SAP HANA system are distributed to different storage controllers. For example, volume `SID1_data_mnt00001` is configured on controller A, and volume `SID1_log_mnt00001` is configured on controller B.



If only one storage controller of an HA pair is used for the SAP HANA systems, data and log volumes can also be stored on the same storage controller.



If the data and log volumes are stored on the same controller, access from the server to the storage must be performed with two different LIFs: one LIF to access the data volume and the other to access the log volume.



For each SAP HANA host, a data volume, a log volume, and a volume for `/hana/shared` are configured. The following table shows an example configuration for single-host SAP HANA systems.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller b
Data, log, and shared volumes for system SID1	Data volume: SID1_data_mnt00001	Shared volume: SID1_shared	–	Log volume: SID1_log_mnt00001
Data, log, and shared volumes for system SID2	–	Log volume: SID2_log_mnt00001	Data volume: SID2_data_mnt00001	Shared volume: SID2_shared
Data, log, and shared volumes for system SID3	Shared volume: SID3_shared	Data volume: SID3_data_mnt00001	Log volume: SID3_log_mnt00001	–
Data, log, and shared volumes for system SID4	Log volume: SID4_log_mnt00001	–	Shared volume: SID4_shared	Data volume: SID4_data_mnt00001

The following table shows an example of the mount point configuration for a single-host system. To place the home directory of the `sidadm` user on the central storage, the `/usr/sap/SID` file system should be mounted from the `SID_shared` volume.

Junction path	Directory	Mount point at HANA host
SID_data_mnt00001		/hana/data/SID/mnt00001
SID_log_mnt00001		/hana/log/SID/mnt00001
SID_shared	usr-sap shared	/usr/sap/SID /hana/shared/

#### Volume configuration for SAP HANA multiple-host systems

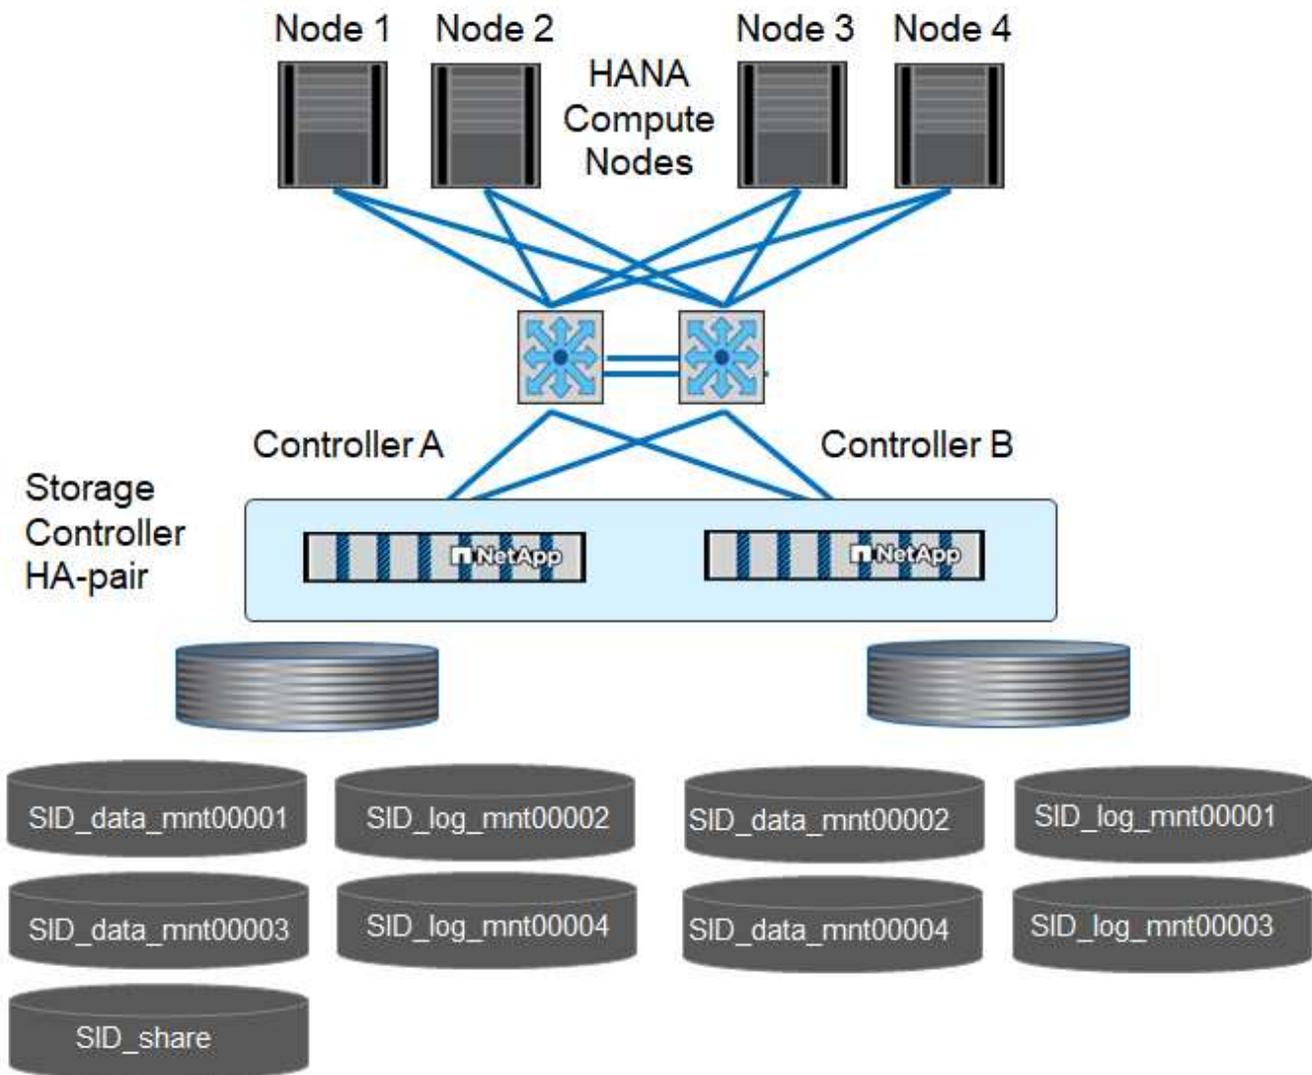
The following figure shows the volume configuration of a 4+1 SAP HANA system. The data and log volumes of each SAP HANA host are distributed to different storage controllers. For example, volume `SID1_data1_mnt00001` is configured on controller A, and volume `SID1_log1_mnt00001` is configured on controller B.



If only one storage controller of an HA pair is used for the SAP HANA system, the data and log volumes can also be stored on the same storage controller.



If the data and log volumes are stored on the same controller, access from the server to the storage must be performed with two different LIFs: one LIF to access the data volume and one to access the log volume.



For each SAP HANA host, a data volume and a log volume are created. The `/hana/shared` volume is used by all hosts of the SAP HANA system. The following table shows an example configuration for a multiple-host SAP HANA system with four active hosts.

Purpose	Aggregate 1 at controller A	Aggregate 2 at controller A	Aggregate 1 at controller B	Aggregate 2 at controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	–	Log volume: SID_log_mnt00001	–
Data and log volumes for node 2	Log volume: SID_log_mnt00002	–	Data volume: SID_data_mnt00002	–
Data and log volumes for node 3	–	Data volume: SID_data_mnt00003	–	Log volume: SID_log_mnt00003
Data and log volumes for node 4	–	Log volume: SID_log_mnt00004	–	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared			

The following table shows the configuration and the mount points of a multiple-host system with four active

SAP HANA hosts. To place the home directories of the `sidadm` user of each host on the central storage, the `/usr/sap/SID` file systems are mounted from the `SID_shared` volume.

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00001	–	/hana/data/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001	–	/hana/log/SID/mnt00001	Mounted at all hosts
SID_data_mnt00002	–	/hana/data/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002	–	/hana/log/SID/mnt00002	Mounted at all hosts
SID_data_mnt00003	–	/hana/data/SID/mnt00003	Mounted at all hosts
SID_log_mnt00003	–	/hana/log/SID/mnt00003	Mounted at all hosts
SID_data_mnt00004	–	/hana/data/SID/mnt00004	Mounted at all hosts
SID_log_mnt00004	–	/hana/log/SID/mnt00004	Mounted at all hosts
SID_shared	shared	/hana/shared/SID	Mounted at all hosts
SID_shared	usr-sap-host1	/usr/sap/SID	Mounted at host 1
SID_shared	usr-sap-host2	/usr/sap/SID	Mounted at host 2
SID_shared	usr-sap-host3	/usr/sap/SID	Mounted at host 3
SID_shared	usr-sap-host4	/usr/sap/SID	Mounted at host 4
SID_shared	usr-sap-host5	/usr/sap/SID	Mounted at host 5

### Volume options

You must verify and set the volume options listed in the following table on all SVMs. For some of the commands, you must switch to the advanced privilege mode within ONTAP.

Action	Command
Disable visibility of Snapshot directory	<code>vol modify -vserver &lt;vserver-name&gt; -volume &lt;volname&gt; -snapdir-access false</code>
Disable automatic Snapshot copies	<code>vol modify -vserver &lt;vserver-name&gt; -volume &lt;volname&gt; -snapshot-policy none</code>
Disable access time update, except of the <code>SID_shared</code> volume	<code>set advanced</code> <code>vol modify -vserver &lt;vserver-name&gt; -volume &lt;volname&gt; -atime-update false</code> <code>set admin</code>

### NFS configuration for NFSv3

The NFS options listed in the following table must be verified and set on all storage controllers. For some of the commands shown in this table, you must switch to the advanced privilege mode.

Action	Command
Enable NFSv3	<code>nfs modify -vserver &lt;vserver-name&gt; v3.0 enabled</code>

Action	Command
Set NFS TCP maximum transfer size to 1MB	set advanced nfs modify -vserver <vserver_name> -tcp-max-xfer -size 1048576 set admin



In shared environments with different workloads set the max NFS TCP transfer size to 262144

#### NFS configuration for NFSv4

The NFS options listed in the following table must be verified and set on all SVMs.

For some of the commands in this table, you must switch to the advanced privilege mode.

Action	Command
Enable NFSv4	nfs modify -vserver <vserver-name> -v4.1 enabled
Set NFS TCP maximum transfer size to 1MB	set advanced nfs modify -vserver <vserver_name> -tcp-max-xfer -size 1048576 set admin
Disable NFSv4 access control lists (ACLs)	nfs modify -vserver <vserver_name> -v4.1-acl disabled
Set NFSv4 domain ID	nfs modify -vserver <vserver_name> -v4-id-domain <domain-name>
Disable NFSv4 read delegation	nfs modify -vserver <vserver_name> -v4.1-read -delegation disabled
Disable NFSv4 write delegation	nfs modify -vserver <vserver_name> -v4.1-write -delegation disabled
Disable NFSv4 numeric ids	nfs modify -vserver <vserver_name> -v4-numeric-ids disabled
Change amount of NFSv4.x session slots optional	set advanced nfs modify -vserver hana -v4.x-session-num-slots <value> set admin



In shared environments with different workloads set the max NFS TCP transfer size to 262144



Please note that disabling numeric ids requires user management, as described in the section [“SAP HANA installation preparations for NFSv4.”](#)



The NFSv4 domain ID must be set to the same value on all Linux servers ( `/etc/idmapd.conf`) and SVMs, as described in the section [“SAP HANA installation preparations for NFSv4.”](#)



pNFS can be enabled and used.

If SAP HANA multiple-host systems with host auto-failover are being used, the failover parameters need to be adjusted within `nameserver.ini` as shown in the following table. Keep the default retry interval of 10 seconds within these sections..

Section within <code>nameserver.ini</code>	Parameter	Value
failover	normal_retries	9
distributed_watchdog	deactivation_retries	11
distributed_watchdog	takeover_retries	9

### Mount volumes to namespace and set export policies

When a volume is created, the volume must be mounted to the namespace. In this document, we assume that the junction path name is the same as the volume name. By default, the volume is exported with the default policy. The export policy can be adapted if required.

### Host setup

All the host-setup steps described in this section are valid for both SAP HANA environments on physical servers and for SAP HANA running on VMware vSphere.

### Configuration parameter for SUSE Linux Enterprise Server

Additional kernel and configuration parameters at each SAP HANA host must be adjusted for the workload generated by SAP HANA.

### SUSE Linux Enterprise Server 12 and 15

Starting with SUSE Linux Enterprise Server 12 SP1, the kernel parameter must be set in a configuration file in the `/etc/sysctl.d` directory. For example, you must create a configuration file with the name `91-NetApp-HANA.conf`.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
sunrpc.tcp_max_slot_table_entries = 128
```



Saptune, included in SLES for SAP OS versions, can be used to set these values. For more information, see [SAP Note 3024346](#) (requires SAP login).

## Configuration parameters for Red Hat Enterprise Linux 7.2 or later

You must adjust additional kernel and configuration parameters at each SAP HANA host for the workload generated by SAP HANA.

Starting with Red Hat Enterprise Linux 7.2, you must set the kernel parameters in a configuration file in the `/etc/sysctl.d` directory. For example, you must create a configuration file with the name `91-NetApp-HANA.conf`.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
sunrpc.tcp_max_slot_table_entries = 128
```



Since Red Hat Enterprise Linux version 8.6, the settings can be also applied by using the RHEL System Roles for SAP (Ansible). See [SAP Note 3024346](#) (requires SAP login).

## Create subdirectories in `/hana/shared` volume



The following examples show an SAP HANA database with `SID=NF2`.

To create the required subdirectories, take one of the following actions:

- For a single-host system, mount the `/hana/shared` volume and create the `shared` and `usr-sap` subdirectories.

```
sapcc-hana-tst-06:/mnt # mount <storage-hostname>:/NF2_shared /mnt/tmp
sapcc-hana-tst-06:/mnt # cd /mnt/tmp
sapcc-hana-tst-06:/mnt/tmp # mkdir shared
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap
sapcc-hana-tst-06:/mnt/tmp # cd ..
sapcc-hana-tst-06:/mnt # umount /mnt/tmp
```

- For a multiple-host system, mount the `/hana/shared` volume and create the `shared` and the `usr-sap` subdirectories for each host.

The example commands show a 2+1 multiple-host HANA system.

```
sapcc-hana-tst-06:/mnt # mount <storage-hostname>:/NF2_shared /mnt/tmp
sapcc-hana-tst-06:/mnt # cd /mnt/tmp
sapcc-hana-tst-06:/mnt/tmp # mkdir shared
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host1
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host2
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host3
sapcc-hana-tst-06:/mnt/tmp # cd ..
sapcc-hana-tst-06:/mnt # umount /mnt/tmp
```

### Create mount points



The following examples show an SAP HANA database with SID=NF2.

To create the required mount point directories, take one of the following actions:

- For a single-host system, create mount points and set the permissions on the database host.

```
sapcc-hana-tst-06:/ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-06:/ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-06:/ # mkdir -p /hana/shared
sapcc-hana-tst-06:/ # mkdir -p /usr/sap/NF2

sapcc-hana-tst-06:/ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/shared
sapcc-hana-tst-06:/ # chmod -R 777 /usr/sap/NF2
```

- For a multiple-host system, create mount points and set the permissions on all worker and standby hosts. The following example commands are for a 2+1 multiple-host HANA system.
  - First worker host:

```

sapcc-hana-tst-06:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-06:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-06:~ # mkdir -p /hana/shared
sapcc-hana-tst-06:~ # mkdir -p /usr/sap/NF2

sapcc-hana-tst-06:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-06:~ # chmod -R 777 /usr/sap/NF2

```

- Second worker host:

```

sapcc-hana-tst-07:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-07:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-07:~ # mkdir -p /hana/shared
sapcc-hana-tst-07:~ # mkdir -p /usr/sap/NF2

sapcc-hana-tst-07:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-07:~ # chmod -R 777 /usr/sap/NF2

```

- Standby host:

```

sapcc-hana-tst-08:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-08:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-08:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-08:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-08:~ # mkdir -p /hana/shared
sapcc-hana-tst-08:~ # mkdir -p /usr/sap/NF2

sapcc-hana-tst-08:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-08:~ # chmod -R 777 /usr/sap/NF2

```

## Mount file systems

Different mount options must be used depending on the NFS version and ONTAP release. The following file systems must be mounted to the hosts:

- /hana/data/SID/mnt0000\*
- /hana/log/SID/mnt0000\*
- /hana/shared
- /usr/sap/SID

The following table shows the NFS versions that you must use for the different files systems for single-host and multiple-host SAP HANA databases.

File systems	SAP HANA single host	SAP HANA multiple hosts
/hana/data/SID/mnt0000*	NFSv3 or NFSv4	NFSv4
/hana/log/SID/mnt0000*	NFSv3 or NFSv4	NFSv4
/hana/shared	NFSv3 or NFSv4	NFSv3 or NFSv4
/usr/sap/SID	NFSv3 or NFSv4	NFSv3 or NFSv4

The following table shows the mount options for the various NFS versions and ONTAP releases. The common parameters are independent of the NFS and ONTAP versions.



SAP LaMa requires the /usr/sap/SID directory to be local. Therefore, don't mount an NFS volume for /usr/sap/SID if you are using SAP LaMa.

For NFSv3, you must switch off NFS locking to avoid NFS lock cleanup operations in case of a software or server failure.

With ONTAP 9, the NFS transfer size can be configured up to 1MB. Specifically, with 40GbE or faster connections to the storage system, you must set the transfer size to 1MB to achieve the expected throughput values.

Common parameter	NFSv3	NFSv4	NFS transfer size with ONTAP 9	NFS transfer size with ONTAP 8
rw, bg, hard, timeo=600, noatime	nfsvers=3,nolock	nfsvers=4.1,lock	rsize=1048576,wsize=262144	rsize=65536,wsize=65536



To improve read performance with NFSv3, NetApp recommends that you use the `nconnect=n` mount option, which is available with SUSE Linux Enterprise Server 12 SP4 or later and RedHat Enterprise Linux (RHEL) 8.3 or later.



Performance tests showed that `nconnect=4` provides good read results for the data volumes. Log writes might benefit from a lower number of sessions, such as `nconnect=2`. Shared volumes may benefit as well from using the 'nconnect' option. Be aware that the first mount from an NFS server (IP address) defines the amount of sessions being used. Further mounts to the same IP address do not change this even if a different value is used for `nconnect`.



Starting with ONTAP 9.8 and SUSE SLES15SP2 or RedHat RHEL 8.4 or higher, NetApp supports the `nconnect` option also for NFSv4.1. For additional information, check the Linux vendor documentation.

If `nconnect` is being used with NFSv4.x the amount of NFSv4.x session slots should be adjusted according to the following rule:

Amount of session slots equals `<nconnect value> x 64`.



At the host this will be adjusted by

```
echo options nfs max_session_slots=<calculated value> >
/etc/modprobe.d/nfsclient.conf
```

followed by a reboot. The server side value must be adjusted as well, set the number of session slots as described in [NFS configuration for NFSv4](#).

The following example shows a single host SAP HANA database with `SID=NF2` using NFSv3 and an NFS transfer size of 1MB for reads and 256k for writes. To mount the file systems during system boot with the `/etc/fstab` configuration file, complete the following steps:

1. Add the required file systems to the `/etc/fstab` configuration file.

```
sapcc-hana-tst-06:/ # cat /etc/fstab
<storage-vif-data01>:/NF2_data_mnt00001 /hana/data/NF2/mnt00001 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-log01>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=2,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data01>:/NF2_shared/usr-sap /usr/sap/NF2 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data01>:/NF2_shared/shared /hana/shared nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
```

2. Run `mount -a` to mount the file systems on all hosts.

The next example shows a multiple-host SAP HANA database with `SID=NF2` using NFSv4.1 for data and log file systems and NFSv3 for the `/hana/shared` and `/usr/sap/NF2` file systems. An NFS transfer size of 1MB for reads and 256k for writes is used.

1. Add the required file systems to the `/etc/fstab` configuration file on all hosts.



The `/usr/sap/NF2` file system is different for each database host. The following example shows `/NF2_shared/usr-sap-host1`.

```

stlrx300s8-5:/ # cat /etc/fstab
<storage-vif-data01>:/NF2_data_mnt00001 /hana/data/NF2/mnt00001 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,no
oatime,lock 0 0
<storage-vif-data02>:/NF2_data_mnt00002 /hana/data/NF2/mnt00002 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-log01>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=2,rsize=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-log02>:/NF2_log_mnt00002 /hana/log/NF2/mnt00002 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=2,rsize=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-data02>:/NF2_shared/usr-sap-host1 /usr/sap/NF2 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data02>:/NF2_shared/shared /hana/shared nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0

```

2. Run `mount -a` to mount the file systems on all hosts.

## SAP HANA installation preparations for NFSv4

NFS version 4 and higher requires user authentication. This authentication can be accomplished by using a central user management tool such as a Lightweight Directory Access Protocol (LDAP) server or with local user accounts. The following sections describe how to configure local user accounts.

The administration users `<sid>adm,<sid>crypt` and the `sapsys` group must be created manually on the SAP HANA hosts and the storage controllers before the installation of the SAP HANA software begins.

### SAP HANA hosts

If it doesn't exist, the `sapsys` group must be created on the SAP HANA host. A unique group ID must be chosen that does not conflict with the existing group IDs on the storage controllers.

The users `<sid>adm` and `<sid>crypt` are created on the SAP HANA host. Unique IDs must be chosen that does not conflict with existing user IDs on the storage controllers.

For a multiple-host SAP HANA system, the user and group IDs must be the same on all SAP HANA hosts. The group and user are created on the other SAP HANA hosts by copying the affected lines in `/etc/group` and `/etc/passwd` from the source system to all other SAP HANA hosts.

For a multiple-host SAP HANA system, the user and group ID must be the same on all SAP HANA hosts. The group and user are created on the other SAP HANA hosts by copying the affected lines in `/etc/group` and `/etc/passwd` from the source system to all other SAP HANA hosts.



The NFSv4 domain must be set to the same value on all Linux servers and SVMs. Set the domain parameter “Domain = <domain\_name>” in file `/etc/idmapd.conf` for the Linux hosts.

Enable and start the NFS idmapd service:

```
systemctl enable nfs-idmapd.service
systemctl start nfs-idmapd.service
```



The latest Linux kernels do not require this step. You can safely ignore warning messages.

### Storage controllers

The user ID and group ID must be the same on the SAP HANA hosts and the storage controllers. The group and user are created by entering the following commands on the storage cluster:

```
vserver services unix-group create -vserver <vserver> -name <group name>
-id <group id>
vserver services unix-user create -vserver <vserver> -user <user name> -id
<user-id> -primary-gid <group id>
```

Additionally, set the group ID of the UNIX user root of the SVM to 0.

```
vserver services unix-user modify -vserver <vserver> -user root -primary
-gid 0
```

### I/O stack configuration for SAP HANA

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage systems used.

NetApp conducted performance tests to define the ideal values. The following table lists the optimal values inferred from the performance tests.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For SAP HANA 1.0 versions up to SPS12, these parameters can be set during the installation of the SAP HANA database, as described in SAP note [2267798: Configuration of the SAP HANA Database During Installation Using hdbparam](#).

Alternatively, the parameters can be set after SAP HANA database installation by using the `hdbparam` framework.

```
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset
fileio.max_parallel_io_requests=128
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset
fileio.async_write_submit_active=on
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset
fileio.async_read_submit=on
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset
fileio.async_write_submit_blocks=all
```

Starting with SAP HANA 2.0, `hdbparam` was deprecated and the parameters were moved to `global.ini`. The parameters can be set using SQL commands or SAP HANA Studio. For more details, see [SAP note 2399079: Elimination of hdbparam in HANA 2](#). The parameters can also be set within the `global.ini` as shown below:

```
nf2adm@stlrx300s8-6: /usr/sap/NF2/SYS/global/hdb/custom/config> cat
global.ini
...
[fileio]
async_read_submit = on
async_write_submit_active = on
max_parallel_io_requests = 128
async_write_submit_blocks = all
...
```

As of SAP HANA 2.0 SPS5, you can use the `setParameter.py` script to set the correct parameters:

```
nf2adm@sapcc-hana-tst-03:/usr/sap/NF2/HDB00/exe/python_support>
python setParameter.py
-set=SYSTEM/global.ini/fileio/max_parallel_io_requests=128
python setParameter.py -set=SYSTEM/global.ini/fileio/async_read_submit=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_active=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_blocks=all
```

### SAP HANA data volume size

As the default, SAP HANA uses only one data volume per SAP HANA service. Due to the maximum file size limitation of the file system, NetApp recommends limiting the maximum data volume size.

To do so automatically, set the following parameter in `global.ini` in the section `[persistence]`:

```
datavolume_stripping = true
datavolume_stripping_size_gb = 8000
```

This creates a new data volume after the 8,000GB limit is reached. [SAP note 240005 question 15](#) provides more information.

## SAP HANA software installation

This section describes how to configure a system for the installation of SAP HANA software on single-host and multiple-host systems.

### Install on a single-host system

SAP HANA software installation does not require any additional preparation for a single-host system.

### Install on a multiple-host system

To install SAP HANA on a multiple-host system, complete the following steps:

1. Using the SAP `hdbclm` installation tool, start the installation by running the following command at one of the worker hosts. Use the `addhosts` option to add the second worker (`sapcc-hana-tst-03`) and the standby host (`sapcc-hana-tst-04`).

```
apcc-hana-tst-02:/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_LCM_LINUX_X86_64 #
./hdbclm --action=install --addhosts=sapcc-hana-tst-03:role=worker,sapcc-
-hana-tst-04:role=standby

SAP HANA Lifecycle Management - SAP HANA Database 2.00.073.00.1695288802
*****

Scanning software locations...
Detected components:
    SAP HANA AFL (incl.PAL,BFL,OFL) (2.00.073.0000.1695321500) in
/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
    SAP HANA Database (2.00.073.00.1695288802) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
    SAP HANA Database Client (2.18.24.1695756995) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_CLIENT_LINUX_X86_64/SAP_HANA_CLIENT/client
    SAP HANA Studio (2.3.75.000000) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
```

73/DATA\_UNITS/HDB\_STUDIO\_LINUX\_X86\_64/studio  
 SAP HANA Local Secure Store (2.11.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/HANA\_LSS\_24\_LINUX\_X86\_64/packages  
 SAP HANA XS Advanced Runtime (1.1.3.230717145654) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_RT\_10\_LINUX\_X86\_64/packages  
 SAP HANA EML AFL (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/HDB\_EML\_AFL\_10\_LINUX\_X86\_64/packages  
 SAP HANA EPM-MDS (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/SAP\_HANA\_EPM-MDS\_10/packages  
 Automated Predictive Library (4.203.2321.0.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/PAAPL4\_H20\_LINUX\_X86\_64/apl-4.203.2321.0-hana2sp03-linux\_x64/installer/packages  
 GUI for HALM for XSA (including product installer) Version 1 (1.015.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACALMPIUI15\_0.zip  
 XSAC FILEPROCESSOR 1.0 (1.000.102) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACFILEPROC00\_102.zip  
 SAP HANA tools for accessing catalog content, data preview, SQL console, etc. (2.015.230503) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSAC\_HRTT\_20/XSACHRTT15\_230503.zip  
 Develop and run portal services for customer applications on XSA (2.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACPORTALSERV07\_0.zip  
 The SAP Web IDE for HANA 2.0 (4.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSAC\_SAP\_WEB\_IDE\_20/XSACSAPWEBIDE07\_0.zip  
 XS JOB SCHEDULER 1.0 (1.007.22) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACSERVICES07\_22.zip  
 SAPUI5 FESV6 XSA 1 - SAPUI5 1.71 (1.071.52) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV671\_52.zip  
 SAPUI5 FESV9 XSA 1 - SAPUI5 1.108 (1.108.5) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV9108\_5.zip  
 SAPUI5 SERVICE BROKER XSA 1 - SAPUI5 Service Broker 1.0 (1.000.4) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5SB00\_4.zip  
 XSA Cockpit 1 (1.001.37) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACXSACOCKPIT01\_37.zip

```
SAP HANA Database version '2.00.073.00.1695288802' will be installed.
```

```
Select additional components for installation:
```

Index	Components	Description
1	all	All components
2	server	No additional components
3	client	Install SAP HANA Database Client version 2.18.24.1695756995
4	lss	Install SAP HANA Local Secure Store version 2.11.0
5	studio	Install SAP HANA Studio version 2.3.75.000000
6	xs	Install SAP HANA XS Advanced Runtime version 1.1.3.230717145654
7	afl	Install SAP HANA AFL (incl.PAL,BFL,OFL) version 2.00.073.0000.1695321500
8	eml	Install SAP HANA EML AFL version 2.00.073.0000.1695321500
9	epmmds	Install SAP HANA EPM-MDS version 2.00.073.0000.1695321500
10	sap_afl_sdk_apl	Install Automated Predictive Library version 4.203.2321.0.0

```
Enter comma-separated list of the selected indices [3,4]: 2,3
```

2. Verify that the installation tool installed all selected components at all worker and standby hosts.

### Adding additional data volume partitions

Starting with SAP HANA 2.0 SPS4, additional data volume partitions can be configured. This allows you to configure two or more volumes for the data volume of an SAP HANA tenant database and scale beyond the size and performance limits of a single volume.



Using two or more individual volumes for the data volume is available for SAP HANA single-host and SAP HANA multiple-host systems. You can add additional data volume partitions at any time.

### Enabling additional data volume partitions

To enable additional data volume partitions, add the following entry within `global.ini` by using SAP HANA Studio or Cockpit in the SYSTEMDB configuration.

```
[customizable_functionalities]
persistence_datavolume_partition_multipath = true
```



Adding the parameter manually to the `global.ini` file requires the restart of the database.

### Volume configuration for single-host SAP HANA systems

The layout of volumes for a single-host SAP HANA system with multiple partitions is like the layout for a system with one data volume partition but with an additional data volume stored on a different aggregate as the log volume and the other data volume. The following table shows an example configuration of an SAP HANA single-host system with two data volume partitions.

Aggregate 1 at controller A	Aggregate 2 at controller A	Aggregate 1 at controller B	Aggregate 2 at controller b
Data volume: SID_data_mnt00001	Shared volume: SID_shared	Data volume: SID_data2_mnt00001	Log volume: SID_log_mnt00001

The following table shows an example of the mount point configuration for a single-host system with two data volume partitions.

Junction path	Directory	Mount point at HANA host
SID_data_mnt00001	–	/hana/data/SID/mnt00001
SID_data2_mnt00001	–	/hana/data2/SID/mnt00001
SID_log_mnt00001	–	/hana/log/SID/mnt00001
SID_shared	usr-sap shared	/usr/sap/SID /hana/shared

You can create the new data volume and mount it to the namespace using either NetApp ONTAP System Manager or the ONTAP CLI.

### Volume configuration for multiple-host SAP HANA systems

The layout of volumes is like the layout for a multiple- host SAP HANA system with one data volume partition but with an additional data volume stored on a different aggregate as log volume and the other data volume. The following table shows an example configuration of an SAP HANA multiple-host system with two data volume partitions.

Purpose	Aggregate 1 at controller A	Aggregate 2 at controller A	Aggregate 1 at controller B	Aggregate 2 at controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	–	Log volume: SID_log_mnt00001	Data2 volume: SID_data2_mnt00001
Data and log volumes for node 2	Log volume: SID_log_mnt00002	Data2 volume: SID_data2_mnt00002	Data volume: SID_data_mnt00002	–

Purpose	Aggregate 1 at controller A	Aggregate 2 at controller A	Aggregate 1 at controller B	Aggregate 2 at controller B
Data and log volumes for node 3	–	Data volume: SID_data_mnt00003	Data2 volume: SID_data2_mnt00003	Log volume: SID_log_mnt00003
Data and log volumes for node 4	Data2 volume: SID_data2_mnt00004	Log volume: SID_log_mnt00004	–	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared	–	–	–

The following table shows an example of the mount point configuration for a single-host system with two data volume partitions.

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00001	–	/hana/data/SID/mnt00001	Mounted at all hosts
SID_data2_mnt00001	–	/hana/data2/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001	–	/hana/log/SID/mnt00001	Mounted at all hosts
SID_data_mnt00002	–	/hana/data/SID/mnt00002	Mounted at all hosts
SID_data2_mnt00002	–	/hana/data2/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002	–	/hana/log/SID/mnt00002	Mounted at all hosts
SID_data_mnt00003	–	/hana/data/SID/mnt00003	Mounted at all hosts
SID_data2_mnt00003		/hana/data2/SID/mnt00003	Mounted at all hosts
SID_log_mnt00003		/hana/log/SID/mnt00003	Mounted at all hosts
SID_data_mnt00004		/hana/data/SID/mnt00004	Mounted at all hosts
SID_data2_mnt00004	–	/hana/data2/SID/mnt00004	Mounted at all hosts
SID_log_mnt00004	–	/hana/log/SID/mnt00004	Mounted at all hosts
SID_shared	shared	/hana/shared/SID	Mounted at all hosts
SID_shared	usr-sap-host1	/usr/sap/SID	Mounted at host 1
SID_shared	usr-sap-host2	/usr/sap/SID	Mounted at host 2
SID_shared	usr-sap-host3	/usr/sap/SID	Mounted at host 3
SID_shared	usr-sap-host4	/usr/sap/SID	Mounted at host 4
SID_shared	usr-sap-host5	/usr/sap/SID	Mounted at host 5

You can create the new data volume and mount it to the namespace using either ONTAP System Manager or the ONTAP CLI.

## Host configuration

In addition to the tasks described in the section "Host Setup," the additional mount points and `fstab` entries for the new additional data volume/s must be created and the new volumes must be mounted.

### 1. Create additional mount points.

- For a single-host system, create mount points and set the permissions on the database host:

```
sapcc-hana-tst-06:/ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-06:/ # chmod -R 777 /hana/data2/SID
```

- For a multiple-host system, create mount points and set the permissions on all worker and standby hosts.

The following example commands are for a 2-plus-1 multiple-host HANA system.

- First worker host:

```
sapcc-hana-tst-06:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-06:~ # chmod -R 777 /hana/data2/SID
```

- Second worker host:

```
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data2/SID
```

- Standby host:

```
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data2/SID
```

### 2. Add the additional file systems to the `/etc/fstab` configuration file on all hosts.

See the following example for a single-host system using NFSv4.1:

```
<storage-vif-data02>:/SID_data2_mnt00001 /hana/data2/SID/mnt00001 nfs
rw, vers=4
minorversion=1,hard,timeo=600,rsiz=1048576,wsiz=262144,bg,noatime,lock
0 0
```



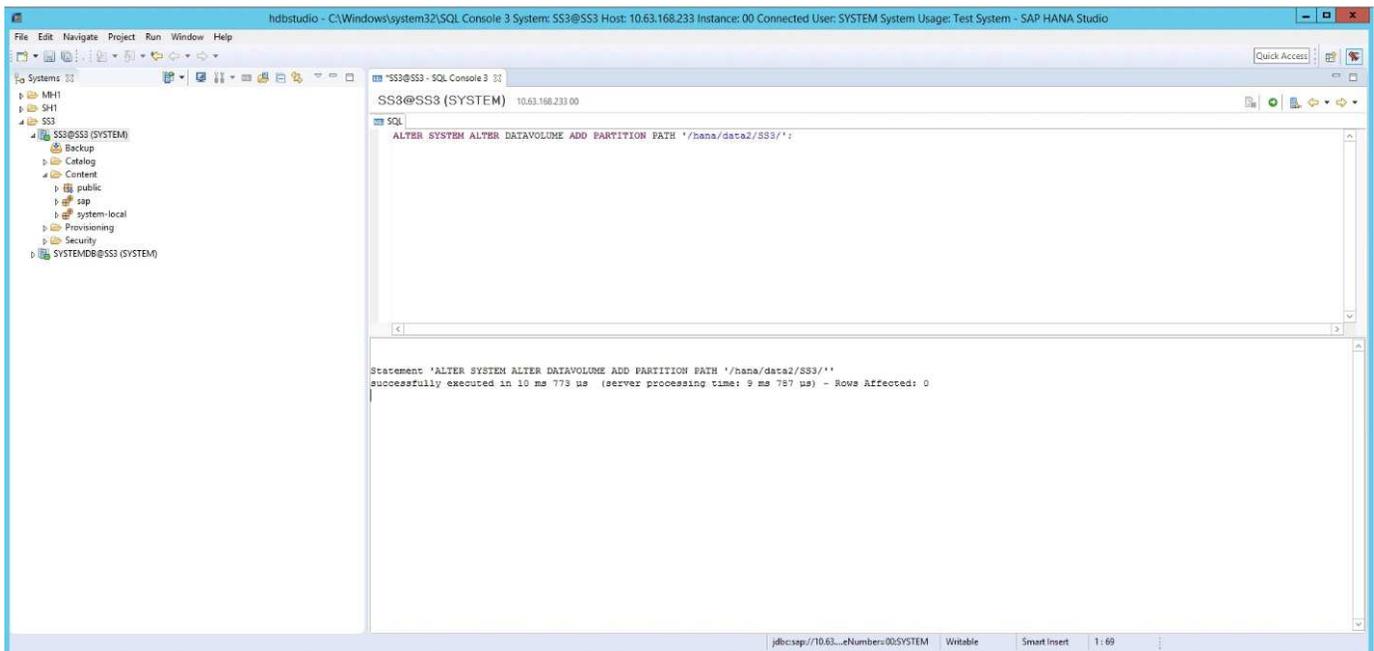
Use a different storage virtual interface for connecting each data volume to ensure that you are using different TCP sessions for each volume or use the nconnect mount option, if available for your OS.

3. Mount the file systems by running the `mount -a` command.

### Adding an additional data volume partition

Execute the following SQL statement against the tenant database to add an additional data volume partition to your tenant database. Use the path to additional volumes:

```
ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION PATH '/hana/data2/SID/';
```



### Where to find additional information

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

- [SAP HANA Software Solutions](#)
- [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP System Copies Using the SnapCenter SAP HANA Plug-In](#)
- [NetApp Documentation Centers](#)

<https://www.netapp.com/support-and-training/documentation/>

- [SAP Certified Enterprise Storage Hardware for SAP HANA](#)

<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/>

- SAP HANA Storage Requirements

<https://www.sap.com/documents/2024/03/146274d3-ae7e-0010-bca6-c68f7e60039b.html>

- SAP HANA Tailored Data Center Integration Frequently Asked Questions

<https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html>

- SAP HANA on VMware vSphere Wiki

[https://help.sap.com/docs/SUPPORT\\_CONTENT/virtualization/3362185751.html](https://help.sap.com/docs/SUPPORT_CONTENT/virtualization/3362185751.html)

- SAP HANA on VMware vSphere Best Practices Guide

[https://www.vmware.com/docs/sap\\_hana\\_on\\_vmware\\_vsphere\\_best\\_practices\\_guide-white-paper](https://www.vmware.com/docs/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper)

## Update history

The following technical changes have been made to this solution since its original publication.

Date	Update summary
October 2015	Initial version
March 2016	Updated capacity sizing Updated mount options for <code>/hana/shared</code> Updated <code>sysctl</code> parameter
February 2017	New NetApp storage systems and disk shelves New features of ONTAP 9 Support for 40GbE New OS releases (SUSE Linux Enterprise Server 12 SP1 and Red Hat Enterprise Linux 7.2) New SAP HANA release
July 2017	Minor updates
September 2018	New NetApp storage systems Support for 100GbE New OS releases (SUSE Linux Enterprise Server 12 SP3 and Red Hat Enterprise Linux 7.4) Additional minor changes SAP HANA 2.0 SPS3
October 2019	New NetApp storage systems and NVMe shelf New OS releases (SUSE Linux Enterprise Server 12 SP4, SUSE Linux Enterprise Server 15, and Red Hat Enterprise Linux 7.6) Max data volume size Minor changes
December 2019	New NetApp storage systems New OS release SUSE Linux Enterprise Server 15 SP1
March 2020	Support for <code>nconnect</code> for NFSv3 New OS release Red Hat Enterprise Linux 8

Date	Update summary
May 2020	Support for multiple data volume partitions available with SAP HANA 2.0 SPS4
June 2020	Additional information about optional functionalities Minor updates
December 2020	Support for nconnect for NFSv4.1 starting with ONTAP 9.8 New OS releases New SAP HANA versions
February 2021	New NetApp storage systems Changes in host network settings Minor changes
April 2021	VMware vSphere-specific information added
September 2022	New OS-Releases
August 2023	New Storage Systems (AFF C-Series)
December 2023	Update of host setup Revised nconnect settings Added information about NFSv4.1 sessions
May 2024	New Storage Systems (AFF A-Series)
September 2024	Minor Updates
November 2024	New Storage Systems
July 2025	Minor updates

# SAP HANA on NetApp ASA Systems with FCP Configuration Guide

## SAP HANA on NetApp ASA Systems with Fibre Channel Protocol

The NetApp ASA product family is certified for use with SAP HANA in TDI projects. This guide provides best practices for SAP HANA on this platform.

Marco Schoen, NetApp

### Introduction

The NetApp ASA A-Series and ASA C-Series product families have been certified for use with SAP HANA in tailored data center integration (TDI) projects.

This guide describes the best practices for the following certified models:

- ASA A20, ASA A30, ASA A50, ASA A70, ASA A90, ASA A1K
- ASA C30

For a complete list of NetApp certified storage solutions for SAP HANA, see the [Certified and supported SAP HANA hardware directory](#).

This document describes ASA configurations that use the Fibre Channel Protocol (FCP).

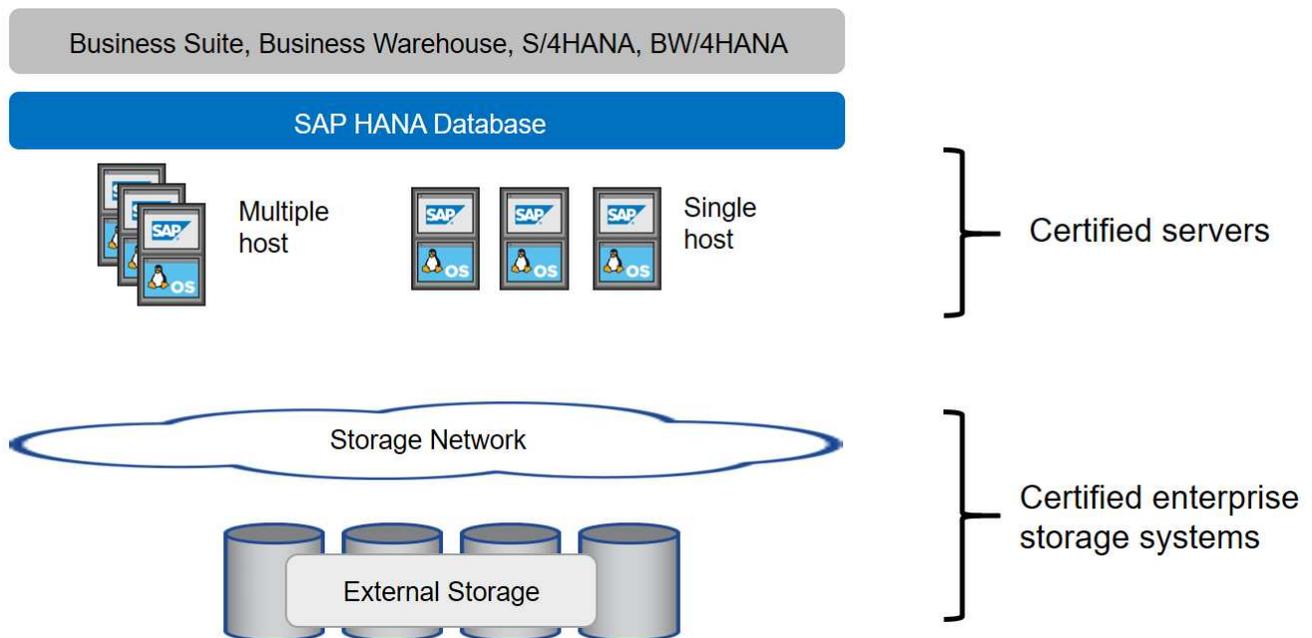


The configuration described in this paper is necessary to achieve the required SAP HANA KPIs and the best performance for SAP HANA. Changing any settings or using features not listed herein might cause performance degradation or unexpected behavior and should only be done if advised by NetApp support.

In an SAP HANA multiple-host environment, the standard SAP HANA storage connector is used to provide fencing in the event of an SAP HANA host failover. Always refer to the relevant SAP notes for operating system configuration guidelines and HANA specific Linux kernel dependencies. For more information, see [SAP Note 2235581 – SAP HANA Supported Operating Systems](#).

### SAP HANA tailored data center integration

NetApp ASA storage systems are certified in the SAP HANA TDI program using FC (SAN) protocols. They can be deployed in any of the current SAP HANA scenarios, such as SAP Business Suite on HANA, S/4HANA, BW/4HANA, or SAP Business Warehouse on HANA in either single-host or multiple-host configurations. Any server that is certified for use with SAP HANA can be combined with NetApp certified storage solutions. The following figure shows an architecture overview.



For more information regarding the prerequisites and recommendations for productive SAP HANA systems, see the following resource:

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

### SAP HANA using VMware vSphere

There are several options to connect storage to virtual machines (VMs). Raw device mappings (RDM), FCP datastores, or VVOL datastores with FCP are supported. For both datastore options, only one SAP HANA data or log volume must be stored within the datastore for productive use cases.

For more information about using vSphere with SAP HANA, see the following links:

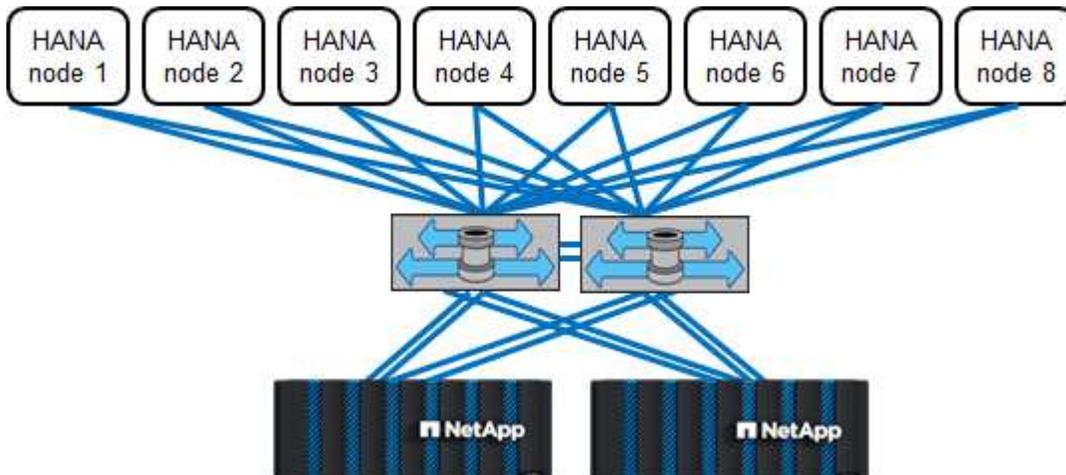
- [SAP HANA on VMware vSphere - Virtualization - Community Wiki](#)
- [SAP HANA on VMware vSphere Best Practices Guide](#)
- [2161991 - VMware vSphere configuration guidelines - SAP ONE Support Launchpad \(Login required\)](#)

## Architecture

SAP HANA hosts are connected to storage controllers using a redundant FCP infrastructure and multipath software. A redundant FCP switch infrastructure is required to provide fault-tolerant SAP HANA host-to-storage connectivity in case of switch or host bus adapter (HBA) failure. Appropriate zoning must be configured at the switch to allow all HANA hosts to reach the required LUNs on the storage controllers.

Different models of the ASA system product family can be mixed and matched at the storage layer to allow for growth and differing performance and capacity needs. The maximum number of SAP HANA hosts that can be attached to the storage system is defined by the SAP HANA performance requirements and the model of NetApp controller used. The number of required disk shelves is only determined by the capacity and performance requirements of the SAP HANA systems.

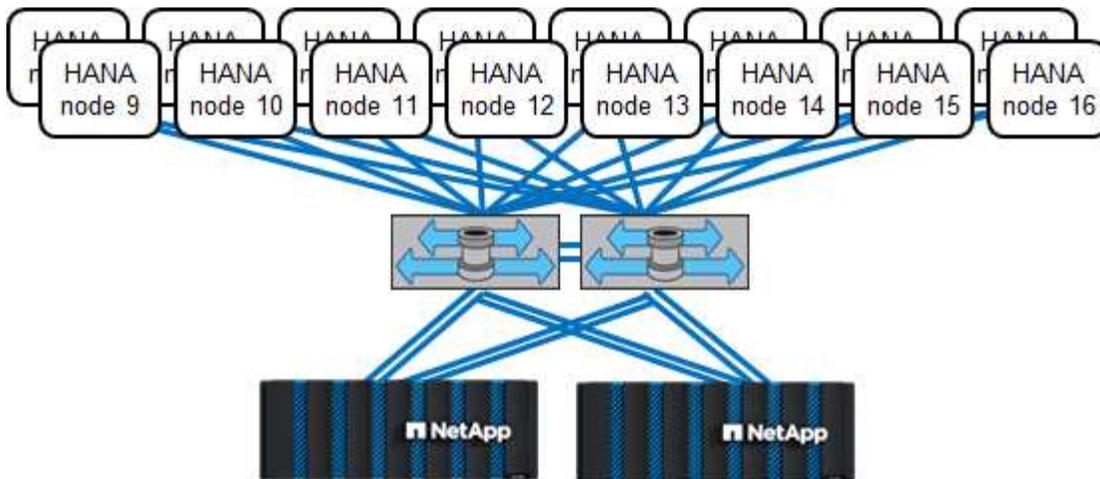
The following figure shows an example configuration with eight SAP HANA hosts attached to a storage HA pair.



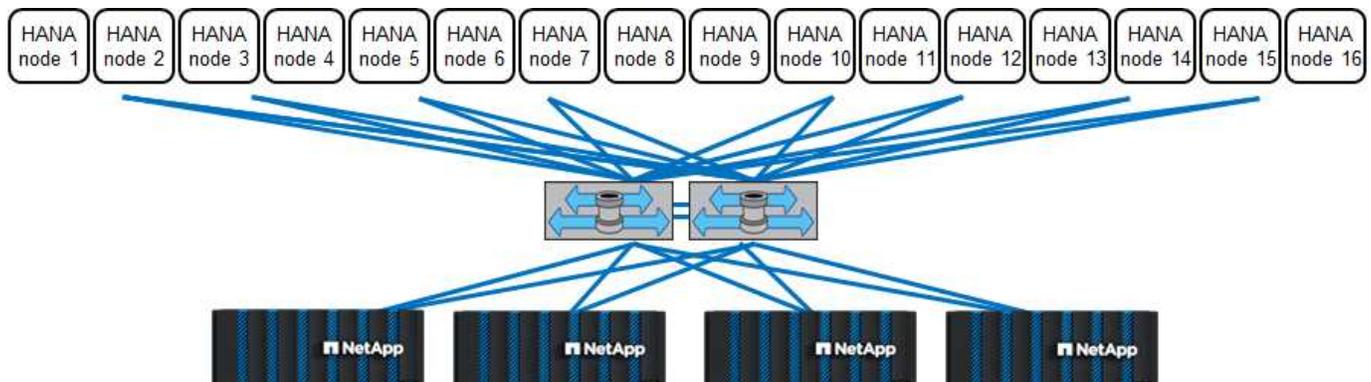
This architecture can be scaled in two dimensions:

- By attaching additional SAP HANA hosts and storage capacity to the existing storage, if the storage controllers provide enough performance to meet the current SAP HANA KPIs
- By adding more storage systems with additional storage capacity for the additional SAP HANA hosts

The following figure shows a configuration example in which more SAP HANA hosts are attached to the storage controllers. In this example, more disk shelves are necessary to meet the capacity and performance requirements of the 16 SAP HANA hosts. Depending on the total throughput requirements, you must add additional FC connections to the storage controllers.



Independent of the deployed ASA system, the SAP HANA landscape can also be scaled by adding any certified storage controllers to meet the desired node density, as shown in the following figure.

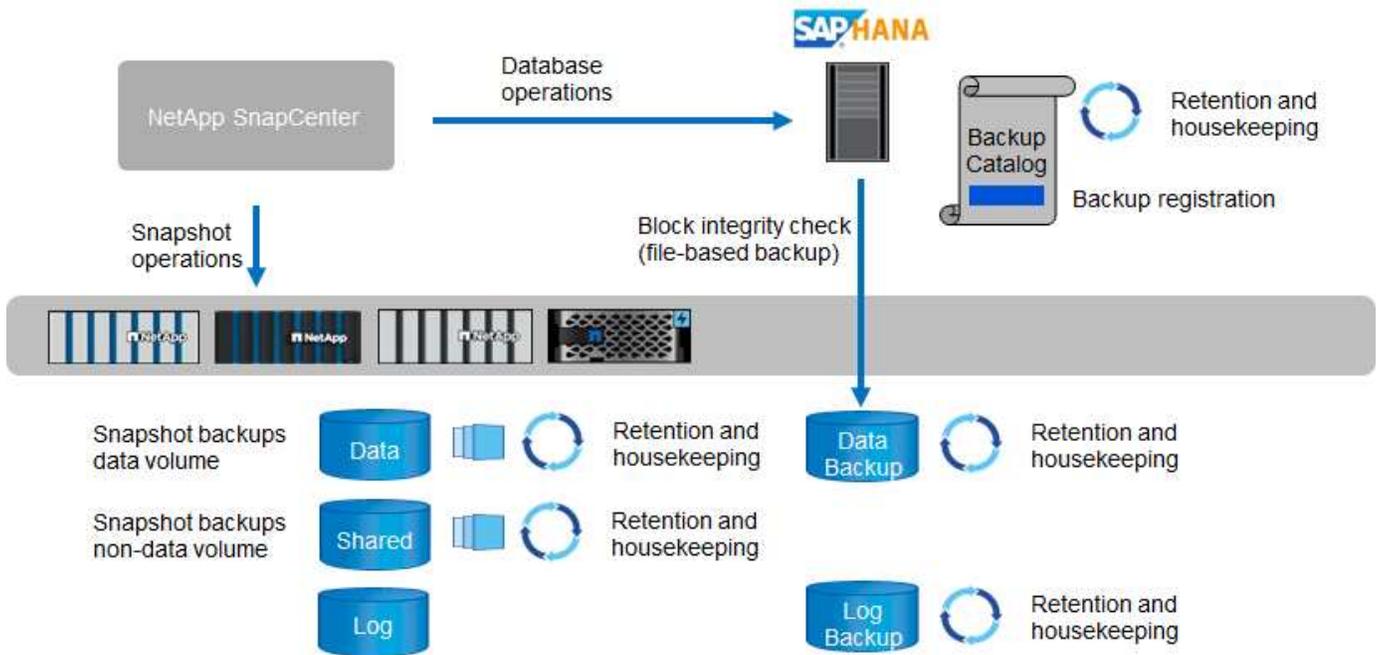


### SAP HANA backup

The ONTAP software present on all NetApp storage controllers provides a built-in mechanism to back up SAP HANA databases while in operation with no effect on performance. Storage-based NetApp Snapshot backups are a fully supported and integrated backup solution available for SAP HANA single containers and for SAP HANA MDC systems with a single tenant or multiple tenants.

Storage-based Snapshot backups are implemented by using the NetApp SnapCenter plug-in for SAP HANA. This allows users to create consistent storage-based Snapshot backups by using the interfaces provided natively by SAP HANA databases. SnapCenter registers each of the Snapshot backups into the SAP HANA backup catalog. Therefore, backups taken by SnapCenter are visible within SAP HANA Studio or Cockpit where they can be selected directly for restore and recovery operations.

NetApp SnapMirror technology allows for Snapshot copies that were created on one storage system to be replicated to a secondary backup storage system that is controlled by SnapCenter. Different backup retention policies can then be defined for each of the backup sets on the primary storage and also for the backup sets on the secondary storage systems. The SnapCenter Plug-in for SAP HANA automatically manages the retention of Snapshot copy-based data backups and log backups, including the housekeeping of the backup catalog. The SnapCenter Plug-in for SAP HANA also allows for the execution of a block integrity check of the SAP HANA database by executing a file-based backup.



Storage-based Snapshot backups provide significant advantages compared to conventional file-based backups. These advantages include, but are not limited to the following:

- Faster backup (a few minutes)
- Reduced RTO due to a much faster restore time on the storage layer (a few minutes) as well as more frequent backups
- No performance degradation of the SAP HANA database host, network, or storage during backup and recovery operations
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes

For detailed information about the SAP HANA backup and recovery solution, see [SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster](#).



At creation of this documents only VMware based VMs using VMDKs as storage are supported by SnapCenter for ASA.

## SAP HANA disaster recovery

SAP HANA disaster recovery can be done either on the database layer by using SAP HANA system replication or on the storage layer by using storage replication technologies. The following section provides an overview of disaster recovery solutions based on storage replication.

For detailed information about the SAP HANA disaster recovery solutions, see [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#).

### Storage replication based on SnapMirror

The following figure shows a three-site disaster recovery solution using synchronous SnapMirror active sync to the local DR datacenter and asynchronous SnapMirror to replicate the data to the remote DR datacenter. SnapMirror active sync enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy (RPO=0 and RTO=0). There is no manual intervention or custom scripting required to trigger a failover with SnapMirror active sync.

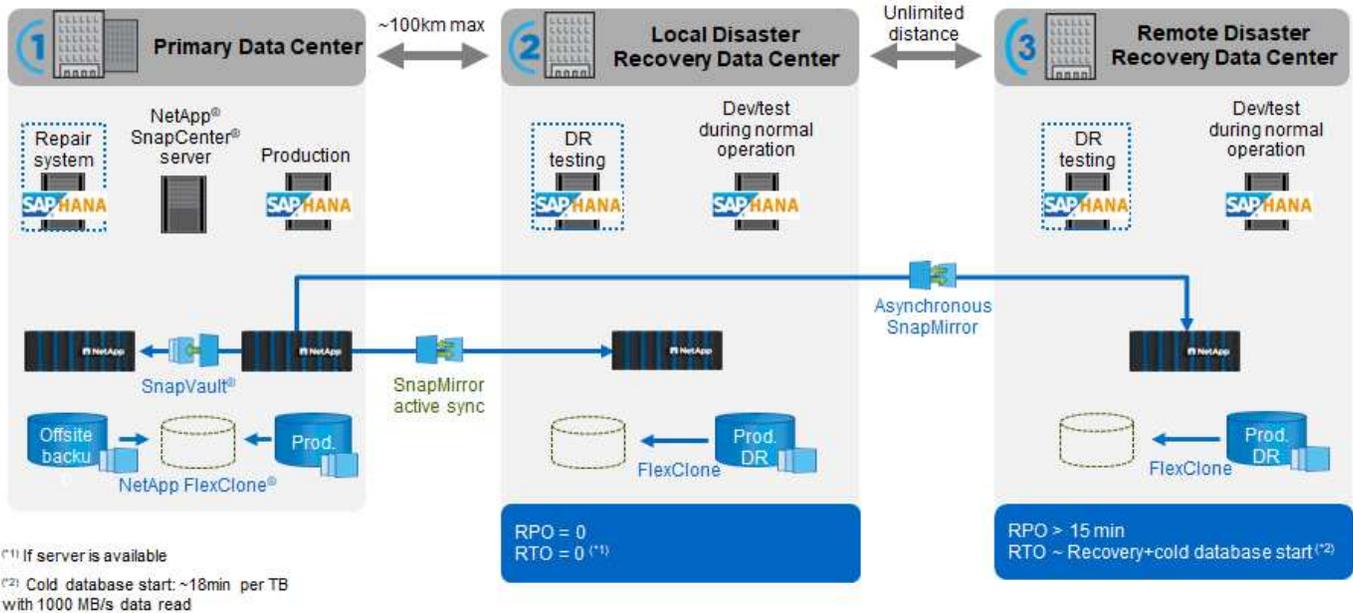
Beginning with ONTAP 9.15.1, SnapMirror active sync supports a symmetric active/active capability. Symmetric active/active enable read and write I/O operations from both copies of a protected LUN with bidirectional synchronous replication so that both LUN copies can serve I/O operations locally.

More details can be found at [SnapMirror active sync overview in ONTAP..](#)

The RTO for the asynchronous SnapMirror replication primarily depends on the time needed to start the HANA database at the DR site and load the data into memory. With the assumption that the data is read with a throughput of 1000MBps, loading 1TB of data would take approximately 18 minutes.

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

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



## Storage sizing

The following section provides an overview of performance and capacity considerations required for sizing a storage system for SAP HANA.



Contact your NetApp or NetApp partner sales representative to support the storage sizing process and to assist you with creating a properly sized storage environment.

## Performance considerations

SAP has defined a static set of storage key performance indicators (KPIs). These KPIs are valid for all production SAP HANA environments independent of the memory size of the database hosts and the applications that use the SAP HANA database. These KPIs are valid for single-host, multiple-host, Business Suite on HANA, Business Warehouse on HANA, S/4HANA, and BW/4HANA environments. Therefore, the current performance sizing approach depends on only the number of active SAP HANA hosts that are attached to the storage system.



Storage performance KPIs are only mandated for production SAP HANA systems, but you can implement them in for all HANA system.

SAP delivers a performance test tool which must be used to validate the storage systems performance for active SAP HANA hosts attached to the storage.

NetApp tested and predefined the maximum number of SAP HANA hosts that can be attached to a specific storage model, while still fulfilling the required storage KPIs from SAP for production-based SAP HANA systems.

The maximum number of SAP HANA hosts that can be run on a disk shelf and the minimum number of SSDs required per SAP HANA host were determined by running the SAP performance test tool. This test does not consider the actual storage capacity requirements of the hosts. You must also calculate the capacity requirements to determine the actual storage configuration needed.

#### **NS224 NVMe shelf**

One NVMe SSDs (data) supports up to 2/5 SAP HANA hosts depending on the specific NVMe disk being used.



Adding more disk shelves does not increase the maximum number of SAP HANA hosts that a storage controller can support.

#### **Mixed workloads**

SAP HANA and other application workloads running on the same storage controller or in the same storage aggregate are supported. However, it is a NetApp best practice to separate SAP HANA workloads from all other application workloads.

You might decide to deploy SAP HANA workloads and other application workloads on either the same storage controller or the same aggregate. If so, you must make sure that adequate performance is available for SAP HANA within the mixed workload environment. NetApp also recommends that you use quality of service (QoS) parameters to regulate the effect these other applications could have on SAP HANA applications and to guarantee throughput for SAP HANA applications.

The SAP HCMT test tool must be used to check if additional SAP HANA hosts can be run on an existing storage controller that is already in use for other workloads. SAP application servers can be safely placed on the same storage controller and/or aggregate as the SAP HANA databases.

#### **Capacity considerations**

A detailed description of the capacity requirements for SAP HANA is in the [SAP Note 1900823](#) white paper.



The capacity sizing of the overall SAP landscape with multiple SAP HANA systems must be determined by using SAP HANA storage sizing tools from NetApp. Contact NetApp or your NetApp partner sales representative to validate the storage sizing process for a properly sized storage environment.

#### **Configuration of performance test tool**

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used. These parameters must also be set for the performance test tool from SAP when the storage performance is being tested with the SAP test tool.

NetApp conducted performance tests to define the optimal values. The following table lists the parameters that must be set within the configuration file of the SAP test tool.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For more information about the configuration of SAP test tool, see [SAP note 1943937](#) for HWCCT (SAP HANA 1.0) and [SAP note 2493172](#) for HCMT/HCOT (SAP HANA 2.0).

The following example shows how variables can be set for the HCMT/HCOT execution plan.

```
...
{
    "Comment": "Log Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "LogAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Data Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "DataAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Log Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "LogAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Data Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "DataAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Log Volume: Controls which blocks are written
```

```

asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
  "Name": "LogAsyncWriteSubmitBlocks",
  "Value": "all",
  "Request": "false"
},
{
  "Comment": "Data Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
  "Name": "DataAsyncWriteSubmitBlocks",
  "Value": "all",
  "Request": "false"
},
{
  "Comment": "Log Volume: Maximum number of parallel I/O requests
per completion queue",
  "Name": "LogExtMaxParallelIoRequests",
  "Value": "128",
  "Request": "false"
},
{
  "Comment": "Data Volume: Maximum number of parallel I/O requests
per completion queue",
  "Name": "DataExtMaxParallelIoRequests",
  "Value": "128",
  "Request": "false"
}, ...

```

These variables must be used for the test configuration. This is usually the case with the predefined execution plans SAP delivers with the HCMT/HCOT tool. The following example for a 4k log write test is from an execution plan.

```

...
{
  "ID": "D664D001-933D-41DE-A904F304AEB67906",
  "Note": "File System Write Test",
  "ExecutionVariants": [
    {
      "ScaleOut": {
        "Port": "${RemotePort}",
        "Hosts": "${Hosts}",
        "ConcurrentExecution": "${FSConcurrentExecution}"
      },
      "RepeatCount": "${TestRepeatCount}",
      "Description": "4K Block, Log Volume 5GB, Overwrite",
      "Hint": "Log",
      "InputVector": {
        "BlockSize": 4096,
        "DirectoryName": "${LogVolume}",
        "FileOverwrite": true,
        "FileSize": 5368709120,
        "RandomAccess": false,
        "RandomData": true,
        "AsyncReadSubmit": "${LogAsyncReadSubmit}",
        "AsyncWriteSubmitActive":
"${LogAsyncWriteSubmitActive}",
        "AsyncWriteSubmitBlocks":
"${LogAsyncWriteSubmitBlocks}",
        "ExtMaxParallelIoRequests":
"${LogExtMaxParallelIoRequests}",
        "ExtMaxSubmitBatchSize": "${LogExtMaxSubmitBatchSize}",
        "ExtMinSubmitBatchSize": "${LogExtMinSubmitBatchSize}",
        "ExtNumCompletionQueues":
"${LogExtNumCompletionQueues}",
        "ExtNumSubmitQueues": "${LogExtNumSubmitQueues}",
        "ExtSizeKernelIoQueue": "${ExtSizeKernelIoQueue}"
      }
    },
    ...
  ]
}

```

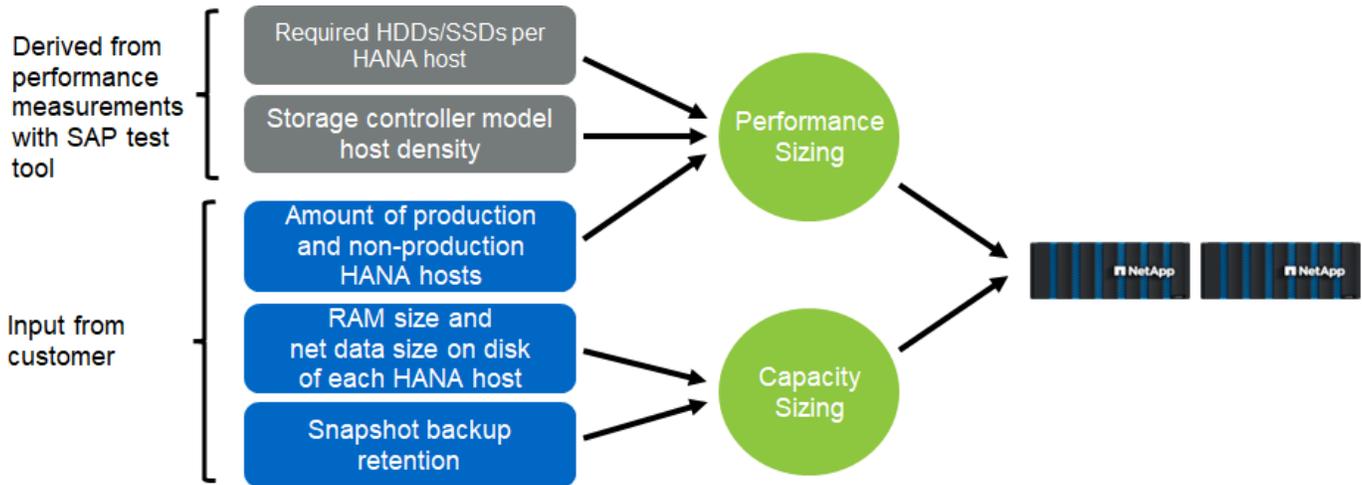
## Storage sizing process overview

The number of disks per HANA host and the SAP HANA host density for each storage model were determined using the SAP HANA test tool.

The sizing process requires details such as the number of production and nonproduction SAP HANA hosts, the RAM size of each host, and the backup retention of the storage-based Snapshot copies. The number of SAP HANA hosts determines the storage controller and the number of disks required.

The size of the RAM, net data size on the disk of each SAP HANA host, and the Snapshot copy backup retention period are used as inputs during capacity sizing.

The following figure summarizes the sizing process.



## Infrastructure setup and configuration

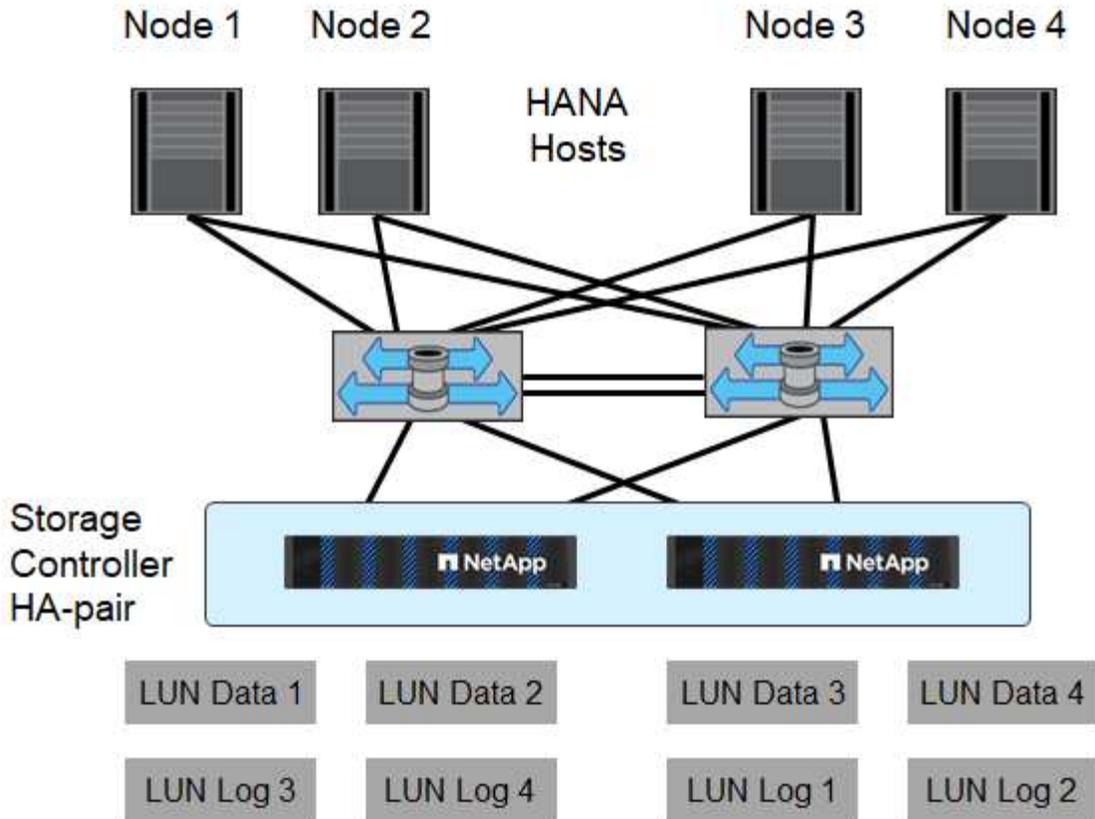
The following sections provide SAP HANA infrastructure setup and configuration guidelines and describes all the steps needed to set up an SAP HANA system. Within these sections, the following example configurations are used:

- HANA system with SID=FC5
  - SAP HANA single and multiple host

### SAN fabric setup

Each SAP HANA server must have a redundant FCP SAN connection with a minimum of 8Gbps bandwidth. For each SAP HANA host attached to a storage controller, at least 8Gbps bandwidth must be configured at the storage controller.

The following figure shows an example with four SAP HANA hosts attached to two storage controllers. Each SAP HANA host has two FCP ports connected to the redundant fabric. At the storage layer, four FCP ports are configured to provide the required throughput for each SAP HANA host.



In addition to the zoning on the switch layer, you must map each LUN on the storage system to the hosts that connect to this LUN. Keep the zoning on the switch simple; that is, define one zone set in which all host HBAs can see all controller HBAs.

### Time synchronization

You must synchronize the time between the storage controllers and the SAP HANA database hosts. To do so, set the same time server for all storage controllers and all SAP HANA hosts.

### Storage controller setup

This section describes the configuration of the NetApp storage system. You must complete the primary installation and setup according to the corresponding Data ONTAP setup and configuration guides.

### Storage efficiency

Inline deduplication, cross-volume inline deduplication, inline compression, and inline compaction are supported with SAP HANA in an SSD configuration.

### Quality of Service

QoS can be used to limit the storage throughput for specific SAP HANA systems or non-SAP applications on a shared controller.

## Production and Dev/Test

One use case would be to limit the throughput of development and test systems so that they cannot influence production systems in a mixed setup.

During the sizing process, you should determine the performance requirements of a nonproduction system. Development and test systems can be sized with lower performance values, typically in the range of 20% to 50% of a production-system KPI as defined by SAP.

Large write I/O has the biggest performance effect on the storage system. Therefore, the QoS throughput limit should be set to a percentage of the corresponding write SAP HANA storage performance KPI values in the data and log volumes.

## Shared Environments

Another use case is to limit the throughput of heavy write workloads, especially to avoid that these workloads have an impact on other latency sensitive write workloads.

In such environments it is best practice to apply a non-shared throughput ceiling QoS group-policy to each LUN within each Storage Virtual Machine (SVM) to restrict the max throughput of each individual storage object to the given value. This reduces the possibility that a single workload can negatively influence other workloads.

To do so, a group-policy needs to be created using the CLI of the ONTAP cluster for each SVM:

```
qos policy-group create -policy-group <policy-name> -vserver <vserver
name> -max-throughput 1000MB/s -is-shared false
```

and applied to each LUN within the SVM. Below is an example to apply the policy group to all existing LUNs within an SVM:

```
lun modify -vserver <vserver name> -path * -qos-policy-group <policy-
name>
```

This needs to be done for every SVM. The name of the QoS police group for each SVM needs to be different. For new LUNs, the policy can be applied directly:

```
lun create -vserver <vserver_name> -path /vol/<volume_name>/<lun_name>
-size <size> -ostype <e.g. linux> -qos-policy-group <policy-name>
```

It is recommended to use 1000MB/s as maximum throughput for a given LUN. If an application requires more throughput, multiple LUNs with LUN striping shall be used to provide the needed bandwidth. This guide provides an example for SAP HANA based on Linux LVM in section [Host Setup](#).



The limit applies also to reads. Therefore use enough LUNs to fulfil the required SLAs for SAP HANA database startup time and for backups.

## Configure storage

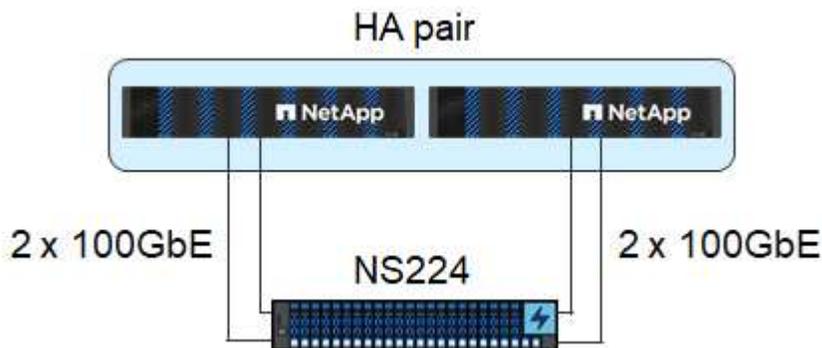
The following overview summarizes the required storage configuration steps. Each step is covered in more detail in the subsequent sections. In this section, we assume that the storage hardware is set up and that the

ONTAP software is already installed. Also, the connection of the storage FCP ports to the SAN fabric must already be in place.

1. Check the correct disk shelf configuration, as described in [NVMe-based disk shelves](#).
2. Create initiator groups (igroups) with worldwide names (WWNs) of HANA servers as described in the section [xref:./bp/saphana-asa-fc-storage-controller-setup.html#initiator-groups](#) [Initiator groups](#).
3. Create LUNs and map them to the servers described in the section [LUN configuration for SAP HANA single-host systems](#) and [LUN configuration for SAP HANA multiple-hosts systems](#)

### NVMe-based disk shelves

Each NS224 NVMe disk shelf is connected with two 100GbE ports per controller, as shown in the following figure. The disks are automatically distributed to both controllers of the HA pair.



### Initiator groups

An igroup can be configured for each server or for a group of servers that require access to a LUN. The igroup configuration requires the worldwide port names (WWPNs) of the servers.

Using the `sanlun` tool, run the following command to obtain the WWPNs of each SAP HANA host:

```
sapcc-hana-tst:~ # sanlun fcp show adapter
/sbin/udevadm
/sbin/udevadm

host0 ..... WWPN:2100000e1e163700
host1 ..... WWPN:2100000e1e163701
```



The `sanlun` tool is part of the NetApp Host Utilities and must be installed on each SAP HANA host. More details can be found in section [Host setup](#).

### Single host

This section describes the configuration of the NetApp storage system specific to SAP HANA single-host systems

## Creating LUNs and mapping LUNs to initiator groups

You can use NetApp ONTAP System Manager to create storage volumes and LUNs and the map them to the igroups of the servers and the ONTAP CLI.

### Creating LUNs and mapping LUNs to initiator groups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a SAP HANA single host system with SID FC5 using LVM and two LUNs per LVM volume group:

1. Create all LUNs.

```
lun create -path FC5_data_mnt00001_1 -size 1t -ostype linux -class regular
lun create -path FC5_data_mnt00001_2 -size 1t -ostype linux -class regular
lun create -path FC5_log_mnt00001_1 -size 260g -ostype linux -class regular
lun create -path FC5_log_mnt00001_2 -size 260g -ostype linux -class regular
lun create -path FC5_shared -size 260g -ostype linux -class regular
```

2. Create the initiator group for all servers belonging to system FC5.

```
lun igroup create -igroup HANA-FC5 -protocol fcp -ostype linux
-initiator 10000090fadcc5fa,10000090fadcc5fb -vserver svml
```

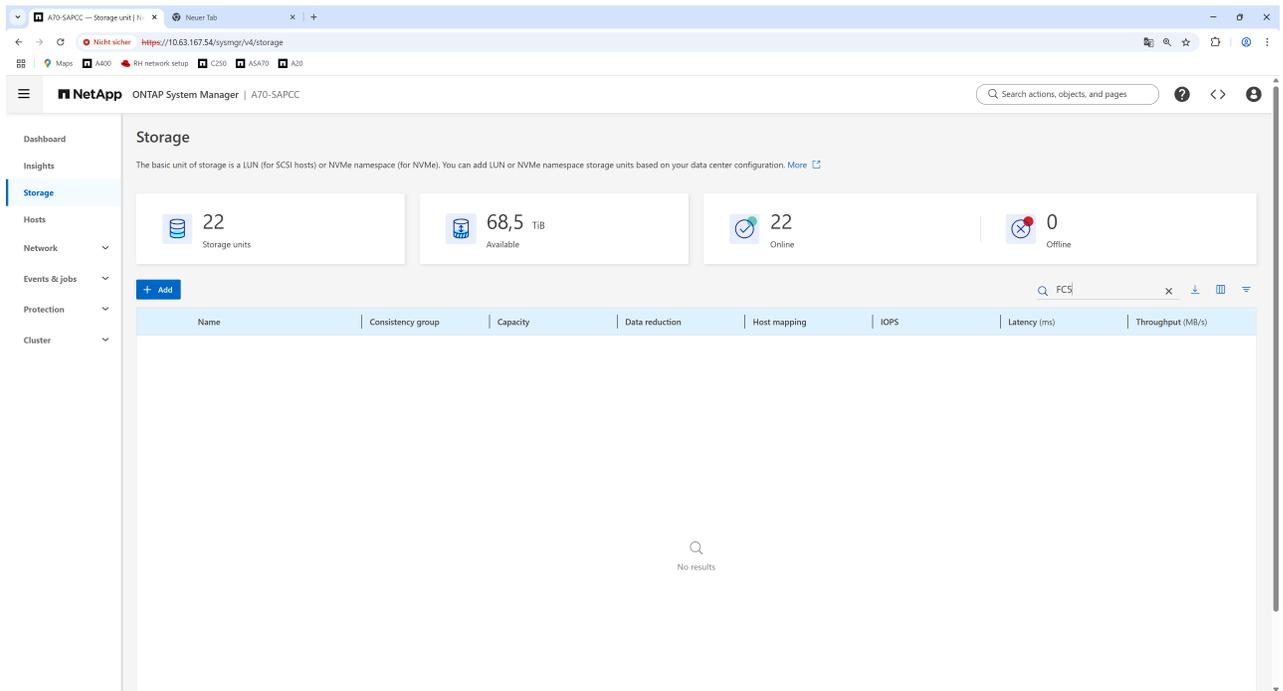
3. Map all LUNs to created initiator group.

```
lun map -path FC5_data_mnt00001_1 -igroup HANA-FC5
lun map -path FC5_data_mnt00001_2 -igroup HANA-FC5
lun map -path FC5_log_mnt00001_1 -igroup HANA-FC5
lun map -path FC5_log_mnt00001_2 -igroup HANA-FC5
lun map -path FC5_shared -igroup HANA-FC5
```

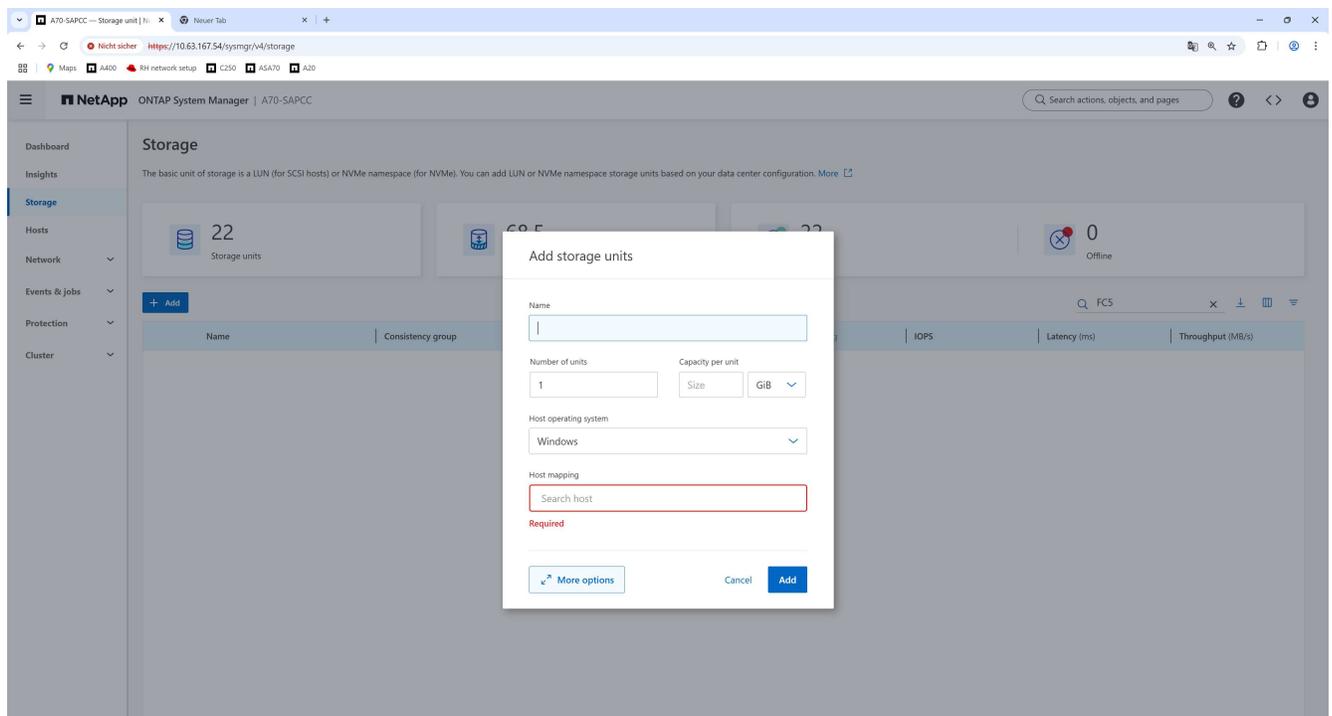
### Creating LUNs and mapping LUNs to initiator groups using the GUI

This section shows an example configuration using ONTAP System Manager for a SAP HANA single host system with SID FC5 using LVM and two LUNs per LVM volume group:

1. Log on to the ONTAP System Manager of your ONTAP Cluster and choose Storage from the left menu.
  - a. Press Add



## 2. Choose More options



## 3. Provide the required information:

- the name of the data LUNs, e.g. FC5\_data\_mnt00001
- the amount of LUNs to be combined with LVM, e.g. 2
- the size of each LUN, e.g. 1000 GB
- choose SCSI (FC or iSCSI)
- choose Linux as Host Operating system

- f. choose `New host` for the `Host mapping` option, provide a name, e.g `FC5_host`, pick or add the desired initiators
- g. Keep `Schedule snapshots` unchecked
- h. press `Add`

NetApp ONTAP System Manager | A70-SAPCC

Search actions, objects, and pages

Dashboard

Insights

**Storage**

Hosts

Network

Events & jobs

Protection

Cluster

### Add storage units

Name  
FC5\_data\_mnt00001

### Storage and optimization

Number of units: 2  
Capacity per unit: 1000 GiB

+ Add a different capacity

Quality of service (QoS): Unlimited

### Host information

Select a connection protocol based on your host and data center configuration.

Connection protocol  
 SCSI (FC or iSCSI)     NVMe

Host operating system  
Linux

Host mapping  
 Existing hosts  
 New host group  
 New hosts

Host Name  
FC5\_Host

FC (2)     iSCSI

Name	Description
<input checked="" type="checkbox"/> 10:00:70:b7:e4:08:94:75	-
<input checked="" type="checkbox"/> 10:00:70:b7:e4:08:94:76	-
<input type="checkbox"/> 10:00:70:b7:e4:0a:e0:cc	-
<input type="checkbox"/> 10:00:70:b7:e4:0a:e0:cd	-
<input type="checkbox"/> 10:00:70:b7:e4:0a:e2:ed	-

+ Add initiator

### Local protection

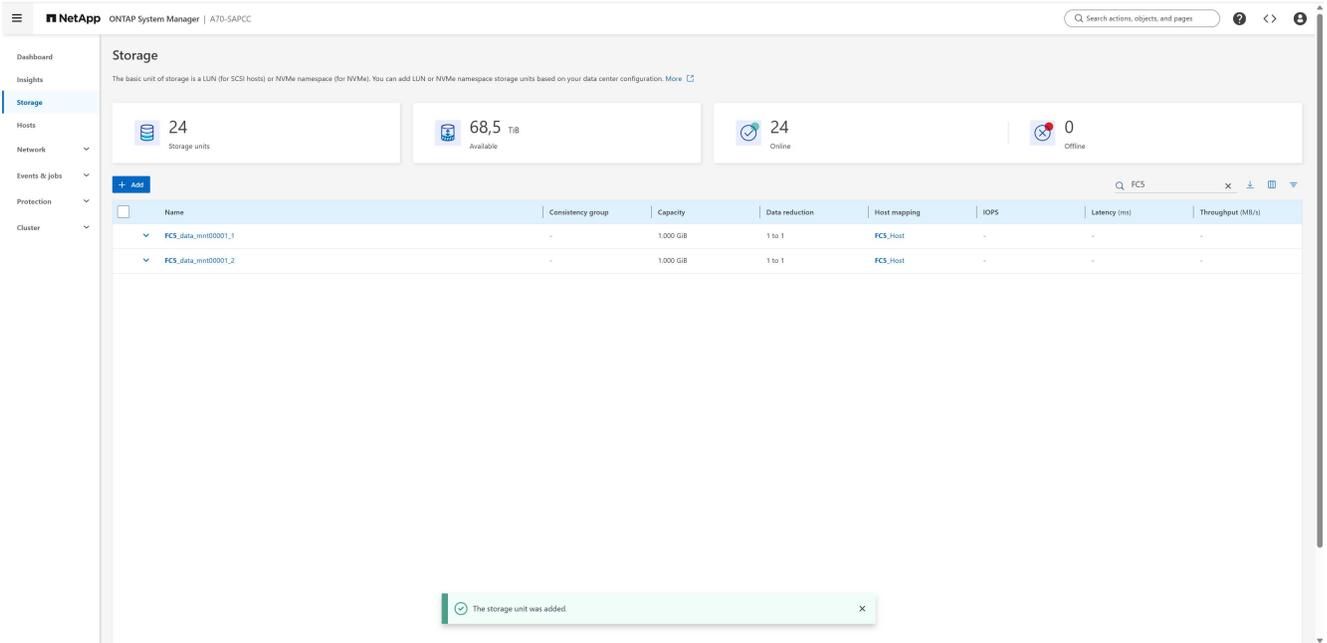
Schedule snapshots

### Remote protection

Replicate to a remote cluster  
SnapMirror copies snapshots to a remote cluster.

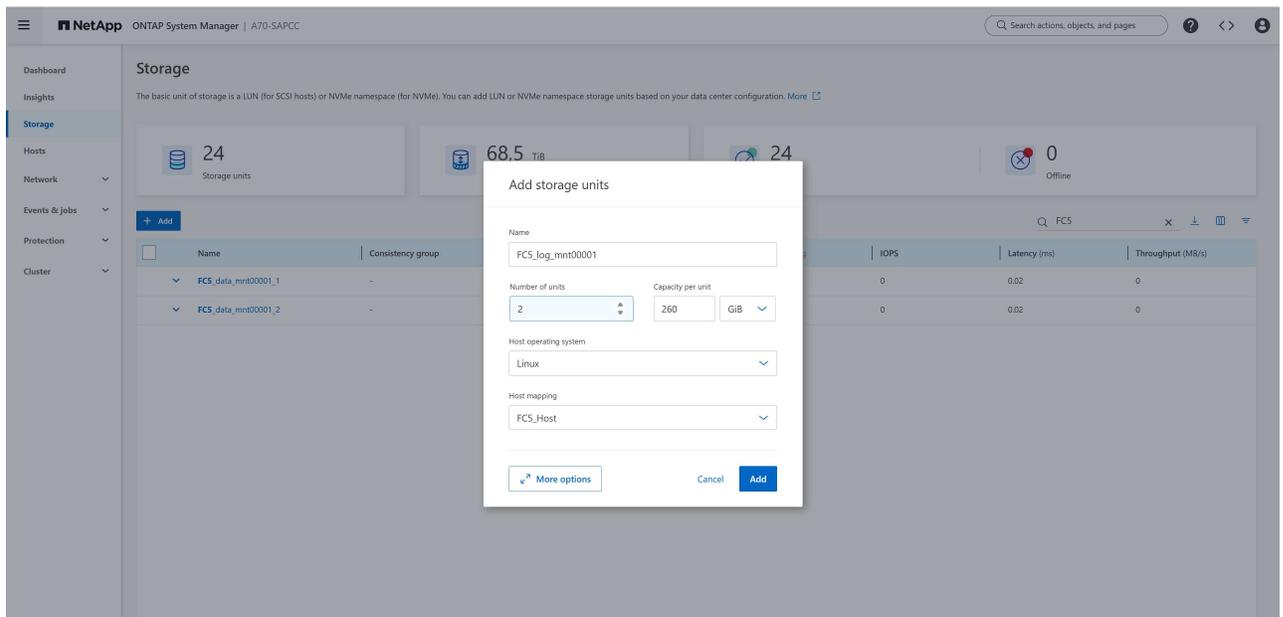
**Add** Cancel

4. After successful creation of the data LUNs create the log LUNs by pressing Add

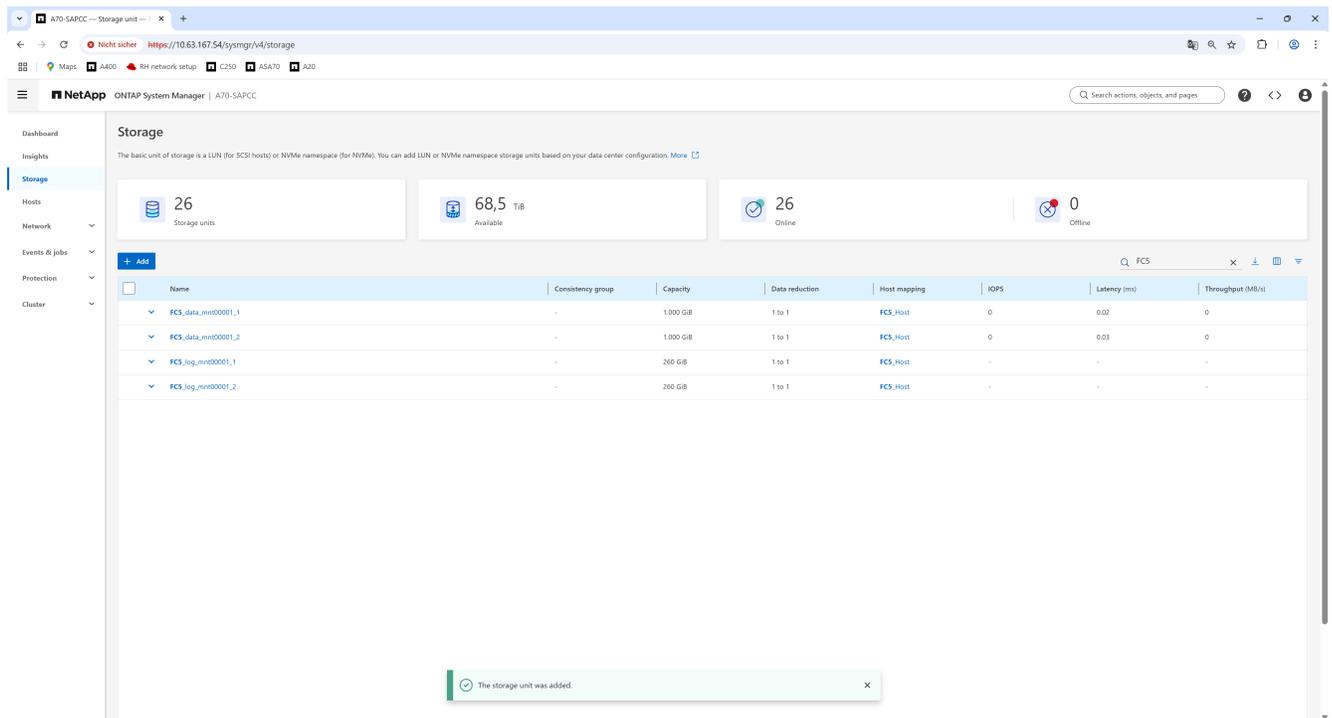


5. Provide the required information:

- the name of the log LUNs, e.g. FC5\_log\_mnt00001
- the amount of LUNs to be combined with LVM, e.g. 2
- the size of each LUN, e.g. 260 GB
- choose Linux as Host Operating system
- choose the previously created mapping FC5\_host for the Host mapping option
- press Add

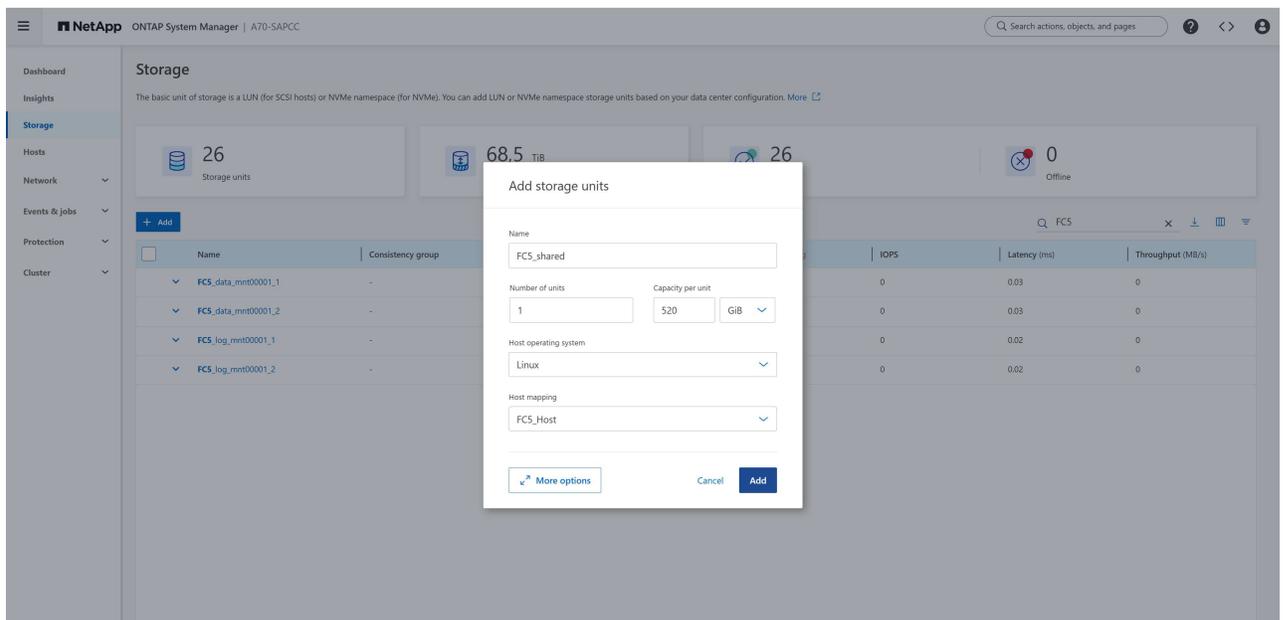


6. After successful creation of the log LUNs create the shared LUN by pressing Add

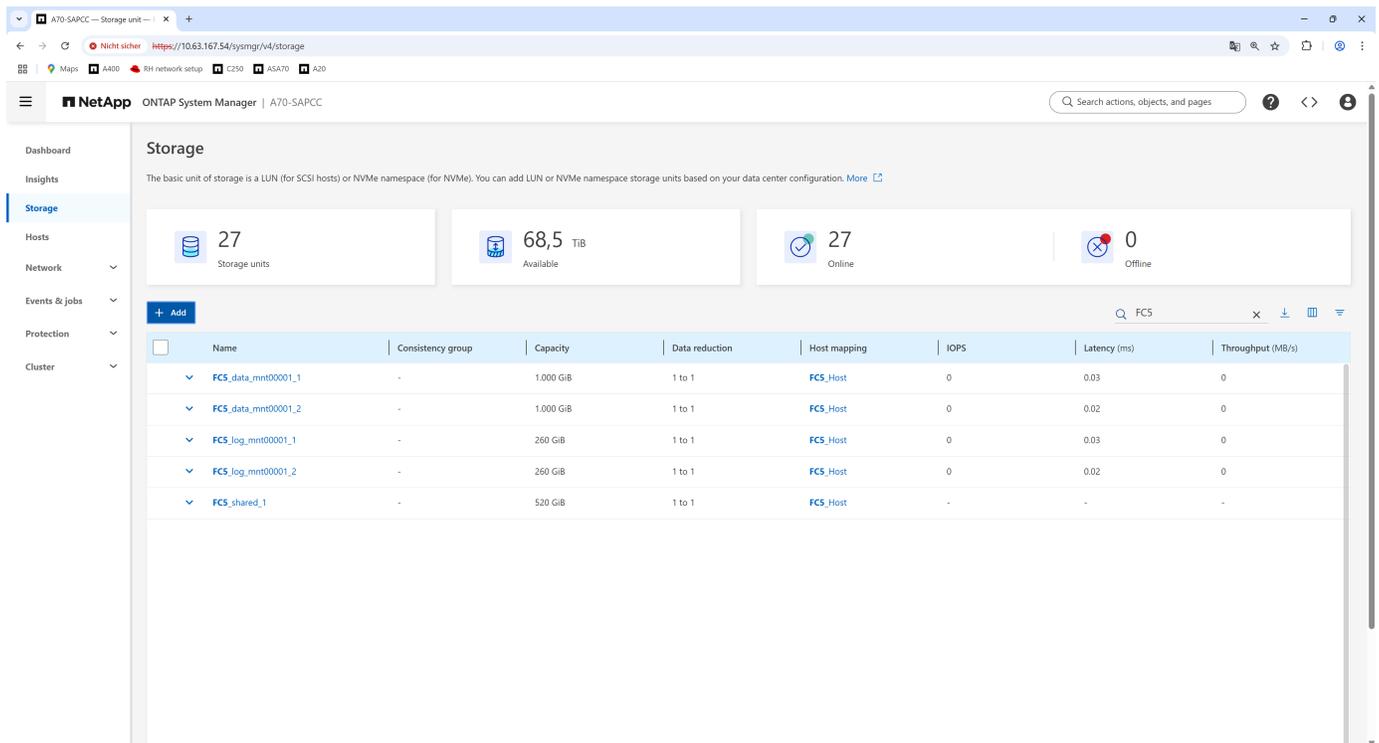


7. Provide the required information:

- the name of the shared LUN, e.g. FC5\_shared
- the amount of LUNs, e.g. 1
- the size of the LUN, e.g. 520 GB
- choose Linux as Host Operating system
- choose the previously created mapping FC5\_host for the Host mapping option
- press Add



All required LUNs for a SAP HANA single-host system have been created.



## Multiple hosts

This section describes the configuration of the NetApp storage system specific to SAP HANA multiple-hosts systems

### Creating LUNs and mapping LUNs to initiator groups

You can use NetApp ONTAP System Manager to create storage volumes and LUNs and the map them to the igroups of the servers and the ONTAP CLI.

### Creating LUNs and mapping LUNs to initiator groups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a 2+1 SAP HANA multiple host system with SID FC5 using LVM and two LUNs per LVM volume group:

1. Create all LUNs.

```

lun create -path FC5_data_mnt00001_1 -size 1t -ostype linux -class
regular
lun create -path FC5_data_mnt00001_2 -size 1t -ostype linux -class
regular
lun create -path FC5_data_mnt00002_1 -size 1t -ostype linux -class
regular
lun create -path FC5_data_mnt00002_2 -size 1t -ostype linux -class
regular
lun create -path FC5_log_mnt00001_1 -size 260g -ostype linux -class
regular
lun create -path FC5_log_mnt00001_2 -size 260g -ostype linux -class
regular
lun create -path FC5_log_mnt00002_1 -size 260g -ostype linux -class
regular
lun create -path FC5_log_mnt00002_2 -size 260g -ostype linux -class
regular

```

## 2. Create the initiator group for all servers belonging to system FC5.

```

lun igroup create -igroup HANA-FC5 -protocol fcp -ostype linux
-initiator
10000090fadcc5fa,10000090fadcc5fb,10000090fadcc5c1,10000090fadcc5c2,1000
0090fadcc5c3,10000090fadcc5c4 -vserver svm1

```

## 3. Map all LUNs to created initiator group.

```

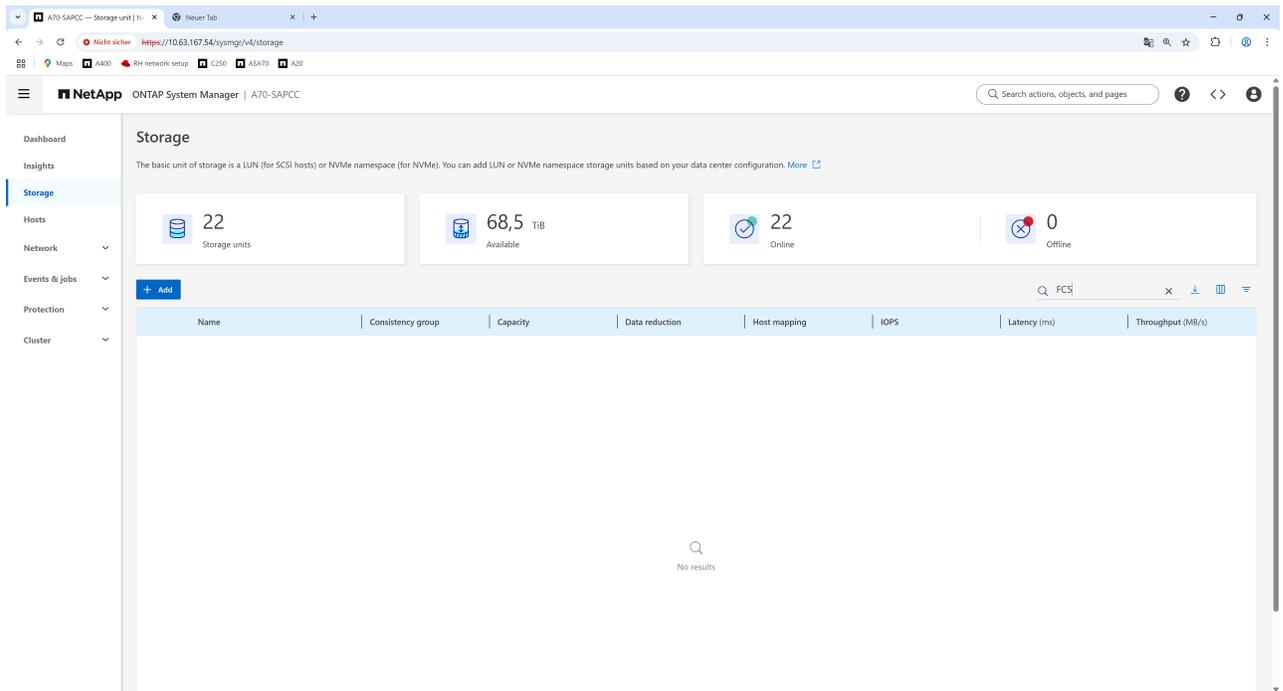
lun map -path FC5_data_mnt00001_1 -igroup HANA-FC5
lun map -path FC5_data_mnt00001_2 -igroup HANA-FC5
lun map -path FC5_data_mnt00002_1 -igroup HANA-FC5
lun map -path FC5_data_mnt00002_2 -igroup HANA-FC5
lun map -path FC5_log_mnt00001_1 -igroup HANA-FC5
lun map -path FC5_log_mnt00001_2 -igroup HANA-FC5
lun map -path FC5_log_mnt00002_1 -igroup HANA-FC5
lun map -path FC5_log_mnt00002_2 -igroup HANA-FC5

```

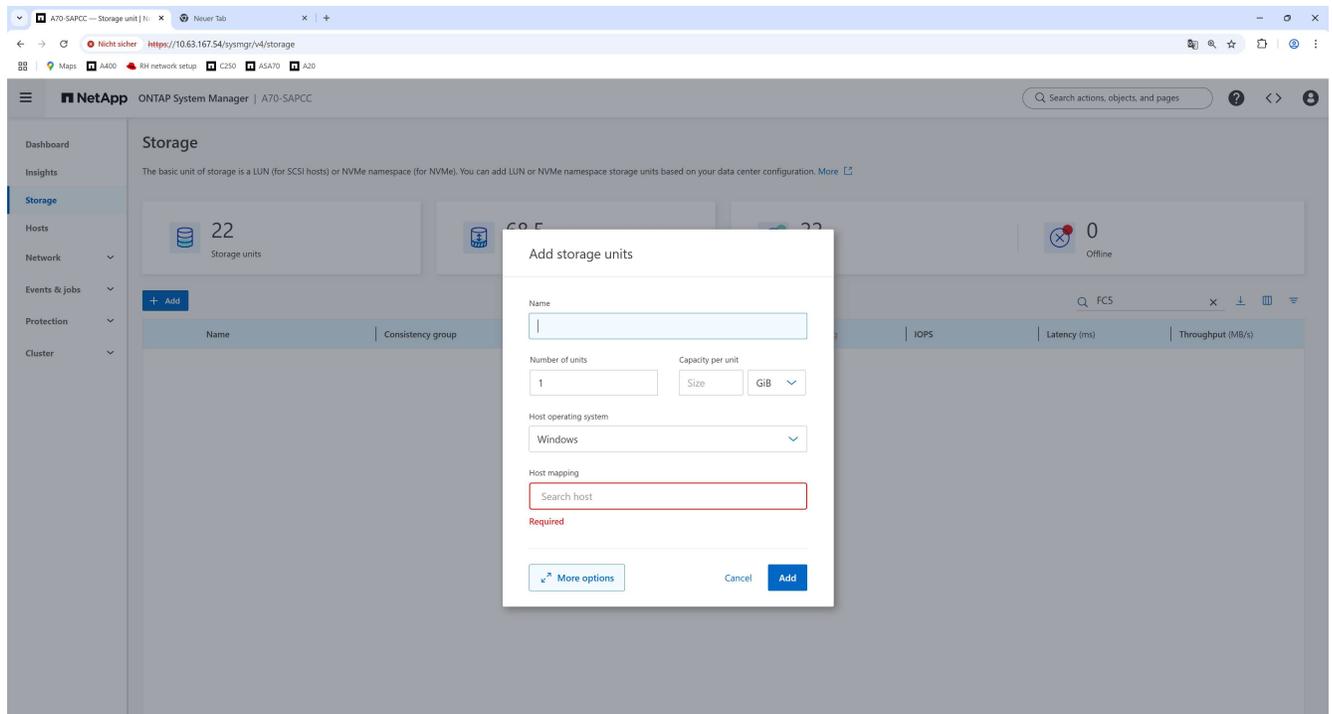
## Creating LUNs and mapping LUNs to initiator groups using the GUI

This section shows an example configuration using ONTAP System Manager for a 2+1 SAP HANA multiple host system with SID FC5 using LVM and two LUNs per LVM volume group:

1. Log on to the ONTAP System Manager of your ONTAP Cluster and choose Storage from the left menu.
  - a. Press Add



## 2. Choose More options



## 3. Provide the required information:

- name of the data LUNs, e.g. FC5\_data\_mnt00001
- the amount of LUNs to be combined with LVM, e.g. 2
- the size of each LUN, e.g. 1000 GB
- choose SCSI (FC or iSCSI)
- choose Linux as Host Operating system

- f. choose `New host` for the `Host mapping` option, provide a name, e.g `FC5_host`, pick or add the desired initiators
- g. Keep `Schedule snapshots` unchecked
- h. press `Add`

NetApp ONTAP System Manager | A70-SAPCC

Search actions, objects, and pages

Dashboard

Insights

**Storage**

Hosts

Network

Events & jobs

Protection

Cluster

### Add storage units

Name  
FC5\_data\_mnt00001

### Storage and optimization

Number of units: 2  
Capacity per unit: 1000 GiB

+ Add a different capacity

Quality of service (QoS): Unlimited

### Host information

Select a connection protocol based on your host and data center configuration.

Connection protocol  
 SCSI (FC or iSCSI)  NVMe

Host operating system  
Linux

Host mapping  
 Existing hosts  
 New host group  
 New hosts

Host Name  
FC5\_Host

FC (2)  iSCSI

Name	Description
<input checked="" type="checkbox"/> 10:00:70:b7:e4:08:94:75	-
<input checked="" type="checkbox"/> 10:00:70:b7:e4:08:94:76	-
<input type="checkbox"/> 10:00:70:b7:e4:0a:e0:cc	-
<input type="checkbox"/> 10:00:70:b7:e4:0a:e0:cd	-
<input type="checkbox"/> 10:00:70:b7:e4:0a:e2:ed	-

+ Add initiator

### Local protection

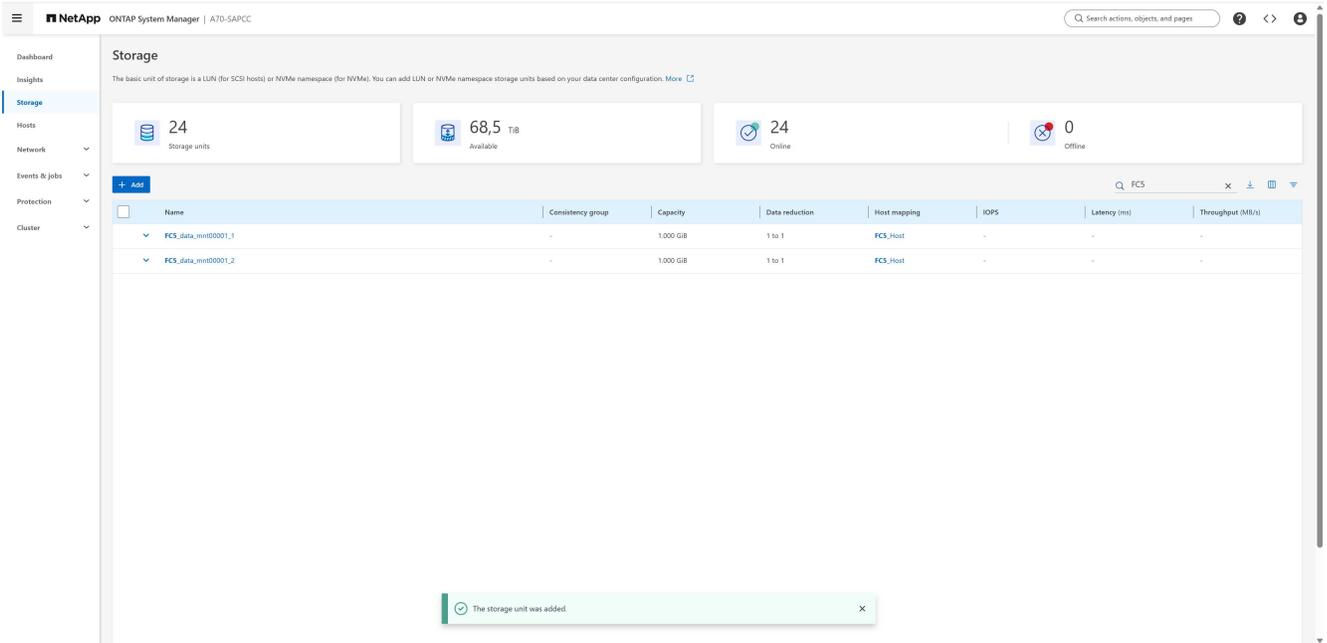
Schedule snapshots

### Remote protection

Replicate to a remote cluster  
SnapMirror copies snapshots to a remote cluster.

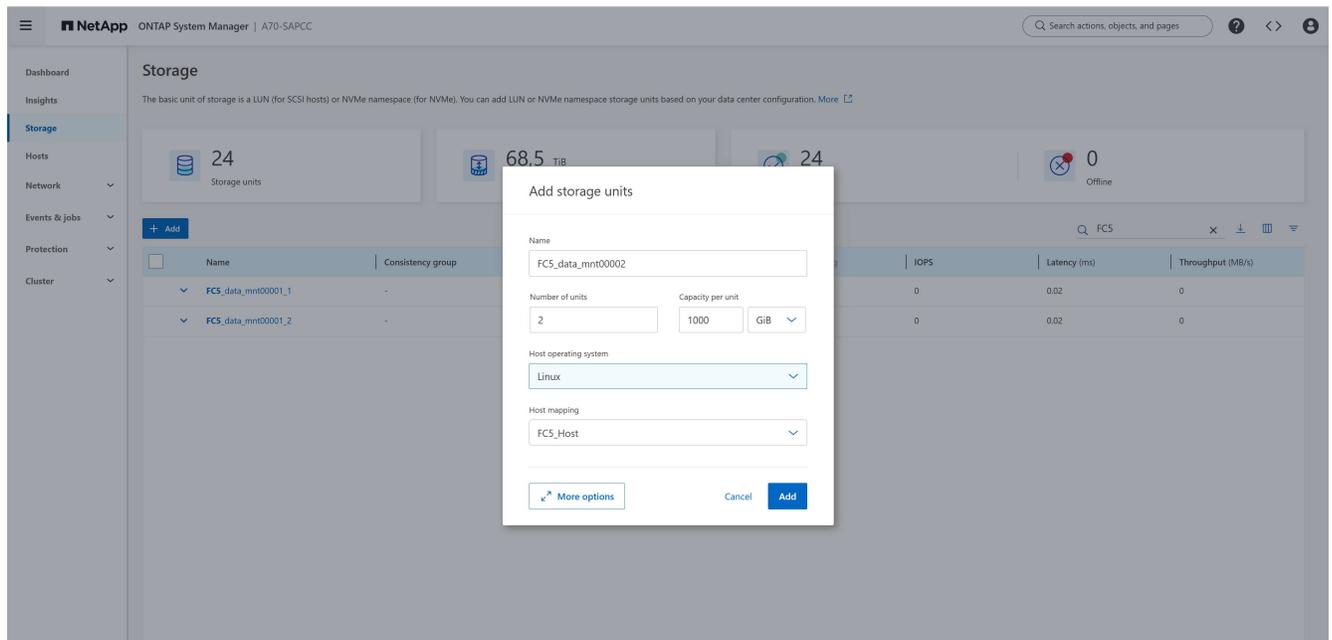
Add Cancel

4. Create the data LUNs for the next worker host by pressing Add



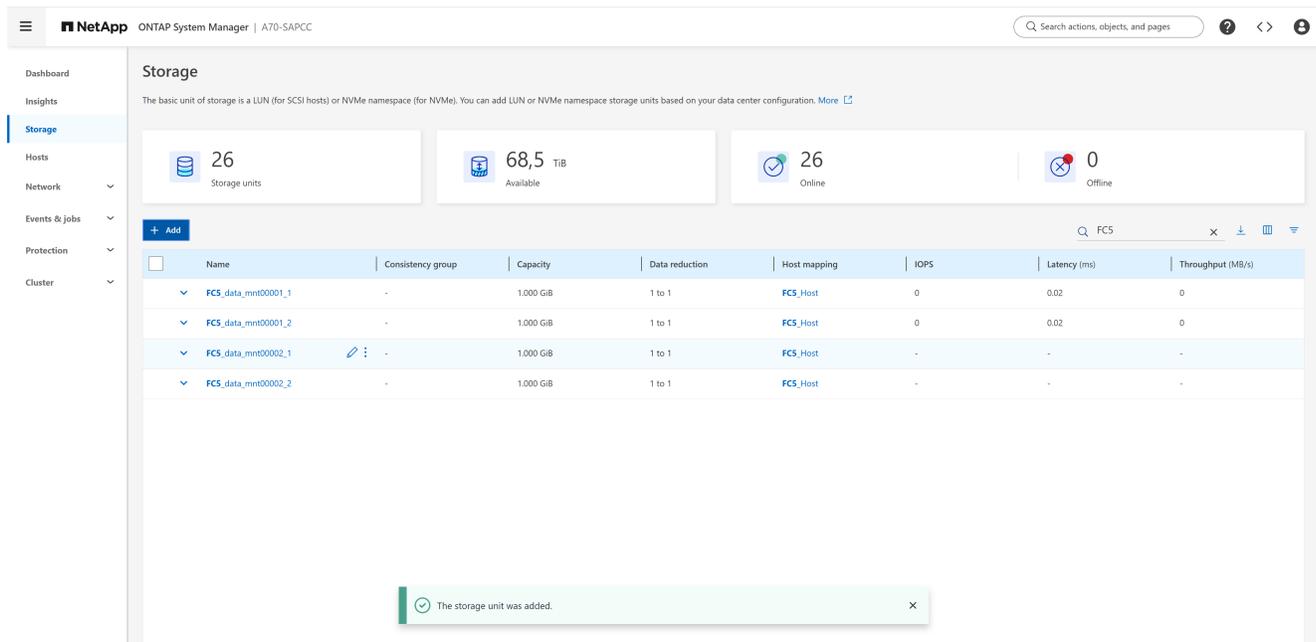
5. Provide the required information:

- the name of the additional data LUNs , e.g. FC5\_data\_mnt00002
- the amount of LUNs to be combined with LVM, e.g. 2
- the size of each LUN, e.g. 1000 GB
- choose Linux as Host Operating system
- choose the previously created mapping FC5\_host for the Host mapping option
- press Add



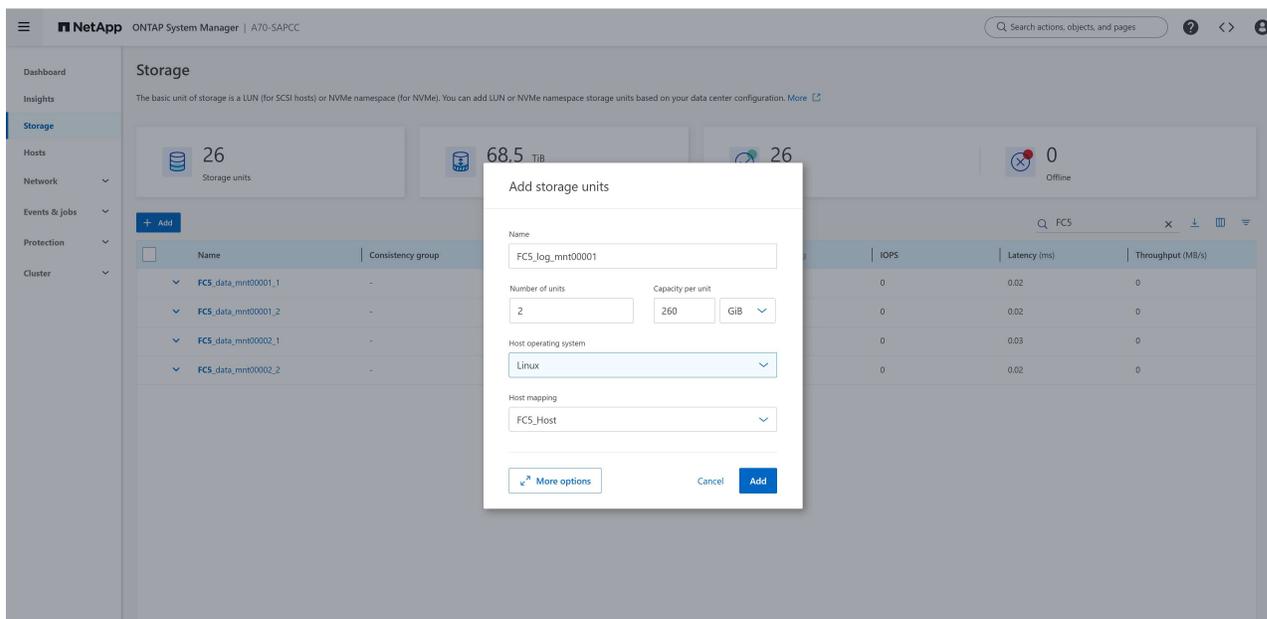
6. Repeat steps 4 and 5 for every additional worker host

7. After successful creation of the data LUNs create the log LUNs by pressing Add

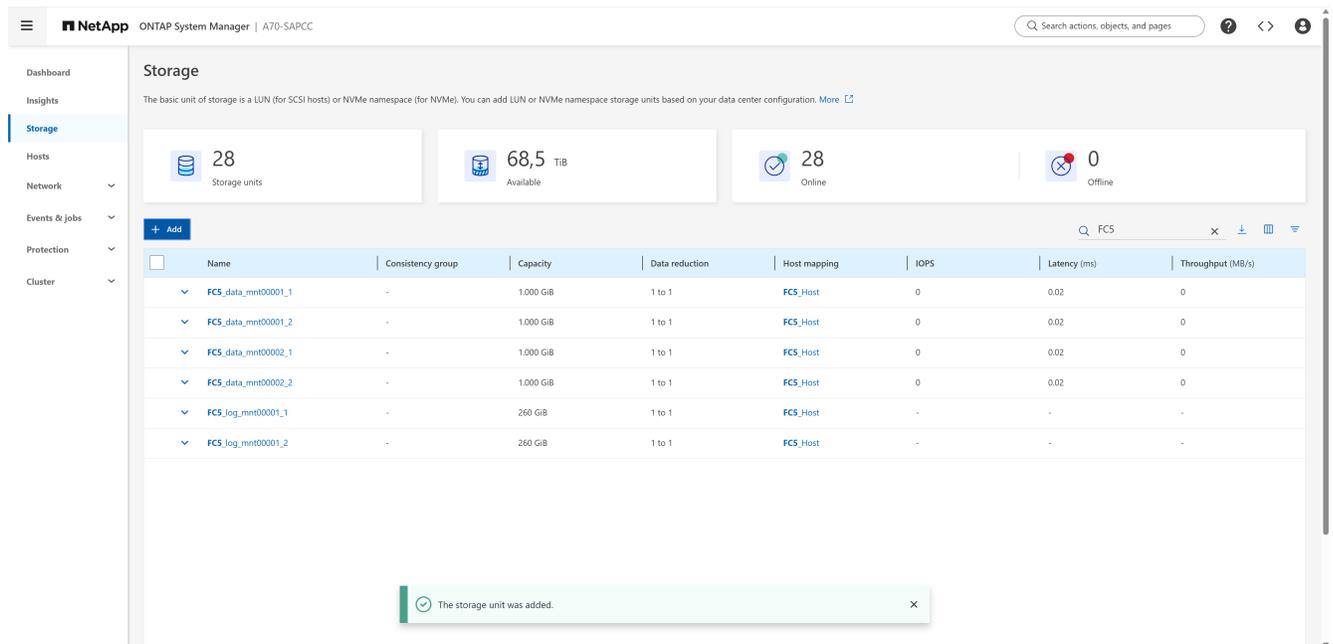


8. Provide the required information:

- the name of the log LUNs to be combined with LVM, e.g. FCS\_log\_mnt00001
- the amount of LUNs to be combined with LVM, e.g. 2
- the size of each LUN, e.g. 260 GB
- choose Linux as Host Operating system
- choose the previously created mapping FCS\_host for the Host mapping option
- press Add

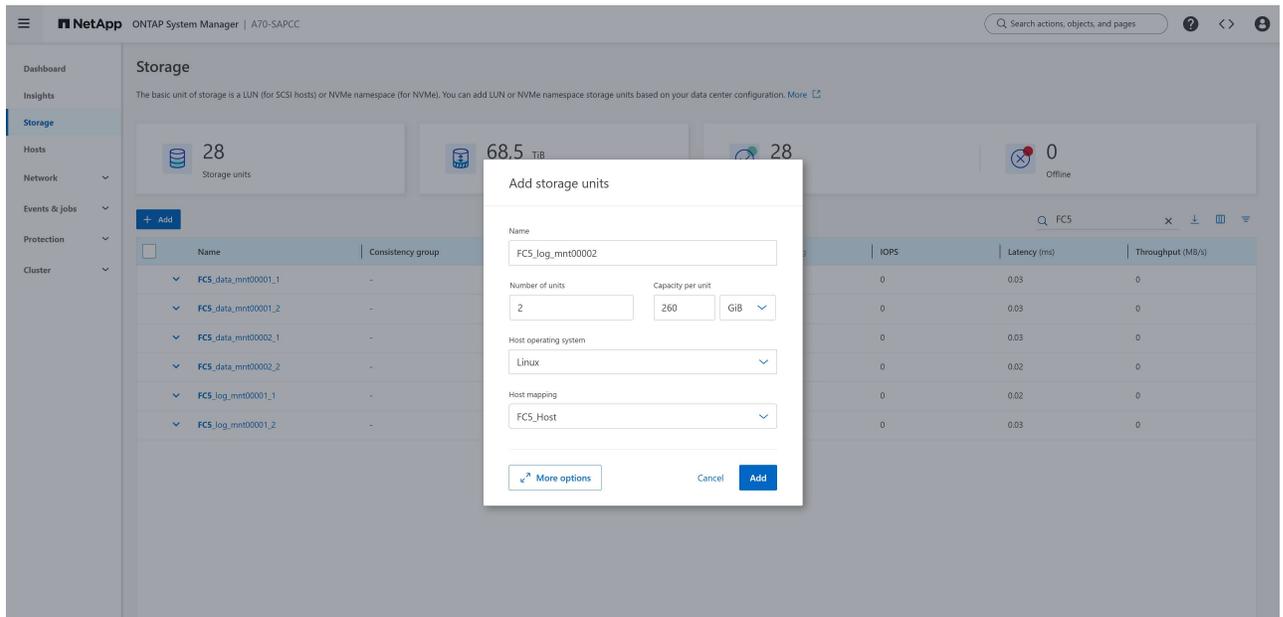


9. Create the log LUNs for the next worker host by pressing Add



10. Provide the required information:

- the name of the additional log LUNs, e.g. FC5\_log\_mnt00002
- the amount of LUNs to be combined with LVM, e.g. 2
- the size of each LUN, e.g. 260 GB
- choose Linux as Host Operating system
- choose the previously created mapping FC5\_host for the Host mapping option
- press Add



11. Repeat steps 9 and 10 for every additional worker host

All required LUNs for a SAP HANA multiple-hosts system have been created.

The screenshot shows the NetApp ONTAP System Manager interface for the Storage section. At the top, there are three summary cards: 'Storage units' (30), 'Available' capacity (68,5 TiB), and 'Online' status (30). Below these is a table of storage units with columns for Name, Consistency group, Capacity, Data reduction, Host mapping, IOPS, Latency (ms), and Throughput (MB/s). A notification at the bottom states 'The storage unit was added.'

Name	Consistency group	Capacity	Data reduction	Host mapping	IOPS	Latency (ms)	Throughput (MB/s)
FCS_data_mnt00001_1	-	1,000 GiB	1 to 1	FCS_Host	0	0.02	0
FCS_data_mnt00001_2	-	1,000 GiB	1 to 1	FCS_Host	0	0.02	0
FCS_data_mnt00002_1	-	1,000 GiB	1 to 1	FCS_Host	0	0.03	0
FCS_data_mnt00002_2	-	1,000 GiB	1 to 1	FCS_Host	0	0.02	0
FCS_log_mnt00001_1	-	260 GiB	1 to 1	FCS_Host	0	0.02	0
FCS_log_mnt00001_2	-	260 GiB	1 to 1	FCS_Host	0	0.03	0
FCS_log_mnt00002_1	-	260 GiB	1 to 1	FCS_Host	-	-	-
FCS_log_mnt00002_2	-	260 GiB	1 to 1	FCS_Host	-	-	-

## SAP HANA storage connector API

A storage connector is required only in multiple-host environments that have failover capabilities. In multiple-host setups, SAP HANA provides high-availability functionality so that an SAP HANA database host can fail over to a standby host.

In this case, the LUNs of the failed host are accessed and used by the standby host. The storage connector is used to make sure that a storage partition can be actively accessed by only one database host at a time.

In SAP HANA multiple-host configurations with NetApp storage, the standard storage connector delivered by SAP is used. The “SAP HANA Fibre Channel Storage Connector Admin Guide” can be found as an attachment to [SAP note 1900823](#).

## Host setup

Before setting up the host, NetApp SAN host utilities must be downloaded from the [NetApp Support](#) site and installed on the HANA servers. The host utility documentation includes information about additional software that must be installed depending on the FCP HBA used.

The documentation also contains information on multipath configurations that are specific to the Linux version used. This document covers the required configuration steps for SLES 12 SP1 or higher and RHEL 7. 2 or later, as described in the [Linux Host Utilities 7.1 Installation and Setup Guide](#).

## Configure multipathing



Steps 1 through 6 must be executed on all worker and standby hosts in an SAP HANA multiple-host configuration.

To configure multipathing, complete the following steps:

1. Run the Linux `rescan-scsi-bus.sh -a` command on each server to discover new LUNs.

- Run the `sanlun lun show` command and verify that all required LUNs are visible. The following example shows the `sanlun lun show` command output for a 2+1 multiple-host HANA system with two data LUNs and two log LUNs. The output shows the LUNs and the corresponding device files, such as LUN FC5\_data\_mnt00001 and the device file `/dev/sdaq`. Each LUN has eight FC paths from the host to the storage controllers.

```

sapcc-hana-tst:~ # sanlun lun show
controller(7mode/E-Series)/                               device
host                lun
vserver(cDOT/FlashRay)    lun-pathname            filename
adapter  protocol  size  product
-----
-----
svm1                FC5_log_mnt00002_2            /dev/sdbb
host21             FCP                500g    cDOT
svm1                FC5_log_mnt00002_1            /dev/sdba
host21             FCP                500g    cDOT
svm1                FC5_log_mnt00001_2            /dev/sdaz
host21             FCP                500g    cDOT
svm1                FC5_log_mnt00001_1            /dev/sday
host21             FCP                500g    cDOT
svm1                FC5_data_mnt00002_2            /dev/sdax
host21             FCP                1t      cDOT
svm1                FC5_data_mnt00002_1            /dev/sdaw
host21             FCP                1t      cDOT
svm1                FC5_data_mnt00001_2            /dev/sdav
host21             FCP                1t      cDOT
svm1                FC5_data_mnt00001_1            /dev/sdau
host21             FCP                1t      cDOT
svm1                FC5_log_mnt00002_2            /dev/sdat
host21             FCP                500g    cDOT
svm1                FC5_log_mnt00002_1            /dev/sdas
host21             FCP                500g    cDOT
svm1                FC5_log_mnt00001_2            /dev/sdar
host21             FCP                500g    cDOT
svm1                FC5_log_mnt00001_1            /dev/sdaq
host21             FCP                500g    cDOT
svm1                FC5_data_mnt00002_2            /dev/sdap
host21             FCP                1t      cDOT
svm1                FC5_data_mnt00002_1            /dev/sdao
host21             FCP                1t      cDOT
svm1                FC5_data_mnt00001_2            /dev/sdan
host21             FCP                1t      cDOT
svm1                FC5_data_mnt00001_1            /dev/sdam
host21             FCP                1t      cDOT
svm1                FC5_log_mnt00002_2            /dev/sdal

```

```

host20      FCP      500g      cDOT
svm1        FC5_log_mnt00002_1      /dev/sdak
host20      FCP      500g      cDOT
svm1        FC5_log_mnt00001_2      /dev/sdaj
host20      FCP      500g      cDOT
svm1        FC5_log_mnt00001_1      /dev/sdai
host20      FCP      500g      cDOT
svm1        FC5_data_mnt00002_2      /dev/sdah
host20      FCP      1t        cDOT
svm1        FC5_data_mnt00002_1      /dev/sdag
host20      FCP      1t        cDOT
svm1        FC5_data_mnt00001_2      /dev/sdaf
host20      FCP      1t        cDOT
svm1        FC5_data_mnt00001_1      /dev/sdae
host20      FCP      1t        cDOT
svm1        FC5_log_mnt00002_2      /dev/sdad
host20      FCP      500g      cDOT
svm1        FC5_log_mnt00002_1      /dev/sdac
host20      FCP      500g      cDOT
svm1        FC5_log_mnt00001_2      /dev/sdab
host20      FCP      500g      cDOT
svm1        FC5_log_mnt00001_1      /dev/sdaa
host20      FCP      500g      cDOT
svm1        FC5_data_mnt00002_2      /dev/sdz
host20      FCP      1t        cDOT
svm1        FC5_data_mnt00002_1      /dev/sdy
host20      FCP      1t        cDOT
svm1        FC5_data_mnt00001_2      /dev/sdx
host20      FCP      1t        cDOT
svm1        FC5_data_mnt00001_1      /dev/sdw
host20      FCP      1t        cDOT

```

3. Run the `multipath -r` and `multipath -ll` command to get the worldwide identifiers (WWIDs) for the device file names.



In this example, there are eight LUNs.

```

sapcc-hana-tst:~ # multipath -r
sapcc-hana-tst:~ # multipath -ll
3600a098038314e63492b59326b4b786d dm-7 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:2 sdaf 65:240 active ready running
  |- 20:0:5:2 sdx 65:112 active ready running

```

```

|- 21:0:4:2 sdav 66:240 active ready running
`- 21:0:6:2 sdan 66:112 active ready running
3600a098038314e63492b59326b4b786e dm-9 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
|- 20:0:4:4 sdah 66:16 active ready running
|- 20:0:5:4 sdz 65:144 active ready running
|- 21:0:4:4 sdax 67:16 active ready running
`- 21:0:6:4 sdap 66:144 active ready running
3600a098038314e63492b59326b4b786f dm-11 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
|- 20:0:4:6 sdaj 66:48 active ready running
|- 20:0:5:6 sdab 65:176 active ready running
|- 21:0:4:6 sdaz 67:48 active ready running
`- 21:0:6:6 sdar 66:176 active ready running
3600a098038314e63492b59326b4b7870 dm-13 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
|- 20:0:4:8 sdal 66:80 active ready running
|- 20:0:5:8 sdad 65:208 active ready running
|- 21:0:4:8 sdbb 67:80 active ready running
`- 21:0:6:8 sdat 66:208 active ready running
3600a098038314e63532459326d495a64 dm-6 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
|- 20:0:4:1 sdae 65:224 active ready running
|- 20:0:5:1 sdw 65:96 active ready running
|- 21:0:4:1 sdau 66:224 active ready running
`- 21:0:6:1 sdam 66:96 active ready running
3600a098038314e63532459326d495a65 dm-8 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
|- 20:0:4:3 sdag 66:0 active ready running
|- 20:0:5:3 sdy 65:128 active ready running
|- 21:0:4:3 sdaw 67:0 active ready running
`- 21:0:6:3 sdao 66:128 active ready running
3600a098038314e63532459326d495a66 dm-10 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active

```

```
|- 20:0:4:5 sdai 66:32 active ready running
|- 20:0:5:5 sdaa 65:160 active ready running
|- 21:0:4:5 sday 67:32 active ready running
`- 21:0:6:5 sdaq 66:160 active ready running
3600a098038314e63532459326d495a67 dm-12 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
|- 20:0:4:7 sdak 66:64 active ready running
|- 20:0:5:7 sdac 65:192 active ready running
|- 21:0:4:7 sdba 67:64 active ready running
`- 21:0:6:7 sdas 66:192 active ready running
```

4. Edit the `/etc/multipath.conf` file and add the WWIDs and alias names.



The example output shows the content of the `/etc/multipath.conf` file, which includes alias names for the four LUNs of a 2+1 multiple-host system. If there is no `multipath.conf` file available, you can create one by running the following command: `multipath -T > /etc/multipath.conf`.

```

sapcc-hana-tst:/ # cat /etc/multipath.conf
multipaths {
    multipath {
        wwid      3600a098038314e63492b59326b4b786d
        alias     svm1-FC5_data_mnt00001_2
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786e
        alias     svm1-FC5_data_mnt00002_2
    }
    multipath {
        wwid      3600a098038314e63532459326d495a64
        alias     svm1-FC5_data_mnt00001_1
    }
    multipath {
        wwid      3600a098038314e63532459326d495a65
        alias     svm1-FC5_data_mnt00002_1
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786f
        alias     svm1-FC5_log_mnt00001_2
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b7870
        alias     svm1-FC5_log_mnt00002_2
    }
    multipath {
        wwid      3600a098038314e63532459326d495a66
        alias     svm1-FC5_log_mnt00001_1
    }
    multipath {
        wwid      3600a098038314e63532459326d495a67
        alias     svm1-FC5_log_mnt00002_1
    }
}

```

5. Run the `multipath -r` command to reload the device map.
6. Verify the configuration by running the `multipath -ll` command to list all the LUNs, alias names, and active and standby paths.



The following example output shows the output of a 2+1 multiple-host HANA system with two data and two log LUNs.

```

sapcc-hana-tst:~ # multipath -ll
svm1-FC5_data_mnt00001_2 (3600a098038314e63492b59326b4b786d) dm-7
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:2 sdaf 65:240 active ready running
  |- 20:0:5:2 sdx 65:112 active ready running
  |- 21:0:4:2 sdav 66:240 active ready running
  `-- 21:0:6:2 sdan 66:112 active ready running
svm1-FC5_data_mnt00002_2 (3600a098038314e63492b59326b4b786e) dm-9
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:4 sdah 66:16 active ready running
  |- 20:0:5:4 sdz 65:144 active ready running
  |- 21:0:4:4 sdax 67:16 active ready running
  `-- 21:0:6:4 sdap 66:144 active ready running
svm1-FC5_data_mnt00001_1 (3600a098038314e63532459326d495a64) dm-6
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:1 sdae 65:224 active ready running
  |- 20:0:5:1 sdw 65:96 active ready running
  |- 21:0:4:1 sdau 66:224 active ready running
  `-- 21:0:6:1 sdam 66:96 active ready running
svm1-FC5_data_mnt00002_1 (3600a098038314e63532459326d495a65) dm-8
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:3 sdag 66:0 active ready running
  |- 20:0:5:3 sdy 65:128 active ready running
  |- 21:0:4:3 sdaw 67:0 active ready running
  `-- 21:0:6:3 sdao 66:128 active ready running
svm1-FC5_log_mnt00001_2 (3600a098038314e63492b59326b4b786f) dm-11
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:6 sdaj 66:48 active ready running
  |- 20:0:5:6 sdab 65:176 active ready running
  |- 21:0:4:6 sdaz 67:48 active ready running
  `-- 21:0:6:6 sdar 66:176 active ready running

```

```

svm1-FC5_log_mnt00002_2 (3600a098038314e63492b59326b4b7870) dm-13
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:8 sdal 66:80 active ready running
  |- 20:0:5:8 sdad 65:208 active ready running
  |- 21:0:4:8 sdbb 67:80 active ready running
  `-- 21:0:6:8 sdat 66:208 active ready running
svm1-FC5_log_mnt00001_1 (3600a098038314e63532459326d495a66) dm-10
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:5 sdai 66:32 active ready running
  |- 20:0:5:5 sdaa 65:160 active ready running
  |- 21:0:4:5 sday 67:32 active ready running
  `-- 21:0:6:5 sdaq 66:160 active ready running
svm1-FC5_log_mnt00002_1 (3600a098038314e63532459326d495a67) dm-12
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:7 sdak 66:64 active ready running
  |- 20:0:5:7 sdac 65:192 active ready running
  |- 21:0:4:7 sdba 67:64 active ready running
  `-- 21:0:6:7 sdas 66:192 active ready running

```

## Single host setup

This chapter describes the setup of an SAP HANA single host.

### LUN configuration for SAP HANA single-host systems

The Linux LVM is being used to increase performance and to address LUN size limitations.

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

Logical volume/LUN	Mount point at SAP HANA host	Note
LV: FC5_data_mnt00001-vol	/hana/data/FC5/mnt00001	Mounted using /etc/fstab entry
LV: FC5_log_mnt00001-vol	/hana/log/FC5/mnt00001	Mounted using /etc/fstab entry
LUN: FC5_shared	/hana/shared/FC5	Mounted using /etc/fstab entry



With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user FC5adm is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN for the `/usr/sap/FC5` directory so that all file systems are on the central storage.

## Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```
pvcreate /dev/mapper/svm1-FC5_data_mnt00001_1
pvcreate /dev/mapper/svm1-FC5_data_mnt00001_2
pvcreate /dev/mapper/svm1-FC5_log_mnt00001_1
pvcreate /dev/mapper/svm1-FC5_log_mnt00001_2
```

2. Create the volume groups for each data and log partition.

```
vgcreate FC5_data_mnt00001 /dev/mapper/svm1-FC5_data_mnt00001_1
/dev/mapper/svm1-FC5_data_mnt00001_2
vgcreate FC5_log_mnt00001 /dev/mapper/svm1-FC5_log_mnt00001_1
/dev/mapper/svm1-FC5_log_mnt00001_2
```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00001
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00001
```

4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```
modprobe dm_mod
pvscan
vgscan
lvscan
```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```
vgchange -a y
```

## Create file systems

Create the XFS file system on all data and log logical volumes and the hana shared LUN.

```
mkfs.xfs FC5_data_mnt00001-vol
mkfs.xfs FC5_log_mnt00001-vol
mkfs.xfs /dev/mapper/svm1-FC5_shared
```



The multiple host example commands show a 2+1 multiple-host HANA system.

## Create mount points

Create the required mount point directories, and set the permissions on the database host:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

## Mount file systems

To mount file systems during system boot using the `/etc/fstab` configuration file, add the required file systems to the `/etc/fstab` configuration file:

```
# cat /etc/fstab
/dev/mapper/hana-FC5_shared /hana/shared xfs defaults 0 0
/dev/mapper/FC5_log_mnt00001-vol /hana/log/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/FC5_data_mnt00001-vol /hana/data/FC5/mnt00001 xfs
relatime,inode64 0 0
```



The XFS file systems for the data and log LUNs must be mounted with the `relatime` and `inode64` mount options.

To mount the file systems, run the `mount -a` command at the host.

## Multiple hosts setup

This chapter describes the setup of a 2+1 SAP HANA multiple host system as example.

## LUN configuration for SAP HANA multiple-hosts systems

The Linux LVM is being used to increase performance and to address LUN size limitations.

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

Logical volume (LV)	Mount point at SAP HANA host	Note
LV: FC5_data_mnt00001-vol	/hana/data/FC5/mnt00001	Mounted using storage connector
LV: FC5_log_mnt00001-vol	/hana/log/FC5/mnt00001	Mounted using storage connector
LV: FC5_data_mnt00002-vol	/hana/data/FC5/mnt00002	Mounted using storage connector
LV: FC5_log_mnt00002-vol	/hana/log/FC5/mnt00002	Mounted using storage connector
External NFS share: FC5_shared	/hana/shared	Mounted at all hosts using NFS and /etc/fstab entry



SAP HANA multiple-host systems require the `/hana/shared` file system connected to all hosts of a system. Usually this is a NFS share provided by an NFS server. It is recommended to use a high available NFS server e.g. such as a NetApp FAS or AFF system. Another option is to use the build-in NFS server of a LINUX host for this.



With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user FC5adm is stored, is on the local disk for each HANA host. In a disaster recovery setup with disk-based replication, NetApp recommends using four additional LUNs for `/usr/sap/FC5` file system each host so that each database host has all its file systems on the central storage.

## Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```
pvcreate /dev/mapper/svm1-FC5_data_mnt00001_1
pvcreate /dev/mapper/svm1-FC5_data_mnt00001_2
pvcreate /dev/mapper/svm1-FC5_data_mnt00002_1
pvcreate /dev/mapper/svm1-FC5_data_mnt00002_2
pvcreate /dev/mapper/svm1-FC5_log_mnt00001_1
pvcreate /dev/mapper/svm1-FC5_log_mnt00001_2
pvcreate /dev/mapper/svm1-FC5_log_mnt00002_1
pvcreate /dev/mapper/svm1-FC5_log_mnt00002_2
```

2. Create the volume groups for each data and log partition.

```

vgcreate FC5_data_mnt00001 /dev/mapper/svm1-FC5_data_mnt00001_1
/dev/mapper/svm1-FC5_data_mnt00001_2
vgcreate FC5_data_mnt00002 /dev/mapper/svm1-FC5_data_mnt00002_1
/dev/mapper/svm1-FC5_data_mnt00002_2
vgcreate FC5_log_mnt00001 /dev/mapper/svm1-FC5_log_mnt00001_1
/dev/mapper/svm1-FC5_log_mnt00001_2
vgcreate FC5_log_mnt00002 /dev/mapper/svm1-FC5_log_mnt00002_1
/dev/mapper/svm1-FC5_log_mnt00002_2

```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```

lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00001
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00001

```

4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```

modprobe dm_mod
pvscan
vgscan
lvscan

```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```

vgchange -a y

```

### Create file systems

Create the XFS file system on all data and log logical volumes.

```

mkfs.xfs FC5_data_mnt00001-vol
mkfs.xfs FC5_data_mnt00002-vol
mkfs.xfs FC5_log_mnt00001-vol
mkfs.xfs FC5_log_mnt00002-vol

```

## Create mount points

Create the required mount point directories, and set the permissions on all worker and standby hosts:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00002
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00002
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

## Mount file systems

To mount the `/hana/shared` file systems during system boot using the `/etc/fstab` configuration file, add the `/hana/shared` file system to the `/etc/fstab` configuration file of each host.

```
sapcc-hana-tst:/ # cat /etc/fstab
<storage-ip>:/hana_shared /hana/shared nfs rw,vers=3,hard,timeo=600,
intr,noatime,nolock 0 0
```



All the data and log file systems are mounted through the SAP HANA storage connector.

To mount the file systems, run the `mount -a` command at each host.

## I/O Stack configuration for SAP HANA

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used.

NetApp conducted performance tests to define the ideal values. The following table lists the optimal values as inferred from the performance tests.

Parameter	Value
<code>max_parallel_io_requests</code>	128
<code>async_read_submit</code>	on
<code>async_write_submit_active</code>	on
<code>async_write_submit_blocks</code>	all

For SAP HANA 1.0 up to SPS12, these parameters can be set during the installation of the SAP HANA database, as described in SAP Note [2267798 – Configuration of the SAP HANA Database during Installation Using `hdbparam`](#).

Alternatively, the parameters can be set after the SAP HANA database installation by using the `hdbparam` framework.

```

FC5adm@sapcc-hana-tst:/usr/sap/FC5/HDB00> hdbparam --paramset
fileio.max_parallel_io_requests=128
FC5adm@sapcc-hana-tst:/usr/sap/FC5/HDB00> hdbparam --paramset
fileio.async_write_submit_active=on
FC5adm@sapcc-hana-tst:/usr/sap/FC5/HDB00> hdbparam --paramset
fileio.async_read_submit=on
FC5adm@sapcc-hana-tst:/usr/sap/FC5/HDB00> hdbparam --paramset
fileio.async_write_submit_blocks=all

```

Starting with SAP HANA 2.0, `hdbparam` is deprecated, and the parameters are moved to the `global.ini` file. The parameters can be set by using SQL commands or SAP HANA Studio. For more details, refer to SAP note [2399079: Elimination of hdbparam in HANA 2](#). The parameters can be also set within the `global.ini` file.

```

FC5adm@sapcc-hana-tst: /usr/sap/FC5/SYS/global/hdb/custom/config> cat
global.ini
...
[fileio]
async_read_submit = on
async_write_submit_active = on
max_parallel_io_requests = 128
async_write_submit_blocks = all
...

```

For SAP HANA 2.0 SPS5 and later, use the `setParameter.py` script to set the correct parameters.

```

fc5adm@sapcc-hana-tst-03:/usr/sap/FC5/HDB00/exe/python_support>
python setParameter.py
-set=SYSTEM/global.ini/fileio/max_parallel_io_requests=128
python setParameter.py -set=SYSTEM/global.ini/fileio/async_read_submit=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_active=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_blocks=all

```

## SAP HANA software installation

This section describes the preparation necessary to install SAP HANA on single-host and multiple-host systems.

### Installation on single-host system

SAP HANA software installation does not require any additional preparation for a single-host system.

## Installation on multiple-host system

Before beginning the installation, create a `global.ini` file to enable use of the SAP storage connector during the installation process. The SAP storage connector mounts the required file systems at the worker hosts during the installation process. The `global.ini` file must be available in a file system that is accessible from all hosts, such as the `/hana/shared` file system.

Before installing SAP HANA software on a multiple-host system, the following steps must be completed:

1. Add the following mount options for the data LUNs and the log LUNs to the `global.ini` file:
  - `relatime` and `inode64` for the data and log file system
2. Add the WWIDs of the data and log partitions. The WWIDs must match the alias names configured in the `/etc/multipath.conf` file.

The following example shows a 2+1 multiple-host setup with `SID=FC5`.

```
sapcc-hana-tst-03:/hana/shared # cat global.ini
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/FC5
basepath_logvolumes = /hana/log/FC5
[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*__prtype = 5
partition_*_data__mountOptions = -o relatime,inode64
partition_*_log__mountOptions = -o relatime,inode64
partition_1_data__lvmname = FC5_data_mnt00001-vol
partition_1_log__lvmname = FC5_log_mnt00001-vol
partition_2_data__lvmname = FC5_data_mnt00002-vol
partition_2_log__lvmname = FC5_log_mnt00002-vol
sapcc-hana-tst-03:/hana/shared #
```

Using the SAP `hdbclm` installation tool, start the installation by running the following command at one of the worker hosts. Use the `addhosts` option to add the second worker (`sapcc-hana-tst-06`) and the standby host (`sapcc-hana-tst-07`).

+



The directory where the prepared `global.ini` file is stored is included with the `storage_cfg` CLI option (`--storage_cfg=/hana/shared`).

+



Depending on the OS version being used, it might be necessary to install Python 2.7 before installing the SAP HANA database.

+

```
./hdblcm --action=install --addhosts=sapcc-hana-tst
-06:role=worker:storage_partition=2,sapcc-hana-tst-07:role=standby
--storage_cfg=/hana/shared/
```

```
AP HANA Lifecycle Management - SAP HANA Database 2.00.073.00.1695288802
*****
```

Scanning software locations...

Detected components:

```
    SAP HANA AFL (incl.PAL,BFL,OFL) (2.00.073.0000.1695321500) in
/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
    SAP HANA Database (2.00.073.00.1695288802) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
    SAP HANA Database Client (2.18.24.1695756995) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_CLIENT_LINUX_X86_64/SAP_HANA_CLIENT/client
    SAP HANA Studio (2.3.75.000000) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_STUDIO_LINUX_X86_64/studio
    SAP HANA Local Secure Store (2.11.0) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HANA_LSS_24_LINUX_X86_64/packages
    SAP HANA XS Advanced Runtime (1.1.3.230717145654) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_RT_10_LINUX_X86_64/packages
    SAP HANA EML AFL (2.00.073.0000.1695321500) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_EML_AFL_10_LINUX_X86_64/packages
    SAP HANA EPM-MDS (2.00.073.0000.1695321500) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/SAP_HANA_EPM-MDS_10/packages
    Automated Predictive Library (4.203.2321.0.0) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/PAAPL4_H20_LINUX_X86_64/apl-
4.203.2321.0-hana2sp03-linux_x64/installer/packages
    GUI for HALM for XSA (including product installer) Version 1 (1.015.0)
in /mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACALMPIUI15_0.zip
    XSAC FILEPROCESSOR 1.0 (1.000.102) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACFILEPROC00_102.zip
    SAP HANA tools for accessing catalog content, data preview, SQL
console, etc. (2.015.230503) in /mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSAC_HRTT_20/XSACHRTT15_230503.zip
    Develop and run portal services for customer applications on XSA
(2.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-
```

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACPORTALSERV07\_0.zip

The SAP Web IDE for HANA 2.0 (4.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSAC\_SAP\_WEB\_IDE\_20/XSACSAPWEBIDE07\_0.zip

XS JOB SCHEDULER 1.0 (1.007.22) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACSERVICES07\_22.zip

SAPUI5 FESV6 XSA 1 - SAPUI5 1.71 (1.071.52) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV671\_52.zip

SAPUI5 FESV9 XSA 1 - SAPUI5 1.108 (1.108.5) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV9108\_5.zip

SAPUI5 SERVICE BROKER XSA 1 - SAPUI5 Service Broker 1.0 (1.000.4) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5SB00\_4.zip

XSA Cockpit 1 (1.001.37) in /mnt/sapcc-share/software/SAP/HANA2SPS7-

73/DATA\_UNITS/XSA\_CONTENT\_10/XSACXSACOCKPIT01\_37.zip

SAP HANA Database version '2.00.073.00.1695288802' will be installed.

Select additional components for installation:

Index	Components	Description
1	all	All components
2	server	No additional components
3	client	Install SAP HANA Database Client version 2.18.24.1695756995
4	lss	Install SAP HANA Local Secure Store version 2.11.0
5	studio	Install SAP HANA Studio version 2.3.75.000000
6	xs	Install SAP HANA XS Advanced Runtime version 1.1.3.230717145654
7	afl	Install SAP HANA AFL (incl.PAL,BFL,OFL) version 2.00.073.0000.1695321500
8	eml	Install SAP HANA EML AFL version 2.00.073.0000.1695321500
9	epmmds	Install SAP HANA EPM-MDS version 2.00.073.0000.1695321500
10	sap_afl_sdk_apl	Install Automated Predictive Library version 4.203.2321.0.0

Enter comma-separated list of the selected indices [3,4]: 2,3

1. Verify that the installation tool installed all selected components at all worker and standby hosts.

## Where to find additional information

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

- [SAP HANA Software Solutions](#)
- [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#)
- [SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- NetApp Documentation Centers

<https://www.netapp.com/support-and-training/documentation/>

- SAP Certified Enterprise Storage Hardware for SAP HANA

<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/>

- SAP HANA Storage Requirements

<https://www.sap.com/documents/2024/03/146274d3-ae7e-0010-bca6-c68f7e60039b.html>

- SAP HANA Tailored Data Center Integration Frequently Asked Questions

<https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html>

- SAP HANA on VMware vSphere Wiki

[https://help.sap.com/docs/SUPPORT\\_CONTENT/virtualization/3362185751.html](https://help.sap.com/docs/SUPPORT_CONTENT/virtualization/3362185751.html)

- SAP HANA on VMware vSphere Best Practices Guide

[https://www.vmware.com/docs/sap\\_hana\\_on\\_vmware\\_vsphere\\_best\\_practices\\_guide-white-paper](https://www.vmware.com/docs/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper)

## Update history

The following technical changes have been made to this solution since its original publication.

Date	Update summary
July 2025	Initial Version

# SAP HANA on NetApp FAS Systems with NFS Configuration Guide

## SAP HANA on NetApp FAS systems with NFS Configuration guide

The NetApp FAS product family has been certified for use with SAP HANA in tailored data center integration (TDI) projects. This guide provides best practices for SAP HANA on this platform with NFS.

Marco Schoen, NetApp

This certification is currently only valid for the following models:

- FAS2750, FAS2820, FAS8300, FAS50, FAS8700, FAS70, FAS9500, FAS90  
A complete list of NetApp certified storage solutions for SAP HANA can be found at the [Certified and Supported SAP HANA Hardware Directory](#).

This document describes the ONTAP configuration requirements for the NFS version 3 (NFSv3) protocol or the NFS version 4 (NFSv4.1) protocol.



Only NFS versions 3 or 4.1 are supported. NFS versions 1, 2, 4.0, and 4.2 aren't supported.



The configuration described in this paper is necessary to achieve the required SAP HANA KPIs and the best performance for SAP HANA. Changing any settings or using features not listed herein might cause performance degradation or unexpected behavior and should only be performed if advised by NetApp support.

The configuration guides for NetApp FAS systems using FCP and for AFF systems using NFS or FC can be found at the following links:

- [SAP HANA on NetApp FAS Systems with FCP](#)
- [SAP HANA on NetApp AFF Systems with NFS](#)
- [SAP HANA on NetApp AFF Systems with FCP](#)
- [SAP HANA on NetApp ASA Systems with FCP](#)

The following table shows the supported combinations for NFS versions, NFS locking, and the required isolation implementations, depending on the SAP HANA database configuration.

For SAP HANA single-host systems or multiple hosts without Host Auto-Failover, NFSv3 and NFSv4 are supported.

For SAP HANA multiple host systems with Host Auto-Failover, NetApp only supports NFSv4, while using NFSv4 locking as an alternative to a server-specific STONITH (SAP HANA HA/DR provider) implementation.

SAP HANA	NFS Version	NFS Locking	SAP HANA HA/DR Provider
SAP HANA single host, multiple hosts without Host Auto-Failover	NFSv3	Off	n/a
	NFSv4	On	n/a

SAP HANA	NFS Version	NFS Locking	SAP HANA HA/DR Provider
SAP HANA multiple hosts with Host Auto-Failover	NFSv3	Off	Server-specific STONITH implementation mandatory
	NFSv4	On	Not required



A server-specific STONITH implementation is not part of this guide. Contact your server vendor for such an implementation.

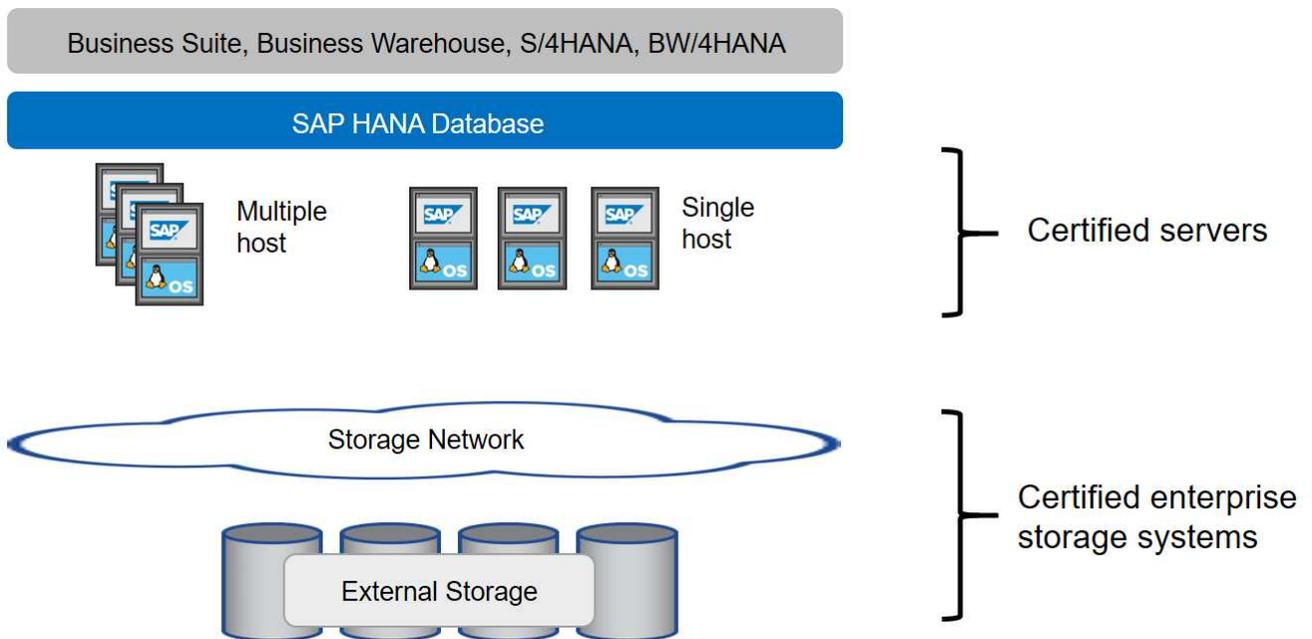
This document covers configuration recommendations for SAP HANA running on physical servers and on virtual servers that use VMware vSphere.



Always refer to the relevant SAP notes for operating system configuration guidelines and HANA-specific Linux kernel dependencies. For more information, see [SAP note 2235581: SAP HANA Supported Operating Systems](#).

### SAP HANA Tailored Data Center integration

NetApp FAS storage controllers are certified in the SAP HANA TDI program using both NFS (NAS) and FC (SAN) protocols. They can be deployed in any of the current SAP HANA scenarios such as SAP Business Suite on HANA, S/4HANA, BW/4HANA, or SAP Business Warehouse on HANA in either single- host or multiple-host configurations. Any server that is certified for use with SAP HANA can be combined with NetApp certified storage solutions. See the following figure for an architecture overview.



For more information regarding the prerequisites and recommendations for production SAP HANA systems, see the following SAP resource:

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

## SAP HANA using VMware vSphere

There are several options to connect the storage to virtual machines (VMs). The preferred one is to connect the storage volumes with NFS directly out of the guest operating system. Using this option, the configuration of hosts and storages do not differ between physical hosts and VMs.

NFS datastores or VVOL datastores with NFS are supported as well. For both options, only one SAP HANA data or log volume must be stored within the datastore for production use cases.

This document describes the recommended setup with direct NFS mounts from the guest OS.

For more information about using vSphere with SAP HANA, see the following links:

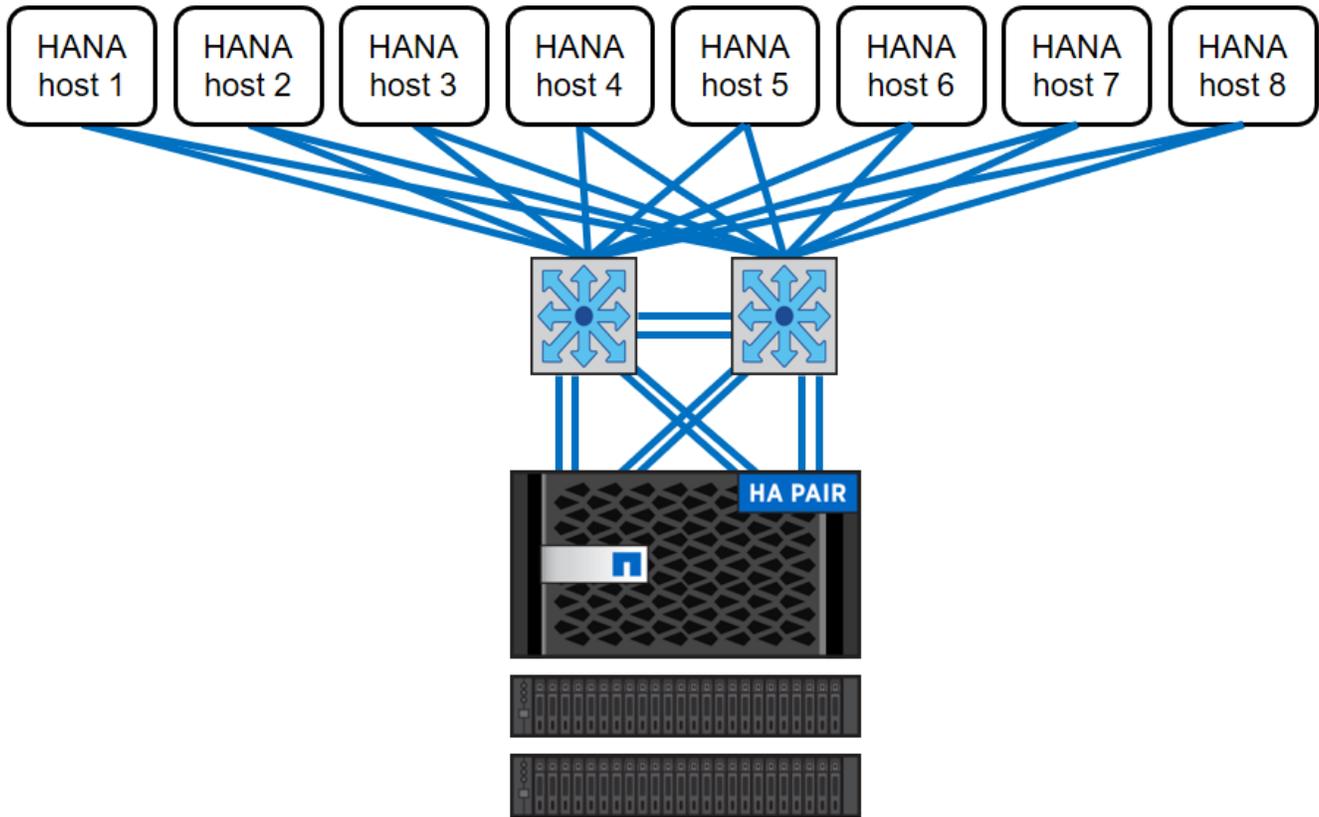
- [SAP HANA on VMware vSphere - Virtualization - Community Wiki](#)
- [SAP HANA on VMware vSphere Best Practices Guide](#)
- [2161991 - VMware vSphere configuration guidelines - SAP ONE Support Launchpad \(Login required\)](#)

## Architecture

SAP HANA hosts are connected to storage controllers by using a redundant 10GbE or faster network infrastructure. Data communication between SAP HANA hosts and storage controllers is based on the NFS protocol.

A redundant switching infrastructure is recommended to provide fault-tolerant SAP HANA host- to- storage connectivity in case of switch or network interface card (NIC) failure. The switches might aggregate individual port performance with port channels in order to appear as a single logical entity at the host level.

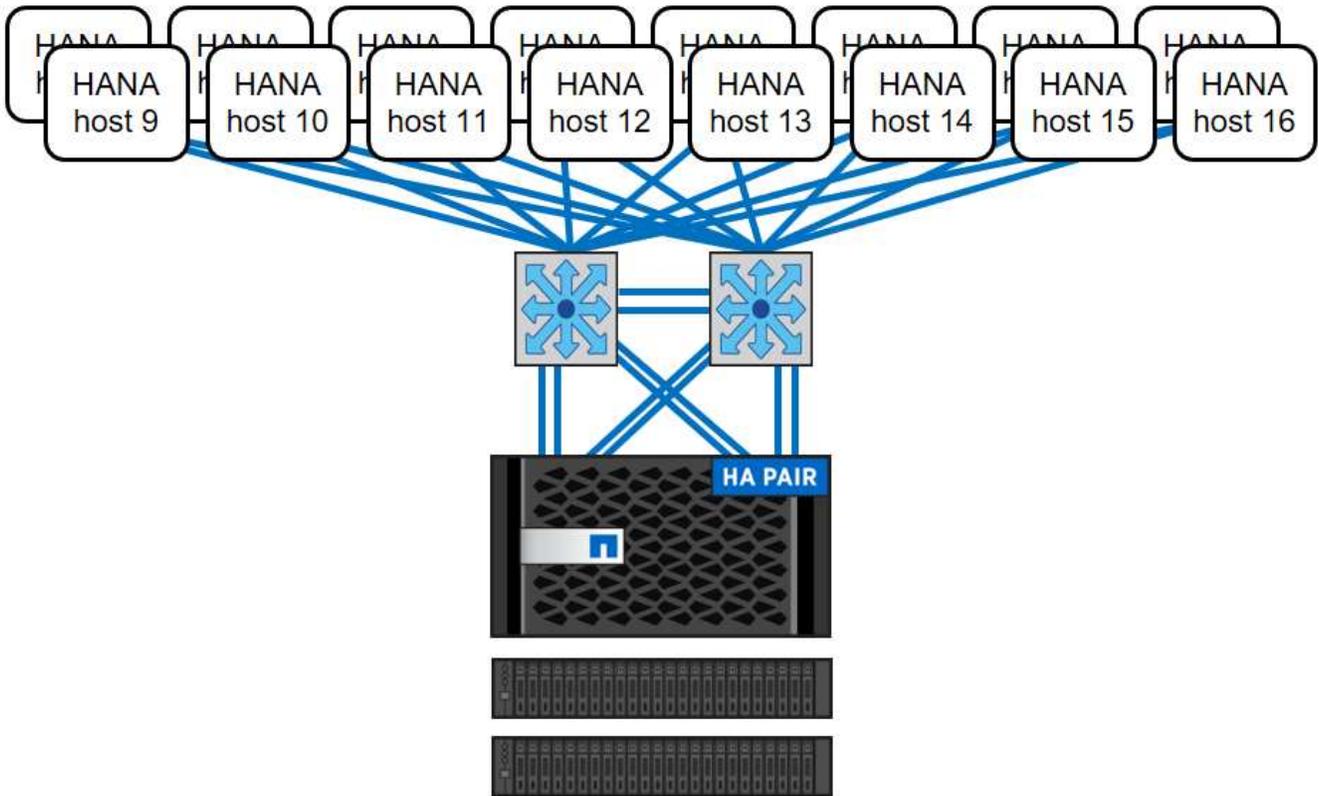
Different models of the FAS system product family can be mixed and matched at the storage layer to allow for growth and differing performance and capacity needs. The maximum number of SAP HANA hosts that can be attached to the storage system is defined by the SAP HANA performance requirements and the model of NetApp controller used. The number of required disk shelves is only determined by the capacity and performance requirements of the SAP HANA systems. The following figure shows an example configuration with eight SAP HANA hosts attached to a storage high availability (HA) pair.



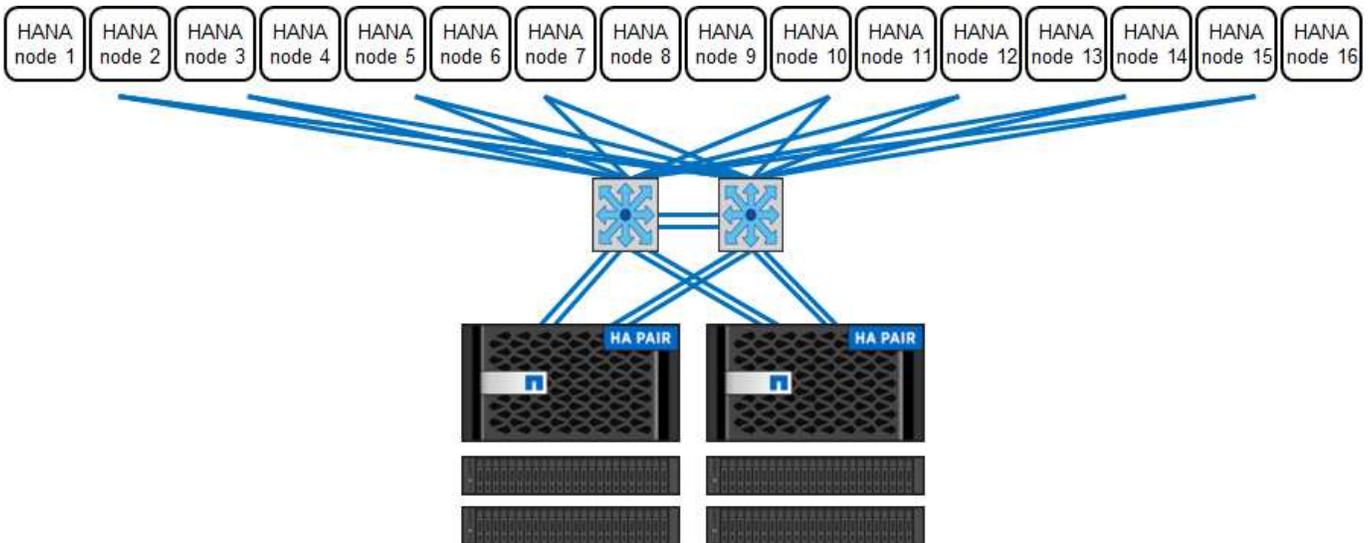
The architecture can be scaled in two dimensions:

- By attaching additional SAP HANA hosts and/or storage capacity to the existing storage, if the storage controllers provide enough performance to meet the current SAP key performance indicators (KPIs)
- By adding more storage systems with additional storage capacity for the additional SAP HANA hosts

The following figure shows an example configuration in which more SAP HANA hosts are attached to the storage controllers. In this example, more disk shelves are necessary to fulfill both the capacity and performance requirements of 16 SAP HANA hosts. Depending on the total throughput requirements, additional 10GbE (or faster) connections to the storage controllers must be added.



Independent of the deployed FAS system, the SAP HANA landscape can also be scaled by adding any of the certified storage controllers to meet the desired node density (the following figure).



### SAP HANA backup

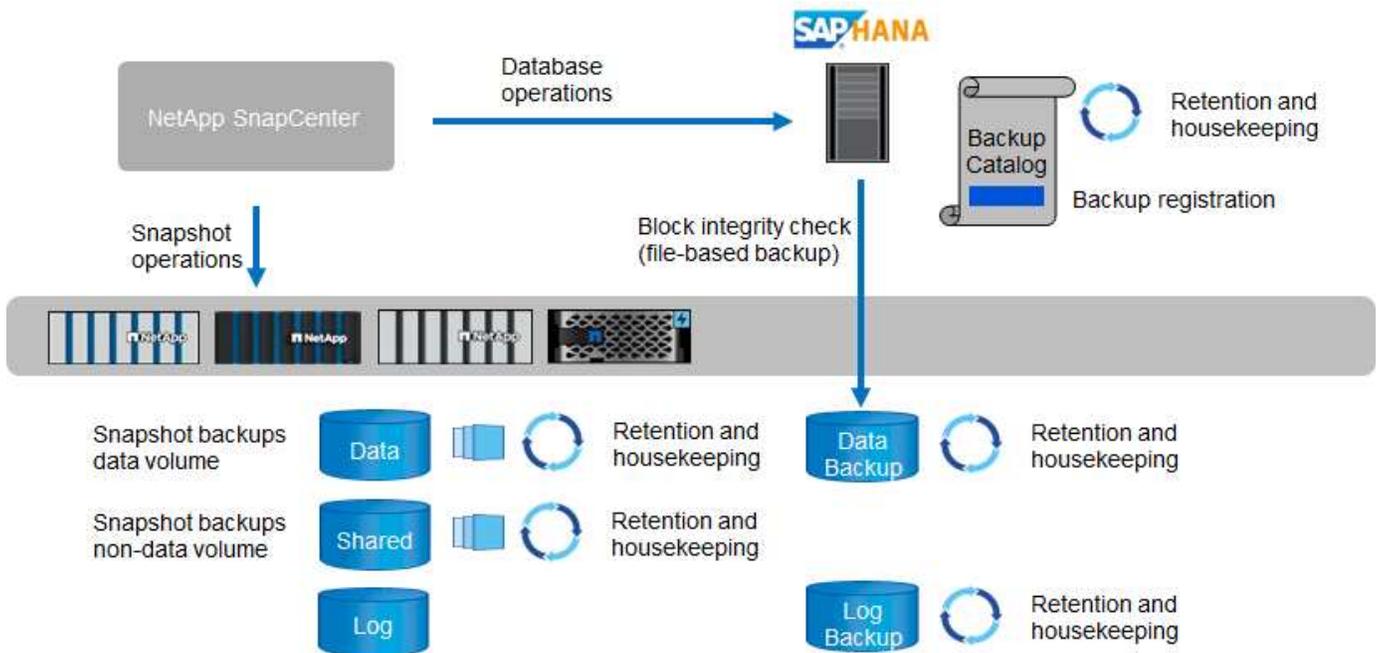
The ONTAP software present on all NetApp storage controllers provides a built-in mechanism to back up SAP HANA databases while in operation with no effect on performance. Storage-based NetApp Snapshot backups are a fully supported and integrated backup solution available for SAP HANA single containers and for SAP HANA Multitenant Database Container (MDC) systems with a single tenant or multiple tenants.

Storage-based Snapshot backups are implemented by using the NetApp SnapCenter plug-in for SAP HANA.

This allows users to create consistent storage-based Snapshot backups by using the interfaces provided natively by SAP HANA databases. SnapCenter registers each of the Snapshot backups into the SAP HANA backup catalog. Therefore, the backups taken by SnapCenter are visible within SAP HANA Studio and Cockpit where they can be selected directly for restore and recovery operations.

NetApp SnapMirror technology allows Snapshot copies that were created on one storage system to be replicated to a secondary backup storage system that is controlled by SnapCenter. Different backup retention policies can then be defined for each of the backup sets on the primary storage and for the backup sets on the secondary storage systems. The SnapCenter Plug-in for SAP HANA automatically manages the retention of Snapshot copy-based data backups and log backups, including the housekeeping of the backup catalog. The SnapCenter Plug-in for SAP HANA also allows the execution of a block integrity check of the SAP HANA database by executing a file-based backup.

The database logs can be backed up directly to the secondary storage by using an NFS mount, as shown in the following figure.



Storage-based Snapshot backups provide significant advantages when compared to conventional file-based backups. These advantages include, but are not limited to, the following:

- Faster backup (a few minutes)
- Reduced recovery time objective (RTO) due to a much faster restore time on the storage layer (a few minutes) as well as more frequent backups
- No performance degradation of the SAP HANA database host, network, or storage during backup and recovery operations
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes

For detailed information about the SAP HANA backup and recovery solution using SnapCenter, see [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

### SAP HANA disaster recovery

SAP HANA disaster recovery can be performed either on the database layer by using SAP HANA system replication or on the storage layer by using storage replication technologies. The following section provides an

overview of disaster recovery solutions based on storage replication.

For detailed information about the SAP HANA disaster recovery solutions, see [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#).

### Storage replication based on SnapMirror

The following figure shows a three-site disaster recovery solution that uses synchronous SnapMirror replication to the local disaster recovery data center and asynchronous SnapMirror to replicate data to the remote disaster recovery data center.

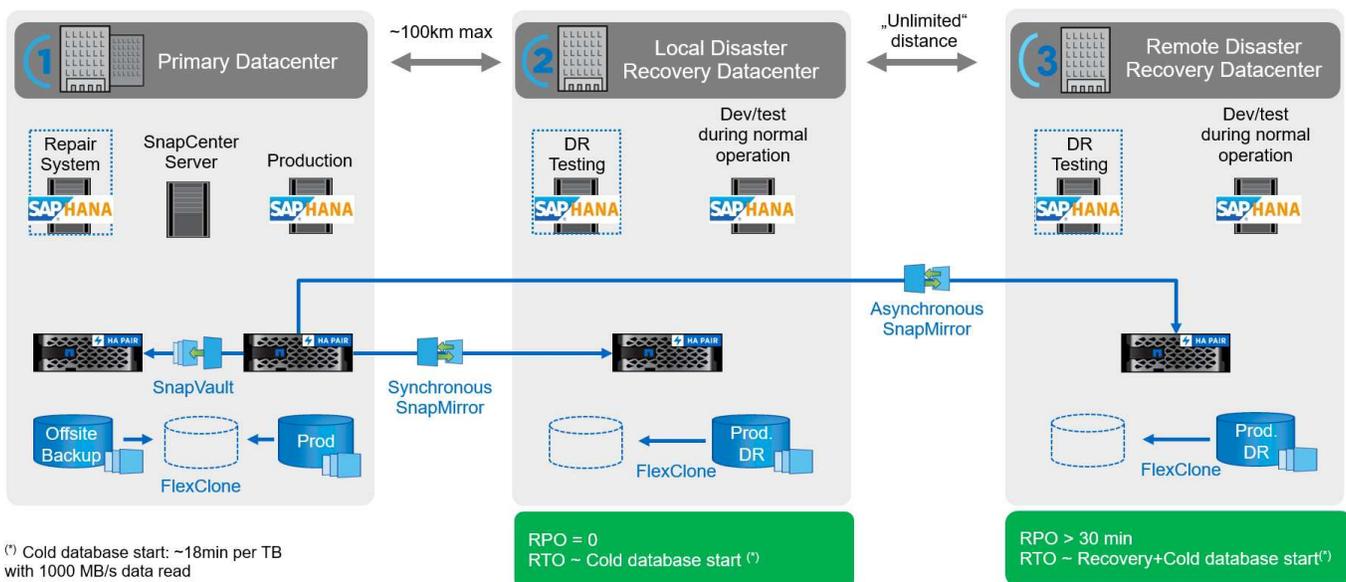
Data replication using synchronous SnapMirror provides an RPO of zero. The distance between the primary and the local disaster recovery data center is limited to around 100km.

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

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

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

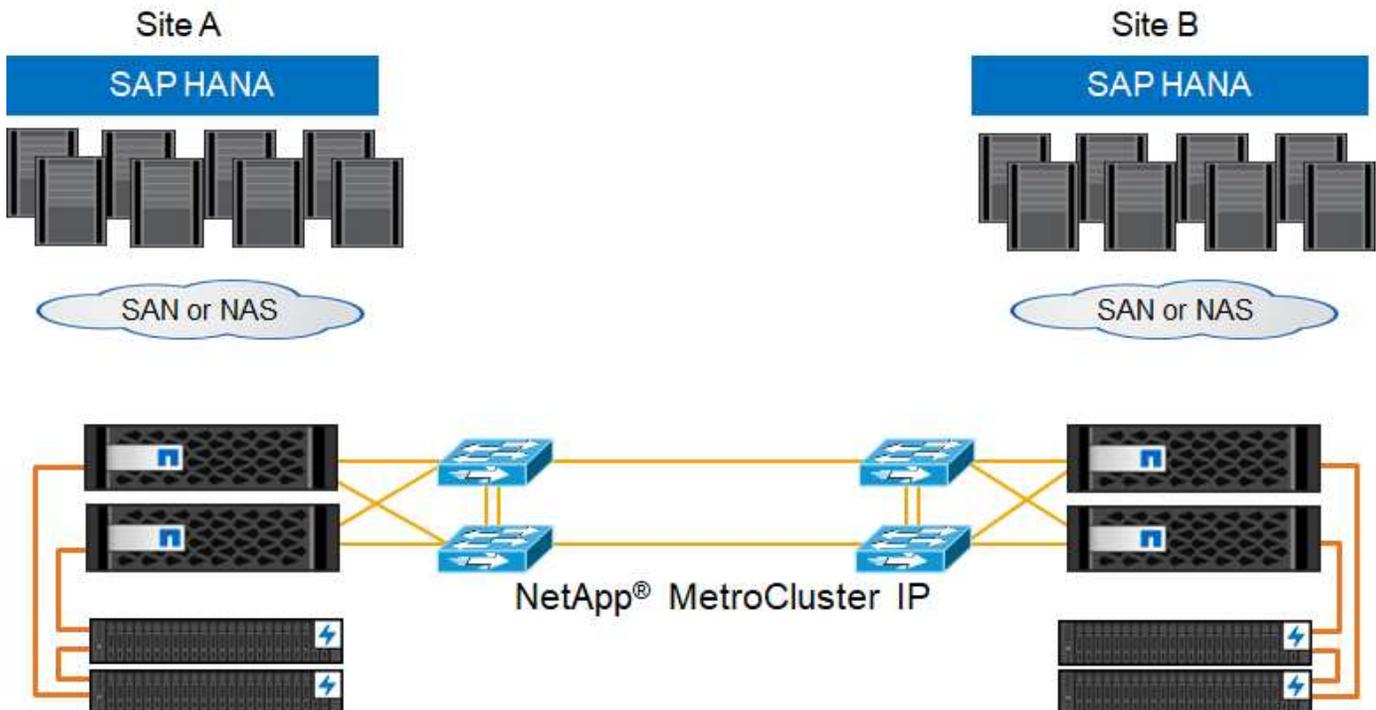
Both replication methods allow to you execute disaster recovery workflow testing without influencing the RPO and RTO. FlexClone volumes are created on the storage and are attached to the disaster recovery testing servers.



Synchronous replication offers StrictSync mode. If the write to secondary storage is not completed for any reason, the application I/O fails, thereby ensuring that the primary and secondary storage systems are identical. Application I/O to the primary resumes only after the SnapMirror relationship returns to InSync status. If the primary storage fails, application I/O can be resumed on the secondary storage after failover, with no loss of data. In StrictSync mode, the RPO is always zero.

## Storage replication based on MetroCluster

The following figure shows a high-level overview of the solution. The storage cluster at each site provides local high availability and is used for the production workload. The data of each site is synchronously replicated to the other location and is available if there is disaster failover.



## Storage sizing

The following section provides an overview of the required performance and capacity considerations needed for sizing a storage system for SAP HANA.



Contact NetApp or your NetApp partner sales representative to assist you in creating a properly sized storage environment.

## Performance considerations

SAP has defined a static set of storage KPIs that are valid for all production SAP HANA environments independent of the memory size of the database hosts and the applications that use the SAP HANA database. These KPIs are valid for single-host, multiple-host, Business Suite on HANA, Business Warehouse on HANA, S/4HANA, and BW/4HANA environments. Therefore, the current performance sizing approach only depends on the number of active SAP HANA hosts that are attached to the storage system.



Storage performance KPIs are only mandated for production SAP HANA systems, but you can implement them in all HANA systems.

SAP delivers a performance test tool used to validate the performance of the storage system for active SAP HANA hosts attached to the storage.

NetApp tested and predefined the maximum number of SAP HANA hosts that can be attached to a specific storage model, while still fulfilling the required storage KPIs from SAP for production-based SAP HANA systems.



The storage controllers of the certified FAS product family can also be used for SAP HANA with other disk types or disk back-end solutions. However, they must be supported by NetApp and fulfill SAP HANA TDI performance KPIs. Examples include NetApp Storage Encryption (NSE) and NetApp FlexArray technology.

This document describes disk sizing for SAS HDDs and solid-state drives (SSDs).

### HDDs

A minimum of 10 data disks (10k RPM SAS) per SAP HANA node is required to fulfill the storage performance KPIs from SAP.



This calculation is independent of the storage controller and disk shelf used as well as the capacity requirements of the database. Adding more disk shelves does not increase the maximum amount of SAP HANA hosts a storage controller can support.

### Solid-state drives

With SSDs, the number of data disks is determined by the SAS connection throughput from the storage controllers to the SSD shelf.

The maximum number of SAP HANA hosts that can be run on a single disk shelf and the minimum number of SSDs required per SAP HANA host were determined by running the SAP performance test tool. This test does not consider the actual storage capacity requirements of the hosts. In addition, you must also calculate the capacity requirements to determine the actual storage configuration needed.

- The 12Gb SAS disk shelf (DS224C) with 24 SSDs supports up to 14 SAP HANA hosts when the disk shelf is connected with 12Gb.
- The 6Gb SAS disk shelf (DS2246) with 24 SSDs supports up to 4 SAP HANA hosts.

The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers.

The following table summarizes the supported number of SAP HANA hosts per disk shelf.

	<b>6Gb SAS shelves (DS2246)fully loaded with 24 SSDs</b>	<b>12Gb SAS shelves (DS224C)fully loaded with 24 SSDs</b>
Maximum number of SAP HANA hosts per disk shelf	4	14



This calculation is independent of the storage controller used. Adding more disk shelves do not increase the maximum amount of SAP HANA hosts a storage controller can support.

### Mixed workloads

SAP HANA and other application workloads running on the same storage controller or in the same storage aggregate are supported. However, it is a NetApp best practice to separate SAP HANA workloads from all other application workloads.

You might decide to deploy SAP HANA workloads and other application workloads on either the same storage controller or the same aggregate. If so, you must make sure that adequate performance is available for SAP HANA within the mixed workload environment. NetApp also recommends that you use quality of service (QoS) parameters to regulate the effect these other applications could have and to guarantee throughput for SAP

HANA applications.

The SAP performance test tool must be used to check if additional SAP HANA hosts can be run on an existing storage controller that is already in use for other workloads. SAP application servers can be safely placed on the same storage controller and/or aggregate as the SAP HANA databases.

### Capacity considerations

A detailed description of the capacity requirements for SAP HANA is in the [SAP Note 1900823](#) attached white paper.



The capacity sizing of the overall SAP landscape with multiple SAP HANA systems must be determined by using SAP HANA storage sizing tools from NetApp. Contact NetApp or your NetApp partner sales representative to validate the storage sizing process for a properly sized storage environment.

### Configuration of performance test tool

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used. These parameters must also be set when storage performance is being tested with the SAP performance test tool.

NetApp conducted performance tests to define the optimal values. The following table lists the parameters that must be set within the configuration file of the SAP performance test tool.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For more information about the configuration of the SAP test tool, see [SAP note 1943937](#) for HWCCT (SAP HANA 1.0) and [SAP note 2493172](#) for HCMT/HCOT (SAP HANA 2.0).

The following example shows how variables can be set for the HCMT/HCOT execution plan.

```
...{
    "Comment": "Log Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "LogAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Data Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "DataAsyncReadSubmit",
    "Value": "on",
```

```

    "Request": "false"
  },
  {
    "Comment": "Log Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "LogAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "DataAsyncWriteSubmitActive",
    "Value": "on",
    "Request": "false"
  },
  {
    "Comment": "Log Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
    "Name": "LogAsyncWriteSubmitBlocks",
    "Value": "all",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
    "Name": "DataAsyncWriteSubmitBlocks",
    "Value": "all",
    "Request": "false"
  },
  {
    "Comment": "Log Volume: Maximum number of parallel I/O requests
per completion queue",
    "Name": "LogExtMaxParallelIoRequests",
    "Value": "128",
    "Request": "false"
  },
  {
    "Comment": "Data Volume: Maximum number of parallel I/O requests
per completion queue",
    "Name": "DataExtMaxParallelIoRequests",
    "Value": "128",
    "Request": "false"
  },
  }, ...

```

These variables must be used for the test configuration. This is usually the case with the predefined execution plans SAP delivers with the HCMT/HCOT tool. The following example for a 4k log write test is from an execution plan.

```
...
  {
    "ID": "D664D001-933D-41DE-A904F304AEB67906",
    "Note": "File System Write Test",
    "ExecutionVariants": [
      {
        "ScaleOut": {
          "Port": "${RemotePort}",
          "Hosts": "${Hosts}",
          "ConcurrentExecution": "${FSConcurrentExecution}"
        },
        "RepeatCount": "${TestRepeatCount}",
        "Description": "4K Block, Log Volume 5GB, Overwrite",
        "Hint": "Log",
        "InputVector": {
          "BlockSize": 4096,
          "DirectoryName": "${LogVolume}",
          "FileOverwrite": true,
          "FileSize": 5368709120,
          "RandomAccess": false,
          "RandomData": true,
          "AsyncReadSubmit": "${LogAsyncReadSubmit}",
          "AsyncWriteSubmitActive":
"${LogAsyncWriteSubmitActive}",
          "AsyncWriteSubmitBlocks":
"${LogAsyncWriteSubmitBlocks}",
          "ExtMaxParallelIoRequests":
"${LogExtMaxParallelIoRequests}",
          "ExtMaxSubmitBatchSize": "${LogExtMaxSubmitBatchSize}",
          "ExtMinSubmitBatchSize": "${LogExtMinSubmitBatchSize}",
          "ExtNumCompletionQueues":
"${LogExtNumCompletionQueues}",
          "ExtNumSubmitQueues": "${LogExtNumSubmitQueues}",
          "ExtSizeKernelIoQueue": "${ExtSizeKernelIoQueue}"
        }
      }, ...
    ]
  }
```

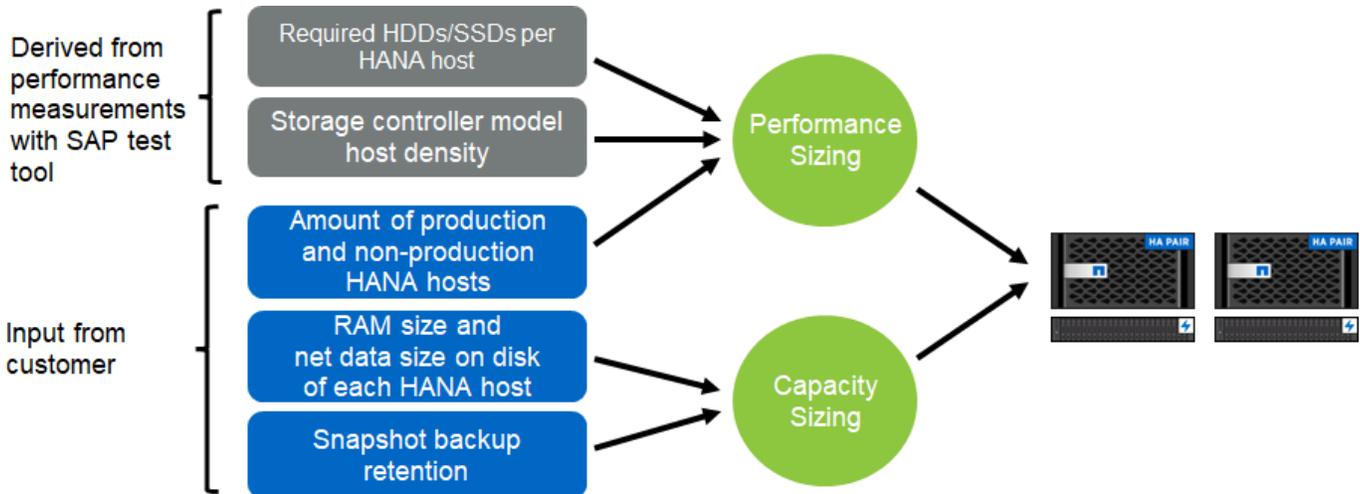
## Storage sizing process overview

The number of disks per HANA host and the SAP HANA host density for each storage model were determined with the SAP performance test tool.

The sizing process requires details such as the number of production and nonproduction SAP HANA hosts, the RAM size of each host, and the backup retention of the storage-based Snapshot copies. The number of SAP HANA hosts determines the storage controller and the number of disks required.

The size of the RAM, net data size on the disk of each SAP HANA host, and the Snapshot copy backup retention period are used as inputs during capacity sizing.

The following figure summarizes the sizing process.



## Infrastructure setup and configuration

### Network setup

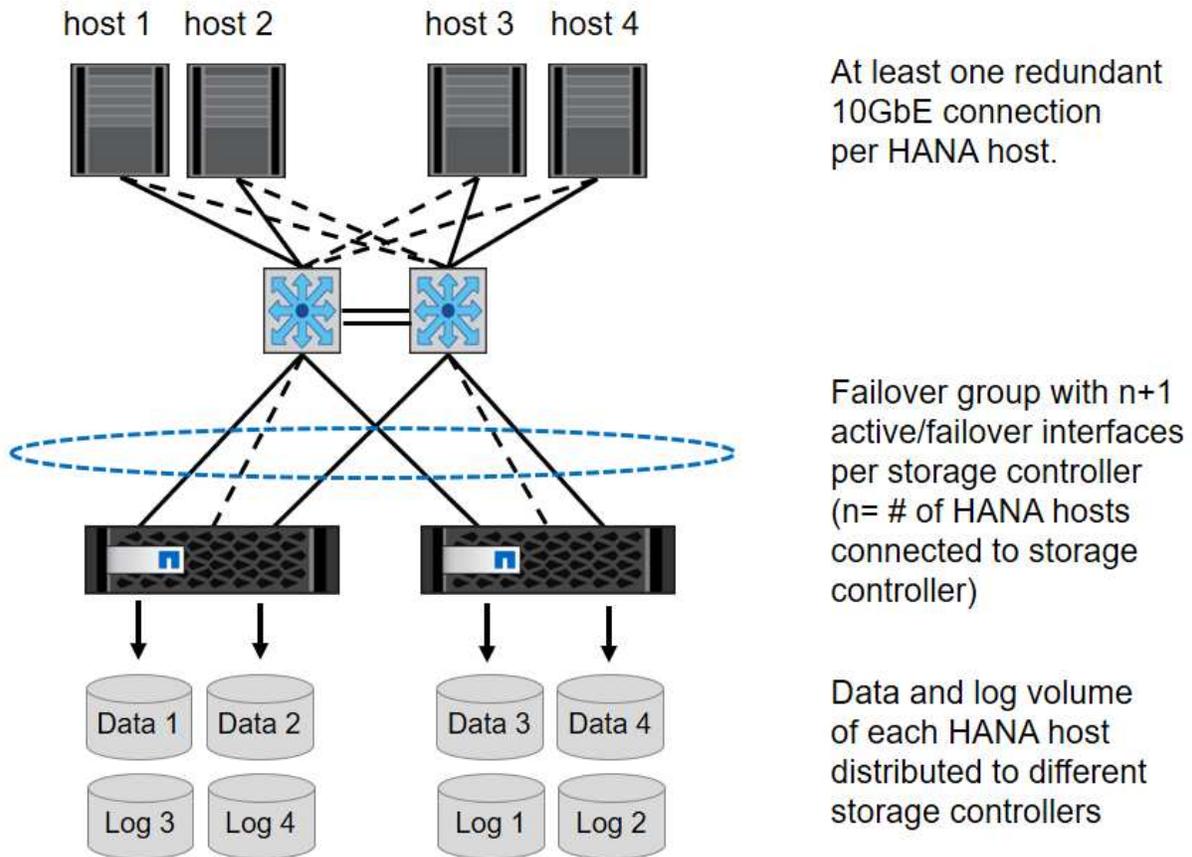
Use the following guidelines when configuring the network:

- A dedicated storage network must be used to connect the SAP HANA hosts to the storage controllers with a 10GbE or faster network.
- Use the same connection speed for storage controllers and SAP HANA hosts. If this is not possible, ensure that the network components between the storage controllers and the SAP HANA hosts are able to handle different speeds. For example, you must provide enough buffer space to allow speed negotiation at the NFS level between storage and hosts. Network components are usually switches, but other components within blade chassis, such as the back plane, must be considered as well.
- Disable flow control on all physical ports used for storage traffic on the storage network switch and host layer.
- Each SAP HANA host must have a redundant network connection with a minimum of 10Gb of bandwidth.
- Jumbo frames with a maximum transmission unit (MTU) size of 9,000 must be enabled on all network components between the SAP HANA hosts and the storage controllers.
- In a VMware setup, dedicated VMXNET3 network adapters must be assigned to each running virtual machine. Check the relevant papers mentioned in the [Introduction](#) for further requirements.
- To avoid interference between each other, use separate network/IO paths for the log and data area.

The following figure shows an example with four SAP HANA hosts attached to a storage controller HA pair using a 10GbE network. Each SAP HANA host has an active-passive connection to the redundant fabric.

At the storage layer, four active connections are configured to provide 10Gb throughput for each SAP HANA host. In addition, one spare interface is configured on each storage controller.

At the storage layer, a broadcast domain with an MTU size of 9000 is configured, and all required physical interfaces are added to this broadcast domain. This approach automatically assigns these physical interfaces to the same failover group. All logical interfaces (LIFs) that are assigned to these physical interfaces are added to this failover group.



In general, it is also possible to use HA interface groups on the servers (bonds) and the storage systems (for example, Link Aggregation Control Protocol [LACP] and ifgroups). With HA interface groups, verify that the load is equally distributed between all interfaces within the group. The load distribution depends on the functionality of the network switch infrastructure.

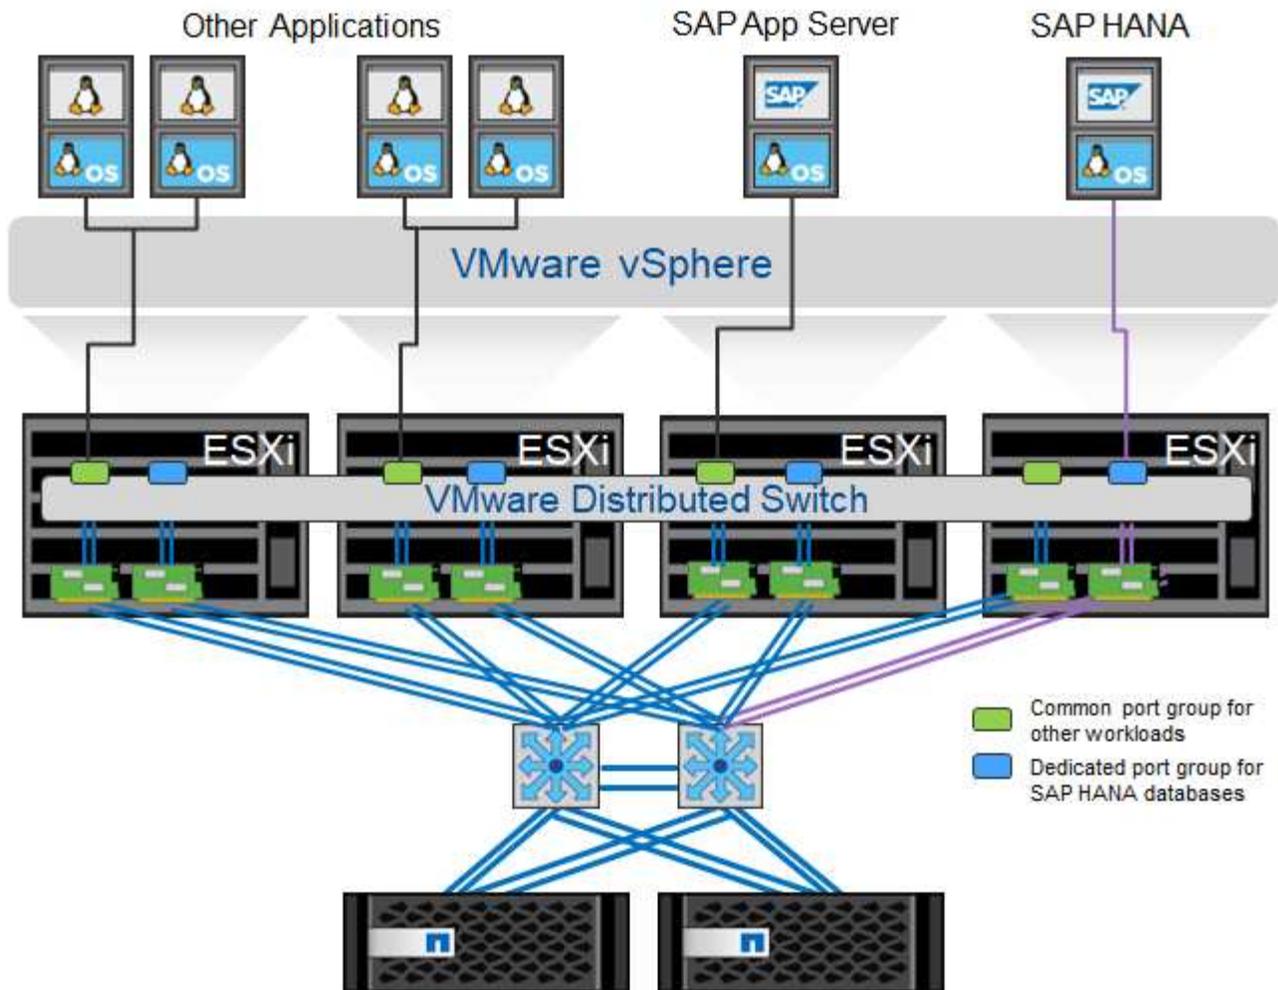


Depending on the number of SAP HANA hosts and the connection speed used, different numbers of active physical ports are needed.

### VMware-specific network setup

Because all data for SAP HANA instances, including performance-critical data and log volumes for the database, is provided through NFS in this solution, proper network design and configuration are crucial. A dedicated storage network is used to separate the NFS traffic from communication and user access traffic between SAP HANA nodes. Each SAP HANA node requires a redundant dedicated network connection with a minimum of 10Gb of bandwidth. Higher bandwidth is also supported. This network must extend end to end from the storage layer through network switching and computing up to the guest operating system hosted on VMware vSphere. In addition to the physical switching infrastructure, a VMware distributed switch (vDS) is used to provide adequate performance and manageability of network traffic at the hypervisor layer.

The following figure provide a network overview.



Each SAP HANA node uses a dedicated port group on the VMware distributed switch. This port group allows for enhanced quality of service (QoS) and dedicated assignment of physical network interface cards (NICs) on the ESX hosts. To use dedicated physical NICs while preserving HA capabilities if there was a NIC failure, the dedicated physical NIC is configured as an active uplink. Additional NICs are configured as standby uplinks in the teaming and failover settings of the SAP HANA port group. In addition, jumbo frames (MTU 9,000) must be enabled end to end on physical and virtual switches. In addition, turn off flow control on all ethernet ports used for storage traffic on servers, switches, and storage systems. The following figure shows an example of such a configuration.



LRO (large receive offload) must be turned off for interfaces used for NFS traffic. For all other network configuration guidelines, see the respective VMware best practices guides for SAP HANA.

t003-HANA-HV1 - Edit Settings

- General
- Advanced
- Security
- Traffic shaping
- VLAN
- Teaming and failover**
- Monitoring
- Traffic filtering and marking
- Miscellaneous

Load balancing:

Network failure detection:

Notify switches:

Failback:

Failover order

↑ ↓

**Active uplinks**

-  dvUplink2

**Standby uplinks**

-  dvUplink1

**Unused uplinks**

### Time synchronization

You must synchronize the time between the storage controllers and the SAP HANA database hosts. To do so, set the same time server for all storage controllers and all SAP HANA hosts.

### Storage controller setup

This section describes the configuration of the NetApp storage system. You must complete the primary installation and setup according to the corresponding ONTAP setup and configuration guides.

### Storage efficiency

Inline deduplication, cross- volume inline deduplication, inline compression, and inline compaction are supported with SAP HANA in an SSD configuration.

Enabling storage efficiency features in an HDD-based configuration is not supported.

### NetApp FlexGroup Volumes

The usage of NetApp FlexGroup Volumes is not supported for SAP HANA. Due to the architecture of SAP HANA the usage of FlexGroup Volumes does not provide any benefit and may results in performance issues.

### NetApp volume and aggregate encryption

The use of NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) are supported with SAP HANA.

### Quality of service

QoS can be used to limit the storage throughput for specific SAP HANA systems or other applications on a shared-use controller. One use case would be to limit the throughput of development and test systems so that

they cannot influence production systems in a mixed setup.

During the sizing process, you should determine the performance requirements of a nonproduction system. Development and test systems can be sized with lower performance values, typically in the range of 20% to 50% of a production-system KPI as defined by SAP.

Starting with ONTAP 9, QoS is configured on the storage volume level and uses maximum values for throughput (MBps) and the amount of I/O (IOPS).

Large write I/O has the biggest performance effect on the storage system. Therefore, the QoS throughput limit should be set to a percentage of the corresponding write SAP HANA storage performance KPI values in the data and log volumes.

### NetApp FabricPool

NetApp FabricPool technology must not be used for active primary file systems in SAP HANA systems. This includes the file systems for the data and log area as well as the `/hana/shared` file system. Doing so results in unpredictable performance, especially during the startup of an SAP HANA system.

Using the “snapshot-only” tiering policy is possible as well as using FabricPool in general at a backup target such as a SnapVault or SnapMirror destination.



Using FabricPool for tiering Snapshot copies at primary storage or using FabricPool at a backup target changes the required time for the restore and recovery of a database or other tasks such as creating system clones or repair systems. Take this into consideration for planning your overall lifecycle- management strategy and check to make sure that your SLAs are still being met while using this function.

FabricPool is a good option for moving log backups to another storage tier. Moving backups affects the time needed to recover an SAP HANA database. Therefore, the option “tiering-minimum-cooling-days” should be set to a value that places log backups, which are routinely needed for recovery, on the local fast storage tier.

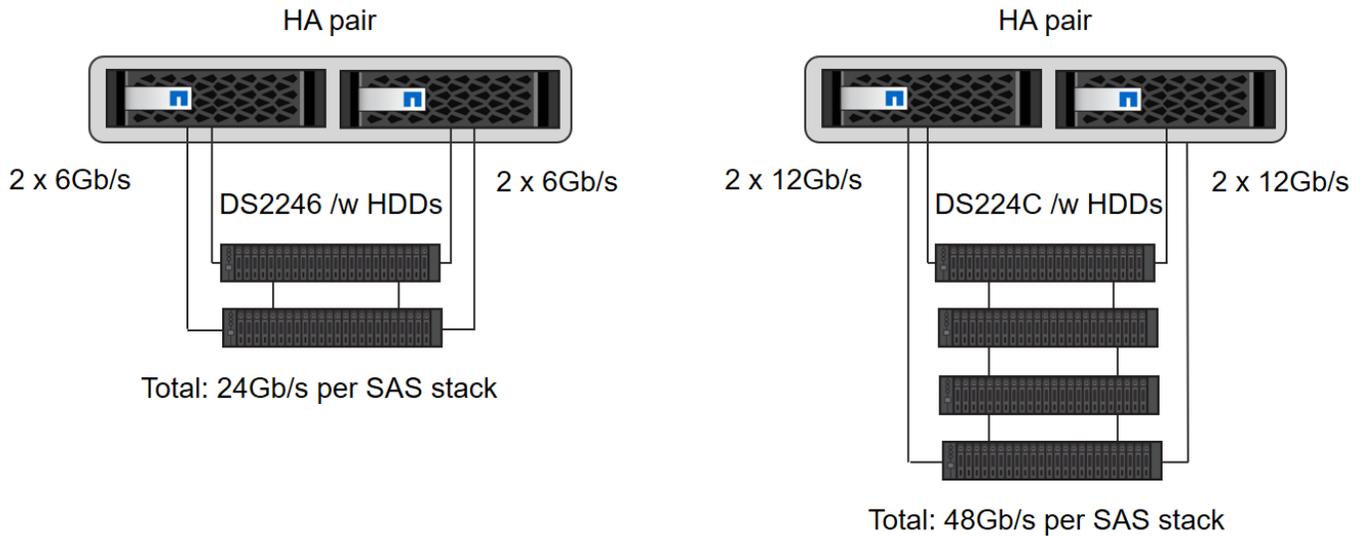
### Storage configuration

The following overview summarizes the required storage configuration steps. Each step is covered in detail in the subsequent sections. In this section, we assume that the storage hardware is set up and that the ONTAP software is already installed. Also, the connections between the storage ports (10GbE or faster) and the network must already be in place.

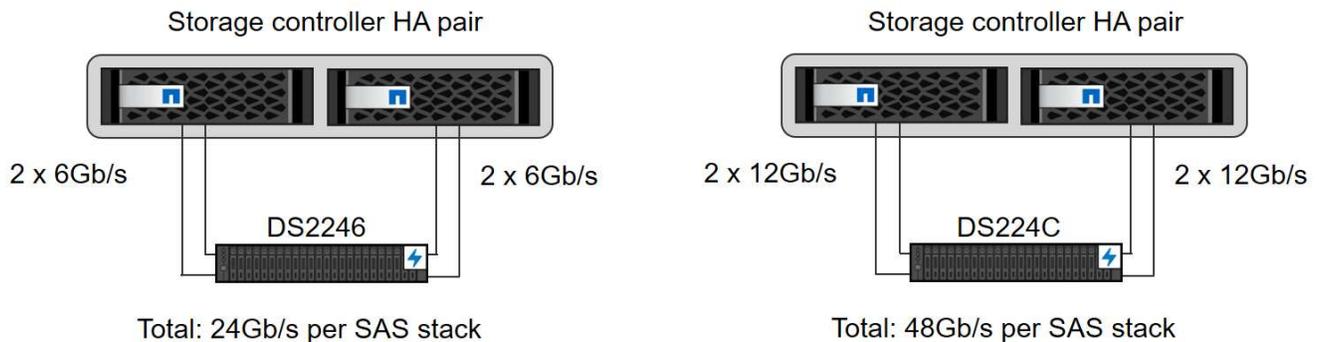
1. Check the correct SAS stack configuration as described in [Disk shelf connection](#).
2. Create and configure the required aggregates as described in [Aggregate configuration](#).
3. Create a storage virtual machine (SVM) as described in [Storage virtual machine configuration](#).
4. Create LIFs as described in [Logical interface configuration](#).
5. Create volumes within the aggregates as described in [Volume configuration for SAP HANA single-host systems](#) and [Volume configuration for SAP HANA multiple-host systems](#).
6. Set the required volume options as described in [Volume options](#).
7. Set the required options for NFSv3 as described in [NFS configuration for NFSv3](#) or for NFSv4 as described in [NFS configuration for NFSv4](#).
8. Mount the volumes to namespace and set export policies as described in [Mount volumes to namespace and set export policies](#).

## Disk shelf connection

With HDDs, a maximum of two DS2246 disk shelves or four DS224C disk shelves can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair.



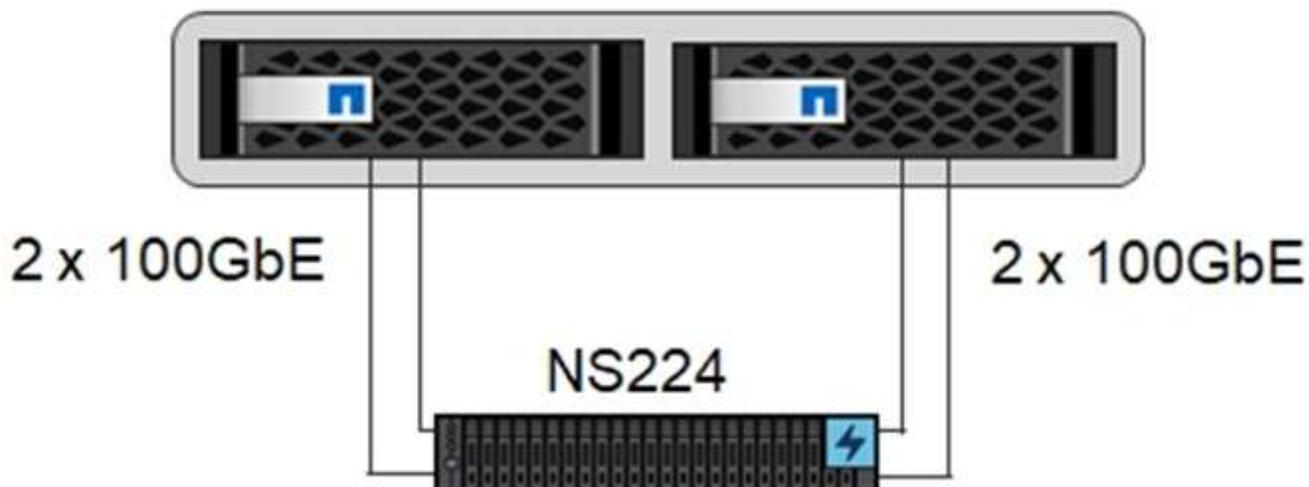
With SSDs, a maximum of one disk shelf can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair. With the DS224C disk shelf, quad-path SAS cabling can also be used, but is not required.



## NVMe (100GbE) disk shelves

Each NS224 NVMe disk shelf is connected with two 100GbE ports per controller, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair.

## Storage controller HA pair

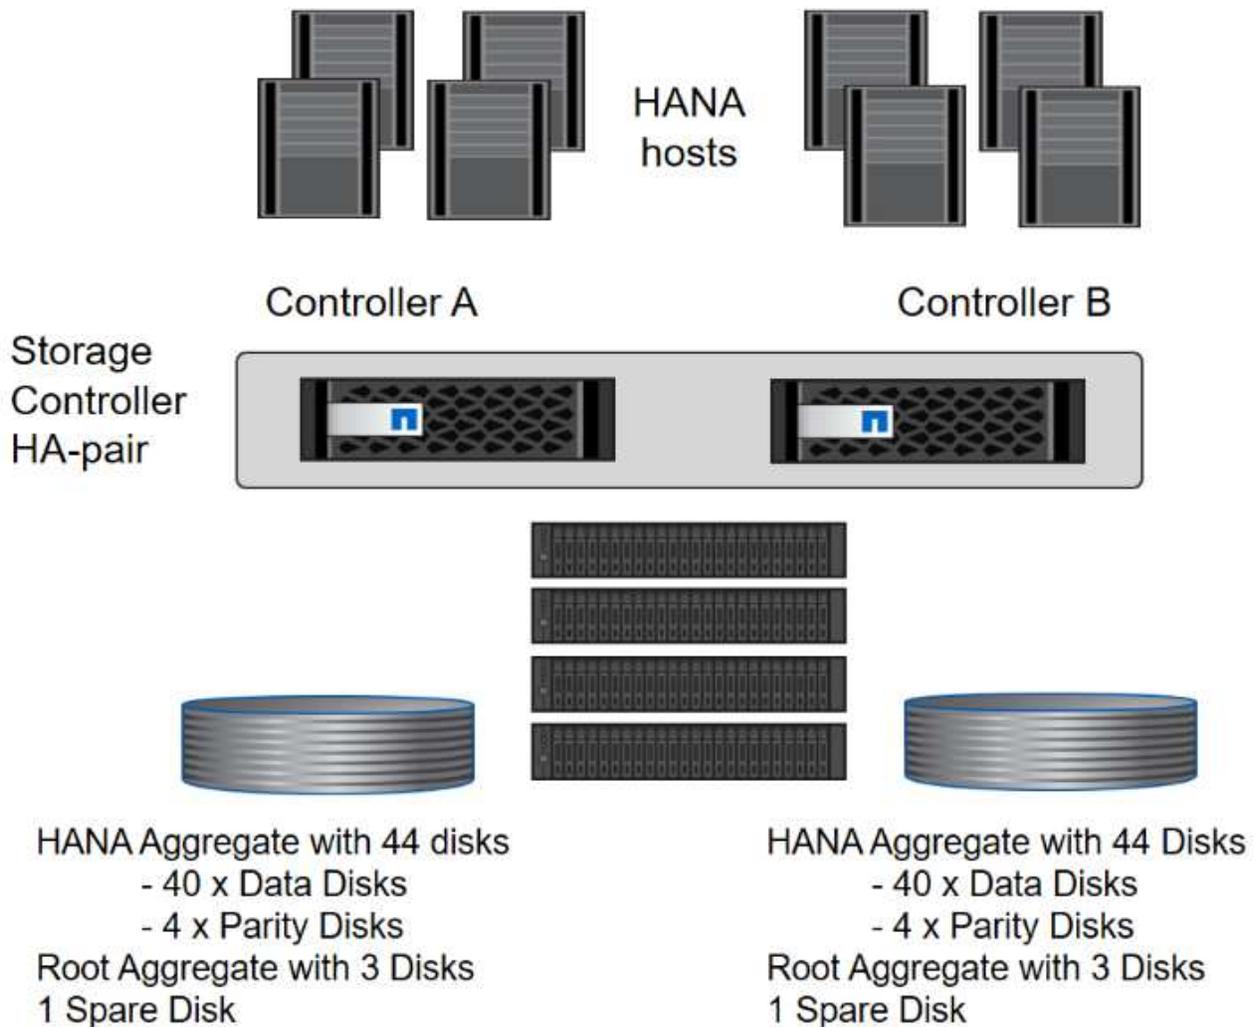


### Aggregate configuration

In general, you must configure two aggregates per controller, independent of the disk shelf or drive technology (SSD or HDD) that is used. For FAS2000 series systems, one data aggregate is enough.

### Aggregate configuration with HDDs

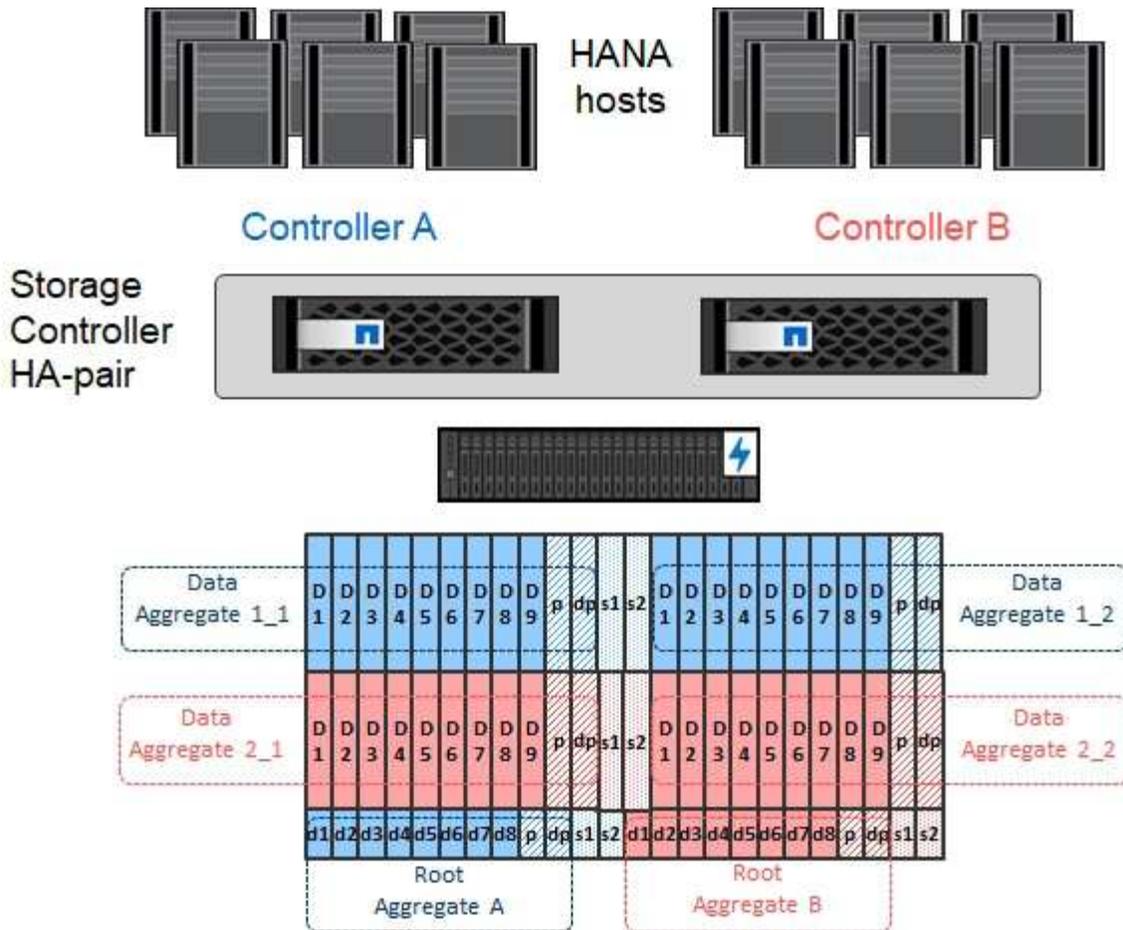
The following figure shows a configuration for eight SAP HANA hosts. Four SAP HANA hosts are attached to each storage controller. Two separate aggregates, one at each storage controller, are configured. Each aggregate is configured with  $4 \times 10 = 40$  data disks (HDDs).



### Aggregate configuration with SDD-only systems

In general, you must configure two aggregates per controller, independent of which disk shelf or disk technology (SSDs or HDDs) is used. For FAS2000 series systems, one data aggregate is enough.

The following figure shows a configuration of 12 SAP HANA hosts running on a 12Gb SAS shelf configured with ADPv2. Six SAP HANA hosts are attached to each storage controller. Four separate aggregates, two at each storage controller, are configured. Each aggregate is configured with 11 disks with nine data and two parity disk partitions. For each controller, two spare partitions are available.



### Storage virtual machine configuration

Multiple SAP landscapes with SAP HANA databases can use a single SVM. An SVM can also be assigned to each SAP landscape, if necessary, in case they are managed by different teams within a company.

If a QoS profile was automatically created and assigned during new SVM creation, remove the automatically created profile from the SVM to provide the required performance for SAP HANA:

```
vserver modify -vserver <svm-name> -qos-policy-group none
```

### Logical interface configuration

For SAP HANA production systems, you must use different LIFs for mounting the data volume and the log volume from the SAP HANA host. Therefore at least two LIFs are required.

The data and log volume mounts of different SAP HANA hosts can share a physical storage network port by using either the same LIFs or by using individual LIFs for each mount.

The maximum number of data and log volume mounts per physical interface are shown in the following table.

Ethernet port speed	10GbE	25GbE	40GbE	100GbE
Maximum number of log or data volume mounts per physical port	3	8	12	30



Sharing one LIF between different SAP HANA hosts might require a remount of data or log volumes to a different LIF. This change avoids performance penalties if a volume is moved to a different storage controller.

Development and test systems can use more data and volume mounts or LIFs on a physical network interface.

For production, development, and test systems, the `/hana/shared` file system can use the same LIF as the data or log volume.

### Volume configuration for SAP HANA single-host systems

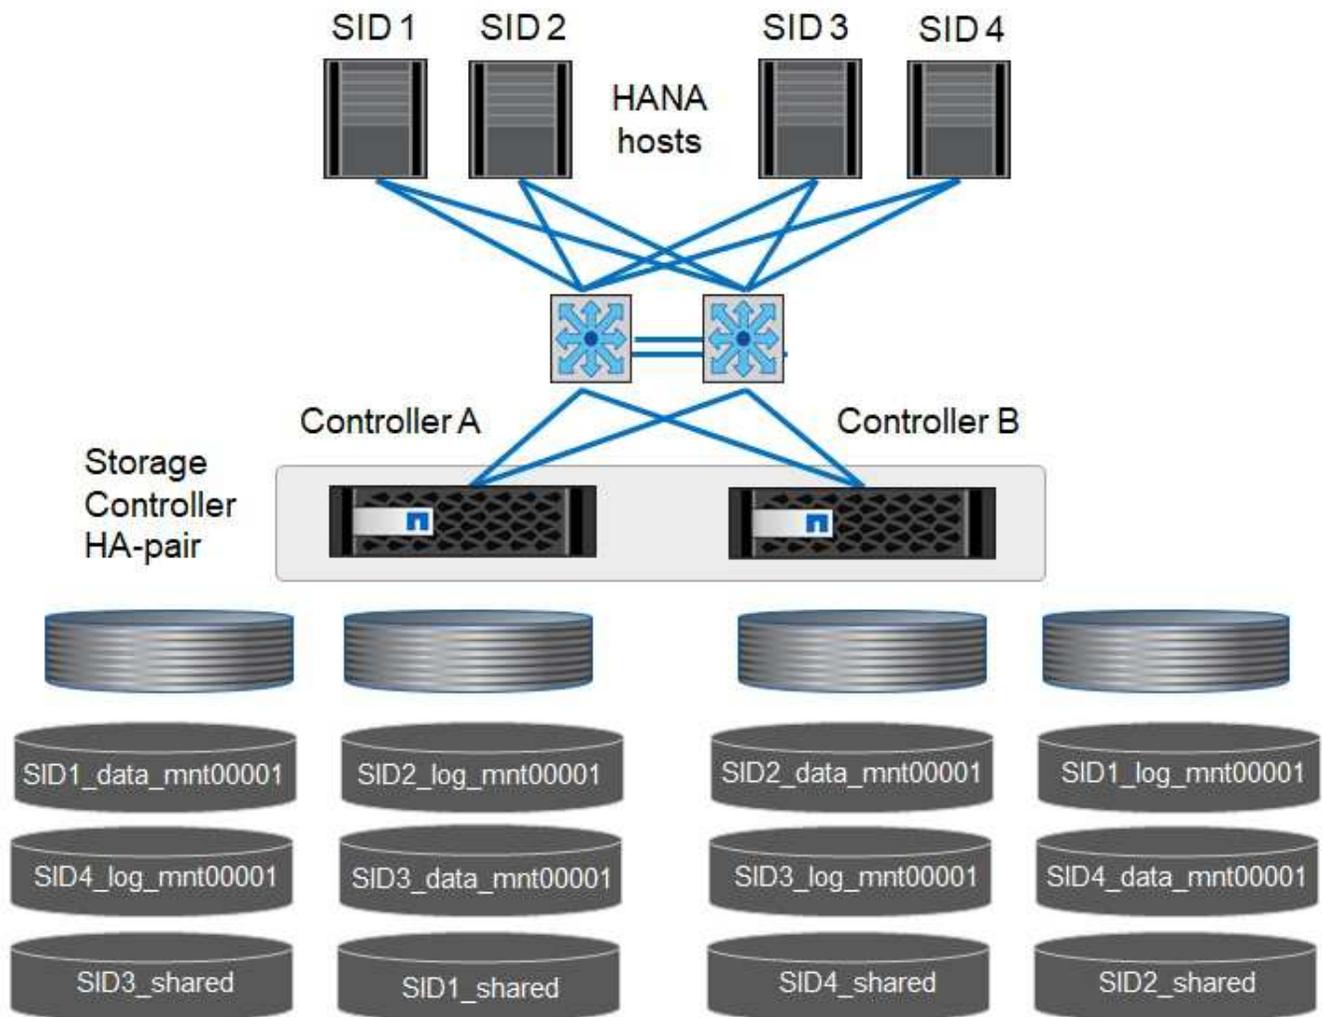
The following figure shows the volume configuration of four single-host SAP HANA systems. The data and log volumes of each SAP HANA system are distributed to different storage controllers. For example, volume `SID1_data_mnt00001` is configured on controller A, and volume `SID1_log_mnt00001` is configured on controller B.



If only one storage controller of an HA pair is used for the SAP HANA systems, data and log volumes can also be stored on the same storage controller.



If the data and log volumes are stored on the same controller, access from the server to the storage must be performed with two different LIFs: one LIF to access the data volume and one to access the log volume.



For each SAP HANA DB host, a data volume, a log volume, and a volume for `/hana/shared` are configured. The following table shows an example configuration for single-host SAP HANA systems.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller b
Data, log, and shared volumes for system SID1	Data volume: SID1_data_mnt00001	Shared volume: SID1_shared	–	Log volume: SID1_log_mnt00001
Data, log, and shared volumes for system SID2	–	Log volume: SID2_log_mnt00001	Data volume: SID2_data_mnt00001	Shared volume: SID2_shared
Data, log, and shared volumes for system SID3	Shared volume: SID3_shared	Data volume: SID3_data_mnt00001	Log volume: SID3_log_mnt00001	–
Data, log, and shared volumes for system SID4	Log volume: SID4_log_mnt00001	–	Shared volume: SID4_shared	Data volume: SID4_data_mnt00001

The following table shows an example of the mount point configuration for a single-host system. To place the home directory of the `sidadm` user on the central storage, the `/usr/sap/SID` file system should be mounted

from the SID\_shared volume.

Junction Path	Directory	Mount point at HANA host
SID_data_mnt00001	–	/hana/data/SID/mnt00001
SID_log_mnt00001	–	/hana/log/SID/mnt00001
SID_shared	usr-sap shared	/usr/sap/SID /hana/shared

#### Volume configuration for SAP HANA multiple-host systems

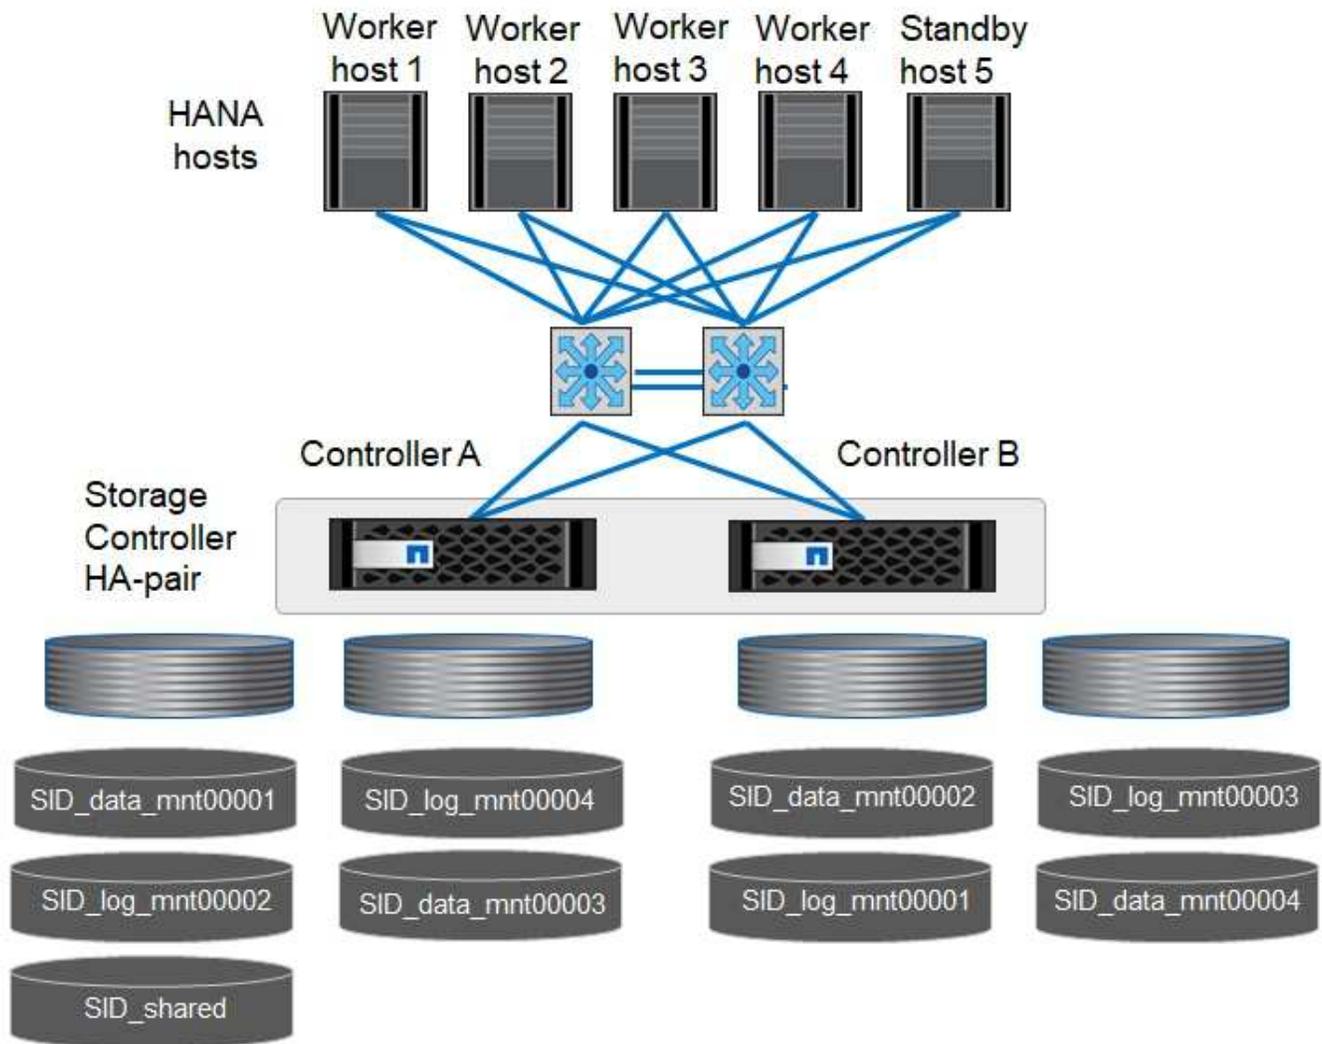
The following figure shows the volume configuration of a 4+1 SAP HANA system. The data and log volumes of each SAP HANA host are distributed to different storage controllers. For example, volume SID1\_data1\_mnt00001 is configured on controller A, and volume SID1\_log1\_mnt00001 is configured on controller B.



If only one storage controller of an HA pair is used for the SAP HANA system, the data and log volumes can also be stored on the same storage controller.



If the data and log volumes are stored on the same controller, access from the server to the storage must be performed with two different LIFs: one to access the data volume and one to access the log volume.



For each SAP HANA host, a data volume and a log volume are created. The `/hana/shared` volume is used by all hosts of the SAP HANA system. The following table shows an example configuration for a multiple-host SAP HANA system with four active hosts.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	–	Log volume: SID_log_mnt00001	–
Data and log volumes for node 2	Log volume: SID_log_mnt00002	–	Data volume: SID_data_mnt00002	–
Data and log volumes for node 3	–	Data volume: SID_data_mnt00003	–	Log volume: SID_log_mnt00003
Data and log volumes for node 4	–	Log volume: SID_log_mnt00004	–	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared	–	–	–

The following table shows the configuration and the mount points of a multiple-host system with four active SAP HANA hosts. To place the home directories of the `sidadm` user of each host on the central storage, the

/usr/sap/SID file systems are mounted from the SID\_shared volume.

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00001	–	/hana/data/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001	–	/hana/log/SID/mnt00001	Mounted at all hosts
SID_data_mnt00002	–	/hana/data/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002	–	/hana/log/SID/mnt00002	Mounted at all hosts
SID_data_mnt00003	–	/hana/data/SID/mnt00003	Mounted at all hosts
SID_log_mnt00003	–	/hana/log/SID/mnt00003	Mounted at all hosts
SID_data_mnt00004	–	/hana/data/SID/mnt00004	Mounted at all hosts
SID_log_mnt00004	–	/hana/log/SID/mnt00004	Mounted at all hosts
SID_shared	shared	/hana/shared/	Mounted at all hosts
SID_shared	usr-sap-host1	/usr/sap/SID	Mounted at host 1
SID_shared	usr-sap-host2	/usr/sap/SID	Mounted at host 2
SID_shared	usr-sap-host3	/usr/sap/SID	Mounted at host 3
SID_shared	usr-sap-host4	/usr/sap/SID	Mounted at host 4
SID_shared	usr-sap-host5	/usr/sap/SID	Mounted at host 5

### Volume options

You must verify and set the volume options listed in the following table on all SVMs. For some of the commands, you must switch to the advanced privilege mode within ONTAP.

Action	Command
Disable visibility of Snapshot directory	vol modify -vserver <vserver-name> -volume <volname> -snapdir-access false
Disable automatic Snapshot copies	vol modify -vserver <vserver-name> -volume <volname> -snapshot-policy none
Disable access time update except of the SID_shared volume	set advanced vol modify -vserver <vserver-name> -volume <volname> -atime-update false set admin

### NFS configuration for NFSv3

The NFS options listed in the following table must be verified and set on all storage controllers.

For some of the commands shown, you must switch to the advanced privilege mode within ONTAP.

Action	Command
Enable NFSv3	nfs modify -vserver <vserver-name> v3.0 enabled

Action	Command
Set NFS TCP maximum transfer size to 1MB	set advanced nfs modify -vserver <vserver_name> -tcp-max-xfer -size 1048576 set admin



In shared environments with different workloads set the max NFS TCP transfer size to 262144

#### NFS configuration for NFSv4

The NFS options listed in the following table must be verified and set on all SVMs.

For some of the commands, you must switch to the advanced privilege mode within ONTAP.

Action	Command
Enable NFSv4	nfs modify -vserver <vserver-name> -v4.1 enabled
Set NFS TCP maximum transfer size to 1MB	set advanced nfs modify -vserver <vserver_name> -tcp-max-xfer -size 1048576 set admin
Disable NFSv4 access control lists (ACLs)	nfs modify -vserver <vserver_name> -v4.1-acl disabled
Set NFSv4 domain ID	nfs modify -vserver <vserver_name> -v4-id-domain <domain-name>
Disable NFSv4 read delegation	nfs modify -vserver <vserver_name> -v4.1-read -delegation disabled
Disable NFSv4 write delegation	nfs modify -vserver <vserver_name> -v4.1-write -delegation disabled
Disable NFSv4 numeric ids	nfs modify -vserver <vserver_name> -v4-numeric-ids disabled
Change amount of NFSv4.x session slots optional	set advanced nfs modify -vserver hana -v4.x-session-num-slots <value> set admin



In shared environments with different workloads set the max NFS TCP transfer size to 262144



Please note that disabling numbering ids requires user management as described in [SAP HANA installation preparations for NFSv4](#).



The NFSv4 domain ID must be set to the same value on all Linux servers (`/etc/idmapd.conf`) and SVMs, as described in [SAP HANA installation preparations for NFSv4](#).



pNFS can be enabled and used.

If SAP HANA multiple-host systems with host auto-failover are being used, the failover parameters need to be adjusted within `nameserver.ini` as shown in the following table. Keep the default retry interval of 10 seconds within these sections.

Section within <code>nameserver.ini</code>	Parameter	Value
failover	<code>normal_retries</code>	9
distributed_watchdog	<code>deactivation_retries</code>	11
distributed_watchdog	<code>takeover_retries</code>	9

### Mount volumes to namespace and set export policies

When a volume is created, the volume must be mounted to the namespace. In this document, we assume that the junction path name is the same as the volume name. By default, the volume is exported with the default policy. The export policy can be adapted if required.

### Host setup

All the steps described in this section are valid for both SAP HANA environments on physical servers and for SAP HANA running on VMware vSphere.

### Configuration parameter for SUSE Linux Enterprise Server

Additional kernel and configuration parameters at each SAP HANA host must be adjusted for the workload generated by SAP HANA.

### SUSE Linux Enterprise Server 12 and 15

Starting with SUSE Linux Enterprise Server (SLES) 12 SP1, the kernel parameter must be set in a configuration file in the `/etc/sysctl.d` directory. For example, a configuration file with the name `91-NetApp-HANA.conf` must be created.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
sunrpc.tcp_max_slot_table_entries = 128
```



Saptune, which is included in SLES for SAP OS versions, can be used to set these values. See [SAP Note 3024346](#) (requires SAP login).

## Configuration parameter for Red Hat Enterprise Linux 7.2 or later

You must adjust additional kernel and configuration parameters at each SAP HANA host for the workload generated by SAP HANA.

Starting with Red Hat Enterprise Linux 7.2, you must set the kernel parameters in a configuration file in the `/etc/sysctl.d` directory. For example, a configuration file with the name `91-NetApp-HANA.conf` must be created.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
sunrpc.tcp_max_slot_table_entries = 128
```



Since Red Hat Enterprise Linux version 8.6, these settings can also be applied by using RHEL System Roles for SAP (Ansible). See [SAP Note 3024346](#) (requires SAP login).

## Create subdirectories in `/hana/shared` volume



The examples show an SAP HANA database with `SID=NF2`.

To create the required subdirectories, take one of the following actions:

- For a single-host system, mount the `/hana/shared` volume and create the `shared` and `usr-sap` subdirectories.

```
sapcc-hana-tst-06:/mnt # mount <storage-hostname>:/NF2_shared /mnt/tmp
sapcc-hana-tst-06:/mnt # cd /mnt/tmp
sapcc-hana-tst-06:/mnt/tmp # mkdir shared
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap
sapcc-hana-tst-06:/mnt/tmp # cd ..
sapcc-hana-tst-06:/mnt # umount /mnt/tmp
```

- For a multiple-host system, mount the `/hana/shared` volume and create the `shared` and the `usr-sap` subdirectories for each host.

The example commands show a 2+1 multiple-host HANA system.

```

sapcc-hana-tst-06:/mnt # mount <storage-hostname>:/NF2_shared /mnt/tmp
sapcc-hana-tst-06:/mnt # cd /mnt/tmp
sapcc-hana-tst-06:/mnt/tmp # mkdir shared
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host1
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host2
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host3
sapcc-hana-tst-06:/mnt/tmp # cd ..
sapcc-hana-tst-06:/mnt # umount /mnt/tmp

```

### Create mount points



The examples show an SAP HANA database with SID=NF2.

To create the required mount point directories, take one of the following actions:

- For a single-host system, create mount points and set the permissions on the database host.

```

sapcc-hana-tst-06:/ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-06:/ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-06:/ # mkdir -p /hana/shared
sapcc-hana-tst-06:/ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/shared
sapcc-hana-tst-06:/ # chmod -R 777 /usr/sap/NF2

```

- For a multiple-host system, create mount points and set the permissions on all worker and standby hosts.

The following example commands are for a 2+1 multiple-host HANA system.

- First worker host:

```

sapcc-hana-tst-06:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-06:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-06:~ # mkdir -p /hana/shared
sapcc-hana-tst-06:~ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-06:~ # chmod -R 777 /usr/sap/NF2

```

- Second worker host:

```

sapcc-hana-tst-07:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-07:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-07:~ # mkdir -p /hana/shared
sapcc-hana-tst-07:~ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-07:~ # chmod -R 777 /usr/sap/NF2

```

- Standby host:

```

sapcc-hana-tst-08:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-08:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-08:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-08:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-08:~ # mkdir -p /hana/shared
sapcc-hana-tst-08:~ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-08:~ # chmod -R 777 /usr/sap/NF2

```

### Mount file systems

Different mount options are used depending on the NFS version and ONTAP release. The following file systems must be mounted to the hosts:

- /hana/data/SID/mnt0000\*
- /hana/log/SID/mnt0000\*
- /hana/shared
- /usr/sap/SID

The following table shows the NFS versions that must be used for the different file systems for single-host and multiple-host SAP HANA databases.

File systems	SAP HANA single host	SAP HANA multiple hosts
/hana/data/SID/mnt0000*	NFSv3 or NFSv4	NFSv4
/hana/log/SID/mnt0000*	NFSv3 or NFSv4	NFSv4
/hana/shared	NFSv3 or NFSv4	NFSv3 or NFSv4
/usr/sap/SID	NFSv3 or NFSv4	NFSv3 or NFSv4

The following table shows the mount options for the various NFS versions and ONTAP releases. The common parameters are independent of the NFS and ONTAP versions.



SAP LaMa requires the /usr/sap/SID directory to be local. Therefore, do not mount an NFS volume for /usr/sap/SID if you are using SAP LaMa.

For NFSv3, you must switch off NFS locking to avoid NFS lock cleanup operations if there is a software or server failure.

With ONTAP 9, the NFS transfer size can be configured up to 1MB. Specifically, with 40GbE or faster connections to the storage system, you must set the transfer size to 1MB to achieve the expected throughput values.

Common parameter	NFSv3	NFSv4	NFS transfer size with ONTAP 9	NFS transfer size with ONTAP 8
rw, bg, hard, timeo=600, noatime,	nfsvers=3,nolock,	nfsvers=4.1,lock	rsize=1048576,wsize=262144,	rsize=65536,wsize=65536,



To improve read performance with NFSv3, NetApp recommends that you use the `nconnect=n` mount option, which is available with SUSE Linux Enterprise Server 12 SP4 or later and RedHat Enterprise Linux (RHEL) 8.3 or later.



Performance tests show that `nconnect=4` provides good read results especially for the data volumes. Log writes might benefit from a lower number of sessions, such as `nconnect=2`. Shared volumes might benefit as well from using the 'nconnect' option. Be aware that the first mount from an NFS server (IP address) defines the amount of sessions being used. Further mounts to the same IP address do not change this even if a different value is used for `nconnect`.



Starting with ONTAP 9.8 and SUSE SLES15SP2 or RedHat RHEL 8.4 or higher, NetApp supports the `nconnect` option also for NFSv4.1.



If `nconnect` is being used with NFSv4.x the amount of NFSv4.x session slots should be adjusted according to the following rule:

Amount of session slots equals `<nconnect value> x 64`.

At the host this will be adjusted by

```
echo options nfs max_session_slots=<calculated value> >
/etc/modprobe.d/nfsclient.conf
```

followed by a reboot. The server side value must be adjusted as well, set the number of session slots as described in [NFS configuration for NFSv4](#).

To mount the file systems during system boot with the `/etc/fstab` configuration file, complete the following steps:

The following example shows a single host SAP HANA database with `SID=NF2` using NFSv3 and an NFS transfer size of 1MB for reads and 256k for writes.

1. Add the required file systems to the `/etc/fstab` configuration file.

```

sapcc-hana-tst-06:/ # cat /etc/fstab
<storage-vif-data01>:/NF2_data_mnt00001 /hana/data/NF2/mnt00001 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-log01>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=2,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data01>:/NF2_shared/usr-sap /usr/sap/NF2 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data01>:/NF2_shared/shared /hana/shared nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0

```

## 2. Run `mount -a` to mount the file systems on all hosts.

The next example shows a multiple-host SAP HANA database with SID=NF2 using NFSv4.1 for data and log file systems and NFSv3 for the `/hana/shared` and `/usr/sap/NF2` file systems. An NFS transfer size of 1MB for reads and 256k for writes is used.

## 1. Add the required file systems to the `/etc/fstab` configuration file on all hosts.



The `/usr/sap/NF2` file system is different for each database host. The following example shows `/NF2_shared/usr-sap-host1`.

```

sapcc-hana-tst-06:/ # cat /etc/fstab
<storage-vif-data01>:/NF2_data_mnt00001 /hana/data/NF2/mnt00001 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-data02>:/NF2_data_mnt00002 /hana/data/NF2/mnt00002 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-log01>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=2,rsize=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-log02>:/NF2_log_mnt00002 /hana/log/NF2/mnt00002 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=2,rsize=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-data02>:/NF2_shared/usr-sap-host1 /usr/sap/NF2 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data02>:/NF2_shared/shared /hana/shared nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsize=1048576,wsiz=262144,bg,noa
time,nolock 0 0

```

2. Run `mount -a` to mount the file systems on all hosts.

## SAP HANA installation preparations for NFSv4

NFS version 4 and higher requires user authentication. This authentication can be accomplished by using a central user management tool such as a Lightweight Directory Access Protocol (LDAP) server or with local user accounts. The following sections describe how to configure local user accounts.

The administration users `<sid>adm`, `<sid>crypt` and the `sapsys` group must be created manually on the SAP HANA hosts and the storage controllers before the installation of the SAP HANA software begins.

### SAP HANA hosts

If it doesn't exist, the `sapsys` group must be created on the SAP HANA host. A unique group ID must be chosen that does not conflict with the existing group IDs on the storage controllers.

The users `<sid>adm` and `<sid>crypt` are created on the SAP HANA host. Unique IDs must be chosen that does not conflict with existing user IDs on the storage controllers.

For a multiple-host SAP HANA system, the user and group IDs must be the same on all SAP HANA hosts. The group and user are created on the other SAP HANA hosts by copying the affected lines in `/etc/group` and `/etc/passwd` from the source system to all other SAP HANA hosts.



The NFSv4 domain must be set to the same value on all Linux servers (`/etc/idmapd.conf`) and SVMs. Set the domain parameter "Domain = <domain-name>" in the file `/etc/idmapd.conf` for the Linux hosts.

Enable and start the NFS IDMAPD service.

```
systemctl enable nfs-idmapd.service
systemctl start nfs-idmapd.service
```



The latest Linux kernels do not require this step. Warning messages can be safely ignored.

### Storage controllers

The user IDs and group ID must be the same on the SAP HANA hosts and the storage controllers. The group and user are created by entering the following commands on the storage cluster:

```
vserver services unix-group create -vserver <vserver> -name <group name>
-id <group id>
vserver services unix-user create -vserver <vserver> -user <user name> -id
<user-id> -primary-gid <group id>
```

Additionally, set the group ID of the UNIX user root of the SVM to 0.

```
vserver services unix-user modify -vserver <vserver> -user root -primary  
-gid 0
```

## I/O stack configuration for SAP HANA

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage systems used.

NetApp conducted performance tests to define the ideal values. The following table lists the optimal values inferred from the performance tests.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For SAP HANA 1.0 versions up to SPS12, these parameters can be set during the installation of the SAP HANA database, as described in SAP note [2267798: Configuration of the SAP HANA Database During Installation Using hdbparam](#).

Alternatively, the parameters can be set after the SAP HANA database installation by using the `hdbparam` framework.

```
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.max_parallel_io_requests=128  
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.async_write_submit_active=on  
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.async_read_submit=on  
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.async_write_submit_blocks=all
```

Starting with SAP HANA 2.0, `hdbparam` has been deprecated, and the parameters have been moved to `global.ini`. The parameters can be set using SQL commands or SAP HANA Studio. For more details, see SAP note [2399079: Elimination of hdbparam in HANA 2](#). You can also set the parameters within `global.ini` as shown in the following text:

```
nf2adm@stlrx300s8-6: /usr/sap/NF2/SYS/global/hdb/custom/config> cat
global.ini
...
[fileio]
async_read_submit = on
async_write_submit_active = on
max_parallel_io_requests = 128
async_write_submit_blocks = all
...
```

Since SAP HANA 2.0 SPS5, the `setParameter.py` script can be used to set the correct parameters:

```
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00/exe/python_support>
python setParameter.py
-set=SYSTEM/global.ini/fileio/max_parallel_io_requests=128
python setParameter.py -set=SYSTEM/global.ini/fileio/async_read_submit=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_active=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_blocks=all
```

### SAP HANA data volume size

As the default, SAP HANA uses only one data volume per SAP HANA service. Due to the maximum file size limitation of the file system, we recommend limiting the maximum data volume size.

To do so automatically, set the following parameter in `global.ini` in the section `[persistence]`:

```
datavolume_stripping = true
datavolume_stripping_size_gb = 8000
```

This creates a new data volume after the 8,000GB limit is reached. [SAP note 240005 question 15](#) provides more information.

### SAP HANA software installation

The following are requirements for software installation for SAP HANA.

#### Install on single-host system

The SAP HANA software installation does not require any additional preparation for a single-host system.

## Install on multiple-host system

To install SAP HANA on a multiple-host system, complete the following steps:

1. Using the SAP `hdbclm` installation tool, start the installation by running the following command at one of the worker hosts. Use the `addhosts` option to add the second worker (`sapcc-hana-tst-03`) and the standby host (`sapcc-hana-tst-04`).

```
apcc-hana-tst-02:/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_LCM_LINUX_X86_64 #
./hdbclm --action=install --addhosts=sapcc-hana-tst-03:role=worker,sapcc-
-hana-tst-04:role=standby

SAP HANA Lifecycle Management - SAP HANA Database 2.00.073.00.1695288802
*****

Scanning software locations...
Detected components:
    SAP HANA AFL (incl.PAL,BFL,OFL) (2.00.073.0000.1695321500) in
/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
    SAP HANA Database (2.00.073.00.1695288802) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
    SAP HANA Database Client (2.18.24.1695756995) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_CLIENT_LINUX_X86_64/SAP_HANA_CLIENT/client
    SAP HANA Studio (2.3.75.000000) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_STUDIO_LINUX_X86_64/studio
    SAP HANA Local Secure Store (2.11.0) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HANA_LSS_24_LINUX_X86_64/packages
    SAP HANA XS Advanced Runtime (1.1.3.230717145654) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_RT_10_LINUX_X86_64/packages
    SAP HANA EML AFL (2.00.073.0000.1695321500) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_EML_AFL_10_LINUX_X86_64/packages
    SAP HANA EPM-MDS (2.00.073.0000.1695321500) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/SAP_HANA_EPM-MDS_10/packages
    Automated Predictive Library (4.203.2321.0.0) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/PAAPL4_H20_LINUX_X86_64/apl-
4.203.2321.0-hana2sp03-linux_x64/installer/packages
    GUI for HALM for XSA (including product installer) Version 1
(1.015.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-
```

```

73/DATA_UNITS/XSA_CONTENT_10/XSACALMPIUI15_0.zip
    XSAC FILEPROCESSOR 1.0 (1.000.102) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACFILEPROC00_102.zip
    SAP HANA tools for accessing catalog content, data preview, SQL
console, etc. (2.015.230503) in /mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSAC_HRTT_20/XSACHRTT15_230503.zip
    Develop and run portal services for customer applications on XSA
(2.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACPORTALSERV07_0.zip
    The SAP Web IDE for HANA 2.0 (4.007.0) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSAC_SAP_WEB_IDE_20/XSACSAPWEBIDE07_0.zip
    XS JOB SCHEDULER 1.0 (1.007.22) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACSERVICES07_22.zip
    SAPUI5 FESV6 XSA 1 - SAPUI5 1.71 (1.071.52) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACUI5FESV671_52.zip
    SAPUI5 FESV9 XSA 1 - SAPUI5 1.108 (1.108.5) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACUI5FESV9108_5.zip
    SAPUI5 SERVICE BROKER XSA 1 - SAPUI5 Service Broker 1.0 (1.000.4) in
/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACUI5SB00_4.zip
    XSA Cockpit 1 (1.001.37) in /mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/XSA_CONTENT_10/XSACXSACOCKPIT01_37.zip

```

SAP HANA Database version '2.00.073.00.1695288802' will be installed.

Select additional components for installation:

Index	Components	Description
1	all	All components
2	server	No additional components
3	client	Install SAP HANA Database Client version 2.18.24.1695756995
4	lss	Install SAP HANA Local Secure Store version 2.11.0
5	studio	Install SAP HANA Studio version 2.3.75.000000
6	xs	Install SAP HANA XS Advanced Runtime version

```

1.1.3.230717145654
 7      | afl          | Install SAP HANA AFL (incl.PAL,BFL,OFL)
version 2.00.073.0000.1695321500
 8      | eml          | Install SAP HANA EML AFL version
2.00.073.0000.1695321500
 9      | epmmds      | Install SAP HANA EPM-MDS version
2.00.073.0000.1695321500
10     | sap_afl_sdk_apl | Install Automated Predictive Library version
4.203.2321.0.0

```

Enter comma-separated list of the selected indices [3,4]: 2,3

2. Verify that the installation tool installed all selected components at all worker and standby hosts.

### Adding additional data volume partitions

Starting with SAP HANA 2.0 SPS4, you can configure additional data volume partitions, which allows you to configure two or more volumes for the data volume of an SAP HANA tenant database. You can also scale beyond the size and performance limits of a single volume.



Using two or more individual volumes for the data volume is available for SAP HANA single-host and multiple-host systems. You can add additional data volume partitions at any time, but doing so might require a restart of the SAP HANA database.

### Enabling additional data volume partitions

1. To enable additional data volume partitions, add the following entry within `global.ini` using SAP HANA Studio or Cockpit in the SYSTEMDB configuration.

```

[customizable_functionalities]
persistence_datavolume_partition_multipath = true

```



Adding the parameter manually to the `global.ini` file requires the restart of the database.

### Volume configuration for a single-host SAP HANA system

The layout of volumes for a single-host SAP HANA system with multiple partitions is like the layout for a system with one data volume partition, but with an additional data volume stored on a different aggregate as the log volume and the other data volume. The following table shows an example configuration of an SAP HANA single-host system with two data volume partitions.

Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller b
Data volume: SID_data_mnt00001	Shared volume: SID_shared	Data volume: SID_data2_mnt00001	Log volume: SID_log_mnt00001

The following table shows an example of the mount point configuration for a single-host system with two data volume partitions.

Junction path	Directory	Mount point at HANA host
SID_data_mnt00001	–	/hana/data/SID/mnt00001
SID_data2_mnt00001	–	/hana/data2/SID/mnt00001
SID_log_mnt00001	–	/hana/log/SID/mnt00001
SID_shared	usr-sap shared	/usr/sap/SID /hana/shared

Create the new data volume and mount it to the namespace using either ONTAP System Manager or the ONTAP cluster command line interface.

#### Volume configuration for multiple-host SAP HANA system

The layout of volumes for a multiple-host SAP HANA system with multiple partitions is like the layout for a system with one data volume partition, but with an additional data volume stored on a different aggregate as the log volume and the other data volume. The following table shows an example configuration of an SAP HANA multiple-host system with two data volume partitions.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	–	Log volume: SID_log_mnt00001	Data2 volume: SID_data2_mnt00001
Data and log volumes for node 2	Log volume: SID_log_mnt00002	Data2 volume: SID_data2_mnt00002	Data volume: SID_data_mnt00002	–
Data and log volumes for node 3	–	Data volume: SID_data_mnt00003	Data2 volume: SID_data2_mnt00003	Log volume: SID_log_mnt00003
Data and log volumes for node 4	Data2 volume: SID_data2_mnt00004	Log volume: SID_log_mnt00004	–	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared	–	–	–

The following table shows an example of the mount point configuration for a single-host system with two data volume partitions.

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00001	–	/hana/data/SID/mnt00001	Mounted at all hosts
SID_data2_mnt00001	–	/hana/data2/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001	–	/hana/log/SID/mnt00001	Mounted at all hosts

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00002	–	/hana/data/SID/mnt00002	Mounted at all hosts
SID_data2_mnt00002	–	/hana/data2/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002	–	/hana/log/SID/mnt00002	Mounted at all hosts
SID_data_mnt00003	–	/hana/data/SID/mnt00003	Mounted at all hosts
SID_data2_mnt00003	–	/hana/data2/SID/mnt00003	Mounted at all hosts
SID_log_mnt00003	–	/hana/log/SID/mnt00003	Mounted at all hosts
SID_data_mnt00004	–	/hana/data/SID/mnt00004	Mounted at all hosts
SID_data2_mnt00004	–	/hana/data2/SID/mnt00004	Mounted at all hosts
SID_log_mnt00004	–	/hana/log/SID/mnt00004	Mounted at all hosts
SID_shared	shared	/hana/shared/SID	Mounted at all hosts
SID_shared	usr-sap-host1	/usr/sap/SID	Mounted at host 1
SID_shared	usr-sap-host2	/usr/sap/SID	Mounted at host 2
SID_shared	usr-sap-host3	/usr/sap/SID	Mounted at host 3
SID_shared	usr-sap-host4	/usr/sap/SID	Mounted at host 4
SID_shared	usr-sap-host5	/usr/sap/SID	Mounted at host 5

Create the new data volume and mount it to the namespace using either ONTAP System Manager or the ONTAP cluster command line interface.

### Host configuration

In addition to the tasks described in the section “[Host setup](#),” you must create the additional mount points and fstab entries for the new additional data volume(s), and you must mount the new volumes.

#### 1. Create additional mount points:

- For a single-host system, create mount points and set the permissions on the database host.

```
sapcc-hana-tst-06:/ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-06:/ # chmod -R 777 /hana/data2/SID
```

- For a multiple-host system, create mount points and set the permissions on all worker and standby hosts. The following example commands are for a 2+1 multiple-host HANA system.
  - First worker host:

```
sapcc-hana-tst-06:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-06:~ # chmod -R 777 /hana/data2/SID
```

▪ **Second worker host:**

```
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data2/SID
```

▪ **Standby host:**

```
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data2/SID
```

2. Add the additional file systems to the `/etc/fstab` configuration file on all hosts. An example for a single-host system using NFSv4.1 is as follows:

```
<storage-vif-data02>:/SID_data2_mnt00001 /hana/data2/SID/mnt00001 nfs
rw,vers=4,
minorversion=1,hard,timeo=600,rsz=1048576,wsz=262144,bg,noatime,lock
0 0
```



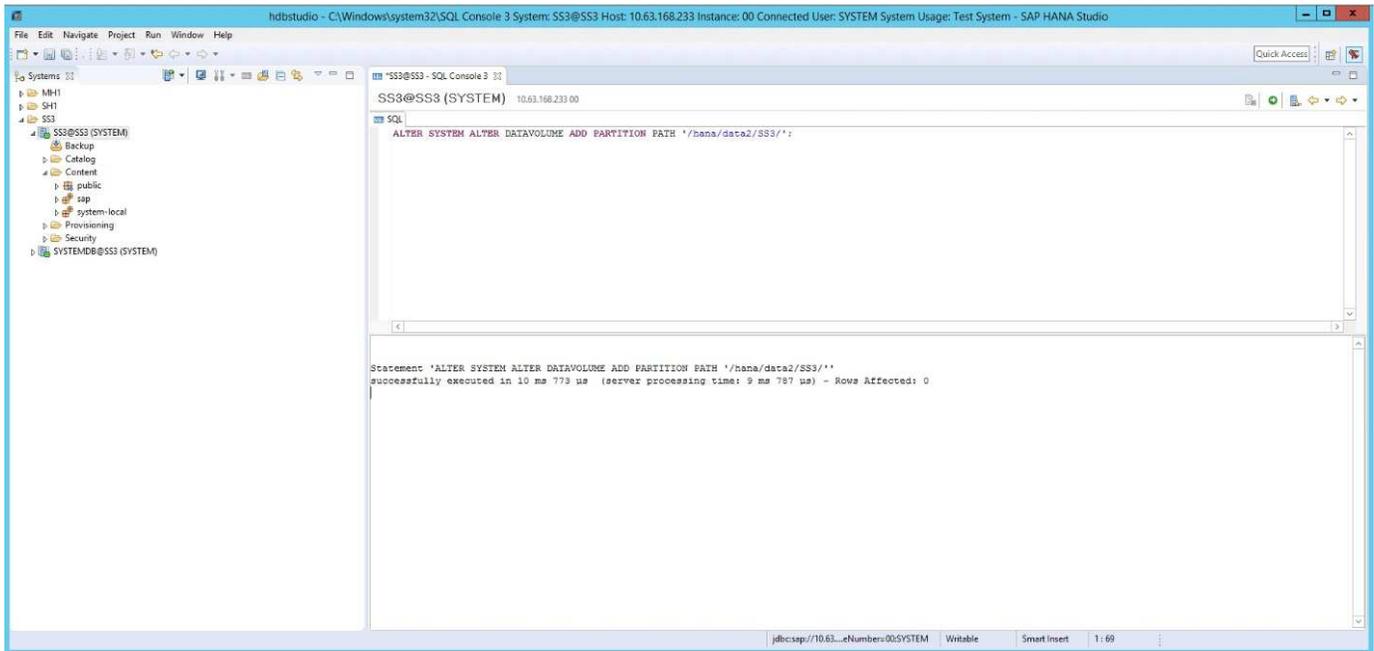
Use a different storage virtual interface for connecting to each data volume to make sure that different TCP sessions are used for each volume. You can also use the `nconnect` mount option if it is available for your OS.

3. To mount the file systems, run the `mount -a` command.

#### **Adding an additional data volume partition**

Execute the following SQL statement against the tenant database to add an additional data volume partition to your tenant database. Use the path to additional volume(s):

```
ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION PATH '/hana/data2/SID/';
```



## Where to find additional information

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

- [SAP HANA Software Solutions](#)
- [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP System Copies Using the SnapCenter SAP HANA Plug-In](#)
- [NetApp Documentation Centers](#)

<https://www.netapp.com/support-and-training/documentation/>

- [SAP Certified Enterprise Storage Hardware for SAP HANA](#)

<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/>

- [SAP HANA Storage Requirements](#)

<https://www.sap.com/documents/2024/03/146274d3-ae7e-0010-bca6-c68f7e60039b.html>

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

<https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html>

- [SAP HANA on VMware vSphere Wiki](#)

[https://help.sap.com/docs/SUPPORT\\_CONTENT/virtualization/3362185751.html](https://help.sap.com/docs/SUPPORT_CONTENT/virtualization/3362185751.html)

- [SAP HANA on VMware vSphere Best Practices Guide](#)

[https://www.vmware.com/docs/sap\\_hana\\_on\\_vmware\\_vsphere\\_best\\_practices\\_guide-white-paper](https://www.vmware.com/docs/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper)

## Update history

The following technical changes have been made to this solution since its original publication.

Date	Update summary
April 2014	Initial version
August 2014	Updated disk sizing selection and added SSD configuration Added Red Hat Enterprise Linux OS configuration Added SAP HANA storage connector information Added information about VMware configuration
November 2014	Updated storage sizing section
January 2015	Updated storage connector API section Updated aggregate and volume configuration
March 2015	Added new STONITH implementation for SAP HANA SPS9 Added compute node setup and HANA installation section
October 2015	Added NFSv4 support for cDOT Updated sysctl parameter Included I/O parameter for SAP HANA and HWVAL > SPS10
March 2016	Updated capacity sizing Updated mount options for /hana/shared Updated sysctl parameter
February 2017	New NetApp storage systems and disk shelves New features of ONTAP 9 Support for 40GbE New OS releases (SUSE Linux Enterprise Server12 SP1 and Red Hat Enterprise Linux 7.2) New SAP HANA release
July 2017	Minor updates
September 2018	New NetApp storage systems New OS releases (SUSE Linux Enterprise Server 12 SP3 and Red Hat Enterprise Linux 7.4) Additional minor changes SAP HANA 2.0 SPS3
September 2019	New OS releases (SUSE Linux Enterprise Server 12 SP4, SUSE Linux Enterprise Server 15, and Red Hat Enterprise Linux 7.6) Max data volume size Minor changes
December 2019	New NetApp storage systems New OS release SUSE Linux Enterprise Server 15 SP1
March 2020	Support of nconnect for NFSv3 New OS release Red Hat Enterprise Linux 8
May 2020	Introduced multiple data partition features available since SAP HANA 2.0 SPS4

Date	Update summary
June 2020	Additional information about optional functionalities Minor updates
December 2020	Support for nconnect for NFSv4.1 starting with ONTAP 9.8 New operating system releases New SAP HANA version
February 2021	Changes in host network settings and other minor changes
April 2021	VMware vSphere-specific information added
September 2022	New OS-Releases
December 2023	Update of host setup Revised nconnect settings Added information about NFSv4.1 sessions
September 2024	New Storage Systems and Minor Updates
February 2025	New Storage System
July 2025	Minor updates

# SAP HANA on FAS Systems with FCP Configuration Guide

## SAP HANA on NetApp FAS Systems with Fibre Channel Protocol Configuration Guide

The NetApp FAS product family has been certified for use with SAP HANA in TDI projects. This guide provides best practices for SAP HANA on this platform for FCP.

Marco Schoen, NetApp

The certification is valid for the following models:

- FAS2750, FAS2820, FAS8300, FAS50, FAS8700, FAS70, FAS9500, FAS90

For a complete list of NetApp's certified storage solutions for SAP HANA, see the [certified and supported SAP HANA hardware directory](#).

This document describes FAS configurations that use the Fibre Channel Protocol (FCP).



The configuration described in this paper is necessary to achieve the required SAP HANA KPIs and the best performance for SAP HANA. Changing any settings or using features not listed herein might result in performance degradation or unexpected behavior and should only be done if advised by NetApp support.

The configuration guides for FAS systems using NFS and NetApp AFF systems can be found using the following links:

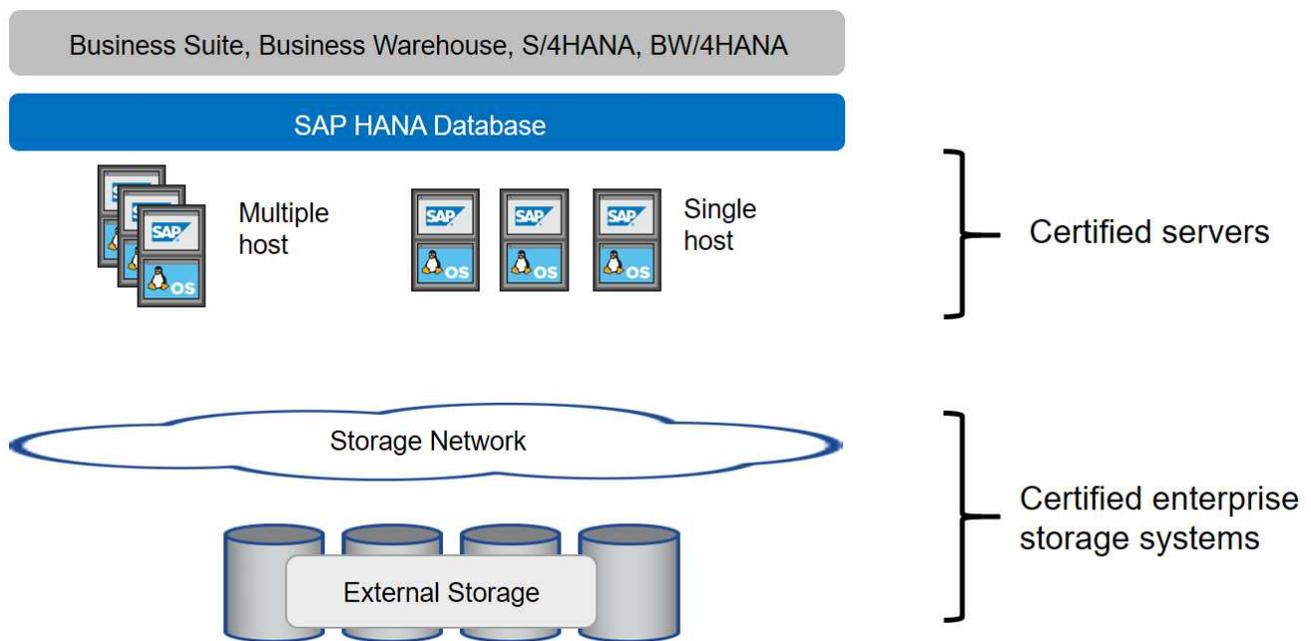
- [SAP HANA on NetApp AFF Systems with FCP](#)
- [SAP HANA on NetApp ASA Systems with FCP](#)

- [SAP HANA on NetApp FAS Systems with NFS](#)
- [SAP HANA on NetApp AFF Systems with NFS](#)

In an SAP HANA multiple-host environment, the standard SAP HANA storage connector is used to provide fencing in the event of an SAP HANA host failover. Refer to the relevant SAP notes for operating system configuration guidelines and HANA-specific Linux kernel dependencies. For more information, see [SAP Note 2235581 – SAP HANA Supported Operating Systems](#).

### SAP HANA tailored data center integration

NetApp FAS storage controllers are certified in the SAP HANA Tailored Data Center Integration (TDI) program using NFS (NAS) and Fibre Channel (SAN) protocols. They can be deployed in any SAP HANA scenario, such as, SAP Business Suite on HANA, S/4HANA, BW/4HANA or SAP Business Warehouse on HANA in single-host or multiple-host configurations. Any server that is certified for use with SAP HANA can be combined with the certified storage solution. See the following figure for an architecture overview.



For more information regarding the prerequisites and recommendations for productive SAP HANA systems, see the following resource:

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

### SAP HANA using VMware vSphere

There are several options for connecting storage to virtual machines (VMs). The preferred one is to connect the storage volumes with NFS directly out of the guest operating system. This option is described in [SAP HANA on NetApp AFF Systems with NFS](#).

Raw device mappings (RDM), FCP datastores, or VVOL datastores with FCP are supported as well. For both datastore options, only one SAP HANA data or log volume must be stored within the datastore for productive use cases.

For more information about using vSphere with SAP HANA, see the following links:

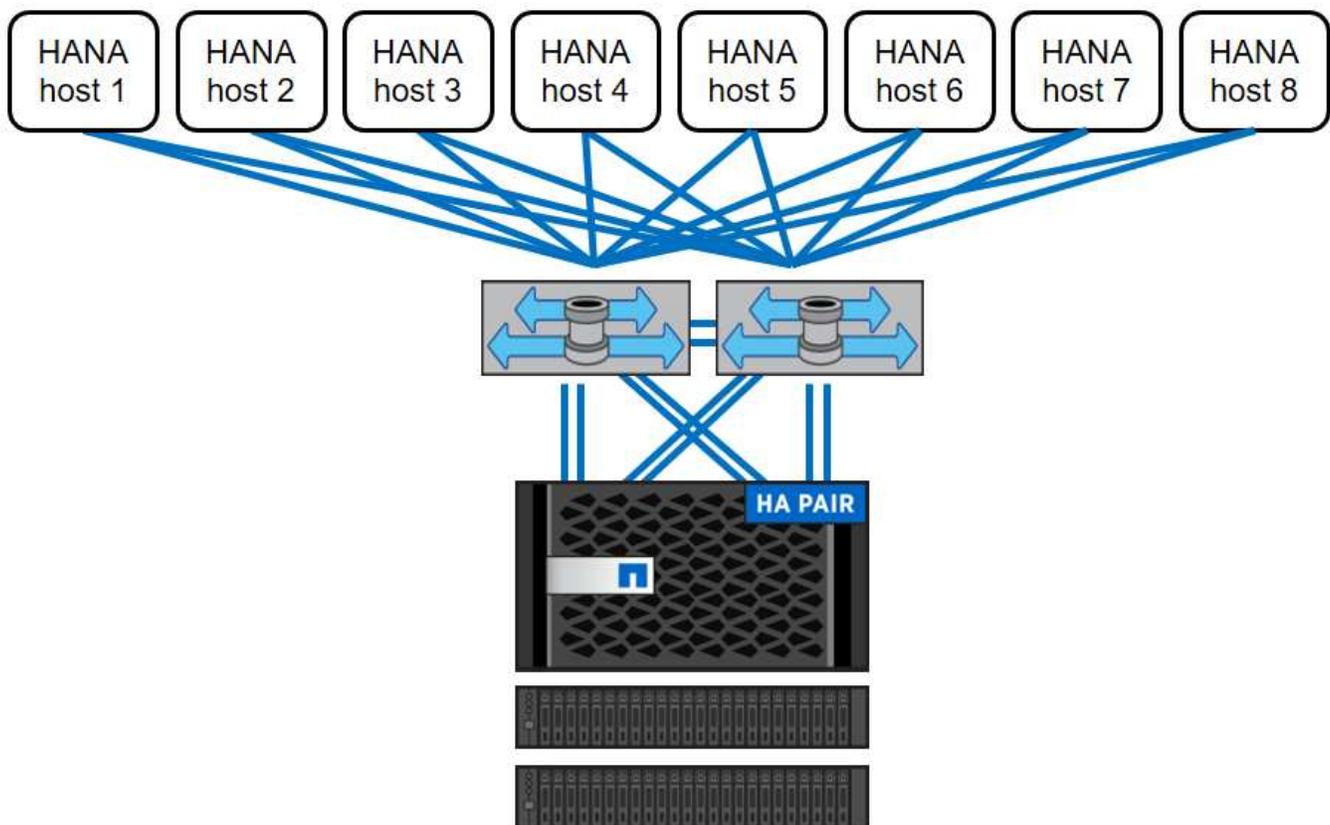
- [SAP HANA on VMware vSphere - Virtualization - Community Wiki](#)
- [SAP HANA on VMware vSphere Best Practices Guide](#)
- [2161991 - VMware vSphere configuration guidelines - SAP ONE Support Launchpad \(Login required\)](#)

## Architecture

SAP HANA hosts are connected to the storage controllers using a redundant FCP infrastructure and multipath software. A redundant FCP switch infrastructure is required to provide fault-tolerant SAP HANA host-to-storage connectivity in case of switch or host bus adapter (HBA) failure. Appropriate zoning must be configured at the switch to allow all HANA hosts to reach the required LUNs on the storage controllers.

Different models of the FAS product family can be used at the storage layer. The maximum number of SAP HANA hosts attached to the storage is defined by the SAP HANA performance requirements. The number of disk shelves required is determined by the capacity and performance requirements of the SAP HANA systems.

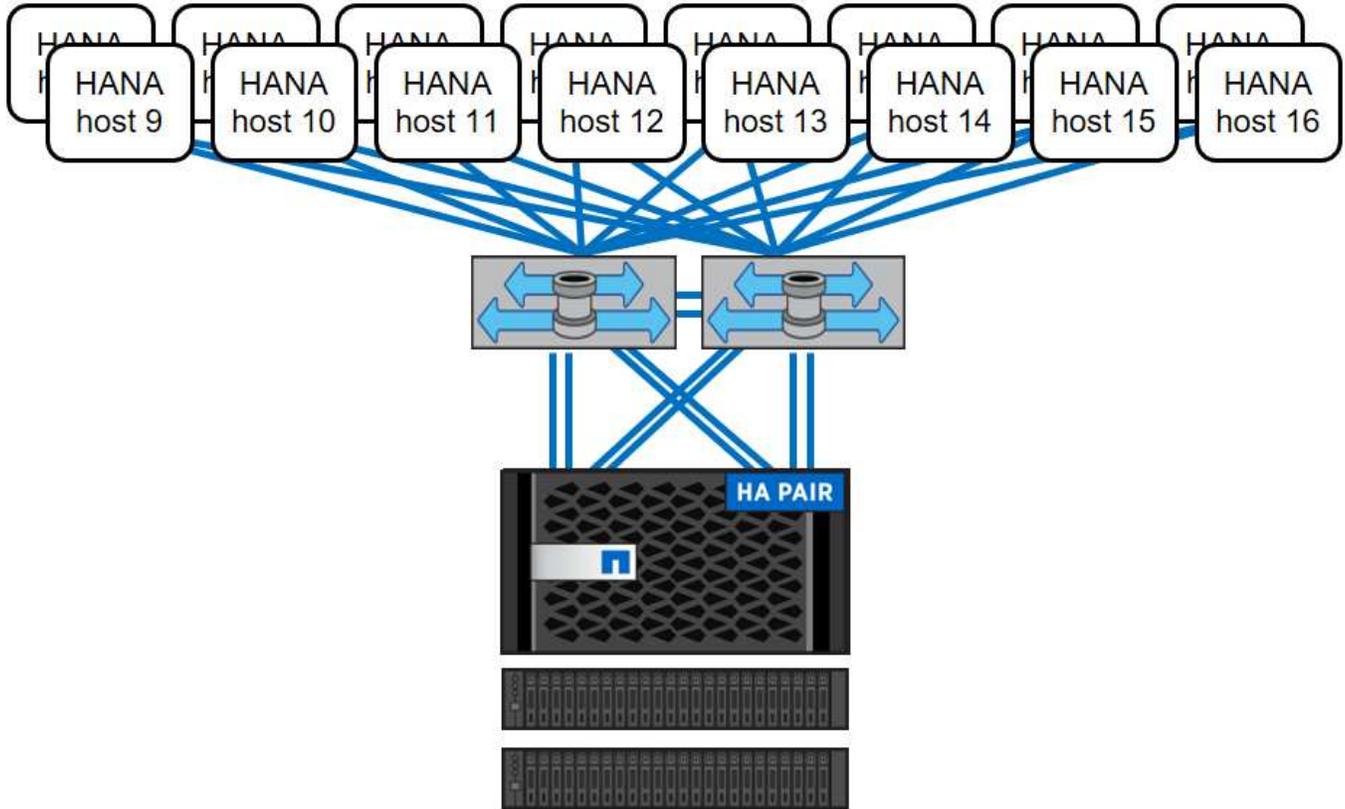
The following figure shows an example configuration with eight SAP HANA hosts attached to a storage HA pair.



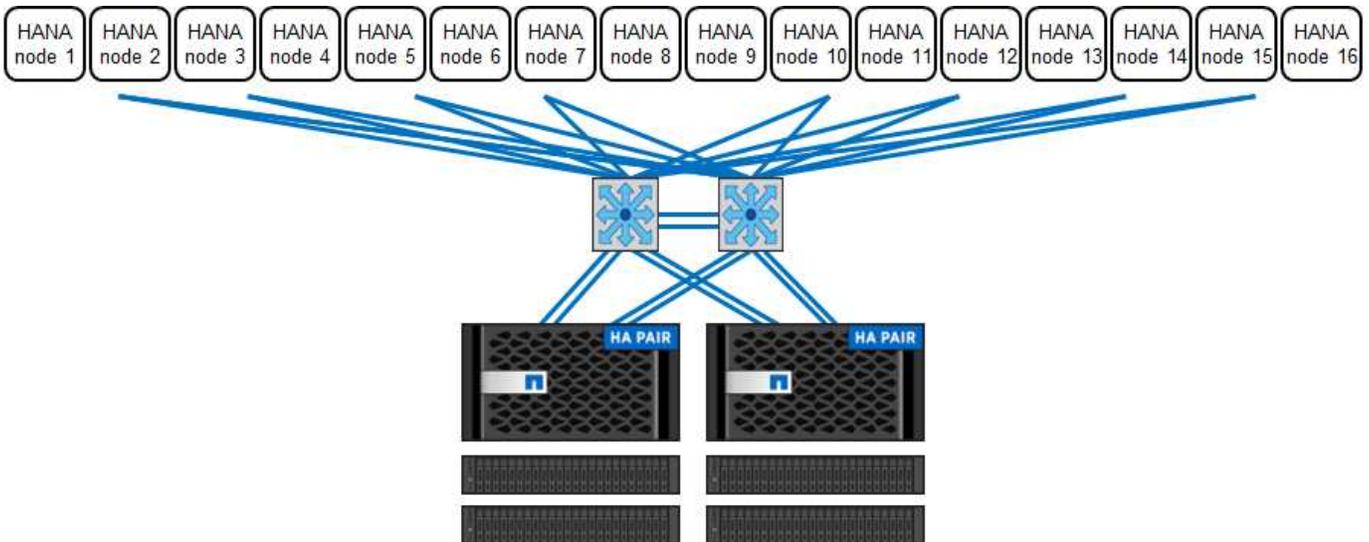
This architecture can be scaled in two dimensions:

- By attaching additional SAP HANA hosts and disk capacity to the storage, assuming that the storage controllers can provide enough performance under the new load to meet key performance indicators (KPIs)
- By adding more storage systems and disk capacity for the additional SAP HANA hosts

The following figure shows a configuration example in which more SAP HANA hosts are attached to the storage controllers. In this example, more disk shelves are necessary to meet the capacity and performance requirements of the 16 SAP HANA hosts. Depending on the total throughput requirements, you must add additional FC connections to the storage controllers.



Independent of the deployed FAS system storage model, the SAP HANA landscape can also be scaled by adding more storage controllers, as shown in the following figure.



### SAP HANA backup

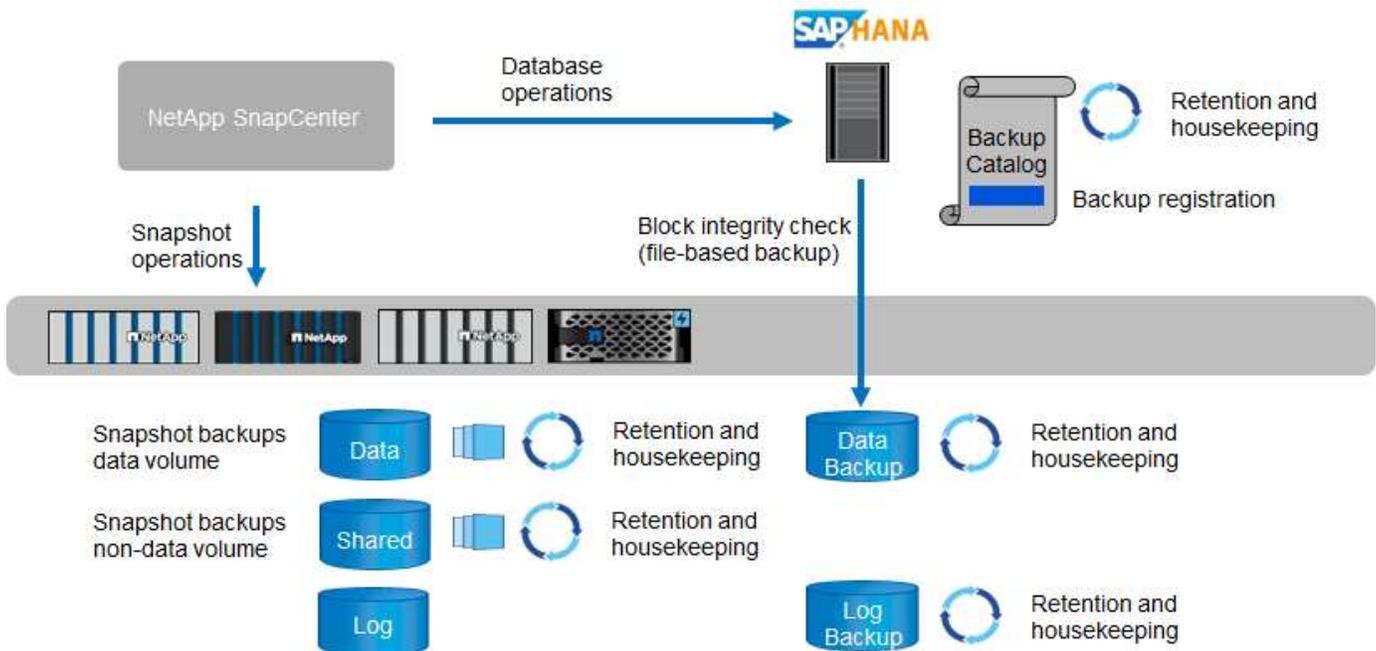
NetApp ONTAP software provides a built-in mechanism to back up SAP HANA databases. Storage-based Snapshot backup is a fully supported and integrated backup solution available for SAP HANA single-container

systems and for SAP HANA MDC single-tenant systems.

Storage-based Snapshot backups are implemented by using the NetApp SnapCenter plug-in for SAP HANA, which enables consistent storage-based Snapshot backups by using the interfaces provided by the SAP HANA database. SnapCenter registers the Snapshot backups in the SAP HANA backup catalog so that the backups are visible within the SAP HANA studio and can be selected for restore and recovery operations.

By using NetApp SnapVault software, the Snapshot copies that were created on the primary storage can be replicated to the secondary backup storage controlled by SnapCenter. Different backup retention policies can be defined for backups on the primary storage and for backups on the secondary storage. The SnapCenter Plug-in for SAP HANA Database manages the retention of Snapshot copy-based data backups and log backups including the housekeeping of the backup catalog. The SnapCenter Plug-in for SAP HANA Database also enables the execution of a block-integrity check of the SAP HANA database by performing a file-based backup.

The database logs can be backed up directly to the secondary storage by using an NFS mount, as shown in the following figure.



Storage-based Snapshot backups provide significant advantages compared to file-based backups. Those advantages include the following:

- Faster backup (few minutes)
- Faster restore on the storage layer (a few minutes)
- No effect on the performance of the SAP HANA database host, network, or storage during backup
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes

For detailed information about the SAP HANA backup and recovery solution using SnapCenter, see [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

### SAP HANA disaster recovery

SAP HANA disaster recovery can be performed on the database layer by using SAP system replication or on the storage layer by using storage-replication technologies. The following section provides an overview of

disaster recovery solutions based on storage replication.

For detailed information about the SAP HANA disaster recovery solution using SnapCenter, see [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#).

### Storage replication based on SnapMirror

The following figure shows a three-site disaster recovery solution, using synchronous SnapMirror replication to the local DR datacenter and asynchronous SnapMirror to replicate data to the remote DR datacenter.

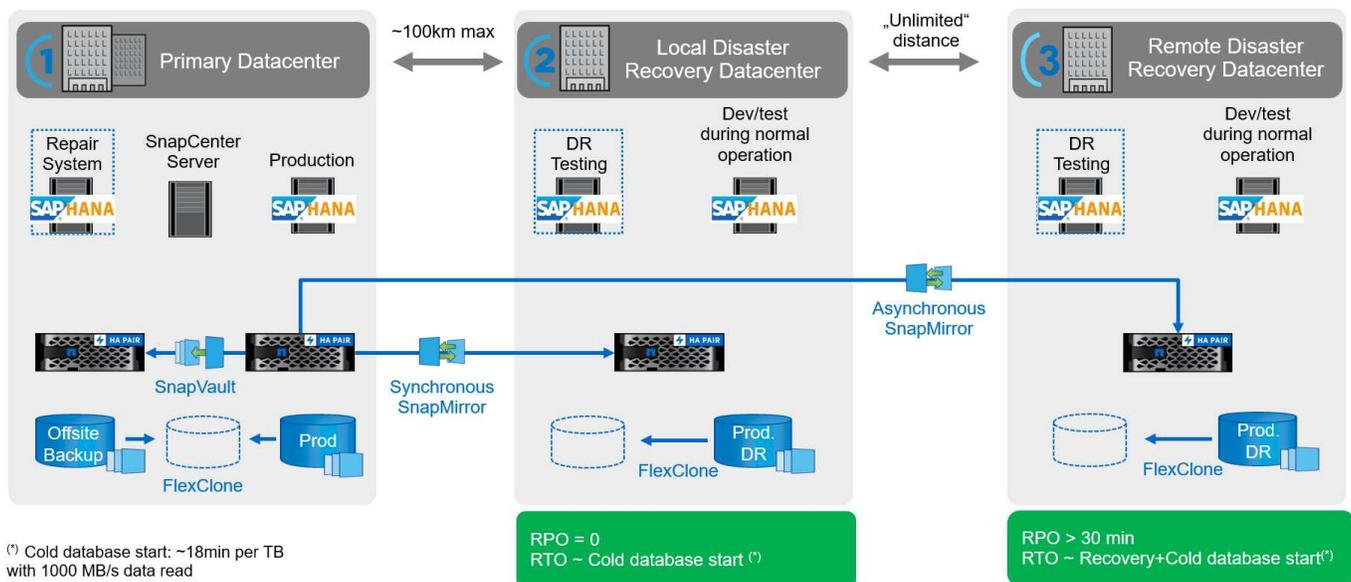
Data replication using synchronous SnapMirror provides an RPO of zero. The distance between the primary and the local DR datacenter is limited to around 100km.

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

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

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

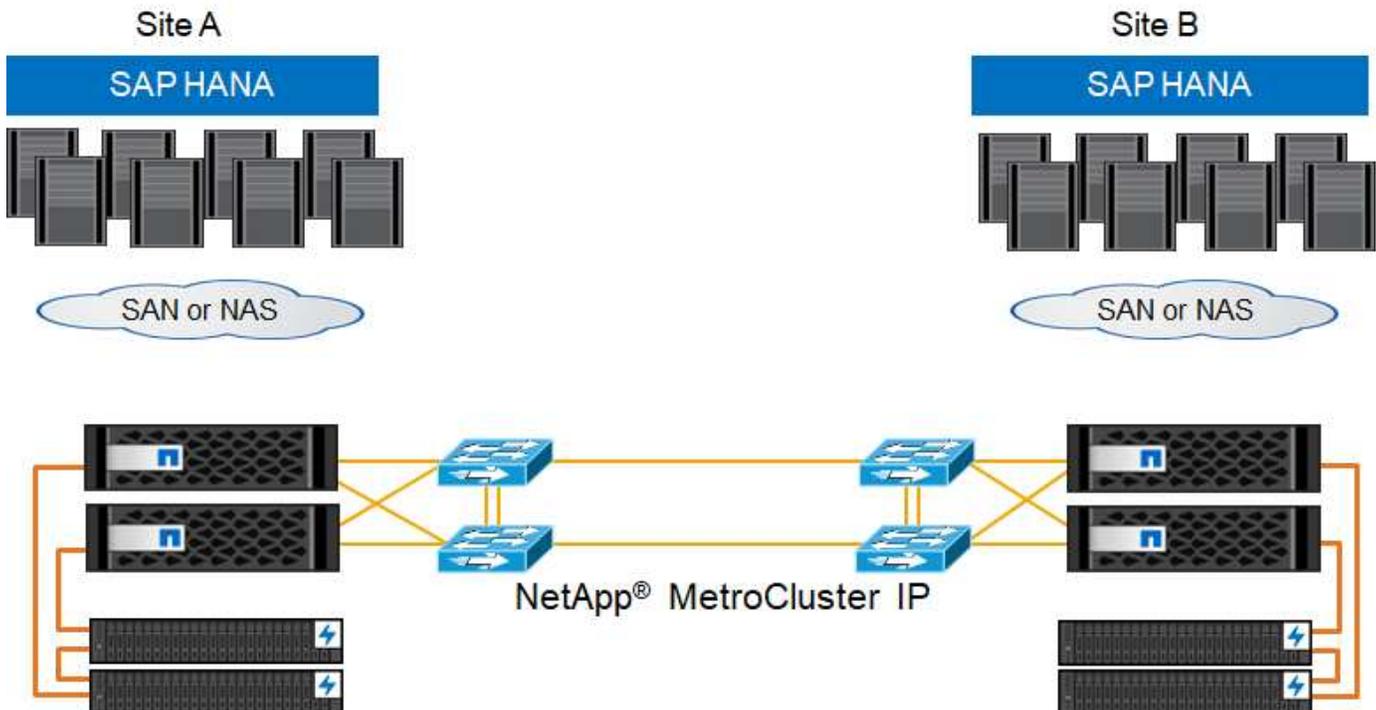
Both replication methods allow to you execute DR workflow testing without influencing the RPO and RTO. FlexClone volumes are created on the storage and are attached to the DR testing servers.



Synchronous replication offers StrictSync mode. If the write to secondary storage is not completed for any reason, the application I/O fails, thereby ensuring that the primary and secondary storage systems are identical. Application I/O to the primary resumes only after the SnapMirror relationship returns to the InSync status. If the primary storage fails, application I/O can be resumed on the secondary storage after failover with no loss of data. In StrictSync mode, the RPO is always zero.

## Storage replication based on NetApp MetroCluster

The following figure shows a high-level overview of the solution. The storage cluster at each site provides local high availability and is used for production workloads. The data at each site is synchronously replicated to the other location and is available in case of disaster failover.



## Storage sizing

The following section provides an overview of performance and capacity considerations for sizing a storage system for SAP HANA.



Contact your NetApp or NetApp partner sales representative to support the storage sizing process and to create a properly sized storage environment.

## Performance considerations

SAP has defined a static set of storage KPIs. These KPIs are valid for all production SAP HANA environments independent of the memory size of the database hosts and the applications that use the SAP HANA database. These KPIs are valid for single-host, multiple-host, Business Suite on HANA, Business Warehouse on HANA, S/4HANA, and BW/4HANA environments. Therefore, the current performance sizing approach depends on only the number of active SAP HANA hosts that are attached to the storage system.



Storage performance KPIs are required only for production SAP HANA systems.

SAP delivers a performance test tool, which must be used to validate the storage performance for active SAP HANA hosts attached to the storage.

NetApp tested and predefined the maximum number of SAP HANA hosts that can be attached to a specific storage model, while still fulfilling the required storage KPIs from SAP for production-based SAP HANA systems.



The storage controllers of the certified FAS product family can also be used for SAP HANA with other disk types or disk back-end solutions, as long as they are supported by NetApp and fulfill SAP HANA TDI performance KPIs. Examples include NetApp Storage Encryption (NSE) and NetApp FlexArray technology.

This document describes disk sizing for SAS hard disk drives and solid-state drives.

### Hard disk drives

A minimum of 10 data disks (10k RPM SAS) per SAP HANA node is required to fulfill the storage performance KPIs from SAP.



This calculation is independent of the storage controller and disk shelf used.

### Solid-state drives

With solid-state drives (SSDs), the number of data disks is determined by the SAS connection throughput from the storage controllers to the SSD shelf.

The maximum number of SAP HANA hosts that can be run on a disk shelf and the minimum number of SSDs required per SAP HANA host were determined by running the SAP performance test tool.

- The 12Gb SAS disk shelf (DS224C) with 24 SSDs supports up to 14 SAP HANA hosts, when the disk shelf is connected with 12Gb.
- The 6Gb SAS disk shelf (DS2246) with 24 SSDs supports up to 4 SAP HANA hosts.

The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers.

The following table summarizes the supported number of SAP HANA hosts per disk shelf.

	<b>6Gb SAS shelves (DS2246) fully loaded with 24 SSDs</b>	<b>12Gb SAS shelves (DS224C) fully loaded with 24 SSDs</b>
Maximum number of SAP HANA hosts per disk shelf	4	14



This calculation is independent of the storage controller used. Adding more disk shelves does not increase the maximum number of SAP HANA hosts that a storage controller can support.

### NS224 NVMe shelf

One NVMe SSDs (data) supports up to 2 SAP HANA hosts.

The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers.

### Mixed workloads

SAP HANA and other application workloads running on the same storage controller or in the same storage aggregate are supported. However, it is a NetApp best practice to separate SAP HANA workloads from all other application workloads.

You might decide to deploy SAP HANA workloads and other application workloads on either the same storage controller or the same aggregate. If so, you must make sure that enough performance is always available for SAP HANA within the mixed workload environment. NetApp also recommends that you use quality of service

(QoS) parameters to regulate the impact these other applications could have on SAP HANA applications.

The SAP HCMT test tool must be used to check if additional SAP HANA hosts can be run on a storage controller that is already used for other workloads. However, SAP application servers can be safely placed on the same storage controller and aggregate as the SAP HANA databases.

## Capacity considerations

A detailed description of the capacity requirements for SAP HANA is in the [SAP Note 1900823](#) white paper.



The capacity sizing of the overall SAP landscape with multiple SAP HANA systems must be determined by using SAP HANA storage sizing tools from NetApp. Contact NetApp or your NetApp partner sales representative to validate the storage sizing process for a properly sized storage environment.

## Configuration of performance test tool

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used. These parameters must also be set for the performance test tool from SAP (fsperf) when the storage performance is tested by using the SAP test tool.

Performance tests were conducted by NetApp to define the optimal values. The following table lists the parameters that must be set within the configuration file of the SAP test tool.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For more information about the configuration of SAP test tool, see [SAP note 1943937](#) for HWCCT (SAP HANA 1.0) and [SAP note 2493172](#) for HCMT/HCOT (SAP HANA 2.0).

The following example shows how variables can be set for the HCMT/HCOT execution plan.

```
...{
    "Comment": "Log Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "LogAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Data Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "DataAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
```

```

    },
    {
        "Comment": "Log Volume: Controls whether write requests can be
submitted asynchronously",
        "Name": "LogAsyncWriteSubmitActive",
        "Value": "on",
        "Request": "false"
    },
    {
        "Comment": "Data Volume: Controls whether write requests can be
submitted asynchronously",
        "Name": "DataAsyncWriteSubmitActive",
        "Value": "on",
        "Request": "false"
    },
    {
        "Comment": "Log Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
        "Name": "LogAsyncWriteSubmitBlocks",
        "Value": "all",
        "Request": "false"
    },
    {
        "Comment": "Data Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
        "Name": "DataAsyncWriteSubmitBlocks",
        "Value": "all",
        "Request": "false"
    },
    {
        "Comment": "Log Volume: Maximum number of parallel I/O requests
per completion queue",
        "Name": "LogExtMaxParallelIoRequests",
        "Value": "128",
        "Request": "false"
    },
    {
        "Comment": "Data Volume: Maximum number of parallel I/O requests
per completion queue",
        "Name": "DataExtMaxParallelIoRequests",
        "Value": "128",
        "Request": "false"
    },
    }, ...

```

These variables must be used for the test configuration. This is usually the case with the predefined execution plans SAP delivers with the HCMT/HCOT tool. The following example for a 4k log write test is from an execution plan.

```
...
  {
    "ID": "D664D001-933D-41DE-A904F304AEB67906",
    "Note": "File System Write Test",
    "ExecutionVariants": [
      {
        "ScaleOut": {
          "Port": "${RemotePort}",
          "Hosts": "${Hosts}",
          "ConcurrentExecution": "${FSConcurrentExecution}"
        },
        "RepeatCount": "${TestRepeatCount}",
        "Description": "4K Block, Log Volume 5GB, Overwrite",
        "Hint": "Log",
        "InputVector": {
          "BlockSize": 4096,
          "DirectoryName": "${LogVolume}",
          "FileOverwrite": true,
          "FileSize": 5368709120,
          "RandomAccess": false,
          "RandomData": true,
          "AsyncReadSubmit": "${LogAsyncReadSubmit}",
          "AsyncWriteSubmitActive":
"${LogAsyncWriteSubmitActive}",
          "AsyncWriteSubmitBlocks":
"${LogAsyncWriteSubmitBlocks}",
          "ExtMaxParallelIoRequests":
"${LogExtMaxParallelIoRequests}",
          "ExtMaxSubmitBatchSize": "${LogExtMaxSubmitBatchSize}",
          "ExtMinSubmitBatchSize": "${LogExtMinSubmitBatchSize}",
          "ExtNumCompletionQueues":
"${LogExtNumCompletionQueues}",
          "ExtNumSubmitQueues": "${LogExtNumSubmitQueues}",
          "ExtSizeKernelIoQueue": "${ExtSizeKernelIoQueue}"
        }
      }, ...
    ]
  }
```

### Storage sizing process overview

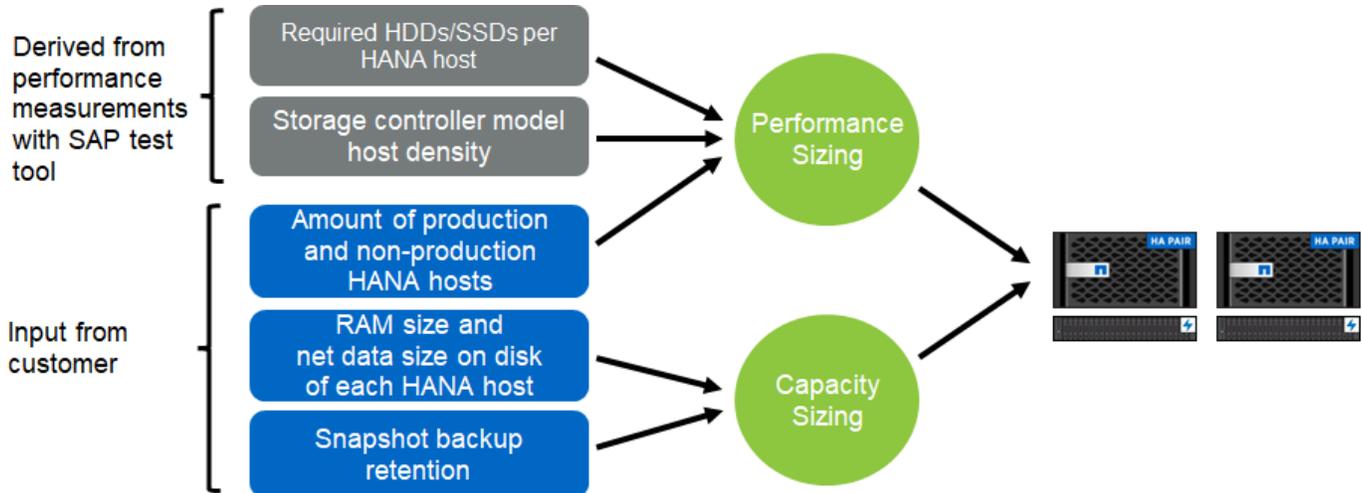
The number of disks per HANA host and the SAP HANA host density for each storage model were determined with the SAP HANA test tool.

The sizing process requires details such as the number of production and nonproduction SAP HANA hosts, the

RAM size of each host, and the backup retention period of the storage-based Snapshot copies. The number of SAP HANA hosts determines the storage controller and the number of disks required.

The size of the RAM, the net data size on the disk of each SAP HANA host, and the Snapshot copy backup retention period are used as inputs during capacity sizing.

The following figure summarizes the sizing process.



## Infrastructure setup and configuration

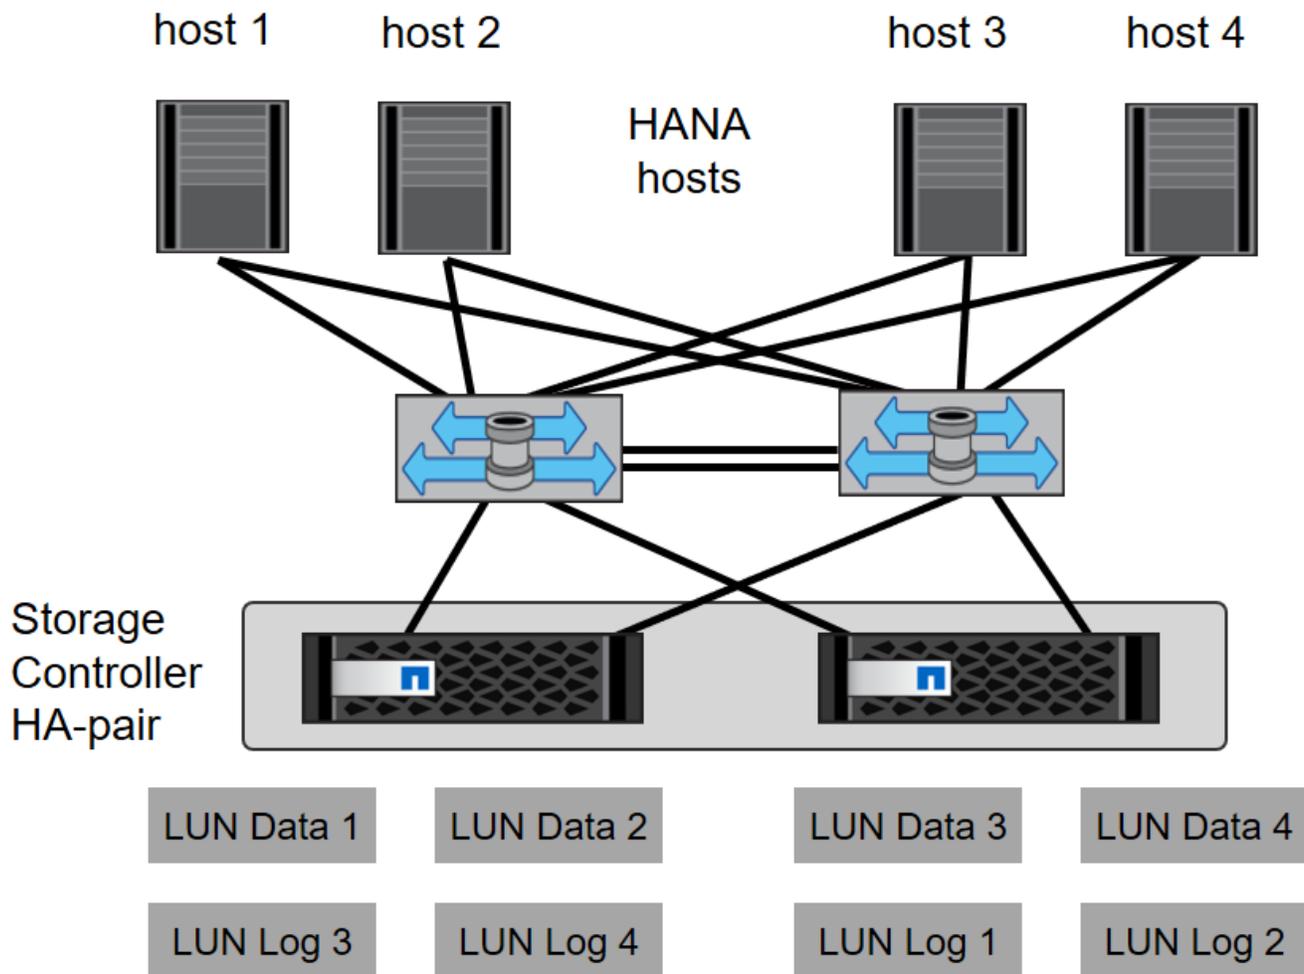
The following sections provide SAP HANA infrastructure setup and configuration guidelines and describes all the steps needed to set up an SAP HANA system. Within these sections, the following example configurations are used:

- HANA system with SID=FC5
  - SAP HANA single and multiple host using Linux logical volume manager (LVM)
  - SAP HANA single host using SAP HANA multiple partitions

### SAN fabric setup

Each SAP HANA server must have a redundant FCP SAN connection with a minimum of 8Gbps bandwidth. For each SAP HANA host attached to a storage controller, at least 8Gbps of bandwidth must be configured at the storage controller.

The following figure shows an example with four SAP HANA hosts attached to two storage controllers. Each SAP HANA host has two FCP ports connected to the redundant fabric. At the storage layer, four FCP ports are configured to provide the required throughput for each SAP HANA host.



In addition to the zoning on the switch layer, you must map each LUN on the storage system to the hosts that connect to this LUN. Keep the zoning on the switch simple; that is, define one zone set in which all host HBAs can see all controller HBAs.

### Time synchronization

You must synchronize the time between the storage controllers and the SAP HANA database hosts. The same time server must be set for all storage controllers and all SAP HANA hosts.

### Storage controller setup

This section describes the configuration of the NetApp storage system. You must complete the primary installation and setup according to the corresponding ONTAP setup and configuration guides.

### Storage efficiency

Inline deduplication, cross- volume inline deduplication, inline compression, and inline compaction are supported with SAP HANA in an SSD configuration.

Enabling the storage efficiency features in an HDD configuration is not supported.

## NetApp FlexGroup Volumes

The usage of NetApp FlexGroup Volumes is not supported for SAP HANA. Due to the architecture of SAP HANA the usage of FlexGroup Volumes does not provide any benefit and may result in performance issues.

## NetApp Volume and Aggregate Encryption

The use of NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) are supported with SAP HANA.

## Quality of Service

QoS can be used to limit the storage throughput for specific SAP HANA systems or non-SAP applications on a shared controller.

## Production and Dev/Test

One use case would be to limit the throughput of development and test systems so that they cannot influence production systems in a mixed setup.

During the sizing process, you should determine the performance requirements of a nonproduction system. Development and test systems can be sized with lower performance values, typically in the range of 20% to 50% of a production-system KPI as defined by SAP.

Large write I/O has the biggest performance effect on the storage system. Therefore, the QoS throughput limit should be set to a percentage of the corresponding write SAP HANA storage performance KPI values in the data and log volumes.

## Shared Environments

Another use case is to limit the throughput of heavy write workloads, especially to avoid that these workloads have an impact on other latency sensitive write workloads.

In such environments it is best practice to apply a non-shared throughput ceiling QoS group-policy to each LUN within each Storage Virtual Machine (SVM) to restrict the max throughput of each individual storage object to the given value. This reduces the possibility that a single workload can negatively influence other workloads.

To do so, a group-policy needs to be created using the CLI of the ONTAP cluster for each SVM:

```
qos policy-group create -policy-group <policy-name> -vserver <vserver
name> -max-throughput 1000MB/s -is-shared false
```

and applied to each LUN within the SVM. Below is an example to apply the policy group to all existing LUNs within an SVM:

```
lun modify -vserver <vserver name> -path * -qos-policy-group <policy-
name>
```

This needs to be done for every SVM. The name of the QoS police group for each SVM needs to be different. For new LUNs, the policy can be applied directly:

```
lun create -vserver <vserver_name> -path /vol/<volume_name>/<lun_name>
-size <size> -ostype <e.g. linux> -qos-policy-group <policy-name>
```

It is recommended to use 1000MB/s as maximum throughput for a given LUN. If an application requires more throughput, multiple LUNs with LUN striping shall be used to provide the needed bandwidth. This guide provides an example for SAP HANA based on Linux LVM in section [Host Setup](#).



The limit applies also to reads. Therefore use enough LUNs to fulfil the required SLAs for SAP HANA database startup time and for backups.

### NetApp FabricPool

NetApp FabricPool technology must not be used for active primary file systems in SAP HANA systems. This includes the file systems for the data and log area as well as the `/hana/shared` file system. Doing so results in unpredictable performance, especially during the startup of an SAP HANA system.

Using the “snapshot-only” tiering policy is possible as well as using FabricPool in general at a backup target such as SnapVault or SnapMirror destination.



Using FabricPool for tiering Snapshot copies at primary storage or using FabricPool at a backup target changes the required time for the restore and recovery of a database or other tasks such as creating system clones or repair systems. Take this into consideration for planning your overall lifecycle- management strategy, and check to make sure that your SLAs are still being met while using this function.

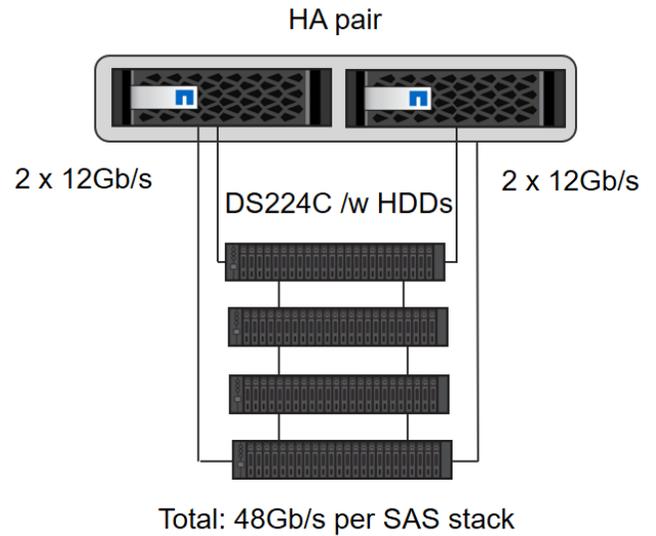
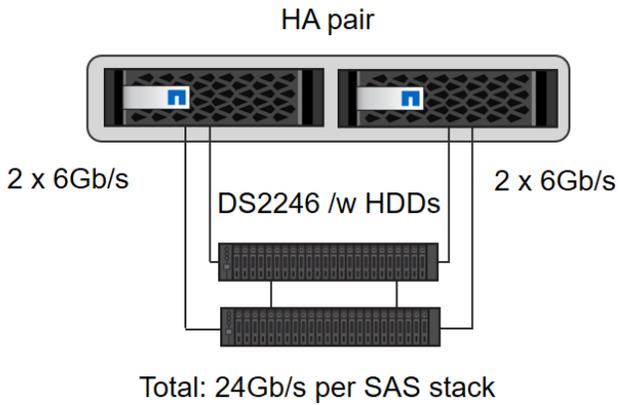
FabricPool is a good option for moving log backups to another storage tier. Moving backups affects the time needed to recover an SAP HANA database. Therefore, the option “tiering-minimum-cooling-days” should be set to a value that places log backups, which are routinely needed for recovery, on the local fast storage tier.

### Configure storage

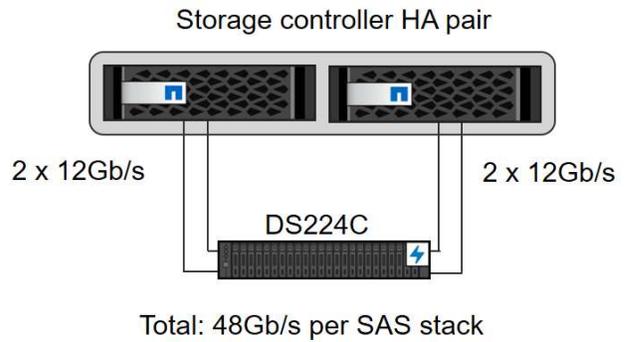
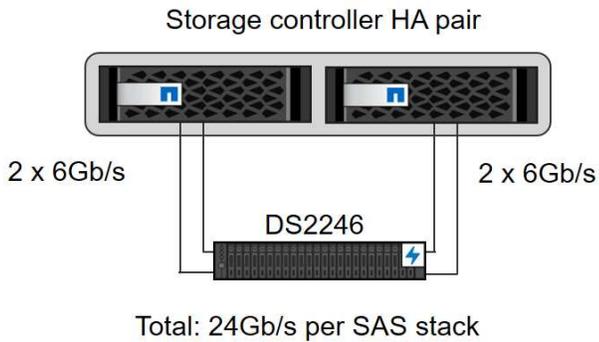
The following overview summarizes the required storage configuration steps. Each step is covered in more detail in the subsequent sections. Before initiating these steps, complete the storage hardware setup, the ONTAP software installation, and the connection of the storage FCP ports to the SAN fabric.

1. Check the correct disk shelf configuration, as described in [Disk shelf connection](#).
2. Create and configure the required aggregates, as described in [Aggregate configuration](#).
3. Create a storage virtual machine (SVM), as described in [Storage virtual machine configuration](#).
4. Create logical interfaces (LIFs), as described in [Logical interface configuration](#).
5. Create initiator groups (igroups) with worldwide names (WWNs) of HANA servers as described in the section [xref:./bp/hana-fas-fc-storage-controller-setup.html#initiator-groups](#) [Initiator groups](#).
6. Create and configure volumes and LUNs within the aggregates as described in the section [Single Host Setup](#) for single hosts or in section [Multiple Host Setup](#) for multiple hosts

With HDDs, a maximum of two DS2246 disk shelves or four DS224C disk shelves can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair.



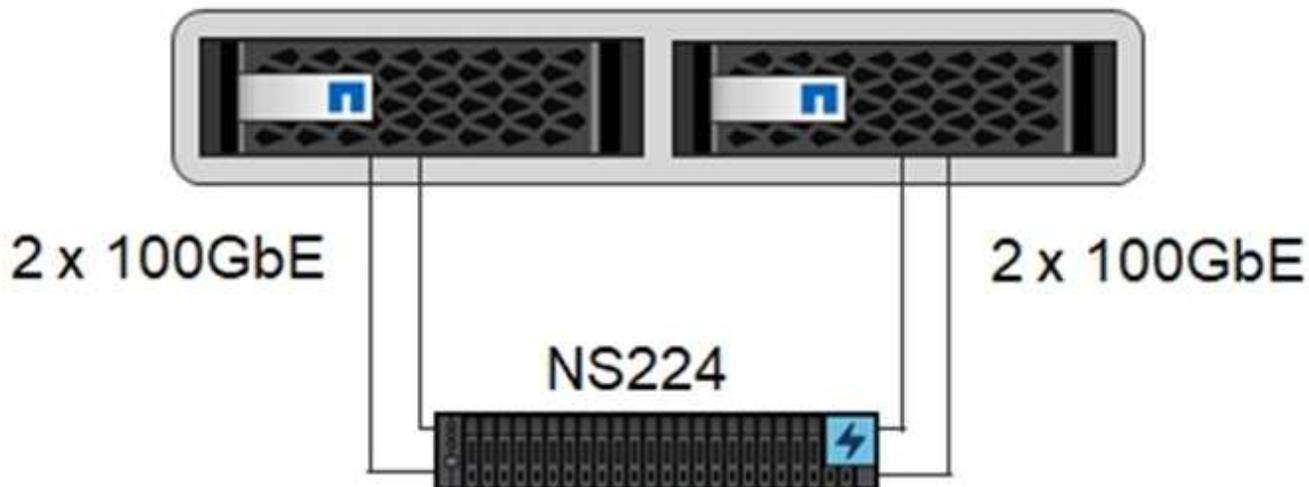
With SSDs, a maximum of one disk shelf can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair. With the DS224C disk shelf, quad-path SAS cabling can also be used but is not required.



### NVMe disk shelves

Each NS224 NVMe disk shelf is connected with two 100GbE ports per controller, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair.

## Storage controller HA pair

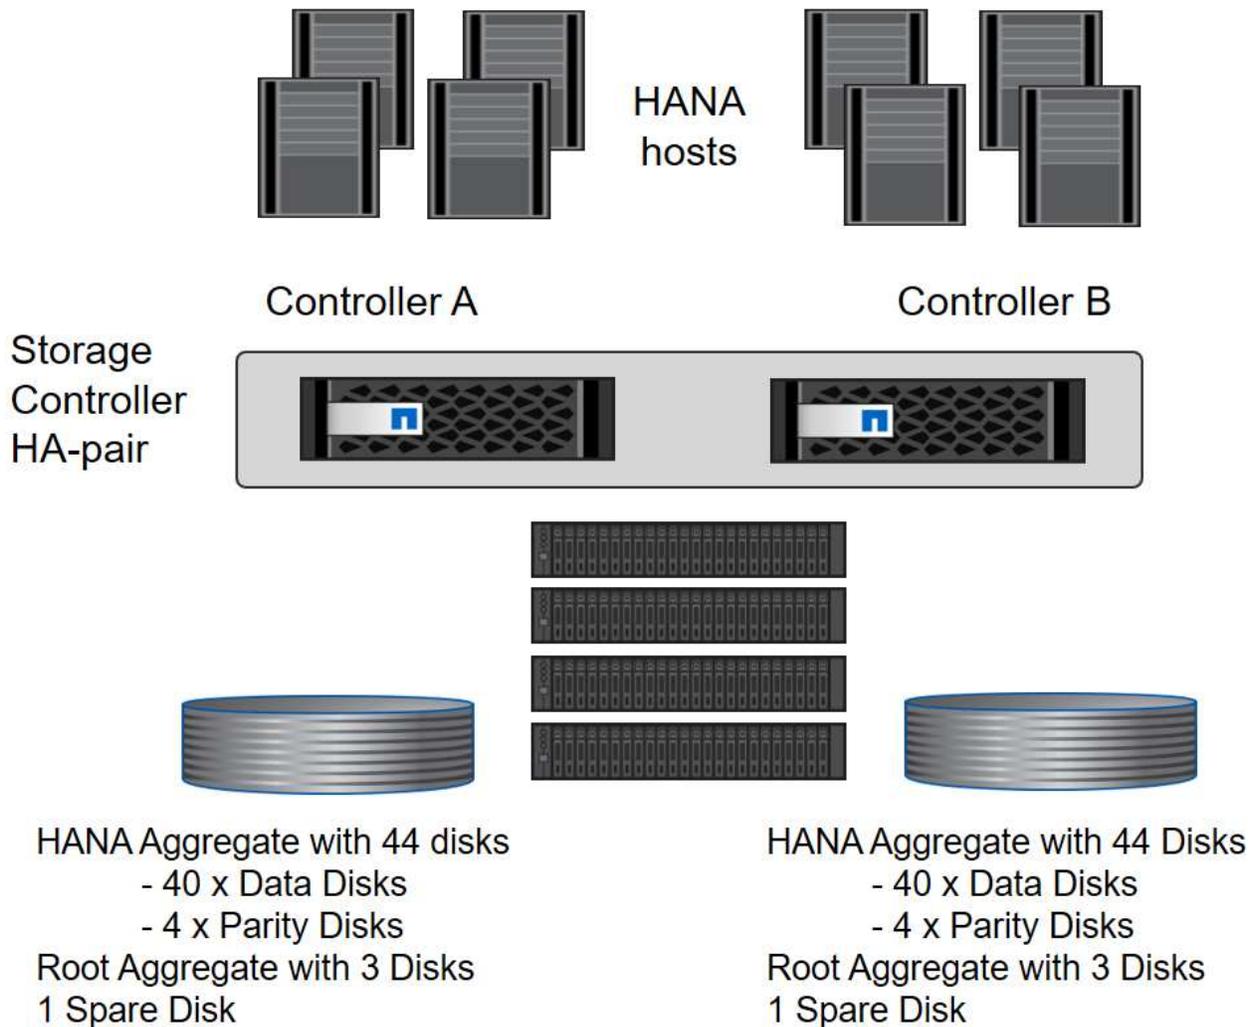


### Aggregate configuration

In general, you must configure two aggregates per controller, independent of which disk shelf or disk technology (SSD or HDD) is used. This step is necessary so that you can use all available controller resources. For FAS 2000 series systems, one data aggregate is sufficient.

### Aggregate configuration with HDDs

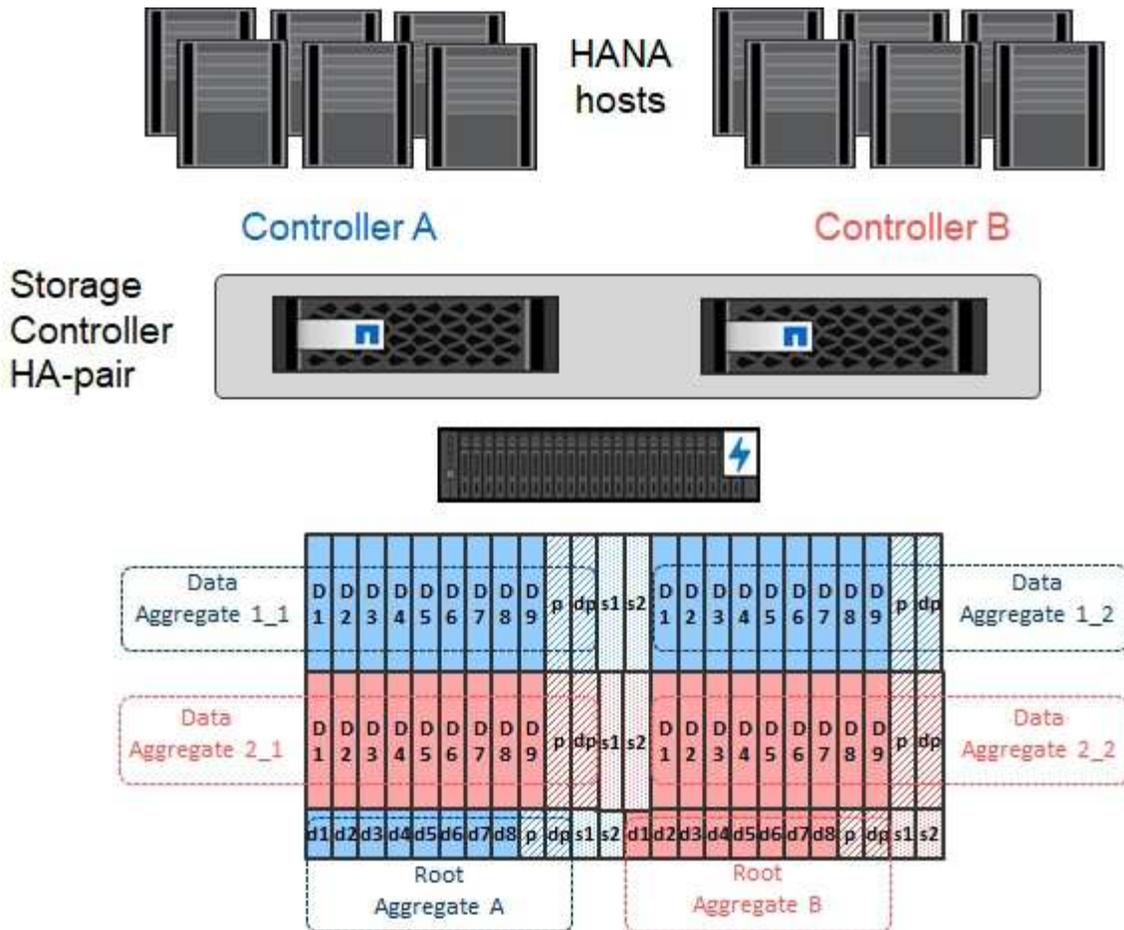
The following figure shows a configuration for eight SAP HANA hosts. Four SAP HANA hosts are attached to each storage controller. Two separate aggregates, one at each storage controller, are configured. Each aggregate is configured with  $4 \times 10 = 40$  data disks (HDDs).



### Aggregate configuration with SDD-only systems

In general, two aggregates per controller must be configured, independently of which disk shelf or disk technology (SSDs or HDDs) is used.

The following figure shows a configuration of 12 SAP HANA hosts running on a 12Gb SAS shelf configured with ADPv2. Six SAP HANA hosts are attached to each storage controller. Four separate aggregates, two at each storage controller, are configured. Each aggregate is configured with 11 disks with nine data and two parity disk partitions. For each controller, two spare partitions are available.



### Storage virtual machine configuration

Multiple-host SAP landscapes with SAP HANA databases can use a single SVM. An SVM can also be assigned to each SAP landscape if necessary in case they are managed by different teams within a company. The screenshots and command outputs in this document use an SVM named `hana`.

### Logical interface configuration

Within the storage cluster configuration, one network interface (LIF) must be created and assigned to a dedicated FCP port. If, for example, four FCP ports are required for performance reasons, four LIFs must be created. The following figure shows a screenshot of the eight LIFs that were configured on the SVM.

NetApp ONTAP System Manager | a400-sapcc

Search actions, objects, and pages

Dashboard

Insights

Storage

**Network**

Overview

Ethernet ports

FC ports

Events & jobs

Protection

Hosts

Cluster

**IPspaces**

+ Add

Cluster	Broadcast domains
Cluster	Cluster
Default	Storage VMs BlueXPDR_SVM1 ,C30-HANA ,TCP-NVME ,abhi-a400 , hana-A400 ,infra-svm ,svm-dietmare-misc ,test_rdma Broadcast domains Default ,NFS ,NFS2 ,rdma ,vlan-data ,vlan-log

**Broadcast domains** [Learn more](#)

+ Add

Cluster	9000 MTU	IPspace: Cluster
a400-sapcc-01	e3a e3b	a400-sapcc-02 e3a e3b
Default	1500 MTU	IPspace: Default a400-sapcc-01 e0M a400-sapcc-02 e0M
NFS	9000 MTU	IPspace: Default a400-sapcc-01 a0a a400-sapcc-02 a0a
NFS2	9000 MTU	IPspace: Default

**Network interfaces** **Subnets**

+ Add

Name	Status	Storage VM	IPspace	Address	Current node	Current port	Portset	Protocols	Throughput (M)
lif_hana_345	✔	hana-A400	20:0b:d0:39:ea:2ef9:41	a400-sapcc-01	1a	FC	0		
lif_hana_965	✔	hana-A400	20:0c:d0:39:ea:2ef9:41	a400-sapcc-01	1b	FC	0		
lif_hana_205	✔	hana-A400	20:0d:d0:39:ea:2ef9:41	a400-sapcc-01	1c	FC	0		
lif_hana_314	✔	hana-A400	20:0e:d0:39:ea:2ef9:41	a400-sapcc-01	1d	FC	0		
lif_hana_908	✔	hana-A400	20:0f:d0:39:ea:2ef9:41	a400-sapcc-02	1a	FC	0		
lif_hana_726	✔	hana-A400	20:10:d0:39:ea:2ef9:41	a400-sapcc-02	1b	FC	0		
lif_hana_521	✔	hana-A400	20:11:d0:39:ea:2ef9:41	a400-sapcc-02	1c	FC	0		
lif_hana_946	✔	hana-A400	20:12:d0:39:ea:2ef9:41	a400-sapcc-02	1d	FC	0		

During SVM creation with ONTAP 9 System Manager, all the required physical FCP ports can be selected, and one LIF per physical port is created automatically.

The following figure depicts the creation of SVM and LIFs with ONTAP System Manager.

NetApp ONTAP System Manager | a400-sapcc

Search actions, objects, and pages

Dashboard

Insights

Storage

- Overview
- Volumes
- LUNs
- NVMe namespaces
- Consistency groups
- Shares
- Qtrees
- Quotas
- Storage VMs
- Tiers

Network

Events & jobs

Protection

Hosts

Cluster

### Add storage VM

Storage VM name: hana

Access protocol: SMB/CIFS, NFS, iSCSI, **FC**, NVMe

Enable FC

Configure FC ports

Nodes	1a	1b	1c	1d
a400-sapcc-01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
a400-sapcc-02	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Storage VM administration

Enable maximum capacity limit  
The maximum capacity that all volumes in this storage VM can allocate. [Learn More](#)

Manage administrator account

User name: vsadmin

Password: .....

Confirm password: .....

Add a network interface for storage VM management.

Node: a400-sapcc-01

IP address: 10.10.10.10

Subnet mask: 255.255.255.0

Save Cancel

## Initiator groups

An igroup can be configured for each server or for a group of servers that require access to a LUN. The igroup configuration requires the worldwide port names (WWPNs) of the servers.

Using the `sanlun` tool, run the following command to obtain the WWPNs of each SAP HANA host:

```
stlrx300s8-6:~ # sanlun fcp show adapter
/sbin/udevadm
/sbin/udevadm

host0 ..... WWPN:2100000e1e163700
host1 ..... WWPN:2100000e1e163701
```



The `sanlun` tool is part of the NetApp Host Utilities and must be installed on each SAP HANA host. More details can be found in section [Host setup](#).

The initiator groups can be created using the CLI of the ONTAP Cluster.

```
lun igroup create -igroup <igroup name> -protocol fcp -ostype linux
-initiator <list of initiators> -vserver <SVM name>
```

## Single host

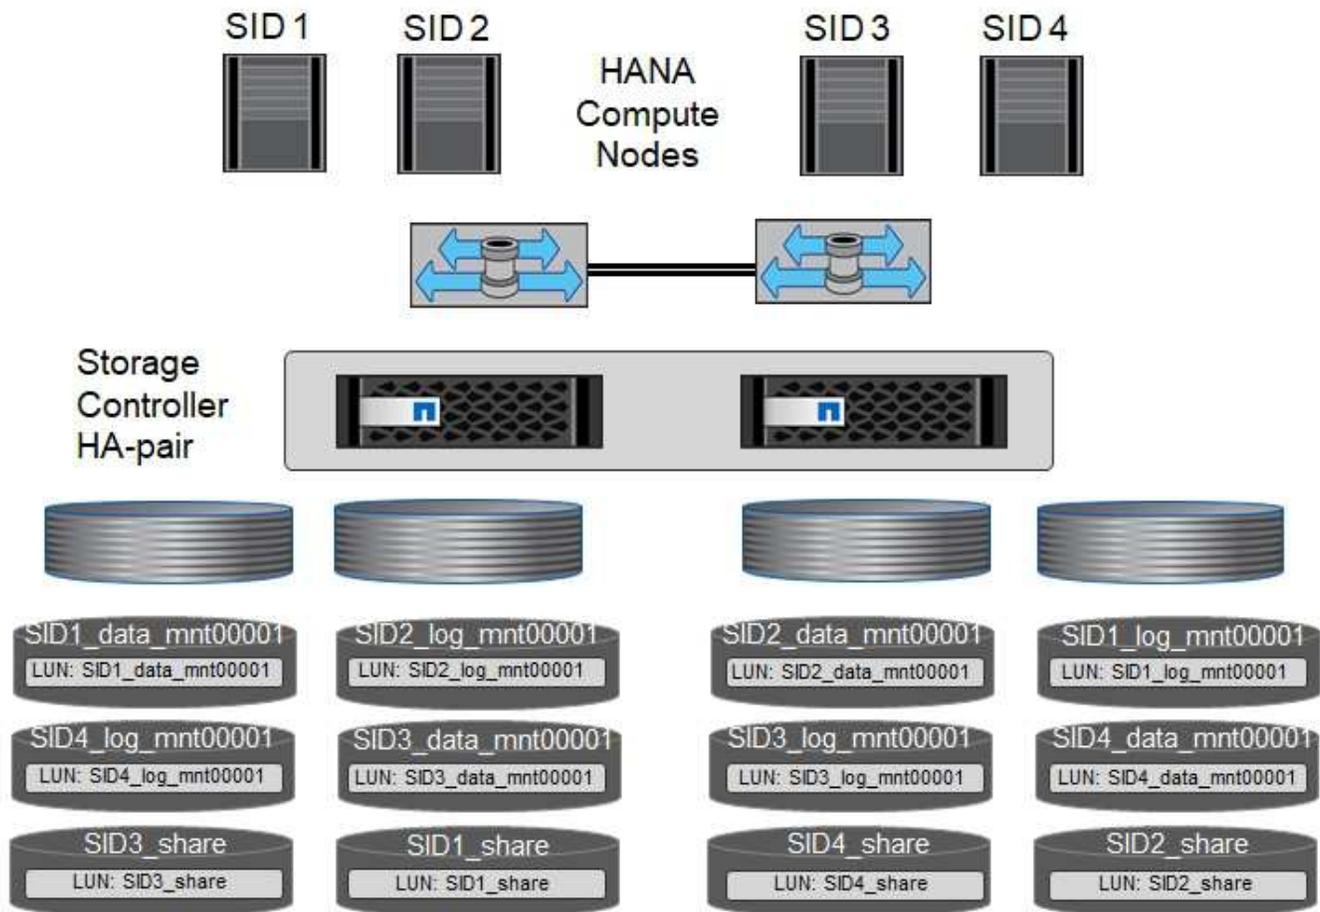
This section describes the configuration of the NetApp storage system specific to SAP HANA single-host systems

### Volume and LUN configuration for SAP HANA single-host systems

The following figure shows the volume configuration of four single-host SAP HANA systems. The data and log volumes of each SAP HANA system are distributed to different storage controllers. For example, volume `SID1_data_mnt00001` is configured on controller A and volume `SID1_log_mnt00001` is configured on controller B. Within each volume, a single LUN is configured.



If only one storage controller of a high-availability (HA) pair is used for the SAP HANA systems, data volumes and log volumes can also be stored on the same storage controller.



For each SAP HANA host, a data volume, a log volume, and a volume for /hana/shared are configured. The following table shows an example configuration with four SAP HANA single-host systems.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data, log, and shared volumes for system SID1	Data volume: SID1_data_mnt00001	Shared volume: SID1_shared	–	Log volume: SID1_log_mnt00001
Data, log, and shared volumes for system SID2	–	Log volume: SID2_log_mnt00001	Data volume: SID2_data_mnt00001	Shared volume: SID2_shared
Data, log, and shared volumes for system SID3	Shared volume: SID3_shared	Data volume: SID3_data_mnt00001	Log volume: SID3_log_mnt00001	–
Data, log, and shared volumes for system SID4	Log volume: SID4_log_mnt00001	–	Shared volume: SID4_shared	Data volume: SID4_data_mnt00001

The next table shows an example of the mount point configuration for a single-host system.

LUN	Mount point at HANA host	Note
SID1_data_mnt00001	/hana/data/SID1/mnt00001	Mounted using /etc/fstab entry

LUN	Mount point at HANA host	Note
SID1_log_mnt00001	/hana/log/SID1/mnt00001	Mounted using /etc/fstab entry
SID1_shared	/hana/shared/SID1	Mounted using /etc/fstab entry



With the described configuration, the `/usr/sap/SID1` directory in which the default home directory of user `SID1adm` is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `SID1_shared` volume for the `/usr/sap/SID1` directory so that all file systems are on the central storage.

## Volume and LUN configuration for SAP HANA single-host systems using Linux LVM

The Linux LVM can be used to increase performance and to address LUN size limitations. The different LUNs of an LVM volume group should be stored within a different aggregate and at a different controller. The following table shows an example for two LUNs per volume group.



It is not necessary to use LVM with multiple LUNs to fulfil the SAP HANA KPIs, but it is recommended

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data, log, and shared volumes for LVM based system	Data volume: SID1_data_mnt00001	Shared volume: SID1_shared Log2 volume: SID1_log2_mnt00001	Data2 volume: SID1_data2_mnt00001	Log volume: SID1_log_mnt00001



With the described configuration, the `/usr/sap/SID1` directory in which the default home directory of user `SID1adm` is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `SID1_shared` volume for the `/usr/sap/SID1` directory so that all file systems are on the central storage.

## Volume options

The volume options listed in the following table must be verified and set on all volumes used for SAP HANA.

Action	ONTAP 9
Disable automatic Snapshot copies	<code>vol modify -vserver &lt;vserver-name&gt; -volume &lt;volname&gt; -snapshot-policy none</code>
Disable visibility of Snapshot directory	<code>vol modify -vserver &lt;vserver-name&gt; -volume &lt;volname&gt; -snapdir-access false</code>

## Creating LUNs, volumes, and mapping LUNs to initiator groups

You can use NetApp ONTAP System Manager to create storage volumes and LUNs and the map them to the igroups of the servers and the ONTAP CLI. This guide describes the usage of the CLI.

## Creating LUNs, volumes, and mapping LUNs to igroups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a SAP HANA single host system with SID FC5 using LVM and two LUNs per LVM volume group:

### 1. Create all necessary volumes.

```
vol create -volume FC5_data_mnt00001 -aggregate aggr1_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00001 -aggregate aggr1_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00001 -aggregate aggr1_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00001 -aggregate aggr1_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_shared -aggregate aggr1_1 -size 512g -state
online -policy default -snapshot-policy none -junction-path /FC5_shared
-encrypt false -space-guarantee none
```

### 2. Create all LUNs.

```
lun create -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
```

### 3. Create the initiator group for all ports belonging to sythe hosts of FC5.

```
lun igroup create -igroup HANA-FC5 -protocol fcp -ostype linux
-initiator 10000090fadcc5fa,10000090fadcc5fb -vserver hana
```

### 4. Map all LUNs to created initiator group.

```
lun map -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -igroup HANA-FC5
```

## Multiple hosts

This section describes the configuration of the NetApp storage system specific to SAP HANA multiple-hosts systems

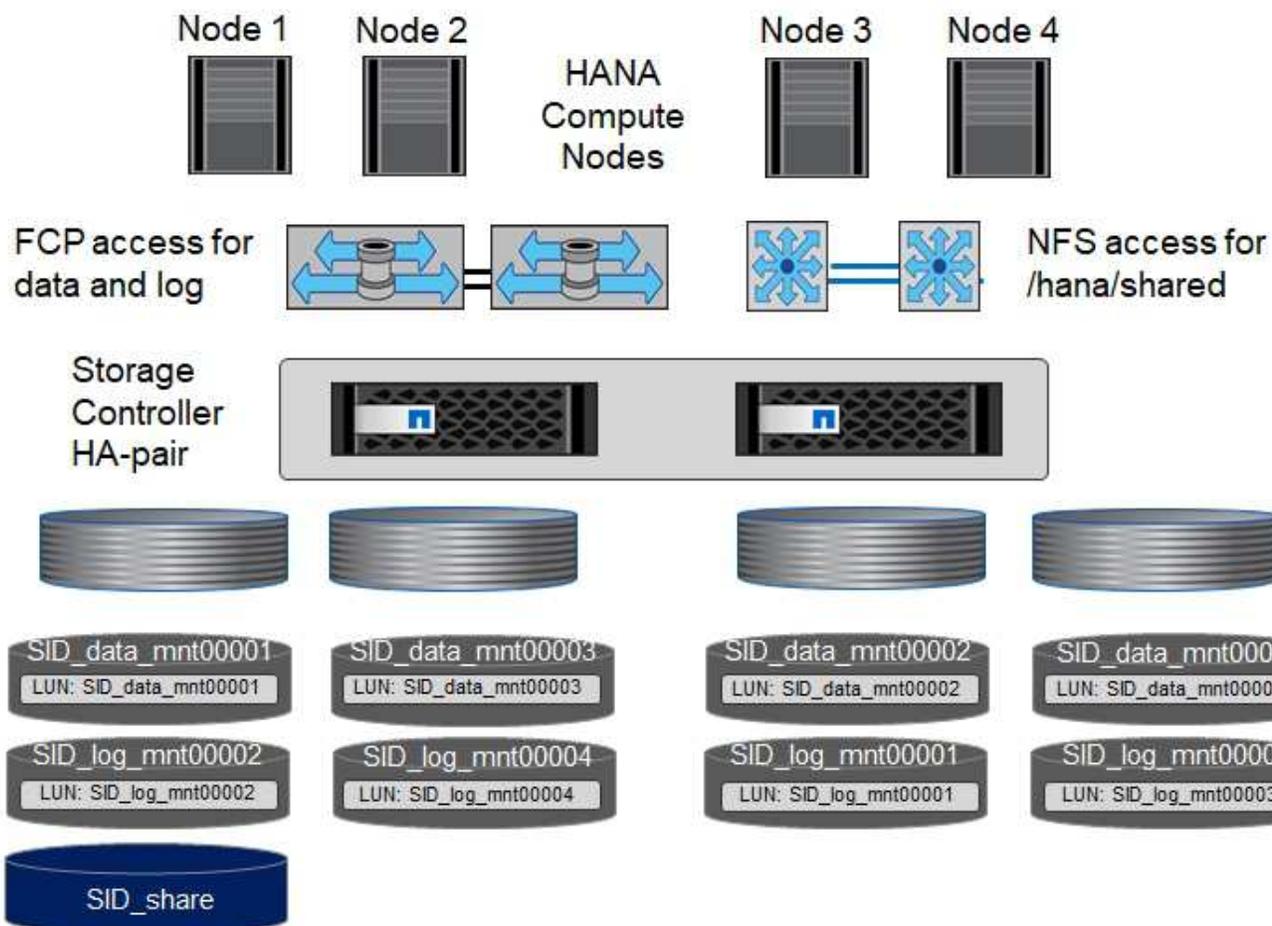
### Volume and LUN configuration for SAP HANA multiple-host systems

The following figure shows the volume configuration of a 4+1 multiple-host SAP HANA system. The data volumes and log volumes of each SAP HANA host are distributed to different storage controllers. For example, the volume `SID_data_mnt00001` is configured on controller A and the volume `SID_log_mnt00001` is configured on controller B. One LUN is configured within each volume.

The `/hana/shared` volume must be accessible by all HANA hosts and is therefore exported by using NFS. Even though there are no specific performance KPIs for the `/hana/shared` file system, NetApp recommends using a 10Gb Ethernet connection.



If only one storage controller of an HA pair is used for the SAP HANA system, data and log volumes can also be stored on the same storage controller.



For each SAP HANA host, a data volume and a log volume are created. The `/hana/shared` volume is used by all hosts of the SAP HANA system. The following figure shows an example configuration for a 4+1 multiple-host SAP HANA system.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	–	Log volume: SID_log_mnt00001	–
Data and log volumes for node 2	Log volume: SID_log_mnt00002	–	Data volume: SID_data_mnt00002	–
Data and log volumes for node 3	–	Data volume: SID_data_mnt00003	–	Log volume: SID_log_mnt00003
Data and log volumes for node 4	–	Log volume: SID_log_mnt00004	–	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared	–	–	–

The next table shows the configuration and the mount points of a multiple-host system with four active SAP HANA hosts.

LUN or Volume	Mount point at SAP HANA host	Note
LUN: SID_data_mnt00001	/hana/data/SID/mnt00001	Mounted using storage connector
LUN: SID_log_mnt00001	/hana/log/SID/mnt00001	Mounted using storage connector
LUN: SID_data_mnt00002	/hana/data/SID/mnt00002	Mounted using storage connector
LUN: SID_log_mnt00002	/hana/log/SID/mnt00002	Mounted using storage connector
LUN: SID_data_mnt00003	/hana/data/SID/mnt00003	Mounted using storage connector
LUN: SID_log_mnt00003	/hana/log/SID/mnt00003	Mounted using storage connector
LUN: SID_data_mnt00004	/hana/data/SID/mnt00004	Mounted using storage connector
LUN: SID_log_mnt00004	/hana/log/SID/mnt00004	Mounted using storage connector
Volume: SID_shared	/hana/shared/SID	Mounted at all hosts using NFS and /etc/fstab entry



With the described configuration, the `/usr/sap/SID` directory in which the default home directory of user `SIDadm` is stored is on the local disk for each HANA host. In a disaster recovery setup with disk-based replication, NetApp recommends creating four additional subdirectories in the `SID_shared` volume for the `/usr/sap/SID` file system so that each database host has all its file systems on the central storage.

### Volume and LUN configuration for SAP HANA multiple-host systems using Linux LVM

The Linux LVM can be used to increase performance and to address LUN size limitations. The different LUNs of an LVM volume group should be stored within a different aggregate and at a different controller. The following table shows an example for two LUNs per volume group for a 2+1 SAP HANA multiple host system.



It is not necessary to use LVM to combine several LUN to fulfil the SAP HANA KPIs, but it is recommended.

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	Log2 volume: SID_log2_mnt00001	Log volume: SID_log_mnt00001	Data2 volume: SID_data2_mnt00001
Data and log volumes for node 2	Log2 volume: SID_log2_mnt00002	Data volume: SID_data_mnt00002	Data2 volume: SID_data2_mnt00002	Log volume: SID_log_mnt00002
Shared volume for all hosts	Shared volume: SID_shared	—	—	—

### Volume options

The volume options listed in the following table must be verified and set on all volumes used for SAP HANA.

Action	ONTAP 9
Disable automatic Snapshot copies	vol modify -vserver <vserver-name> -volume <volname> -snapshot-policy none
Disable visibility of Snapshot directory	vol modify -vserver <vserver-name> -volume <volname> -snapdir-access false

## Creating LUNs, volumes, and mapping LUNs to initiator groups

You can use NetApp ONTAP System Manager to create storage volumes and LUNs and the map them to the igroups of the servers and the ONTAP CLI. This guide describes the usage of the CLI.

## Creating LUNs, volumes, and mapping LUNs to igroups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a 2+1 SAP HANA multiple host system with SID FC5 using LVM and two LUNs per LVM volume group.

1. Create all necessary volumes.

```
vol create -volume FC5_data_mnt00001 -aggregate aggr1_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00002 -aggregate aggr2_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00001 -aggregate aggr1_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data_mnt00002 -aggregate aggr2_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00001 -aggregate aggr1_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00002 -aggregate aggr2_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00001 -aggregate aggr1_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00002 -aggregate aggr2_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_shared -aggregate aggr1_1 -size 512g -state
online -policy default -snapshot-policy none -junction-path /FC5_shared
-encrypt false -space-guarantee none
```

## 2. Create all LUNs.

```
lun create -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data_mnt00002/FC5_data_mnt00002 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00002/FC5_data2_mnt00002 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log_mnt00002/FC5_log_mnt00002 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00002/FC5_log2_mnt00002 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
```

## 3. Create the igroup for all servers belonging to system FC5.

```
lun igroup create -igroup HANA-FC5 -protocol fcp -ostype linux
-initiator 10000090fadcc5fa,10000090fadcc5fb,
10000090fadcc5c1,10000090fadcc5c2, 10000090fadcc5c3,10000090fadcc5c4
-vserver hana
```

## 4. Map all LUNs to the created igroup.

```
lun map -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_data_mnt00002/FC5_data_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00002/FC5_data2_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt00002/FC5_log_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00002/FC5_log2_mnt00002 -igroup HANA-FC5
```

## SAP HANA storage connector API

A storage connector is required only in multiple-host environments that have failover capabilities. In multiple-host setups, SAP HANA provides high-availability functionality so that an SAP HANA database host can fail over to a standby host. In this case, the LUNs of the failed host are accessed and used by the standby host. The storage connector is used to make sure that a storage partition can be actively accessed by only one database host at a time.

In SAP HANA multiple-host configurations with NetApp storage, the standard storage connector delivered by SAP is used. The “SAP HANA FC Storage Connector Admin Guide” can be found as an attachment to [SAP note 1900823](#).

## Host setup

Before setting up the host, NetApp SAN Host Utilities must be downloaded from the [NetApp Support](#) site and installed on the HANA servers. The Host Utility documentation includes information about additional software that must be installed depending on the FCP HBA used.

The documentation also contains information about multipath configurations that are specific to the Linux version used. This document covers the required configuration steps for SLES 15 and Red Hat Enterprise Linux 7.6 or higher, as described in the [Linux Host Utilities 7.1 Installation and Setup Guide](#).

## Configure multipathing



Steps 1 to 6 must be performed on all worker and standby hosts in the SAP HANA multiple-host configuration.

To configure multipathing, complete the following steps:

1. Run the Linux `rescan-scsi-bus.sh -a` command on each server to discover new LUNs.
2. Run the `sanlun lun show` command and verify that all required LUNs are visible. The following example shows the `sanlun lun show` command output for a 2+1 multiple-host HANA system with two data LUNs and two log LUNs. The output shows the LUNs and the corresponding device files, such as LUN

SS3\_data\_mnt00001 and the device file /dev/sdag. Each LUN has eight FC paths from the host to the storage controllers.

```

sapcc-hana-tst:~ # sanlun lun show
controller(7mode/E-Series)/                               device
host                lun
vserver(cDOT/FlashRay)    lun-pathname                filename
adapter    protocol    size    product
-----
svm1                FC5_log2_mnt00002                /dev/sdbb
host21    FCP        500g    cDOT
svm1                FC5_log_mnt00002                /dev/sdba
host21    FCP        500g    cDOT
svm1                FC5_log2_mnt00001                /dev/sdaz
host21    FCP        500g    cDOT
svm1                FC5_log_mnt00001                /dev/sday
host21    FCP        500g    cDOT
svm1                FC5_data2_mnt00002                /dev/sdax
host21    FCP        1t        cDOT
svm1                FC5_data_mnt00002                /dev/sdaw
host21    FCP        1t        cDOT
svm1                FC5_data2_mnt00001                /dev/sdav
host21    FCP        1t        cDOT
svm1                FC5_data_mnt00001                /dev/sdau
host21    FCP        1t        cDOT
svm1                FC5_log2_mnt00002                /dev/sdat
host21    FCP        500g    cDOT
svm1                FC5_log_mnt00002                /dev/sdas
host21    FCP        500g    cDOT
svm1                FC5_log2_mnt00001                /dev/sdar
host21    FCP        500g    cDOT
svm1                FC5_log_mnt00001                /dev/sdaq
host21    FCP        500g    cDOT
svm1                FC5_data2_mnt00002                /dev/sdap
host21    FCP        1t        cDOT
svm1                FC5_data_mnt00002                /dev/sdao
host21    FCP        1t        cDOT
svm1                FC5_data2_mnt00001                /dev/sdan
host21    FCP        1t        cDOT
svm1                FC5_data_mnt00001                /dev/sdam
host21    FCP        1t        cDOT
svm1                FC5_log2_mnt00002                /dev/sdal
host20    FCP        500g    cDOT
svm1                FC5_log_mnt00002                /dev/sdak
host20    FCP        500g    cDOT

```

```

svm1          FC5_log2_mnt00001          /dev/sdaj
host20        FCP          500g          cDOT
svm1          FC5_log_mnt00001          /dev/sdai
host20        FCP          500g          cDOT
svm1          FC5_data2_mnt00002        /dev/sdah
host20        FCP          1t          cDOT
svm1          FC5_data_mnt00002        /dev/sdag
host20        FCP          1t          cDOT
svm1          FC5_data2_mnt00001        /dev/sdaf
host20        FCP          1t          cDOT
svm1          FC5_data_mnt00001        /dev/sdae
host20        FCP          1t          cDOT
svm1          FC5_log2_mnt00002        /dev/sdad
host20        FCP          500g          cDOT
svm1          FC5_log_mnt00002        /dev/sdac
host20        FCP          500g          cDOT
svm1          FC5_log2_mnt00001        /dev/sdab
host20        FCP          500g          cDOT
svm1          FC5_log_mnt00001        /dev/sdaa
host20        FCP          500g          cDOT
svm1          FC5_data2_mnt00002        /dev/sdz
host20        FCP          1t          cDOT
svm1          FC5_data_mnt00002        /dev/sdy
host20        FCP          1t          cDOT
svm1          FC5_data2_mnt00001        /dev/sdx
host20        FCP          1t          cDOT
svm1          FC5_data_mnt00001        /dev/sdw
host20        FCP          1t          cDOT

```

3. Run the `multipath -r` and `multipath -ll` command to get the worldwide identifiers (WWIDs) for the device file names.



In this example, there are eight LUNs.

```

sapcc-hana-tst:~ # multipath -r
sapcc-hana-tst:~ # multipath -ll
3600a098038314e63492b59326b4b786d dm-7 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:2 sdaf 65:240 active ready running
  |- 20:0:5:2 sdx 65:112 active ready running
  |- 21:0:4:2 sdav 66:240 active ready running
  `-- 21:0:6:2 sdan 66:112 active ready running
3600a098038314e63492b59326b4b786e dm-9 NETAPP,LUN C-Mode

```

```

size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:4 sdah 66:16 active ready running
  |- 20:0:5:4 sdz 65:144 active ready running
  |- 21:0:4:4 sdax 67:16 active ready running
  `- 21:0:6:4 sdap 66:144 active ready running
3600a098038314e63492b59326b4b786f dm-11 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:6 sdaj 66:48 active ready running
  |- 20:0:5:6 sdab 65:176 active ready running
  |- 21:0:4:6 sdaz 67:48 active ready running
  `- 21:0:6:6 sdar 66:176 active ready running
3600a098038314e63492b59326b4b7870 dm-13 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:8 sdal 66:80 active ready running
  |- 20:0:5:8 sdad 65:208 active ready running
  |- 21:0:4:8 sdbb 67:80 active ready running
  `- 21:0:6:8 sdat 66:208 active ready running
3600a098038314e63532459326d495a64 dm-6 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:1 sdae 65:224 active ready running
  |- 20:0:5:1 sdw 65:96 active ready running
  |- 21:0:4:1 sdau 66:224 active ready running
  `- 21:0:6:1 sdam 66:96 active ready running
3600a098038314e63532459326d495a65 dm-8 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:3 sdag 66:0 active ready running
  |- 20:0:5:3 sdy 65:128 active ready running
  |- 21:0:4:3 sdaw 67:0 active ready running
  `- 21:0:6:3 sdao 66:128 active ready running
3600a098038314e63532459326d495a66 dm-10 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:5 sdai 66:32 active ready running
  |- 20:0:5:5 sdaa 65:160 active ready running
  |- 21:0:4:5 sday 67:32 active ready running

```

```
`- 21:0:6:5 sdaq 66:160 active ready running
3600a098038314e63532459326d495a67 dm-12 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
|- 20:0:4:7 sdak 66:64 active ready running
|- 20:0:5:7 sdac 65:192 active ready running
|- 21:0:4:7 sdba 67:64 active ready running
`- 21:0:6:7 sdas 66:192 active ready running
```

4. Edit the `/etc/multipath.conf` file and add the WWIDs and alias names.



The example output shows the content of the `/etc/multipath.conf` file, which includes alias names for the four LUNs of a 2+1 multiple-host system. If there is no `multipath.conf` file available, you can create one by running the following command: `multipath -T > /etc/multipath.conf`.

```

sapcc-hana-tst:/ # cat /etc/multipath.conf
multipaths {
    multipath {
        wwid      3600a098038314e63492b59326b4b786d
        alias     svm1-FC5_data2_mnt00001
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786e
        alias     svm1-FC5_data2_mnt00002
    }
    multipath {
        wwid      3600a098038314e63532459326d495a64
        alias     svm1-FC5_data_mnt00001
    }
    multipath {
        wwid      3600a098038314e63532459326d495a65
        alias     svm1-FC5_data_mnt00002
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786f
        alias     svm1-FC5_log2_mnt00001
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b7870
        alias     svm1-FC5_log2_mnt00002
    }
    multipath {
        wwid      3600a098038314e63532459326d495a66
        alias     svm1-FC5_log_mnt00001
    }
    multipath {
        wwid      3600a098038314e63532459326d495a67
        alias     svm1-FC5_log_mnt00002
    }
}

```

5. Run the `multipath -r` command to reload the device map.
6. Verify the configuration by running the `multipath -ll` command to list all the LUNs, alias names, and active and standby paths.



The following example output shows the output of a 2+1 multiple-host HANA system with two data and two log LUNs.

```

sapcc-hana-tst:~ # multipath -ll
hsvm1-FC5_data2_mnt00001 (3600a098038314e63492b59326b4b786d) dm-7
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:2 sdaf 65:240 active ready running
  |- 20:0:5:2 sdx 65:112 active ready running
  |- 21:0:4:2 sdav 66:240 active ready running
  `-- 21:0:6:2 sdan 66:112 active ready running
svm1-FC5_data2_mnt00002 (3600a098038314e63492b59326b4b786e) dm-9
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:4 sdah 66:16 active ready running
  |- 20:0:5:4 sdz 65:144 active ready running
  |- 21:0:4:4 sdax 67:16 active ready running
  `-- 21:0:6:4 sdap 66:144 active ready running
svm1-FC5_data_mnt00001 (3600a098038314e63532459326d495a64) dm-6
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:1 sdae 65:224 active ready running
  |- 20:0:5:1 sdw 65:96 active ready running
  |- 21:0:4:1 sdau 66:224 active ready running
  `-- 21:0:6:1 sdam 66:96 active ready running
svm1-FC5_data_mnt00002 (3600a098038314e63532459326d495a65) dm-8
NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:3 sdag 66:0 active ready running
  |- 20:0:5:3 sdy 65:128 active ready running
  |- 21:0:4:3 sdaw 67:0 active ready running
  `-- 21:0:6:3 sdao 66:128 active ready running
svm1-FC5_log2_mnt00001 (3600a098038314e63492b59326b4b786f) dm-11
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:6 sdaj 66:48 active ready running
  |- 20:0:5:6 sdab 65:176 active ready running
  |- 21:0:4:6 sdaz 67:48 active ready running
  `-- 21:0:6:6 sdar 66:176 active ready running

```

```

svm1-FC5_log2_mnt00002 (3600a098038314e63492b59326b4b7870) dm-13
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:8 sdal 66:80 active ready running
  |- 20:0:5:8 sdad 65:208 active ready running
  |- 21:0:4:8 sdbb 67:80 active ready running
  `-- 21:0:6:8 sdat 66:208 active ready running
svm1-FC5_log_mnt00001 (3600a098038314e63532459326d495a66) dm-10
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:5 sdai 66:32 active ready running
  |- 20:0:5:5 sdaa 65:160 active ready running
  |- 21:0:4:5 sday 67:32 active ready running
  `-- 21:0:6:5 sdaq 66:160 active ready running
svm1-FC5_log_mnt00002 (3600a098038314e63532459326d495a67) dm-12
NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 20:0:4:7 sdak 66:64 active ready running
  |- 20:0:5:7 sdac 65:192 active ready running
  |- 21:0:4:7 sdba 67:64 active ready running
  `-- 21:0:6:7 sdas 66:192 active ready running

```

## Single host setup

This chapter describes the setup of an SAP HANA single host using Linux LVM.

### LUN configuration for SAP HANA single-host systems

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

Logical volume/LUN	Mount point at SAP HANA host	Note
LV: FC5_data_mnt0000-vol	/hana/data/FC51/mnt00001	Mounted using /etc/fstab entry
LV: FC5_log_mnt00001-vol	/hana/log/FC5/mnt00001	Mounted using /etc/fstab entry
LUN: FC5_shared	/hana/shared/FC5	Mounted using /etc/fstab entry



With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user `FC5adm` is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `FC5_shared` volume for the `/usr/sap/FC5` directory so that all file systems are on the central storage.

## Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```
pvcreate /dev/mapper/hana-FC5_data_mnt00001
pvcreate /dev/mapper/hana-FC5_data2_mnt00001
pvcreate /dev/mapper/hana-FC5_log_mnt00001
pvcreate /dev/mapper/hana-FC5_log2_mnt00001
```

2. Create the volume groups for each data and log partition.

```
vgcreate FC5_data_mnt00001 /dev/mapper/hana-FC5_data_mnt00001
/dev/mapper/hana-FC5_data2_mnt00001
vgcreate FC5_log_mnt00001 /dev/mapper/hana-FC5_log_mnt00001
/dev/mapper/hana-FC5_log2_mnt00001
```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00001
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00001
```

4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```
modprobe dm_mod
pvscan
vgscan
lvscan
```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```
vgchange -a y
```

## Create file systems

Create the XFS file system on all data and log logical volumes and the hana shared LUN.

```
mkfs.xfs FC5_data_mnt00001-vol
mkfs.xfs FC5_log_mnt00001-vol
mkfs.xfs /dev/mapper/svml-FC5_shared
```

## Create mount points

Create the required mount point directories, and set the permissions on the database host:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

## Mount file systems

To mount file systems during system boot using the `/etc/fstab` configuration file, add the required file systems to the `/etc/fstab` configuration file:

```
# cat /etc/fstab
/dev/mapper/hana-FC5_shared /hana/shared xfs defaults 0 0
/dev/mapper/FC5_log_mnt00001-vol /hana/log/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/FC5_data_mnt00001-vol /hana/data/FC5/mnt00001 xfs
relatime,inode64 0 0
```



The XFS file systems for the data and log LUNs must be mounted with the `relatime` and `inode64` mount options.

To mount the file systems, run the `mount -a` command at the host.

## Multiple hosts setup

This chapter describes the setup of a 2+1 SAP HANA multiple host system as example.

### LUN configuration for SAP HANA multiple-hosts systems

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

Logical volume (LV) or volume	Mount point at SAP HANA host	Note
LV: FC5_data_mnt00001-vol	/hana/data/FC5/mnt00001	Mounted using storage connector
LV: FC5_log_mnt00001-vol	/hana/log/FC5/mnt00001	Mounted using storage connector
LV: FC5_data_mnt00002-vol	/hana/data/FC5/mnt00002	Mounted using storage connector
LV: FC5_log_mnt00002-vol	/hana/log/FC5/mnt00002	Mounted using storage connector
Volume: FC5_shared	/hana/shared	Mounted at all hosts using NFS and /etc/fstab entry



With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user FC5adm is stored, is on the local disk for each HANA host. In a disaster recovery setup with disk-based replication, NetApp recommends creating four additional subdirectories in the `FC5_shared` volume for the `/usr/sap/FC5` file system so that each database host has all its file systems on the central storage.

## Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```
pvccreate /dev/mapper/hana-FC5_data_mnt00001
pvccreate /dev/mapper/hana-FC5_data2_mnt00001
pvccreate /dev/mapper/hana-FC5_data_mnt00002
pvccreate /dev/mapper/hana-FC5_data2_mnt00002
pvccreate /dev/mapper/hana-FC5_log_mnt00001
pvccreate /dev/mapper/hana-FC5_log2_mnt00001
pvccreate /dev/mapper/hana-FC5_log_mnt00002
pvccreate /dev/mapper/hana-FC5_log2_mnt00002
```

2. Create the volume groups for each data and log partition.

```
vgcreate FC5_data_mnt00001 /dev/mapper/hana-FC5_data_mnt00001
/dev/mapper/hana-FC5_data2_mnt00001
vgcreate FC5_data_mnt00002 /dev/mapper/hana-FC5_data_mnt00002
/dev/mapper/hana-FC5_data2_mnt00002
vgcreate FC5_log_mnt00001 /dev/mapper/hana-FC5_log_mnt00001
/dev/mapper/hana-FC5_log2_mnt00001
vgcreate FC5_log_mnt00002 /dev/mapper/hana-FC5_log_mnt00002
/dev/mapper/hana-FC5_log2_mnt00002
```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00001
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00001
```

#### 4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```
modprobe dm_mod
pvscan
vgscan
lvscan
```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```
vgchange -a y
```

#### Create file systems

Create the XFS file system on all data and log logical volumes.

```
mkfs.xfs FC5_data_mnt00001-vol
mkfs.xfs FC5_data_mnt00002-vol
mkfs.xfs FC5_log_mnt00001-vol
mkfs.xfs FC5_log_mnt00002-vol
```

#### Create mount points

Create the required mount point directories, and set the permissions on all worker and standby hosts:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00002
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00002
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

## Mount file systems

To mount the `/hana/shared` file systems during system boot using the `/etc/fstab` configuration file, add the `/hana/shared` file system to the `/etc/fstab` configuration file of each host.

```
sapcc-hana-tst:/ # cat /etc/fstab
<storage-ip>:/hana_shared /hana/shared nfs rw,vers=3,hard,timeo=600,
intr,noatime,nolock 0 0
```



All the data and log file systems are mounted through the SAP HANA storage connector.

To mount the file systems, run the `mount -a` command at each host.

## I/O stack configuration for SAP HANA

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used.

NetApp conducted performance tests to define the ideal values. The following table lists the optimal values as inferred from the performance tests.

Parameter	Value
<code>max_parallel_io_requests</code>	128
<code>async_read_submit</code>	on
<code>async_write_submit_active</code>	on
<code>async_write_submit_blocks</code>	all

For SAP HANA 1.0 up to SPS12, these parameters can be set during the installation of the SAP HANA database as described in SAP Note [2267798 – Configuration of the SAP HANA Database during Installation Using hdbparam](#).

Alternatively, the parameters can be set after the SAP HANA database installation using the `hdbparam` framework.

```
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.max_parallel_io_requests=128
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_write_submit_active=on
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_read_submit=on
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_write_submit_blocks=all
```

Starting with SAP HANA 2.0, `hdbparam` is deprecated and the parameters have been moved to the `global.ini` file. The parameters can be set by using SQL commands or SAP HANA Studio. For more

information, see SAP Note [2399079 - Elimination of hdbparam in HANA 2](#). The parameters can be also set within the `global.ini` file.

```
SS3adm@stlrx300s8-6:/usr/sap/SS3/SYS/global/hdb/custom/config> cat
global.ini
...
[fileio]
async_read_submit = on
async_write_submit_active = on
max_parallel_io_requests = 128
async_write_submit_blocks = all
...
```

With SAP HANA 2.0 SPS5 and later, you can use the `setParameter.py` script to set the parameters mentioned above.

```
fc5adm@sapcc-hana-tst-03:/usr/sap/FC5/HDB00/exe/python_support>
python setParameter.py
-set=SYSTEM/global.ini/fileio/max_parallel_io_requests=128
python setParameter.py -set=SYSTEM/global.ini/fileio/async_read_submit=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_active=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_blocks=all
```

## SAP HANA software installation

Below are the requirements for SAP HANA software installation.

### Install on single-host system

SAP HANA software installation does not require any additional preparation for a single-host system.

### Install on multiple-host system



The following installation procedure is based on SAP HANA 1.0 SPS12 or later.

Before beginning the installation, create a `global.ini` file to enable use of the SAP storage connector during the installation process. The SAP storage connector mounts the required file systems at the worker hosts during the installation process. The `global.ini` file must be available in a file system that is accessible from all hosts, such as the `/hana/shared/SID` file system.

Before installing SAP HANA software on a multiple-host system, the following steps must be completed:

1. Add the following mount options for the data LUNs and the log LUNs to the `global.ini` file:
  - `relatime` and `inode64` for the data and log file system

2. Add the WWIDs of the data and log partitions. The WWIDs must match the alias names configured in the `/etc/multipath.conf` file.

The following output shows an example of a 2+1 multiple-host setup in which the system identifier (SID) is SS3.

```
stlrx300s8-6:~ # cat /hana/shared/global.ini
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/SS3
basepath_logvolumes = /hana/log/SS3
[storage]
ha_provider = hdb_ha.fcClient
partition_*_*_prtype = 5
partition_*_data__mountoptions = -o relatime,inode64
partition_*_log__mountoptions = -o relatime,inode64,nobarrier
partition_1_data__wwid = hana-SS3_data_mnt00001
partition_1_log__wwid = hana-SS3_log_mnt00001
partition_2_data__wwid = hana-SS3_data_mnt00002
partition_2_log__wwid = hana-SS3_log_mnt00002
[system_information]
usage = custom
[trace]
ha_fcclient = info
stlrx300s8-6:~ #
```

If LVM is used, the needed configuration is different. The example below shows a 2+1 multiple-host setup with SID=FC5.

```

sapcc-hana-tst-03:/hana/shared # cat global.ini
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/FC5
basepath_logvolumes = /hana/log/FC5
[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*_prtype = 5
partition_*_data__mountOptions = -o relatime,inode64
partition_*_log__mountOptions = -o relatime,inode64
partition_1_data__lvmname = FC5_data_mnt00001-vol
partition_1_log__lvmname = FC5_log_mnt00001-vol
partition_2_data__lvmname = FC5_data_mnt00002-vol
partition_2_log__lvmname = FC5_log_mnt00002-vol
sapcc-hana-tst-03:/hana/shared #

```

Using the SAP `hdblcm` installation tool, start the installation by running the following command at one of the worker hosts. Use the `addhosts` option to add the second worker (`sapcc-hana-tst-06`) and the standby host (`sapcc-hana-tst-07`).

The directory where the prepared the `global.ini` file has been stored is included with the `storage_cfg` CLI option (`--storage_cfg=/hana/shared`).

Depending on the OS version being used, it might be necessary to install python 2.7 before installing the SAP HANA database.

```

/hdblcm --action=install --addhosts=sapcc-hana-tst
-06:role=worker:storage_partition=2,sapcc-hana-tst-07:role=standby
--storage_cfg=/hana/shared/

```

```

AP HANA Lifecycle Management - SAP HANA Database 2.00.073.00.1695288802
*****

```

Scanning software locations...

Detected components:

```

    SAP HANA AFL (incl.PAL,BFL,OFL) (2.00.073.0000.1695321500) in
/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
    SAP HANA Database (2.00.073.00.1695288802) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
    SAP HANA Database Client (2.18.24.1695756995) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_CLIENT_LINUX_X86_64/SAP_HANA_CLIENT/client
    SAP HANA Studio (2.3.75.000000) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_STUDIO_LINUX_X86_64/studio

```

SAP HANA Local Secure Store (2.11.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/HANA\_LSS\_24\_LINUX\_X86\_64/packages  
SAP HANA XS Advanced Runtime (1.1.3.230717145654) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_RT\_10\_LINUX\_X86\_64/packages  
SAP HANA EML AFL (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/HDB\_EML\_AFL\_10\_LINUX\_X86\_64/packages  
SAP HANA EPM-MDS (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/SAP\_HANA\_EPM-MDS\_10/packages  
Automated Predictive Library (4.203.2321.0.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/PAAPL4\_H20\_LINUX\_X86\_64/apl-4.203.2321.0-hana2sp03-linux\_x64/installer/packages  
GUI for HALM for XSA (including product installer) Version 1 (1.015.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACALMPIUI15\_0.zip  
XSAC FILEPROCESSOR 1.0 (1.000.102) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACFILEPROC00\_102.zip  
SAP HANA tools for accessing catalog content, data preview, SQL console, etc. (2.015.230503) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSAC\_HRTT\_20/XSACHRTT15\_230503.zip  
Develop and run portal services for customer applications on XSA (2.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACPORTALSERV07\_0.zip  
The SAP Web IDE for HANA 2.0 (4.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSAC\_SAP\_WEB\_IDE\_20/XSACSAPWEBIDE07\_0.zip  
XS JOB SCHEDULER 1.0 (1.007.22) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACSERVICES07\_22.zip  
SAPUI5 FESV6 XSA 1 - SAPUI5 1.71 (1.071.52) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV671\_52.zip  
SAPUI5 FESV9 XSA 1 - SAPUI5 1.108 (1.108.5) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5FESV9108\_5.zip  
SAPUI5 SERVICE BROKER XSA 1 - SAPUI5 Service Broker 1.0 (1.000.4) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACUI5SB00\_4.zip  
XSA Cockpit 1 (1.001.37) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA\_UNITS/XSA\_CONTENT\_10/XSACXSACOCKPIT01\_37.zip

SAP HANA Database version '2.00.073.00.1695288802' will be installed.

Select additional components for installation:

Index	Components	Description
1	all	All components
2	server	No additional components
3	client	Install SAP HANA Database Client version 2.18.24.1695756995
4	lss	Install SAP HANA Local Secure Store version 2.11.0
5	studio	Install SAP HANA Studio version 2.3.75.000000
6	xs	Install SAP HANA XS Advanced Runtime version 1.1.3.230717145654
7	afl	Install SAP HANA AFL (incl.PAL,BFL,OFL) version 2.00.073.0000.1695321500
8	eml	Install SAP HANA EML AFL version 2.00.073.0000.1695321500
9	epmmds	Install SAP HANA EPM-MDS version 2.00.073.0000.1695321500
10	sap_afl_sdk_apl	Install Automated Predictive Library version 4.203.2321.0.0

Enter comma-separated list of the selected indices [3,4]: 2,3

Verify that the installation tool installed all selected components at all worker and standby hosts.

### Adding additional data volume partitions for SAP HANA single-host systems

Starting with SAP HANA 2.0 SPS4, additional data volume partitions can be configured. This feature allows you to configure two or more LUNs for the data volume of an SAP HANA tenant database and to scale beyond the size and performance limits of a single LUN.



It is not necessary to use multiple partitions to fulfil the SAP HANA KPIs. A single LUN with a single partition fulfils the required KPIs.



Using two or more individual LUNs for the data volume is only available for SAP HANA single-host systems. The SAP storage connector required for SAP HANA multiple-host systems does only support one device for the data volume.

You can add more data volume partitions at any time but it might require a restart of the SAP HANA database.

## Enabling additional data volume partitions

To enable additional data volume partitions, complete the following steps:

1. Add the following entry within the `global.ini` file:

```
[customizable_functionalities]
persistence_datavolume_partition_multipath = true
```

2. Restart the database to enable the feature. Adding the parameter through the SAP HANA Studio to the `global.ini` file by using the Systemdb configuration prevents the restart of the database.

## Volume and LUN configuration

The layout of volumes and LUNs is similar to the layout of a single host with one data volume partition, but with an additional data volume and LUN stored on a different aggregate as log volume and the other data volume. The following table shows an example configuration of an SAP HANA single-host systems with two data volume partitions.

Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller B
Data volume: SID_data_mnt00001	Shared volume: SID_shared	Data volume: SID_data2_mnt00001	Log volume: SID_log_mnt00001

The next table shows an example of the mount point configuration for a single-host system with two data volume partitions.

LUN	Mount point at HANA host	Note
SID_data_mnt00001	/hana/data/SID/mnt00001	Mounted using <code>/etc/fstab</code> entry
SID_data2_mnt00001	/hana/data2/SID/mnt00001	Mounted using <code>/etc/fstab</code> entry
SID_log_mnt00001	/hana/log/SID/mnt00001	Mounted using <code>/etc/fstab</code> entry
SID_shared	/hana/shared/SID	Mounted using <code>/etc/fstab</code> entry

Create the new data LUNs by using either ONTAP System Manager or the ONTAP CLI.

## Host configuration

To configure a host, complete the following steps:

1. Configure multipathing for the additional LUNs, as described in section 0.
2. Create the XFS file system on each additional LUN belonging to the HANA system.

```
stlrx300s8-6:/ # mkfs.xfs /dev/mapper/hana-FC5_data2_mnt00001
```

3. Add the additional file system/s to the `/etc/fstab` configuration file.



The XFS file systems for the data LUN must be mounted with the `relatime` and `inode64` mount options. The XFS file systems for the log LUN must be mounted with the `relatime`, `inode64`, and `nobarrier` mount options.

```
stlrx300s8-6:/ # cat /etc/fstab
/dev/mapper/hana-FC5_shared /hana/shared xfs defaults 0 0
/dev/mapper/hana-FC5_log_mnt00001 /hana/log/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/hana-FC5_data_mnt00001 /hana/data/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/hana-FC5_data2_mnt00001 /hana/data2/FC5/mnt00001 xfs
relatime,inode64 0 0
```

4. Create the mount points and set the permissions on the database host.

```
stlrx300s8-6:/ # mkdir -p /hana/data2/FC5/mnt00001
stlrx300s8-6:/ # chmod -R 777 /hana/data2/FC5
```

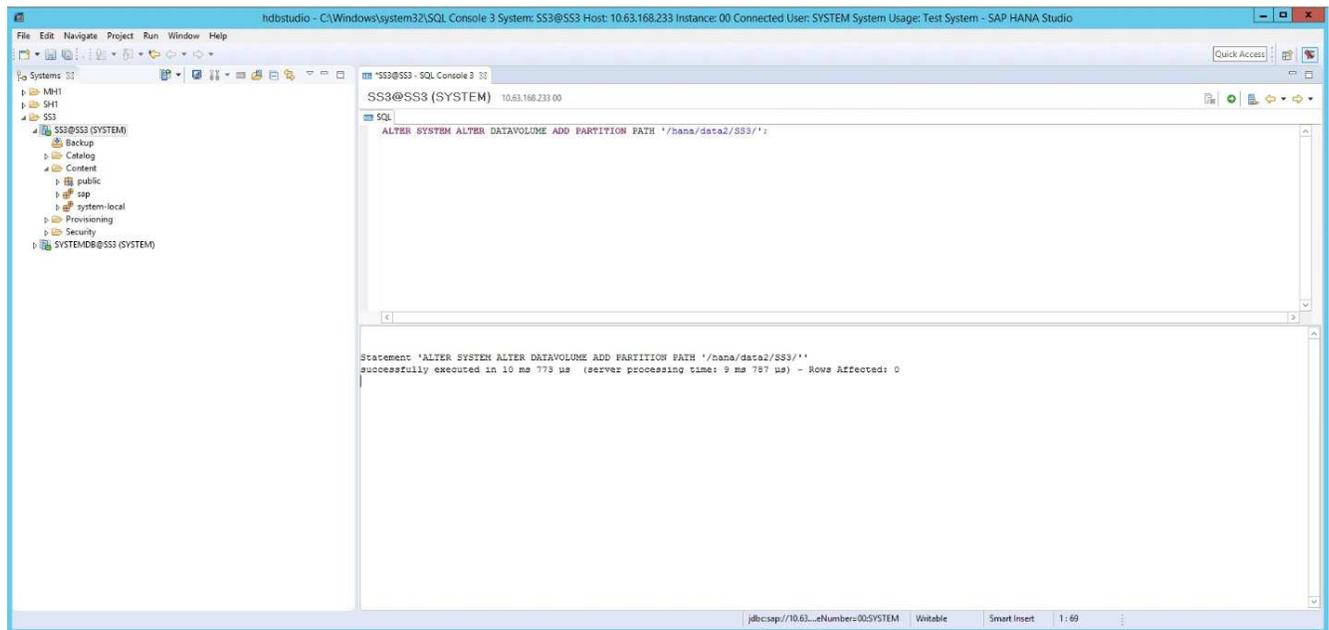
5. To mount the file systems, run the `mount -a` command.

#### Adding an additional datavolume partition

To add an additional datavolume partition to your tenant database, complete the following step:

1. Execute the following SQL statement against the tenant database. Each additional LUN can have a different path.

```
ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION PATH '/hana/data2/SID/';
```



## Where to find additional information

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

- [SAP HANA Software Solutions](#)
- [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP System Copies Using the SnapCenter SAP HANA Plug-In](#)
- [NetApp Documentation Centers](#)

<https://www.netapp.com/support-and-training/documentation/>

- [SAP Certified Enterprise Storage Hardware for SAP HANA](#)

<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/>

- [SAP HANA Storage Requirements](#)

<https://www.sap.com/documents/2024/03/146274d3-ae7e-0010-bca6-c68f7e60039b.html>

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

<https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html>

- [SAP HANA on VMware vSphere Wiki](#)

[https://help.sap.com/docs/SUPPORT\\_CONTENT/virtualization/3362185751.html](https://help.sap.com/docs/SUPPORT_CONTENT/virtualization/3362185751.html)

- [SAP HANA on VMware vSphere Best Practices Guide](#)

[https://www.vmware.com/docs/sap\\_hana\\_on\\_vmware\\_vsphere\\_best\\_practices\\_guide-white-paper](https://www.vmware.com/docs/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper)

## Update history

The following technical changes have been made to this solution since its original publication.

Date	Update summary
February 2015	Initial version
October 2015	Included I/O parameters for SAP HANA and HWVAL SPS 10 and later
February 2016	Updated capacity sizing
February 2017	New NetApp storage systems and disk shelves New features of ONTAP 9 New OS releases (SLES12 SP1 and Red Hat Enterprise Linux 7.2) New SAP HANA release
July 2017	Minor updates
September 2018	New NetApp storage systems New OS releases (SLES12 SP3 and Red Hat Enterprise Linux 7.4) Additional minor updates SAP HANA 2.0 SPS3
September 2019	New OS releases Minor updates
April 2020	Introduced multiple data partition features available since SAP HANA 2.0 SPS4
June 2020	Additional information about optional functionalities Minor updates
February 2021	Linux LVM support New NetApp storage systems New OS releases (SLES15SP2, RHEL 8)
April 2021	VMware vSphere-specific information added
September 2022	New OS-Releases
September 2024	New Storage Systems
February 2025	New Storage System
July 2025	Minor updates

# Backup, Restore and Disaster Recovery

## SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter

### TR-4926: SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter

This technical report provides best practices for SAP HANA data protection on Amazon FSx for NetApp ONTAP and NetApp SnapCenter. This document covers SnapCenter concepts, configuration recommendations, and operation workflows, including configuration, backup operations, and restore and recovery operations.

Author: Nils Bauer, NetApp

Companies today require continuous, uninterrupted availability for their SAP applications. They expect consistent performance levels in the face of ever-increasing volumes of data and the need for routine maintenance tasks, such as system backups. Performing backups of SAP databases is a critical task and can have a significant performance impact on the production SAP system.

Backup windows are shrinking while the amount of data to be backed up is increasing. Therefore, it is difficult to find a time when you can perform backups with minimal effect on business processes. The time needed to restore and recover SAP systems is a concern because downtime for SAP production and nonproduction systems must be minimized to reduce cost to the business.

### Backup and recovery using Amazon FSx for ONTAP

You can use NetApp Snapshot technology to create database backups in minutes.

The time needed to create a Snapshot copy is independent of the size of the database because a Snapshot copy does not move any physical data blocks on the storage platform. In addition, the use of Snapshot technology has no performance effect on the live SAP system. Therefore, you can schedule the creation of Snapshot copies without considering peak dialog or batch activity periods. SAP and NetApp customers typically schedule multiple online Snapshot backups during the day; for example, every six hours is common. These Snapshot backups are typically kept for three to five days on the primary storage system before being removed or tiered to cheaper storage for long term retention.

Snapshot copies also provide key advantages for restore and recovery operations. NetApp SnapRestore technology enables the restoration of an entire database or, alternatively, just a portion of a database to any point in time, based on the currently available Snapshot copies. Such restore processes are finished in a few seconds, independent of the size of the database. Because several online Snapshot backups can be created during the day, the time needed for the recovery process is significantly reduced relative to a traditional once per day backup approach. Because you can perform a restore with a Snapshot copy that is at most only a few hours old (rather than up to 24 hours), fewer transaction logs must be applied during forward recovery. Therefore, the RTO is reduced to several minutes rather than the several hours required for conventional streaming backups.

Snapshot copy backups are stored on the same disk system as the active online data. Therefore, NetApp recommends using Snapshot copy backups as a supplement rather than a replacement for backups to a secondary location. Most restore and recovery actions are managed by using SnapRestore on the primary storage system. Restores from a secondary location are only necessary if the primary storage system

containing the Snapshot copies is damaged. You can also use the secondary location if it is necessary to restore a backup that is no longer available on the primary location.

A backup to a secondary location is based on Snapshot copies created on the primary storage. Therefore, the data is read directly from the primary storage system without generating load on the SAP database server. The primary storage communicates directly with the secondary storage and replicates the backup data to the destination by using the NetApp SnapVault feature.

SnapVault offers significant advantages when compared to traditional backups. After an initial data transfer, in which all data has been transferred from the source to the destination, all subsequent backups copy only move the changed blocks to the secondary storage. Therefore, the load on the primary storage system and the time needed for a full backup are significantly reduced. Because SnapVault stores only the changed blocks at the destination, any additional full database backups consume significantly less disk space.

### Runtime of Snapshot backup and restore operations

The following figure shows a customer's HANA Studio using Snapshot backup operations. The image shows that the HANA database (approximately 4TB in size) is backed up in 1 minute and 20 seconds by using Snapshot backup technology and more than 4 hours with a file-based backup operation.

The largest part of the overall backup workflow runtime is the time needed to execute the HANA backup save point operation, and this step is dependent on the load on the HANA database. The storage Snapshot backup itself always finishes in a couple of seconds.

Stat...	Started	Duration	Size	Backup Ty...	Destinati...
●	Jan 11, 2022 10:26:59 AM	00h 01m 17s	4.51 TB	Data Back...	Snapshot
●	Jan 11, 2022 8:40:02 AM	00h 27m 11s	4.51 TB	Data Back...	Snapshot
●	Jan 11, 2022 1:00:58 AM	04h 05m 39s	3.82 TB	Data Back...	File
●	Jan 9, 2022 4:40:03 PM	00h 01m 23s	4.51 TB	Data Back...	Snapshot
●	Jan 9, 2022 8:00:02 AM	02h 39m 04s	3.82 TB	Data Back...	File
●	Jan 9, 2022 12:40:03 AM	00h 01m 18s	4.51 TB	Data Back...	Snapshot
●	Jan 8, 2022 4:40:03 PM	00h 01m 18s	4.51 TB	Data Back...	Snapshot
●	Jan 8, 2022 8:40:03 AM	00h 01m 22s	4.51 TB	Data Back...	Snapshot
●	Jan 8, 2022 12:40:03 AM	00h 01m 19s	4.51 TB	Data Back...	Snapshot
●	Jan 7, 2022 4:40:03 PM	00h 01m 19s	4.51 TB	Data Back...	Snapshot
●	Jan 7, 2022 8:40:02 AM	00h 01m 19s	4.51 TB	Data Back...	Snapshot
●	Jan 7, 2022 12:40:02 AM	00h 01m 20s	4.51 TB	Data Back...	Snapshot
●	Jan 6, 2022 4:40:02 PM	00h 01m 18s	4.51 TB	Data Back...	Snapshot
●	Jan 6, 2022 8:40:03 AM	00h 01m 17s	4.51 TB	Data Back...	Snapshot
●	Jan 6, 2022 12:40:03 AM	00h 01m 19s	4.51 TB	Data Back...	Snapshot
●	Jan 5, 2022 4:40:03 PM	00h 01m 19s	4.51 TB	Data Back...	Snapshot

File-based backup: **4 hours 05 min**

(~270 MB/s throughput)

04h 05m 39s	3.82 TB	Data Back...	File
-------------	---------	--------------	------

Snapshot backup: **1 min 20 sec**

00h 01m 18s	4.51 TB	Data Back...	Snapshot
00h 01m 22s	4.51 TB	Data Back...	Snapshot
00h 01m 19s	4.51 TB	Data Back...	Snapshot

**Backup runtime reduced by 99%**

### Recovery time objective comparison

This section provides a recovery time objective (RTO) comparison of file-based and storage-based Snapshot backups. The RTO is defined by the sum of the time needed to restore, recover, and then start the database.

#### Time needed to restore database

With a file-based backup, the restore time depends on the size of the database and backup infrastructure, which defines the restore speed in megabytes per second. For example, if the infrastructure supports a restore operation at a speed of 250MBps, it takes approximately 4.5 hours to restore a database 4TB in size on the persistence.

With storage Snapshot copy backups, the restore time is independent of the size of the database and is always

in the range of a couple of seconds.

### Time needed to start database

The database start time depends on the size of the database and the time needed to load the data into memory. In the following examples, it is assumed that the data can be loaded with 1000MBps. Loading 4TB into memory takes around 1hour and 10 minutes. The start time is the same for a file-based and Snapshot based restore and recovery operations.

### Time needed to recover database

The recovery time depends on the number of logs that must be applied after the restore. This number is determined by the frequency at which data backups are taken.

With file-based data backups, the backup schedule is typically once per day. A higher backup frequency is normally not possible, because the backup degrades production performance. Therefore, in the worst case, all the logs that were written during the day must be applied during forward recovery.

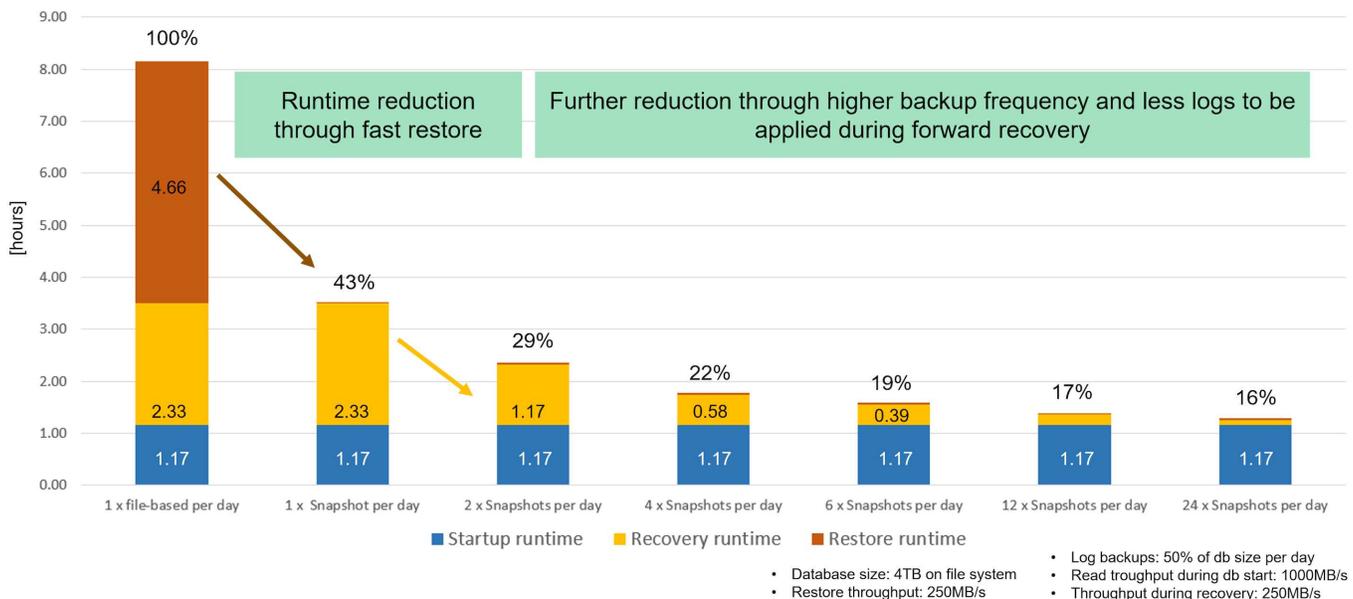
Snapshot backups are typically scheduled with a higher frequency because they do not influence the performance of the SAP HANA database. For example, if Snapshot backups are scheduled every six hours, the recovery time would be, in the worst case, one-fourth of the recovery time for a file-based backup (6 hours / 24 hours = .25).

The following figure shows a comparison of restore and recovery operations with a daily file-based backup and Snapshot backups with different schedules.

The first two bars show that even with a single Snapshot backup per day, the restore and recovery is reduced to 43% due to the speed of the restore operation from a Snapshot backup. If multiple Snapshot backups per day are created, the runtime can be reduced further because less logs need to be applied during forward recovery.

The following figure also shows that four to six Snapshot backups per day makes the most sense, because a higher frequency does not have a big influence on the overall runtime anymore.

## Restore and Recovery of a 4TB HANA Database (8TB RAM)



## Use cases and values of accelerated backup and cloning operations

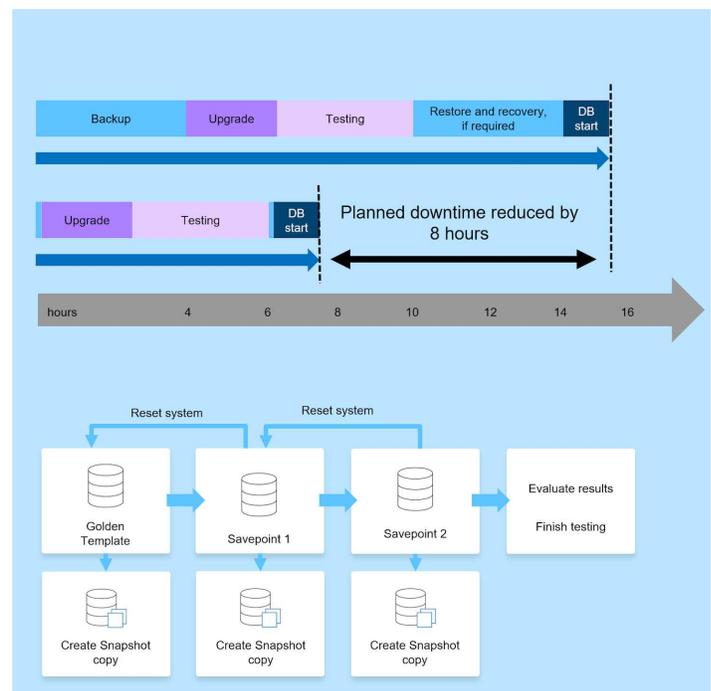
Executing backups is a critical part of any data protection strategy. Backups are scheduled on a regular basis to ensure that you can recover from system failures. This is the most obvious use case, but there are also other SAP lifecycle management tasks, where accelerating backup and recovery operations is crucial.

SAP HANA system upgrade is an example of where an on-demand backup before the upgrade and a possible restore operation if the upgrade fails has a significant impact on the overall planned downtime. With the example of a 4TB database, you can reduce the planned downtime by 8 hours by using the Snapshot-based backup and restore operations.

Another use case example would be a typical test cycle, where testing must be done over multiple iterations with different data sets or parameters. When leveraging the fast backup and restore operations, you can easily create save points within your test cycle and reset the system to any of these previous save points if a test fails or needs to be repeated. This enables testing to finish earlier or enables more testing at the same time and improves test results.

## Use Cases for Backup and Recovery Operations

- Accelerate HANA system upgrade operations
  - Fast on-demand backup before HANA system upgrade
  - Fast restore operation in case of an upgrade failure
  - Reduction of planned downtime
- Accelerate test cycles
  - Fast creation of savepoints after a successful step
  - Fast reset of system to any savepoint
  - Repeat step until successful



When Snapshot backups have been implemented, they can be used to address multiple other use cases, which require copies of a HANA database. With FSx for ONTAP, you can create a new volume based on the content of any available Snapshot backup. The runtime of this operation is a few seconds, independent of the size of the volume.

The most popular use case is the SAP System Refresh, where data from the production system needs to be copied to the test or QA system. By leveraging the FSx for ONTAP cloning feature, you can provision the volume for the test system from any Snapshot copy of the production system in a matter of seconds. The new volume then must be attached to the test system and the HANA database recovered.

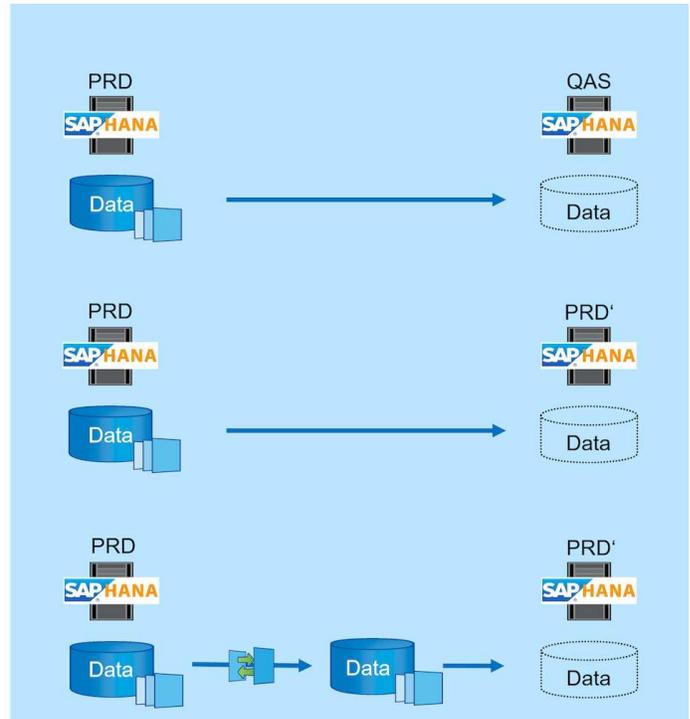
The second use case is the creation of a repair system, which is used to address a logical corruption in the production system. In this case, an older Snapshot backup of the production system is used to start a repair system, which is an identical clone of the production system with the data before the corruption occurred. The repair system is then used to analyze the problem and export the required data before it was corrupted.

The last use case is the ability to run a disaster recover failover test without stopping the replication and therefore without influencing RTO and recovery point objective (RPO) of the disaster recovery setup. When

FSx for ONTAP NetApp SnapMirror replication is used to replicate the data to the disaster recovery site, the production Snapshot backups are available at the disaster recovery site as well and can then be used to create a new volume for disaster recover testing.

## Use Cases for Cloning Operations

- SAP System Refresh
  - Fast creation of a new volume based on a production Snapshot backup
  - Attach volume to the test system and recover HANA database with SID change
  
- Repair System creation to address logical corruption
  - Fast creation of a new volume based on a production Snapshot backup
  - Attach volume to the repair system and recover HANA database w/o SID change
  
- Disaster Recovery testing
  - Combined with SnapMirror Replication
  - Attach storage clone from a replicated production Snapshot backup to a DR test system



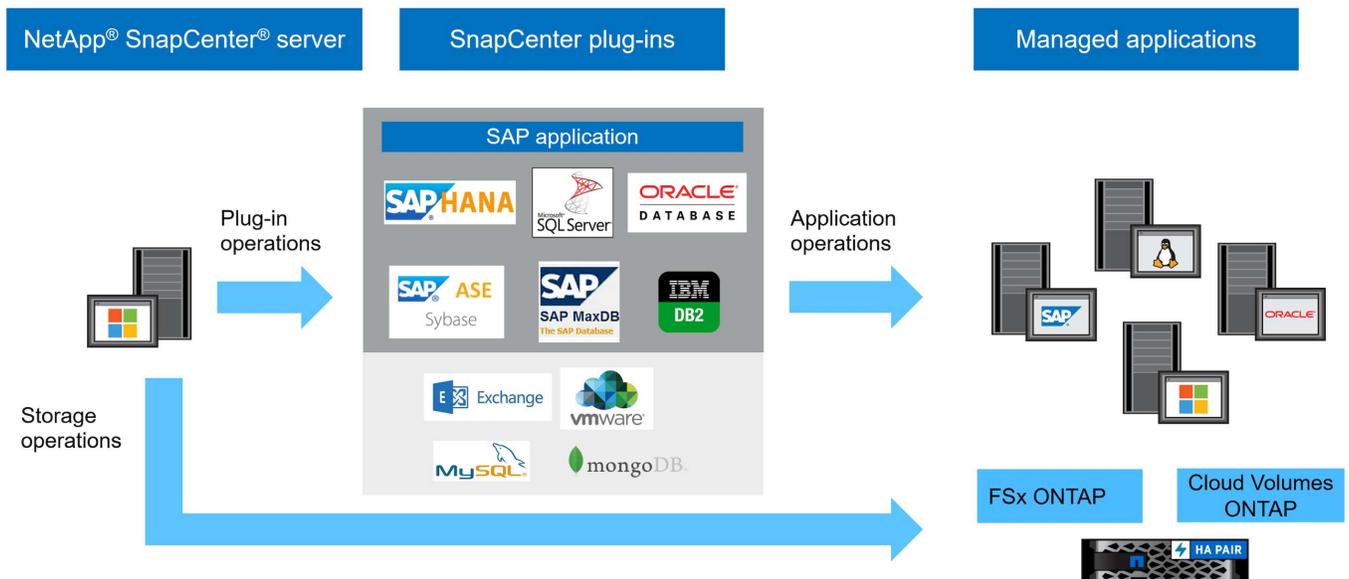
## SnapCenter architecture

SnapCenter is a unified, scalable platform for application-consistent data protection. SnapCenter provides centralized control and oversight, while delegating the ability for users to manage application-specific backup, restore, and clone jobs. With SnapCenter, database and storage administrators learn a single tool to manage backup, restore, and cloning operations for a variety of applications and databases.

SnapCenter manages data across endpoints in the data fabric powered by NetApp. You can use SnapCenter to replicate data between on-premises environments; between on-premises environments and the cloud; and between private, hybrid, or public clouds.

## SnapCenter components

SnapCenter includes the SnapCenter Server, the SnapCenter Plug-In Package for Windows, and the SnapCenter Plug-In Package for Linux. Each package contains plug-ins to SnapCenter for various applications and infrastructure components.



## SnapCenter SAP HANA backup solution

The SnapCenter backup solution for SAP HANA covers the following areas:

- Backup operations, scheduling, and retention management
  - SAP HANA data backup with storage-based Snapshot copies
  - Non-data volume backup with storage-based Snapshot copies (for example, /hana/shared)
  - Database block integrity checks using a file-based backup
  - Replication to an off-site backup or disaster recovery location
- Housekeeping of the SAP HANA backup catalog
  - For HANA data backups (Snapshot and file-based)
  - For HANA log backups
- Restore and recovery operations
  - Automated restore and recovery
  - Single tenant restore operations for SAP HANA (MDC) systems

Database data file backups are executed by SnapCenter in combination with the plug-in for SAP HANA. The plug-in triggers the SAP HANA database backup save point so that the Snapshot copies, which are created on the primary storage system, are based on a consistent image of the SAP HANA database.

SnapCenter enables the replication of consistent database images to an off-site backup or disaster recovery location by using SnapVault or the SnapMirror feature. Typically, different retention policies are defined for backups at primary and at the off-site backup storage. SnapCenter handles the retention at primary storage, and ONTAP handles the retention at the off-site backup storage.

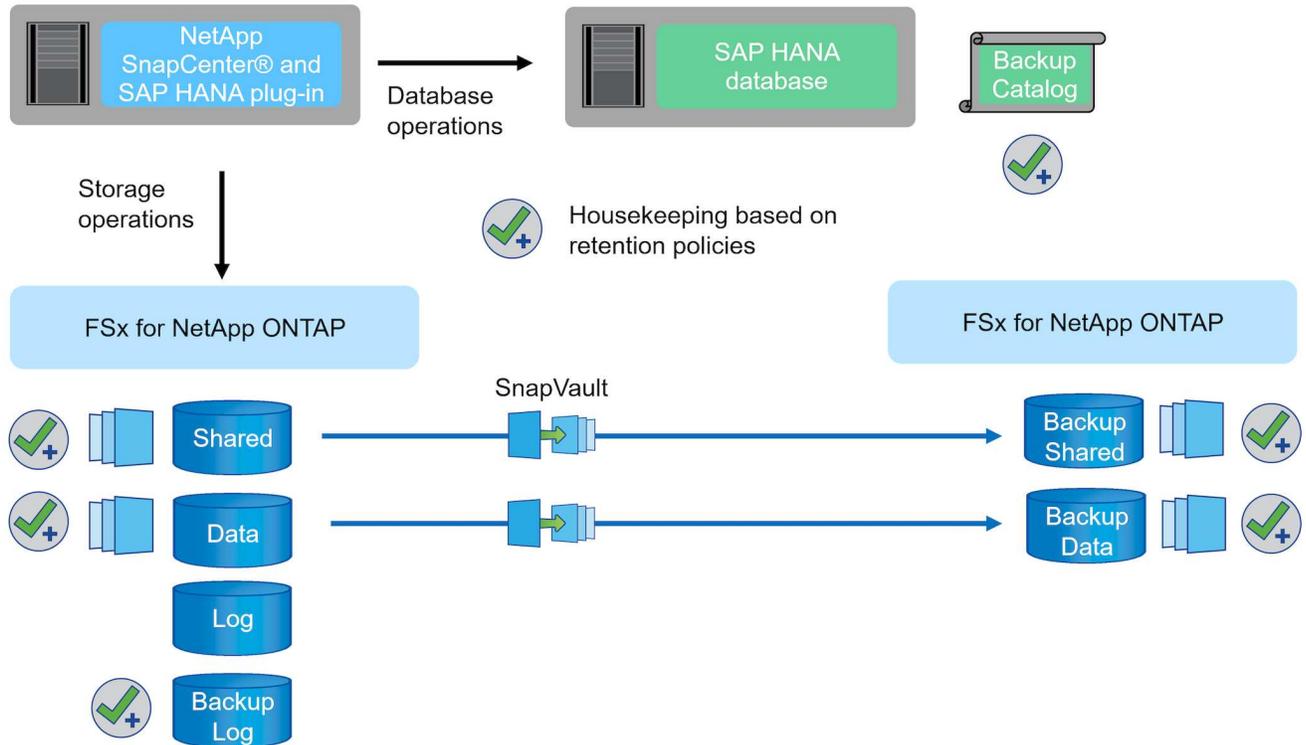
To allow a complete backup of all SAP HANA-related resources, SnapCenter also enables you to back up all non-data volumes by using the SAP HANA plug-in with storage-based Snapshot copies. You can schedule non-data volumes independently from the database data backup to enable individual retention and protection policies.

SAP recommends combining storage-based Snapshot backups with a weekly file-based backup to execute a

block integrity check. You can execute the block integrity check from within SnapCenter. Based on your configured retention policies, SnapCenter manages the housekeeping of data file backups at the primary storage, log file backups, and the SAP HANA backup catalog.

SnapCenter handles the retention at primary storage, while FSx for ONTAP manages secondary backup retention.

The following figure shows an overview of the SnapCenter backup and retention management operations.



When executing a storage-based Snapshot backup of the SAP HANA database, SnapCenter performs the following tasks:

1. Creates an SAP HANA backup save point to create a consistent image on the persistence layer.
2. Creates a storage-based Snapshot copy of the data volume.
3. Registers the storage-based Snapshot back up in the SAP HANA backup catalog.
4. Releases the SAP HANA backup save point.
5. Executes a SnapVault or SnapMirror update for the data volume, if configured.
6. Deletes storage Snapshot copies at the primary storage based on the defined retention policies.
7. Deletes SAP HANA backup catalog entries if the backups do not exist anymore at the primary or off-site backup storage.
8. Whenever a backup has been deleted based on the retention policy or manually, SnapCenter also deletes all log backups that are older than the oldest data backup. Log backups are deleted on the file system and in the SAP HANA backup catalog.

### Scope of this document

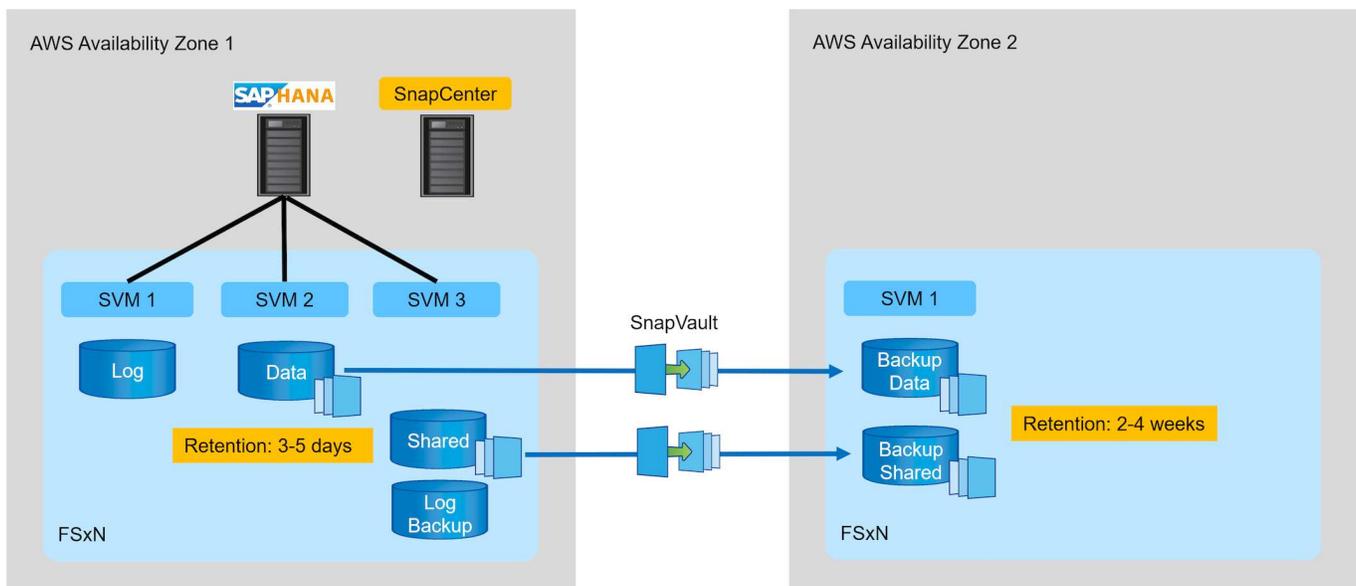
This document describes the most common SnapCenter configuration option for an SAP HANA MDC single host system with a single tenant on FSx for ONTAP. Other configuration options are possible and, in some

cases, required for specific SAP HANA systems, for example, for a multiple host system. For a detailed description about other configuration options, see [SnapCenter concepts and best practices \(netapp.com\)](#).

In this document, we use the Amazon Web Services (AWS) console and the FSx for ONTAP CLI to execute the required configuration steps on the storage layer. You can also use NetApp Cloud Manager to manage FSx for ONTAP, but this is out of scope for this document. For information about using NetApp Cloud Manager for FSx for ONTAP, see [Learn about Amazon FSx for ONTAP \(netapp.com\)](#).

## Data protection strategy

The following figure shows a typical backup architecture for SAP HANA on FSx for ONTAP. The HANA system is located in the AWS availability zone 1 and is using an FSx for ONTAP file system within the same availability zone. Snapshot backup operations are executed for the data and the shared volume of the HANA database. In addition to the local Snapshot backups, which are kept for 3-5 days, backups are also replicated to an offsite storage for longer term retention. The offsite backup storage is a second FSx for ONTAP file system located in a different AWS availability zone. Backups of the HANA data and shared volume are replicated with SnapVault to the second FSx for ONTAP file system and are kept for 2-3 weeks.



Before configuring SnapCenter, the data protection strategy must be defined based on the RTO and RPO requirements of the various SAP systems.

A common approach is to define system types such as production, development, test, or sandbox systems. All SAP systems of the same system type typically have the same data protection parameters.

The following parameters must be defined:

- How often should a Snapshot backup be executed?
- How long should Snapshot copy backups be kept on the primary storage system?
- How often should a block integrity check be executed?
- Should the primary backups be replicated to an off-site backup site?
- How long should the backups be kept at the off-site backup storage?

The following table shows an example of data protection parameters for the system types: production, development, and test. For the production system, a high backup frequency has been defined, and the backups are replicated to an off-site backup site once per day. The test systems have lower requirements and

no replication of the backups.

Parameters	Production systems	Development systems	Test systems
Backup frequency	Every 6 hours	Every 6 hours	Every 6 hours
Primary retention	3 days	3 days	3 days
Block integrity check	Once per week	Once per week	No
Replication to off-site backup site	Once per day	Once per day	No
Off-site backup retention	2 weeks	2 weeks	Not applicable

The following table shows the policies that must be configured for the data protection parameters.

Parameters	Policy LocalSnap	Policy LocalSnapAndSnapVault	Policy BlockIntegrityCheck
Backup type	Snapshot based	Snapshot based	File based
Schedule frequency	Hourly	Daily	Weekly
Primary retention	Count = 12	Count = 3	Count = 1
SnapVault replication	No	Yes	Not applicable

The policy `LocalSnapshot` is used for the production, development, and test systems to cover the local Snapshot backups with a retention of two days.

In the resource protection configuration, the schedule is defined differently for the system types:

- Production: Schedule every 4 hours.
- Development: Schedule every 4 hours.
- Test: Schedule every 4 hours.

The policy `LocalSnapAndSnapVault` is used for the production and development systems to cover the daily replication to the off-site backup storage.

In the resource protection configuration, the schedule is defined for production and development:

- Production: Schedule every day.
- Development: Schedule every day. The policy `BlockIntegrityCheck` is used for the production and development systems to cover the weekly block integrity check by using a file-based backup.

In the resource protection configuration, the schedule is defined for production and development:

- Production: Schedule every week.
- Development: Schedule every week.

For each individual SAP HANA database that uses the off-site backup policy, you must configure a protection relationship on the storage layer. The protection relationship defines which volumes are replicated and the retention of backups at the off-site backup storage.

With the following example, for each production and development system, a retention of two weeks is defined at the off-site backup storage.

In this example, protection policies and retention for SAP HANA database resources and non- data volume resources are not different.

### Example lab setup

The following lab setup was used as an example configuration for the rest of this document.

HANA system PFX:

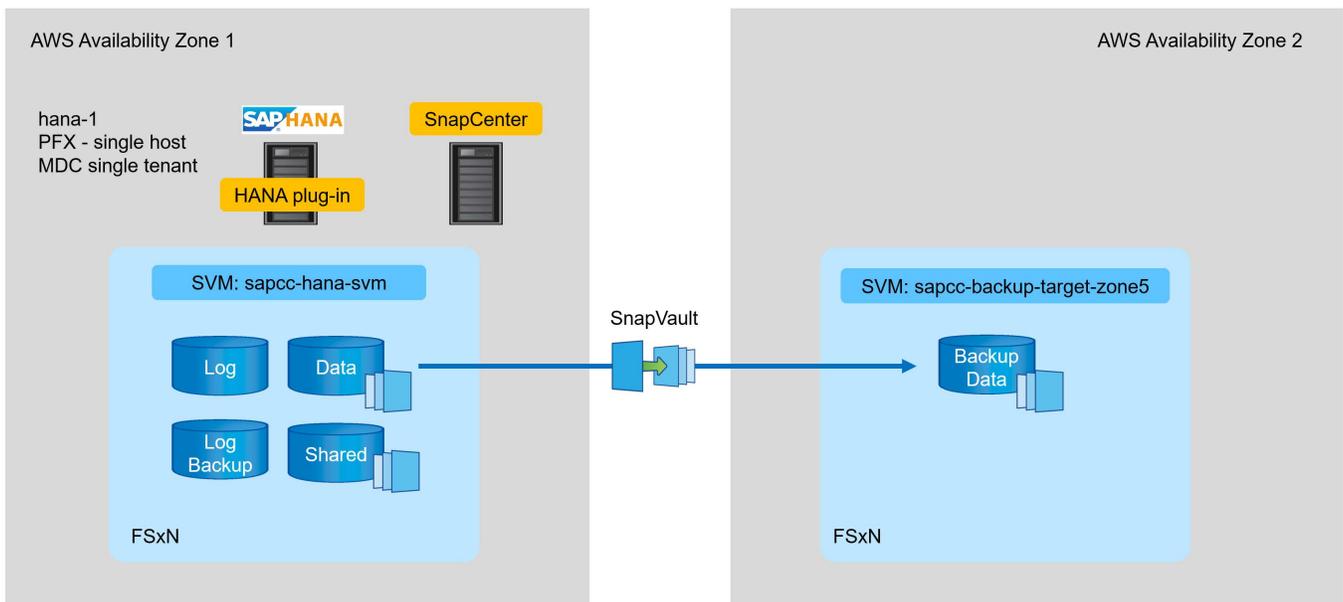
- Single host MDC system with a single tenant
- HANA 2.0 SPS 6 revision 60
- SLES for SAP 15SP3

SnapCenter:

- Version 4.6
- HANA and Linux plug-in deployed on a HANA database host

FSx for ONTAP file systems:

- Two FSx for ONTAP file systems with a single storage virtual machine (SVM)
- Each FSx for ONTAP system in a different AWS availability zone
- HANA data volume replicated to the second FSx for ONTAP file system



### SnapCenter configuration

You must perform the steps in this section for base SnapCenter configuration and the protection of the HANA resource.

## Overview configuration steps

You must perform the following steps for base SnapCenter configuration and the protection of the HANA resource. Each step is described in detail in the following chapters.

1. Configure SAP HANA backup user and hdbuserstore key. Used to access the HANA database with the hdbsql client.
2. Configure storage in SnapCenter. Credentials to access the FSx for ONTAP SVMs from SnapCenter
3. Configure credentials for plug-in deployment. Used to automatically deploy and install the required SnapCenter plug-ins on the HANA database host.
4. Add HANA host to SnapCenter. Deploys and installs the required SnapCenter plug-ins.
5. Configure policies. Defines the backup operation type (Snapshot, file), retentions, as well as optional Snapshot backup replication.
6. Configure HANA resource protection. Provide hdbuserstore key and attach policies and schedules to the HANA resource.

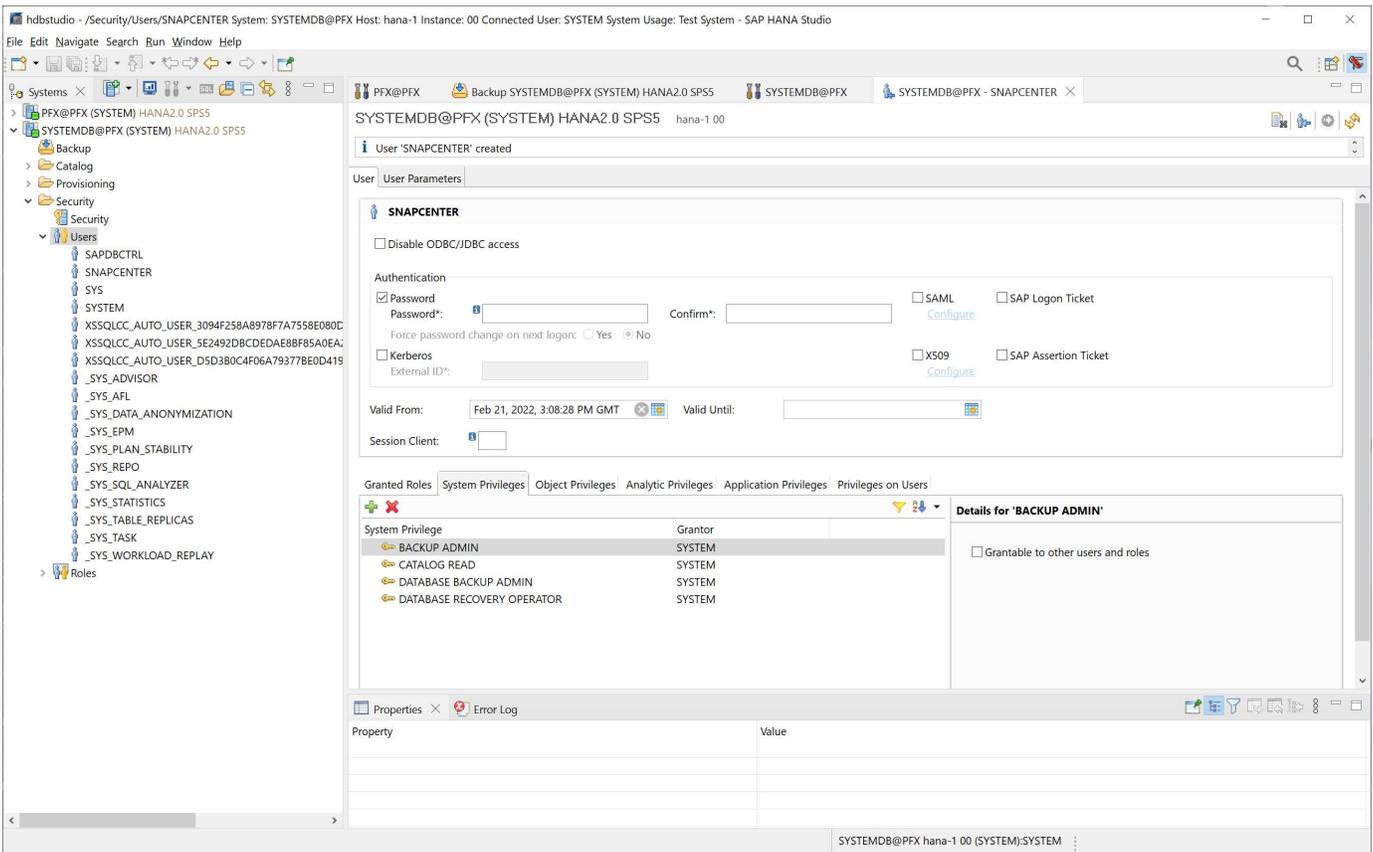
## SAP HANA backup user and hdbuserstore configuration

NetApp recommends configuring a dedicated database user in the HANA database to run the backup operations with SnapCenter. In the second step, an SAP HANA user store key is configured for this backup user, and this user store key is used in the configuration of the SnapCenter SAP HANA plug-in.

The following figure shows the SAP HANA Studio through which you can create the backup user

The required privileges are changed with the HANA 2.0 SPS5 release: backup admin, catalog read, database backup admin, and database recovery operator. For earlier releases, backup admin and catalog read are sufficient.

For an SAP HANA MDC system, you must create the user in the system database because all backup commands for the system and the tenant databases are executed by using the system database.



The following command is used for the user store configuration with the <sid>adm user:

```
hdbuserstore set <key> <host>:<port> <database user> <password>
```

SnapCenter uses the <sid>adm user to communicate with the HANA database. Therefore, you must configure the user store key by using the <'sid>adm` user on the database host. Typically, the SAP HANA hdbsql client software is installed together with the database server installation. If this is not the case, you must install the hdbclient first.

In an SAP HANA MDC setup, port 3<instanceNo>13 is the standard port for SQL access to the system database and must be used in the hdbuserstore configuration.

For an SAP HANA multiple-host setup, you must configure user store keys for all hosts. SnapCenter tries to connect to the database by using each of the provided keys and can therefore operate independently of a failover of an SAP HANA service to a different host. In our lab setup, we configured a user store key for the user pfxadm for our system PFX, which is a single host HANA MDC system with a single tenant.

```
pfxadm@hana-1:/usr/sap/PFX/home> hdbuserstore set PFXKEY hana-1:30013
SNAPCENTER <password>
Operation succeed.
```

```

pfxadm@hana-1:/usr/sap/PFX/home> hdbuserstore list
DATA FILE      : /usr/sap/PFX/home/.hdb/hana-1/SSFS_HDB.DAT
KEY FILE       : /usr/sap/PFX/home/.hdb/hana-1/SSFS_HDB.KEY
ACTIVE RECORDS : 7
DELETED RECORDS : 0
KEY PFXKEY
  ENV : hana-1:30013
  USER: SNAPCENTER
KEY PFXSAPDBCTRL
  ENV : hana-1:30013
  USER: SAPDBCTRL
Operation succeed.

```

You can check the access to the HANA system database that uses the key with the `hdbsql` command.

```

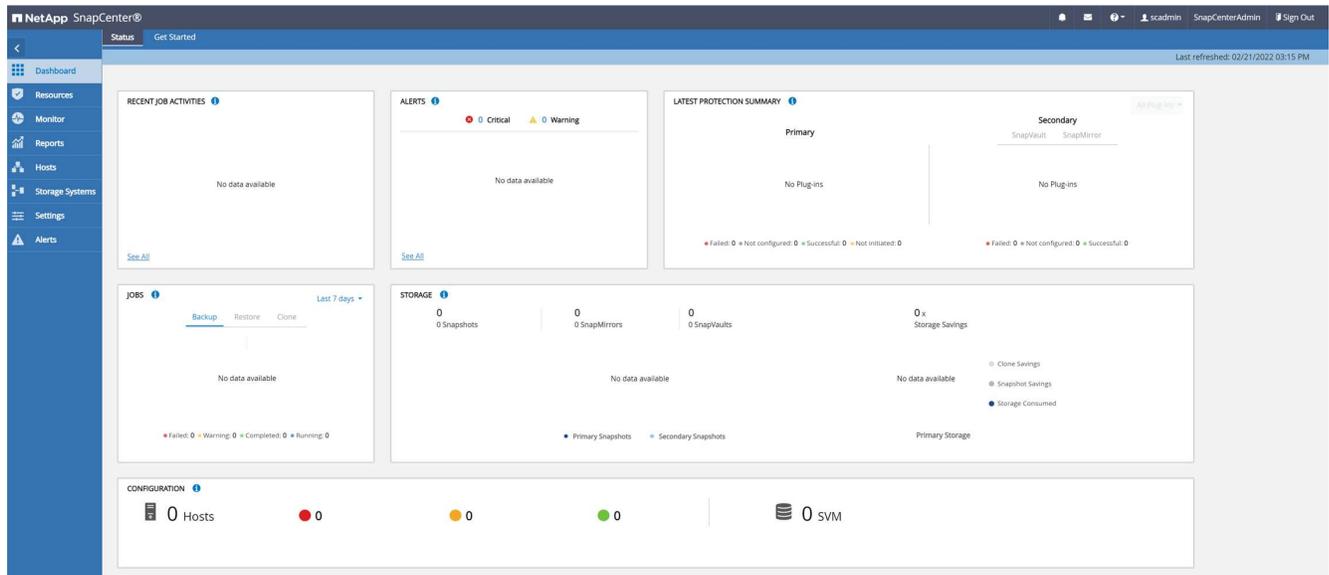
pfxadm@hana-1:/usr/sap/PFX/home> hdbsql -U PFXKEY
Welcome to the SAP HANA Database interactive terminal.
Type: \h for help with commands
      \q to quit
hdbsql SYSTEMDB=>

```

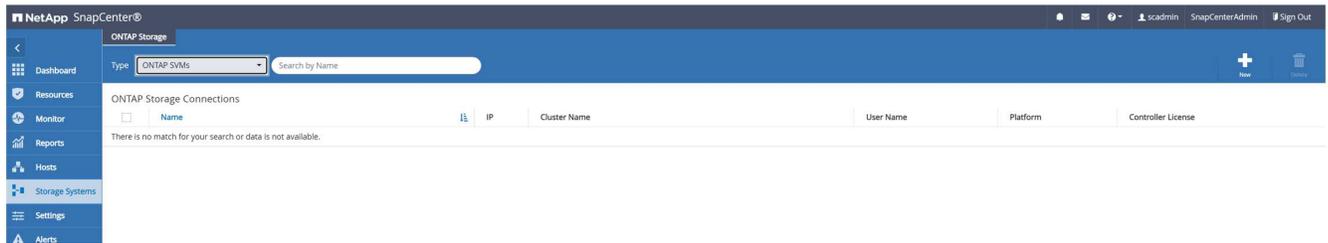
## Configure storage

Follow these steps to configure storage in SnapCenter.

1. In the SnapCenter UI, select Storage Systems.

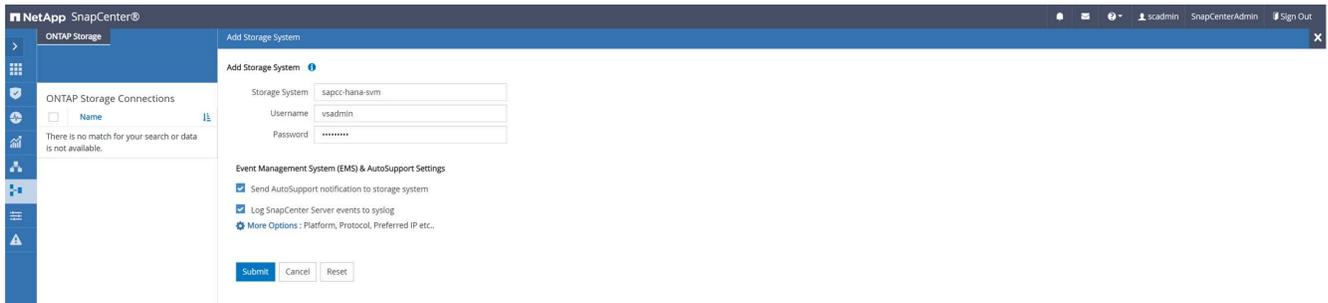


You can select the storage system type, which can be ONTAP SVMs or ONTAP Clusters. In the following example, SVM management is selected.

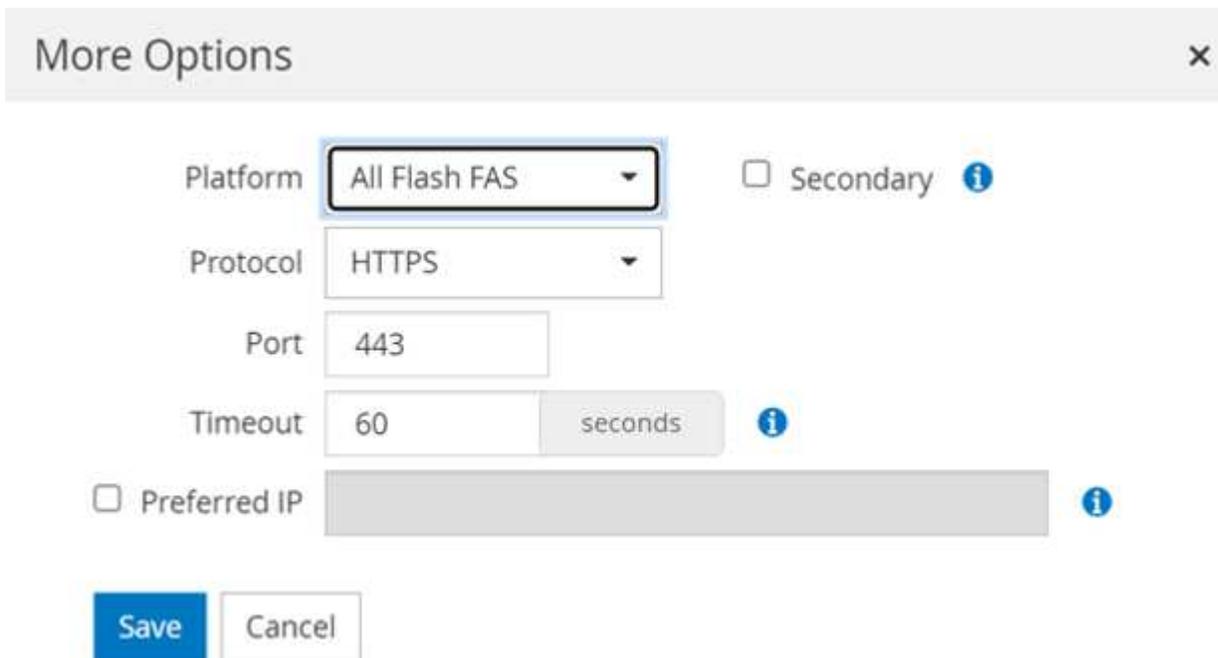


- To add a storage system and provide the required host name and credentials, click New.

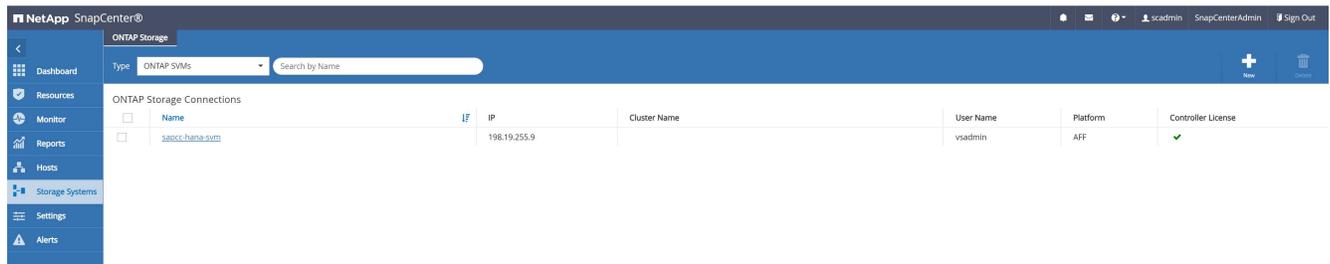
The SVM user is not required to be the vsadmin user, as shown in the following figure. Typically, a user is configured on the SVM and assigned the required permissions to execute backup and restore operations. For information about required privileges, see [SnapCenter Installation Guide](#) in the section titled “Minimum ONTAP privileges required”.



- To configure the storage platform, click More Options.
- Select All Flash FAS as the storage system to ensure that the license, which is part of FSx for ONTAP, is available for SnapCenter.



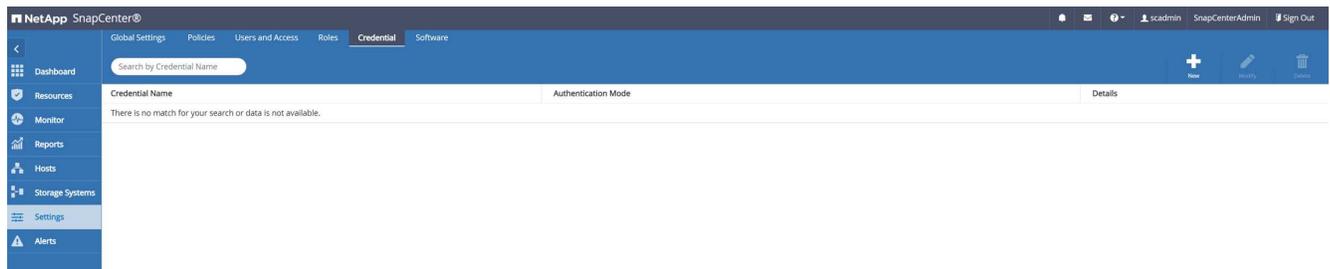
The SVM `sapcc-hana-svm` is now configured in SnapCenter.



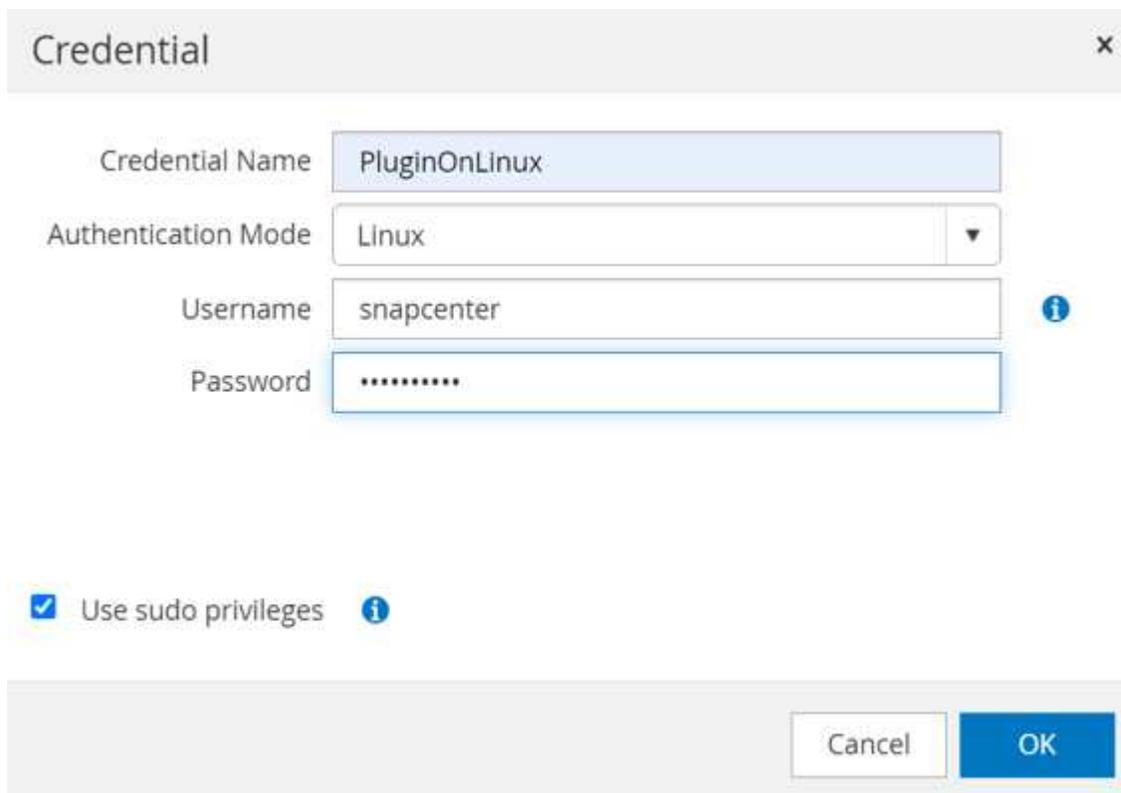
## Create credentials for plugin deployment

To enable SnapCenter to deploy the required plug-ins on the HANA hosts, you must configure user credentials.

1. Go to Settings, select Credentials, and click New.



2. In the lab setup, we configured a new user, `snapcenter`, on the HANA host that is used for the plug-in deployment. You must enable sudo privileges, as shown in the following figure.



```
hana-1:/etc/sudoers.d # cat /etc/sudoers.d/90-cloud-init-users
# Created by cloud-init v. 20.2-8.48.1 on Mon, 14 Feb 2022 10:36:40 +0000
# User rules for ec2-user
ec2-user ALL=(ALL) NOPASSWD:ALL
# User rules for snapcenter user
snapcenter ALL=(ALL) NOPASSWD:ALL
hana-1:/etc/sudoers.d #
```

## Add a SAP HANA host

When adding an SAP HANA host, SnapCenter deploys the required plug-ins on the database host and executes auto discovery operations.

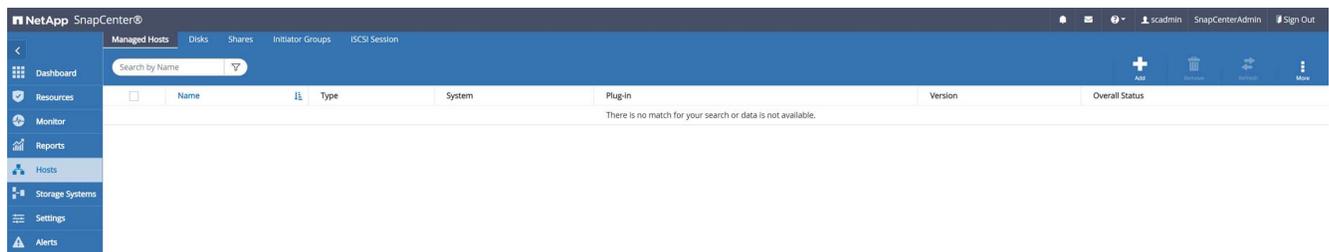
The SAP HANA plug-in requires Java 64-bit version 1.8. Java must be installed on the host before the host is added to SnapCenter.

```
hana-1:/etc/ssh # java -version
openjdk version "1.8.0_312"
OpenJDK Runtime Environment (IcedTea 3.21.0) (build 1.8.0_312-b07 suse-
3.61.3-x86_64)
OpenJDK 64-Bit Server VM (build 25.312-b07, mixed mode)
hana-1:/etc/ssh #
```

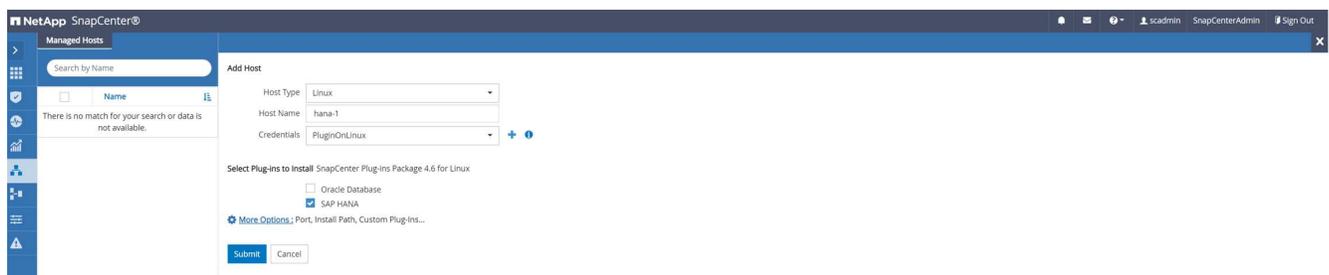
OpenJDK or Oracle Java is supported with SnapCenter.

To add the SAP HANA host, follow these steps:

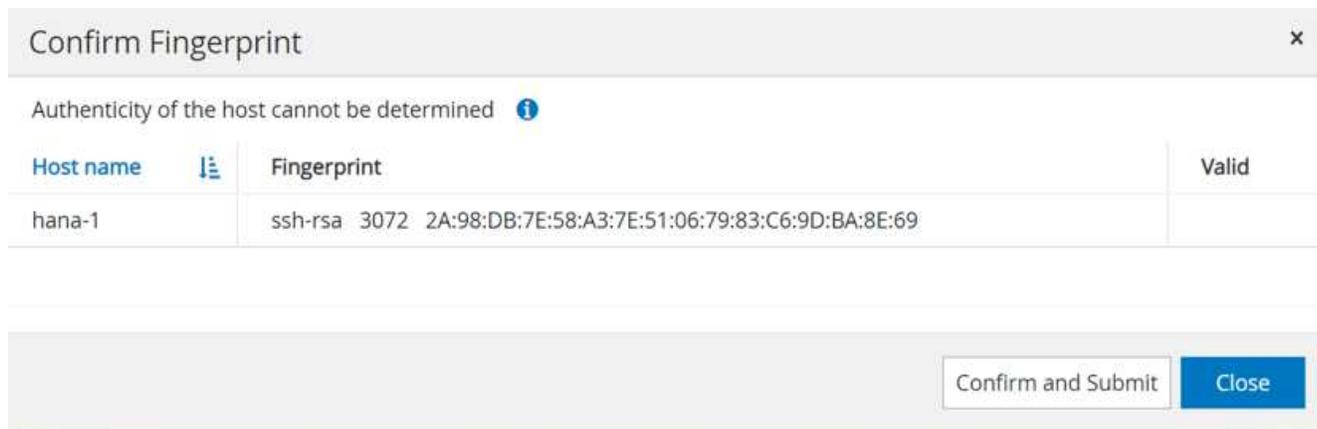
1. From the host tab, click Add.



2. Provide host information and select the SAP HANA plug-in to be installed. Click Submit.

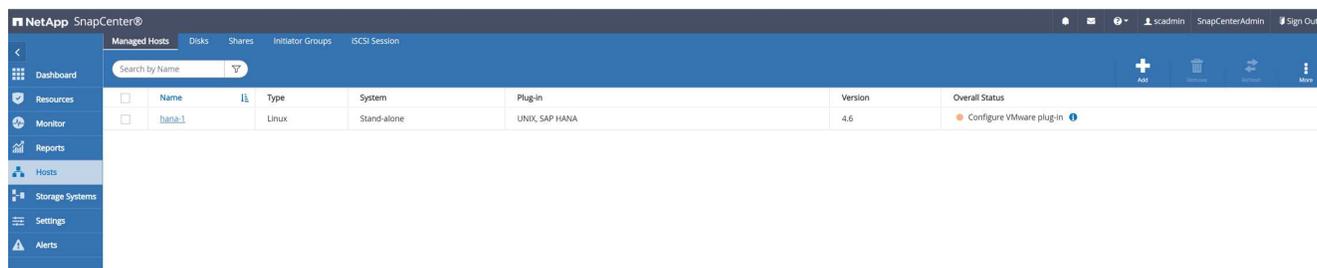


### 3. Confirm the fingerprint.



The installation of the HANA and the Linux plug-in starts automatically. When the installation is finished, the status column of the host shows Configure VMware Plug-in. SnapCenter detects if the SAP HANA plug-in is installed on a virtualized environment. This might be a VMware environment or an environment at a public cloud provider. In this case, SnapCenter displays a warning to configure the hypervisor.

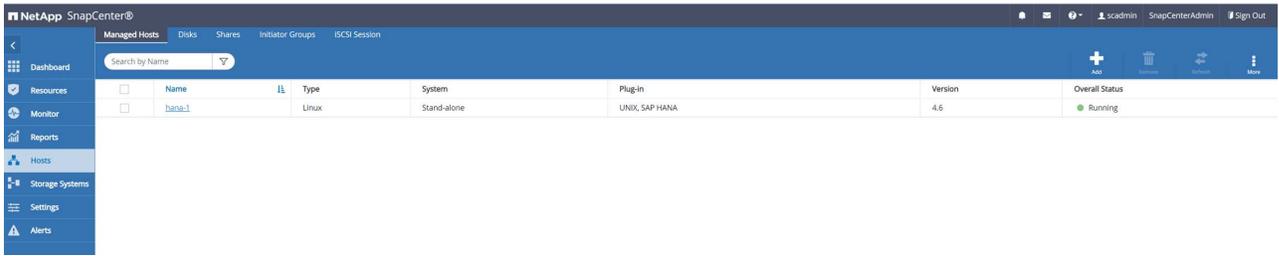
You can remove the warning message by using the following steps.



- From the Settings tab, select Global Settings.
- For the hypervisor settings, select VMs Have iSCSI Direct Attached Disks or NFS For All the Hosts and update the settings.



The screen now shows the Linux plug-in and the HANA plug-in with the status Running.



## Configure policies

Policies are usually configured independently of the resource and can be used by multiple SAP HANA databases.

A typical minimum configuration consists of the following policies:

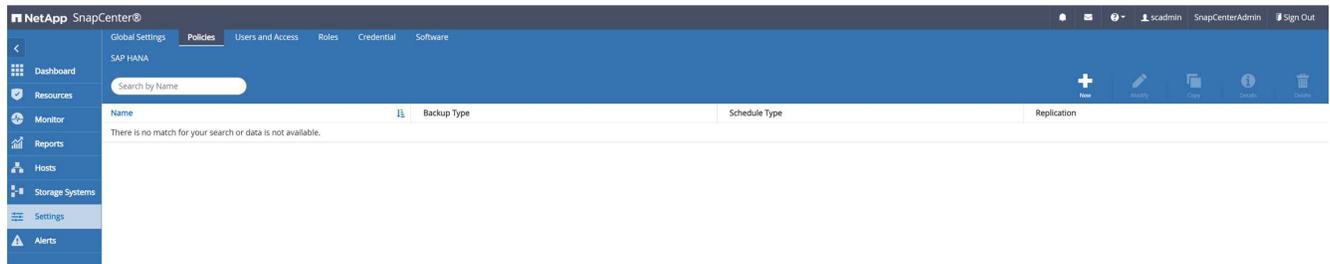
- Policy for hourly backups without replication: LocalSnap.
- Policy for weekly block integrity check using a file-based backup: BlockIntegrityCheck.

The following sections describe the configuration of these policies.

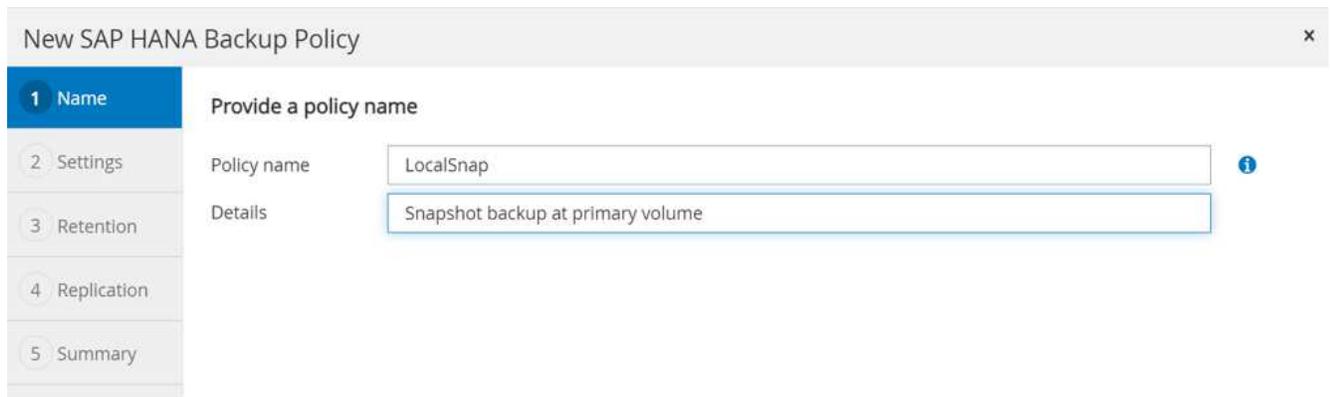
### Policy for Snapshot backups

Follow these steps to configure Snapshot backup policies.

1. Go to Settings > Policies and click New.



2. Enter the policy name and description. Click Next.



3. Select backup type as Snapshot Based and select Hourly for schedule frequency.

The schedule itself is configured later with the HANA resource protection configuration.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Select backup settings

Backup Type  Snapshot Based  File-Based i

Schedule Frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

4. Configure the retention settings for on-demand backups.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Retention settings

Hourly retention settings

Total Snapshot copies to keep  i

Keep Snapshot copies for  days

5. Configure the replication options. In this case, no SnapVault or SnapMirror update is selected.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Select secondary replication options i

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label  i

Error retry count  i

New SAP HANA Backup Policy ✕

1 Name	<b>Summary</b>	
2 Settings	Policy name	LocalSnap
3 Retention	Details	Snapshot backup at primary volume
4 Replication	Backup Type	Snapshot Based Backup
5 Summary	Schedule Type	Hourly
	Hourly backup retention	Total backup copies to retain : 7
	Replication	none

The new policy is now configured.

NetApp SnapCenter®

Global Settings Policies Users and Access Roles Credential Software

SAP HANA

Search by Name

Name	Backup Type	Schedule Type	Replication
LocalSnap	Data Backup	Hourly	

### Policy for block integrity check

Follow these steps to configure the block integrity check policy.

1. Go to Settings > Policies and click New.
2. Enter the policy name and description. Click Next.

New SAP HANA Backup Policy ✕

1 Name	<b>Provide a policy name</b>	
2 Settings	Policy name	<input style="width: 400px;" type="text" value="BlockIntegrityCheck"/>
3 Retention	Details	<input style="width: 400px;" type="text" value="Check HANA DB blocks using file-based backup"/>
4 Replication		
5 Summary		

3. Set the backup type to File-Based and schedule frequency to Weekly. The schedule itself is configured later with the HANA resource protection configuration.

### New SAP HANA Backup Policy ✕

1 Name

2 Settings

3 Retention

4 Summary

#### Select backup settings

Backup Type     Snapshot Based     File-Based i

#### Schedule Frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

4. Configure the retention settings for on-demand backups.

### New SAP HANA Backup Policy ✕

1 Name

2 Settings

3 Retention

4 Summary

#### Retention settings

Weekly retention settings

Total backup copies to keep     i

Keep backup copies for     days

5. On the Summary page, click Finish.

### New SAP HANA Backup Policy ✕

1 Name

2 Settings

3 Retention

4 Summary

#### Summary

Policy name	BlockIntegrityCheck
Details	Check HANA DB blocks using file-based backup
Backup Type	File-Based Backup
Schedule Type	Weekly
Weekly backup retention	Total backup copies to retain : 1

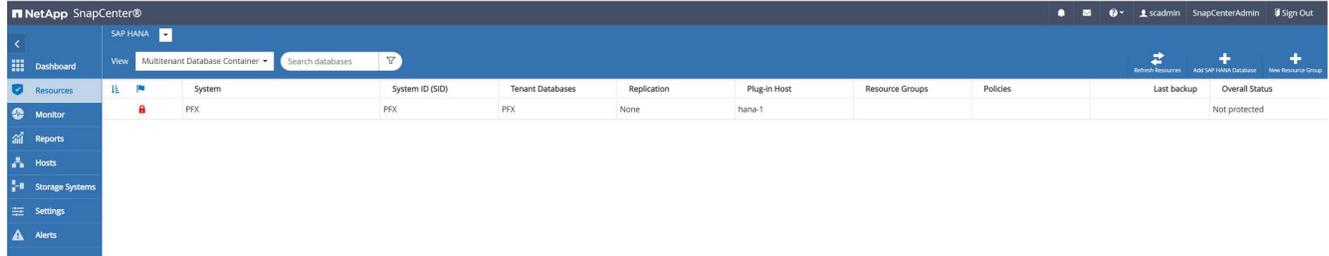
Name	Backup Type	Schedule Type	Replication
BlockIntegrityCheck	File Based Backup	Weekly	
LocalSnap	Data Backup	Hourly	

## Configure and protect a HANA resource

After the plug-in installation, the automatic discovery process of the HANA resource starts automatically. In the Resources screen, a new resource is created, which is marked as locked with the red padlock icon. To configure and protect the new HANA resource, follow these steps:

1. Select and click the resource to continue the configuration.

You can also trigger the automatic discovery process manually within the Resources screen by clicking Refresh Resources.



2. Provide the userstore key for the HANA database.

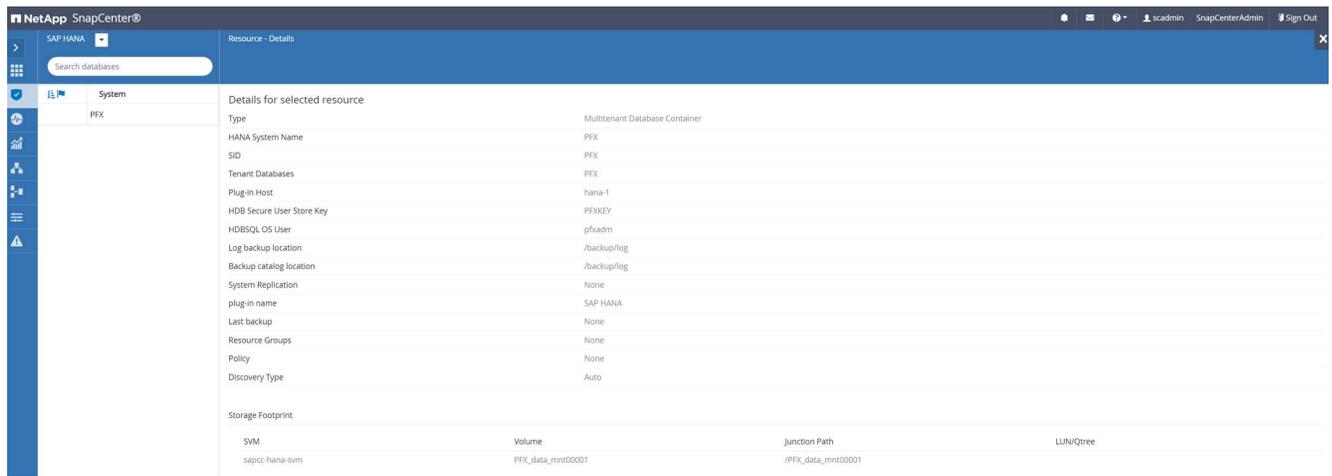
### Configure Database

Plug-in host: hana-1

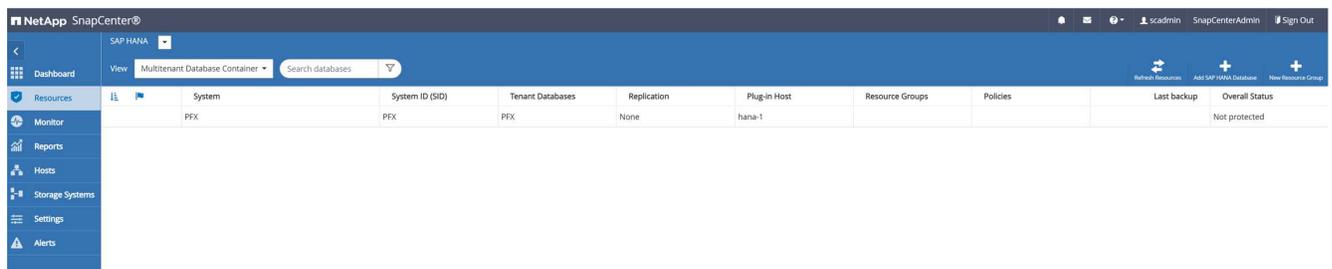
HDBSQL OS User: pfxadm

HDB Secure User Store Key:

The second level automatic discovery process starts in which tenant data and storage footprint information is discovered.

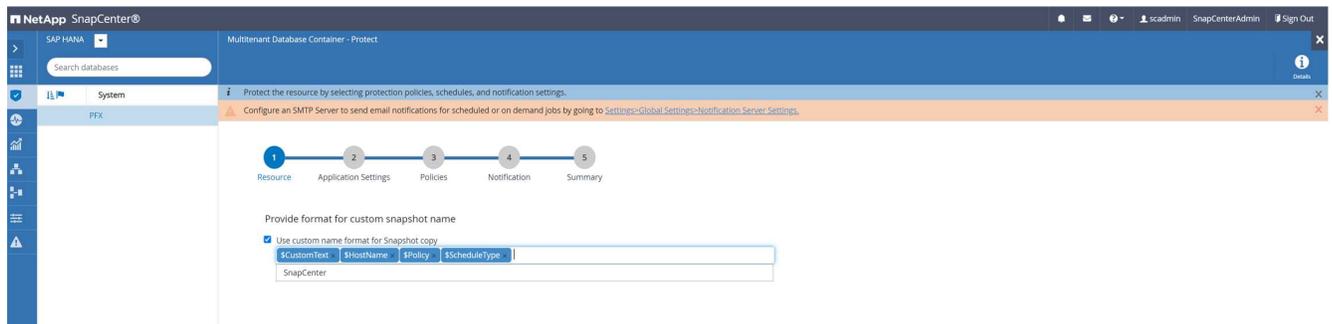


3. From the Resources tab, double click the resource to configure the resource protection.

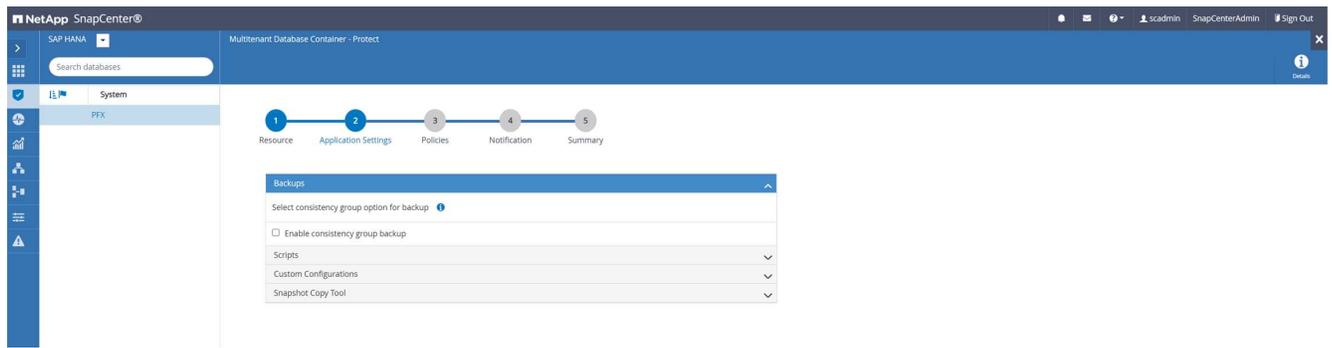


4. Configure a custom name format for the Snapshot copy.

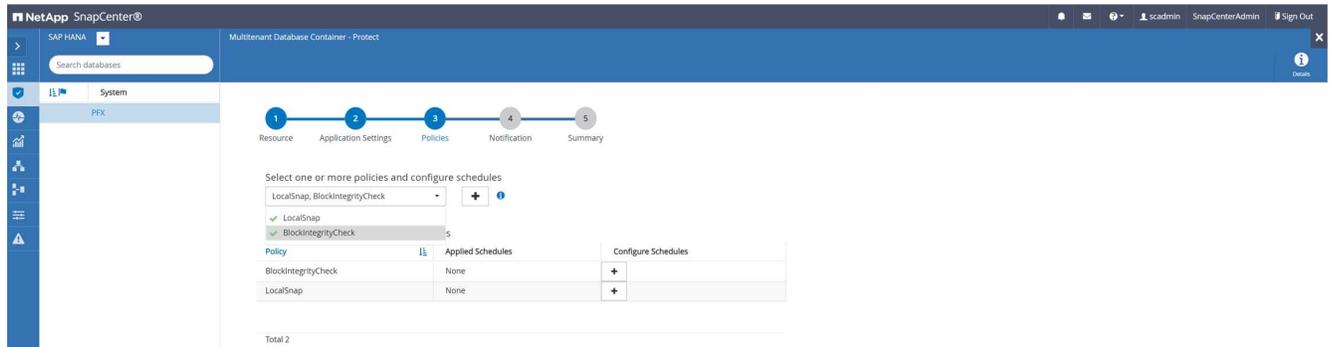
NetApp recommends using a custom Snapshot copy name to easily identify which backups have been created with which policy and schedule type. By adding the schedule type in the Snapshot copy name, you can distinguish between scheduled and on-demand backups. The `schedule` name string for on-demand backups is empty, while scheduled backups include the string `Hourly`, `Daily`, or `Weekly`.



5. No specific setting needs to be made on the Application Settings page. Click Next.



6. Select the policies to be added to the resource.



7. Define the schedule for the block integrity check policy.

In this example, it is set for once per week.

## Add schedules for policy BlockIntegrityCheck



### Weekly

Start date

02/22/2022 12:00 pm



Expires on

03/22/2022 12:00 pm



Days

Sunday

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday



The schedules are triggered in the SnapCenter Server time zone.



Cancel

OK

8. Define the schedule for the local Snapshot policy.

In this example, it is set for every 6 hours.

# Modify schedules for policy LocalSnap



## Hourly

Start date

02/22/2022 02:00 pm



Expires on

04/28/2022 11:57 am



Repeat every

6

hours

0

mins



The schedules are triggered in the SnapCenter Server time zone.



Cancel

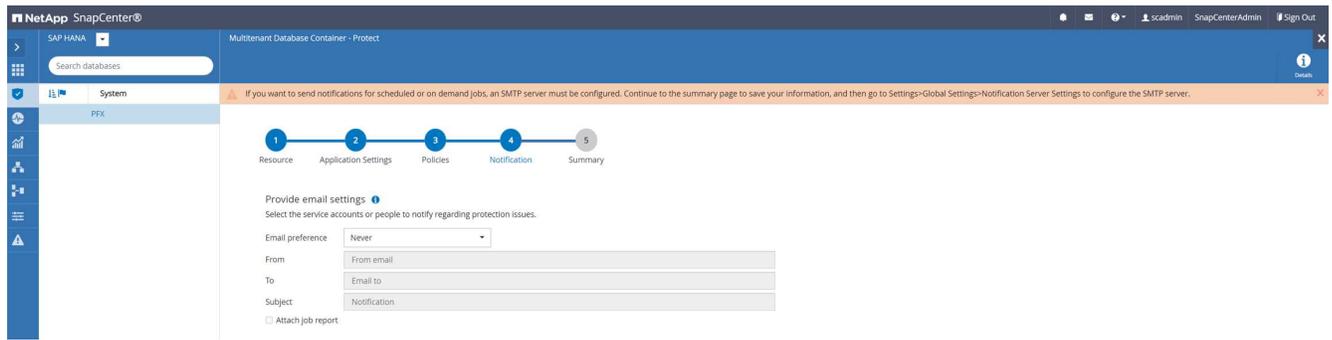
OK

The screenshot shows the NetApp SnapCenter interface for configuring policies. The breadcrumb trail is: Resource > Application Settings > Policies > Notification > Summary. The 'Policies' step is active. A dropdown menu shows 'LocalSnap, BlockIntegrityCheck' with a plus icon. Below, a table titled 'Configure schedules for selected policies' shows the configuration for two policies:

Policy	Applied Schedules	Configure Schedules
BlockIntegrityCheck	Weekly; Run on days: Sunday	
LocalSnap	Hourly; Repeat every 6 hours	

Total 2

9. Provide information about the email notification.



The HANA resource configuration is now completed, and you can execute backups.



## SnapCenter backup operations

You can create an on-demand Snapshot backup and an on-demand block integrity check operation.

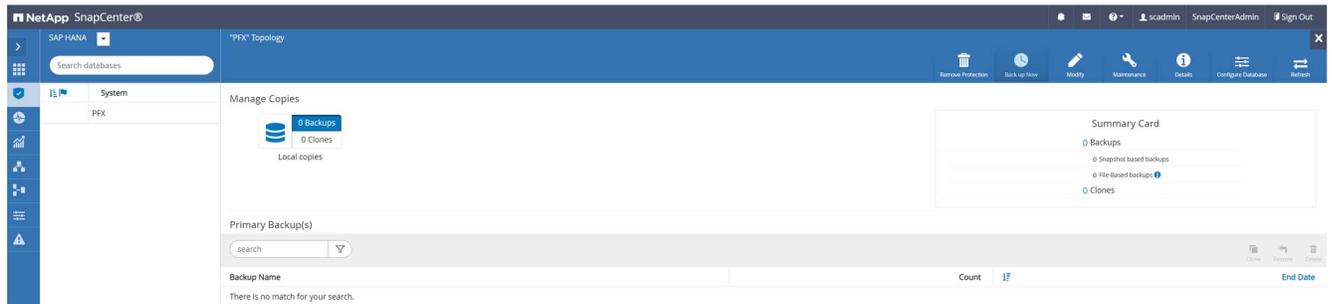
### Create an on-demand Snapshot backup

Follow these steps to create on-demand Snapshot backups.

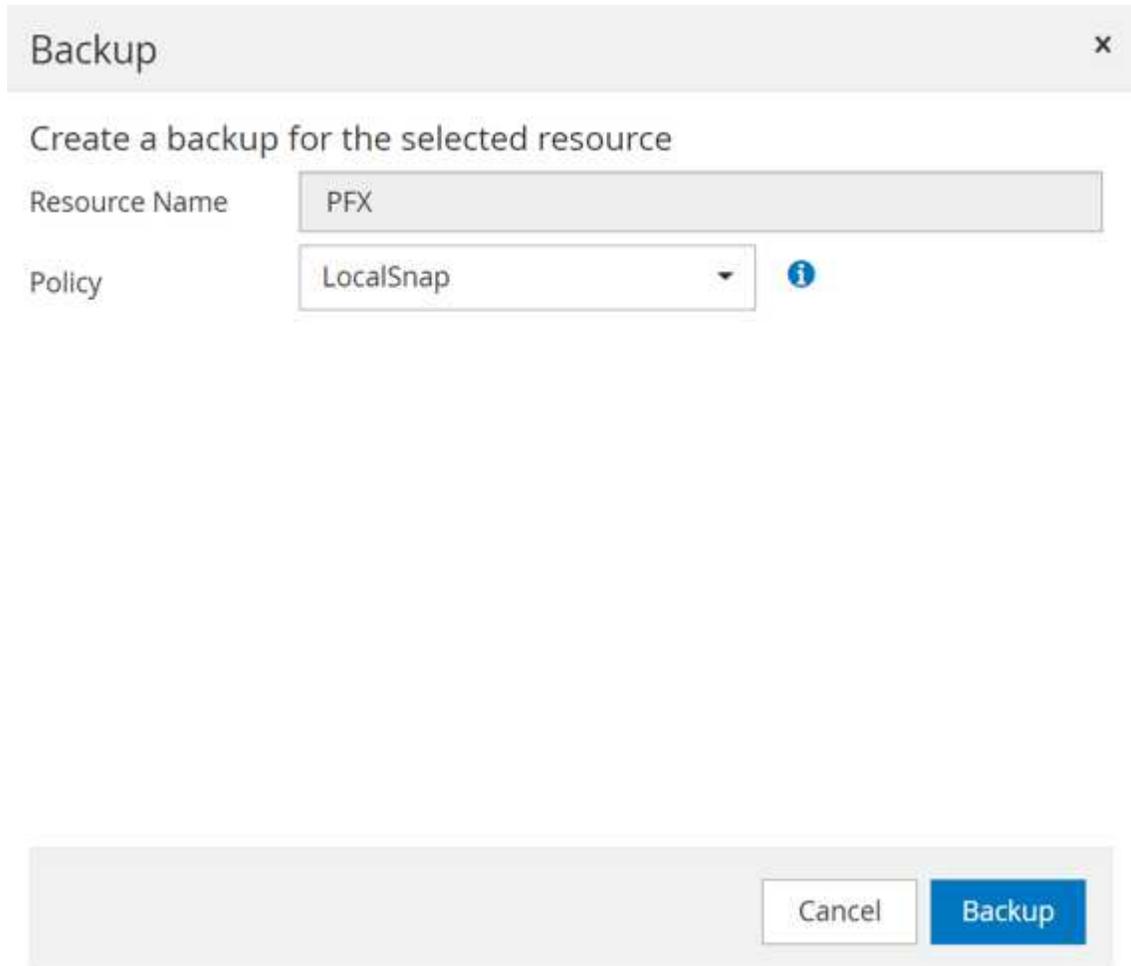
1. In the Resource view, select the resource and double-click the line to switch to the Topology view.

The Resource Topology view provides an overview of all available backups that have been created by using SnapCenter. The top area of this view displays the backup topology showing the backups on the primary storage (local copies) and, if available, on the off-site backup storage (vault copies).

2. In the top row, select the Back up Now icon to start an on-demand backup.



- From the drop-down list, select the backup policy `LocalSnap`, and then click `Backup` to start the on-demand backup.



## Confirmation



The policy selected for the on-demand backup is associated with a backup schedule and the on-demand backups will be retained based on the retention settings specified for the schedule type. Do you want to continue?

Yes

No

A log of the previous five jobs is shown in the Activity area at the bottom of the Topology view.

4. The job details are shown when clicking the job's activity line in the Activity area. You can open a detailed job log by clicking View Logs

Job Details x

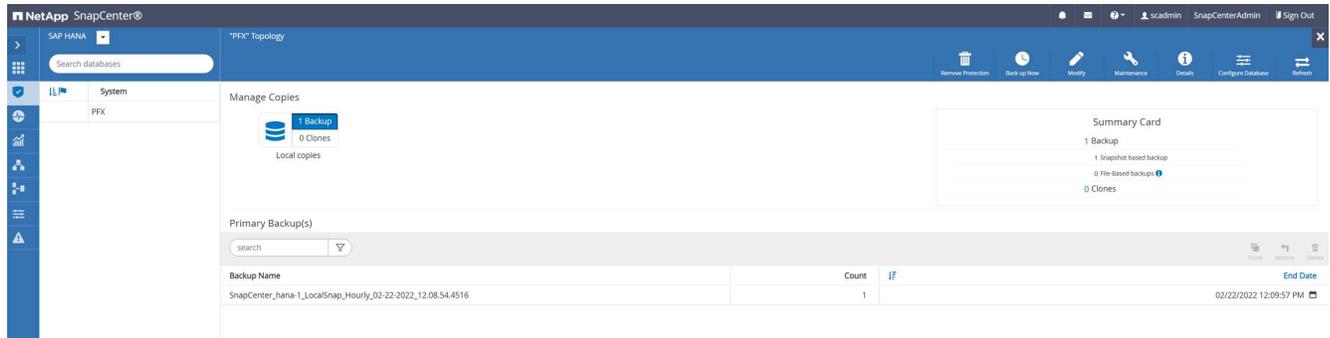
Backup of Resource Group 'hana-1\_hana\_MDC\_PFX' with policy 'LocalSnap'

- ✓ ▾ Backup of Resource Group 'hana-1\_hana\_MDC\_PFX' with policy 'LocalSnap'
- ✓ ▾ hana-1
  - ✓ Backup
    - ✓ ▶ Validate Dataset Parameters
    - ✓ ▶ Validate Plugin Parameters
    - ✓ ▶ Complete Application Discovery
    - ✓ ▶ Initialize Filesystem Plugin
    - ✓ ▶ Discover Filesystem Resources
    - ✓ ▶ Validate Retention Settings
    - ✓ ▶ Quiesce Application
    - ✓ ▶ Quiesce Filesystem
    - ✓ ▶ Create Snapshot
    - ✓ ▶ UnQuiesce Filesystem
    - ✓ ▶ UnQuiesce Application
    - ✓ ▶ Get Snapshot Details
    - ✓ ▶ Get Filesystem Meta Data
    - ✓ ▶ Finalize Filesystem Plugin
    - ✓ ▶ Collect Autosupport data
    - ✓ ▶ Register Backup and Apply Retention
    - ✓ ▶ Register Snapshot attributes
    - ✓ ▶ Application Clean-Up
    - ✓ ▶ Data Collection
    - ✓ ▶ Agent Finalize Workflow

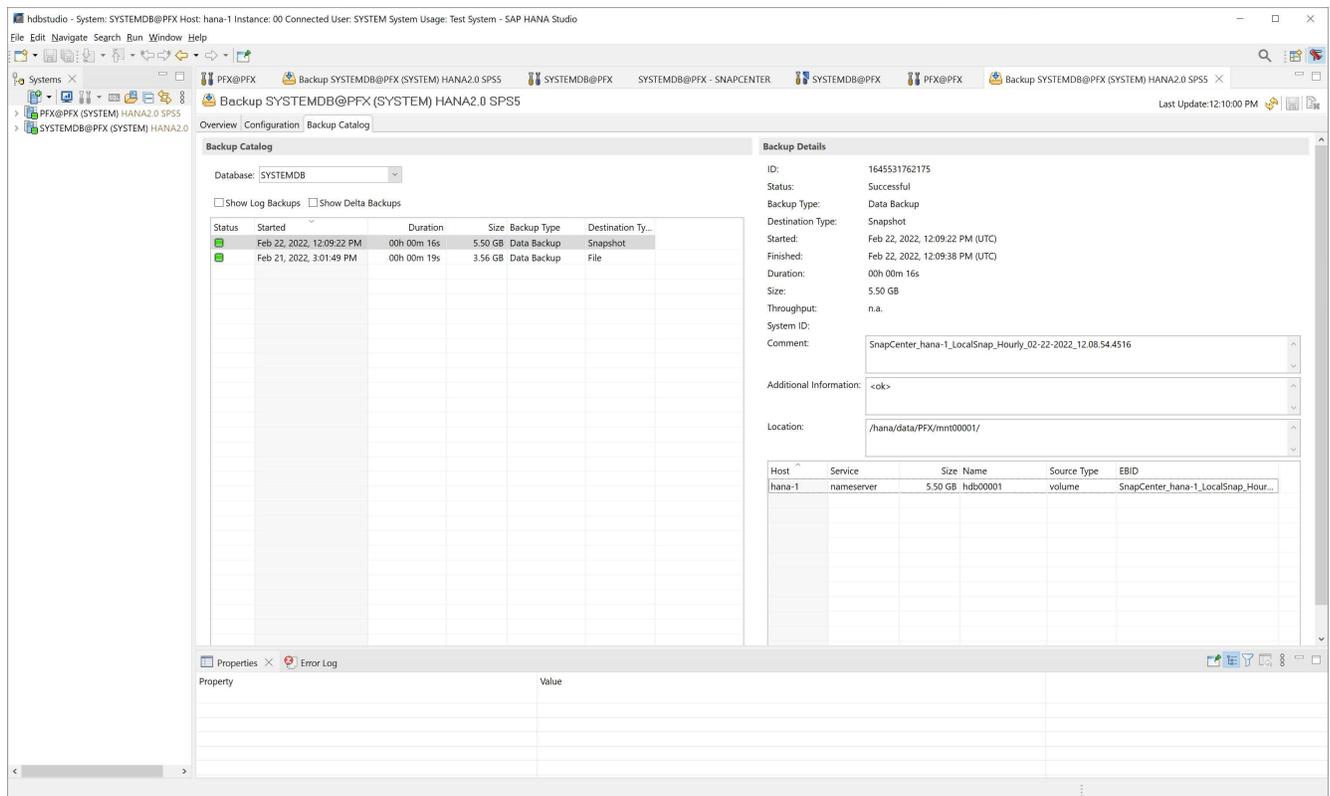
Task Name: Backup Start Time: 02/22/2022 12:08:58 PM End Time: 02/22/2022 12:10:21 PM

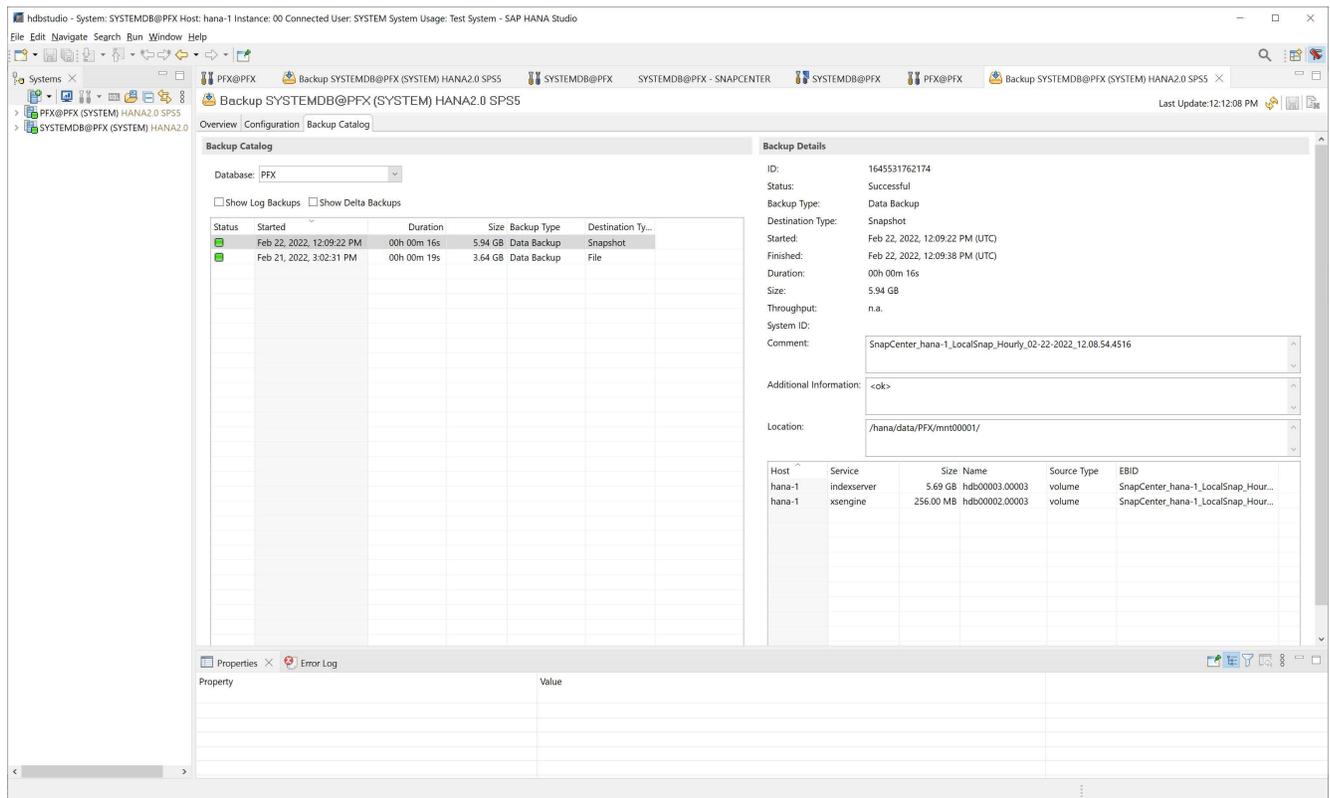
When the backup is finished, a new entry is shown in the topology view. The backup names follow the same naming convention as the Snapshot name defined in the section [“Configure and protect a HANA resource”](#).

You must close and reopen the topology view to see the updated backup list.



In the SAP HANA backup catalog, the SnapCenter backup name is stored as a Comment field as well as External Backup ID (EBID). This is shown in the following figure for the system database and in the next figure for the tenant database PFX.





On the FSx for ONTAP file system, you can list the Snapshot backups by connecting to the console of the SVM.

```

sapcc-hana-svm::> snapshot show -volume PFX_data_mnt00001
---Blocks---
Vserver   Volume      Snapshot                                     Size Total%
Used%
-----
sapcc-hana-svm
          PFX_data_mnt00001
          SnapCenter_hana-1_LocalSnap_Hourly_02-22-
2022_12.08.54.4516
                                           126.6MB    0%
2%
sapcc-hana-svm::>

```

### Create an on-demand block integrity check operation

An on-demand block integrity check operation is executed in the same way as a Snapshot backup job, by selecting the policy BlockIntegrityCheck. When scheduling backups using this policy, SnapCenter creates a standard SAP HANA file backup for the system and tenant databases.

## Backup



Create a backup for the selected resource

Resource Name

PFX

Policy

BlockIntegrityCheck



Cancel

Backup

## Job Details



Backup of Resource Group 'hana-1\_hana\_MDC\_PFX' with policy 'BlockIntegrityCheck'

✓ ▾ Backup of Resource Group 'hana-1\_hana\_MDC\_PFX' with policy 'BlockIntegrityCheck'

✓ ▾ hana-1

✓ ▾ File-Based Backup

- ✓ ▶ Validate Plugin Parameters
- ✓ ▶ Start File-Based Backup
- ✓ ▶ Check File-Based Backup
- ✓ ▶ Register Backup and Apply Retention
- ✓ ▶ Data Collection

**i** Task Name: File-Based Backup Start Time: 02/22/2022 12:55:21 PM End Time: 02/22/2022 12:56:36 PM

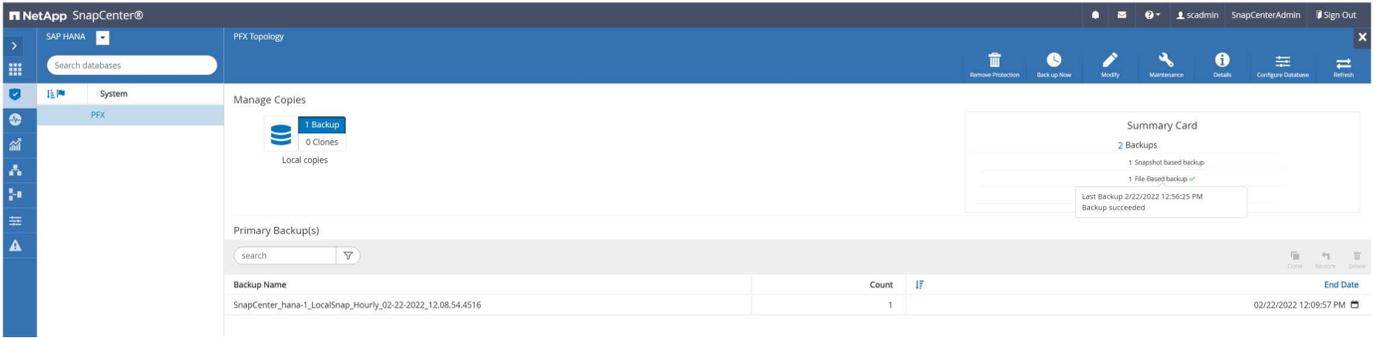
View Logs

Cancel Job

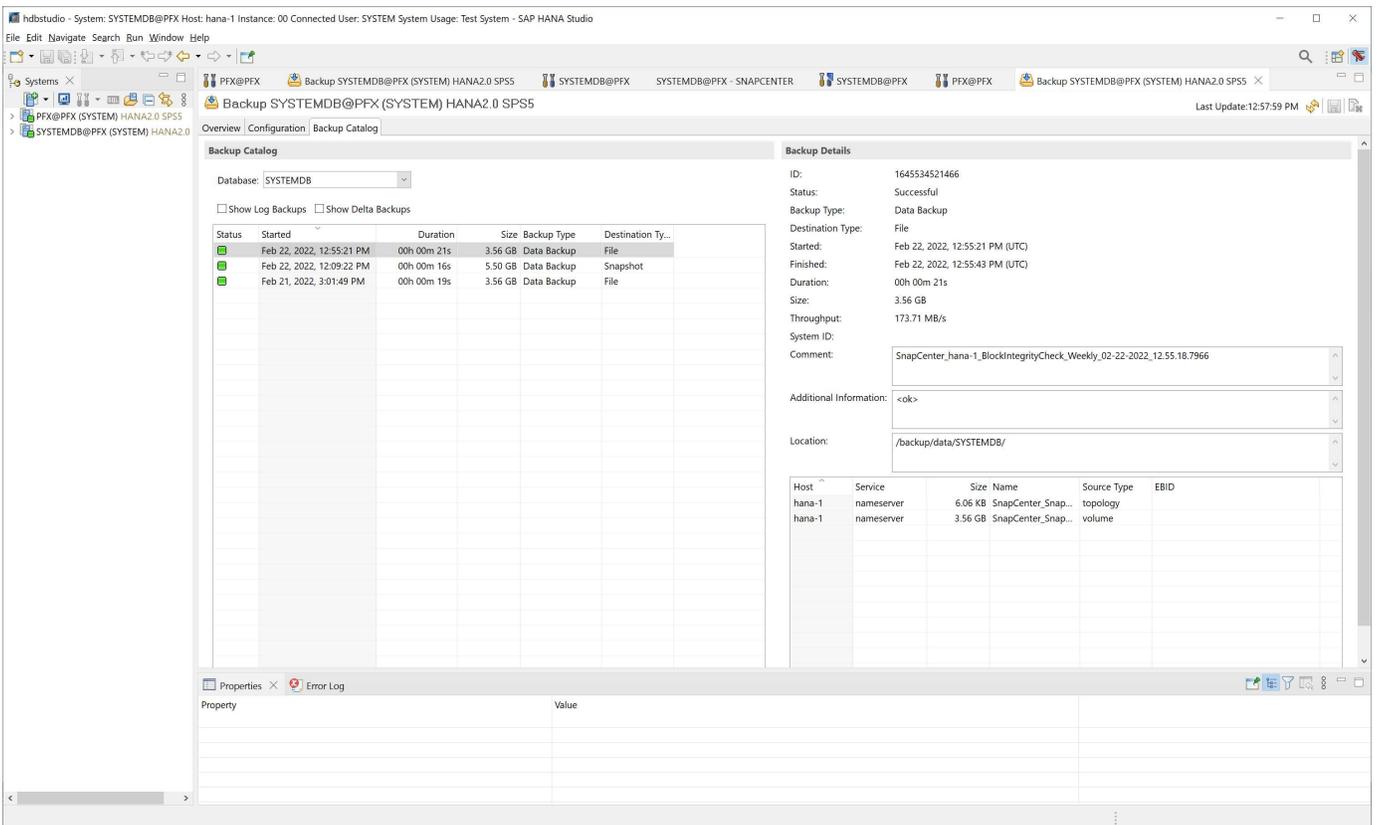
Close

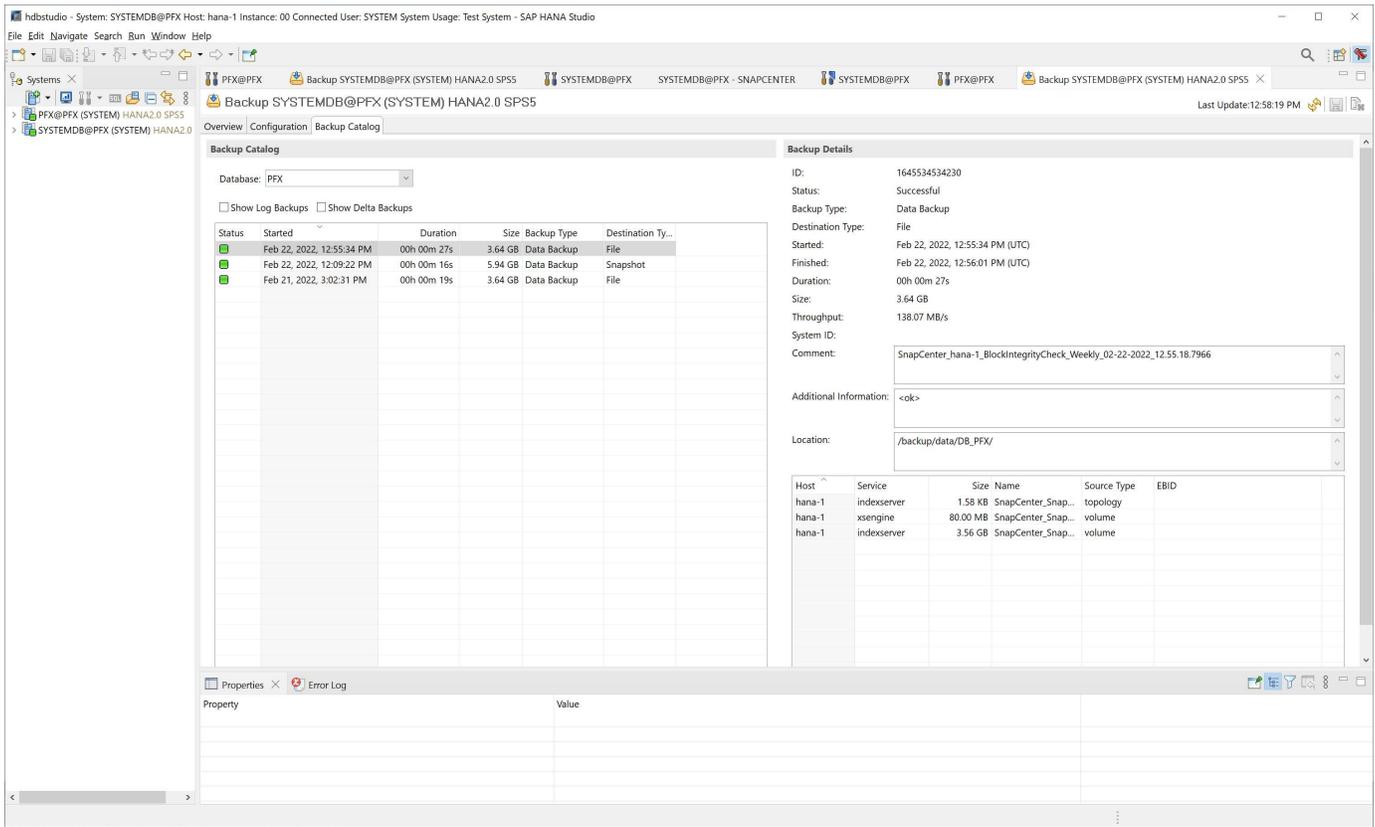
SnapCenter does not display the block integrity check in the same manner as Snapshot copy-based backups.

Instead, the summary card shows the number of file-based backups and the status of the previous backup.



The SAP HANA backup catalog shows entries for both the system and the tenant databases. The following figures show the SnapCenter block integrity check in the backup catalog of the system and the tenant database.





A successful block integrity check creates standard SAP HANA data backup files. SnapCenter uses the backup path that has been configured with the HANA database for file-based data backup operations.

```

hana-1:~ # ls -al /backup/data/*
/backup/data/DB_PFX:
total 7665384
drwxr-xr-- 2 pfxadm sapsys      4096 Feb 22 12:56 .
drwxr-xr-x 4 pfxadm sapsys      4096 Feb 21 15:02 ..
-rw-r----- 1 pfxadm sapsys    155648 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_0_1
-rw-r----- 1 pfxadm sapsys    83894272 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_2_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_3_1
-rw-r----- 1 pfxadm sapsys     155648 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_0_1
-rw-r----- 1 pfxadm sapsys    83894272 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_2_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 22 12:56
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_3_1
/backup/data/SYSTEMDB:
total 7500880
drwxr-xr-- 2 pfxadm sapsys      4096 Feb 22 12:55 .
drwxr-xr-x 4 pfxadm sapsys      4096 Feb 21 15:02 ..
-rw-r----- 1 pfxadm sapsys    159744 Feb 21 15:01
COMPLETE_DATA_BACKUP_databackup_0_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_1_1
-rw-r----- 1 pfxadm sapsys     159744 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_0_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_1_1
hana-1:~ #

```

## Backup of non-data volumes

The backup of non-data volumes is an integrated part of the SnapCenter and the SAP HANA plug-in.

Protecting the database data volume is sufficient to restore and recover the SAP HANA database to a given point in time, provided that the database installation resources, and the required logs are still available.

To recover from situations where other non-data files must be restored, NetApp recommends developing an additional backup strategy for non-data volumes to augment the SAP HANA database backup. Depending on

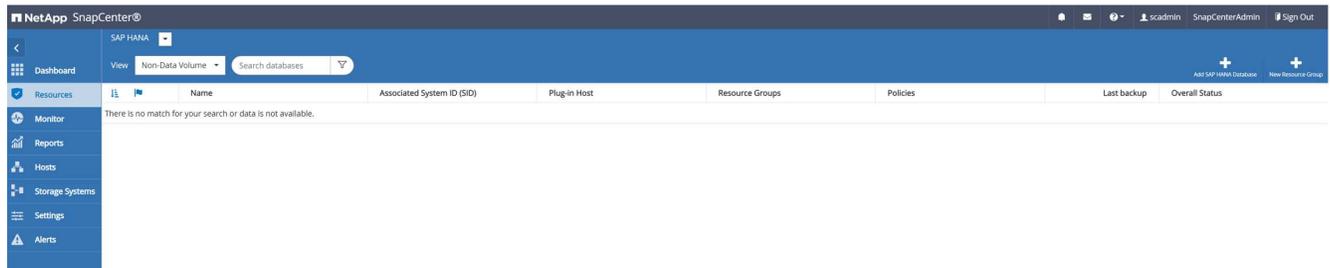
your specific requirements, the backup of non-data volumes might differ in scheduling frequency and retention settings, and you should consider how frequently non-data files are changed. For instance, the HANA volume `/hana/shared` contains executables but also SAP HANA trace files. While executables only change when the SAP HANA database is upgraded, the SAP HANA trace files might need a higher backup frequency to support analyzing problem situations with SAP HANA.

SnapCenter non-data volume backup enables Snapshot copies of all relevant volumes to be created in a few seconds with the same space efficiency as SAP HANA database backups. The difference is that there is no SQL communication with SAP HANA database required.

### Configure non-data volume resources

Follow these steps to configure non-data volume resources:

1. From the Resources tab, select Non-Data-Volume and click Add SAP HANA Database.



2. In step one of the Add SAP HANA Database dialog, in the Resource Type list, select Non- data Volumes. Specify a name for the resource and the associated SID and the SAP HANA plug-in host that you want to use for the resource, then click Next.

Add SAP HANA Database ×

- 1 Name
- 2 Storage Footprint
- 3 Summary

### Provide Resource Details

Resource Type	Non-data Volume <span style="float: right;">▼</span>
Resource Name	PFX-Shared-Volume
Associated SID	PFX <span style="float: right;">?</span>
Plug-in Host	hana-1 <span style="float: right;">▼</span> <span style="float: right;">?</span>

PreviousNext

3. Add the SVM and the storage volume as storage footprint, then click Next.

Add SAP HANA Database x

1 Name

**2 Storage Footprint**

3 Summary

### Provide Storage Footprint Details

Storage Type  ONTAP

Add Storage Footprint x

Storage System

Select one or more volumes and if required their associated Qtrees and LUNs

Volume name	LUNs or Qtrees
<input type="text" value="PFX_shared"/>	<input type="text" value="Default is 'None' or type to find"/>

4. To save the settings, in the summary step, click Finish.

### Add SAP HANA Database

- 1 Name
- 2 Storage Footprint
- 3 Summary

#### Summary

Resource Type	Non-data Volume
Resource Name	PFX-Shared-Volume
Associated SID	PFX
Plug-in Host	hana-1

#### Storage Footprint

Storage System	Volume	LUN/Qtree
sapcc-hana-svm	PFX_shared	

Previous
Finish

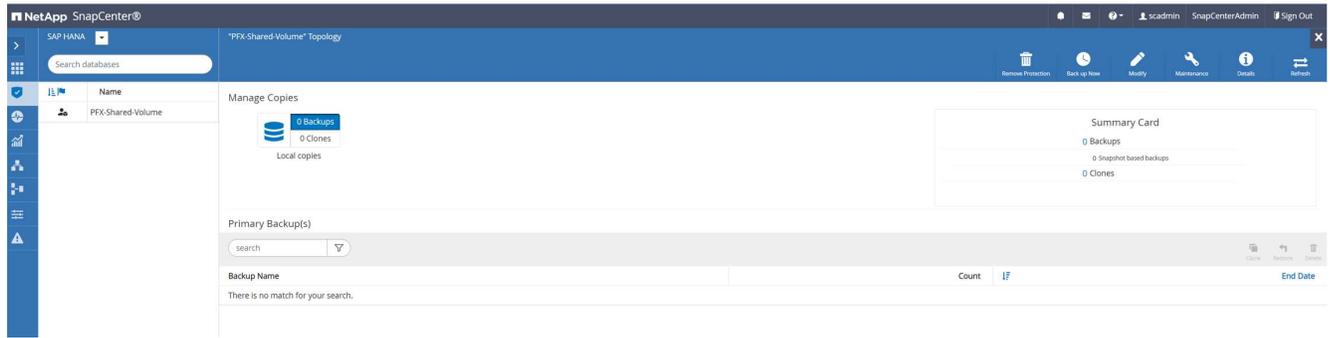
The new non-data volume is now added to SnapCenter. Double click the new resource to execute the resource protection.

The screenshot shows the NetApp SnapCenter interface. At the top, there's a navigation bar with 'SAP HANA' selected. Below it, a table lists resources. The table has columns for Name, Associated System ID (SID), Plug-in Host, Resource Groups, Policies, Last backup, and Overall Status. One resource is listed: PFX-Shared-Volume, associated with SID PFX and Plug-in Host hana-1. The Overall Status is 'Not protected'.

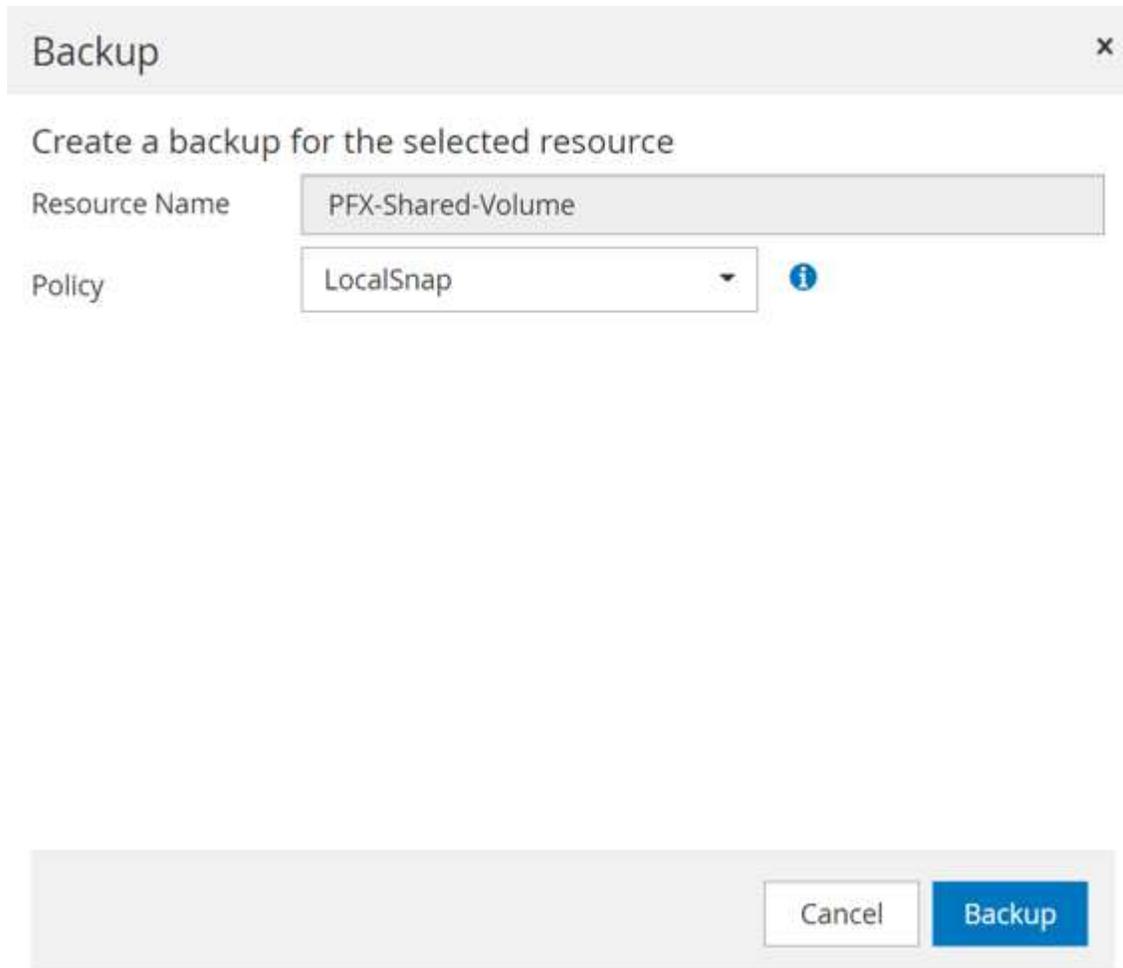
Name	Associated System ID (SID)	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
PFX-Shared-Volume	PFX	hana-1				Not protected

The resource protection is done in the same way as described before with a HANA database resource.

5. You can now execute a backup by clicking on Backup Now.



6. Select the policy and start the backup operation.



The SnapCenter job log shows the individual workflow steps.

## Job Details



Backup of Resource Group 'hana-1\_hana\_NonDataVolume\_PFX\_PFX-Shared-Volume' with policy 'LocalSnap'

✓ ▾ Backup of Resource Group 'hana-1\_hana\_NonDataVolume\_PFX\_PFX-Shared-Volume' with policy 'LocalSnap'

✓ ▾ hana-1

✓ ▾ Backup

- ✓ ▶ Validate Dataset Parameters
- ✓ ▶ Validate Plugin Parameters
- ✓ ▶ Validate Retention Settings
- ✓ ▶ Create Snapshot
- ✓ ▶ Get Snapshot Details
- ✓ ▶ Collect Autosupport data
- ✓ ▶ Register Backup and Apply Retention
- ✓ ▶ Register Snapshot attributes
- ✓ ▶ Data Collection
- ✓ ▶ Agent Finalize Workflow

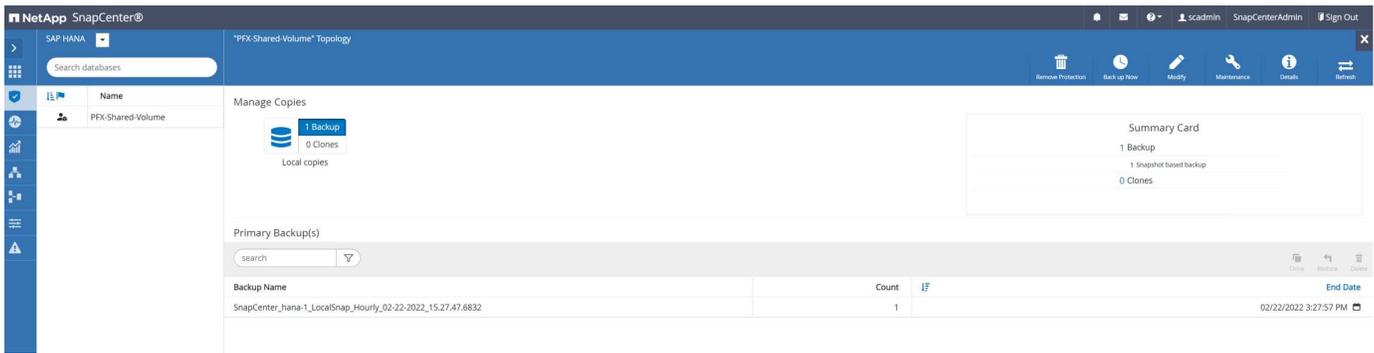
**i** Task Name: Backup Start Time: 02/22/2022 3:27:48 PM End Time:

View Logs

Cancel Job

Close

The new backup is now visible in the resource view of the non- data volume resource.



## Restore and recover

With SnapCenter, automated restore and recovery operations are supported for HANA single host MDC systems with a single tenant. For multiple-host systems or MDC systems with multiple tenants, SnapCenter only executes the restore operation and you must perform the recovery manually.

You can execute an automated restore and recovery operation with the following steps:

1. Select the backup to be used for the restore operation.
2. Select the restore type. Select Complete Restore with Volume Revert or without Volume Revert.
3. Select the recovery type from the following options:
  - To most recent state
  - Point in time
  - To specific data backup
  - No recovery

The selected recovery type is used for the recovery of the system and the tenant database.

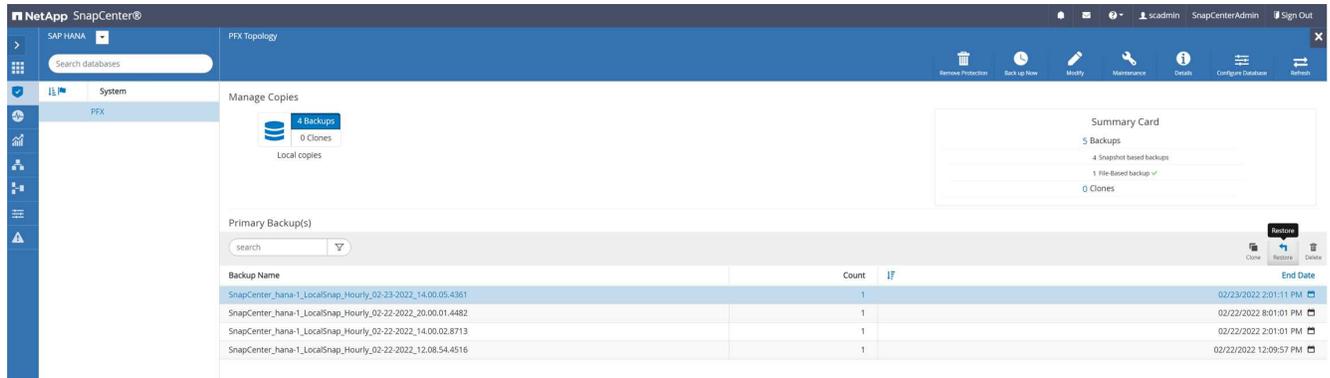
Next, SnapCenter performs the following operations:

1. It stops the HANA database.
2. It restores the database. Depending on the selected restore type, different operations are executed.
  - If Volume Revert is selected, then SnapCenter unmounts the volume, restores the volume by using volume-based SnapRestore on the storage layer, and mounts the volume.
  - If Volume Revert is not selected, then SnapCenter restores all files by using single file SnapRestore operations on the storage layer.
3. It recovers the database:
  - a. By recovering the system database
  - b. recovering the tenant database
  - c. starting the HANA database

If No Recovery is selected, SnapCenter exits, and you must perform the restore operation for the system and the tenant database manually.

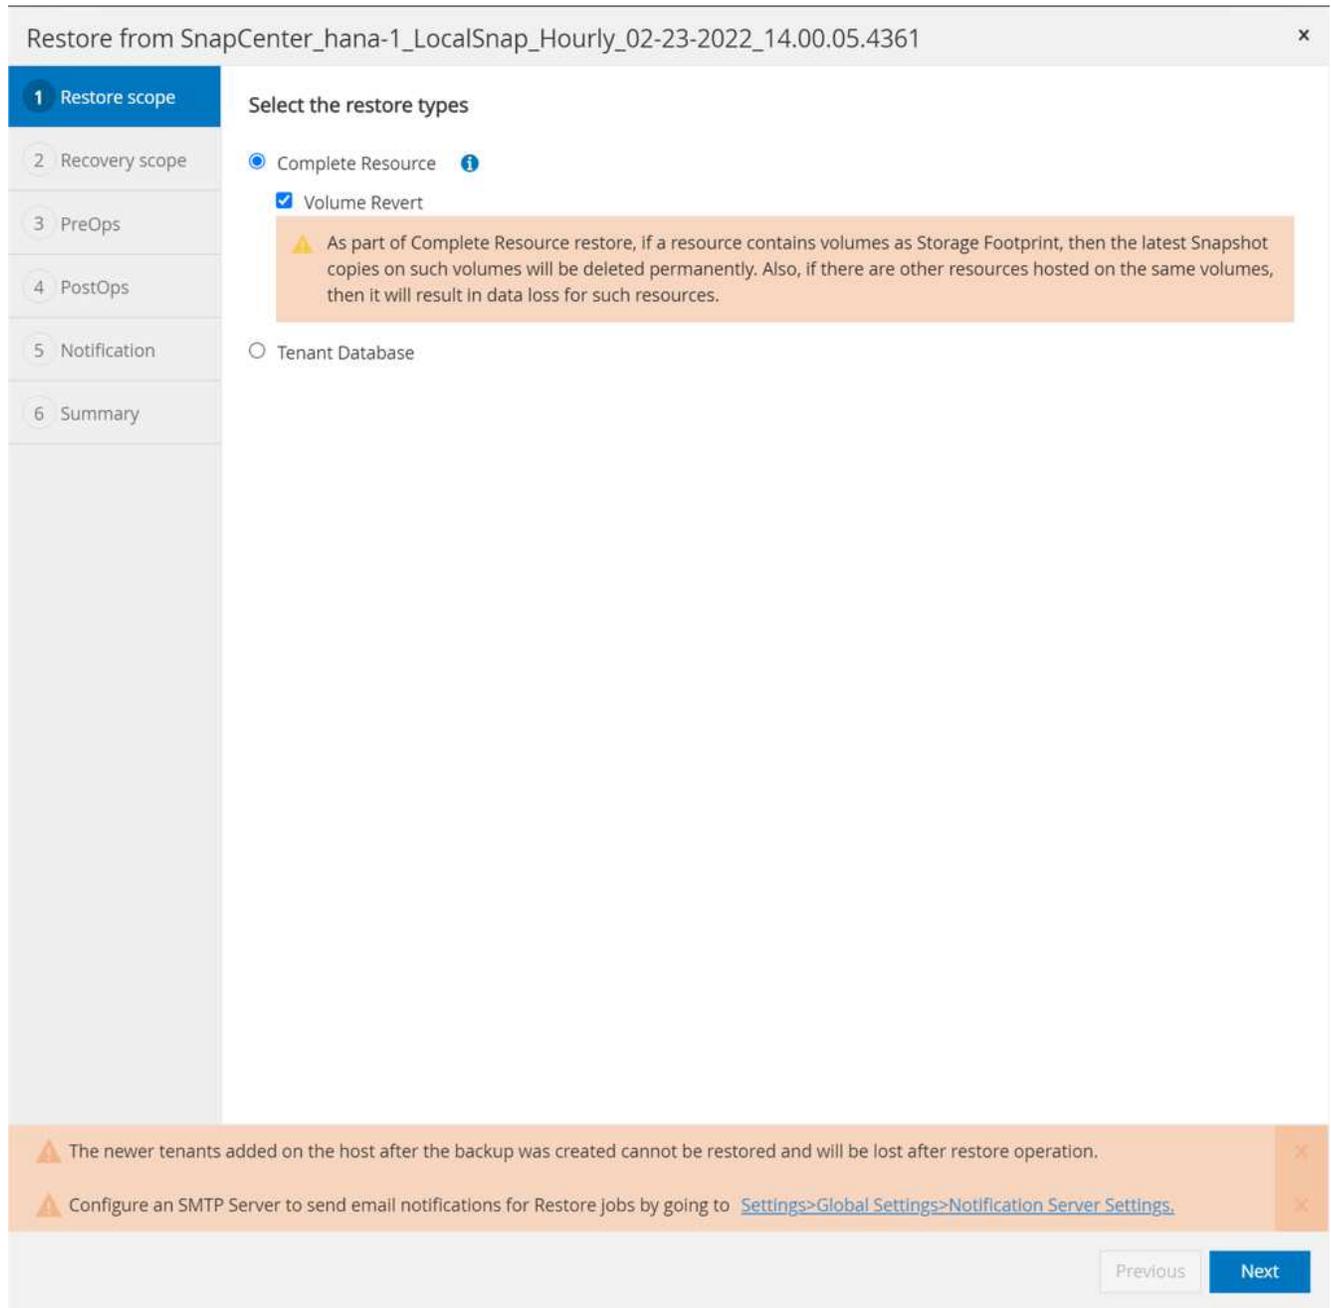
To perform a manual restore operation, follow these steps:

1. Select a backup in SnapCenter to be used for the restore operation.



2. Select the restore scope and type.

The standard scenario for HANA MDC single tenant systems is to use complete resource with volume revert. For a HANA MDC system with multiple tenants, you might want to restore only a single tenant. For more information about the single tenant restore, see [Restore and recovery \(netapp.com\)](https://netapp.com).



3. Select Recovery Scope and provide the location for log backup and catalog backup.

SnapCenter uses the default path or the changed paths in the HANA global.ini file to pre-populate the log and catalog backup locations.

Restore from SnapCenter\_hana-1\_LocalSnap\_Hourly\_02-23-2022\_14.00.05.4361 ×

- 1 Restore scope
- 2 Recovery scope**
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

**Recover database files using**

- Recover to most recent state ?
- Recover to point in time ?
- Recover to specified data backup ?
- No recovery ?

**Specify log backup locations** ?

Add

**Specify backup catalog location** ?

⚠ Recovery options are applicable to both system database and tenant database. ×

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

Previous Next

4. Enter the optional pre-restore commands.

Restore from SnapCenter\_hana-1\_LocalSnap\_Hourly\_02-23-2022\_14.00.05.4361 ×

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps**
- 4 PostOps
- 5 Notification
- 6 Summary

Enter optional commands to run before performing a restore operation ?

Pre restore command

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

Previous Next

5. Enter the optional post-restore commands.

Restore from SnapCenter\_hana-1\_LocalSnap\_Hourly\_02-23-2022\_14.00.05.4361 ×

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps**
- 5 Notification
- 6 Summary

Enter optional commands to run after performing a restore operation ?

Post restore command

[Previous](#) [Next](#)

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

6. To start the restore and recovery operation, click Finish.

Restore from SnapCenter\_hana-1\_LocalSnap\_Hourly\_02-23-2022\_14.00.05.4361
×

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

### Summary

Backup Name	SnapCenter_hana-1_LocalSnap_Hourly_02-23-2022_14.00.05.4361
Backup date	02/23/2022 2:01:11 PM
Restore scope	Complete Resource with Volume Revert
Recovery scope	Recover to most recent state
Log backup locations	/backup/log
Backup catalog location	/backup/log
Pre restore command	
Post restore command	
Send email	No

⚠ If you want to send notifications for Restore Jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous
Finish

SnapCenter executes the restore and recovery operation. This example shows the job details of the restore and recovery job.

## Job Details



### Restore 'hana-1\hana\MDC\PFX'

- ✓ ▾ Restore 'hana-1\hana\MDC\PFX'
- ✓ ▾ hana-1
  - ✓ ▾ Restore
    - ✓ ▶ Validate Plugin Parameters
    - ✓ ▾ Pre Restore Application
      - ✓ ▶ Stopping HANA instance
    - ✓ ▶ Filesystem Pre Restore
    - ✓ ▾ Restore Filesystem
    - ✓ ▶ Filesystem Post Restore
    - ✓ ▾ Recover Application
      - ✓ ▶ Recovering system database
      - ✓ ▶ Checking HDB services status
      - ✓ ▶ Recovering tenant database 'PFX'
      - ✓ ▶ Starting HANA instance
    - ✓ ▶ Clear Catalog on Server
    - ✓ ▶ Application Clean-Up
    - ✓ ▶ Data Collection
    - ✓ ▶ Agent Finalize Workflow

**i** Task Name: Recover Application Start Time: 02/23/2022 2:07:31 PM End Time:

View Logs

Cancel Job

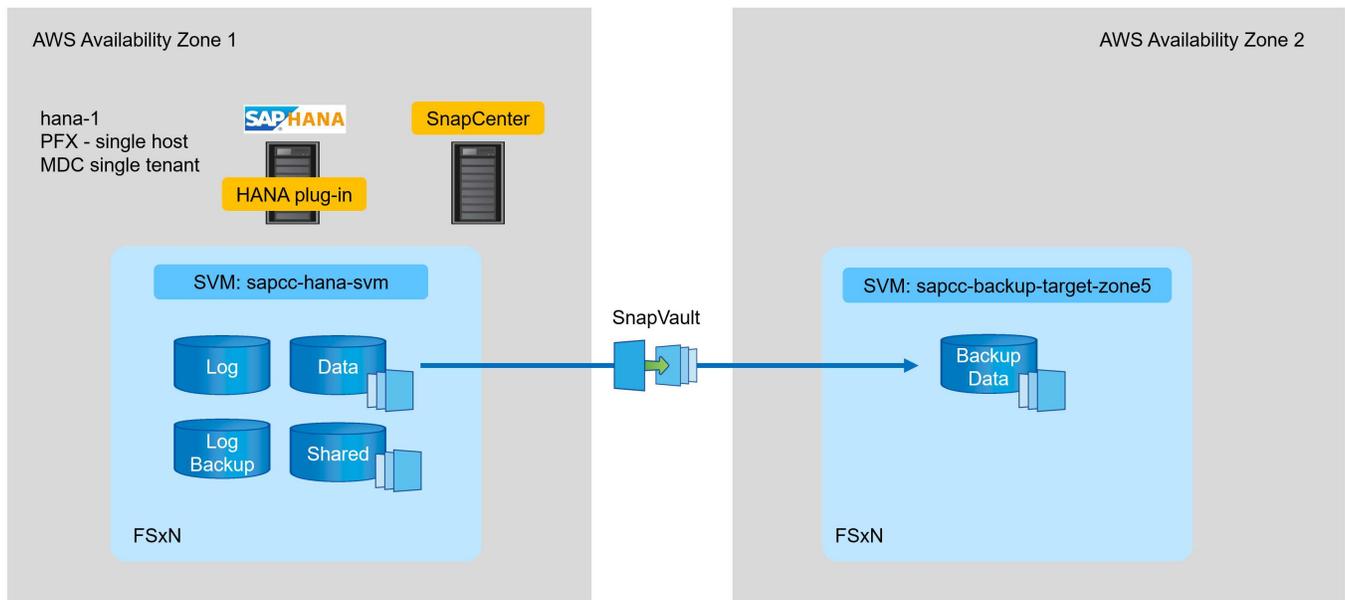
Close

## Backup replication with SnapVault

### Overview - Backup replication with SnapVault

In our lab setup, we use a second FSX for ONTAP file system in a second AWS availability zone to showcase the backup replication for the HANA data volume.

As discussed in chapter “[Data protection strategy](#)”, the replication target must be a second FSx for ONTAP file system in another availability zone to be protected from a failure of the primary FSx for ONTAP file system. Also, the HANA shared volume should be replicated to the secondary FSx for ONTAP file system.



### Overview of configuration steps

There are a couple of configuration steps that you must execute on the FSx for ONTAP layer. You can do this either with NetApp Cloud Manager or the FSx for ONTAP command line.

1. Peer FSx for ONTAP file systems. FSx for ONTAP file systems must be peered to allow replication between each other.
2. Peer SVMs. SVMs must be peered to allow replication between each other.
3. Create a target volume. Create a volume at the target SVM with volume type `DP`. Type `DP` is required to be used as a replication target volume.
4. Create a SnapMirror policy. This is used to create a policy for replication with type `vault`.
  - a. Add a rule to policy. The rule contains the SnapMirror label and the retention for backups at the secondary site. You must configure the same SnapMirror label later in the SnapCenter policy so that SnapCenter creates Snapshot backups at the source volume containing this label.
5. Create a SnapMirror relationship. Defines the replication relationship between the source and target volume and attaches a policy.
6. Initialize SnapMirror. This starts the initial replication in which the complete source data is transferred to the target volume.

When volume replication configuration is complete, you must configure the backup replication in SnapCenter

as follows:

1. Add the target SVM to SnapCenter.
2. Create a new SnapCenter policy for Snapshot backup and SnapVault replication.
3. Add the policy to HANA resource protection.
4. You can now execute backups with the new policy.

The following chapters describe the individual steps in more detail.

### Configure replication relationships on FSx for ONTAP file systems

You can find additional information about SnapMirror configuration options in the ONTAP documentation at [SnapMirror replication workflow \(netapp.com\)](https://netapp.com).

- Source FSx for ONTAP file system: FsxId00fa9e3c784b6abbb
- Source SVM: sapcc-hana-svm
- Target FSx for ONTAP file system: FsxId05f7f00af49dc7a3e
- Target SVM: sapcc-backup-target-zone5

### Peer FSx for ONTAP file systems

```
FsxId00fa9e3c784b6abbb::> network interface show -role intercluster
      Logical      Status      Network      Current      Current
Is
Vserver      Interface  Admin/Oper  Address/Mask      Node      Port
Home
-----
----
FsxId00fa9e3c784b6abbb
      inter_1      up/up      10.1.1.57/24
FsxId00fa9e3c784b6abbb-01
true
      inter_2      up/up      10.1.2.7/24
FsxId00fa9e3c784b6abbb-02
true
      e0e
      e0e
2 entries were displayed.
```

```

FsxId05f7f00af49dc7a3e::> network interface show -role intercluster
          Logical      Status      Network      Current      Current
Is
Vserver   Interface  Admin/Oper  Address/Mask  Node         Port
Home
-----
----
FsxId05f7f00af49dc7a3e
          inter_1      up/up      10.1.2.144/24
FsxId05f7f00af49dc7a3e-01
                                     e0e

true
          inter_2      up/up      10.1.2.69/24
FsxId05f7f00af49dc7a3e-02
                                     e0e

true
2 entries were displayed.

```

```

FsxId05f7f00af49dc7a3e::> cluster peer create -address-family ipv4 -peer
-addr 10.1.1.57, 10.1.2.7
Notice: Use a generated passphrase or choose a passphrase of 8 or more
characters. To ensure the authenticity of the peering relationship, use a
phrase or sequence of characters that would be hard to guess.
Enter the passphrase:
Confirm the passphrase:
Notice: Now use the same passphrase in the "cluster peer create" command
in the other cluster.

```



peer-addr are cluster IPs of the destination cluster.

```

FsxId00fa9e3c784b6abbb::> cluster peer create -address-family ipv4 -peer
-addr 10.1.2.144, 10.1.2.69
Notice: Use a generated passphrase or choose a passphrase of 8 or more
characters. To ensure the authenticity of the peering relationship, use a
phrase or sequence of characters that would be hard to guess.
Enter the passphrase:
Confirm the passphrase:
FsxId00fa9e3c784b6abbb::>
FsxId00fa9e3c784b6abbb::> cluster peer show
Peer Cluster Name          Cluster Serial Number Availability
Authentication
-----
FsxId05f7f00af49dc7a3e    1-80-000011             Available          ok

```

### Peer SVMs

```

FsxId05f7f00af49dc7a3e::> vserver peer create -vserver sapcc-backup-
target-zone5 -peer-vserver sapcc-hana-svm -peer-cluster
FsxId00fa9e3c784b6abbb -applications snapmirror
Info: [Job 41] 'vserver peer create' job queued

```

```

FsxId00fa9e3c784b6abbb::> vserver peer accept -vserver sapcc-hana-svm
-peer-vserver sapcc-backup-target-zone5
Info: [Job 960] 'vserver peer accept' job queued

```

```

FsxId05f7f00af49dc7a3e::> vserver peer show
Peer          Peer          Peering
Remote
Vserver      Vserver      State      Peer Cluster      Applications
Vserver
-----
sapcc-backup-target-zone5
peer-source-cluster
peered      FsxId00fa9e3c784b6abbb
snapmirror
sapcc-hana-svm

```

### Create a target volume

You must create the target volume with the type `DP` to flag it as a replication target.

```
FsxId05f7f00af49dc7a3e::> volume create -vserver sapcc-backup-target-zone5
-volume PFX_data_mnt00001 -aggregate aggr1 -size 100GB -state online
-policy default -type DP -autosize-mode grow_shrink -snapshot-policy none
-foreground true -tiering-policy all -anti-ransomware-state disabled
[Job 42] Job succeeded: Successful
```

### Create a SnapMirror policy

The SnapMirror policy and the added rule define the retention and the Snapmirror label to identify Snapshots that should be replicated. When creating the SnapCenter policy later, you must use the same label.

```
FsxId05f7f00af49dc7a3e::> snapmirror policy create -policy snapcenter-
policy -tries 8 -transfer-priority normal -ignore-atime false -restart
always -type vault -vserver sapcc-backup-target-zone5
```

```
FsxId05f7f00af49dc7a3e::> snapmirror policy add-rule -vserver sapcc-
backup-target-zone5 -policy snapcenter-policy -snapmirror-label
snapcenter -keep 14
```

```
FsxId00fa9e3c784b6abbb::> snapmirror policy showVserver Policy
```

Policy Number	Transfer					
Name	Name	Type	Of Rules	Tries	Priority	Comment
-----						
FsxId00fa9e3c784b6abbb	snapcenter-policy	vault	1	8	normal	-
	SnapMirror Label: snapcenter					Keep: 14
						Total Keep: 14

### Create SnapMirror relationship

Now the relation between the source and target volume is defined as well as the type XDP and the policy we created earlier.

```
FsxId05f7f00af49dc7a3e::> snapmirror create -source-path sapcc-hana-
svm:PFX_data_mnt00001 -destination-path sapcc-backup-target-
zone5:PFX_data_mnt00001 -vserver sapcc-backup-target-zone5 -throttle
unlimited -identity-preserve false -type XDP -policy snapcenter-policy
Operation succeeded: snapmirror create for the relationship with
destination "sapcc-backup-target-zone5:PFX_data_mnt00001".
```

## Initialize SnapMirror

With this command, the initial replication starts. This is a full transfer of all data from the source volume to the target volume.

```
FsxId05f7f00af49dc7a3e::> snapmirror initialize -destination-path sapcc-
backup-target-zone5:PFX_data_mnt00001 -source-path sapcc-hana-
svm:PFX_data_mnt00001
Operation is queued: snapmirror initialize of destination "sapcc-backup-
target-zone5:PFX_data_mnt00001".
```

You can check the status of the replication with the `snapmirror show` command.

```
FsxId05f7f00af49dc7a3e::> snapmirror show

Progress
Source          Destination Mirror Relationship Total
Last
Path           Type Path           State Status           Progress Healthy
Updated
-----
-----
sapcc-hana-svm:PFX_data_mnt00001
                XDP  sapcc-backup-target-zone5:PFX_data_mnt00001
                Uninitialized
                Transferring 1009MB true
02/24 12:34:28
```

```
FsxId05f7f00af49dc7a3e::> snapmirror show

Progress
Source          Destination Mirror Relationship Total
Last
Path           Type Path           State Status           Progress Healthy
Updated
-----
-----
sapcc-hana-svm:PFX_data_mnt00001
                XDP  sapcc-backup-target-zone5:PFX_data_mnt00001
                Snapmirrored
                Idle           - true -
```

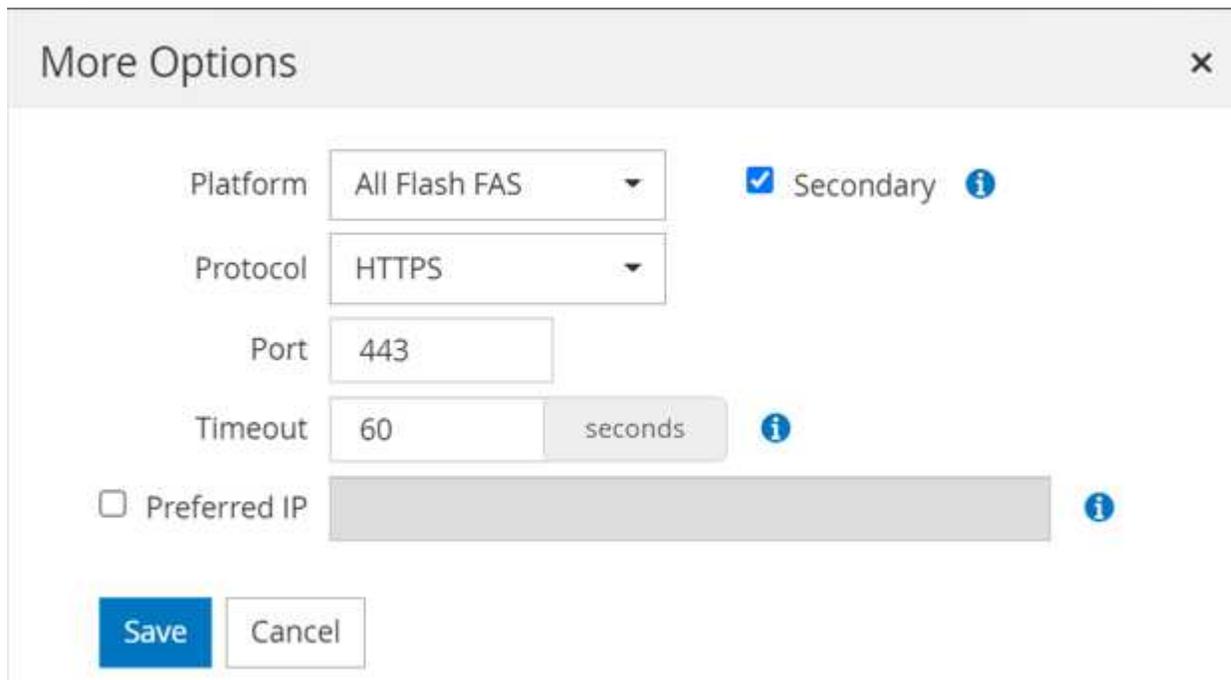
## Add a backup SVM to SnapCenter

To add a backup SVM to SnapCenter, follow these steps:

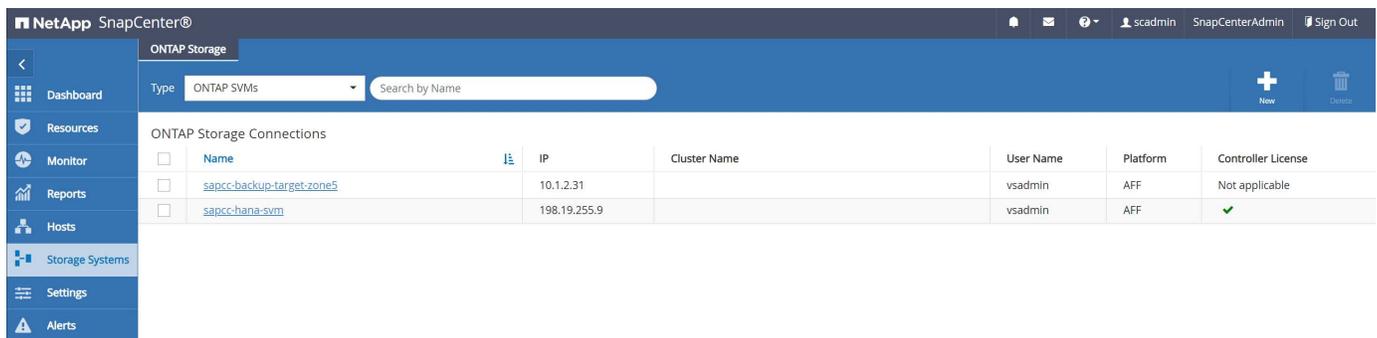
1. Configure the SVM where the SnapVault target volume is located in SnapCenter.



2. On the More Options window, select All Flash FAS as the platform and select Secondary.



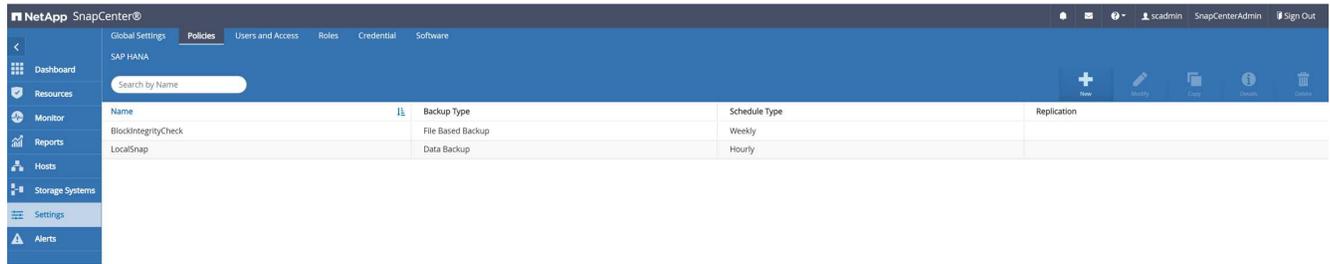
The SVM is now available in SnapCenter.



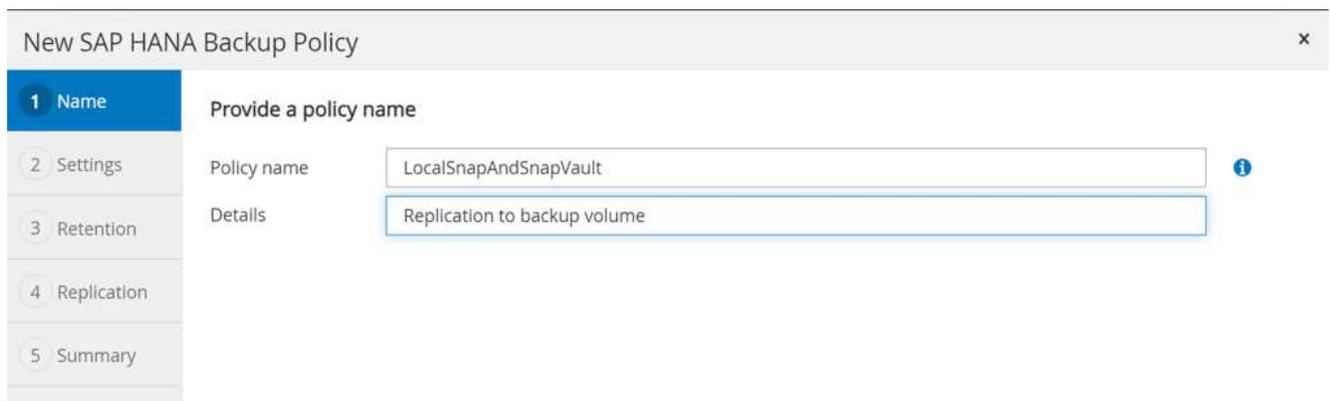
## Create a new SnapCenter policy for backup replication

You must configure a policy for the backup replication as follows:

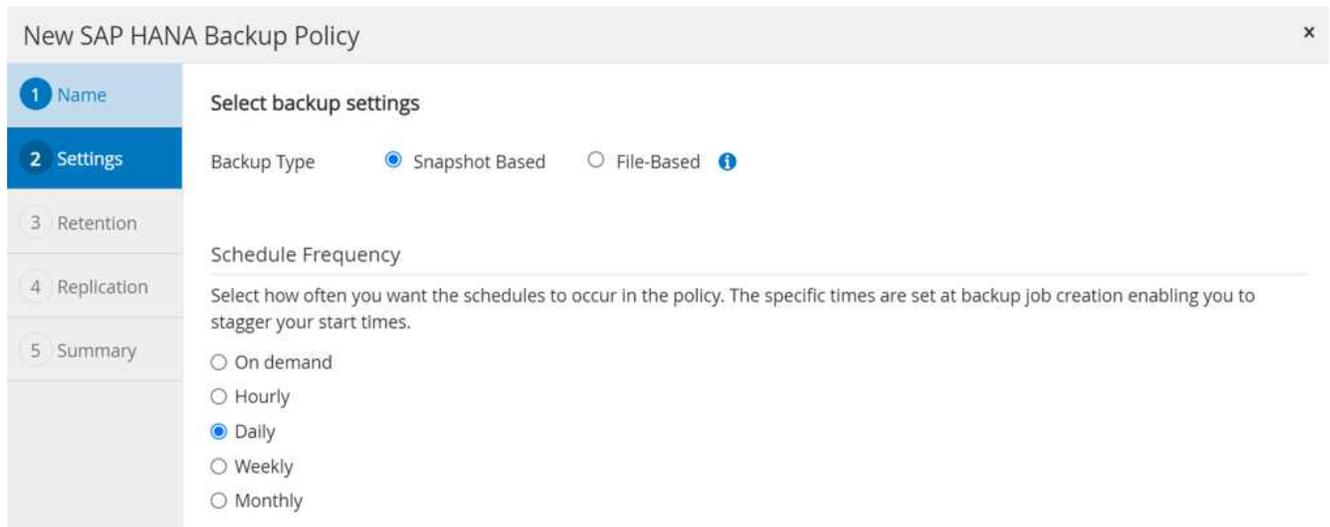
1. Provide a name for the policy.



2. Select Snapshot backup and a schedule frequency. Daily is typically used for backup replication.



3. Select the retention for the Snapshot backups.



This is the retention for the daily Snapshot backups taken at the primary storage. The retention for secondary backups at the SnapVault target has already been configured previously using the add rule command at the ONTAP level. See “Configure replication relationships on FSx for ONTAP file systems” (xref).

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Retention settings

Daily retention settings

Total Snapshot copies to keep  i

Keep Snapshot copies for  days

4. Select the Update SnapVault field and provide a custom label.

This label must match the SnapMirror label provided in the `add rule` command at ONTAP level.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Select secondary replication options i

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label  i

Error retry count  i

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Summary

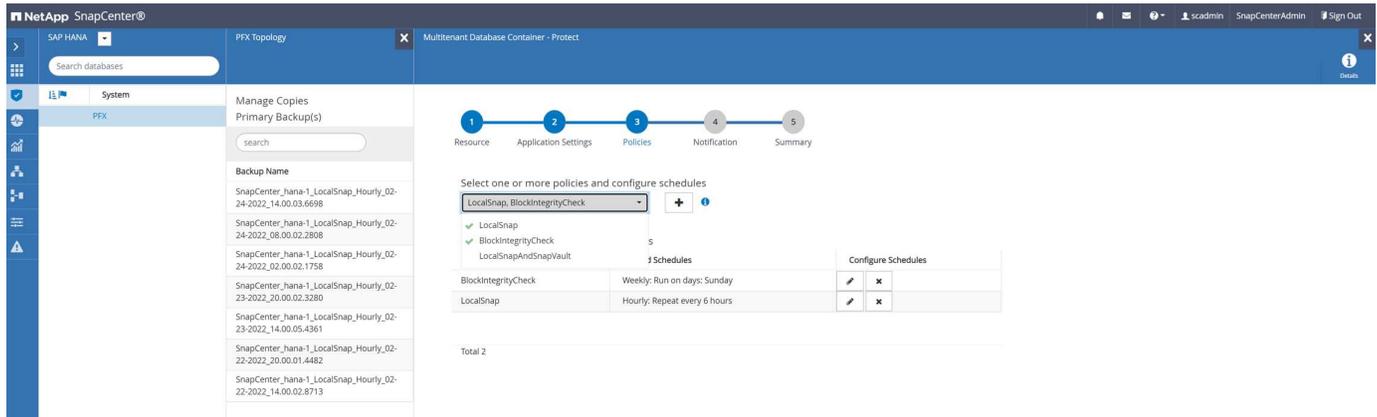
Policy name	LocalSnapAndSnapVault
Details	Replication to backup volume
Backup Type	Snapshot Based Backup
Schedule Type	Daily
Daily backup retention	Total backup copies to retain : 3
Replication	SnapVault enabled , Secondary policy label: Custom Label : snapcenter , Error retry count: 3

The new SnapCenter policy is now configured.

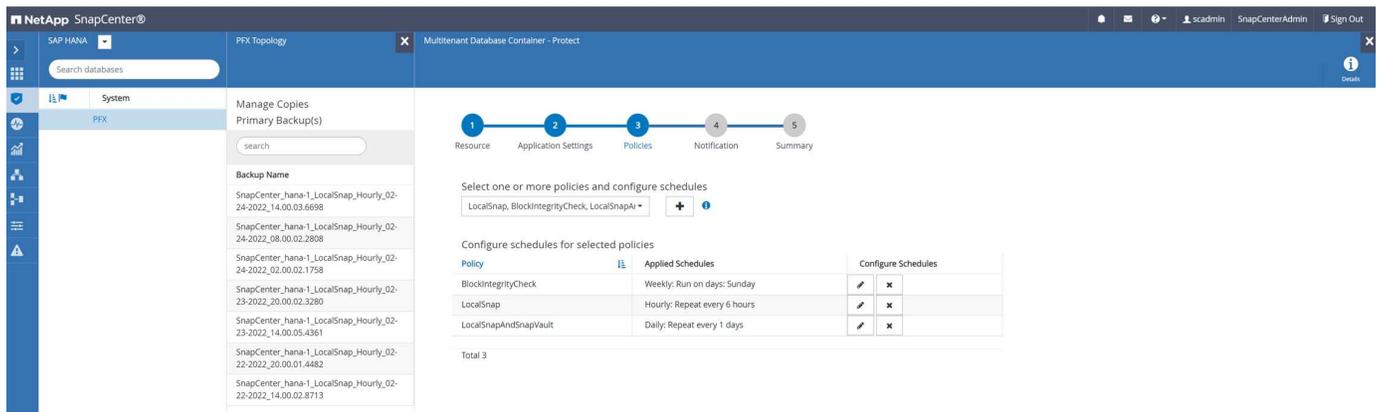
Name	Backup Type	Schedule Type	Replication
BlockIntegrityCheck	File Based Backup	Weekly	
LocalSnap	Data Backup	Hourly	
LocalSnapAndSnapVault	Data Backup	Daily	SnapVault

## Add a policy to resource protection

You must add the new policy to the HANA resource protection configuration, as shown in the following figure.



A daily schedule is defined in our setup.



## Create a backup with replication

A backup is created in the same way as with a local Snapshot copy.

To create a backup with replication, select the policy that includes the backup replication and click Backup.

Backup ×

Create a backup for the selected resource

Resource Name

Policy  i

Within the SnapCenter job log, you can see the Secondary Update step, which initiates a SnapVault update operation. Replication changed blocks from the source volume to the target volume.

Job Details

Backup of Resource Group 'hana-1\_hana\_MDC\_PFX' with policy 'LocalSnapAndSnapVault'

- ✓ ▾ Backup of Resource Group 'hana-1\_hana\_MDC\_PFX' with policy 'LocalSnapAndSnapVault'
- ✓ ▾ hana-1
  - ✓ ▾ Backup
    - ▶ Validate Dataset Parameters
    - ▶ Validate Plugin Parameters
    - ▶ Complete Application Discovery
    - ▶ Initialize Filesystem Plugin
    - ▶ Discover Filesystem Resources
    - ▶ Validate Retention Settings
    - ▶ Quiesce Application
    - ▶ Quiesce Filesystem
    - ▶ Create Snapshot
    - ▶ UnQuiesce Filesystem
    - ▶ UnQuiesce Application
    - ▶ Get Snapshot Details
    - ▶ Get Filesystem Meta Data
    - ▶ Finalize Filesystem Plugin
    - ▶ Collect Autosupport data
    - ▶ Secondary Update
    - ▶ Register Backup and Apply Retention
    - ▶ Register Snapshot attributes
    - ▶ Application Clean-Up
    - ▶ Data Collection
    - ▶ Agent Finalize Workflow
  - ✓ ▾ ( Job 49 ) SnapVault update

Task Name: Secondary Update Start Time: 02/24/2022 3:14:37 PM End Time: 02/24/2022 3:14:46 PM

View Logs Cancel Job Close

On the FSx for ONTAP file system, a Snapshot on the source volume is created using the SnapMirror label,

snapcenter, as configured in the SnapCenter policy.

```
FsxId00fa9e3c784b6abbb::> snapshot show -vserver sapcc-hana-svm -volume
PFX_data_mnt00001 -fields snapmirror-label
vserver          volume          snapshot
snapmirror-label
-----
-----
-----
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_03-31-
2022_13.10.26.5482 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_03-31-
2022_14.00.05.2023 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-05-
2022_08.00.06.3380 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-05-
2022_14.00.01.6482 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-14-
2022_20.00.05.0316 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-28-
2022_08.00.06.3629 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-28-
2022_14.00.01.7275 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-
1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853

snapcenter
8 entries were displayed.
```

At the target volume, a Snapshot copy with the same name is created.

```
FsxId05f7f00af49dc7a3e::> snapshot show -vserver sapcc-backup-target-zone5
-volume PFX_data_mnt00001 -fields snapmirror-label
vserver          volume          snapshot
snapmirror-label
-----
-----
-----
sapcc-backup-target-zone5 PFX_data_mnt00001 SnapCenter_hana-
1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853 snapcenter
FsxId05f7f00af49dc7a3e::>
```

The new Snapshot backup is also listed in the HANA backup catalog.

### Backup Catalog

Database: SYSTEMDB

Show Log Backups  Show Delta Backups

Status	Started	Duration	Size	Backup Type	Destination Ty...
✓	Apr 28, 2022, 4:22:06 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
✓	Apr 28, 2022, 2:00:26 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
✓	Apr 28, 2022, 8:00:35 AM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
✓	Apr 15, 2022, 5:00:44 PM	00h 06m 59s	5.50 GB	Data Backup	Snapshot
✓	Apr 14, 2022, 8:00:32 PM	00h 00m 16s	5.50 GB	Data Backup	Snapshot
✓	Apr 5, 2022, 2:00:29 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
✓	Apr 5, 2022, 8:00:39 AM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
✓	Mar 31, 2022, 2:00:29 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
✓	Mar 31, 2022, 1:10:57 PM	00h 00m 16s	5.50 GB	Data Backup	Snapshot
✓	Feb 22, 2022, 12:55:21 PM	00h 00m 21s	3.56 GB	Data Backup	File

### Backup Details

ID: 1651162926424

Status: Successful

Backup Type: Data Backup

Destination Type: Snapshot

Started: Apr 28, 2022, 4:22:06 PM (UTC)

Finished: Apr 28, 2022, 4:22:21 PM (UTC)

Duration: 00h 00m 15s

Size: 5.50 GB

Throughput: n.a.

System ID:

Comment: SnapCenter\_hana-1\_LocalSnapAndSnapVault\_Daily\_04-28-2022\_16.21.41.5853

Additional Information: <ok>

Location: /hana/data/PFX/mnt00001/

Host	Service	Size	Name	Source Type	EBID
hana-1	nameserver	5.50 GB	hdb00001	volume	SnapCent...

In SnapCenter, you can list the replicated backups by clicking Vault Copies in the topology view.

**Summary Card**

- 10 Backups
- 9 Snapshot based backups
- 1 File-based backup
- 0 Clones

Backup Name	Count	IF	End Date
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853	1		04/28/2022 4:22:40 PM

## Restore and recover from secondary storage

To restore and recover from secondary storage, follow these steps:

To retrieve the list of all the backups on the secondary storage, in the SnapCenter Topology view, click Vault Copies, then select a backup and click Restore.

**Summary Card**

- 10 Backups
- 9 Snapshot based backups
- 1 File-based backup
- 0 Clones

Backup Name	Count	IF	End Date
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853	1		04/28/2022 4:22:40 PM

The restore dialog shows the secondary locations.

Restore from SnapCenter\_hana-1\_LocalSnapAndSnapVault\_Daily\_04-28-2022\_16.21.41.5853 ×

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Select the restore types

Complete Resource ?

Tenant Database

Choose archive location

sapcc-hana-svm:PFX\_data\_mnt00001 sapcc-backup-target-zone5:PFX\_data\_mnt00

⚠ The newer tenants added on the host after the backup was created cannot be restored and will be lost after restore operation. ×

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

Previous Next

Further restore and recovery steps are identical to those previously covered for a Snapshot backup at the primary storage.

## Where to find additional information

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

- FSx for NetApp ONTAP user guide — What is Amazon FSx for NetApp ONTAP?

<https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/what-is-fsx-ontap.html>

- SnapCenter resources page

<https://www.netapp.com/us/documentation/snapcenter-software.aspx>

- SnapCenter Software documentation

<https://docs.netapp.com/us-en/snapcenter/index.html>

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

[Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)

- TR-4719: SAP HANA System Replication — Backup and Recovery with SnapCenter

[Backup and Recovery with SnapCenter](#)

## Version history

Version	Date	Document version history
Version 1.0	May 2022	Initial release.

# SAP HANA backup and recovery with SnapCenter

## TR-4614: SAP HANA backup and recovery with SnapCenter

Companies today require continuous, uninterrupted availability for their SAP applications. They expect consistent performance levels in the face of ever-increasing volumes of data and the need for routine maintenance tasks such as system backups. Performing backups of SAP databases is a critical task and can have a significant performance effect on the production SAP system.

Author: Nils Bauer, NetApp

Backup windows are shrinking, while the amount of data to be backed up is increasing. Therefore, it is difficult to find a time when backups can be performed with minimal effect on business processes. The time needed to restore and recover SAP systems is a concern, because downtime for SAP production and nonproduction systems must be minimized to reduce data loss and cost to the business.

The following points summarize the challenges facing SAP backup and recovery:

- **Performance effects on production SAP systems.** Typically, traditional copy-based backups create a significant performance drain on production SAP systems because of the heavy loads placed on the database server, the storage system, and the storage network.
- **Shrinking backup windows.** Conventional backups can only be made when few dialog or batch activities are in process on the SAP system. The scheduling of backups becomes more difficult when SAP systems are in use around the clock.
- **Rapid data growth.** Rapid data growth and shrinking backup windows require ongoing investment in backup infrastructure. In other words, you must procure more tape drives, additional backup disk space, and faster backup networks. You must also cover the ongoing expense of storing and managing these tape assets. Incremental or differential backups can address these issues, but this arrangement results in a very slow, cumbersome, and complex restore process that is harder to verify. Such systems usually increase recovery time objective (RTO) and recovery point objective (RPO) times in ways that are not acceptable to

the business.

- **Increasing cost of downtime.** Unplanned downtime of an SAP system typically affects business finances. A significant part of any unplanned downtime is consumed by the requirement to restore and recover the SAP system. Therefore, the desired RTO dictates the design of the backup and recovery architecture.
- **Backup and recovery time for SAP upgrade projects.** The project plan for an SAP upgrade includes at least three backups of the SAP database. These backups significantly reduce the time available for the upgrade process. The decision to proceed is generally based on the amount of time required to restore and recover the database from the previously created backup. Rather than just restoring a system to its previous state, a rapid restore provides more time to solve problems that might occur during an upgrade.

## The NetApp solution

NetApp Snapshot technology can be used to create database backups in minutes. The time needed to create a Snapshot copy is independent of the size of the database because a Snapshot copy does not move any physical data blocks on the storage platform. In addition, the use of Snapshot technology has no performance effect on the live SAP system because the NetApp Snapshot technology does not move or copy data blocks when the Snapshot copy is created or when data in the active file system is changed. Therefore, the creation of Snapshot copies can be scheduled without considering peak dialog or batch activity periods. SAP and NetApp customers typically schedule multiple online Snapshot backups during the day; for example, every four hours is common. These Snapshot backups are typically kept for three to five days on the primary storage system before being removed.

Snapshot copies also provide key advantages for restore and recovery operations. NetApp SnapRestore data recovery software enables the restore of an entire database or, alternatively, a portion of a database to any point in time, based on the available Snapshot copies. Such restore processes are finished in a few minutes, independent of the size of the database. Because several online Snapshot backups are created during the day, the time needed for the recovery process is significantly reduced relative to a traditional backup approach. Because a restore can be performed with a Snapshot copy that is only a few hours old (rather than up to 24 hours), fewer transaction logs must be applied. Therefore, the RTO is reduced to several minutes rather than the several hours required for conventional single-cycle tape backups.

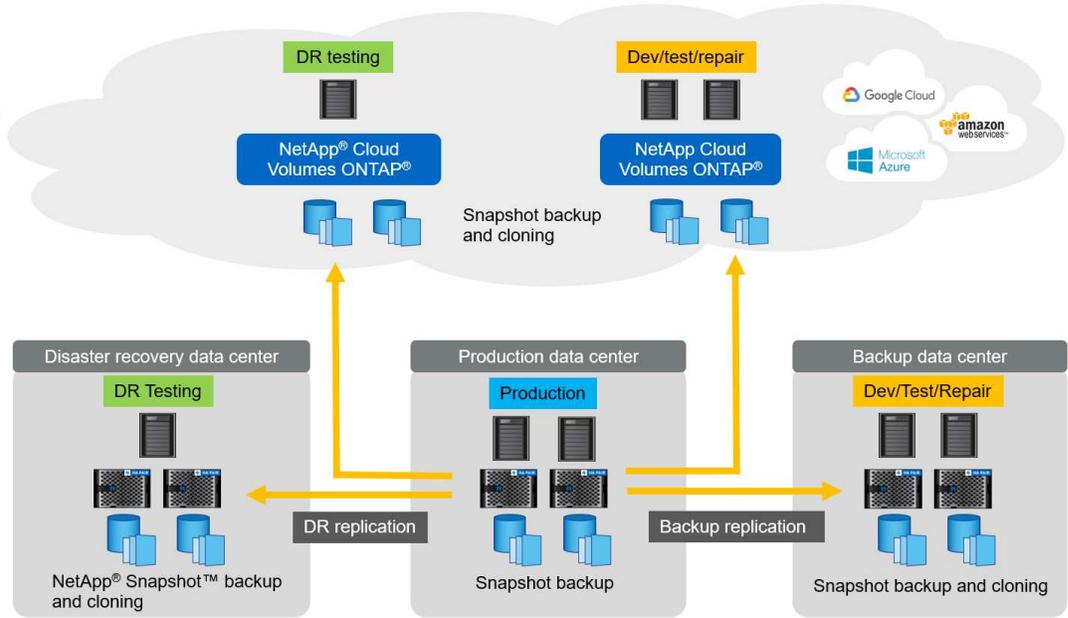
Snapshot copy backups are stored on the same disk system as the active online data. Therefore, NetApp recommends using Snapshot copy backups as a supplement rather than a replacement for backups to a secondary location. Most restore and recovery actions are handled by using SnapRestore on the primary storage system. Restores from a secondary location are only necessary if the primary storage system containing the Snapshot copies is damaged. The secondary location can also be used if it is necessary to restore a backup that is no longer available from a Snapshot copy: a month-end backup, for example.

A backup to a secondary location is based on Snapshot copies created on the primary storage. Therefore, the data is read directly from the primary storage system without generating load on the SAP database server. The primary storage communicates directly with the secondary storage and sends the backup data to the destination by using a NetApp SnapVault disk-to-disk backup.

SnapVault offers significant advantages when compared to traditional backups. After an initial data transfer, in which all data has been transferred from the source to the destination, all subsequent backups copy only the changed blocks to the secondary storage. Therefore, the load on the primary storage system and the time needed for a full backup are significantly reduced. Because SnapVault stores only the changed blocks at the destination, a full database backup requires less disk space.

The solution can also be seamlessly extended to a hybrid cloud operation model. Data replication for disaster recovery or offsite backup purposes can be done from on-premises NetApp ONTAP systems to Cloud Volumes ONTAP instances running in the cloud. You can use SnapCenter as a central tool to manage the data protection and data replication, independent if the SAP HANA system run on-premises or in the cloud. The following figure shows an overview of the backup solution.

Central data protection management for SAP HANA running on-premise or in the cloud



### Runtime of Snapshot backups

The next screenshot shows a customer's HANA Studio running SAP HANA on NetApp storage. The customer is using Snapshot copies to back up the HANA database. The image shows that the HANA database (approximately 2.3TB in size) is backed up in 2 minutes and 11 seconds by using Snapshot backup technology.



The largest part of the overall backup workflow runtime is the time needed to execute the HANA backup savepoint operation, and this step is dependent on the load on the HANA database. The storage Snapshot backup itself always finishes in a couple of seconds.

Status	Started	Duration	Size	Backup Type	Destination
Success	Jun 28, 2017 6:19:11	00h 02m 11s	2.30 TB	Snapshot	...

**Backup Details**

ID: 1498623551457

Status: Successful

Backup Type: Data Backup

Destination Type: Snapshot

Started: Jun 28, 2017 6:19:11 AM (Europe/Berlin)

Finished: Jun 28, 2017 6:21:22 AM (Europe/Berlin)

**Duration: 00h 02m 11s**

**Size: 2.30 TB**

Throughput: n.a.

System ID: SC-PROD\_0100\_20170628061902

Comment: SC-PROD\_0100\_20170628061902

### Recovery time objective comparison

This section provides an RTO comparison of file-based and storage-based Snapshot backups. The RTO is defined by the sum of the time needed to restore the database and the time needed to start and recover the database.

#### Time needed to restore database

With a file-based backup, the restore time depends on the size of the database and backup infrastructure, which defines the restore speed in megabytes per second. For example, if the infrastructure supports a restore

operation at a speed of 250MBps, it takes approximately 1 hour and 10 minutes to restore a database 1TB in size.

With storage Snapshot copy backups, the restore time is independent of the size of the database and is in the range of a couple of seconds when the restore can be performed from primary storage. A restore from secondary storage is only required in the case of a disaster when the primary storage is no longer available.

#### Time needed to start database

The database start time depends on the size of the row and column store. For the column store, the start time also depends on how much data is preloaded during the database start. In the following examples, we assume that the start time is 30 minutes. The start time is the same for a file-based restore and recovery and a restore and recovery based on Snapshot.

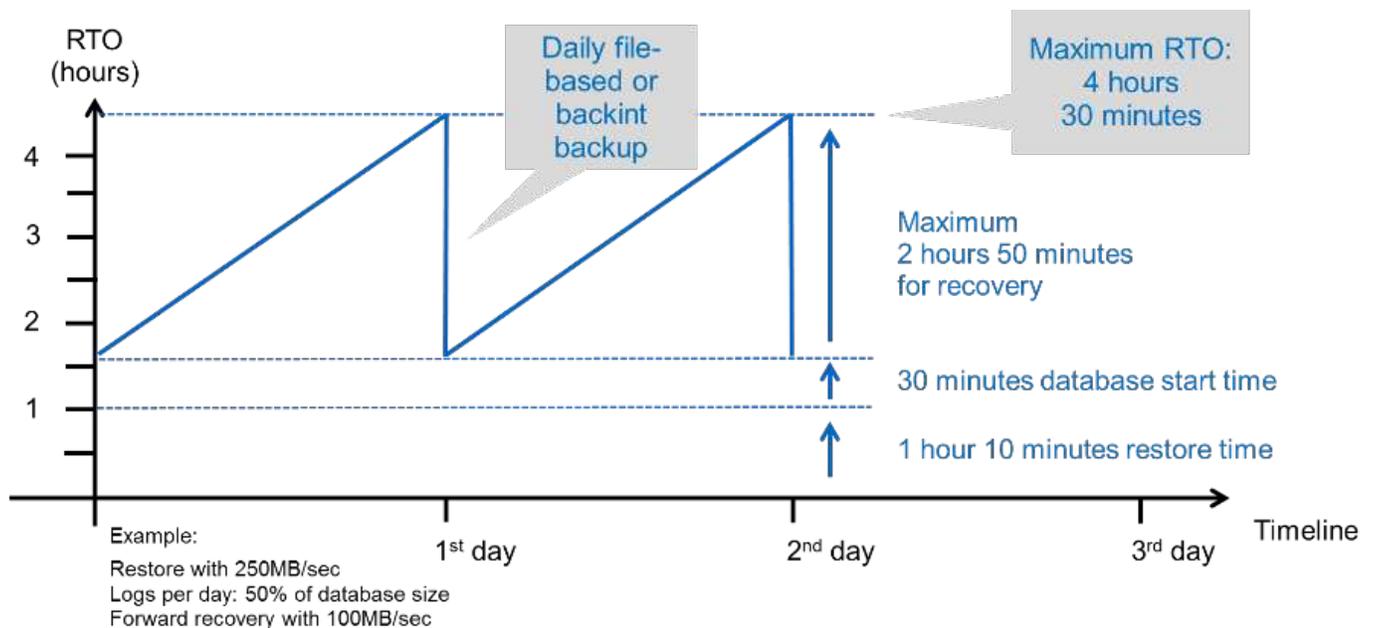
#### Time needed to recover database

The recovery time depends on the number of logs that must be applied after the restore. This number is determined by the frequency at which data backups are taken.

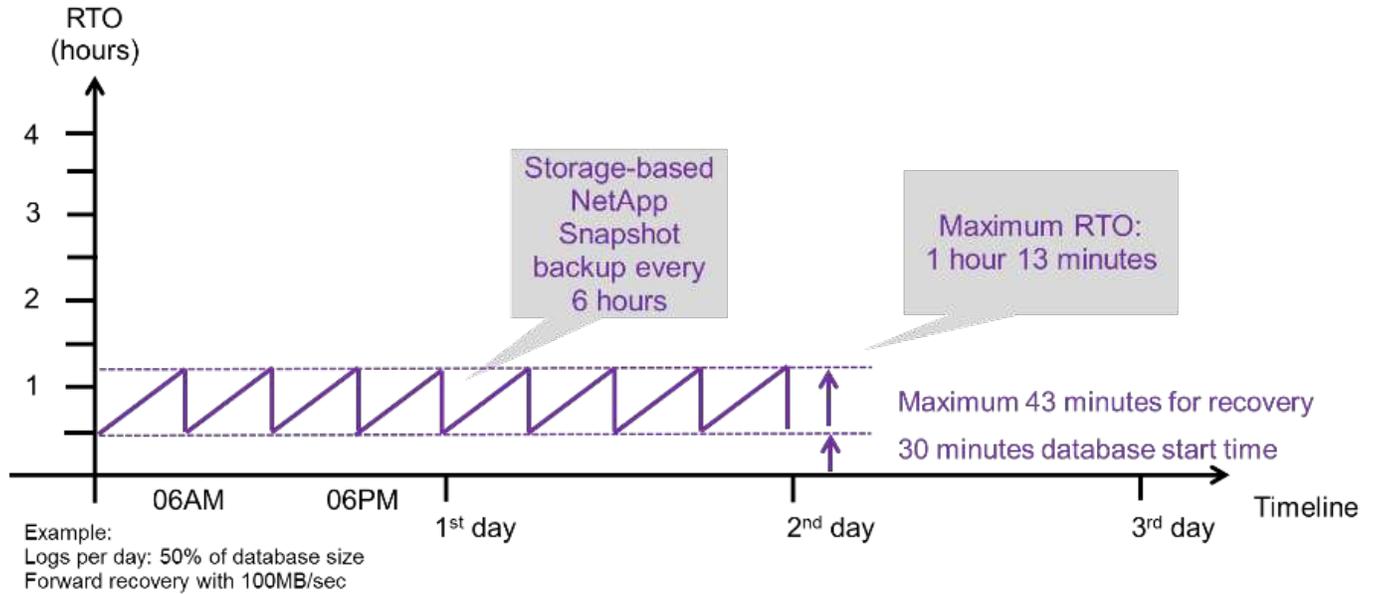
With file-based data backups, the backup schedule is typically once per day. A higher backup frequency is normally not possible, because the backup degrades production performance. Therefore, in the worst case, all the logs that were written during the day must be applied during forward recovery.

Storage Snapshot copy data backups are typically scheduled with a higher frequency because they do not influence the performance of the SAP HANA database. For example, if Snapshot copy backups are scheduled every six hours, the recovery time would be, in the worst case, one-fourth of the recovery time for a file-based backup (6 hours / 24 hours =  $\frac{1}{4}$ ).

The following figure shows an RTO example for a 1TB database when file-based data backups are used. In this example, a backup is taken once per day. The RTO differs depending on when the restore and recovery were performed. If the restore and recovery were performed immediately after a backup was taken, the RTO is primarily based on the restore time, which is 1 hour and 10 minutes in the example. The recovery time increased to 2 hours and 50 minutes when restore and recovery were performed immediately before the next backup was taken, and the maximum RTO was 4 hours and 30 minutes.

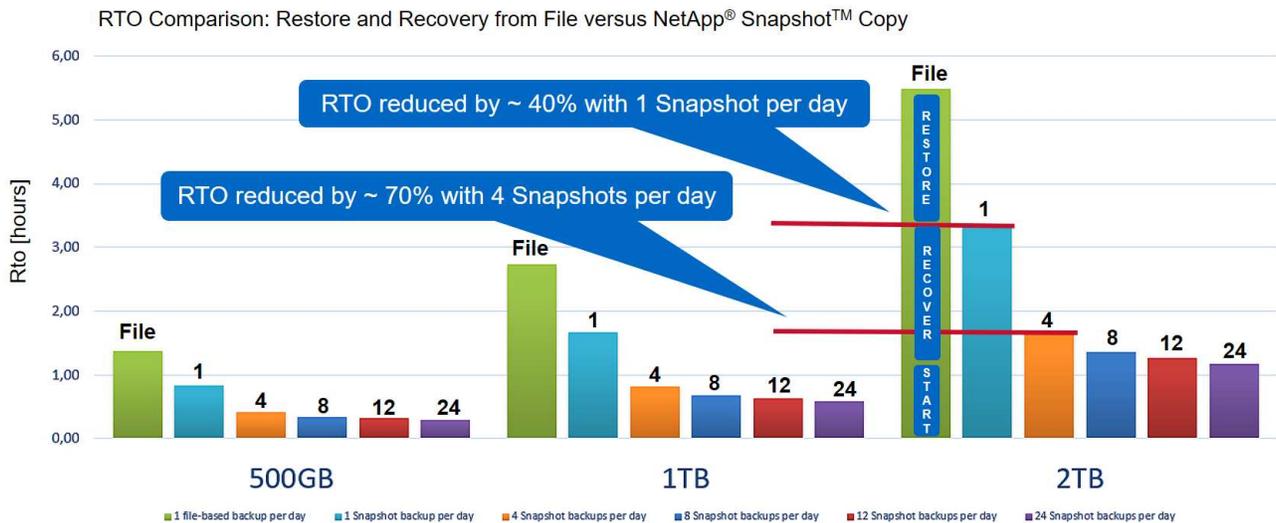


The following figure shows an RTO example for a 1TB database when Snapshot backups are used. With storage-based Snapshot backups, the RTO only depends on the database start time and the forward recovery time because the restore is completed in a few seconds, independent of the size of the database. The forward recovery time also increases depending on when the restore and recovery are done, but due to the higher frequency of backups (every six hours in this example), the forward recovery time is 43 minutes at most. In this example, the maximum RTO is 1 hour and 13 minutes.



The following figure shows an RTO comparison of file-based and storage-based Snapshot backups for different database sizes and different frequencies of Snapshot backups. The green bar shows the file-based backup. The other bars show Snapshot copy backups with different backup frequencies.

With a single Snapshot copy data backup per day, the RTO is already reduced by 40% when compared to a file-based data backup. The reduction increases to 70% when four Snapshot backups are taken per day. The figure also shows that the curve goes flat if you increase the Snapshot backup frequency to more than four to six Snapshot backups per day. Our customers therefore typically configure four to six Snapshot backups per day.





The graph shows the HANA server RAM size. The database size in memory is calculated to be half of the server RAM size.



The restore and recovery time is calculated based on the following assumptions. The database can be restored at 250MBps. The number of log files per day is 50% of the database size. For example, a 1TB database creates 500MB of log files per day. A recovery can be performed at 100MBps.

## SnapCenter architecture

SnapCenter is a unified, scalable platform for application-consistent data protection. SnapCenter provides centralized control and oversight, while delegating the ability for users to manage application-specific backup, restore, and clone jobs. With SnapCenter, database and storage administrators learn a single tool to manage backup, restore, and cloning operations for a variety of applications and databases.

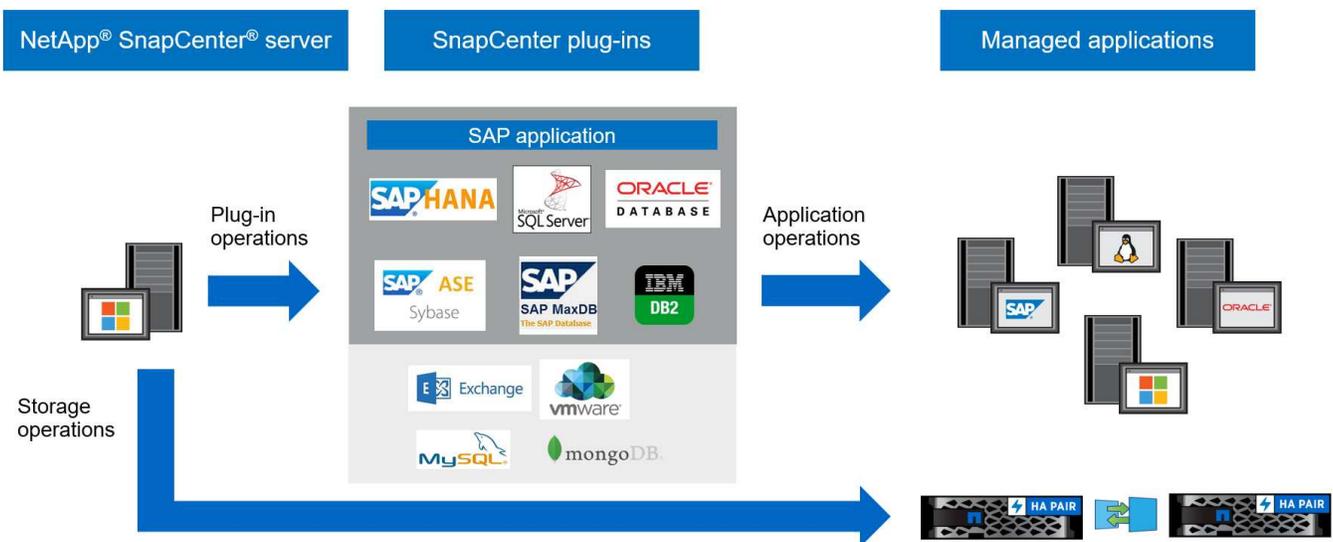
SnapCenter manages data across endpoints in the data fabric powered by NetApp. You can use SnapCenter to replicate data between on-premises environments; between on-premises environments and the cloud; and between private, hybrid, or public clouds.

## SnapCenter components

SnapCenter includes the SnapCenter Server, the SnapCenter Plug-In Package for Windows, and the SnapCenter Plug-Ins Package for Linux. Each package contains plug-ins to SnapCenter for various applications and infrastructure components.

The SnapCenter custom plug-ins enable you to create your own plug-ins and protect your application using the same SnapCenter interface.

The following figure depicts SnapCenter components.



## SnapCenter SAP HANA backup solution

This section lists the components, supported SAP HANA releases and configurations,

and SnapCenter 4.6 enhancements used in this solution.

## Solution components

The SnapCenter backup solution for SAP HANA covers the following areas:

- SAP HANA data backup with storage-based Snapshot copies:
  - Backup scheduling
  - Retention management
  - Housekeeping of the SAP HANA backup catalog
- Non-data volume (for example, `/hana/shared`) backup with storage-based Snapshot copies:
  - Backup scheduling
  - Retention management
- Replication to an off-site backup or disaster recovery location:
  - SAP HANA data Snapshot backups
  - Non-data volumes
  - Retention management configured at off-site backup storage
  - Housekeeping of the SAP HANA backup catalog
- Database block integrity checks using a file-based backup:
  - Backup scheduling
  - Retention management
  - Housekeeping of the SAP HANA backup catalog
- Retention management of HANA database log backup:
  - Retention management based on data backup retention
  - Housekeeping of the SAP HANA backup catalog
- Automatic discovery of HANA databases
- Automated restore and recovery
- Single-tenant restore operations with SAP HANA multitenant database container (MDC) systems

Database data file backups are executed by SnapCenter in combination with the plug-in for SAP HANA. The plug-in triggers an SAP HANA database backup save point so that the Snapshot copies, which are created on the primary storage system, are based on a consistent image of the SAP HANA database.

SnapCenter enables the replication of consistent database images to an off-site backup or disaster recovery location by using SnapVault or the NetApp SnapMirror. feature. Typically, different retention policies are defined for backups at primary and at the off-site backup storage. SnapCenter handles the retention at primary storage, and ONTAP handles the retention at the off-site backup storage.

To allow a complete backup of all SAP HANA-related resources, SnapCenter also allows you to back up all non- data volumes using the SAP HANA plug-in with storage-based Snapshot copies. Non-data volumes can be scheduled independently from the database data backup to enable individual retention and protection policies.

The SAP HANA database automatically executes log backups. Depending on the recovery point objectives, there are several options for the storage location of the log backups:

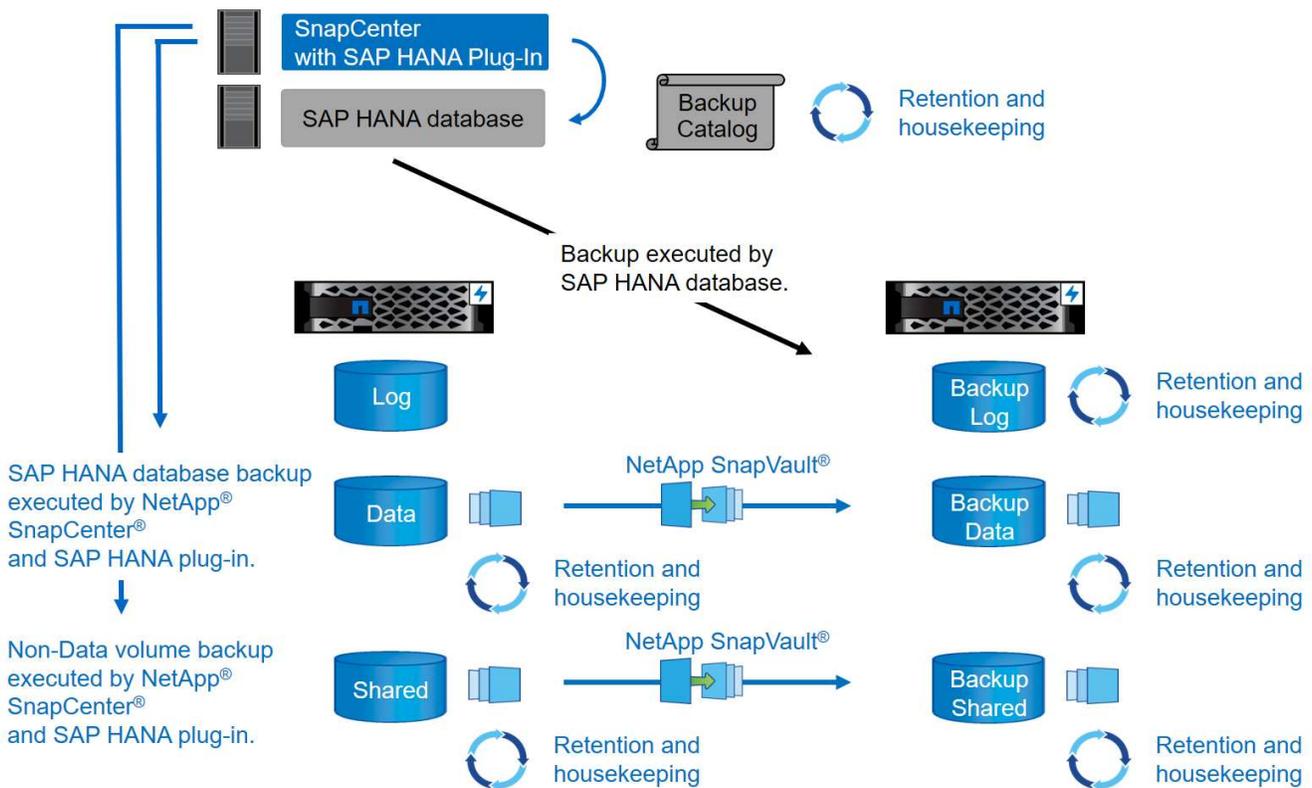
- The log backup is written to a storage system that synchronously mirrors the data to a second location with NetApp MetroCluster high-availability (HA) and disaster recovery storage software.
- The log backup destination can be configured on the same primary storage system and then replicated synchronously or asynchronously to a secondary storage with SnapMirror.
- The log backup destination can be configured on the same off-site backup storage in which the database backups are replicated with SnapVault. With this configuration, the off-site backup storage has availability requirements like those of the primary storage so that log backups can be written to the off-site backup storage.

SAP recommends combining storage-based Snapshot backups with a weekly file-based backup to execute a block integrity check. The block integrity check can be executed from within SnapCenter. Based on your configurable retention policies, SnapCenter manages the housekeeping of data file backups at the primary storage, log file backups, and the SAP HANA backup catalog.



SnapCenter handles the retention at primary storage, while ONTAP manages secondary backup retention.

The following figure shows an overview of the database and log backup configuration, where the log backups are written to an NFS mount of the off-site backup storage.



When executing a storage-based Snapshot backup of non-data volumes, SnapCenter performs the following tasks:

1. Creation of a storage Snapshot copy of the non-data volume.
2. Execution of a SnapVault or SnapMirror update for the data volume, if configured.
3. Deletion of storage Snapshot copies at the primary storage based on the defined retention policy.

When executing a storage-based Snapshot backup of the SAP HANA database, SnapCenter performs the

following tasks:

1. Creation of an SAP HANA backup save point to create a consistent image on the persistence layer.
2. Creation of a storage Snapshot copy of the data volume.
3. Registration of the storage Snapshot back up in the SAP HANA backup catalog.
4. Release of the SAP HANA backup save point.
5. Execution of a SnapVault or SnapMirror update for the data volume, if configured.
6. Deletion of storage Snapshot copies at the primary storage based on the defined retention policy.
7. Deletion of SAP HANA backup catalog entries if the backups do not exist anymore at the primary or off-site backup storage.
8. Whenever a backup has been deleted based on the retention policy or manually, SnapCenter deletes all log backups that are older than the oldest data backup. Log backups are deleted on the file system and in the SAP HANA backup catalog.

### **Supported SAP HANA releases and configurations**

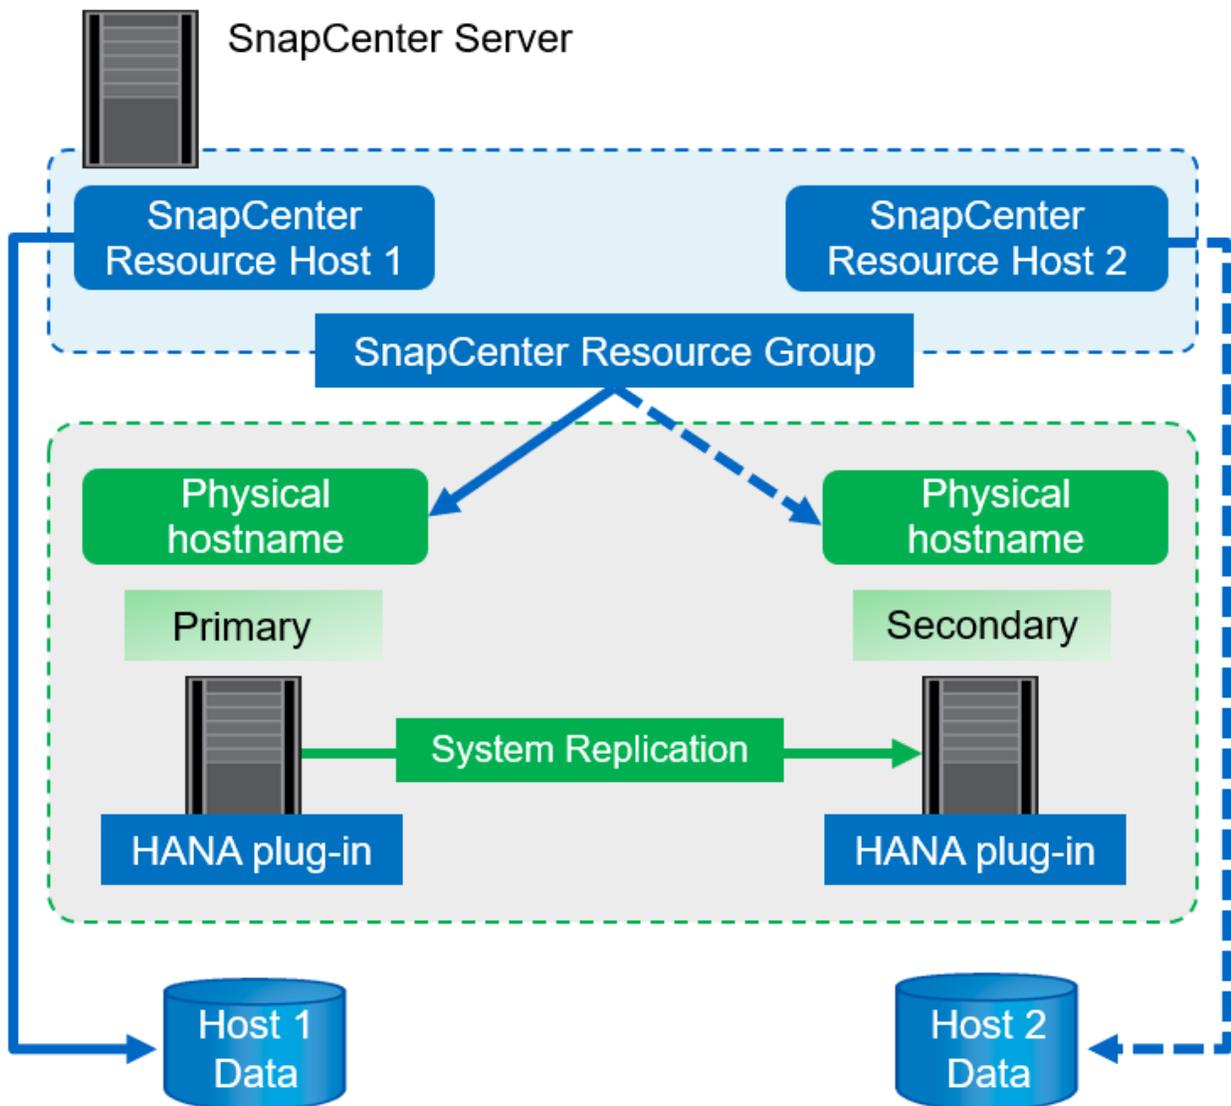
SnapCenter supports SAP HANA single-host and multiple-host configurations using NFS- or FC-attached NetApp storage systems (AFF and FAS), as well as SAP HANA systems running on Cloud Volumes ONTAP at AWS, Azure, the Google Cloud Platform, and AWS FSx ONTAP using NFS.

SnapCenter supports the following SAP HANA architectures and releases:

- SAP HANA single container: SAP HANA 1.0 SPS12
- SAP HANA multitenant-database container (MDC) single tenant: SAP HANA 2.0 SPS3 and later
- SAP HANA multitenant-database container (MDC) multiple tenants: SAP HANA 2.0 SPS4 and later

### **SnapCenter 4.6 enhancements**

Starting with version 4.6, SnapCenter supports auto-discovery of HANA systems configured in a HANA System Replication relationship. Each host is configured using its physical IP address (host name) and its individual data volume on the storage layer. The two SnapCenter resources are combined in a resource group, SnapCenter automatically identifies which host is primary or secondary, and it then executes the required backup operations accordingly. Retention management for Snapshot and file-based backups created with SnapCenter is performed across both hosts to ensure that old backups are also deleted at the current secondary host. The following figure shows a high-level overview. A detailed description of the configuration and operation of HANA System Replication-enabled HANA systems in SnapCenter can be found in [TR-4719 SAP HANA System Replication, Backup and Recovery with SnapCenter](#).



## SnapCenter concepts and best practices

This section describes SnapCenter concepts and best practices as they relate to SAP HANA resource configuration and deployment.

### SAP HANA resource configuration options and concepts

With SnapCenter, SAP HANA database resource configuration can be performed with two different approaches.

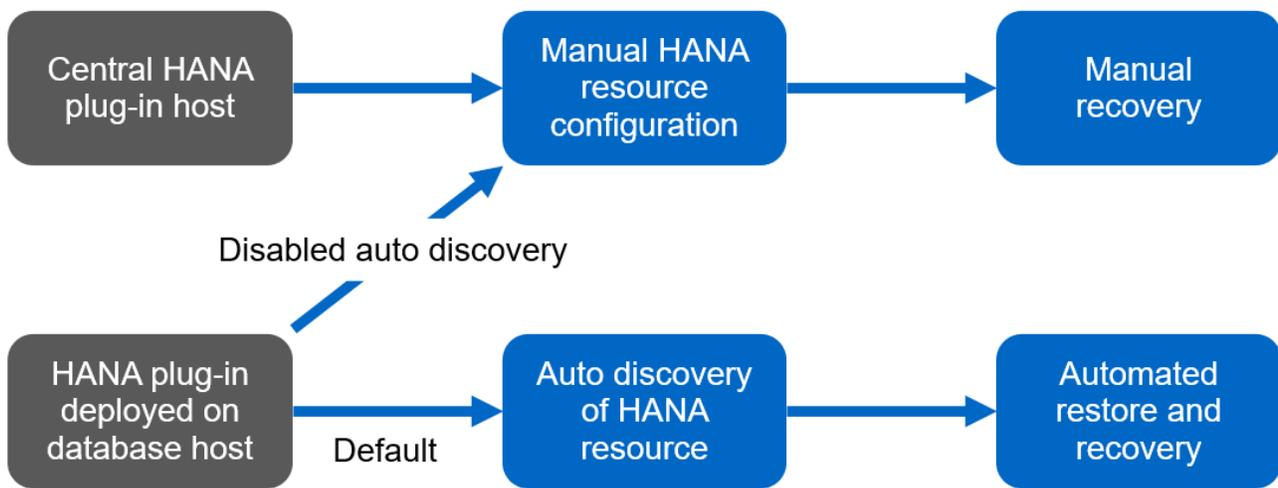
- **Manual resource configuration.** HANA resource and storage footprint information must be provided manually.
- **Automatic discovery of HANA resources.** Automatic discovery simplifies the configuration of HANA databases in SnapCenter and enables automated restore and recovery.

It is important to understand that only HANA database resources in SnapCenter that have been automatically discovered are enabled for automated restore and recovery. HANA database resources that are configured manually in SnapCenter must be recovered manually after a restore operation in SnapCenter.

On the other hand, automatic discovery with SnapCenter is not supported for all HANA architectures and infrastructure configurations. Therefore, HANA landscapes might require a mixed approach in which some HANA systems (HANA multiple host systems) require manual resource configuration and all others can be configured using automatic discovery.

Automatic discovery and automated restore and recovery depend on the ability to execute OS commands on the database host. Examples of this are file system and storage footprint discovery, and unmount, mount, or LUN discovery operations. These operations are executed with the SnapCenter Linux plug-in, which is automatically deployed together with the HANA plug-in. Therefore, it is prerequisite to deploy the HANA plug-in on the database host to enable automatic discovery as well as automated restore and recovery. It is also possible to disable the auto discovery after the deployment of the HANA plug-in on the database host. In this instance, the resource will be a manually configured resource.

The following figure summarizes the dependencies. More details on the HANA deployment options are covered in the section “Deployment options for the SAP HANA plug-in.”



**i** The HANA and Linux plug-ins are currently only available for Intel-based systems. If the HANA databases are running on IBM Power Systems, a central HANA plug-in host must be used.

**Supported HANA architectures for automatic discovery and automated recovery**

With SnapCenter, automatic discovery and automated restore and recovery is supported for most HANA configurations with the exception that HANA multiple host systems require a manual configuration.

The following table shows supported HANA configurations for automatic discovery.

HANA plug-in installed on:	HANA architecture	HANA system configuration	Infrastructure
HANA database host	Single host	<ul style="list-style-type: none"> <li>• HANA single container</li> <li>• SAP HANA multitenant database containers (MDC) with single or multiple tenants</li> <li>• HANA System Replication</li> </ul>	<ul style="list-style-type: none"> <li>• Bare metal with NFS</li> <li>• Bare metal with XFS and FC with or without Linux Logical Volume Manager (LVM)</li> <li>• VMware with direct OS NFS mounts</li> </ul>



HANA MDC systems with multiple tenants are supported for automatic discovery, but not for automated restore and recovery with the current SnapCenter release.

### Supported HANA architectures for manual HANA resource configuration

Manual configuration of HANA resources is supported for all HANA architectures; however, it requires a central HANA plug-in host. The central plug-in host can be the SnapCenter server itself or a separate Linux or Windows host.



When the HANA plug-in is deployed on the HANA database host, by default, the resource is auto discovered. Auto discovery can be disabled for individual hosts, so that the plug-in can be deployed; for example, on a database host with activated HANA System Replication and a SnapCenter release < 4.6, where auto discovery is not supported. For more information, see the section [“Disable auto discovery on the HANA plug-in host.”](#)

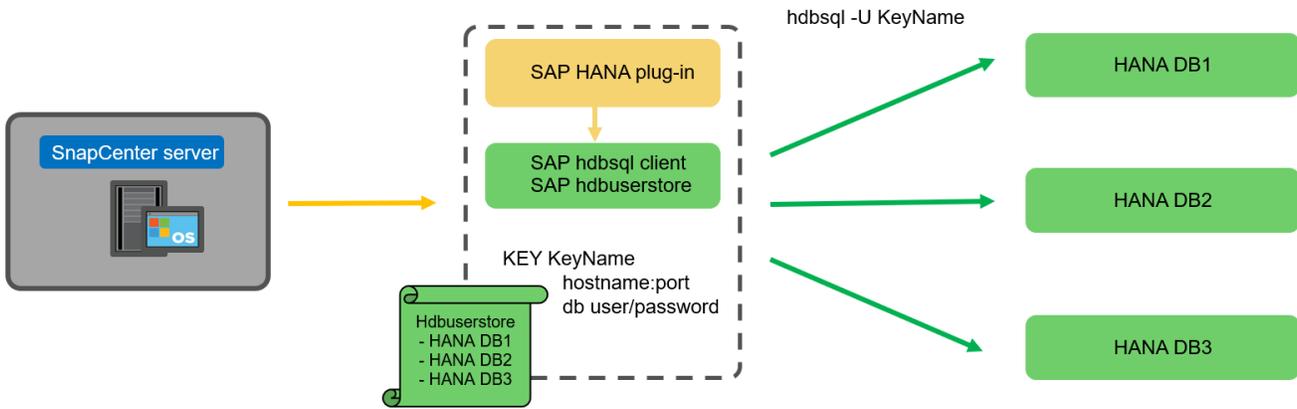
The following table shows supported HANA configurations for manual HANA resource configuration.

HANA Plug-In installed on:	HANA architecture	HANA system configuration	Infrastructure
Central plug-in host (SnapCenter Server or separate Linux host)	Single or multiple host	<ul style="list-style-type: none"> <li>• HANA single container</li> <li>• HANA MDC with single or multiple tenants</li> <li>• HANA System Replication</li> </ul>	<ul style="list-style-type: none"> <li>• Bare metal with NFS</li> <li>• Bare metal with XFS and FC with or without Linux LVM</li> <li>• VMware with direct OS NFS mounts</li> </ul>

### Deployment options for the SAP HANA plug-in

The following figure shows the logical view and the communication between the SnapCenter Server and the SAP HANA databases.

The SnapCenter Server communicates through the SAP HANA plug-in with the SAP HANA databases. The SAP HANA plug-in uses the SAP HANA hdbsql client software to execute SQL commands to the SAP HANA databases. The SAP HANA hdbuserstore is used to provide the user credentials, the host name, and the port information to access the SAP HANA databases.



The SAP HANA plug-in and the SAP hdbsql client software, which include the hdbuserstore configuration tool, must be installed together on the same host.

The host can be the SnapCenter Server itself, a separate central plug-in host, or the individual SAP HANA database hosts.

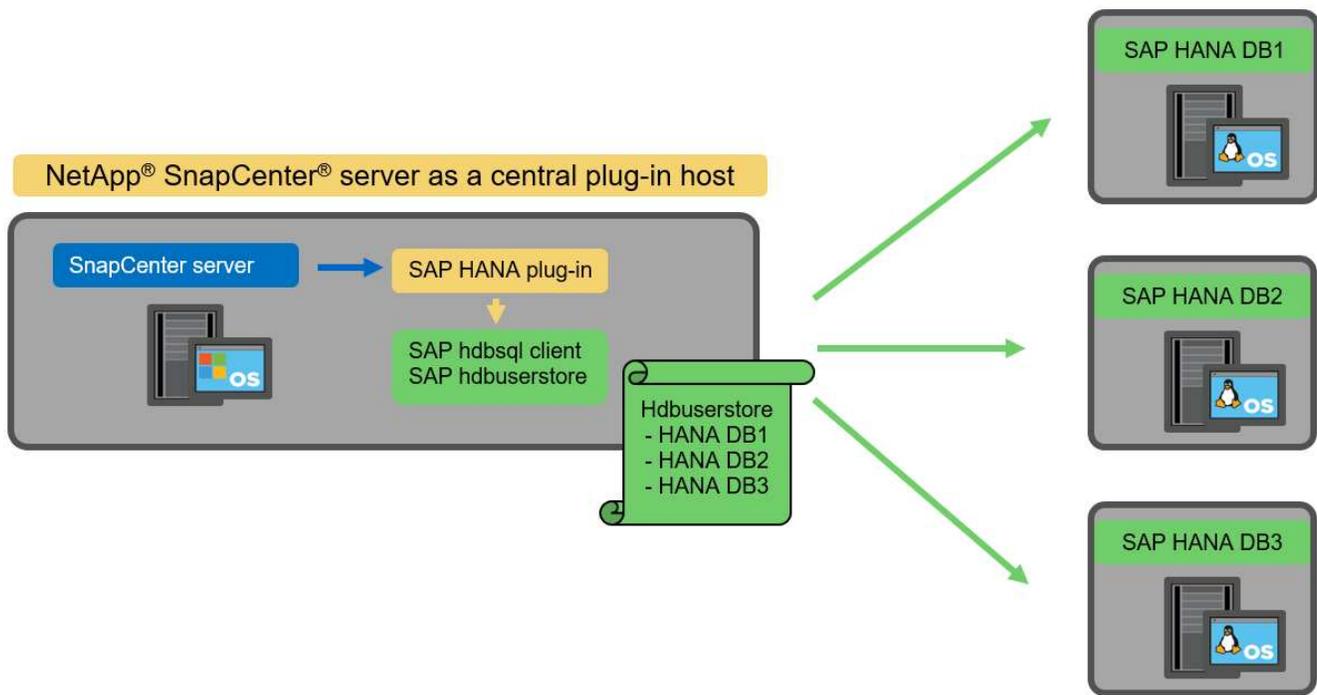
### SnapCenter server high availability

SnapCenter can be set up in a two-node HA configuration. In such a configuration, a load balancer (for example, F5) is used in an active/passive mode using a virtual IP address pointing to the active SnapCenter host. The SnapCenter repository (the MySQL database) is replicated by SnapCenter between the two hosts so that the SnapCenter data is always in-sync.

SnapCenter server HA is not supported if the HANA plug-in is installed on the SnapCenter server. If you plan to set up SnapCenter in an HA configuration, do not install the HANA plug-in on the SnapCenter server. More details on SnapCenter HA can be found at this [NetApp Knowledge Base page](#).

### SnapCenter server as a central HANA plug-in host

The following figure shows a configuration in which the SnapCenter Server is used as a central plug-in host. The SAP HANA plug-in and the SAP hdbsql client software are installed on the SnapCenter Server.



Since the HANA plug-in can communicate with the managed HANA databases using the hdbclient through the network, you do not need to install any SnapCenter components on the individual HANA database hosts. SnapCenter can protect the HANA databases by using a central HANA plug-in host on which all userstore keys are configured for the managed databases.

On the other hand, enhanced workflow automation for automatic discovery, automation of restore and recovery, as well as SAP system refresh operations require SnapCenter components to be installed on the database host. When using a central HANA plug-in host, these features are not available.

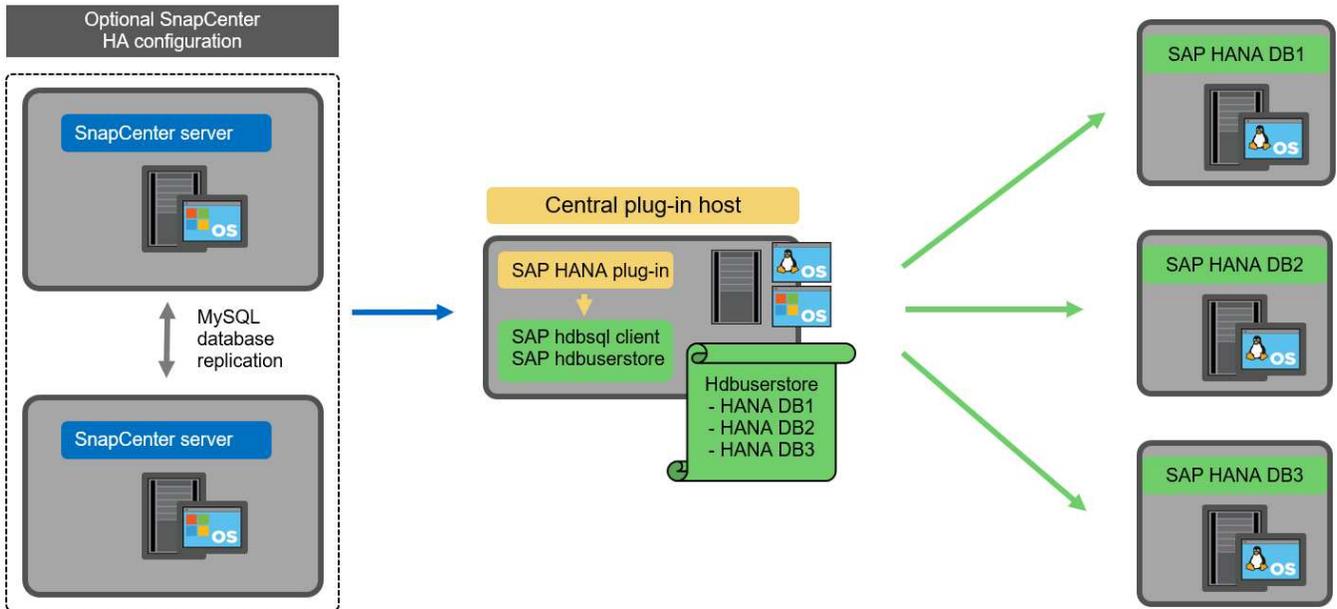
Also, high availability of the SnapCenter server using the in-built HA feature cannot be used when the HANA plug-in is installed on the SnapCenter server. High availability can be achieved using VMware HA if the SnapCenter server is running in a VM within a VMware cluster.

#### Separate host as a central HANA plug-in host

The following figure shows a configuration in which a separate Linux host is used as a central plug-in host. In this case, the SAP HANA plug-in and the SAP hdbsql client software are installed on the Linux host.



The separate central plug-in host can also be a Windows host.

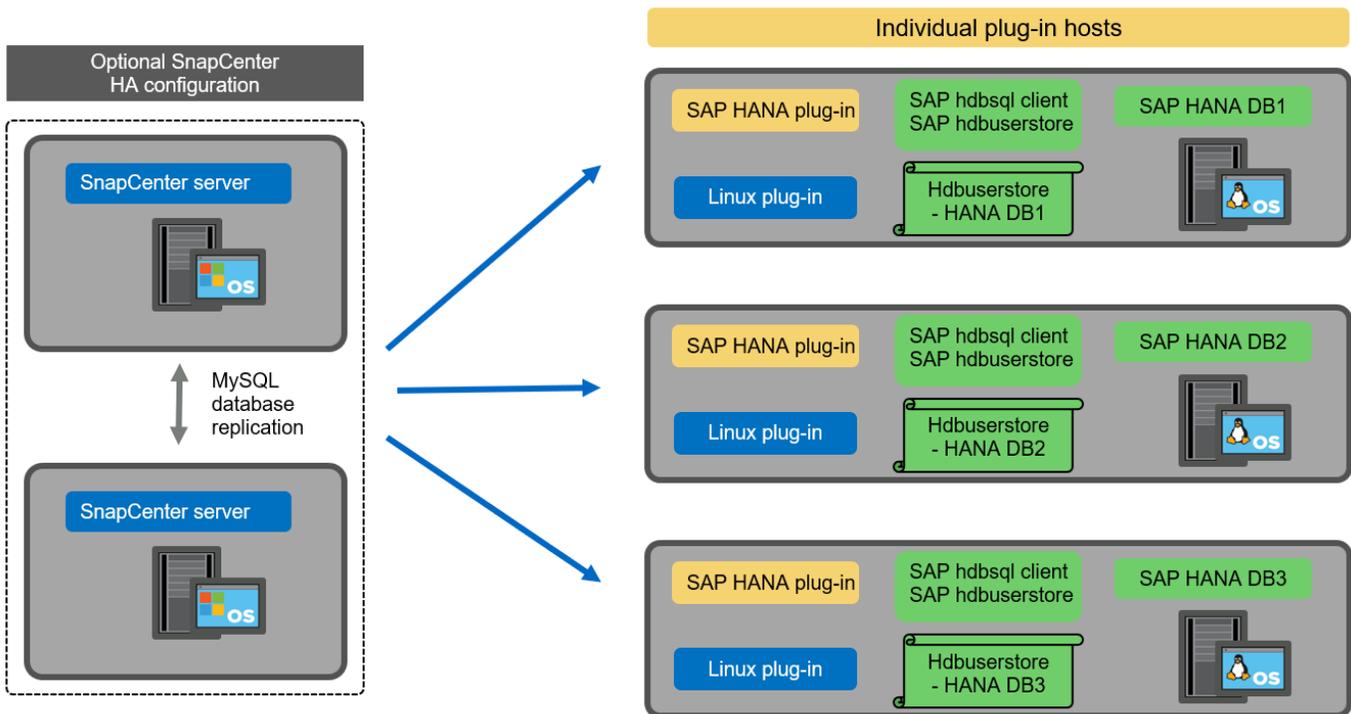


The same restriction regarding feature availability described in the previous section also applies for a separate central plug-in host.

However, with this deployment option the SnapCenter server can be configured with the in-built HA functionality. The central plug-in host must also be HA, for example, by using a Linux cluster solution.

#### HANA plug-in deployed on individual HANA database hosts

The following figure shows a configuration in which the SAP HANA plug-in is installed on each SAP HANA database host.



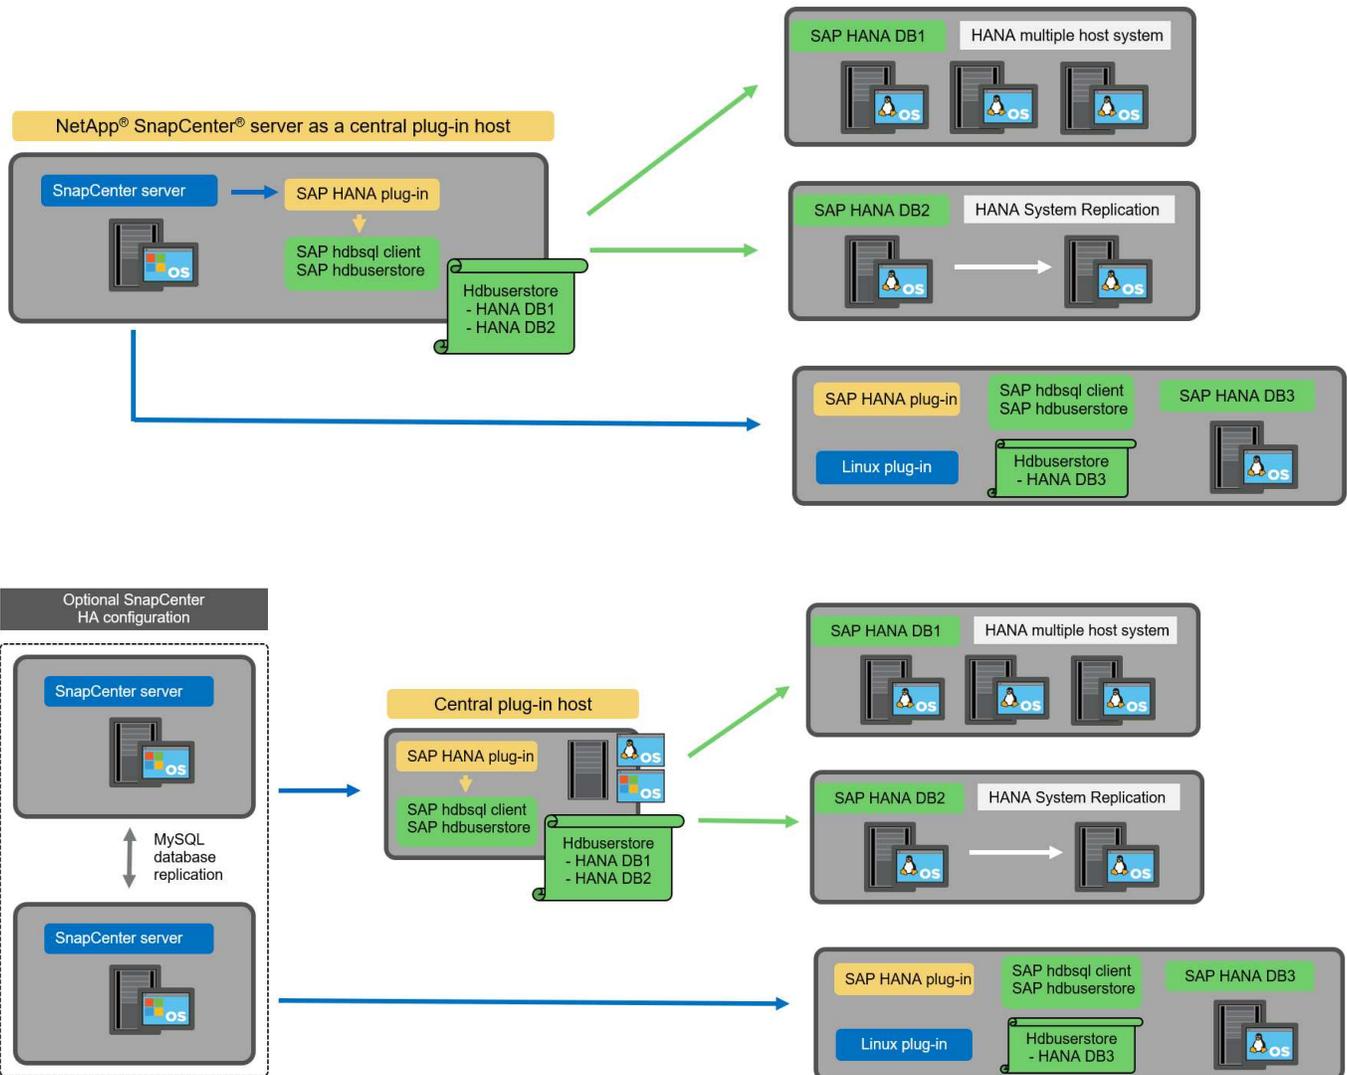
When the HANA plug-in is installed on each individual HANA database host, all features, such as automatic discovery and automated restore and recovery, are available. Also, the SnapCenter server can be set up in an HA configuration.

### Mixed HANA plug-in deployment

As discussed at the beginning of this section, some HANA system configurations, such as multiple-host systems, require a central plug-in host. Therefore, most SnapCenter configurations require a mixed deployment of the HANA plug-in.

NetApp recommends that you deploy the HANA plug-in on the HANA database host for all HANA system configurations that are supported for automatic discovery. Other HANA systems, such as multiple-host configurations, should be managed with a central HANA plug-in host.

The following two figures show mixed plug-in deployments either with the SnapCenter server or a separate Linux host as a central plug-in host. The only difference between these two deployments is the optional HA configuration.



### Summary and recommendations

In general, NetApp recommends that you deploy the HANA plug-in on each SAP HANA host to enable all

available SnapCenter HANA features and to enhance workflow automation.



The HANA and Linux plug-ins are currently only available for Intel- based systems. If the HANA databases are running on IBM Power Systems, a central HANA plug-in host must be used.

For HANA configurations in which automatic discovery is not supported, such as HANA multiple-host configurations, an additional central HANA plug-in host must be configured. The central plug-in host can be the SnapCenter server if VMware HA can be leveraged for SnapCenter HA. If you plan to use the SnapCenter in-build HA capability, use a separate Linux plug-in host.

The following table summarizes the different deployment options.

Deployment option	Dependencies
Central HANA plug-in host Plug-in installed on SnapCenter server	Pros: * Single HANA plug-in, central HDB user store configuration * No SnapCenter software components required on individual HANA database hosts * Support of all HANA architectures Cons: * Manual resource configuration * Manual recovery * No single tenant restore support * Any Pre- and post-script steps are executed on the central plug-in host * In-build SnapCenter high availability not supported * Combination of SID and tenant name must be unique across all managed HANA databases * Log backup retention management enabled/disabled for all managed HANA databases
Central HANA plug-in host Plug-in installed on separate Linux or Windows server	Pros: * Single HANA plug-in, central HDB user store configuration * No SnapCenter software components required on individual HANA database hosts * Support of all HANA architectures * In-build SnapCenter high availability supported Cons: * Manual resource configuration * Manual recovery * No single tenant restore support * Any Pre- and post-script steps are executed on the central plug-in host * Combination of SID and tenant name must be unique across all managed HANA databases * Log backup retention management enabled/disabled for all managed HANA databases

Deployment option	Dependencies
Individual HANA plug-in host Plug-in installed on HANA database server	<p>Pros:</p> <ul style="list-style-type: none"> <li>* Automatic discovery of HANA resources</li> <li>* Automated restore and recovery</li> <li>* Single tenant restore</li> <li>* Pre- and post-script automation for SAP system refresh</li> <li>* In-build SnapCenter high availability supported</li> <li>* Log backup retention management can be enabled/disabled for each individual HANA database</li> </ul> <p>Cons:</p> <ul style="list-style-type: none"> <li>* Not supported for all HANA architectures. Additional central plug-in host required, for HANA multiple host systems.</li> <li>* HANA plug-in must be deployed on each HANA database hosts</li> </ul>

### Data protection strategy

Before configuring SnapCenter and the SAP HANA plug-in, the data protection strategy must be defined based on the RTO and RPO requirements of the various SAP systems.

A common approach is to define system types such as production, development, test, or sandbox systems. All SAP systems of the same system type typically have the same data protection parameters.

The parameters that must be defined are:

- How often should a Snapshot backup be executed?
- How long should Snapshot copy backups be kept on the primary storage system?
- How often should a block integrity check be executed?
- Should the primary backups be replicated to an off-site backup site?
- How long should the backups be kept at the off-site backup storage?

The following table shows an example of data protection parameters for the system type's production, development, and test. For the production system, a high backup frequency has been defined, and the backups are replicated to an off-site backup site once per day. The test systems have lower requirements and no replication of the backups.

Parameters	Production systems	Development systems	Test systems
Backup frequency	Every 4 hours	Every 4 hours	Every 4 hours
Primary retention	2 days	2 days	2 days
Block integrity check	Once per week	Once per week	No
Replication to off-site backup site	Once per day	Once per day	No
Off-site backup retention	2 weeks	2 weeks	Not applicable

The following table shows the policies that must be configured for the data protection parameters.

Parameters	PolicyLocalSnap	PolicyLocalSnapAndSnapVault	PolicyBlockIntegrityCheck
Backup type	Snapshot based	Snapshot based	File based
Schedule frequency	Hourly	Daily	Weekly
Primary retention	Count = 12	Count = 3	Count = 1
SnapVault replication	No	Yes	Not applicable

The policy `LocalSnapshot` is used for the production, development, and test systems to cover the local Snapshot backups with a retention of two days.

In the resource protection configuration, the schedule is defined differently for the system types:

- **Production.** Schedule every 4 hours.
- **Development.** Schedule every 4 hours.
- **Test.** Schedule every 4 hours.

The policy `LocalSnapAndSnapVault` is used for the production and development systems to cover the daily replication to the off-site backup storage.

In the resource protection configuration, the schedule is defined for production and development:

- **Production.** Schedule every day.
- **Development.** Schedule every day.

The policy `BlockIntegrityCheck` is used for the production and development systems to cover the weekly block integrity check using a file-based backup.

In the resource protection configuration, the schedule is defined for production and development:

- **Production.** Schedule every week.
- **Development.** Schedule every week.

For each individual SAP HANA database that uses the off-site backup policy, a protection relationship must be configured on the storage layer. The protection relationship defines which volumes are replicated and the retention of backups at the off-site backup storage.

With our example, for each production and development system, a retention of two weeks is defined at the off-site backup storage.



In our example, protection policies and retention for SAP HANA database resources and non-data volume resources are not different.

## Backup operations

SAP introduced the support of Snapshot backups for MDC multiple tenant systems with HANA 2.0 SPS4. SnapCenter supports Snapshot backup operations of HANA MDC systems with multiple tenants. SnapCenter also supports two different restore operations of a HANA MDC system. You can either restore the complete system, the System DB and all tenants, or you can restore just a single tenant. There are some pre-requisites to enable SnapCenter to execute these operations.

In an MDC System, the tenant configuration is not necessarily static. Tenants can be added or tenants can be deleted. SnapCenter cannot rely on the configuration that is discovered when the HANA database is added to SnapCenter. SnapCenter must know which tenants are available at the point in time the backup operation is executed.

To enable a single tenant restore operation, SnapCenter must know which tenants are included in each Snapshot backup. In addition, it must know which files and directories belong to each tenant included in the Snapshot backup.

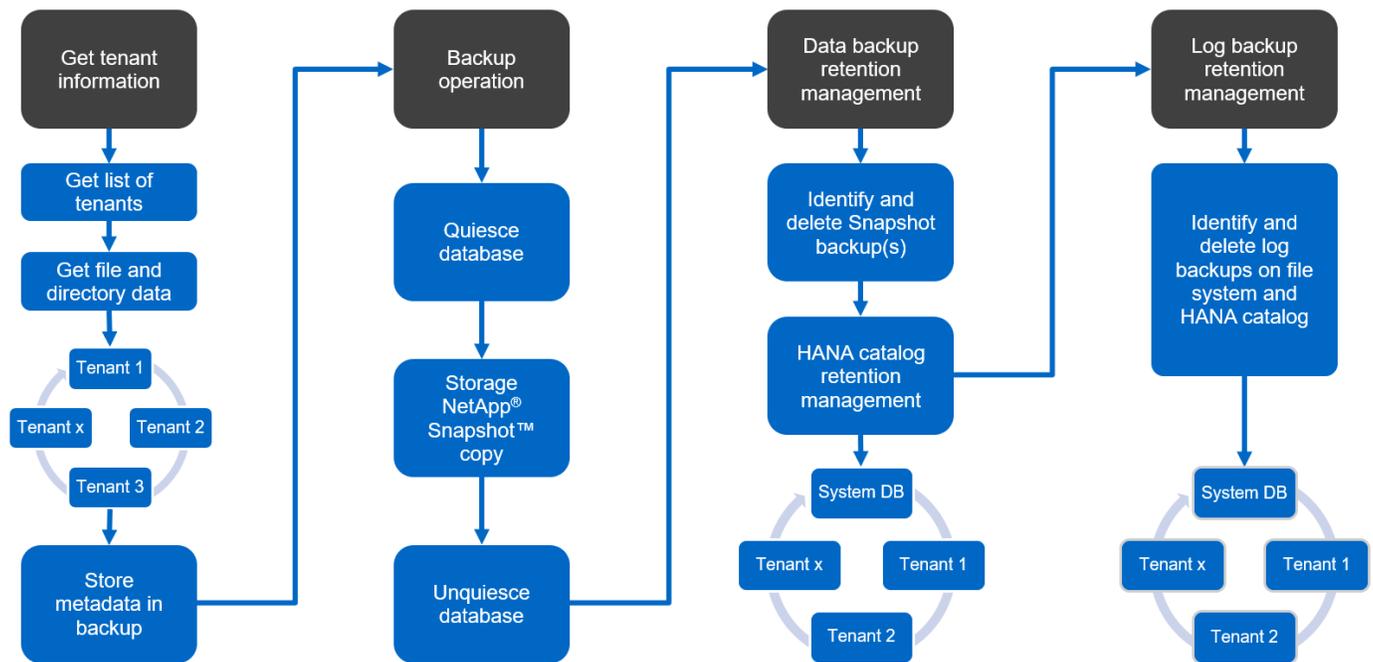
Therefore, with each backup operation, the first step in the workflow is to get the tenant information. This includes the tenant names and the corresponding file and directory information. This data must be stored in the Snapshot backup metadata in order to be able to support a single tenant restore operation. The next step is the Snapshot backup operation itself. This step includes the SQL command to trigger the HANA backup savepoint, the storage Snapshot backup, and the SQL command to close the Snapshot operation. By using the close command, the HANA database updates the backup catalog of the system DB and each tenant.



SAP does not support Snapshot backup operations for MDC systems when one or more tenants are stopped.

For the retention management of data backups and the HANA backup catalog management, SnapCenter must execute the catalog delete operations for the system database and all tenant databases that were identified in the first step. In the same way for the log backups, the SnapCenter workflow must operate on each tenant that was part of the backup operation.

The following figure shows an overview of the backup workflow.



### Backup workflow for Snapshot backups of the HANA database

SnapCenter backs up the SAP HANA database in the following sequence:

1. SnapCenter reads the list of tenants from the HANA database.
2. SnapCenter reads the files and directories for each tenant from the HANA database.

3. Tenant information is stored in the SnapCenter metadata for this backup operation.
4. SnapCenter triggers an SAP HANA global synchronized backup save point to create a consistent database image on the persistence layer.



For an SAP HANA MDC single or multiple tenant system, a synchronized global backup save point for the system database, and for each tenant database is created.

5. SnapCenter creates storage Snapshot copies for all data volumes configured for the resource. In our example of a single-host HANA database, there is only one data volume. With an SAP HANA multiple-host database, there are multiple data volumes.
6. SnapCenter registers the storage Snapshot backup in the SAP HANA backup catalog.
7. SnapCenter deletes the SAP HANA backup save point.
8. SnapCenter starts a SnapVault or SnapMirror update for all configured data volumes in the resource.



This step is only executed if the selected policy includes a SnapVault or SnapMirror replication.

9. SnapCenter deletes the storage Snapshot copies and the backup entries in its database as well as in the SAP HANA backup catalog based on the retention policy defined for backups at the primary storage. HANA backup catalog operations are done for the system database and all tenants.



If the backup is still available at the secondary storage, the SAP HANA catalog entry is not deleted.

10. SnapCenter deletes all log backups on the file system and in the SAP HANA backup catalog that are older than the oldest data backup identified in the SAP HANA backup catalog. These operations are done for the system database and all tenants.



This step is only executed if log backup housekeeping is not disabled.

### **Backup workflow for block integrity check operations**

SnapCenter executes the block integrity check in the following sequence:

1. SnapCenter reads the list of tenants from the HANA database.
2. SnapCenter triggers a file-based backup operation for the system database and each tenant.
3. SnapCenter deletes file-based backups in its database, on the file system, and in the SAP HANA backup catalog based on the retention policy defined for block integrity check operations. Backup deletion on the file system and HANA backup catalog operations are done for the system database and all tenants.
4. SnapCenter deletes all log backups on the file system and in the SAP HANA backup catalog that are older than the oldest data backup identified in the SAP HANA backup catalog. These operations are done for the system database and all tenants.



This step is only executed if log backup housekeeping is not disabled.

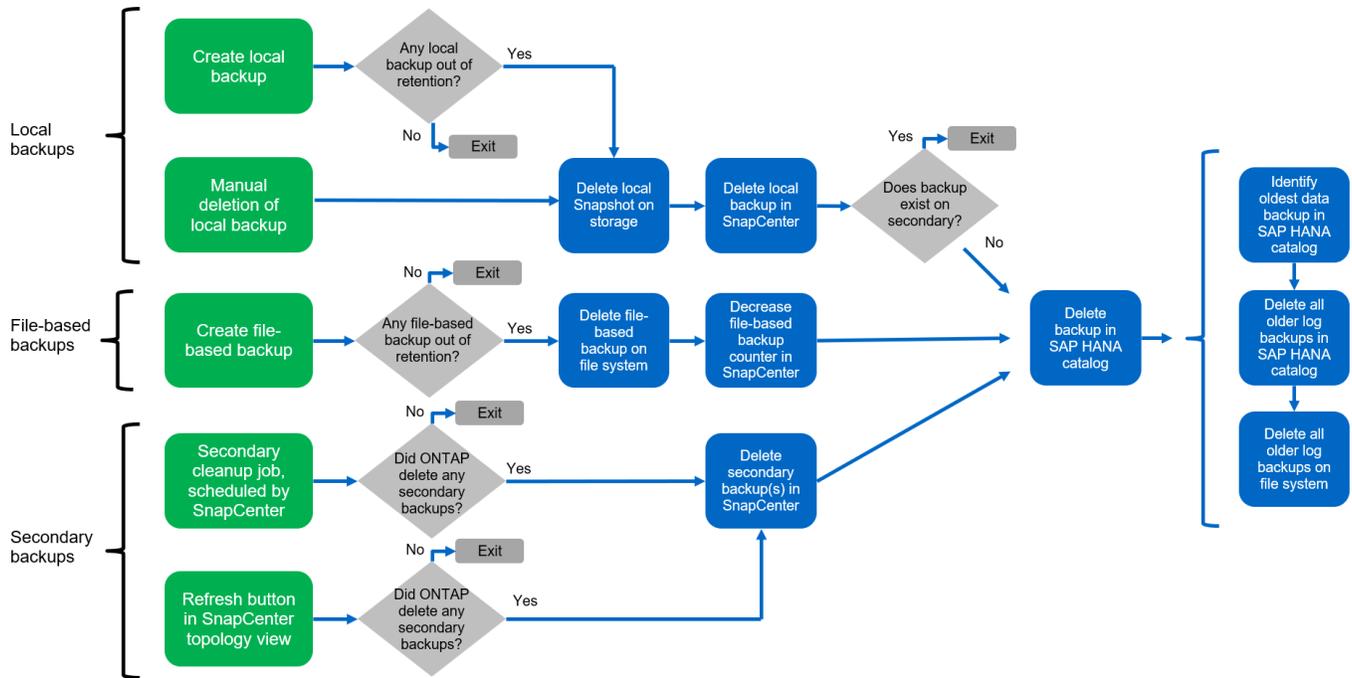
### **Backup retention management and housekeeping of data and log backups**

The data backup retention management and log backup housekeeping can be divided into five main areas,

including retention management of:

- Local backups at the primary storage
- File-based backups
- Backups at the secondary storage
- Data backups in the SAP HANA backup catalog
- Log backups in the SAP HANA backup catalog and the file system

The following figure provides an overview of the different workflows and the dependencies of each operation. The following sections describe the different operations in detail.



### Retention management of local backups at the primary storage

SnapCenter handles the housekeeping of SAP HANA database backups and non-data volume backups by deleting Snapshot copies on the primary storage and in the SnapCenter repository according to a retention defined in the SnapCenter backup policy.

Retention management logic is executed with each backup workflow in SnapCenter.



Be aware that SnapCenter handles retention management individually for both scheduled and on-demand backups.

Local backups at the primary storage can also be deleted manually in SnapCenter.

### Retention management of file-based backups

SnapCenter handles the housekeeping of file-based backups by deleting the backups on the file system according to a retention defined in the SnapCenter backup policy.

Retention management logic is executed with each backup workflow in SnapCenter.



Be aware that SnapCenter handles retention management individually for scheduled or on-demand backups.

### Retention management of backups at the secondary storage

The retention management of backups at the secondary storage is handled by ONTAP based on the retention defined in the ONTAP protection relationship.

To synchronize these changes on the secondary storage in the SnapCenter repository, SnapCenter uses a scheduled cleanup job. This cleanup job synchronizes all secondary storage backups with the SnapCenter repository for all SnapCenter plug-ins and all resources.

The cleanup job is scheduled once per week by default. This weekly schedule results in a delay with deleting backups in SnapCenter and SAP HANA Studio when compared with the backups that have already been deleted at the secondary storage. To avoid this inconsistency, customers can change the schedule to a higher frequency, for example, once per day.



The cleanup job can also be triggered manually for an individual resource by clicking the refresh button in the topology view of the resource.

For details about how to adapt the schedule of the cleanup job or how to trigger a manual refresh, refer to the section [“Change scheduling frequency of backup synchronization with off-site backup storage.”](#)

### Retention management of data backups within the SAP HANA backup catalog

When SnapCenter has deleted any backup, local Snapshot or file based, or has identified the backup deletion at the secondary storage, this data backup is also deleted in the SAP HANA backup catalog.

Before deleting the SAP HANA catalog entry for a local Snapshot backup at the primary storage, SnapCenter checks if the backup still exists at the secondary storage.

### Retention management of log backups

The SAP HANA database automatically creates log backups. These log backup runs create backup files for each individual SAP HANA service in a backup directory configured in SAP HANA.

Log backups older than the latest data backup are no longer required for forward recovery and can therefore be deleted.

SnapCenter handles the housekeeping of log file backups on the file system level as well as in the SAP HANA backup catalog by executing the following steps:

1. SnapCenter reads the SAP HANA backup catalog to get the backup ID of the oldest successful file-based or Snapshot backup.
2. SnapCenter deletes all log backups in the SAP HANA catalog and the file system that are older than this backup ID.



SnapCenter only handles housekeeping for backups that have been created by SnapCenter. If additional file-based backups are created outside of SnapCenter, you must make sure that the file-based backups are deleted from the backup catalog. If such a data backup is not deleted manually from the backup catalog, it can become the oldest data backup, and older log backups are not deleted until this file-based backup is deleted.



Even though a retention is defined for on-demand backups in the policy configuration, the housekeeping is only done when another on-demand backup is executed. Therefore, on-demand backups typically must be deleted manually in SnapCenter to make sure that these backups are also deleted in the SAP HANA backup catalog and that log backup housekeeping is not based on an old on-demand backup.

Log backup retention management is enabled by default. If required, it can be disabled as described in the section [“Disable auto discovery on the HANA plug-in host.”](#)

## Capacity requirements for Snapshot backups

You must consider the higher block change rate on the storage layer relative to the change rate with traditional databases. Due to the HANA table merge process of the column store, the complete table is written to disk, not just the changed blocks.

Data from our customer base shows a daily change rate between 20% and 50% if multiple Snapshot backups are taken during the day. At the SnapVault target, if the replication is done only once per day, the daily change rate is typically smaller.

## Restore and recovery operations

### Restore operations with SnapCenter

From the HANA database perspective, SnapCenter supports two different restore operations.

- **Restore of the complete resource.** All data of the HANA system is restored. If the HANA system contains one or more tenants, the data of the system database and the data of all tenants are restored.
- **Restore of a single tenant.** Only the data of the selected tenant is restored.

From the storage perspective, the above restore operations must be executed differently depending on the used storage protocol (NFS or Fibre Channel SAN), the configured data protection (primary storage with or without offsite backup storage), and the selected backup to be used for the restore operation (restore from primary or offsite backup storage).

### Restore of complete resource from primary storage

When restoring the complete resource from primary storage, SnapCenter supports two different ONTAP features to execute the restore operation. You can choose between the following two features:

- **Volume-based SnapRestore.** A volume based SnapRestore reverts the content of the storage volume to the state of the selected Snapshot backup.
  - Volume Revert check box available for auto discovered resources using NFS.
  - Complete Resource radio button for manual configured resources.
- **File-based SnapRestore.** A file-based SnapRestore, also known as Single File SnapRestore, restores all individual files (NFS), or all LUNs (SAN).
  - Default restore method for auto discovered resources. Can be changed using the Volume revert check box for NFS.
  - File-level radio button for manual configured resources.

The following table provides a comparison of the different restore methods.

	Volume-based SnapRestore	File-based SnapRestore
Speed of restore operation	Very fast, independent of the volume size	Very fast restore operation but uses background copy job on the storage system, which blocks the creation of new Snapshot backups
Snapshot backup history	Restore to an older Snapshot backup, removes all newer Snapshot backups.	No influence
Restore of directory structure	Directory structure is also restored	NFS: Only restores the individual files, not the directory structure. If the directory structure is also lost, it must be created manually before executing the restore operation SAN: Directory structure is also restored
Resource configured with replication to offsite backup storage	A volume-based restore cannot be done to a Snapshot copy backup that is older than the Snapshot copy used for SnapVault synchronization	Any Snapshot backup can be selected

#### Restore of complete resource from offsite backup storage

A restore from the offsite backup storage is always executed using a SnapVault restore operation where all files or all LUNs of the storage volume are overwritten with the content of the Snapshot backup.

#### Restore of a single tenant

Restoring a single tenant requires a file-based restore operation. Depending on the used storage protocol, different restore workflows are executed by SnapCenter.

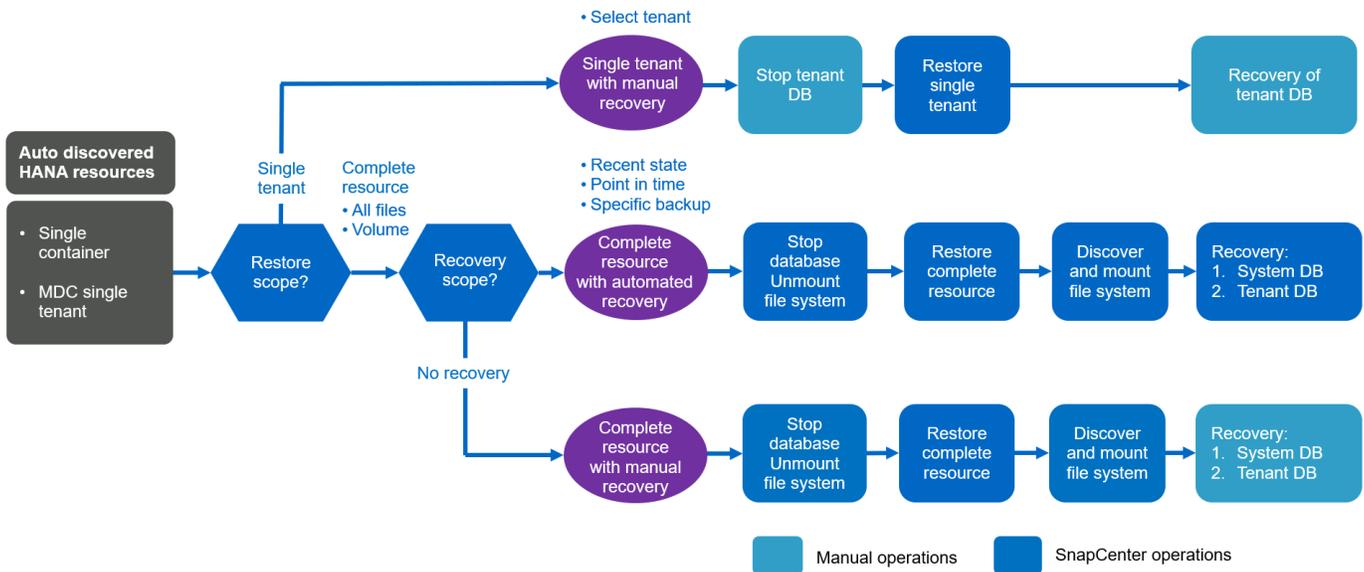
- NFS:
  - Primary storage. File-based SnapRestore operations are executed for all files of the tenant database.
  - Offsite backup storage: SnapVault restore operations are executed for all files of the tenant database.
- SAN:
  - Primary storage. Clone and connect the LUN to the database host and copy all files of the tenant database.
  - Offsite backup storage. Clone and connect the LUN to the database host and copy all files of the tenant database.

#### Restore and recovery of auto-discovered HANA single container and MDC single tenant systems

HANA single container and HANA MDC single tenant systems that have been auto discovered are enabled for automated restore and recovery with SnapCenter. For these HANA systems, SnapCenter supports three different restore and recovery workflows, as shown in the following figure:

- **Single tenant with manual recovery.** If you select a single tenant restore operation, SnapCenter lists all tenants that are included in the selected Snapshot backup. You must stop and recover the tenant database manually. The restore operation with SnapCenter is done with single file SnapRestore operations for NFS, or clone, mount, copy operations for SAN environments.

- **Complete resource with automated recovery.** If you select a complete resource restore operation and automated recovery, the complete workflow is automated with SnapCenter. SnapCenter supports up to recent state, point in time, or to specific backup recovery operations. The selected recovery operation is used for the system and the tenant database.
- **Complete resource with manual recovery.** If you select No Recovery, SnapCenter stops the HANA database and executes the required file system (unmount, mount) and restore operations. You must recover the system and tenant database manually.

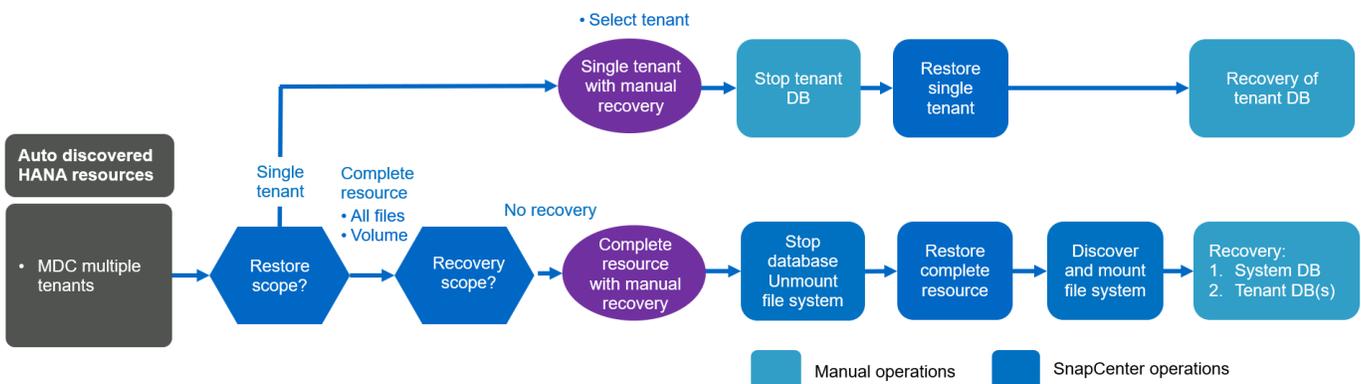


### Restore and recovery of automatically discovered HANA MDC multiple tenant systems

Even though HANA MDC systems with multiple tenants can be automatically discovered, automated restore and recovery is not supported with the current SnapCenter release. For MDC systems with multiple tenants, SnapCenter supports two different restore and recovery workflows, as shown in the following figure:

- Single tenant with manual recovery
- Complete resource with manual recovery

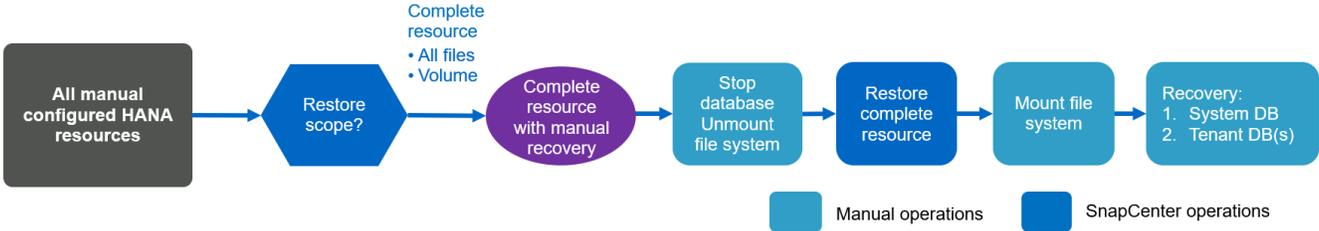
The workflows are the same as described in the previous section.



### Restore and recovery of manual configured HANA resources

Manual configured HANA resources are not enabled for automated restore and recovery. Also, for MDC systems with single or multiple tenants, a single tenant restore operation is not supported.

For manual configured HANA resources, SnapCenter only supports manual recovery, as shown in the following figure. The workflow for manual recovery is the same as described in the previous sections.



**Summary restore and recovery operations**

The following table summarizes the restore and recovery operations depending on the HANA resource configuration in SnapCenter.

SnapCenter resource configuration	Restore and recovery options	Stop HANA database	Unmount before, mount after restore operation	Recovery operation
Auto discovered Single container MDC single tenant	<ul style="list-style-type: none"> <li>Complete resource with either</li> <li>Default (all files)</li> <li>Volume revert (NFS from primary storage only)</li> <li>Automated recovery selected</li> </ul>	Automated with SnapCenter	Automated with SnapCenter	Automated with SnapCenter
	<ul style="list-style-type: none"> <li>Complete resource with either</li> <li>Default (all files)</li> <li>Volume revert (NFS from primary storage only)</li> <li>No recovery selected</li> </ul>	Automated with SnapCenter	Automated with SnapCenter	Manual
	<ul style="list-style-type: none"> <li>Tenant restore</li> </ul>	Manual	Not required	Manual

SnapCenter resource configuration	Restore and recovery options	Stop HANA database	Unmount before, mount after restore operation	Recovery operation
Auto discovered MDC multiple tenants	<ul style="list-style-type: none"> <li>• Complete resource with either</li> <li>• Default (all files)</li> <li>• Volume revert (NFS from primary storage only)</li> <li>• Automated recovery not supported</li> </ul>	Automated with SnapCenter	Automated with SnapCenter	Manual
	<ul style="list-style-type: none"> <li>• Tenant restore</li> </ul>	Manual	Not required	Manual
All manual configured resources	<ul style="list-style-type: none"> <li>• Complete resource (= Volume revert, available for NFS and SAN from primary storage only)</li> <li>• File level (all files)</li> <li>• Automated recovery not supported</li> </ul>	Manual	Manual	Manual

## Lab setup used for this report

The lab setup used for this technical report includes five different SAP HANA configurations:

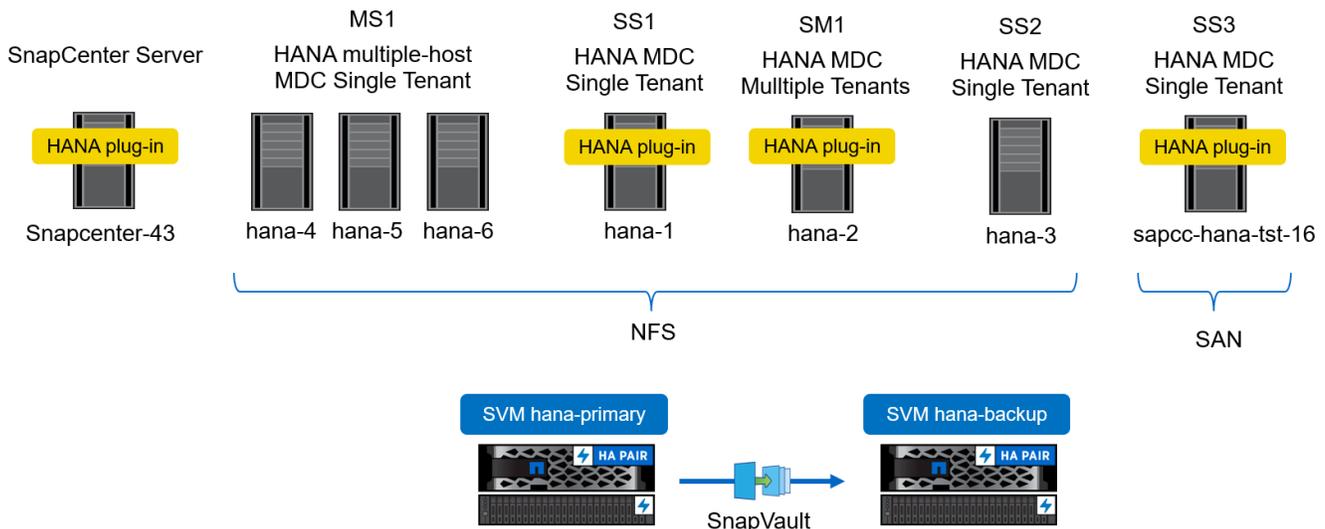
- **MS1.**
  - SAP HANA multiple-host MDC single tenant system
  - Managed with a central plug-in host (SnapCenter server)
  - Uses NFS as storage protocol
- **SS1.**
  - SAP HANA single-host MDC single tenant system
  - Auto discovered with HANA plug-in installed on HANA database host
  - Uses NFS as storage protocol
- **SM1.**

- SAP HANA single-host MDC multiple tenant system
- Auto discovered with HANA plug-in installed on HANA database host
- Uses NFS as storage protocol
- **SS2.**
  - SAP HANA single-host MDC single tenant system
  - Managed with a central plug-in host (SnapCenter Server)
  - Uses NFS as storage protocol
- **SS3.**
  - SAP HANA single-host MDC single tenant system
  - Auto discovered with HANA plug-in installed on HANA database host
  - Uses Fibre Channel SAN as storage protocol

The following sections describe the complete configuration and the backup, restore, and recovery workflows. The description covers local Snapshot backups as well as replication to backup storage using SnapVault. The storage virtual machines (SVMs) are `hana-primary` for the primary storage and `hana-backup` for the off-site backup storage.

The SnapCenter Server is used as a central HANA plug-in host for the HANA systems MS1 and SS2.

The following figure shows the lab setup.

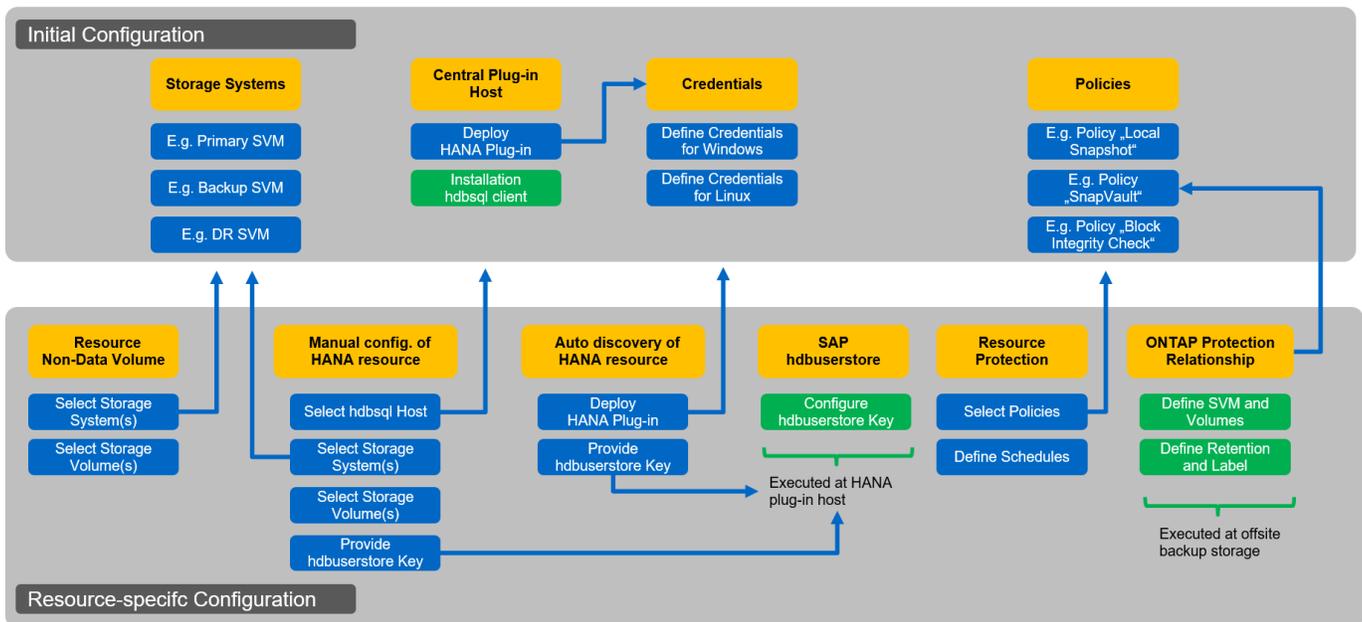


## SnapCenter configuration

The SnapCenter configuration can be separated into two main areas:

- **Initial configuration.** Covers generic configurations, independent of an individual SAP HANA database. Configurations such as storage systems, central HANA plug-in hosts, and policies, which are selected when executing the resource-specific configurations.
- **Resource-specific configuration.** Covers SAP HANA system-specific configurations and must be done for each SAP HANA database.

The following figure provides an overview of the configuration components and their dependencies. The green boxes show configuration steps that must be done outside of SnapCenter; the blue boxes show the steps that are done using the SnapCenter GUI.



With the initial configuration, the following components are installed and configured:

- **Storage system.** Credential configuration for all SVMs that are used by the SAP HANA systems: typically, primary, off-site backup, and disaster recovery storage.
- **Credentials.** Configuration of credentials used to deploy the SAP HANA plug-in on the hosts.
- **Hosts (for central HANA plug-in hosts).** Deployment of SAP HANA plug-in. Installation of the SAP HANA hdbclient software on the host. The SAP hdbclient software must be installed manually.
- **Policies.** Configuration of backup type, retention, and replication. Typically, at least one policy for local Snapshot copies, one for SnapVault replication, and one for file-based backup is required.



Storage cluster credentials can also be configured instead of individual SVM credentials.

The resource-specific configuration must be performed for each SAP HANA database and includes the following configurations:

- SAP HANA non-data volume resource configuration:
  - Storage systems and volumes
- SAP hdbuserstore key configuration:
  - The SAP hdbuserstore key configuration for the specific SAP HANA database must be performed either on the central plug-in host, or on the HANA database host, depending on where the HANA plug-in is deployed.
- Auto discovered SAP HANA database resources:
  - Deployment of SAP HANA plug-in on database host
  - Provide hdbuserstore key
- Manual SAP HANA database resource configuration:

- SAP HANA database SID, plug-in host, hdbuserstore key, storage systems and volumes
- Resource protection configuration:
  - Selection of required policies
  - Definition of schedules for each policy
- ONTAP data protection configuration:
  - Only required if the backups should be replicated to an off-site backup storage.
  - Definition of relationship and retention.

## **Initial SnapCenter configuration**

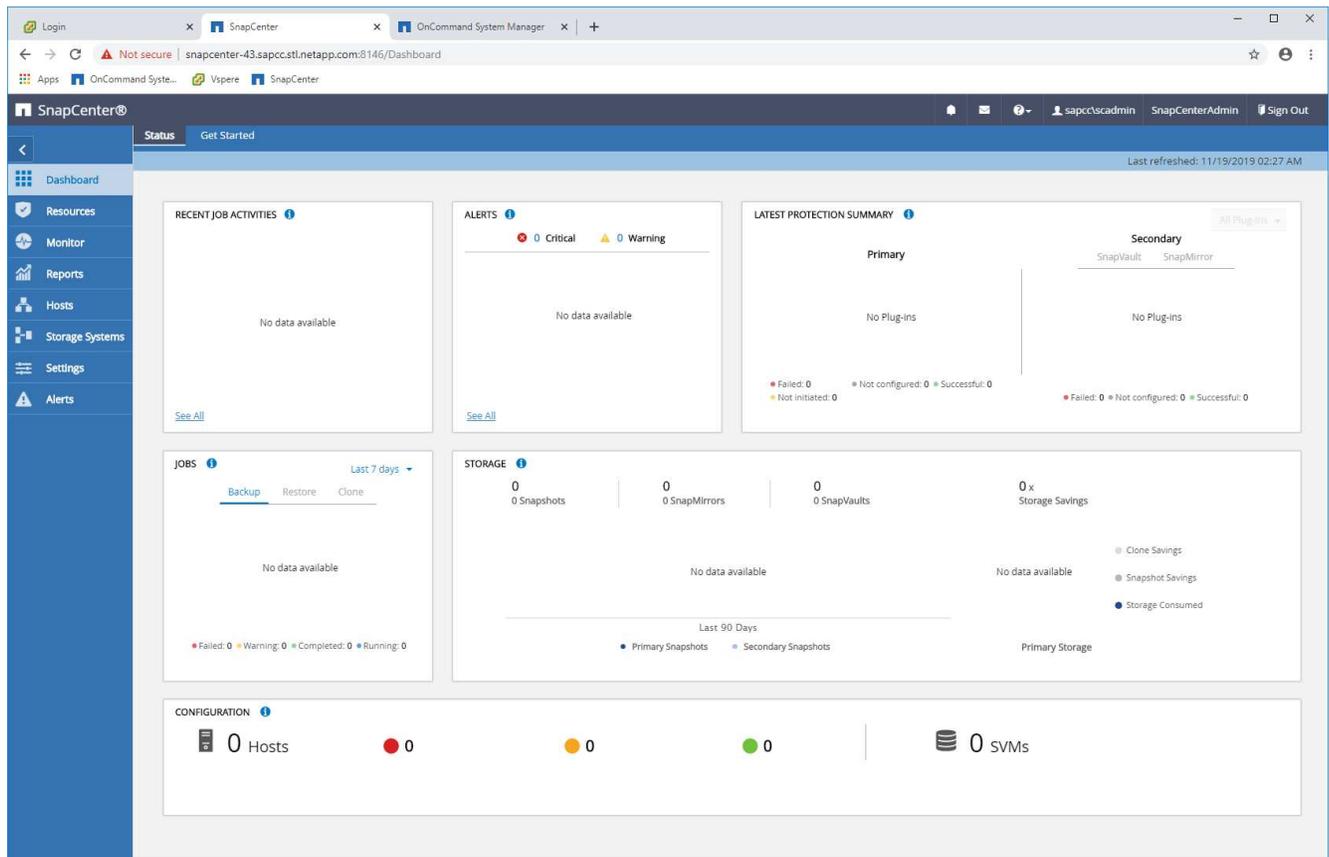
Initial configuration includes the following steps:

1. Storage system configuration
2. Credentials configuration for plug-in installation
3. For a central HANA plug-in host:
  - a. Host configuration and SAP HANA plug-in deployment
  - b. SAP HANA hdbsql client software installation and configuration
4. Policies configuration

The following sections describe the initial configuration steps.

### **Storage system configuration**

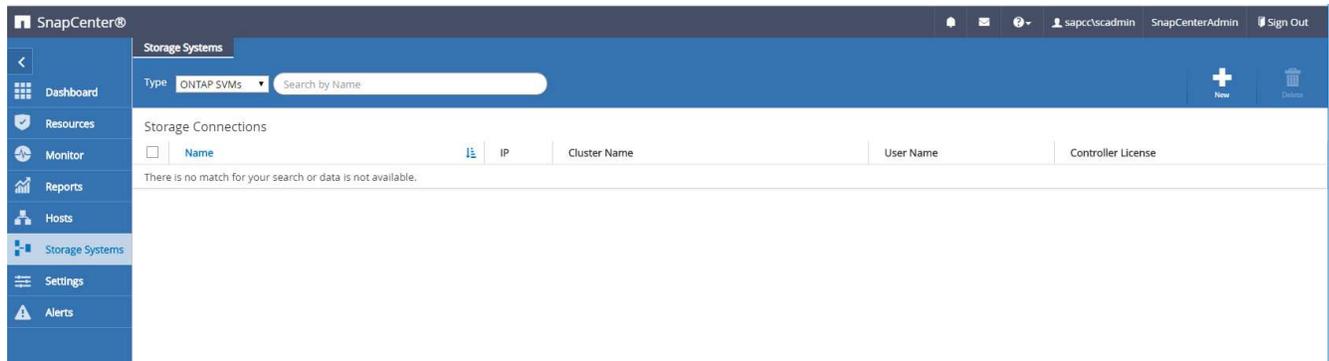
1. Log in to the SnapCenter Server GUI.



## 2. Select Storage Systems.



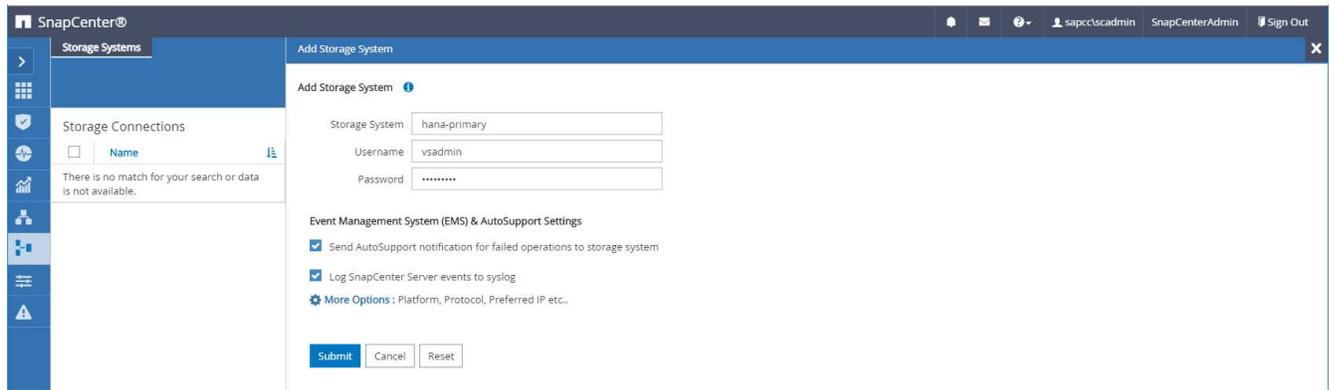
In the screen, you can select the storage system type, which can be ONTAP SVMs or ONTAP Clusters. If you configure the storage systems on SVM level, you need to have a management LIF configured for each SVM. As an alternative, you can use a SnapCenter management access on cluster level. SVM management is used in the following example.



## 3. Click New to add a storage system and provide the required host name and credentials.



The SVM user is not required to be the vsadmin user, as shown in the screenshot. Typically, a user is configured on the SVM and assigned the required permissions to execute backup and restore operations. Details on required privileges can be found in the [SnapCenter Installation Guide](#) in the section titled "Minimum ONTAP privileges required".

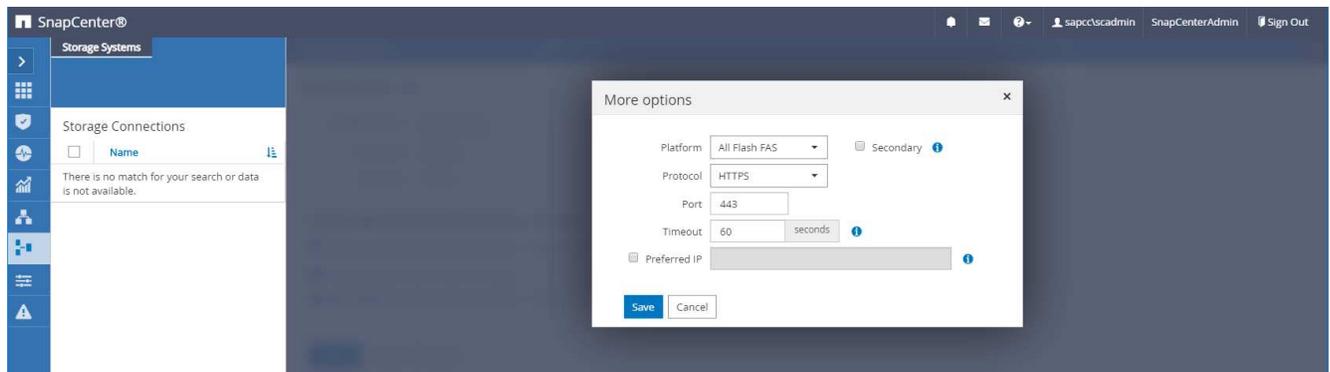


4. Click More Options to configure the storage platform.

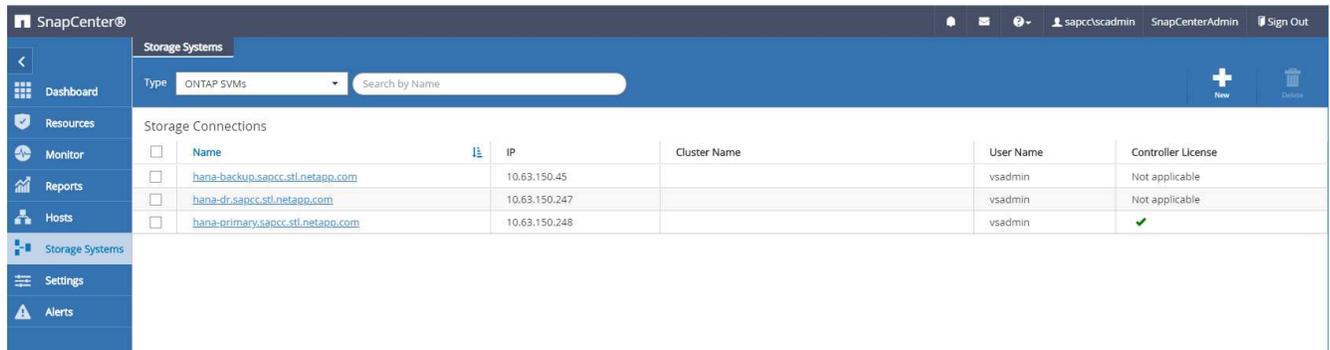
Storage platform can be FAS, AFF, ONTAP Select, or Cloud Volumes ONTAP.



For a system used as a SnapVault or SnapMirror target, select the Secondary icon.

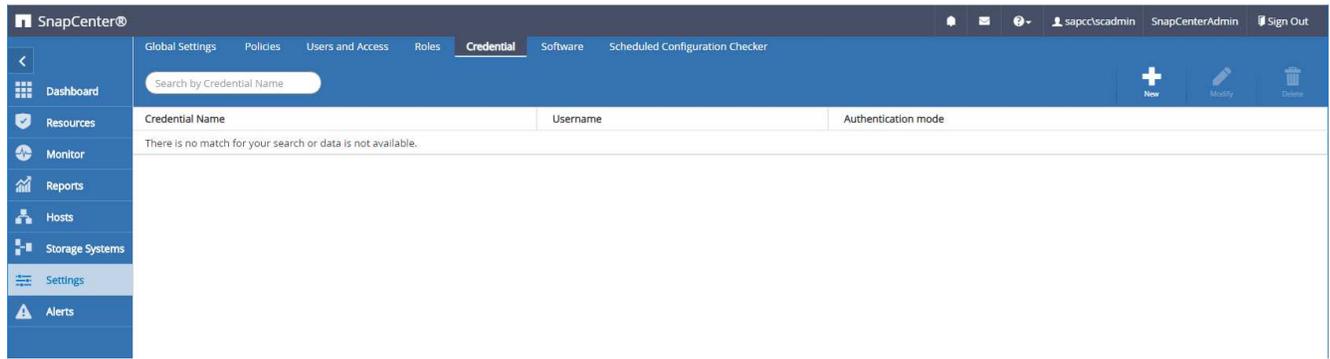


5. Add additional storage systems as required. In our example, an additional offsite backup storage and a disaster recovery storage has been added.



## Credentials configuration

1. Go to Settings, select Credentials, and click New.



2. Provide the credentials for the user that are used for plug-in installations on Linux systems.

### Credential

Provide information for the Credential you want to add

Credential Name	<input type="text" value="InstallPluginOnLinux"/>
Username	<input type="text" value="root"/> 
Password	<input type="password" value="....."/>
Authentication	<input type="text" value="Linux"/> 

Use sudo privileges 

3. Provide the credentials for the user that are used for plug-in installations on Windows systems.

✕
Credential

Provide information for the Credential you want to add

Credential Name

Username

Password

Authentication

The following figure shows the configured credentials.

Credential Name	Username	Authentication mode
InstallPluginOnLinux	root	Linux
InstallPluginOnWindows	sapcc\scadmin	Windows

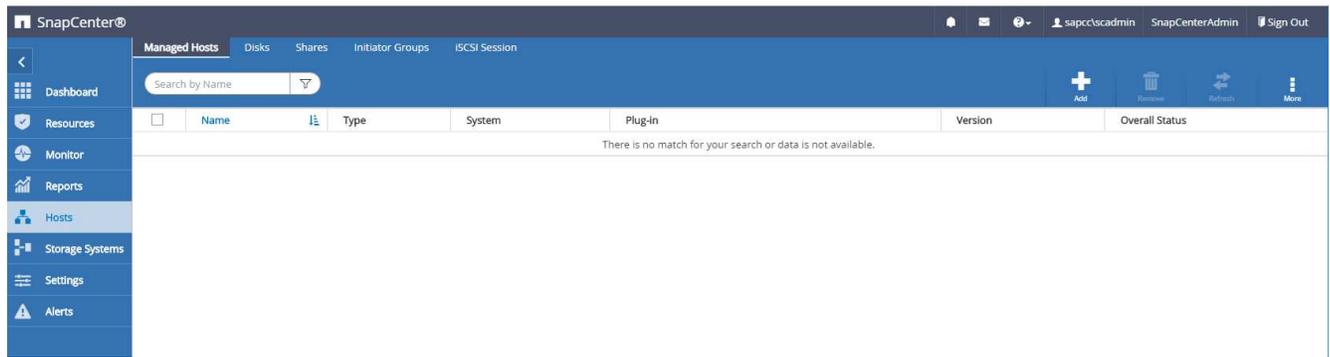
### SAP HANA plug-in installation on a central plug-in host

In the lab setup, the SnapCenter Server is also used as a central HANA plug-in host. The Windows host on which SnapCenter Server runs is added as a host, and the SAP HANA plug-in is installed on the Windows host.

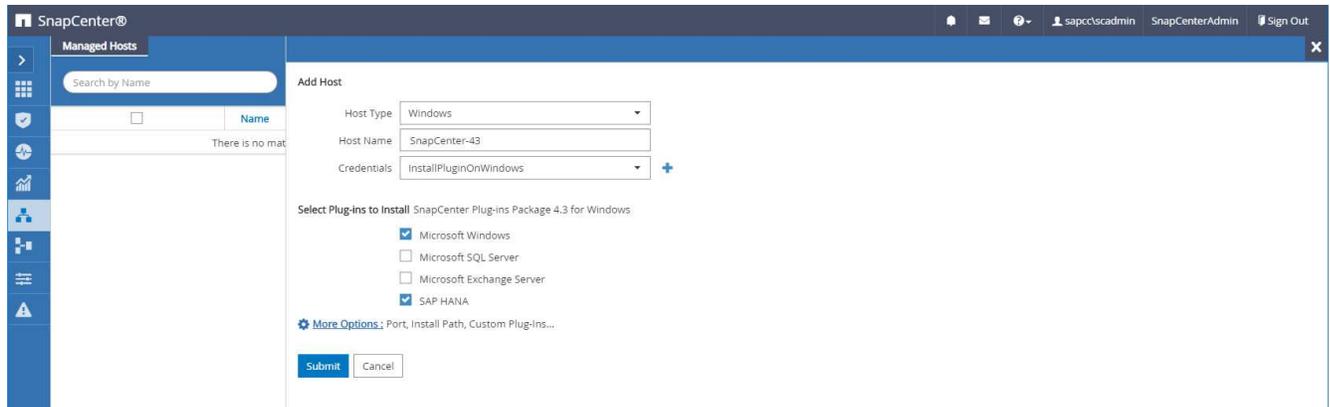


The SAP HANA plug-in requires Java 64-bit version 1.8. Java needs to be installed on the host before the SAP HANA plug-in is deployed.

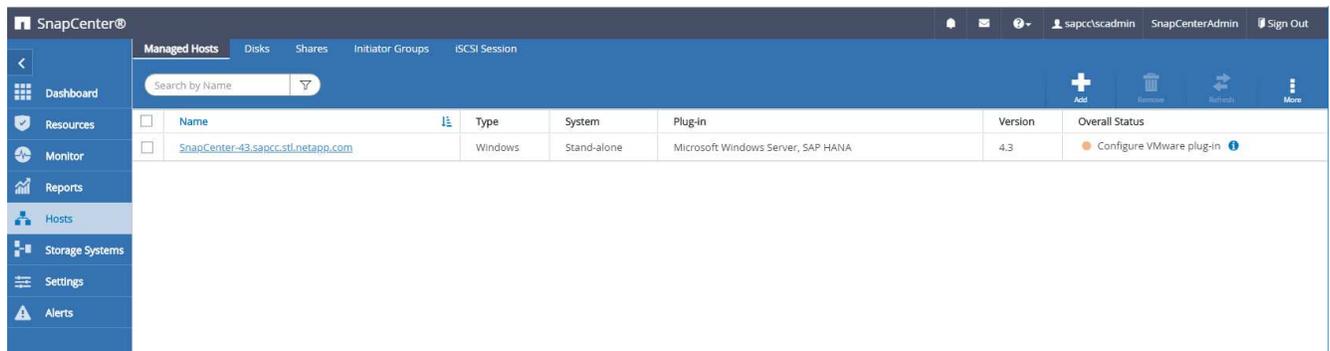
1. Go to Hosts and click Add.



2. Provide the required host information. Click Submit.



The following figure shows all the configured hosts after the HANA plug-in is deployed.



### SAP HANA hdbsql client software installation and configuration

The SAP HANA hdbsql client software must be installed on the same host on which the SAP HANA plug-in is installed. The software can be downloaded from the [SAP Support Portal](#).

The HDBSQL OS user configured during the resource configuration must be able to run the hdbsql executable. The path to the hdbsql executable must be configured in the `hana.properties` file.

- Windows:

```
C:\More C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in
Creator\etc\hana.properties
HANA_HDBSQL_CMD=C:\\Program Files\\sap\\hdbclient\\hdbsql.exe
```

- Linux:

```
cat /opt/NetApp/snapcenter/scc/etc/hana.properties
HANA_HDBSQL_CMD=/usr/sap/hdbclient/hdbsql
```

## Policy configuration

As discussed in the section “[Data protection strategy](#),” policies are usually configured independently of the resource and can be used by multiple SAP HANA databases.

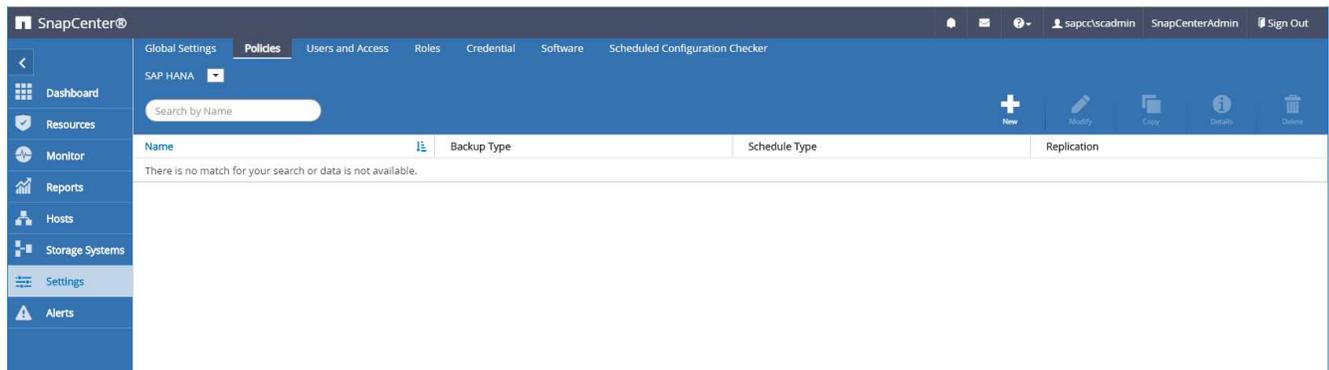
A typical minimum configuration consists of the following policies:

- Policy for hourly backups without replication: LocalSnap
- Policy for daily backups with SnapVault replication: LocalSnapAndSnapVault
- Policy for weekly block integrity check using a file-based backup: BlockIntegrityCheck

The following sections describe the configuration of these three policies.

### Policy for hourly Snapshot backups

1. Go to Settings > Policies and click New.



2. Enter the policy name and description. Click Next.

New SAP HANA Backup Policy ×

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

**Provide a policy name**

Policy name

Description

3. Select backup type as Snapshot Based and select Hourly for schedule frequency.

New SAP HANA Backup Policy ×

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

**Select backup settings**

Backup Type  Snapshot Based  File-Based ?

Schedule Frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

None

Hourly

Daily

Weekly

Monthly

4. Configure the retention settings for on-demand backups.

New SAP HANA Backup Policy ×

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

**Retention settings**

On demand backup retention settings ^

Backup retention settings ?

Total Snapshot copies to keep

Keep Snapshot copies for  days

Hourly retention settings v

5. Configure the retention settings for scheduled backups.

New SAP HANA Backup Policy x

- 1 Name
- 2 Settings
- 3 Retention
- 4 Replication
- 5 Summary

### Retention settings

On demand backup retention settings v

Hourly retention settings ^

Total Snapshot copies to keep  i

Keep Snapshot copies for

6. Configure the replication options. In this case, no SnapVault or SnapMirror update is selected.

New SAP HANA Backup Policy x

- 1 Name
- 2 Settings
- 3 Retention
- 4 Replication
- 5 Summary

### Select secondary replication options i

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label  i

Error retry count  i

7. On the Summary page, click Finish.

New SAP HANA Backup Policy x

- 1 Name
- 2 Settings
- 3 Retention
- 4 Replication
- 5 Summary

### Summary

Policy name	LocalSnap
Description	Snapshot backup at primary storage
Backup Type	Snapshot Based Backup
Schedule Type	Hourly
On demand backup retention	Total backup copies to retain : 2
Hourly backup retention	Total backup copies to retain : 12
Replication	none

#### Policy for daily Snapshot backups with SnapVault replication

1. Go to Settings > Policies and click New.
2. Enter the policy name and description. Click Next.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

**Provide a policy name**

Policy name

Description

3. Set the backup type to Snapshot Based and the schedule frequency to Daily.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

**Select backup settings**

Backup Type  Snapshot Based  File-Based i

**Schedule Frequency**

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

None  
 Hourly  
 Daily  
 Weekly  
 Monthly

4. Configure the retention settings for on-demand backups.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

**Retention settings**

On demand backup retention settings ^

Backup retention settings i

Total Snapshot copies to keep

Keep Snapshot copies for  days

Daily retention settings v

5. Configure the retention settings for scheduled backups.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Retention settings

On demand backup retention settings v

Daily retention settings ^

Total Snapshot copies to keep  i

Keep Snapshot copies for  days

6. Select Update SnapVault after creating a local Snapshot copy.



The secondary policy label must be the same as the SnapMirror label in the data protection configuration on the storage layer. See the section [“Configuration of data protection to off-site backup storage.”](#)

Modify SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

### Select secondary replication options i

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label  i

Error retry count  i

Previous
Next

7. On the Summary page, click Finish.

x
New SAP HANA Backup Policy

- 1 Name
- 2 Settings
- 3 Retention
- 4 Replication
- 5 Summary

### Summary

Policy name	LocalSnapAndSnapVault
Description	Local Snapshot backup replicated to backup storage
Backup Type	Snapshot Based Backup
Schedule Type	Daily
On demand backup retention	Total backup copies to retain : 3
Daily backup retention	Total backup copies to retain : 3
Replication	SnapVault enabled , Secondary policy label: Daily , Error retry count: 3

Previous
Finish

### Policy for Weekly Block Integrity Check

1. Go to Settings > Policies and click New.
2. Enter the policy name and description. Click Next.

x
New SAP HANA Backup Policy

- 1 Name
- 2 Settings
- 3 Retention
- 4 Replication
- 5 Summary

### Provide a policy name

Policy name	BlockIntegrityCheck <span style="float: right; font-size: 0.8em;">i</span>
Description	Block integrity check using file based backup

3. Set the backup type to File-Based and schedule frequency to Weekly.

New SAP HANA Backup Policy ×

1 Name

2 Settings

3 Retention

4 Summary

### Select backup settings

Backup Type  Snapshot Based  File-Based ?

---

### Schedule Frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

None

Hourly

Daily

Weekly

Monthly

4. Configure the retention settings for on-demand backups.

New SAP HANA Backup Policy ×

1 Name

2 Settings

3 Retention

4 Summary

### Retention settings

On demand backup retention settings
^

Backup retention settings ?

Total backup copies to keep

Keep backup copies for  days

Weekly retention settings
∨

5. Configure the retention settings for scheduled backups.

New SAP HANA Backup Policy ×

1 Name

2 Settings

3 Retention

4 Summary

### Retention settings

On demand backup retention settings
^

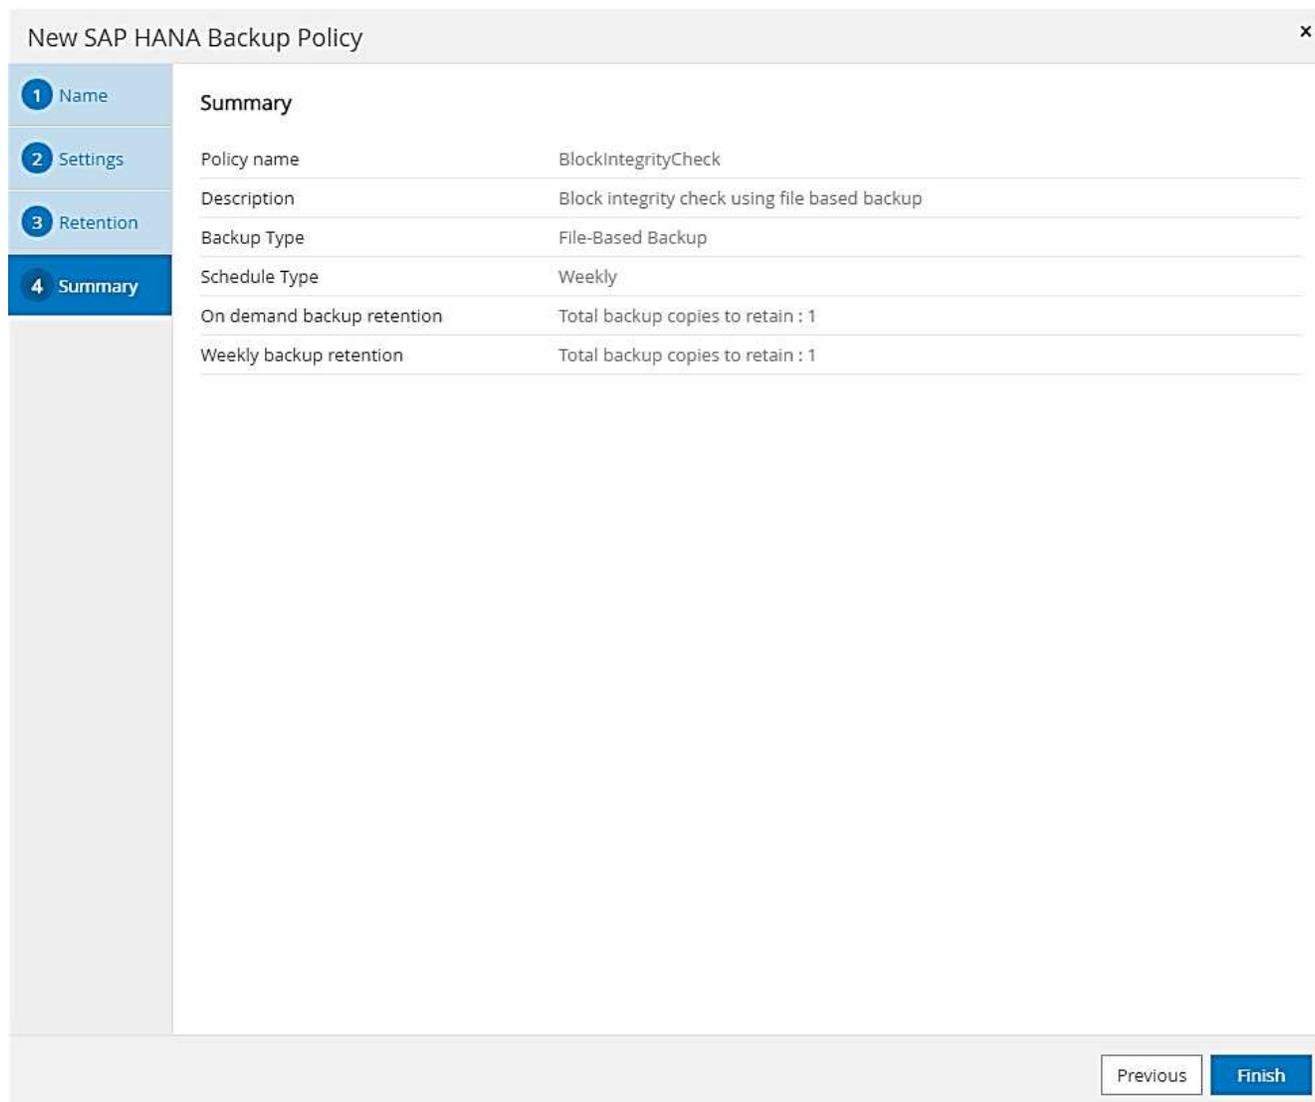
Backup retention settings ?

Total backup copies to keep

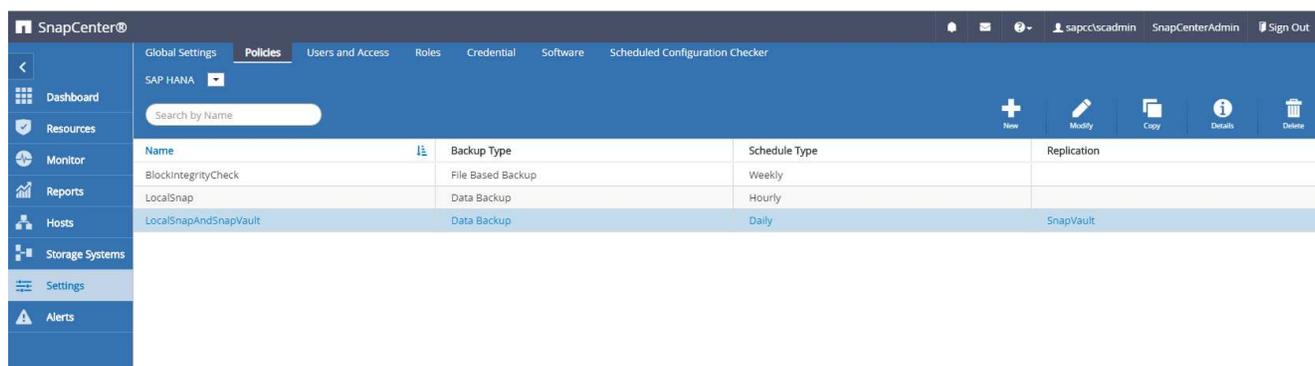
Keep backup copies for  days

Weekly retention settings
∨

6. On the Summary page, click Finish.



The following figure shows a summary of the configured policies.



## SnapCenter resource-specific configuration for SAP HANA database backups

This section describes the configuration steps for two example configurations.

- **SS2.**
  - Single-host SAP HANA MDC single-tenant system using NFS for storage access

- The resource is manually configured in SnapCenter.
- The resource is configured to create local Snapshot backups and perform block integrity checks for the SAP HANA database using a weekly file-based backup.

- **SS1.**

- Single-host SAP HANA MDC single-tenant system using NFS for storage access
- The resource is auto-discovered with SnapCenter.
- The resource is configured to create local Snapshot backups, replicate to an off-site backup storage using SnapVault, and perform block integrity checks for the SAP HANA database using a weekly file-based backup.

The differences for a SAN-attached, a single-container, or a multiple-host system are reflected in the corresponding configuration or workflow steps.

### **SAP HANA backup user and hdbuserstore configuration**

NetApp recommends configuring a dedicated database user in the HANA database to run the backup operations with SnapCenter. In the second step, an SAP HANA user store key is configured for this backup user, and this user store key is used in the configuration of the SnapCenter SAP HANA plug-in.

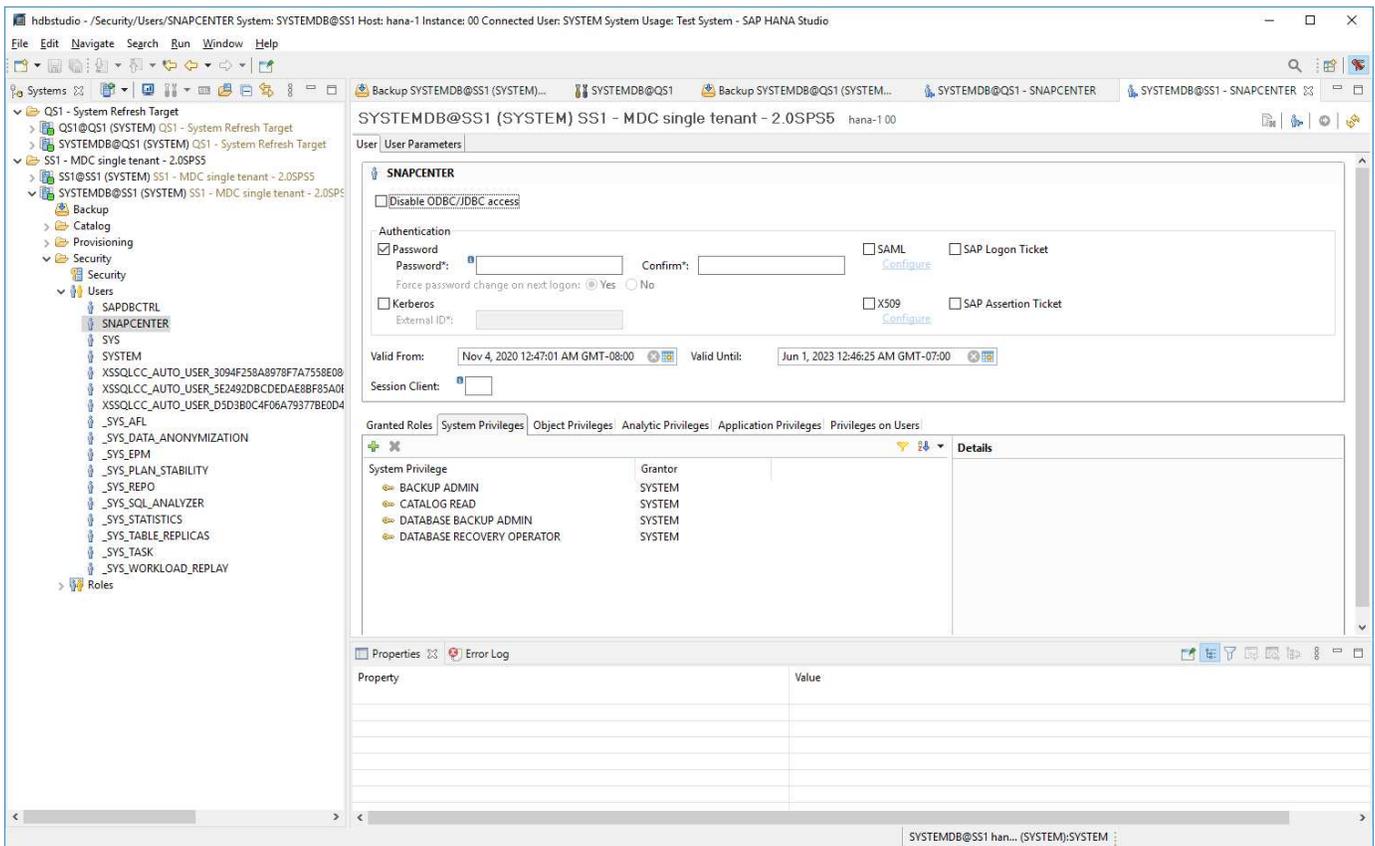
The following figure shows the SAP HANA Studio through which the backup user can be created.



The required privileges were changed with the HANA 2.0 SPS5 release: backup admin, catalog read, database backup admin, and database recovery operator. For earlier releases, backup admin and catalog read are sufficient.



For an SAP HANA MDC system, the user must be created in the system database because all backup commands for the system and the tenant databases are executed using the system database.



At the HANA plug-in host, on which the SAP HANA plug-in and the SAP hdbsql client are installed, a userstore key must be configured.

#### Userstore configuration on the SnapCenter server used as a central HANA plug-in host

If the SAP HANA plug-in and the SAP hdbsql client are installed on Windows, the local system user executes the hdbsql commands and is configured by default in the resource configuration. Because the system user is not a logon user, the user store configuration must be done with a different user and the `-u <User>` option.

```
hdbuserstore.exe -u SYSTEM set <key> <host>:<port> <database user>
<password>
```



The SAP HANA hdbclient software must be first installed on the Windows host.

#### Userstore configuration on a separate Linux host used as a Central HANA plug-in host

If the SAP HANA plug-in and SAP hdbsql client are installed on a separate Linux host, the following command is used for the user store configuration with the user defined in the resource configuration:

```
hdbuserstore set <key> <host>:<port> <database user> <password>
```



The SAP HANA hdbclient software must be first installed on the Linux host.

## Userstore configuration on the HANA database host

If the SAP HANA plug-in is deployed on the HANA database host, the following command is used for the user store configuration with the <sid>adm user:

```
hdbuserstore set <key> <host>:<port> <database user> <password>
```



SnapCenter uses the <sid>adm user to communicate with the HANA database. Therefore, the user store key must be configured using the `<sid>adm` user on the database host.



Typically, the SAP HANA hdbsql client software is installed together with the database server installation. If this is not the case, the hdbclient must be installed first.

## Userstore configuration depending on HANA system architecture

In an SAP HANA MDC single-tenant setup, port 3<instanceNo>13 is the standard port for SQL access to the system database and must be used in the hdbuserstore configuration.

For an SAP HANA single-container setup, port 3<instanceNo>15 is the standard port for SQL access to the index server and must be used in the hdbuserstore configuration.

For an SAP HANA multiple-host setup, user store keys for all hosts must be configured. SnapCenter tries to connect to the database using each of the provided keys and can therefore operate independently of a failover of an SAP HANA service to a different host.

## Userstore configuration examples

In the lab setup, a mixed SAP HANA plug-in deployment is used. The HANA plug-in is installed on the SnapCenter Server for some HANA systems and deployed on the individual HANA database servers for other systems.

### SAP HANA system SS1, MDC single tenant, instance 00

The HANA plug-in has been deployed on the database host. Therefore, the key must be configured on the database host with the user ss1adm.

```

hana-1:/ # su - ssladm
ssladm@hana-1:/usr/sap/SS1/HDB00>
ssladm@hana-1:/usr/sap/SS1/HDB00>
ssladm@hana-1:/usr/sap/SS1/HDB00> hdbuserstore set SS1KEY hana-1:30013
SnapCenter password
ssladm@hana-1:/usr/sap/SS1/HDB00> hdbuserstore list
DATA FILE      : /usr/sap/SS1/home/.hdb/hana-1/SSFS_HDB.DAT
KEY FILE       : /usr/sap/SS1/home/.hdb/hana-1/SSFS_HDB.KEY
KEY SS1KEY
  ENV : hana-1:30013
  USER: SnapCenter
KEY SS1SAPDBCTRLSS1
  ENV : hana-1:30015
  USER: SAPDBCTRL
ssladm@hana-1:/usr/sap/SS1/HDB00>

```

### SAP HANA system MS1, multiple-host MDC single tenant, instance 00

For HANA multiple host systems, a central plug-in host is required, in our setup we used the SnapCenter Server. Therefore, the user store configuration must be done on the SnapCenter Server.

```

hdbuserstore.exe -u SYSTEM set MS1KEYHOST1 hana-4:30013 SNAPCENTER
password
hdbuserstore.exe -u SYSTEM set MS1KEYHOST2 hana-5:30013 SNAPCENTER
password
hdbuserstore.exe -u SYSTEM set MS1KEYHOST3 hana-6:30013 SNAPCENTER
password
C:\Program Files\sap\hdbclient>hdbuserstore.exe -u SYSTEM list
DATA FILE      : C:\ProgramData\.hdb\SNAPCENTER-43\S-1-5-18\SSFS_HDB.DAT
KEY FILE       : C:\ProgramData\.hdb\SNAPCENTER-43\S-1-5-18\SSFS_HDB.KEY
KEY MS1KEYHOST1
  ENV : hana-4:30013
  USER: SNAPCENTER
KEY MS1KEYHOST2
  ENV : hana-5:30013
  USER: SNAPCENTER
KEY MS1KEYHOST3
  ENV : hana-6:30013
  USER: SNAPCENTER
KEY SS2KEY
  ENV : hana-3:30013
  USER: SNAPCENTER
C:\Program Files\sap\hdbclient>

```

## Configuration of data protection to off-site backup storage

The configuration of the data protection relation as well as the initial data transfer must be executed before replication updates can be managed by SnapCenter.

The following figure shows the configured protection relationship for the SAP HANA system SS1. With our example, the source volume `SS1_data_mnt00001` at the SVM `hana-primary` is replicated to the SVM `hana-backup` and the target volume `SS1_data_mnt00001_dest`.



The schedule of the relationship must be set to None, because SnapCenter triggers the SnapVault update.

The screenshot displays the OnCommand System Manager interface. The main window shows a table of Volume Relationships. A row is highlighted with a blue border, showing the following details:

Source Storage Vi...	Source Volume	Destination Volume	Destination Stora...	Is Healthy	Object ...	Rela...	Transf...	Relationship Type	Lag Time	Policy Name	Policy Type
hana-primary	SS1_data_mnt00001	SS1_data_mnt00001_dest	hana-backup	Yes	Volume	Snapmi...	Idle	Asynchronous V...	21 hrs(s)...	SnapCenterVault	Asynchronous Vault

Below the table, the configuration details for the selected relationship are shown:

Source Location:	hana-primary:SS1_data_...	Is Healthy:	Yes	Transfer Status:	Idle
Destination Location:	hana-backup:SS1_data_m...	Relationship State:	Snapmirrored	Current Transfer Type:	None
Source Cluster:	a700-marco	Network Compression Ratio:	Not Applicable	Current Transfer Error:	None
Destination Cluster:	a700-marco	Transfer Schedule:	None	Current Transfer Progress:	None
Data Transfer Rate:	Unlimited			Last Transfer Error:	None
Lag Time:	21 hr(s) 23 min(s)			Last Transfer Type:	Update
				Latest Snapshot Timestamp:	11/26/2019 11:03:53
				Latest Snapshot Copy:	SnapCenter_LocalSnapAndSnapVault_Daily_11-26-2019_08.17.01.8979

The following figure shows the protection policy. The protection policy used for the protection relationship defines the SnapMirror label, as well as the retention of backups at the secondary storage. In our example, the used label is `Daily`, and the retention is set to 5.



The SnapMirror label in the policy being created must match the label defined in the SnapCenter policy configuration. For details, refer to [Policy for daily Snapshot backups with SnapVault replication](#).



The retention for backups at the off-site backup storage is defined in the policy and controlled by ONTAP.

The screenshot displays the OnCommand System Manager interface. The main content area is titled 'Volume Relationships' and contains a table with columns: Source Storage Volume, Source Volume, Destination Volume, Destination Storage, Is Healthy, Object, Relationship, Transf., Relationship Type, Lag Time, Policy Name, and Policy Type. A single row is visible with values: hana-primary, SS1\_data\_mnt00001, SS1\_data\_mnt00001\_dest, hana-backup, Yes, Volume, Snapmi..., Idle, Asynchronous V..., 21 hrs(s)..., SnapCenterVault, Asynchronous Vault.

Below the table, the 'Policy Name' is 'SnapCenterVault'. Under 'Comments', there is a table with columns: Label, Number of Copies, and Matching Snapshot copy Schedules in Source Volume. One row is highlighted with a blue border, showing 'Daily', 5, and 'Source does not have any schedules with this label'.

At the bottom, there are tabs for 'Details', 'Policy Details', and 'Snapshot Copies'.

## Manual HANA resource configuration

This section describes the manual configuration of the SAP HANA resources SS2 and MS1.

- SS2 is a single-host MDC single-tenant system
- MS1 is a multiple-host MDC single-tenant system.
  1. From the Resources tab, select SAP HANA and click Add SAP HANA Database.
  2. Enter the information for configuring the SAP HANA database and click Next.

Select the resource type in our example, Multitenant Database Container.



For a HANA single container system, the resource type Single Container must be selected. All the other configuration steps are identical.

For our SAP HANA system, the SID is SS2.

The HANA plug-in host in our example is the SnapCenter Server.

The hdbuserstore key must match the key that was configured for the HANA database SS2. In our example it is SS2KEY.

Add SAP HANA Database x

1 Name

2 Storage Footprint

3 Summary

**Provide Resource Details**

Resource Type	Multitenant Database Container
HANA System Name	SS2 - HANA 20 SPS4 MDC Single Tenant
SID	SS2 <span style="float: right;">i</span>
Plug-in Host	SnapCenter-43.sapcc.stl.netapp.com <span style="float: right;">i</span>
HDB Secure User Store Keys	SS2KEY <span style="float: right;">i</span>
HDBSQL OS User	SYSTEM <span style="float: right;">i</span>



For an SAP HANA multiple-host system, the hdbuserstore keys for all hosts must be included, as shown in the following figure. SnapCenter will try to connect with the first key in the list, and will continue with the other case, in case the first key does not work. This is required to support HANA failover in a multiple-host system with worker and standby hosts.

Modify SAP HANA Database x

1 Name

2 Storage Footprint

3 Summary

**Provide Resource Details**

Resource Type	Multitenant Database Container
HANA System Name	MS1 - Multiple Hosts MDC Single Tenant
SID	MS1 <span style="float: right;">i</span>
Plug-in Host	SnapCenter-43.sapcc.stl.netapp.com <span style="float: right;">i</span>
HDB Secure User Store Keys	MS1KEYHOST1,MS1KEYHOST2,MS1KEYHOST3 <span style="float: right;">i</span>
HDBSQL OS User	SYSTEM <span style="float: right;">i</span>

3. Select the required data for the storage system (SVM) and volume name.

Add SAP HANA Database x

1 Name

2 Storage Footprint

3 Summary

**Provide Storage Footprint Details**

Add Storage Footprint x

Storage System hana-primary.sapcc.stl.netapp.com

Select one or more volumes and if required their associated Qtrees and LUNs

Volume name	LUNs or Qtrees
SS2_data_mnt00001	Default is 'None' or type to find

Save



For a Fibre Channel SAN configuration, the LUN needs to be selected as well.



For an SAP HANA multiple-host system, all data volumes of the SAP HANA system must be selected, as shown in the following figure.

The screenshot shows the 'Add SAP HANA Database' configuration window with the 'Storage Footprint' step selected. The 'Storage System' is set to 'hana-primary.sapcc.stl.netapp.com'. Below, there are two rows for selecting volumes and LUNs or Qtrees. The first row has 'MS1\_data\_mnt00001' selected for the volume name and 'Default is 'None' or type to find' for the LUNs or Qtrees. The second row has 'MS1\_data\_mnt00002' selected for the volume name and 'Default is 'None' or type to find' for the LUNs or Qtrees. A 'Save' button is located at the bottom right of the configuration area.

The summary screen of the resource configuration is shown.

- Click Finish to add the SAP HANA database.

The screenshot shows the 'Add SAP HANA Database' configuration window with the 'Summary' step selected. The summary table is as follows:

Summary		
Resource Type	Multitenant Database Container	
HANA System Name	SS2 - HANA 20 SP54 MDC Single Tenant	
SID	SS2	
Plug-in Host	SnapCenter-43.sapcc.stl.netapp.com	
HDB Secure User Store Keys	SS2KEY	
HDBSQL OS User	SYSTEM	
Storage Footprint		
Storage System	Volume	LUN/Qtrees
hana-primary.sapcc.stl.netapp.com	SS2_data_mnt00001	

- When resource configuration is finished, perform the configuration of resource protection as described in the section [Resource protection configuration](#).

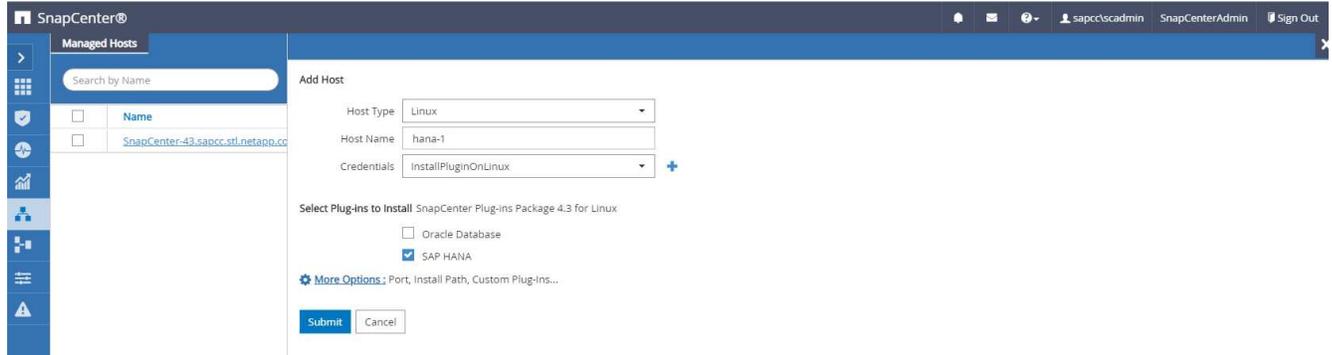
### Automatic discovery of HANA databases

This section describes the automatic discovery of the SAP HANA resource SS1 (single host MDC single tenant system with NFS). All the described steps are identical for a HANA single container, HANA MDC multiple tenants' systems, and a HANA system using Fibre Channel SAN.

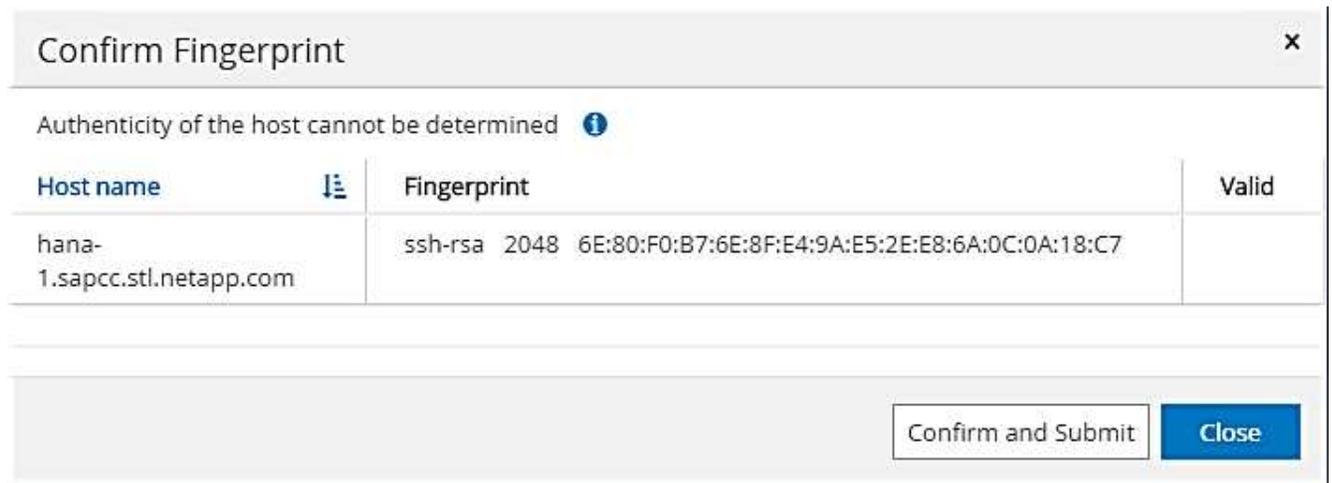


The SAP HANA plug-in requires Java 64-bit version 1.8. Java must be installed on the host before the SAP HANA plug-in is deployed.

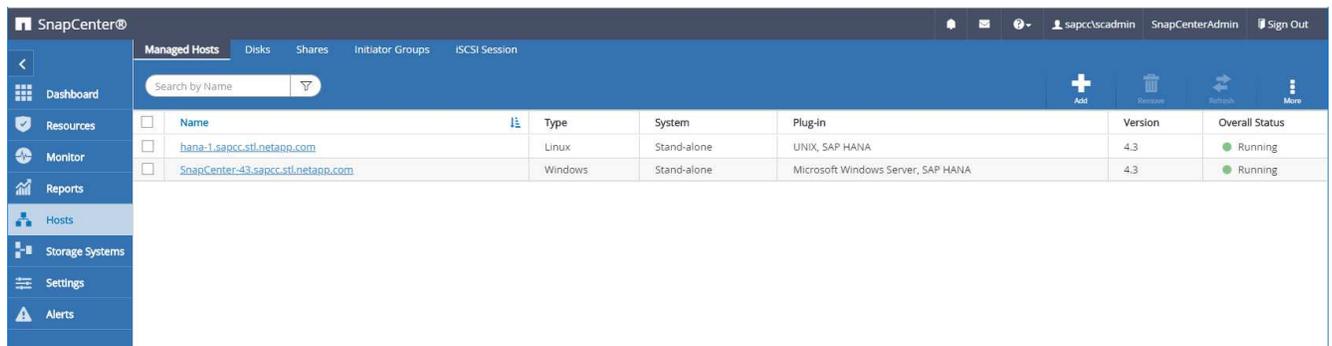
1. From the host tab, click Add.
2. Provide host information and select the SAP HANA plug-in to be installed. Click Submit.



3. Confirm the fingerprint.



The installation of the HANA plug-in and the Linux plug-in starts automatically. When the installation is finished, the status column of the host shows Running. The screen also shows that the Linux plug-in is installed together with the HANA plug-in.

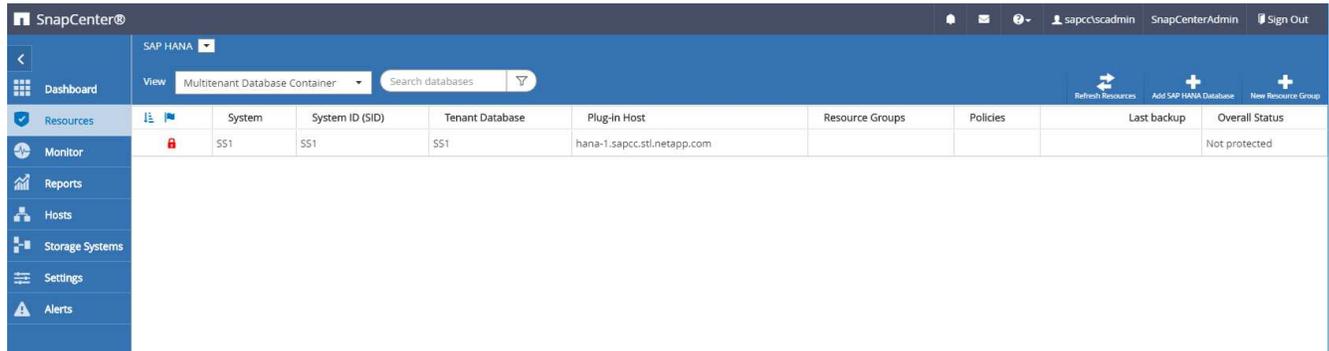


After the plug-in installation, the automatic discovery process of the HANA resource starts automatically. In the Resources screen, a new resource is created, which is marked as locked with the red padlock icon.

4. Select and click on the resource to continue the configuration.



You can also trigger the automatic discovery process manually within the Resources screen, by clicking Refresh Resources.



5. Provide the userstore key for the HANA database.



The second level automatic discovery process starts in which tenant data and storage footprint information is discovered.

6. Click Details to review the HANA resource configuration information in the resource topology view.

**Manage Copies**

Local copies: 17 Backups, 0 Clones

Vault copies: 5 Backups, 0 Clones

**Summary Card**

- 24 Backups
- 22 Snapshot based backups
- 2 File-Based backups ✓
- 0 Clones

Backup Name	Count	IF	End Date
SnapCenter_LocalSnap_Hourly_11-27-2019_02.30.01.1788	1		11/27/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_11-26-2019_22.30.01.0413	1		11/26/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_11-26-2019_18.30.01.0738	1		11/26/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_11-26-2019_14.30.01.0340	1		11/26/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_11-26-2019_10.30.01.0649	1		11/26/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_11-26-2019_08.17.01.8979	1		11/26/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_11-26-2019_06.30.01.0003	1		11/26/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_11-26-2019_02.30.00.9915	1		11/26/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_11-25-2019_22.30.01.0536	1		11/25/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_11-25-2019_18.30.01.0250	1		11/25/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_11-25-2019_14.30.01.0151	1		11/25/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_11-25-2019_10.30.00.9895	1		11/25/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_11-25-2019_08.17.01.8577	1		11/25/2019 8:17:55 AM
SnapCenter_LocalSnap_Hourly_11-24-2019_06.30.00.9717	1		11/24/2019 6:30:55 AM
Total 17			

Activity: The 5 most recent jobs are displayed. 5 Completed, 0 Warnings, 0 Failed, 0 Canceled, 0 Running, 0 Queued.

**Resource - Details**

Details for selected resource

Type	Multitenant Database Container
HANA System Name	SS1
SID	SS1
Tenant Database	SS1
Plug-in Host	hana-1.sapcc.stl.netapp.com
HDB Secure User Store Keys	SS1KEY
HDB SQL OS User	ssladm
plug-in name	SAP HANA
Last backup	11/27/2019 2:30:55 AM (Completed)
Resource Groups	hana-1_sapcc_stl_netapp_com_hana_MDC_SS1
Policy	BlockIntegrityCheck_LocalSnap_LocalSnapAndSnapVault
Discovery Type	Auto

**Storage Footprint**

SVM	Volume	Junction Path	LUN/Qtree
hana-primary.sapcc.stl.netapp.com	SS1_data_mnt00001	/SS1_data_mnt00001	

Activity: The 5 most recent jobs are displayed. 5 Completed, 0 Warnings, 0 Failed, 0 Canceled, 1 Running, 0 Queued.

When the resource configuration is finished, the resource protection configuration must be executed as described in the following section.

## Resource protection configuration

This section describes the resource protection configuration. The resource protection configuration is the same, whether the resource has been auto discovered or configured manually. It is also identical for all HANA architectures, single or multiple hosts, single container, or MDC systems.

1. From the Resources tab, double-click the resource.
2. Configure a custom name format for the Snapshot copy.



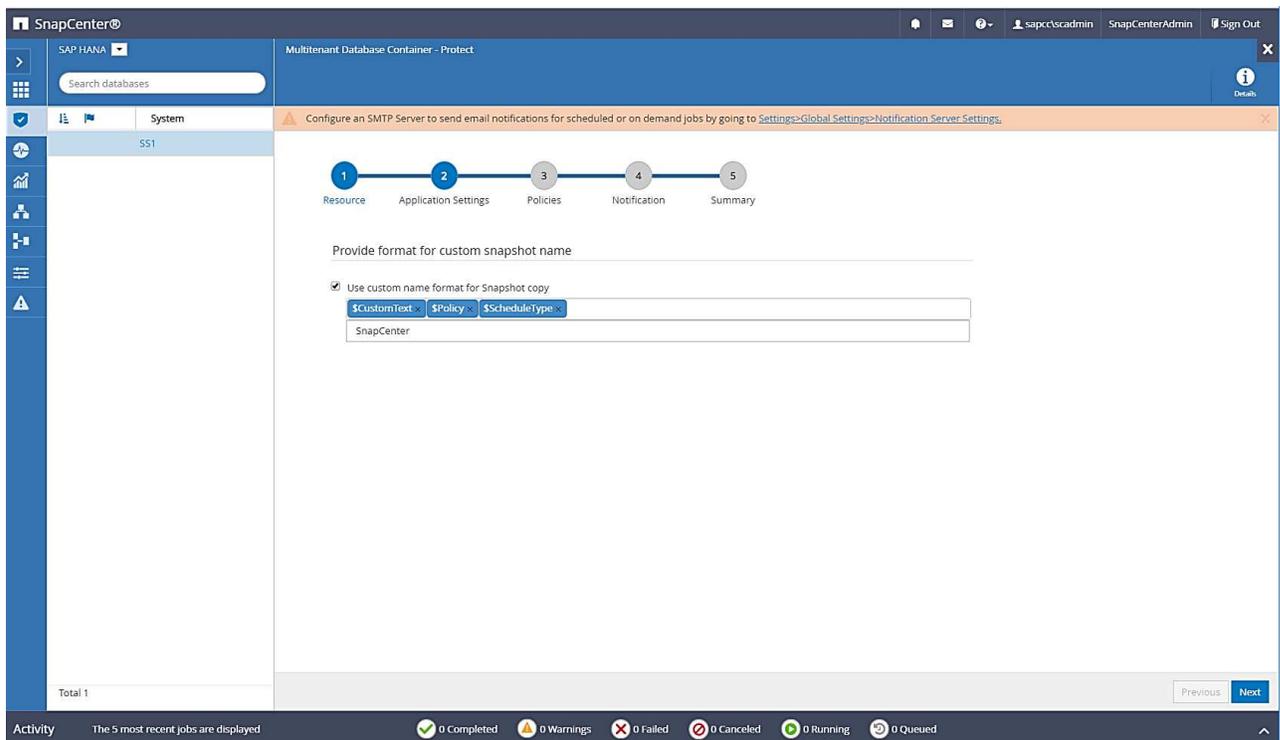
NetApp recommends using a custom Snapshot copy name to easily identify which backups have been created with which policy and schedule type. By adding the schedule type in the Snapshot copy name, you can distinguish between scheduled and on-demand backups. The `schedule_name` string for on-demand backups is empty, while scheduled backups include the string `Hourly`, `Daily`, or `Weekly`.

In the configuration shown in the following figure, the backup and Snapshot copy names have the following format:

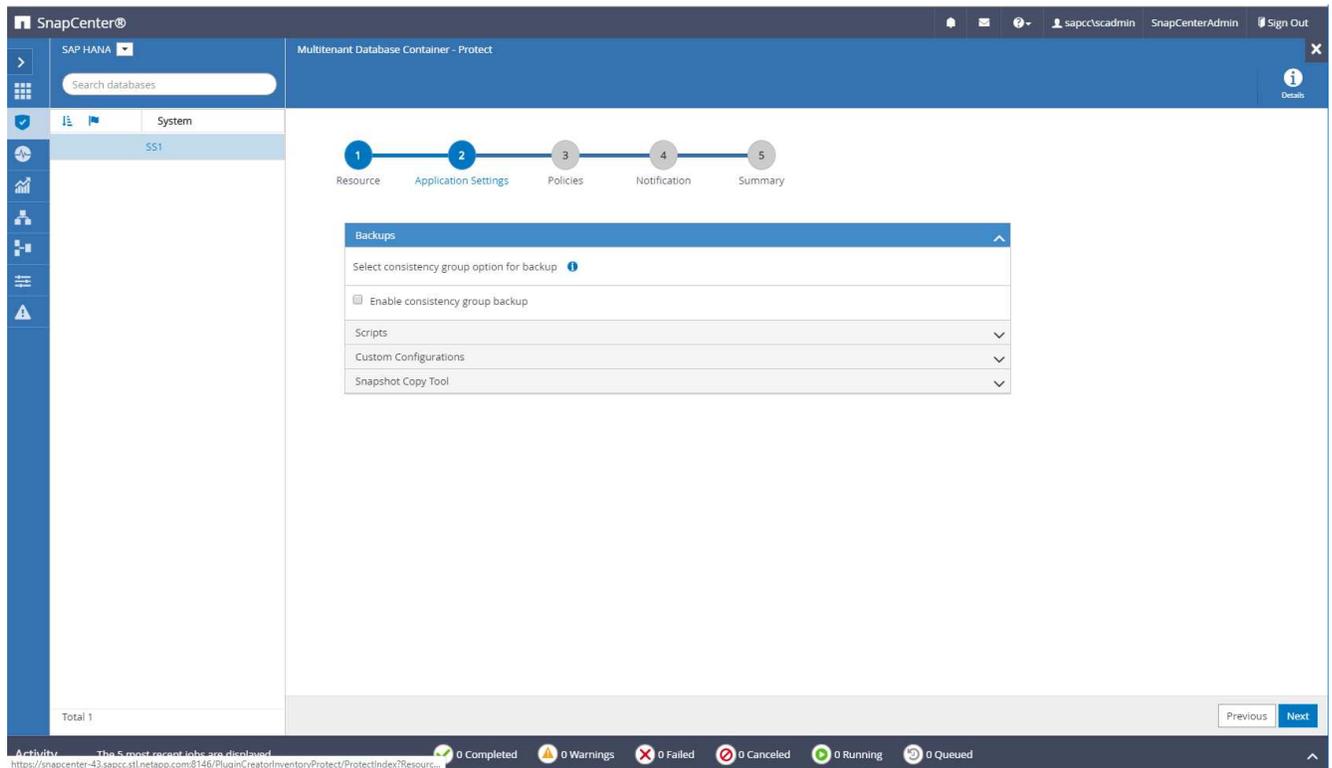
- Scheduled hourly backup: `SnapCenter_LocalSnap_Hourly_<time_stamp>`
- Scheduled daily backup: `SnapCenter_LocalSnapAndSnapVault_Daily_<time_stamp>`
- On-demand hourly backup: `SnapCenter_LocalSnap_<time_stamp>`
- On-demand daily backup: `SnapCenter_LocalSnapAndSnapVault_<time_stamp>`



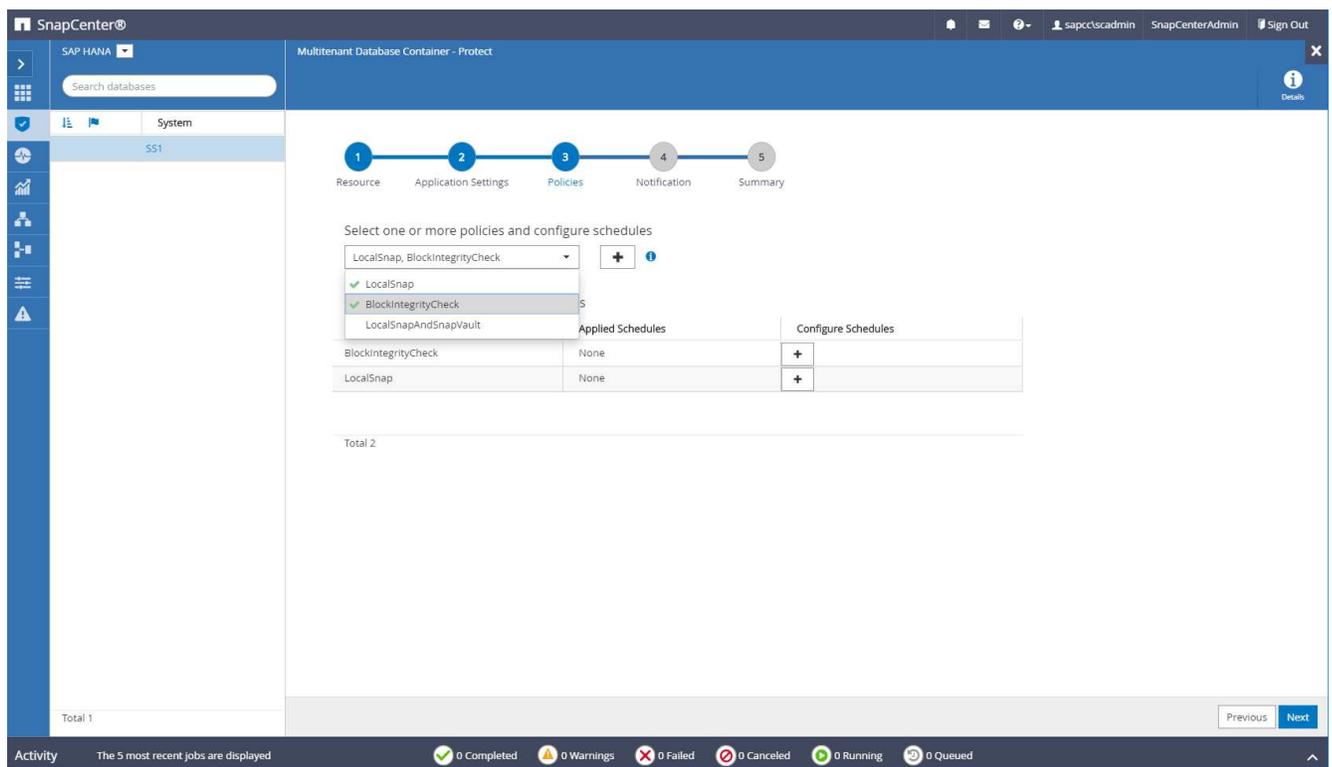
Even though a retention is defined for on-demand backups in the policy configuration, the housekeeping is only done when another on-demand backup is executed. Therefore, on-demand backups must typically be deleted manually in SnapCenter to make sure that these backups are also deleted in the SAP HANA backup catalog and that the log backup housekeeping is not based on an old on-demand backup.



3. No specific setting needs to be made on the Application Settings page. Click Next.



4. Select the policies to add to the resource.



5. Define the schedule for the LocalSnap policy (in this example, every four hours).

Add schedules for policy LocalSnap x

**Hourly**

Start date  🕒

Expires on  📅

Repeat every  hours  mins

**i** The schedules are triggered in the SnapCenter Server time zone. x

6. Define the schedule for the LocalSnapAndSnapVault policy (in this example, once per day).

Modify schedules for policy LocalSnapAndSnapVault ✕

**Daily**

Start date  

Expires on  

Repeat every

**i** The schedules are triggered in the SnapCenter Server time zone. ✕

7. Define the schedule for the block integrity check policy (in this example, once per week).

Add schedules for policy BlockIntegrityCheck ✕

**Weekly**

Start date  

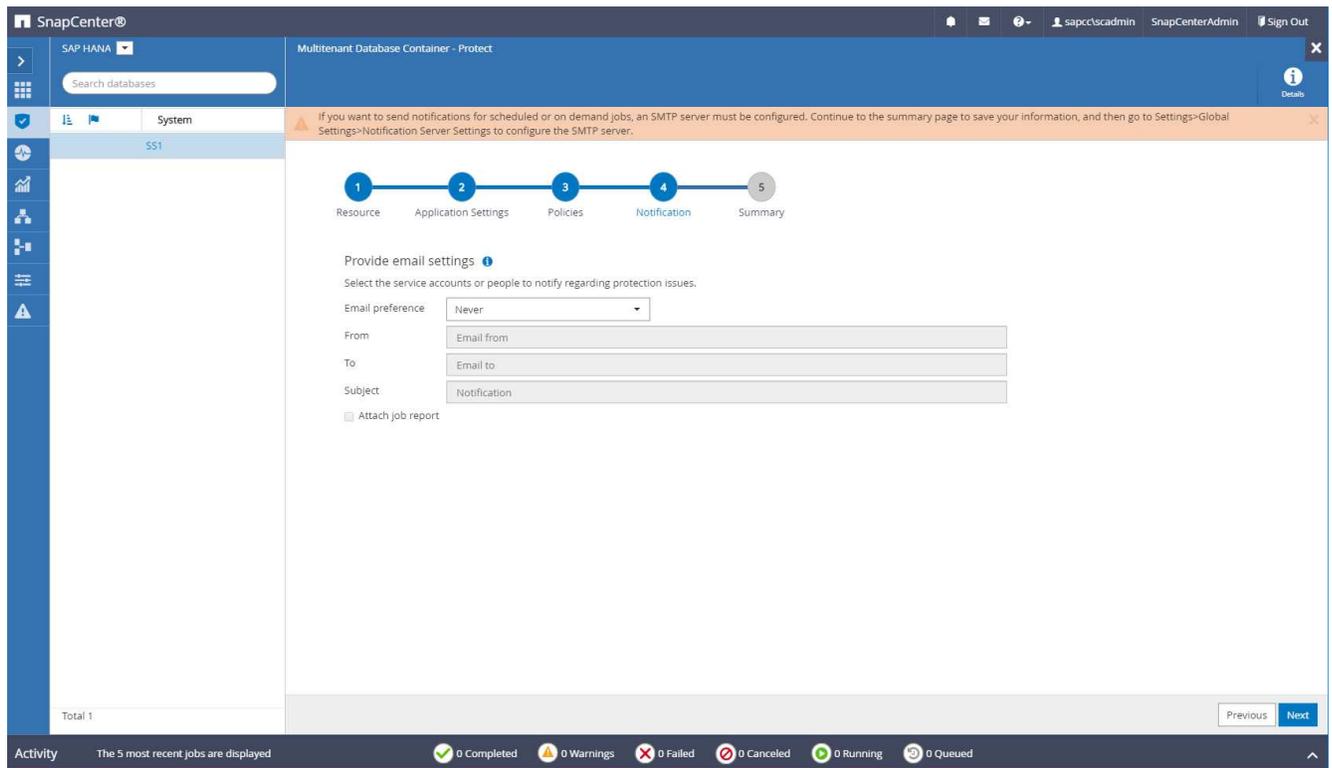
Expires on  

Days  

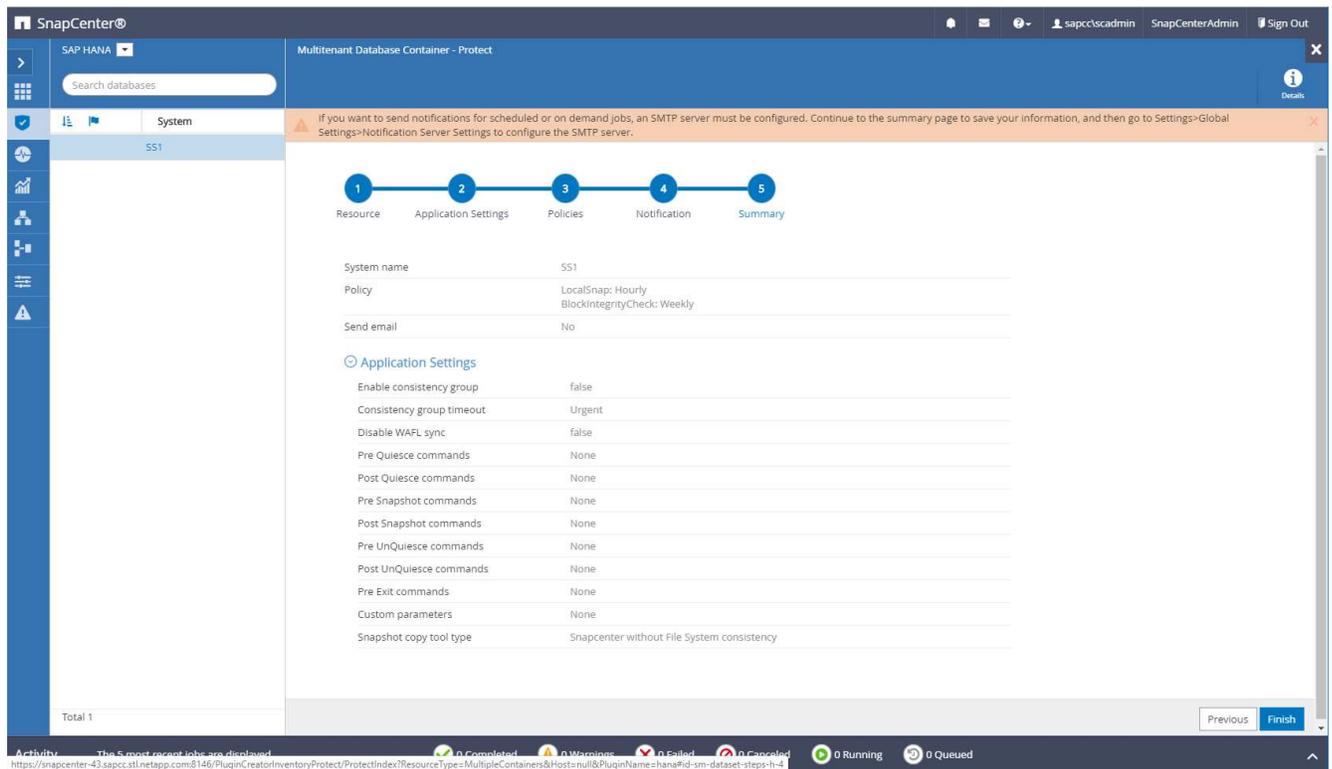
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

 The schedules are triggered in the SnapCenter Server time zone. ✕

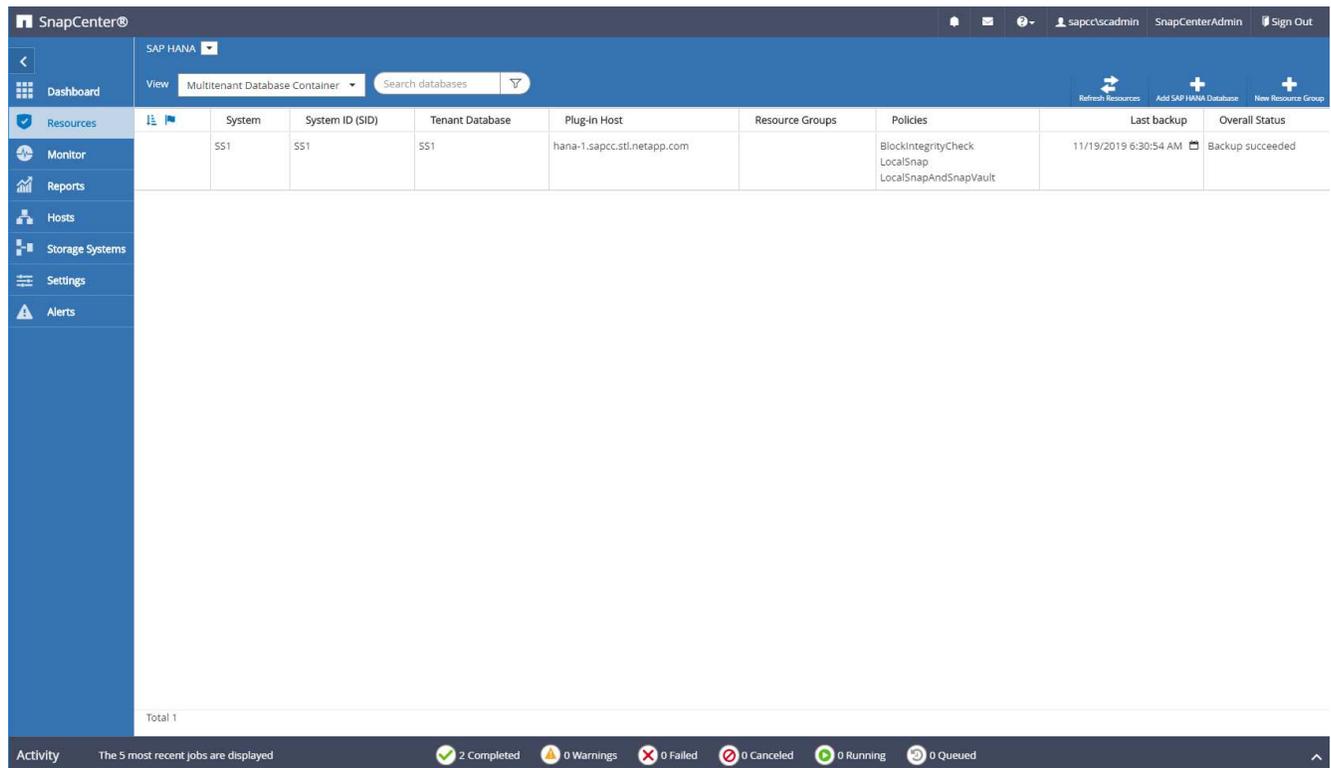
8. Provide information about the email notification.



9. On the Summary page, click Finish.



10. On-demand backups can now be created on the topology page. The scheduled backups are executed based on the configuration settings.



## Additional configuration steps for Fibre Channel SAN environments

Depending on the HANA release and the HANA plug-in deployment, additional configuration steps are required for environments in which the SAP HANA systems are using Fibre Channel and the XFS file system.



These additional configuration steps are only required for HANA resources, which are configured manually in SnapCenter. It is also only required for HANA 1.0 releases and HANA 2.0 releases up to SPS2.

When a HANA backup save point is triggered by SnapCenter in SAP HANA, SAP HANA writes Snapshot ID files for each tenant and database service as a last step (for example, `/hana/data/SID/mnt00001/hdb00001/snapshot_databackup_0_1`). These files are part of the data volume on the storage and are therefore part of the storage Snapshot copy. This file is mandatory when performing a recovery in a situation in which the backup is restored. Due to metadata caching with the XFS file system on the Linux host, the file is not immediately visible at the storage layer. The standard XFS configuration for metadata caching is 30 seconds.



With HANA 2.0 SPS3, SAP changed the write operation of these Snapshot ID files to synchronously so that metadata caching is not a problem.



With SnapCenter 4.3, if the HANA plug-in is deployed on the database host, the Linux plug-in executes a file system flush operation on the host before the storage Snapshot is triggered. In this case, the metadata caching is not a problem.

In SnapCenter, you must configure a `postquiesce` command that waits until the XFS metadata cache is flushed to the disk layer.

The actual configuration of the metadata caching can be checked by using the following command:

```
stlrx300s8-2:/ # sysctl -A | grep xfssyncd_centiseecs
fs.xfs.xfssyncd_centiseecs = 3000
```

NetApp recommends using a wait time that is twice the value of the `fs.xfs.xfssyncd_centiseecs` parameter. Because the default value is 30 seconds, set the sleep command to 60 seconds.

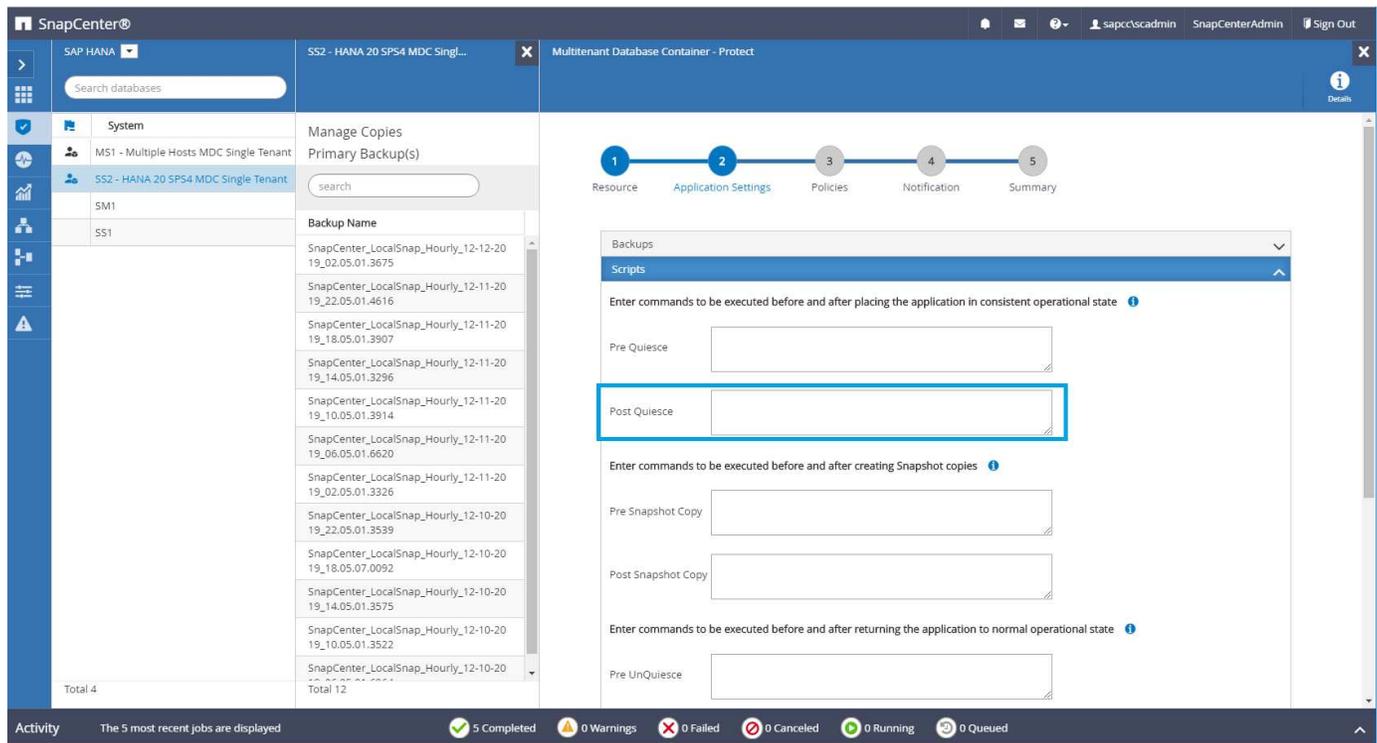
If the SnapCenter server is used as a central HANA plug-in host, a batch file can be used. The batch file must have the following content:

```
@echo off
waitfor AnyThing /t 60 2>NUL
Exit /b 0
```

The batch file can be saved, for example, as `C:\Program Files\NetApp\Wait60Sec.bat`. In the resource protection configuration, the batch file must be added as Post Quiesce command.

If a separate Linux host is used as a central HANA plug-in host, you must configure the command `/bin/sleep 60` as the Post Quiesce command in the SnapCenter UI.

The following figure shows the Post Quiesce command within the resource protection configuration screen.



## SnapCenter resource-specific configuration for non-data volume backups

The backup of non-data volumes is an integrated part of the SAP HANA plug-in. Protecting the database data volume is sufficient to restore and recover the SAP HANA database to a given point in time, provided that the database installation resources and

the required logs are still available.

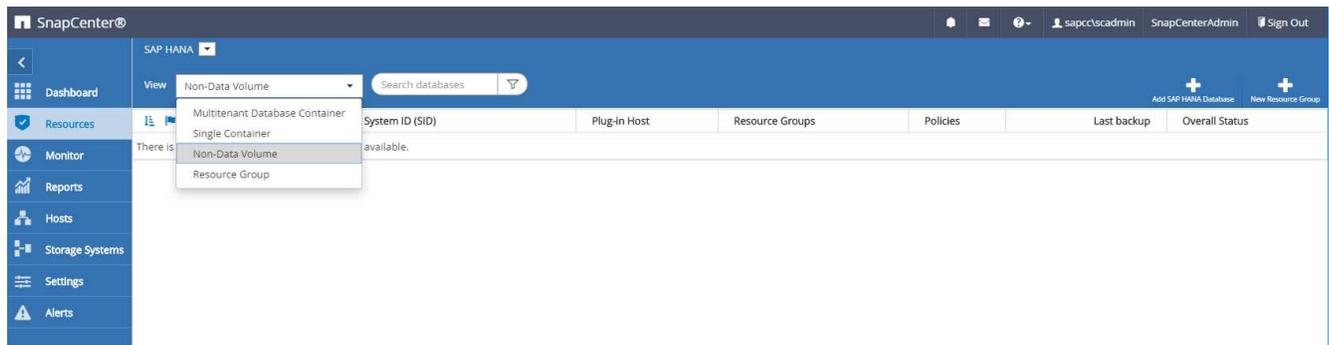
To recover from situations where other non-data files must be restored, NetApp recommends developing an additional backup strategy for non-data volumes to augment the SAP HANA database backup. Depending on your specific requirements, the backup of non-data volumes might differ in scheduling frequency and retention settings, and you should consider how frequently non-data files are changed. For instance, the HANA volume `/hana/shared` contains executables but also SAP HANA trace files. While executables only change when the SAP HANA database is upgraded, the SAP HANA trace files might need a higher backup frequency to support analyzing problem situations with SAP HANA.

SnapCenter non-data volume backup enables Snapshot copies of all relevant volumes to be created in a few seconds with the same space efficiency as SAP HANA database backups. The difference is that there is no SQL communication with SAP HANA database required.

## Configuration of non-data volume resources

In this example, we want to protect the non-data volumes of the SAP HANA database SS1.

1. From the Resource tab, select Non-Data-Volume and click Add SAP HANA Database.



2. In step one of the Add SAP HANA Database dialog, in the Resource Type list, select Non-data Volumes. Specify a name for the resource and the associated SID and the SAP HANA plug-in host you want to use for the resource, then click Next.

Add SAP HANA Database ×

**1 Name**

**2 Storage Footprint**

3 Summary

**Provide Resource Details**

Resource Type	<input type="text" value="Non-data Volumes"/>
Resource Name	<input type="text" value="SS1-Shared-Volume"/>
Associated SID	<input type="text" value="SS1"/> ⓘ
Plug-in Host	<input type="text" value="hana-1.sapcc.stl.netapp.com"/> ⓘ

3. Add the SVM and the storage volume as storage footprint, then click Next.

Add SAP HANA Database

1 Name

2 Storage Footprint

3 Summary

Provide Storage Footprint Details

Add Storage Footprint

Storage System: hana-primary.sapcc.stl.netapp.com

Select one or more volumes and if required their associated Qtrees and LUNs

Volume name: SS1\_shared

LUNs or Qtrees: Default is 'None' or type to find

Save

Previous Next

4. In the summary step, click Finish to save the settings.
5. Repeat these steps for all the required non-data volumes.
6. Continue with the protection configuration of the new resource.



Data protection for a non- data volume resources is identical to the workflow for SAP HANA database resources and can be defined on an individual resource level.

The following figure shows the list of the configured non-data volume resources.

SnapCenter®

SAP HANA

View: Non-Data Volume

Search databases

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		Backup not run

## Resource groups

Resource groups are a convenient way to define the protection of multiple resources that require the same protection policies and schedule. Single resources that are part of a resource group can still be protected on an individual level.

Resource groups provide the following features:

- You can add one or more resources to a resource group. All resources must belong to the same SnapCenter plug-in.
- Protection can be defined on a resource group level. All resources in the resource group use the same policy and schedule when protected.
- All backups in the SnapCenter repository and the storage Snapshot copies have the same name defined in the resource protection.
- Restore operations are applied on a single resource level, not as part of a resource group.
- When using SnapCenter to delete the backup of a resource that was created on a resource group level, this backup is deleted for all resources in the resource group. Deleting the backup includes deleting the backup from the SnapCenter repository as well as deleting the storage Snapshot copies.
- The main use case for resource groups is when a customer wants to use backups created with SnapCenter for system cloning with SAP Landscape Management. This is described in the next section.

## Using SnapCenter together with SAP landscape management

With SAP Landscape Management (SAP LaMa), customers can manage complex SAP system landscapes in on-premises data centers as well as in systems that are running in the cloud. SAP LaMa, together with NetApp Storage Services Connector (SSC), can execute storage operations such as cloning and replication for SAP system clone, copy, and refresh use cases using Snapshot and FlexClone technology. This allows you to completely automate an SAP system copy based on storage cloning technology while also including the required SAP postprocessing. For more details about NetApp's solutions for SAP LaMa, refer to [TR-4018: Integrating NetApp ONTAP Systems with SAP Landscape Management](#).

NetApp SSC and SAP LaMa can create on-demand Snapshot copies directly using NetApp SSC, but they can also utilize Snapshot copies that have been created using SnapCenter. To utilize SnapCenter backups as the basis for system clone and copy operations with SAP LaMa, the following prerequisites must be met:

- SAP LaMa requires that all volumes be included in the backup; this includes SAP HANA data, log and shared volumes.
- All storage Snapshot names must be identical.
- Storage Snapshot names must start with VCM.



In normal backup operations, NetApp does not recommend including the log volume. If you restore the log volume from a backup, it overwrites the last active redo logs and prevents the recovery of the database to the last recent state.

SnapCenter resource groups meet all these requirements. Three resources are configured in SnapCenter: one resource each for the data volume, the log volume, and the shared volume. The resources are put into a resource group, and the protection is then defined on the resource group level. In the resource group protection, the custom Snapshot name must be defined with VCM at the beginning.

## Database backups

In SnapCenter, database backups are typically executed using the schedules defined within the resource protection configuration of each HANA database.

On-demand database backup can be performed by using either the SnapCenter GUI, a PowerShell command line, or REST APIs.

### Identifying SnapCenter backups in SAP HANA Studio

The SnapCenter resource topology shows a list of backups created using SnapCenter. The following figure shows the backups available on the primary storage and highlights the most recent backup.

The screenshot displays the SnapCenter interface for the SS1 topology. On the left, a sidebar shows the system hierarchy: SAP HANA, System, MS1 - Multiple Hosts MDC Single Tenant, SS2 - HANA 20 SP54 MDC Single Tenant, SM1, and SS1. The main area shows 'Manage Copies' with 15 Backups (0 Clones) for Local copies and 5 Backups (0 Clones) for Vault copies. A 'Summary Card' on the right indicates 21 Backups: 20 Snapshot based backups, 1 File-Based backup, and 0 Clones. Below this is a table of 'Primary Backup(s)'. The table has columns for Backup Name, Count, and End Date. The most recent backup is highlighted with a blue box.

Backup Name	Count	End Date
SnapCenter_LocalSnap_Hourly_12-03-2019_02.30.01.5053	1	12/03/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_22.30.01.4925	1	12/02/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_18.30.01.3834	1	12/02/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_14.30.01.3366	1	12/02/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_10.30.01.4510	1	12/02/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-02-2019_08.17.01.9273	1	12/02/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_06.30.01.3164	1	12/02/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_02.30.01.3555	1	12/02/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_12-01-2019_22.30.01.3859	1	12/01/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_12-01-2019_18.30.01.3834	1	12/01/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_12-01-2019_14.30.01.3255	1	12/01/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_12-01-2019_10.30.01.2508	1	12/01/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-01-2019_08.17.01.9654	1	12/01/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_12-01-2019_06.30.01.2968	1	12/01/2019 6:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_11-30-2019_08.17.01.8590	1	11/30/2019 8:17:55 AM

When performing a backup using storage Snapshot copies for an SAP HANA MDC system, a Snapshot copy of the data volume is created. This data volume contains the data of the system database as well as the data of all tenant databases. To reflect this physical architecture, SAP HANA internally performs a combined backup of the system database as well as all tenant databases whenever SnapCenter triggers a Snapshot backup. This results in multiple separate backup entries in the SAP HANA backup catalog: one for the system database and one for each tenant database.



For SAP HANA single-container systems, the database volume contains only the single database, and there is only one entry in SAP HANA's backup catalog.

In the SAP HANA backup catalog, the SnapCenter backup name is stored as a `Comment` field as well as `External Backup ID (EBID)`. This is shown in the following screenshot for the system database and in the screenshot after that for the tenant database SS1. Both figures highlight the SnapCenter backup name stored in the comment field and EBID.



The HANA 2.0 SPS4 (revision 40 and 41) release always shows a backup size of zero for Snapshot-based backups. This was fixed with revision 42. For more information, see the SAP Note <https://launchpad.support.sap.com/#/notes/2795010>.

The screenshot shows the SAP HANA Studio Backup Catalog for SYSTEMDB@SS1. The 'Backup Catalog' tab is active, displaying a table of backup operations. The 'Backup Details' pane on the right shows information for a specific backup, including its ID, status, type, and location. A table in the 'Additional Information' section lists backup details for the host 'hana-1'.

Status	Started	Duration	Size	Backup Type	Destinatio...
Success	Dec 3, 2019 2:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 10:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 6:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 2:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 10:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 8:17:24 ...	00h 00m 13s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 6:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 2:30:24 ...	00h 00m 13s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 10:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 6:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 2:30:24 ...	00h 00m 13s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 10:30:24 ...	00h 00m 13s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 8:17:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 6:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Nov 30, 2019 8:17:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Nov 30, 2019 6:00:04 ...	00h 00m 03s	1.48 GB	Data Backup	File
Success	Nov 29, 2019 8:17:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Nov 28, 2019 8:17:25 ...	00h 00m 13s	0 B	Data Backup	Snapshot

Backup Details:

- ID: 1575369024442
- Status: Successful
- Backup Type: Data Backup
- Destination Type: Snapshot
- Started: Dec 3, 2019 2:30:24 AM (America/Los\_Angeles)
- Finished: Dec 3, 2019 2:30:38 AM (America/Los\_Angeles)
- Duration: 00h 00m 14s
- Size: 0 B
- Throughput: n.a.
- System ID: SnapCenter\_LocalSnap\_Hourly\_12-03-2019\_02.30.01.5053
- Comment: SnapCenter\_LocalSnap\_Hourly\_12-03-2019\_02.30.01.5053
- Location: /hana/data/SS1/mnt00001/

Host	Service	Name	EBID
hana-1	nameserver	hdb00001	SnapCenter_LocalSnap_Hourly_12-03-2019_02.30.01.5053

The screenshot shows the SAP HANA Studio Backup Catalog for SS1. The 'Backup Catalog' tab is active, displaying a table of backup operations. The 'Backup Details' pane on the right shows information for a specific backup, including its ID, status, type, and location. A table in the 'Additional Information' section lists backup details for the host 'hana-1'.

Status	Started	Duration	Size	Backup Type	Destinatio...
Success	Dec 3, 2019 2:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 10:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 6:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 2:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 10:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 8:17:24 ...	00h 00m 13s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 6:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 2, 2019 2:30:24 ...	00h 00m 13s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 10:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 6:30:23 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 2:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 10:30:24 ...	00h 00m 13s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 8:17:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Dec 1, 2019 6:30:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Nov 30, 2019 8:17:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Nov 30, 2019 6:00:10 ...	00h 00m 03s	1.67 GB	Data Backup	File
Success	Nov 29, 2019 8:17:24 ...	00h 00m 14s	0 B	Data Backup	Snapshot
Success	Nov 28, 2019 8:17:25 ...	00h 00m 13s	0 B	Data Backup	Snapshot

Backup Details:

- ID: 1575369024443
- Status: Successful
- Backup Type: Data Backup
- Destination Type: Snapshot
- Started: Dec 3, 2019 2:30:24 AM (America/Los\_Angeles)
- Finished: Dec 3, 2019 2:30:38 AM (America/Los\_Angeles)
- Duration: 00h 00m 14s
- Size: 0 B
- Throughput: n.a.
- System ID: SnapCenter\_LocalSnap\_Hourly\_12-03-2019\_02.30.01.5053
- Comment: SnapCenter\_LocalSnap\_Hourly\_12-03-2019\_02.30.01.5053
- Location: /hana/data/SS1/mnt00001/

Host	Service	Name	EBID
hana-1	indexserver	hdb00003...	SnapCenter_LocalSnap_Hourly_12-03-2019_02.30.01.5053
hana-1	xsengine	hdb00002...	SnapCenter_LocalSnap_Hourly_12-03-2019_02.30.01.5053



SnapCenter is only aware of its own backups. Additional backups created, for example, with SAP HANA Studio, are visible in the SAP HANA catalog but not in SnapCenter.

## Identifying SnapCenter backups on the storage systems

To view the backups on the storage layer, use NetApp OnCommand System Manager and select the database volume in the SVM—Volume view. The lower Snapshot Copies tab displays the Snapshot copies of the volume. The following screenshot shows the available backups for the database volume SS1\_data\_mnt00001 at the primary storage. The highlighted backup is the backup shown in SnapCenter and SAP HANA Studio in the previous images and has the same naming convention.

The screenshot shows the OnCommand System Manager interface for the volume SS1\_data\_mnt00001. The 'Snapshots Copies' tab is active, displaying a table of snapshot copies. The highlighted row is:

Status	State	Snapshot Name	Date Time	Total Size	Application Dependency
Normal	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_12-01-2019_08.17.01.9654	Dec/01/2019 11:03:44	106.27 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-02-2019_06.30.01.3164	Dec/02/2019 09:16:42	74.76 MB	None
Normal	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_12-02-2019_08.17.01.9273	Dec/02/2019 11:03:43	17.21 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-02-2019_10.30.01.4510	Dec/02/2019 13:16:42	39.11 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-02-2019_14.30.01.3366	Dec/02/2019 17:16:42	87.53 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-02-2019_18.30.01.3834	Dec/02/2019 21:16:41	95.67 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-02-2019_22.30.01.4925	Dec/03/2019 01:16:41	29.86 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-03-2019_02.30.01.5053	Dec/03/2019 05:16:41	43.81 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-03-2019_06.30.01.4088	Dec/03/2019 09:16:40	49.46 MB	None
Normal	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_12-03-2019_08.17.01.9180	Dec/03/2019 11:03:41	77.14 MB	snapmirror
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-03-2019_10.30.01.4554	Dec/03/2019 13:16:40	42.12 MB	None
Normal	-NA-	SnapCenter_LocalSnap_Hourly_12-03-2019_14.30.01.3902	Dec/03/2019 17:16:40	57.42 MB	None

The following screenshot shows the available backups for the replication target volume hana\_SA1\_data\_mnt00001\_dest at the secondary storage system.

The screenshot shows the OnCommand System Manager interface for the volume SS1\_data\_mnt00001\_dest. The 'Snapshots Copies' tab is active, displaying a table of snapshot copies. The highlighted row is:

Status	State	Snapshot Name	Date Time	Total Size	Application Dependency
Normal	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_11-29-2019_08.17.01.8567	Nov/29/2019 11:03:48	113.34 MB	None
Normal	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_11-30-2019_08.17.01.8590	Nov/30/2019 11:03:46	87.69 MB	None
Normal	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_12-01-2019_08.17.01.9654	Dec/01/2019 11:03:44	108.67 MB	None
Normal	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_12-02-2019_08.17.01.9273	Dec/02/2019 11:03:43	102 MB	None
Busy	-NA-	SnapCenter_LocalSnapAndSnapVault_Daily_12-03-2019_08.17.01.9180	Dec/03/2019 11:03:41	176 KB	busy

## On-demand database backup at primary storage

1. In the resource view, select the resource and double-click the line to switch to the topology view.

The resource topology view provides an overview of all available backups that have been created using SnapCenter. The top area of this view displays the backup topology, showing the backups on the primary storage (local copies) and, if available, on the off-site backup storage (vault copies).

Backup Name	Count	End Date
SnapCenter_LocalSnap_Hourly_12-03-2019_02:30:01.5053	1	12/03/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_22:30:01.4925	1	12/02/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_18:30:01.3834	1	12/02/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_14:30:01.3366	1	12/02/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_10:30:01.4510	1	12/02/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-02-2019_08:17:01.9273	1	12/02/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_06:30:01.3164	1	12/02/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_02:30:01.3555	1	12/02/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_12-01-2019_22:30:01.3859	1	12/01/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_12-01-2019_18:30:01.3834	1	12/01/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_12-01-2019_14:30:01.3255	1	12/01/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_12-01-2019_10:30:01.2508	1	12/01/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-01-2019_08:17:01.9654	1	12/01/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_12-01-2019_06:30:01.2968	1	12/01/2019 6:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_11-30-2019_08:17:01.8590	1	11/30/2019 8:17:55 AM

2. In the top row, select the Back up Now icon to start an on-demand backup. From the drop-down list, select the backup policy LocalSnap and then click Backup to start the on-demand backup.

### Backup x

Create a backup for the selected resource

Resource Name

Policy  i

This starts the backup job. A log of the previous five jobs is shown in the Activity area below the topology view. When the backup is finished, a new entry is shown in the topology view. The backup names follow the same naming convention as the Snapshot name defined in the section [“Resource protection configuration.”](#)



You must close and reopen the topology view to see the updated backup list.

The screenshot displays the SAP Center web interface for managing SAP HANA backups. The main content area is titled 'Manage Copies' and shows a diagram with 'Local copies' (16 Backups, 0 Clones) and 'Vault copies' (5 Backups, 0 Clones). A 'Summary Card' on the right provides a high-level overview: 22 Backups, 21 Snapshot based backups, 1 File Based backup, and 0 Clones.

Below the diagram is a table of 'Primary Backup(s)'. The first row is highlighted with a blue box:

Backup Name	Count	IF	End Date
SnapCenter_LocalSnap_12-03-2019_06:37:50.1491	1		12/03/2019 6:38:44 AM
SnapCenter_LocalSnap_Hourly_12-03-2019_06:30:01.4088	1		12/03/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_12-03-2019_02:30:01.5053	1		12/03/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_22:30:01.4925	1		12/02/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_18:30:01.3834	1		12/02/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_14:30:01.3366	1		12/02/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_12-02-2019_10:30:01.4510	1		12/02/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-02-2019_08:17:01.9273	1		12/02/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_06:30:01.3164	1		12/02/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_12-02-2019_02:30:01.3555	1		12/02/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_12-01-2019_22:30:01.3859	1		12/01/2019 10:30:55 PM
Total 4			
Total 16			

At the bottom, the 'Activity' section shows the 5 most recent jobs, all of which are 'Completed':

- 2 minutes ago: Backup of Resource Group 'hana-1\_sapcc\_stl\_netapp\_com\_hana\_MDC\_S51' with policy 'LocalSnap' - Completed
- 10 minutes ago: Backup of Resource Group 'hana-1\_sapcc\_stl\_netapp\_com\_hana\_MDC\_S51' with policy 'LocalSnap' - Completed
- 12 minutes ago: Backup of Resource Group 'hana-2\_sapcc\_stl\_netapp\_com\_hana\_MDC\_SM1' with policy 'LocalSnap' - Completed
- 35 minutes ago: Backup of Resource Group 'SnapCenter-43\_sapcc\_stl\_netapp\_com\_hana\_MDC\_S52' with policy 'LocalSnap' - Completed
- 3 hours ago: Backup of Resource Group 'SnapCenter-43\_sapcc\_stl\_netapp\_com\_hana\_MDC\_MS1' with policy 'LocalSnap' - Completed

3. The job details are shown when clicking the job's activity line in the Activity area. You can open a detailed job log by clicking View Logs.

Job Details
✕

Backup of Resource Group 'hana-1\_sapcc\_stl\_netapp\_com\_hana\_MDC\_SS1' with policy 'LocalSnap'

- ✓ ▾ Backup of Resource Group 'hana-1\_sapcc\_stl\_netapp\_com\_hana\_MDC\_SS1' with policy 'LocalSnap'
- ✓ ▾ hana-1.sapcc.stl.netapp.com
- ✓ ▾ Backup
- ✓ ▶ Validate Dataset Parameters
- ✓ ▶ Validate Plugin Parameters
- ✓ ▶ Complete Application Discovery
- ✓ ▶ Initialize Filesystem Plugin
- ✓ ▶ Discover Filesystem Resources
- ✓ ▶ Validate Retention Settings
- ✓ ▶ Quiesce Application
- ✓ ▶ Quiesce Filesystem
- ✓ ▶ Create Snapshot
- ✓ ▶ UnQuiesce Filesystem
- ✓ ▶ UnQuiesce Application
- ✓ ▶ Get Snapshot Details
- ✓ ▶ Get Filesystem Meta Data
- ✓ ▶ Finalize Filesystem Plugin
- ✓ ▶ Collect Autosupport data
- ✓ ▶ Register Backup and Apply Retention
- ✓ ▶ Register Snapshot attributes

**i** Task Name: Backup Start Time: 12/03/2019 6:37:51 AM End Time: 12/03/2019 6:39:03 AM

View Logs
Cancel Job
Close

4. In SAP HANA Studio, the new backup is visible in the backup catalog. The same backup name in SnapCenter is also used in the comment and the EBID field in the backup catalog.

#### On-demand database backups with SnapVault replication

1. In the resource view, select the resource and double-click the line to switch to the topology view.
2. In the top row, select the Backup Now icon to start an on-demand backup. From the drop-down list, select the backup policy LocalSnapAndSnapVault, then click Backup to start the on-demand backup.

Backup ×

Create a backup for the selected resource

Resource Name

Policy  i

3. The job details are shown when clicking the job's activity line in the Activity area.

Job Details x

Backup of Resource Group 'hana-1\_sapcc\_stl\_netapp\_com\_hana\_MDC\_SS1' with policy 'LocalSnapAndSnapVault'

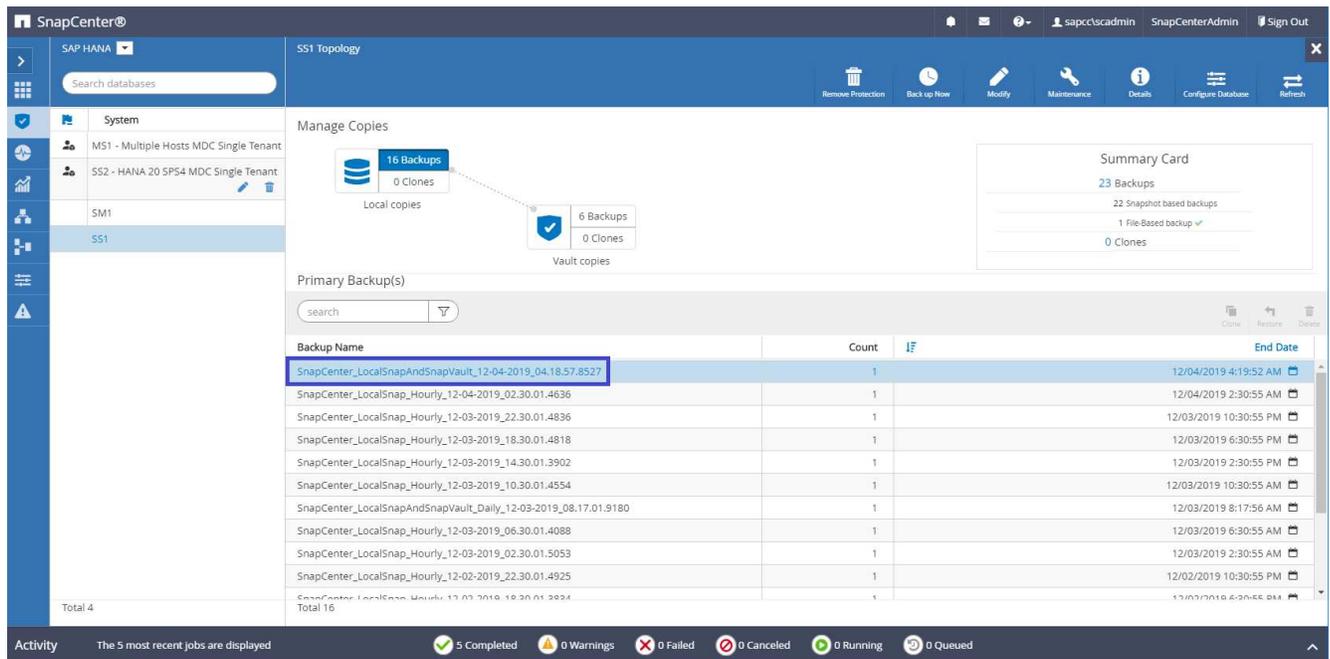
- ✓ ▶ Quiesce Application
- ✓ ▶ Quiesce Filesystem
- ✓ ▶ Create Snapshot
- ✓ ▶ UnQuiesce Filesystem
- ✓ ▶ UnQuiesce Application
- ✓ ▶ Get Snapshot Details
- ✓ ▶ Get Filesystem Meta Data
- ✓ ▶ Finalize Filesystem Plugin
- ✓ ▶ Collect Autosupport data
- ✓ ▶ Secondary Update
- ✓ ▶ Register Backup and Apply Retention
- ✓ ▶ Register Snapshot attributes
- ✓ ▶ Application Clean-Up
- ✓ ▶ Data Collection
- ✓ ▶ Agent Finalize Workflow
- ✓ ▶ ( Job 1031 ) SnapVault update

**i** Task Name: ( Job 1031 ) SnapVault update Start Time: 12/04/2019 4:19:55 AM End Time: 12/04/2019 4:20:55 AM

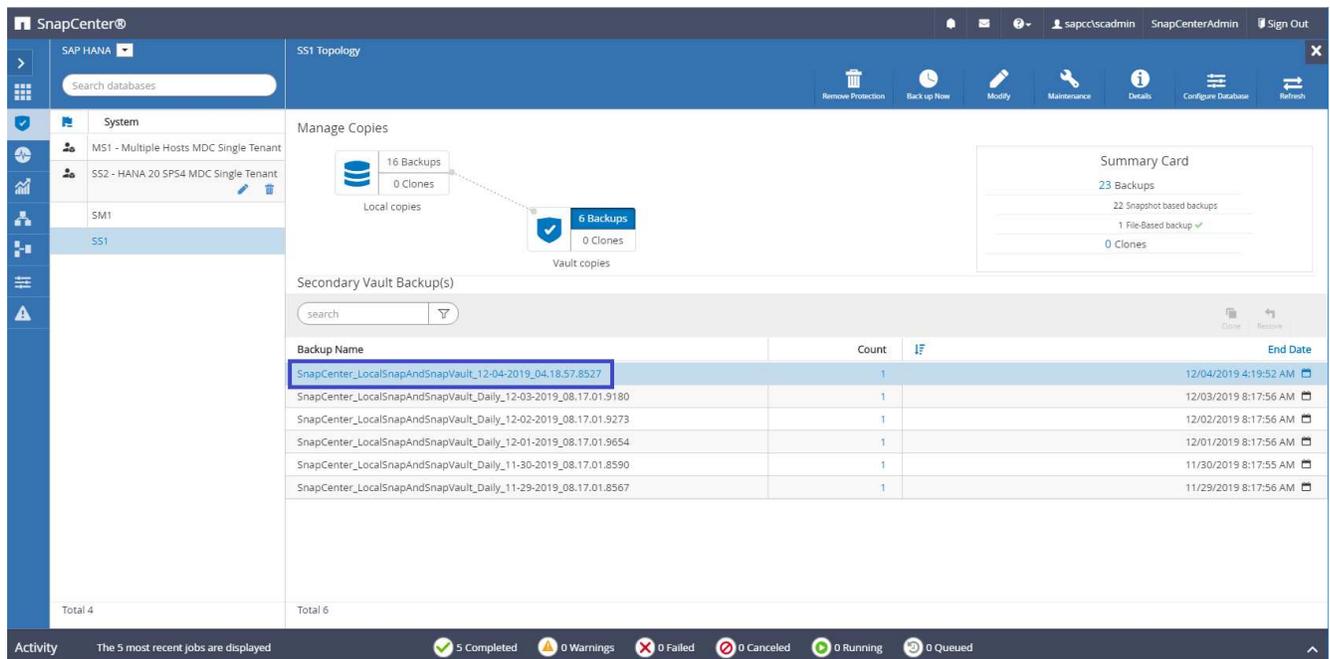
4. When the backup is finished, a new entry is shown in the topology view. The backup names follow the same naming convention as the Snapshot name defined in the section [“Resource protection configuration.”](#)



You must close and reopen the topology view to see the updated backup list.



5. By selecting Vault copies, backups at the secondary storage are shown. The name of the replicated backup is identical to the backup name at the primary storage.



6. In SAP HANA Studio, the new backup is visible in the backup catalog. The same backup name in SnapCenter is also used in the comment and the EBID field in the backup catalog.

## Block integrity check

SAP recommends combining storage-based Snapshot backups with a weekly file-based backup to execute a block integrity check. SnapCenter supports the execution of a block integrity check by using a policy in which file-based backup is selected as the backup type.

When scheduling backups using this policy, SnapCenter creates a standard SAP HANA file backup for the system and tenant databases.

SnapCenter does not display the block integrity check in the same manner as Snapshot copy-based backups. Instead, the summary card shows the number of file-based backups and the status of the previous backup.

The screenshot shows the SnapCenter interface for managing SAP HANA backups. The main content area is titled 'Manage Copies' and shows a visual representation of backup copies: 15 Local copies and 5 Vault copies. A 'Summary Card' is highlighted with a blue box, displaying the following information:

- 22 Backups
- 20 Snapshot based backups
- 2 File-Based backups
- Last Backup 11/23/2019 6:00:59 AM
- Backup succeeded

Below the summary card is a table of backup names, counts, and end dates:

Backup Name	Count	End Date
SnapCenter_LocalSnap_Hourly_11-28-2019_06:30:01.1132	1	11/28/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_11-28-2019_02:30:01.1496	1	11/28/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_11-27-2019_22:30:01.1582	1	11/27/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_11-27-2019_18:30:01.0949	1	11/27/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_11-27-2019_14:30:01.1670	1	11/27/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_11-27-2019_10:30:01.0579	1	11/27/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_11-27-2019_08:17:01.9215	1	11/27/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_11-27-2019_06:30:01.0767	1	11/27/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_11-27-2019_02:30:01.1788	1	11/27/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_11-26-2019_22:30:01.0413	1	11/26/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_11-26-2019_18:30:01.0738	1	11/26/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_11-26-2019_14:30:01.0340	1	11/26/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_11-26-2019_10:30:01.0649	1	11/26/2019 10:30:55 AM
SnapCenter_LocalSnapAndSnapVault_Daily_11-26-2019_08:17:01.8979	1	11/26/2019 8:17:56 AM
Total 15		

At the bottom of the interface, the activity bar shows the following status:

- 5 Completed
- 0 Warnings
- 0 Failed
- 0 Canceled
- 0 Running
- 0 Queued

A block integrity check backup cannot be deleted using the SnapCenter UI, but it can be deleted using PowerShell commands.

```
PS C:\Users\scadmin> Get-SmBackupReport -Resource SS1
SmBackupId           : 9
SmJobId              : 42
StartDateTime        : 11/19/2019 8:26:32 AM
EndDateTime          : 11/19/2019 8:27:33 AM
Duration              : 00:01:00.7652030
CreatedDateTime      : 11/19/2019 8:27:24 AM
Status                : Completed
ProtectionGroupName  : hana-1_sapcc_stl_netapp_com_hana_MDC_SS1
SmProtectionGroupId  : 1
PolicyName           : BlockIntegrityCheck
SmPolicyId           : 5
BackupName           : SnapCenter_BlockIntegrityCheck_11-19-
2019_08.26.33.2913
VerificationStatus   : NotApplicable
VerificationStatuses :
SmJobError           :
BackupType           : SCC_BACKUP
CatalogingStatus     : NotApplicable
CatalogingStatuses  :
ReportDataCreatedDateTime :
PluginCode           : SCC
PluginName           : hana
JobTypeId            : 0
JobHost              :
```

```
PS C:\Users\scadmin> Remove-SmBackup -BackupIds 9
```

```
Remove-SmBackup
```

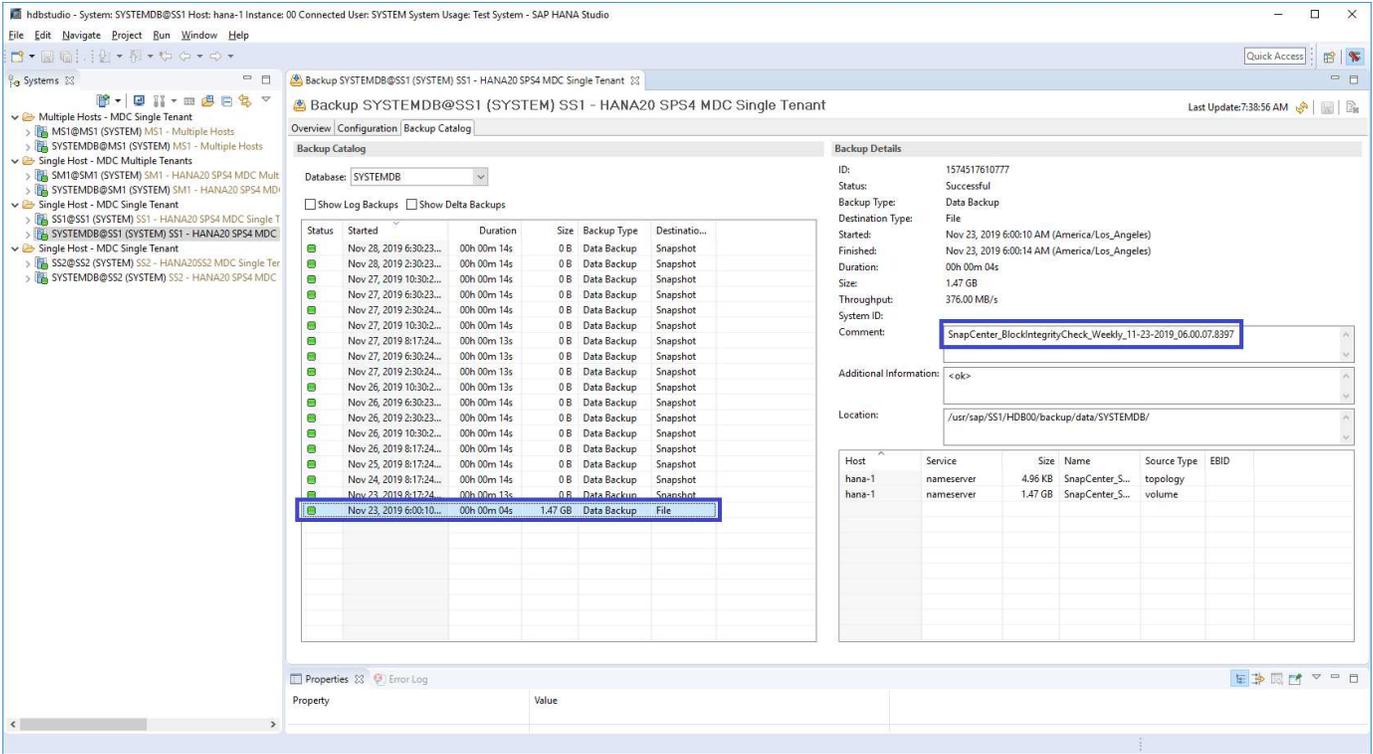
```
Are you sure want to remove the backup(s).
```

```
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help
(default is "Y"): y
```

```
BackupResult : {}
Result       : SMCOREContracts.SMResult
TotalCount   : 0
DisplayCount : 0
Context      :
Job          : SMCOREContracts.SmJob
```

```
PS C:\Users\scadmin>
```

The SAP HANA backup catalog shows entries for both the system and the tenant databases. The following figure shows a SnapCenter block integrity check in the backup catalog of the system database.



A successful block integrity check creates standard SAP HANA data backup files. SnapCenter uses the backup path that has been configured in the HANA database for file-based data backup operations.

```

hana-1:/usr/sap/SS1/HDB00/backup/data # ls -al *
DB_SS1:
total 1710840
drwxr-xr-- 2 ssladm sapsys      4096 Nov 28 10:25 .
drwxr-xr-- 4 ssladm sapsys      4096 Nov 19 05:11 ..
-rw-r----- 1 ssladm sapsys    155648 Nov 23 08:46
SnapCenter_SnapCenter_BlockIntegrityCheck_Weekly_11-23-
2019_06.00.07.8397_databackup_0_1
-rw-r----- 1 ssladm sapsys    83894272 Nov 23 08:46
SnapCenter_SnapCenter_BlockIntegrityCheck_Weekly_11-23-
2019_06.00.07.8397_databackup_2_1
-rw-r----- 1 ssladm sapsys 1660952576 Nov 23 08:46
SnapCenter_SnapCenter_BlockIntegrityCheck_Weekly_11-23-
2019_06.00.07.8397_databackup_3_1
SYSTEMDB:
total 1546340
drwxr-xr-- 2 ssladm sapsys      4096 Nov 28 10:24 .
drwxr-xr-- 4 ssladm sapsys      4096 Nov 19 05:11 ..
-rw-r----- 1 ssladm sapsys    159744 Nov 23 08:46
SnapCenter_SnapCenter_BlockIntegrityCheck_Weekly_11-23-
2019_06.00.07.8397_databackup_0_1
-rw-r----- 1 ssladm sapsys 1577066496 Nov 23 08:46
SnapCenter_SnapCenter_BlockIntegrityCheck_Weekly_11-23-
2019_06.00.07.8397_databackup_1_1

```

## Restore and recovery

The following sections describe the restore and recovery workflows of three different scenarios and example configurations.

- Automated restore and recovery:
  - Auto discovered HANA system SS1
  - SAP HANA single host, MDC single tenant system using NFS
- Single-tenant restore and recovery:
  - Auto discovered HANA system SM1
  - SAP HANA single host, MDC multiple tenant system using NFS
- Restore with manual recovery:
  - Manual configured HANA system SS2
  - SAP HANA single host, MDC multiple tenant system using NFS

In the following sections, the differences between SAP HANA single host and multiple hosts and Fibre Channel SAN attached HANA systems are highlighted.

The examples show SAP HANA Studio as a tool to execute manual recovery. You can also use SAP HANA

Cockpit or HANA SQL statements.

## Automated restore and recovery

With SnapCenter 4.3, automated restore and recovery operations are supported for HANA single container or MDC single tenant systems that have been auto discovered by SnapCenter.

You can execute an automated restore and recovery operation with the following steps:

1. Select the backup to be used for the restore operation. The backup can be selected from the following storage options:
  - Primary storage
  - Offsite backup storage (SnapVault target)
2. Select the restore type. Select Complete Restore with Volume Revert or without Volume Revert.



The Volume Revert option is only available for restore operations from primary storage and if the HANA database is using NFS as the storage protocol.

3. Select the recovery type from the following options:
  - To most recent state
  - Point in time
  - To specific data backup
  - No recovery



The selected recovery type is used for the recovery of the system and the tenant database.

Next, SnapCenter performs the following operations:

1. It stops the HANA database.
2. It restores the database.

Depending on the selected restore type and the used storage protocol, different operations are executed.

- If NFS and Volume Revert are selected, then SnapCenter unmounts the volume, restores the volume using volume-based SnapRestore on the storage layer, and mounts the volume.
  - If NFS is selected and Volume Revert is not selected, SnapCenter restores all files using single-file SnapRestore operations on the storage layer.
  - If Fibre Channel SAN is selected, then SnapCenter unmounts the LUN(s), restores the LUN(s) using single file SnapRestore operations on the storage layer, and discovers and mounts the LUN(s).
3. It recovers the database:
    - a. It recovers the system database.
    - b. It recovers the tenant database.

Or, for HANA single container systems, the recovery is executed in a single step:

- c. It starts the HANA database.



If No Recovery is selected, SnapCenter exits and the recovery operation for the system and the tenant database must be done manually.

This section provides the steps for the automated restore and recovery operation of the auto discovered HANA system SS1 (SAP HANA single host, MDC single tenant system using NFS).

1. Select a backup in SnapCenter to be used for the restore operation.



You can select restore from primary or from offsite backup storage.

**Primary Backup(s)**

Backup Name	Count	IF	End Date
SnapCenter_LocalSnap_Hourly_12-05-2019_22.30.01.5385	1		12/05/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_12-05-2019_18.30.01.5244	1		12/05/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_12-05-2019_14.30.01.6022	1		12/05/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_12-05-2019_10.30.01.5450	1		12/05/2019 10:30:56 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-05-2019_08.17.02.0191	1		12/05/2019 8:17:56 AM
SnapCenter_LocalSnap_Hourly_12-05-2019_06.30.01.5487	1		12/05/2019 6:30:55 AM
SnapCenter_LocalSnap_Hourly_12-05-2019_02.30.01.5470	1		12/05/2019 2:30:55 AM
SnapCenter_LocalSnap_Hourly_12-04-2019_22.30.01.5182	1		12/04/2019 10:30:55 PM
SnapCenter_LocalSnap_Hourly_12-04-2019_18.30.01.5249	1		12/04/2019 6:30:55 PM
SnapCenter_LocalSnap_Hourly_12-04-2019_14.30.01.5069	1		12/04/2019 2:30:55 PM
SnapCenter_LocalSnap_Hourly_12-04-2019_10.30.01.5300	1		12/04/2019 10:30:55 AM
Total 16			

**Secondary Vault Backup(s)**

Backup Name	Count	IF	End Date
SnapCenter_LocalSnapAndSnapVault_Daily_12-05-2019_08.17.02.0191	1		12/05/2019 8:17:56 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-04-2019_08.17.01.9976	1		12/04/2019 8:17:56 AM
SnapCenter_LocalSnapAndSnapVault_12-04-2019_04.18.57.8527	1		12/04/2019 4:19:52 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-03-2019_08.17.01.9180	1		12/03/2019 8:17:56 AM
SnapCenter_LocalSnapAndSnapVault_Daily_12-02-2019_08.17.01.9273	1		12/02/2019 8:17:56 AM
Total 5			

2. Select the restore scope and type.

The following three screenshots show the restore options for restore from primary with NFS, restore from

secondary with NFS, and restore from primary with Fibre Channel SAN.

The restore type options for restore from primary storage.



The Volume Revert option is only available for restore operations from primary with NFS.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.30.01.5385

1 Restore scope

2 Recovery scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Select the restore types

- Complete Resource ?
  - Volume Revert
- Tenant Database

⚠ As part of Complete Resource restore, if a resource contains volumes as Storage Footprint, then the latest Snapshot copies on such volumes will be deleted permanently. Also, if there are other resources hosted on the same volumes, then it will result in data loss for such resources.

⚠ The newer tenants added on the host after the backup was created cannot be restored and will be lost after restore operation.

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

Previous Next

The restore type options for restore from offsite backup storage.

Restore from SnapCenter\_LocalSnapAndSnapVault\_Daily\_12-05-2019\_08.17.02.0191

1 Restore scope

2 Recovery scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Select the restore types

- Complete Resource ?
- Tenant Database

Choose archive location

hana-primary.sapcc.stf.netapp.com:SS1\_data\_mre00001

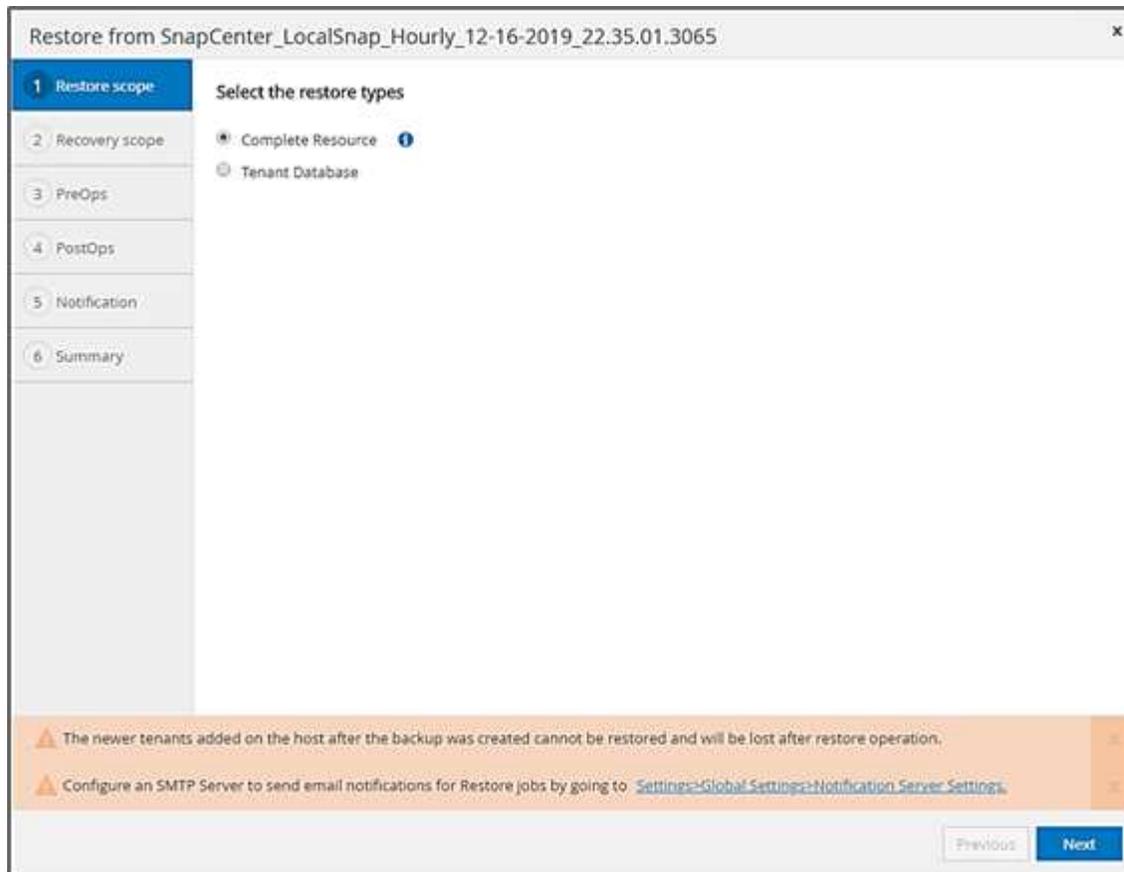
hana-backup.sapcc.stf.netapp.com:SS1\_dat

⚠ The newer tenants added on the host after the backup was created cannot be restored and will be lost after restore operation.

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

Previous Next

The restore type options for restore from primary storage with Fibre Channel SAN.



3. Select Recovery Scope and provide the location for log backup and catalog backup.



SnapCenter uses the default path or the changed paths in the HANA global.ini file to pre-populate the log and catalog backup locations.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.30.01.5385 x

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

### Recover database files using

- Recover to most recent state i
- Recover to point in time i
- Recover to specified data backup i
- No recovery i

### Specify log backup locations i

+ Add

### Specify backup catalog location i

⚠ Recovery options are applicable to both system database and tenant database. x

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

Previous Next

4. Enter the optional prerestore commands.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.30.01.5385 ×

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps**
- 4 PostOps
- 5 Notification
- 6 Summary

Enter optional commands to run before performing a restore operation ?

Pre restore command

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

Previous Next

5. Enter the optional post-restore commands.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.30.01.5385 ×

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps**
- 5 Notification
- 6 Summary

Enter optional commands to run after performing a restore operation ⓘ

Post restore command

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

Previous Next

6. Enter the optional email settings.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.30.01.5385 ×

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification**
- 6 Summary

**Provide email settings** ⓘ

Email preference:

From:

To:

Subject:

Attach Job Report

**⚠** If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server. ×

7. To start the restore operation, click Finish.

x
Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.30.01.5385

<div style="background-color: #0070c0; color: white; padding: 5px; border-radius: 5px; margin-bottom: 5px;">1 Restore scope</div> <div style="background-color: #0070c0; color: white; padding: 5px; border-radius: 5px; margin-bottom: 5px;">2 Recovery scope</div> <div style="background-color: #0070c0; color: white; padding: 5px; border-radius: 5px; margin-bottom: 5px;">3 PreOps</div> <div style="background-color: #0070c0; color: white; padding: 5px; border-radius: 5px; margin-bottom: 5px;">4 PostOps</div> <div style="background-color: #0070c0; color: white; padding: 5px; border-radius: 5px; margin-bottom: 5px;">5 Notification</div> <div style="background-color: #0070c0; color: white; padding: 5px; border-radius: 5px; margin-bottom: 5px;"><b>6 Summary</b></div>	<h3>Summary</h3> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Backup Name</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;">SnapCenter_LocalSnap_Hourly_12-05-2019_22.30.01.5385</td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Backup date</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;">12/05/2019 10:30:55 PM</td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Restore scope</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Complete Resource with Volume Revert</td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Recovery scope</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Recover to most recent state</td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Log backup locations</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;">/mnt/log-backup</td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Backup catalog location</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;">/mnt/log-backup</td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Pre restore command</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;"></td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Post restore command</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;"></td> </tr> <tr> <td style="border-bottom: 1px solid #ccc; padding: 5px;">Send email</td> <td style="border-bottom: 1px solid #ccc; padding: 5px;">No</td> </tr> </table>	Backup Name	SnapCenter_LocalSnap_Hourly_12-05-2019_22.30.01.5385	Backup date	12/05/2019 10:30:55 PM	Restore scope	Complete Resource with Volume Revert	Recovery scope	Recover to most recent state	Log backup locations	/mnt/log-backup	Backup catalog location	/mnt/log-backup	Pre restore command		Post restore command		Send email	No
Backup Name	SnapCenter_LocalSnap_Hourly_12-05-2019_22.30.01.5385																		
Backup date	12/05/2019 10:30:55 PM																		
Restore scope	Complete Resource with Volume Revert																		
Recovery scope	Recover to most recent state																		
Log backup locations	/mnt/log-backup																		
Backup catalog location	/mnt/log-backup																		
Pre restore command																			
Post restore command																			
Send email	No																		

▲ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server. x

Previous
Finish

8. SnapCenter executes the restore and recovery operation. This example shows the job details of the restore and recovery job.

Restore 'hana-1.sapcc.stl.netapp.com\hana\MDC\SS1'

- ✓ ▼ Restore 'hana-1.sapcc.stl.netapp.com\hana\MDC\SS1'
- ✓ ▼ hana-1.sapcc.stl.netapp.com
  - ✓ ▼ Restore
    - ✓ ▼ Validate Plugin Parameters
    - ✓ ▼ Pre Restore Application
      - ▶ Stopping HANA instance
    - ✓ ▼ Filesystem Pre Restore
      - ▶ Determining the restore mechanism
      - ▶ Deporting file systems and associated entities
    - ▶ Restore Filesystem
    - ✓ ▼ Filesystem Post Restore
      - ▶ Building file systems and associated entities
    - ✓ ▼ Recover Application
      - ▶ Recovering system database
      - ▶ Checking HDB services status
      - ▶ Recovering tenant database 'SS1'
      - ▶ Starting HANA instance
    - ▶ Clear Catalog on Server
    - ▶ Application Clean-Up
    - ▶ Data Collection
    - ▶ Agent Finalize Workflow

**i** Task Name: Recover Application Start Time: 12/06/2019 7:26:11 AM End Time: 12/06/2019 7:28:46 AM

[View Logs](#)[Cancel Job](#)[Close](#)

### Single-tenant restore and recovery operation

With SnapCenter 4.3, single-tenant restore operations are supported for HANA MDC systems with a single tenant or with multiple tenants that have been auto-discovered by SnapCenter.

You can perform a single-tenant restore and recovery operation with the following steps:

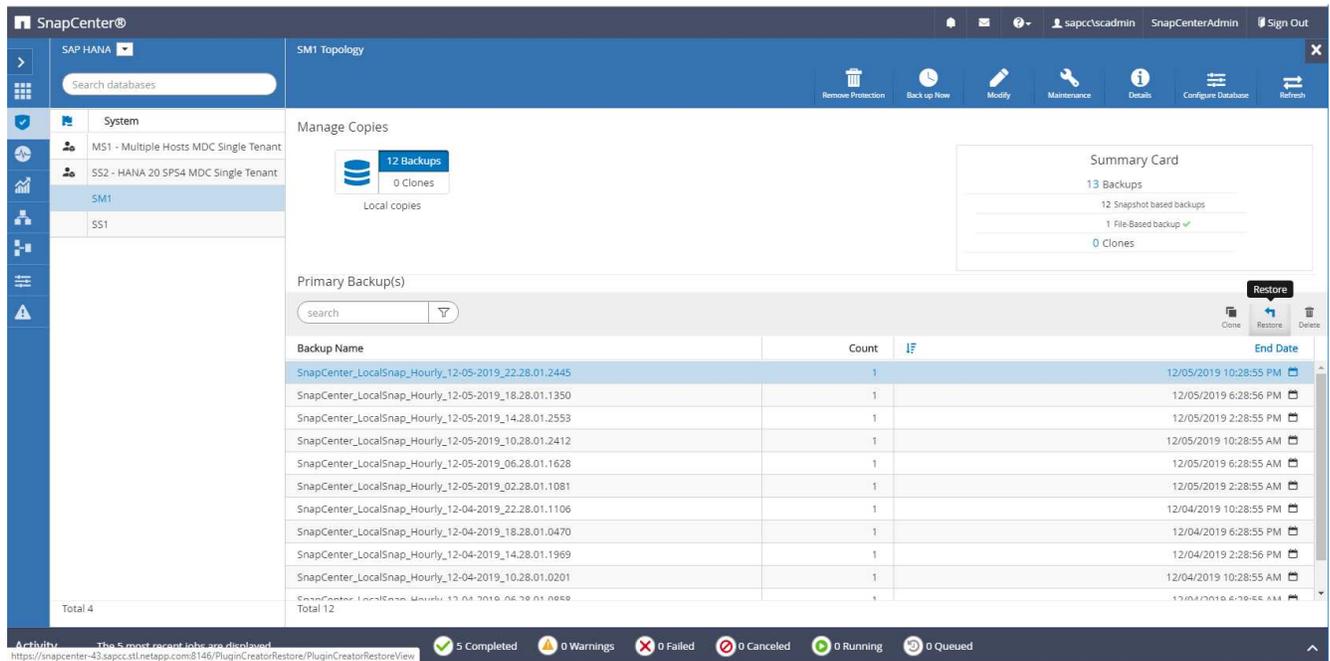
1. Stop the tenant to be restored and recovered.
2. Restore the tenant with SnapCenter.
  - For a restore from primary storage, SnapCenter executes the following operations:
    - **NFS.** Storage Single File SnapRestore operations for all files of the tenant database.
    - **SAN.** Clone and connect the LUN to the database host, and copy all files of the tenant database.
  - For a restore from secondary storage, SnapCenter executes the following operations:
    - **NFS.** Storage SnapVault Restore operations for all files of the tenant database
    - **SAN.** Clone and connect the LUN to the database host, and copy all files of the tenant database
3. Recover the tenant with HANA Studio, Cockpit, or SQL statement.

This section provides the steps for the restore and recovery operation from the primary storage of the auto-discovered HANA system SM1 (SAP HANA single-host, MDC multiple-tenant system using NFS). From the user input perspective, the workflows are identical for a restore from secondary or a restore in a Fibre Channel SAN setup.

1. Stop the tenant database.

```
smladm@hana-2:/usr/sap/SM1/HDB00> hdbsql -U SYSKEY
Welcome to the SAP HANA Database interactive terminal.
Type:  \h for help with commands
       \q to quit
hdbsql=>
hdbsql SYSTEMDB=> alter system stop database tenant2;
0 rows affected (overall time 14.215281 sec; server time 14.212629 sec)
hdbsql SYSTEMDB=>
```

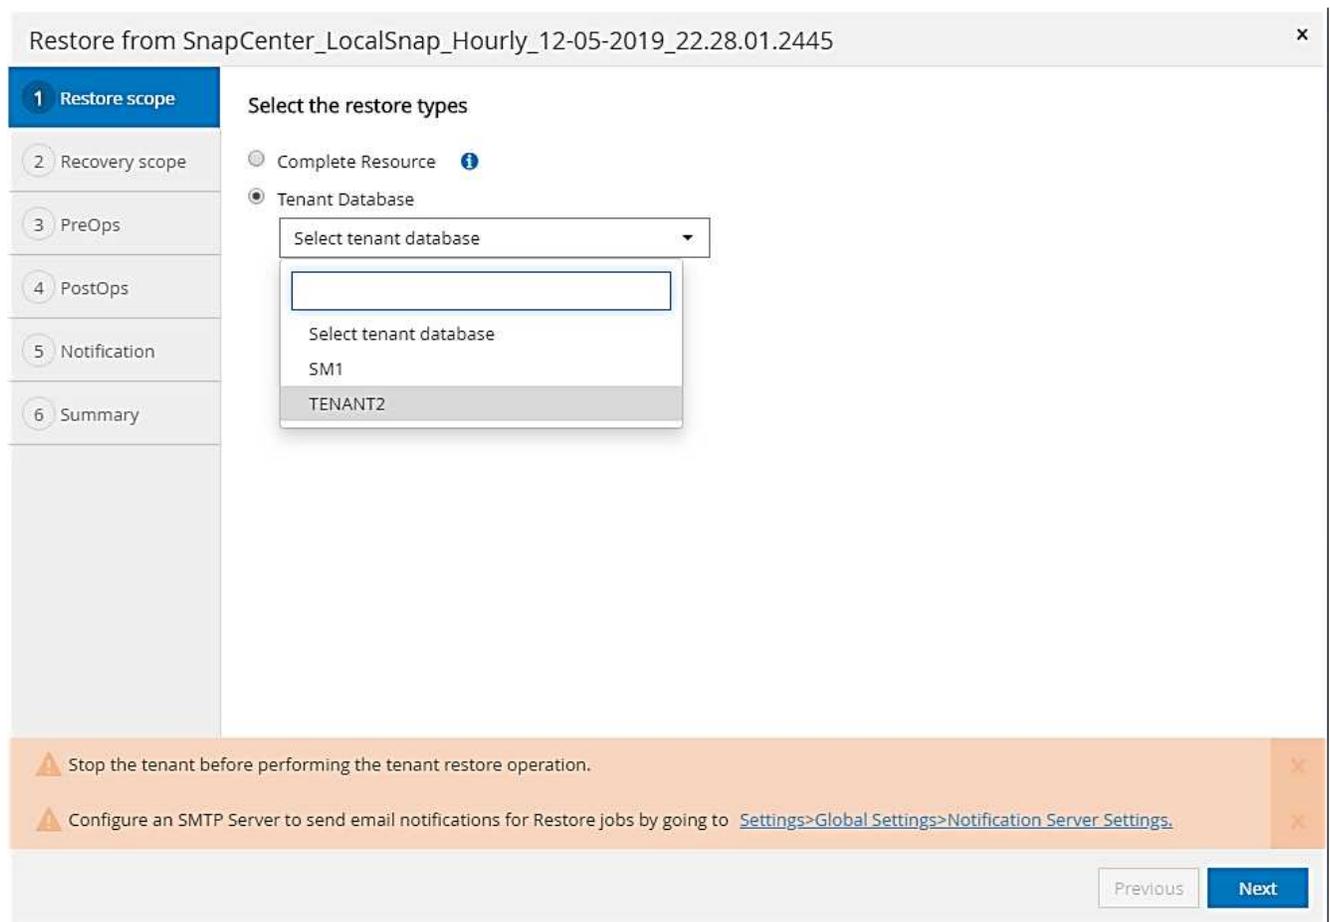
2. Select a backup in SnapCenter to be used for the restore operation.



3. Select the tenant to be restored.



SnapCenter shows a list of all tenants that are included in the selected backup.



Single-tenant recovery is not supported with SnapCenter 4.3. No Recovery is preselected and cannot be changed.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.28.01.2445 ×

- 1 Restore scope
- 2 Recovery scope**
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

**Recover database files using**

- Recover to most recent state ⓘ
- Recover to point in time ⓘ
- Recover to specified data backup ⓘ
- No recovery ⓘ

**Recovery scope**

Recovery of an multitenant database container with multiple tenants is not supported ×

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

Previous Next

4. Enter the optional prerestore commands.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.28.01.2445 x

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps**
- 4 PostOps
- 5 Notification
- 6 Summary

Enter optional commands to run before performing a restore operation ?

Pre restore command

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

Previous Next

5. Enter optional post-restore commands.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.28.01.2445 x

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps**
- 5 Notification
- 6 Summary

**Enter optional commands to run after performing a restore operation** ⓘ

Post restore command

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

Previous Next

6. Enter the optional email settings.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.28.01.2445 ✕

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification**
- 6 Summary

**Provide email settings** ⓘ

Email preference:

From:

To:

Subject:

Attach Job Report

**⚠** If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server. ✕

7. To start the restore operation, click Finish.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.28.01.2445 x

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

### Summary

Backup Name	SnapCenter_LocalSnap_Hourly_12-05-2019_22.28.01.2445
Backup date	12/05/2019 10:28:55 PM
Restore scope	Restore tenant database 'TENANT2'
Recovery scope	No recovery
Pre restore command	
Post restore command	
Send email	No

▲ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server. x

Previous
Finish

The restore operation is executed by SnapCenter. This example shows the job details of the restore job.

## Restore 'hana-2.sapcc.stl.netapp.com\hana\MDC\SM1'

✓ ▼ Restore 'hana-2.sapcc.stl.netapp.com\hana\MDC\SM1'

✓ ▼ hana-2.sapcc.stl.netapp.com

✓ ▼ Restore

✓ ▶ Validate Plugin Parameters

✓ ▶ Pre Restore Application

✓ ▶ Filesystem Pre Restore

✓ ▶ Restore Filesystem

✓ ▶ Filesystem Post Restore

✓ ▶ Recover Application

✓ ▶ Application Clean-Up

✓ ▶ Data Collection

✓ ▶ Agent Finalize Workflow

**i** Task Name: Restore Start Time: 12/06/2019 1:10:40 AM End Time: 12/06/2019 1:12:04 AM

View Logs

Cancel Job

Close



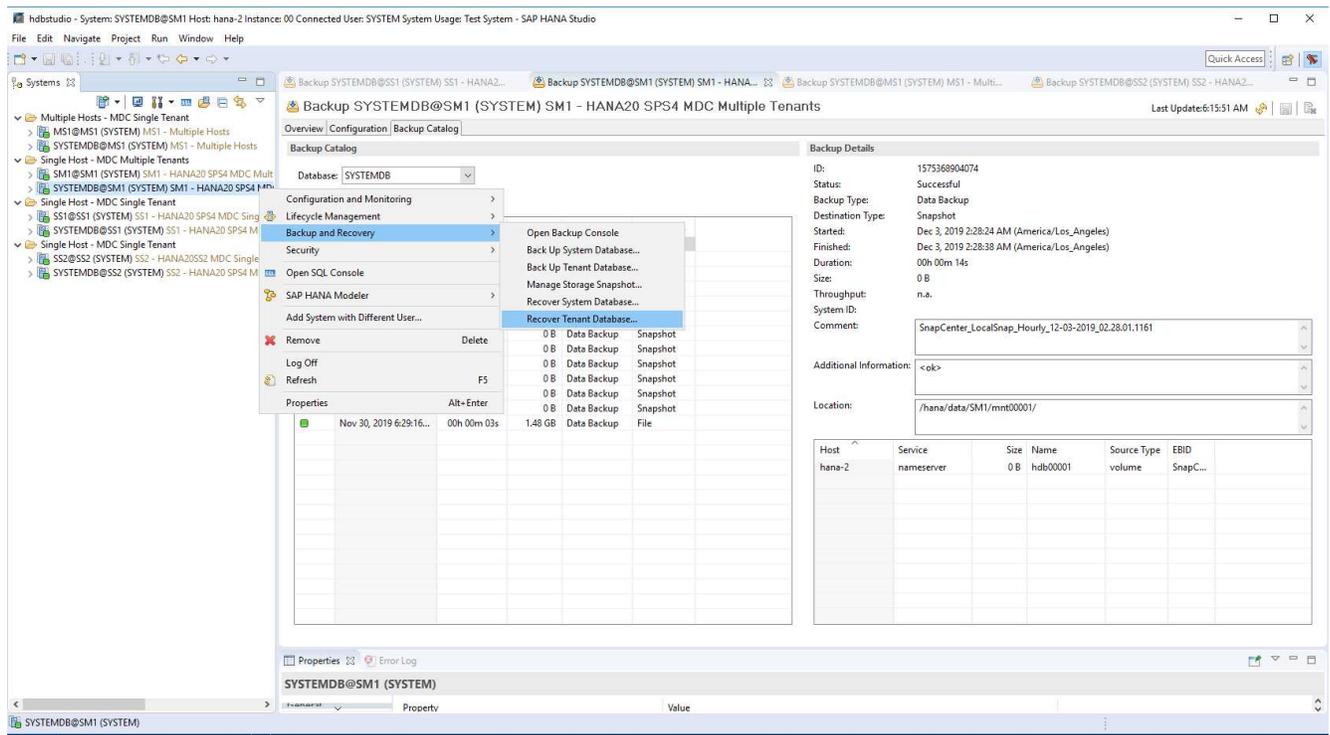
When the tenant restore operation is finished, only the tenant relevant data is restored. On the file system of the HANA database host, the restored data file and the Snapshot backup ID file of the tenant is available.

```

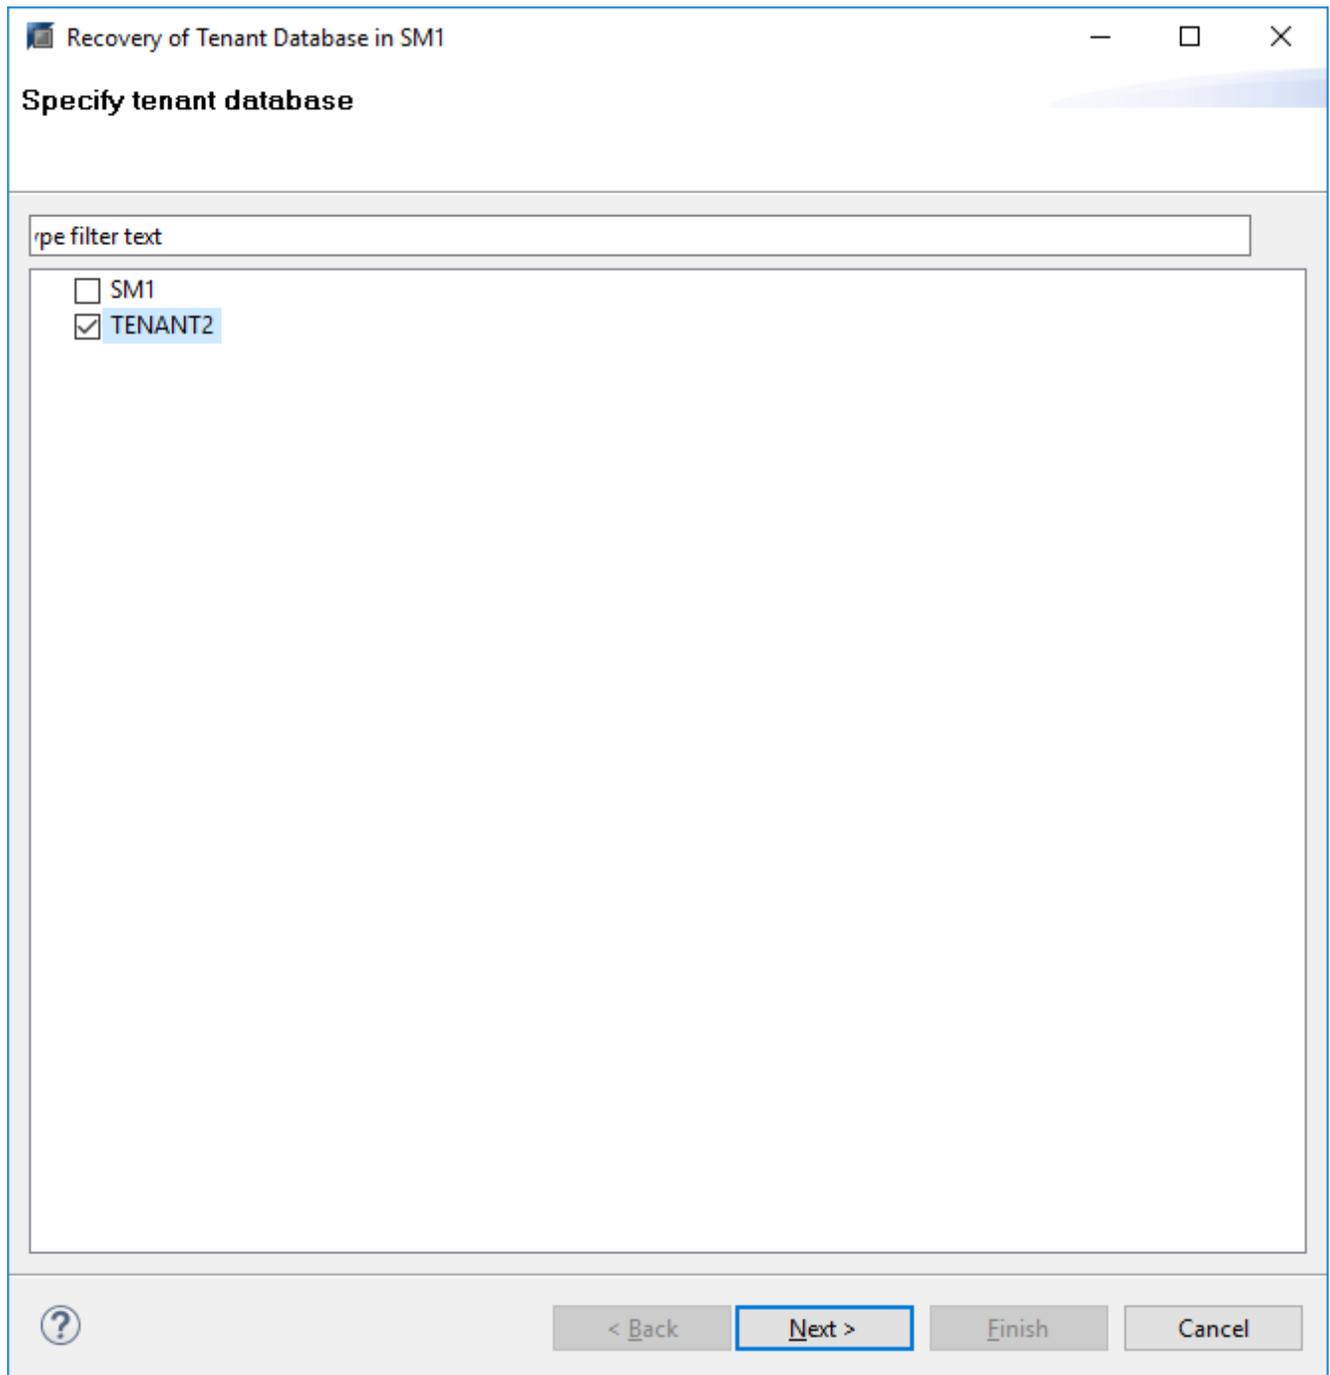
smladm@hana-2:/usr/sap/SM1/HDB00> ls -al /hana/data/SM1/mnt00001/*
-rw-r--r-- 1 smladm sapsys 17 Dec 6 04:01
/hana/data/SM1/mnt00001/nameserver.lck
/hana/data/SM1/mnt00001/hdb00001:
total 3417776
drwxr-x--- 2 smladm sapsys 4096 Dec 6 01:14 .
drwxr-x--- 6 smladm sapsys 4096 Nov 20 09:35 ..
-rw-r----- 1 smladm sapsys 3758096384 Dec 6 03:59 datavolume_0000.dat
-rw-r----- 1 smladm sapsys 0 Nov 20 08:36
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r----- 1 smladm sapsys 36 Nov 20 08:37 landscape.id
/hana/data/SM1/mnt00001/hdb00002.00003:
total 67772
drwxr-xr-- 2 smladm sapsys 4096 Nov 20 08:37 .
drwxr-x--- 6 smladm sapsys 4096 Nov 20 09:35 ..
-rw-r--r-- 1 smladm sapsys 201441280 Dec 6 03:59 datavolume_0000.dat
-rw-r--r-- 1 smladm sapsys 0 Nov 20 08:37
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
/hana/data/SM1/mnt00001/hdb00002.00004:
total 3411836
drwxr-xr-- 2 smladm sapsys 4096 Dec 6 03:57 .
drwxr-x--- 6 smladm sapsys 4096 Nov 20 09:35 ..
-rw-r--r-- 1 smladm sapsys 3758096384 Dec 6 01:14 datavolume_0000.dat
-rw-r--r-- 1 smladm sapsys 0 Nov 20 09:35
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r----- 1 smladm sapsys 155648 Dec 6 01:14
snapshot_databackup_0_1
/hana/data/SM1/mnt00001/hdb00003.00003:
total 3364216
drwxr-xr-- 2 smladm sapsys 4096 Dec 6 01:14 .
drwxr-x--- 6 smladm sapsys 4096 Nov 20 09:35 ..
-rw-r--r-- 1 smladm sapsys 3758096384 Dec 6 03:59 datavolume_0000.dat
-rw-r--r-- 1 smladm sapsys 0 Nov 20 08:37
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
smladm@hana-2:/usr/sap/SM1/HDB00>

```

## 8. Start the recovery with HANA Studio.



9. Select the tenant.



10. Select the recovery type.

Recovery of Tenant Database in SM1

### Specify Recovery Type

Select a recovery type.

Recover the database to its most recent state <sup>1</sup>

Recover the database to the following point in time <sup>1</sup>

Date:   Time:

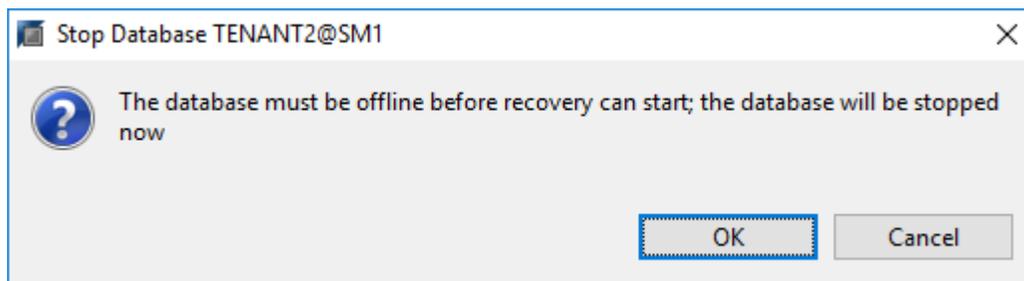
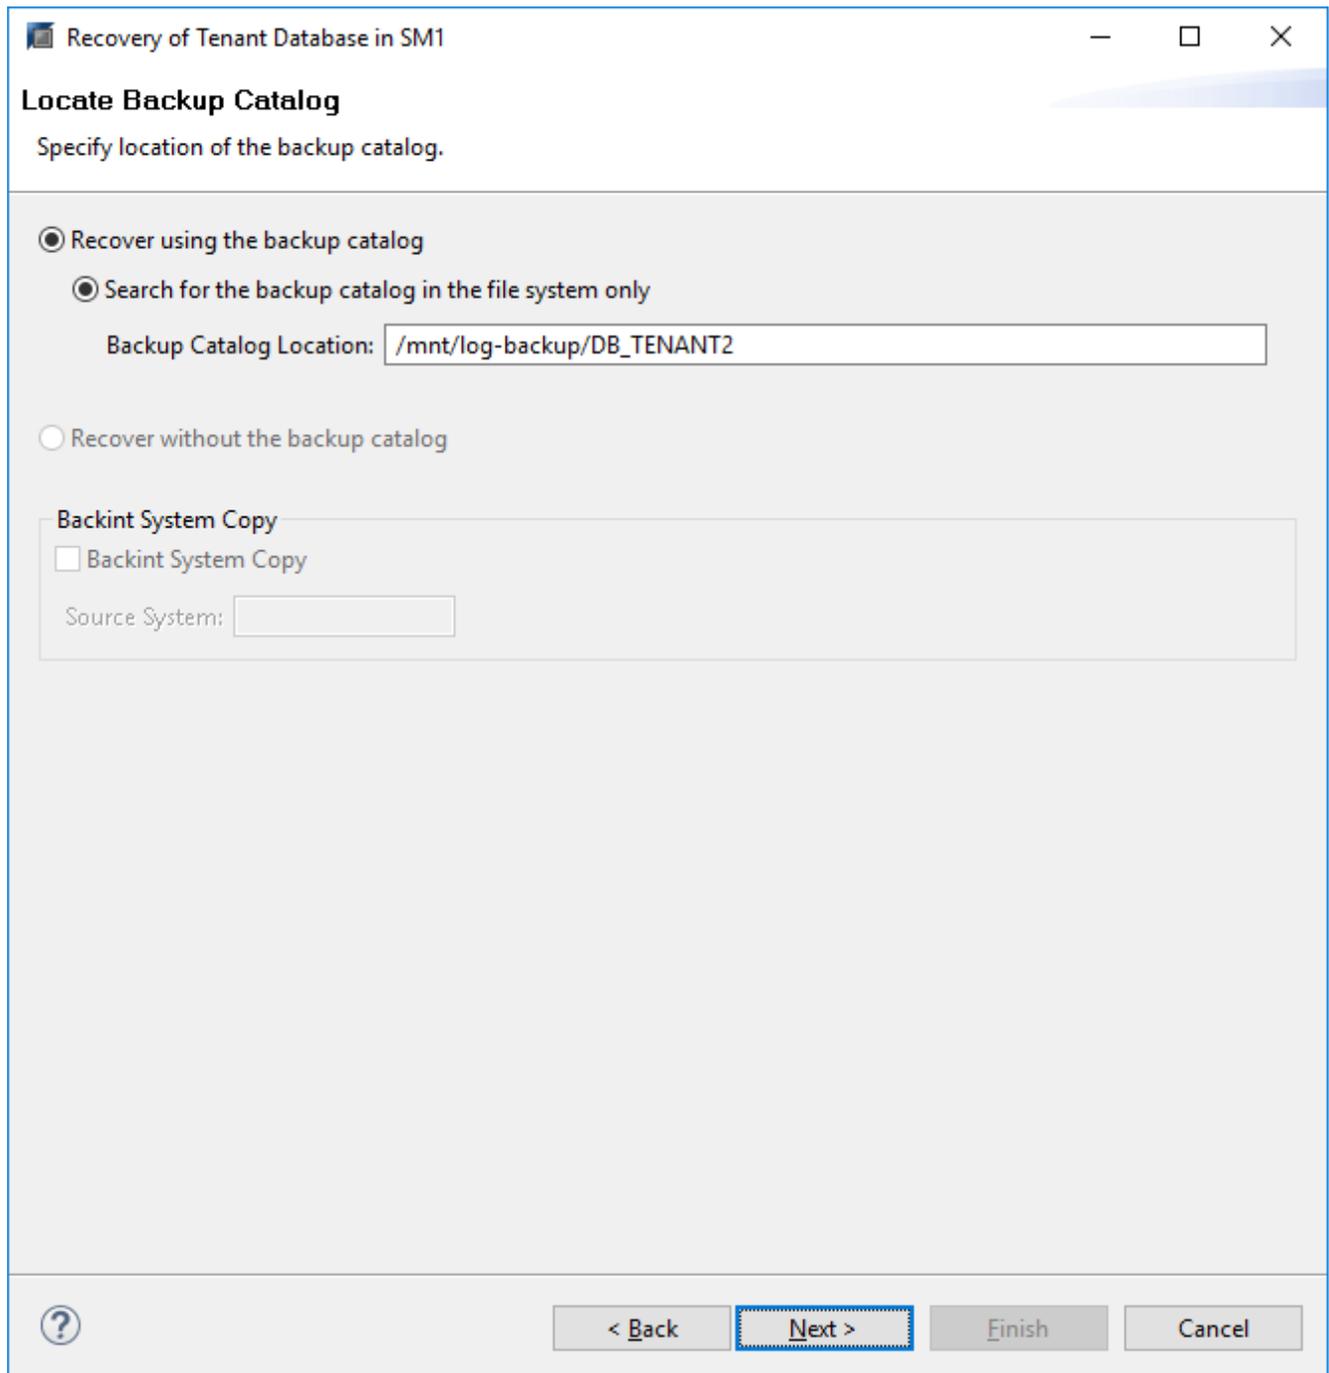
Select Time Zone:  

 System Time Used (GMT): 2019-12-06 09:18:31

Recover the database to a specific data backup <sup>1</sup>



11. Provide the backup catalog location.



Within the backup catalog, the restored backup is highlighted with a green icon. The external backup ID shows the backup name that was previously selected in SnapCenter.

12. Select the entry with the green icon and click Next.

Recovery of Tenant Database in SM1

### Select a Backup

Select a backup to recover the SAP HANA database

**Selected Point in Time**  
Database will be recovered to its most recent state.

**Backups**  
The overview shows backups that were recorded in the backup catalog as successful. The backup at the top is estimated to have the shortest recovery time.

Start Time	Location	Backup Prefix	A...
2019-12-05 22:28:24	/hana/data/SM1	SNAPSHOT	●
2019-12-05 18:28:24	/hana/data/SM1	SNAPSHOT	⊗
2019-12-05 14:28:23	/hana/data/SM1	SNAPSHOT	⊗
2019-12-05 10:28:24	/hana/data/SM1	SNAPSHOT	⊗
2019-12-05 06:28:23	/hana/data/SM1	SNAPSHOT	⊗
2019-12-05 02:28:23	/hana/data/SM1	SNAPSHOT	⊗
2019-12-04 22:28:24	/hana/data/SM1	SNAPSHOT	⊗
2019-12-04 18:28:23	/hana/data/SM1	SNAPSHOT	⊗
2019-12-04 14:28:25	/hana/data/SM1	SNAPSHOT	⊗
2019-12-04 10:28:24	/hana/data/SM1	SNAPSHOT	⊗

Refresh Show More

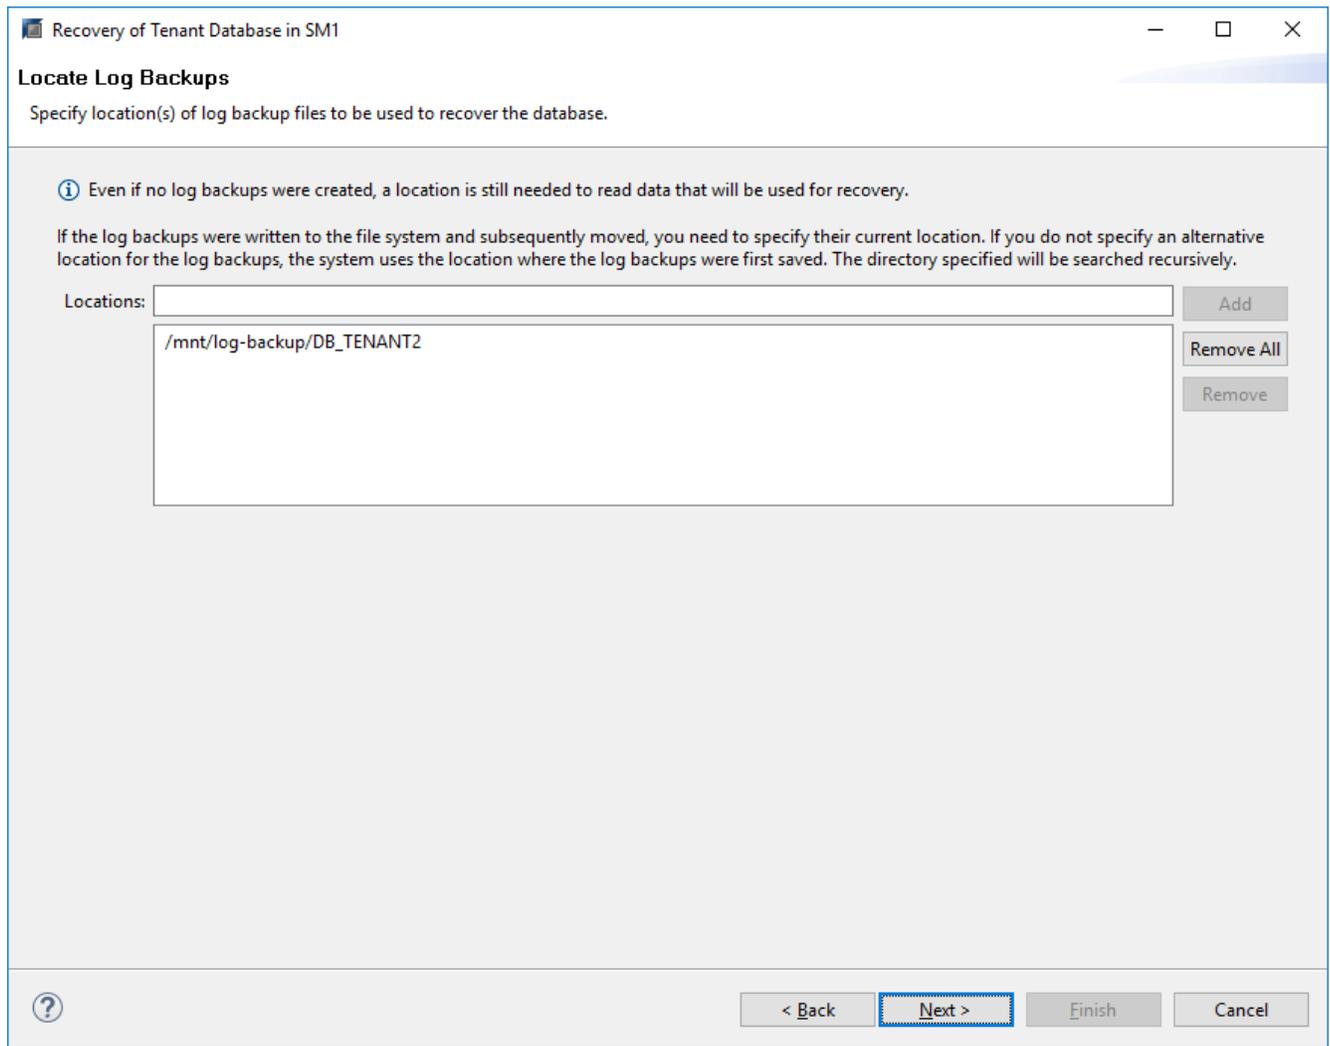
**Details of Selected Item**

Start Time: 2019-12-05 22:28:24 Destination Type: SNAPSHOT Source System: TENANT2@SM1  
 Size: 0 B Backup ID: 1575613704345 External Backup ID: SnapCenter\_LocalSnap\_Hourly\_12-05-2019\_22.28.01.2445  
 Backup Name: /hana/data/SM1  
 Alternative Location:

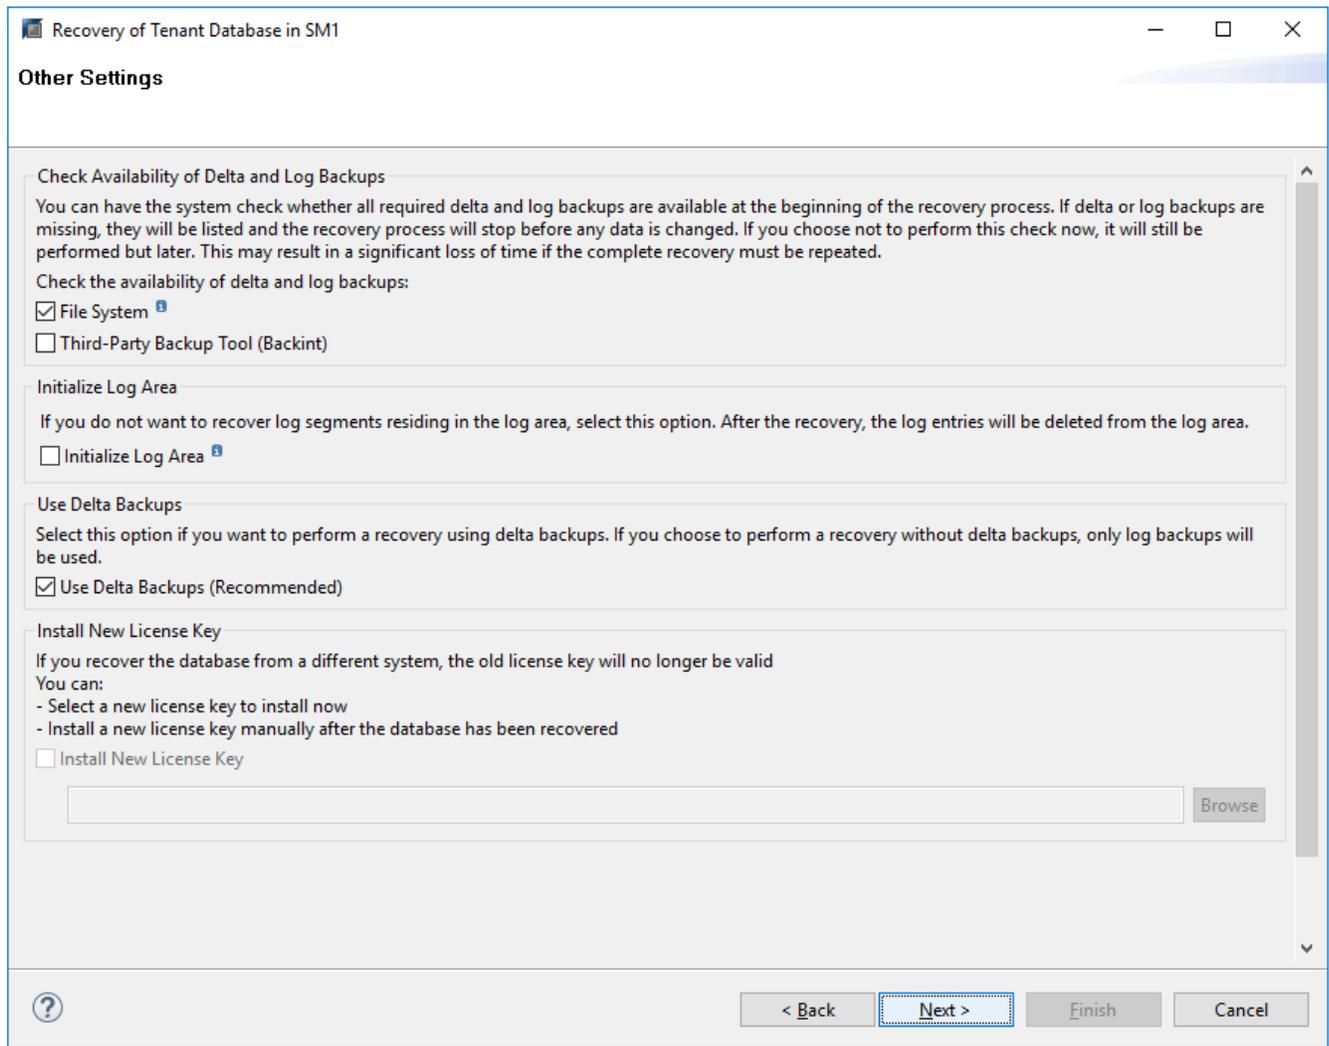
Check Availability

< Back Next > Finish Cancel

13. Provide the log backup location.



14. Select the other settings as required.



15. Start the tenant recovery operation.

Recovery of Tenant Database in SM1

### Review Recovery Settings

Review the recovery settings and choose 'Finish' to start the recovery. You can modify the recovery settings by choosing 'Back'.

**Database Information**

Database:	TENANT2@SM1
Host:	hana-2
Version:	2.00.040.00.1553674765

**Recovery Definition**

Recovery Type:	Snapshot (Point-in-Time Recovery (Until Now))
----------------	---

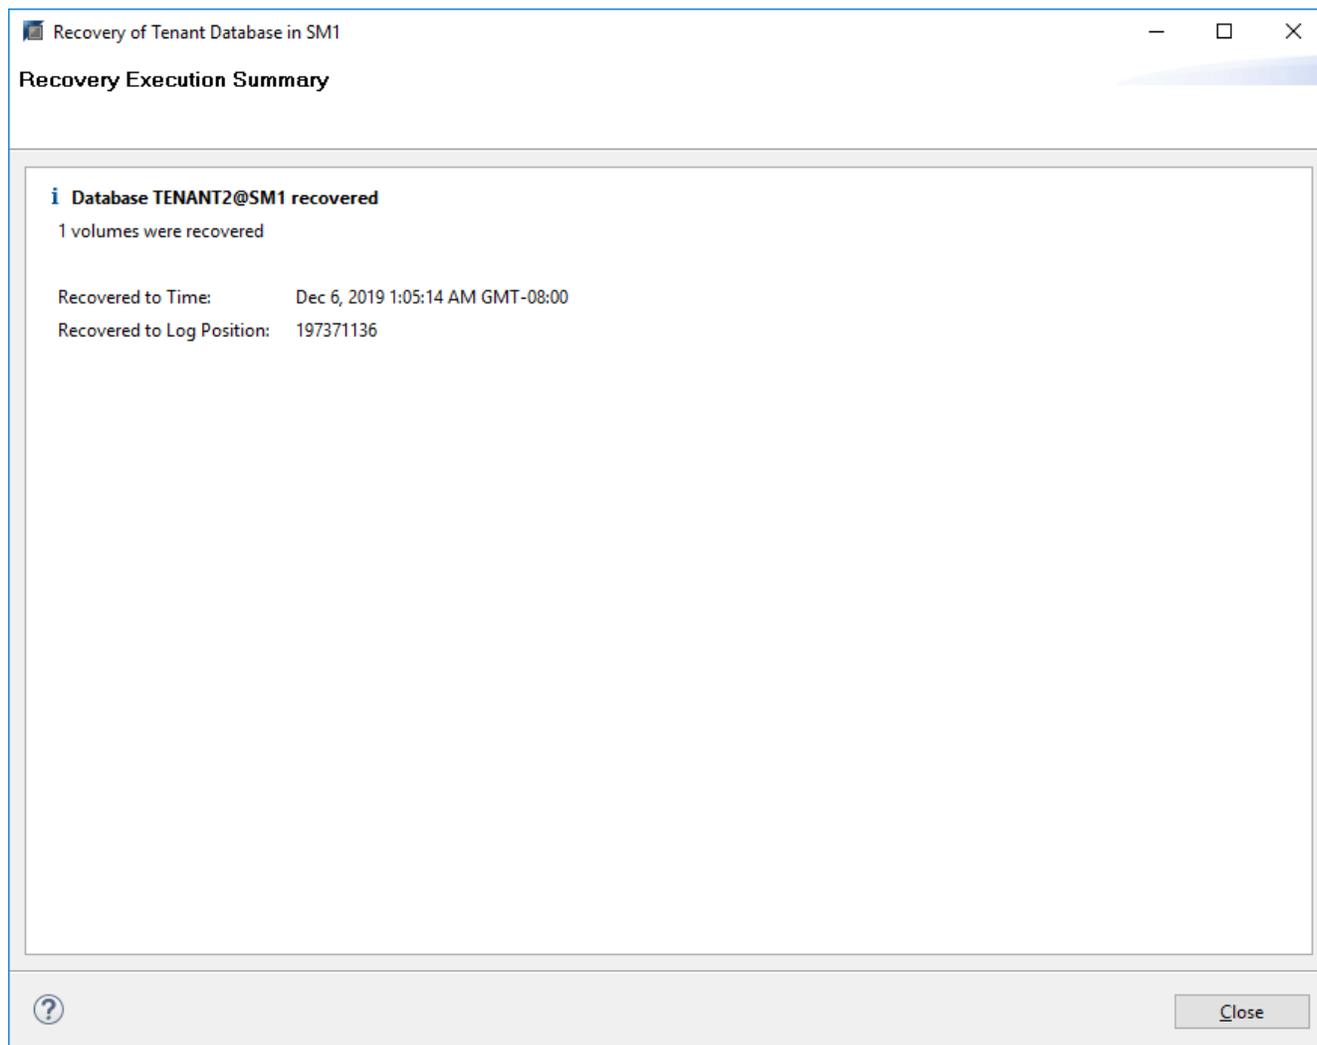
**Configuration File Handling**

 Caution

To recover customer-specific configuration changes, you may need to make the changes manually in the target system.  
More Information: SAP HANA Administration Guide

Show SQL Statement





### Restore with manual recovery

To restore and recover an SAP HANA MDC single-tenant system using SAP HANA Studio and SnapCenter, complete the following steps:

1. Prepare the restore and recovery process with SAP HANA Studio:
  - a. Select Recover System Database and confirm shutdown of the SAP HANA system.
  - b. Select the recovery type and the log backup location.
  - c. The list of data backups is shown. Select Backup to see the external backup ID.
2. Perform the restore process with SnapCenter:
  - a. In the topology view of the resource, select Local Copies to restore from primary storage or Vault Copies if you want to restore from an off-site backup storage.
  - b. Select the SnapCenter backup that matches the external backup ID or comment field from SAP HANA Studio.
  - c. Start the restore process.



If a volume-based restore from primary storage is chosen, the data volumes must be unmounted from all SAP HANA database hosts before the restore and mounted again after the restore process is finished.



In an SAP HANA multiple-host setup with FC, the unmount and mount operations are executed by the SAP HANA name server as part of the shutdown and startup process of the database.

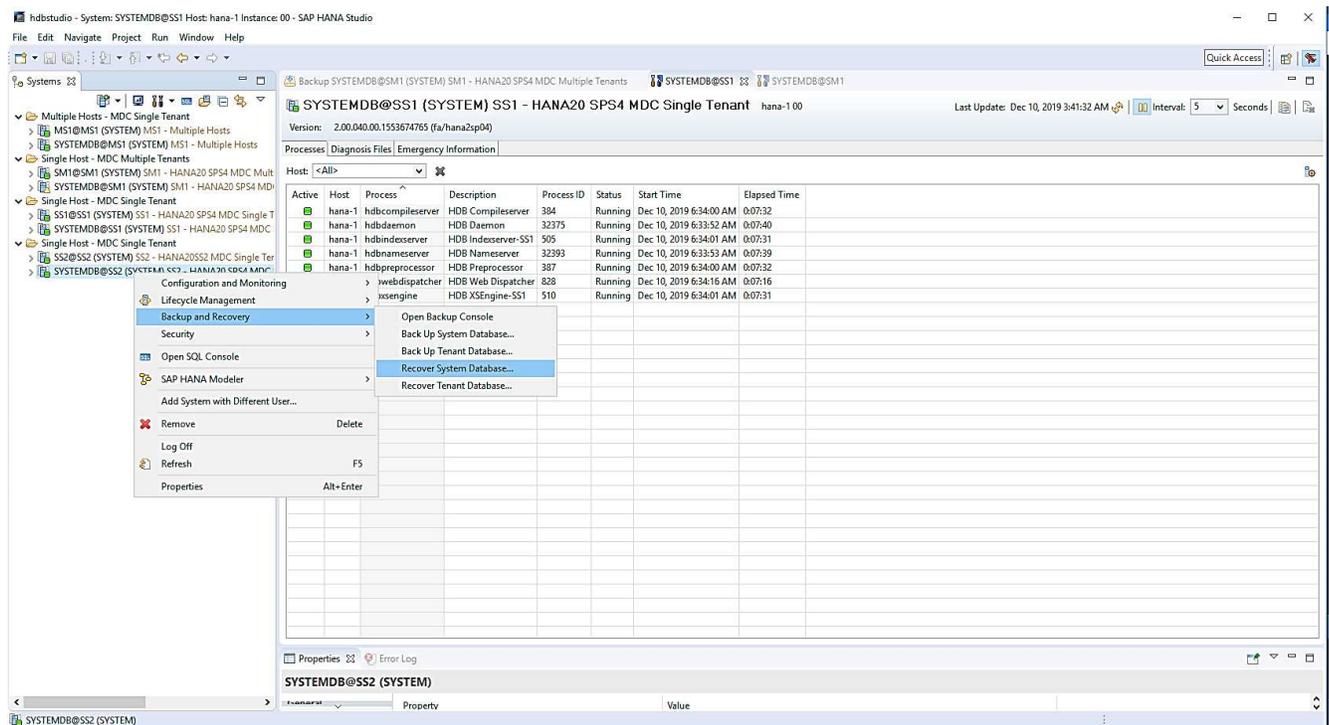
3. Run the recovery process for the system database with SAP HANA Studio:
  - a. Click Refresh from the backup list and select the available backup for recovery (indicated with a green icon).
  - b. Start the recovery process. After the recovery process is finished, the system database is started.
4. Run the recovery process for the tenant database with SAP HANA Studio:
  - a. Select Recover Tenant Database and select the tenant to be recovered.
  - b. Select the recovery type and the log backup location.

A list of data backups displays. Because the data volume has already been restored, the tenant backup is indicated as available (in green).

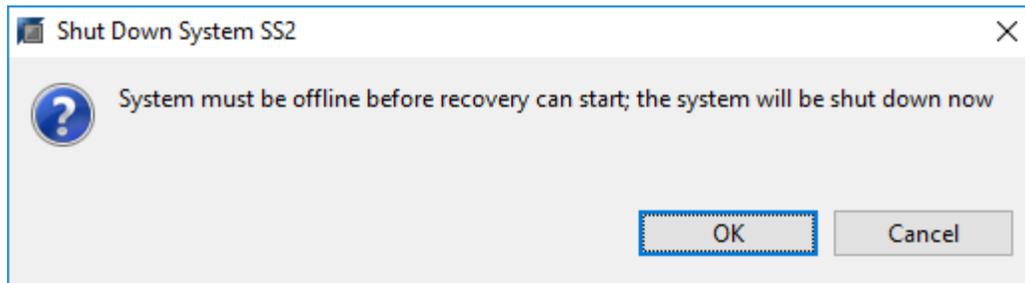
- c. Select this backup and start the recovery process. After the recovery process is finished, the tenant database is started automatically.

The following section describes the steps of the restore and recovery operations of the manually configured HANA system SS2 (SAP HANA single host, MDC multiple tenant system using NFS).

1. In SAP HANA Studio, select the Recover System Database option to start the recovery of the system database.

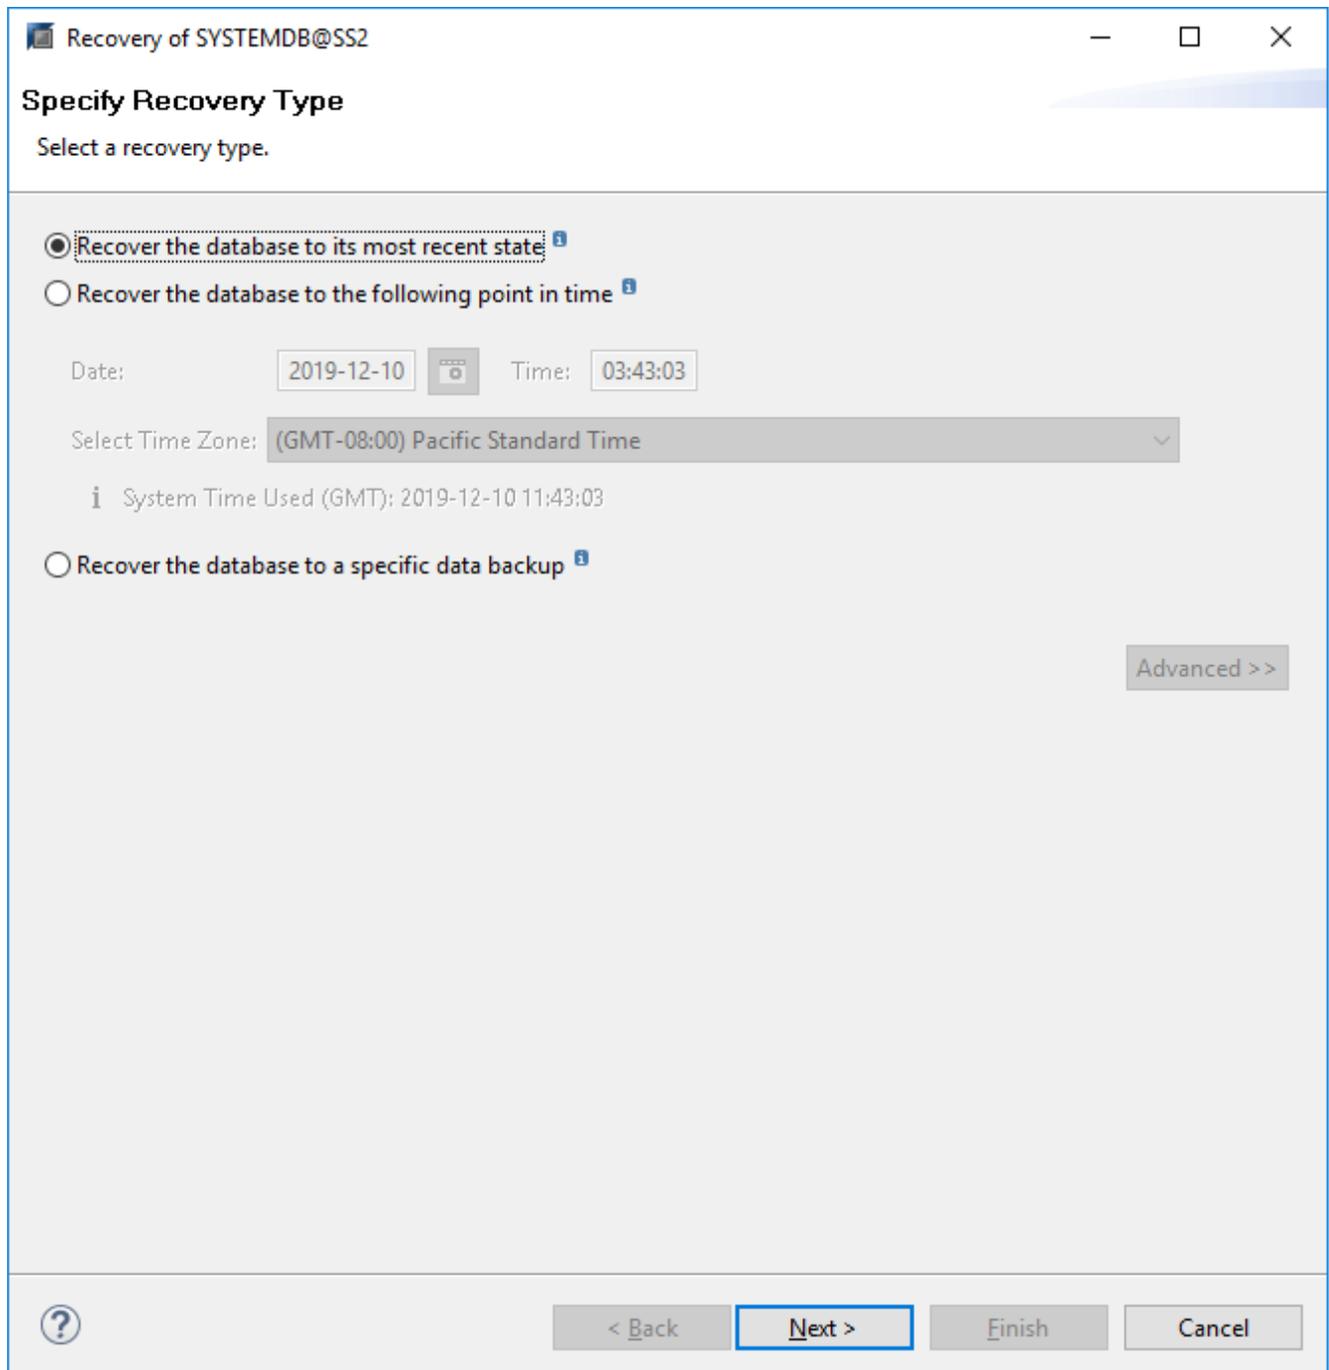


2. Click OK to shut down the SAP HANA database.



The SAP HANA system shuts down and the recovery wizard is started.

3. Select the recovery type and click Next.



4. Provide the location of the backup catalog and click Next.

The screenshot shows a Windows-style dialog box titled "Recovery of SYSTEMDB@SS2". The main heading is "Locate Backup Catalog" with the instruction "Specify location of the backup catalog." Below this, there are three radio button options: "Recover using the backup catalog" (selected), "Search for the backup catalog in the file system only" (selected), and "Recover without the backup catalog". A text box labeled "Backup Catalog Location:" contains the path "/mnt/log-backup/SYSTEMDB". Below the radio buttons is a section titled "Backint System Copy" with a checkbox "Backint System Copy" (unchecked) and a text box "Source System:". At the bottom, there are four buttons: a help icon (?), "< Back", "Next >" (highlighted with a blue dashed border), "Finish", and "Cancel".

5. A list of available backups displays based on the content of the backup catalog. Choose the required backup and note the external backup ID: in our example, the most recent backup.

Recovery of SYSTEMDB@SS2

### Select a Backup

To recover this snapshot, it must be available in the data area.

**Selected Point in Time**  
Database will be recovered to its most recent state.

**Backups**  
The overview shows backups that were recorded in the backup catalog as successful. The backup at the top is estimated to have the shortest recovery time.

Start Time	Location	Backup Prefix	Available
2019-12-10 02:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 22:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 18:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 14:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 10:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 06:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 02:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 22:05:07	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 18:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 14:05:08	/hana/data/SS2	SNAPSHOT	✘

Refresh Show More

**Details of Selected Item**

Start Time: 2019-12-10 02:05:08 Destination Type: SNAPSHOT Source System: SYSTEMDB@SS2  
 Size: 0 B Backup ID: 1575972308584 External Backup ID: SnapCenter\_LocalSnap\_Hourly\_12-10-2019\_02.05.01.3757  
 Backup Name: /hana/data/SS2  
 Alternative Location:

Check Availability

< Back Next > Finish Cancel

6. Unmount all data volumes.

```
umount /hana/data/SS2/mnt00001
```

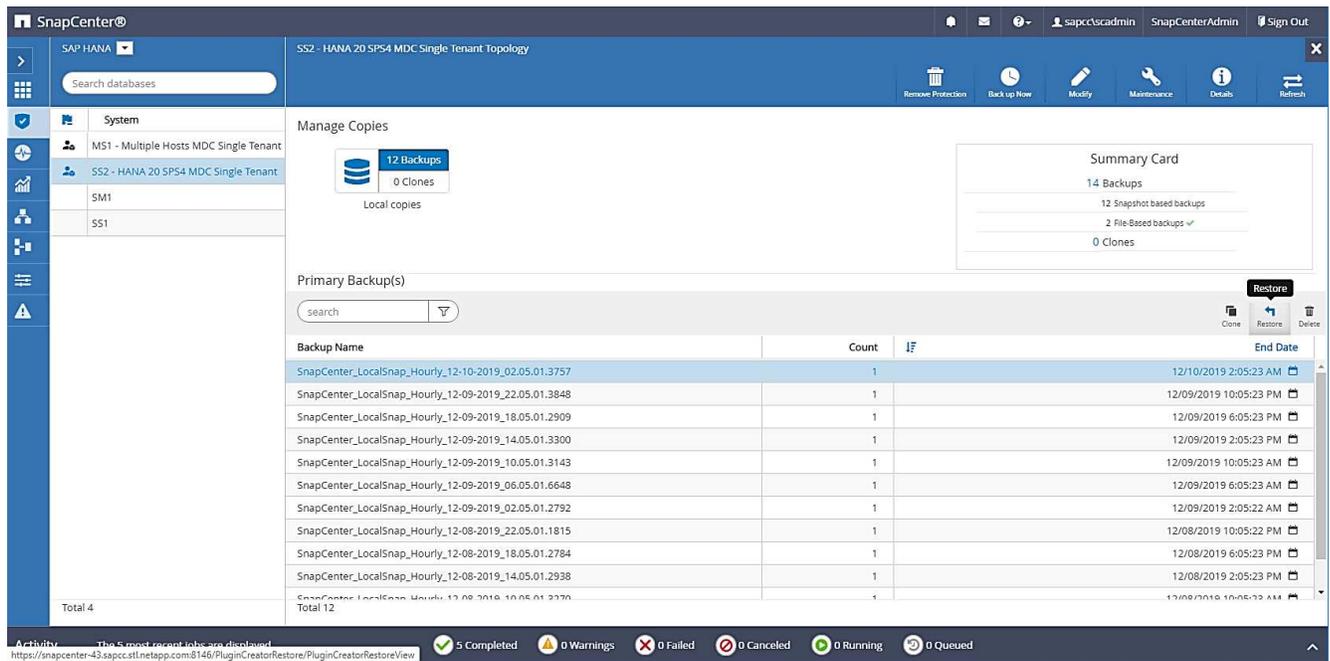


For an SAP HANA multiple host system with NFS, all data volumes on each host must be unmounted.



In an SAP HANA multiple-host setup with FC, the unmount operation is executed by the SAP HANA name server as a part of the shutdown process.

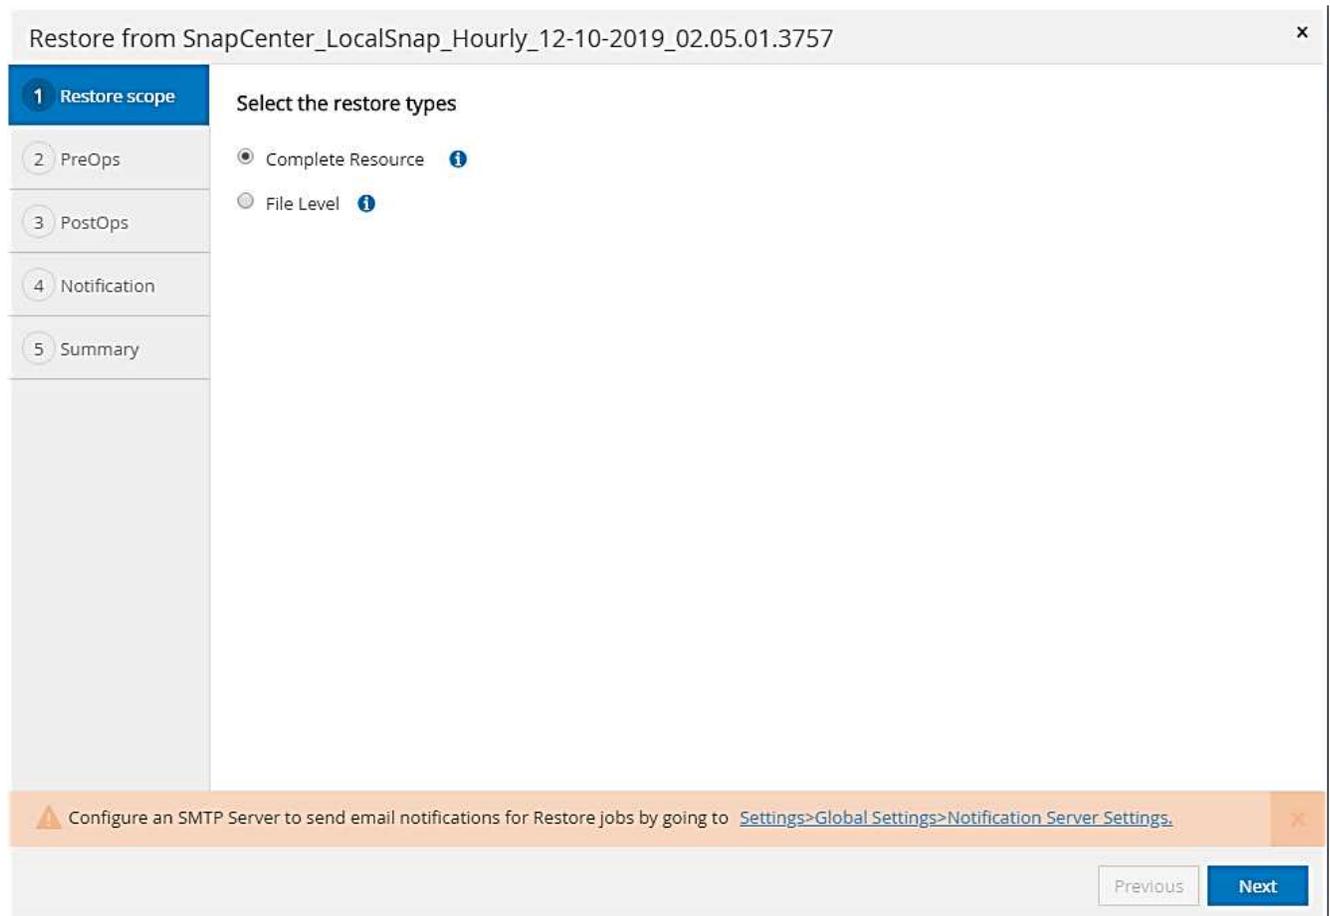
7. From the SnapCenter GUI, select the resource topology view and select the backup that should be restored; in our example, the most recent primary backup. Click the Restore icon to start the restore.



The SnapCenter restore wizard starts.

8. Select the restore type Complete Resource or File Level.

Select Complete Resource to use a volume-based restore.



9. Select File Level and All to use a single-file SnapRestore operation for all files.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-10-2019\_02.05.01.3757

1 Restore scope

2 PreOps

3 PostOps

4 Notification

5 Summary

Select the restore types

Complete Resource

File Level

Select files to restore

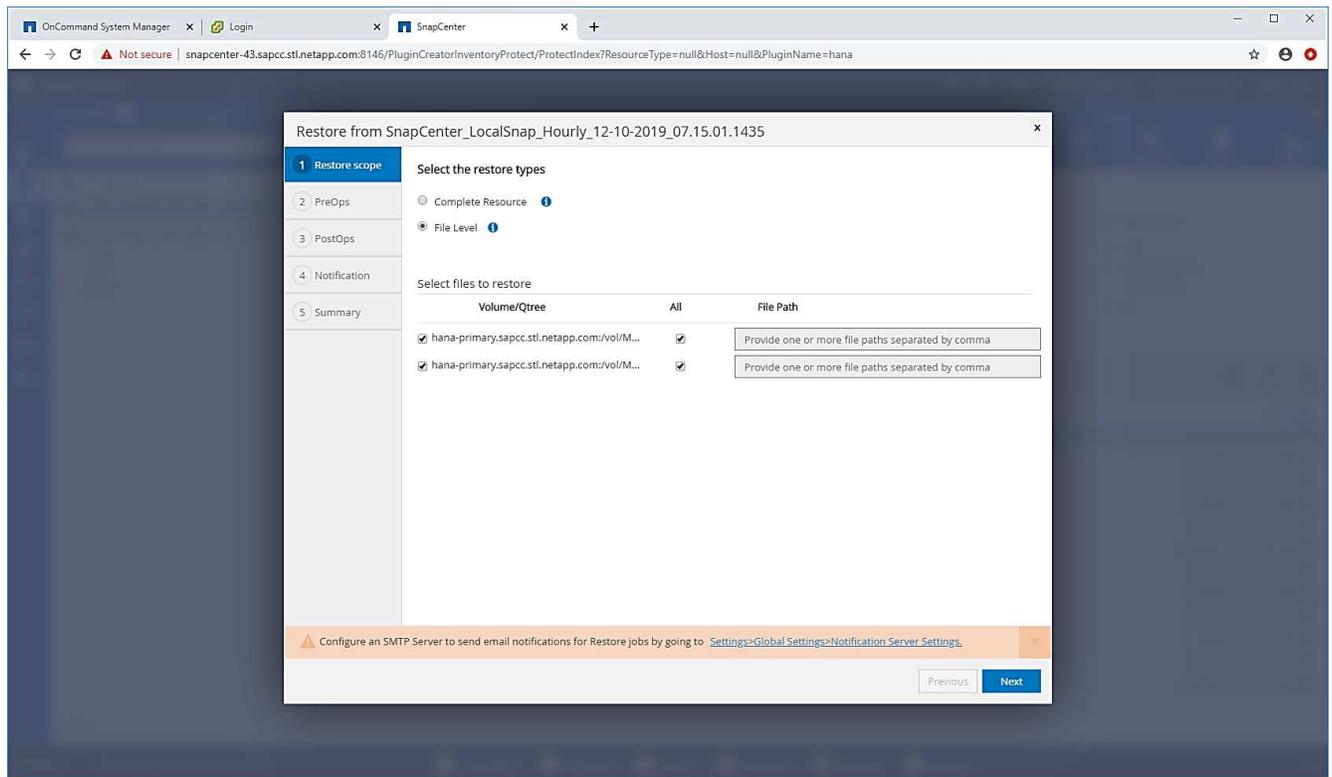
Volume/Qtree	All	File Path
<input checked="" type="checkbox"/> hana-primary.sapcc.stl.netapp.com:/vol/SS...	<input checked="" type="checkbox"/>	Provide one or more file paths separated by comma

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

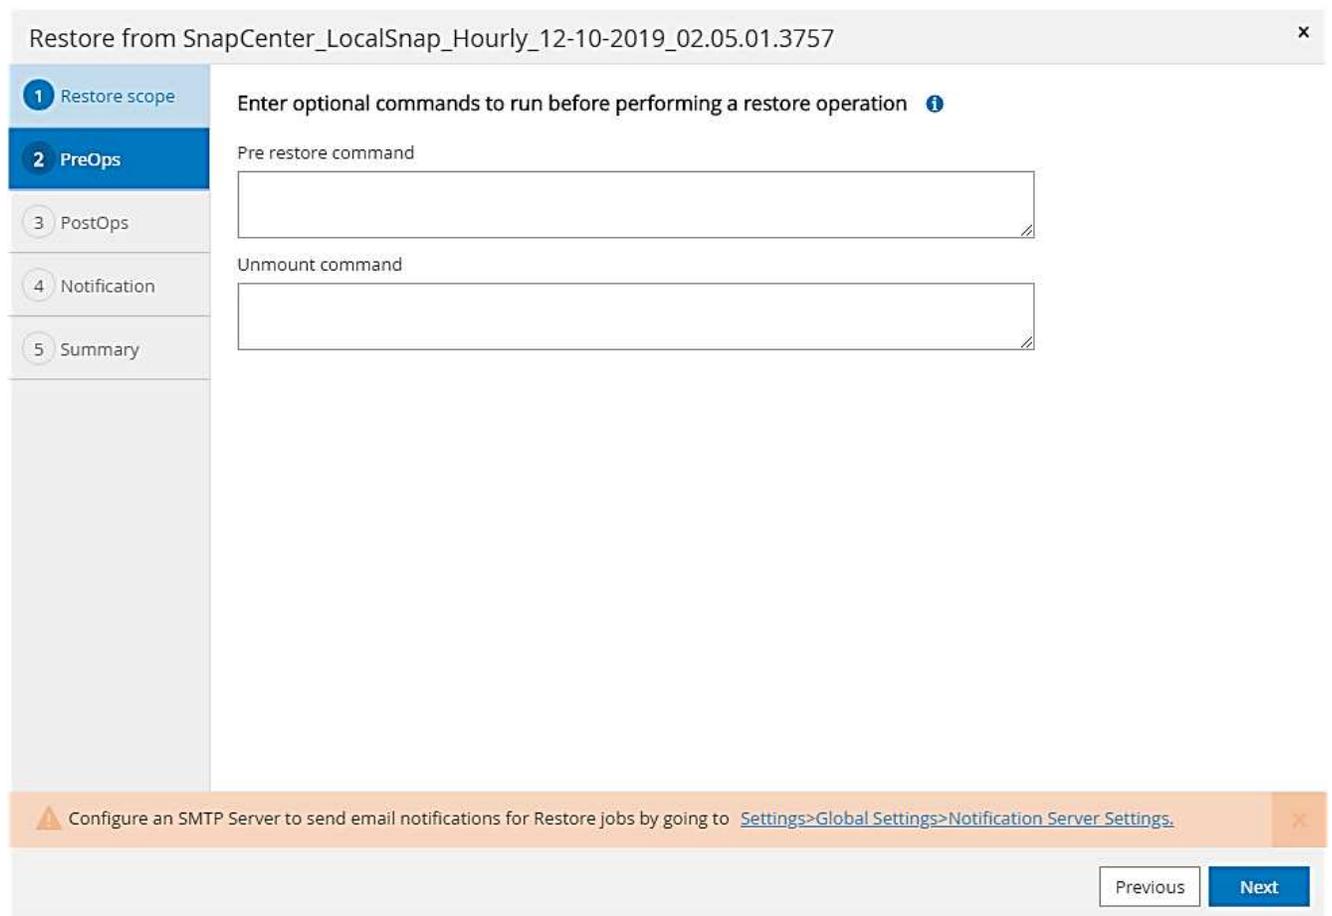
Previous Next



For a file-level restore of a SAP HANA multiple host system, select all the volumes.



10. (Optional) Specify the commands that should be executed from the SAP HANA plug-in running on the central HANA plug-in host. Click Next.



11. Specify the optional commands and click Next.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-10-2019\_02.05.01.3757

1 Restore scope

2 PreOps

3 PostOps

4 Notification

5 Summary

Enter optional commands to run after performing a restore operation ⓘ

Mount command

Post restore command

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

Previous Next

12. Specify the notification settings so that SnapCenter can send a status email and job log. Click Next.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-10-2019\_02.05.01.3757 ×

- 1 Restore scope
- 2 PreOps
- 3 PostOps
- 4 Notification**
- 5 Summary

### Provide email settings ?

Email preference

From

To

Subject

Attach Job Report

⚠ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server. ×

13. Review the summary and click Finish to start the restore.

Restore from SnapCenter\_LocalSnap\_Hourly\_12-10-2019\_02.05.01.3757 x

- 1 Restore scope
- 2 PreOps
- 3 PostOps
- 4 Notification
- 5 Summary**

### Summary

Backup Name	SnapCenter_LocalSnap_Hourly_12-10-2019_02.05.01.3757
Backup date	12/10/2019 2:05:23 AM
Restore scope	Complete Resource
Pre restore command	
Unmount command	
Mount command	
Post restore command	
Send email	No

⚠ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server. x

14. The restore job starts, and the job log can be displayed by double-clicking the log line in the activity pane.

### Job Details ✕

Restore 'SnapCenter-43.sapcc.stl.netapp.com\hana\MDC\SS2'

- ✓ ▼ Restore 'SnapCenter-43.sapcc.stl.netapp.com\hana\MDC\SS2'
- ✓ ▼ SnapCenter-43.sapcc.stl.netapp.com
  - ✓ ▼ Restore
    - ✓ ▶ Validate Plugin Parameters
    - ✓ ▶ Pre Restore Application
    - ✓ ▶ File or Volume Restore
    - ✓ ▶ Recover Application
    - ✓ ▶ Clear Catalog on Server
    - ✓ ▶ Application Clean-Up
    - ✓ ▶ Data Collection
  - ✓ ▼ Agent Finalize Workflow

**i** Task Name: Agent Finalize Workflow Start Time: 12/10/2019 3:47:30 AM End Time: 12/10/2019 3:47:35 AM

15. Wait until the restore process completes. On each database host, mount all data volumes. In our example, only one volume must be remounted on the database host.

```
mount /hana/data/SP1/mnt00001
```

16. Go to SAP HANA Studio and click Refresh to update the list of available backups. The backup that was restored with SnapCenter is shown with a green icon in the list of backups. Select the backup and click Next.

Recovery of SYSTEMDB@SS2

### Select a Backup

Select a backup to recover the SAP HANA database

**Selected Point in Time**  
Database will be recovered to its most recent state.

**Backups**  
The overview shows backups that were recorded in the backup catalog as successful. The backup at the top is estimated to have the shortest recovery time.

Start Time	Location	Backup Prefix	Available
2019-12-10 02:05:08	/hana/data/SS2	SNAPSHOT	●
2019-12-09 22:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 18:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 14:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 10:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 06:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 02:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 22:05:07	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 18:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 14:05:08	/hana/data/SS2	SNAPSHOT	✘

Refresh Show More

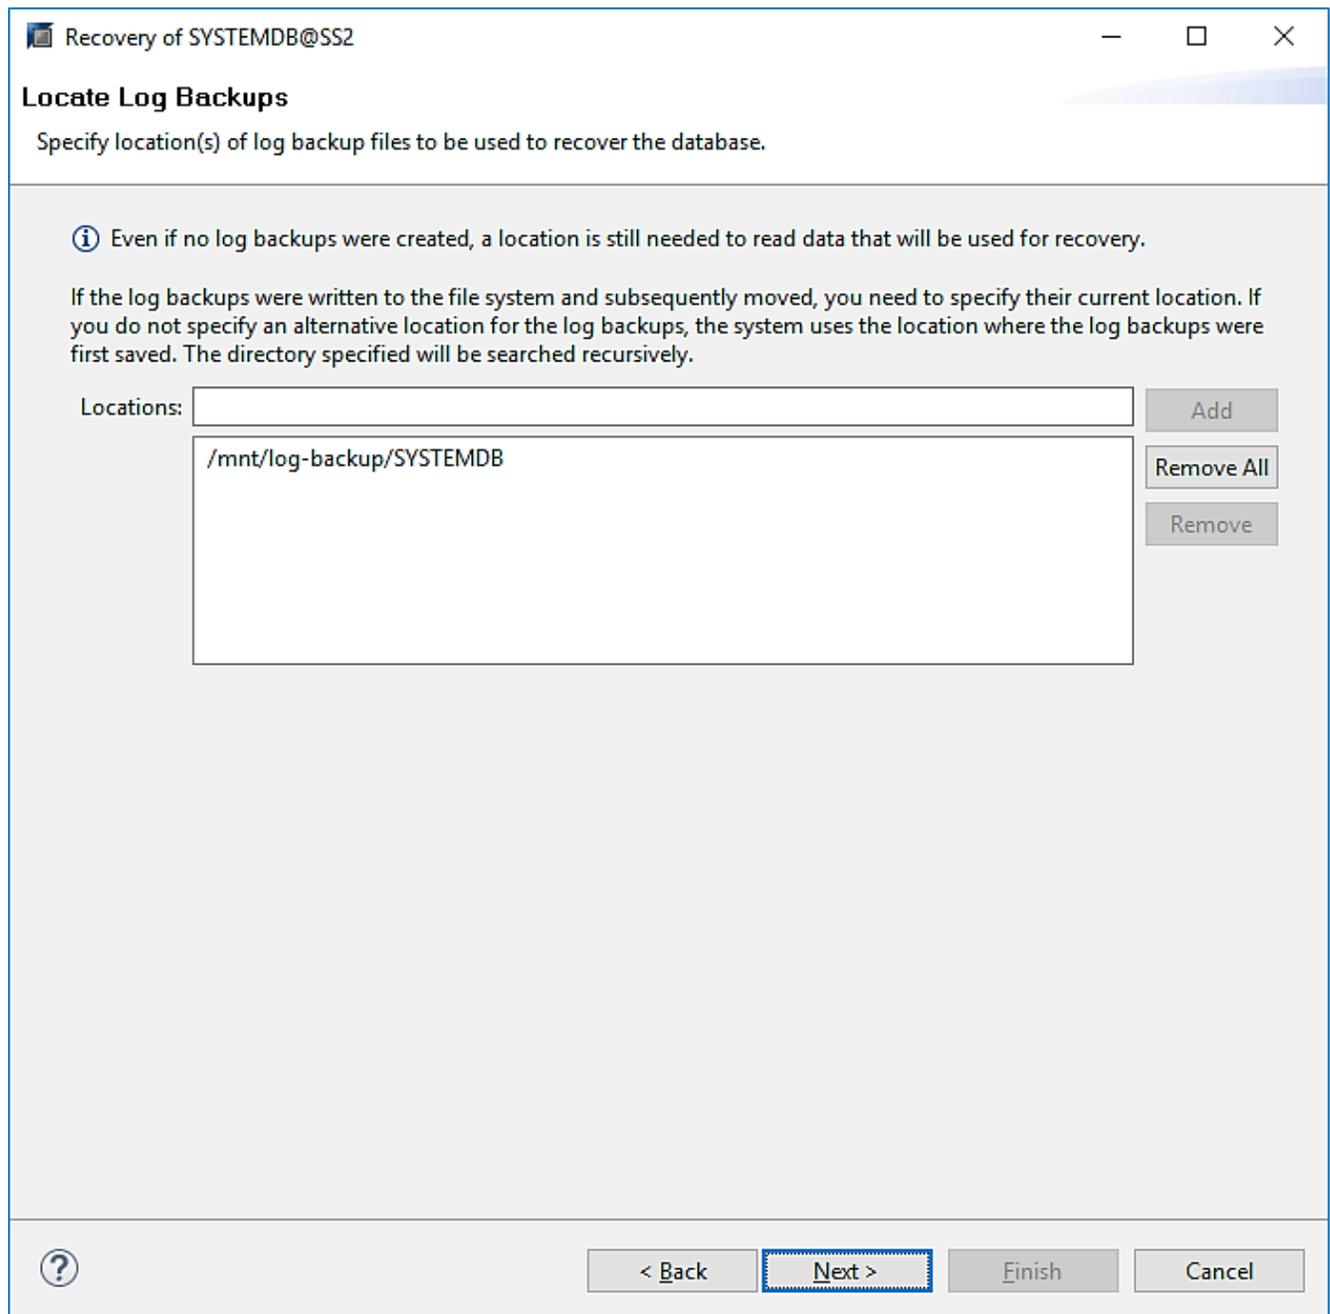
**Details of Selected Item**

Start Time: 2019-12-10 02:05:08 Destination Type: SNAPSHOT Source System: SYSTEMDB@SS2  
 Size: 0 B Backup ID: 1575972308584 External Backup ID: SnapCenter\_LocalSnap\_Hourly\_12-10-2019\_02.05.01.3757  
 Backup Name: /hana/data/SS2  
 Alternative Location:

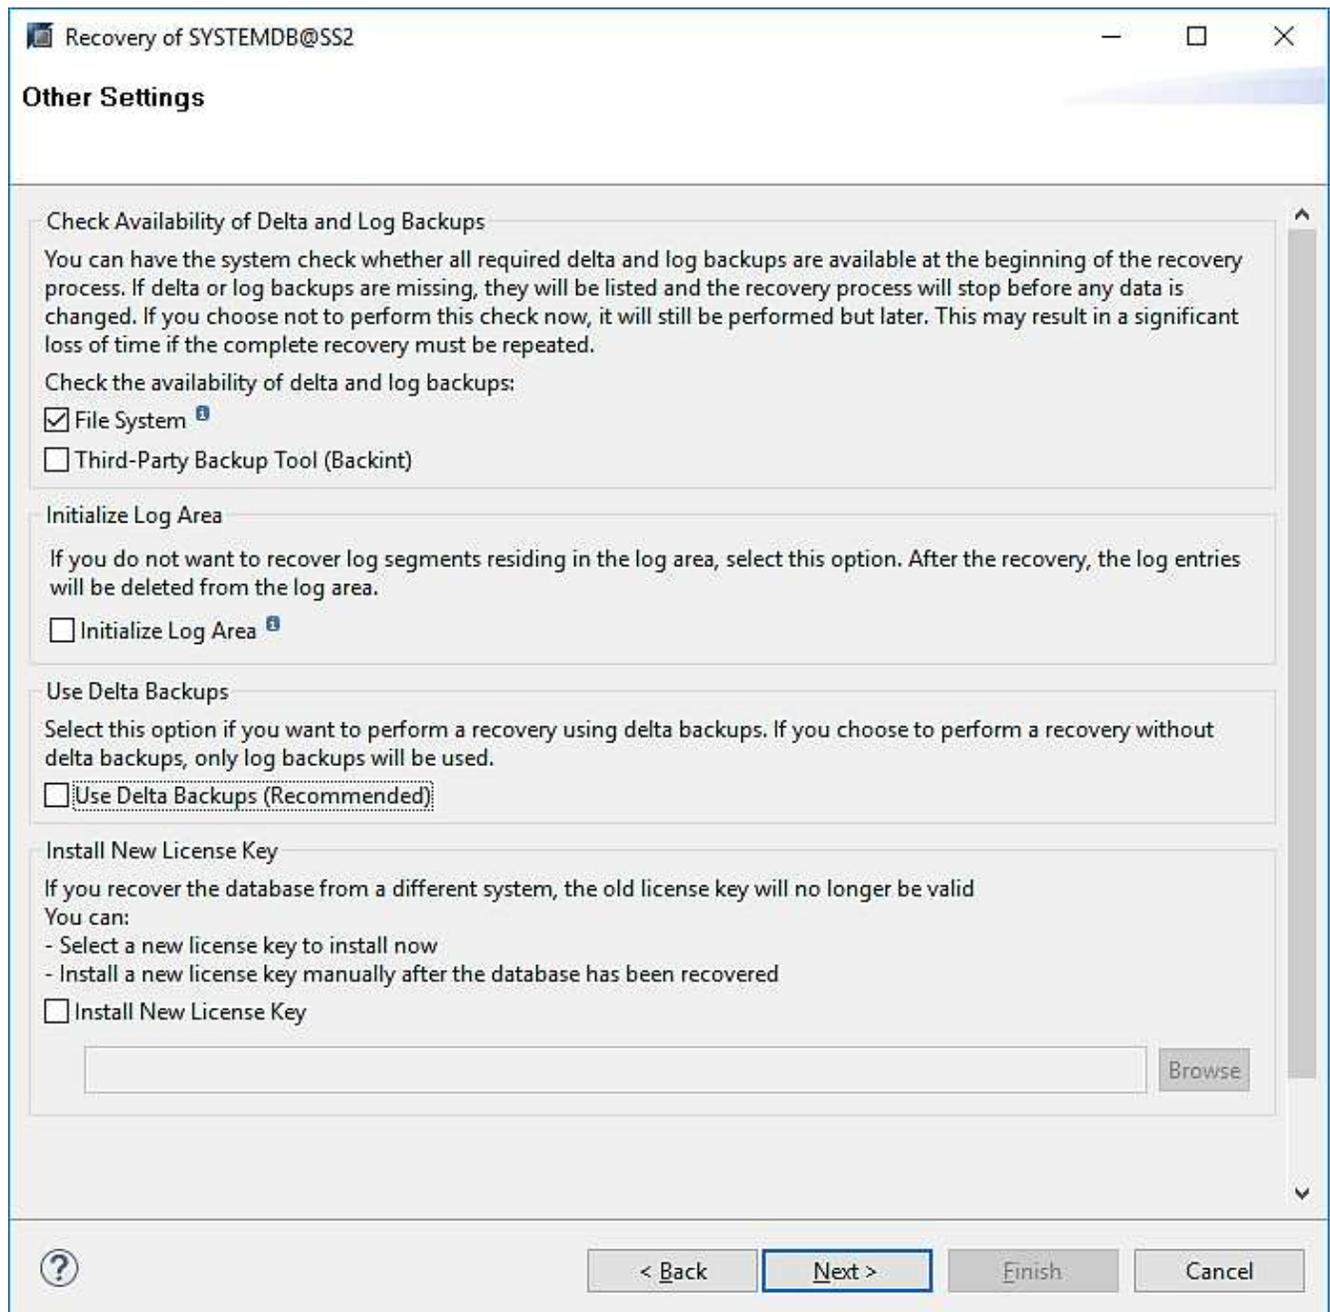
Check Availability

? < Back Next > Finish Cancel

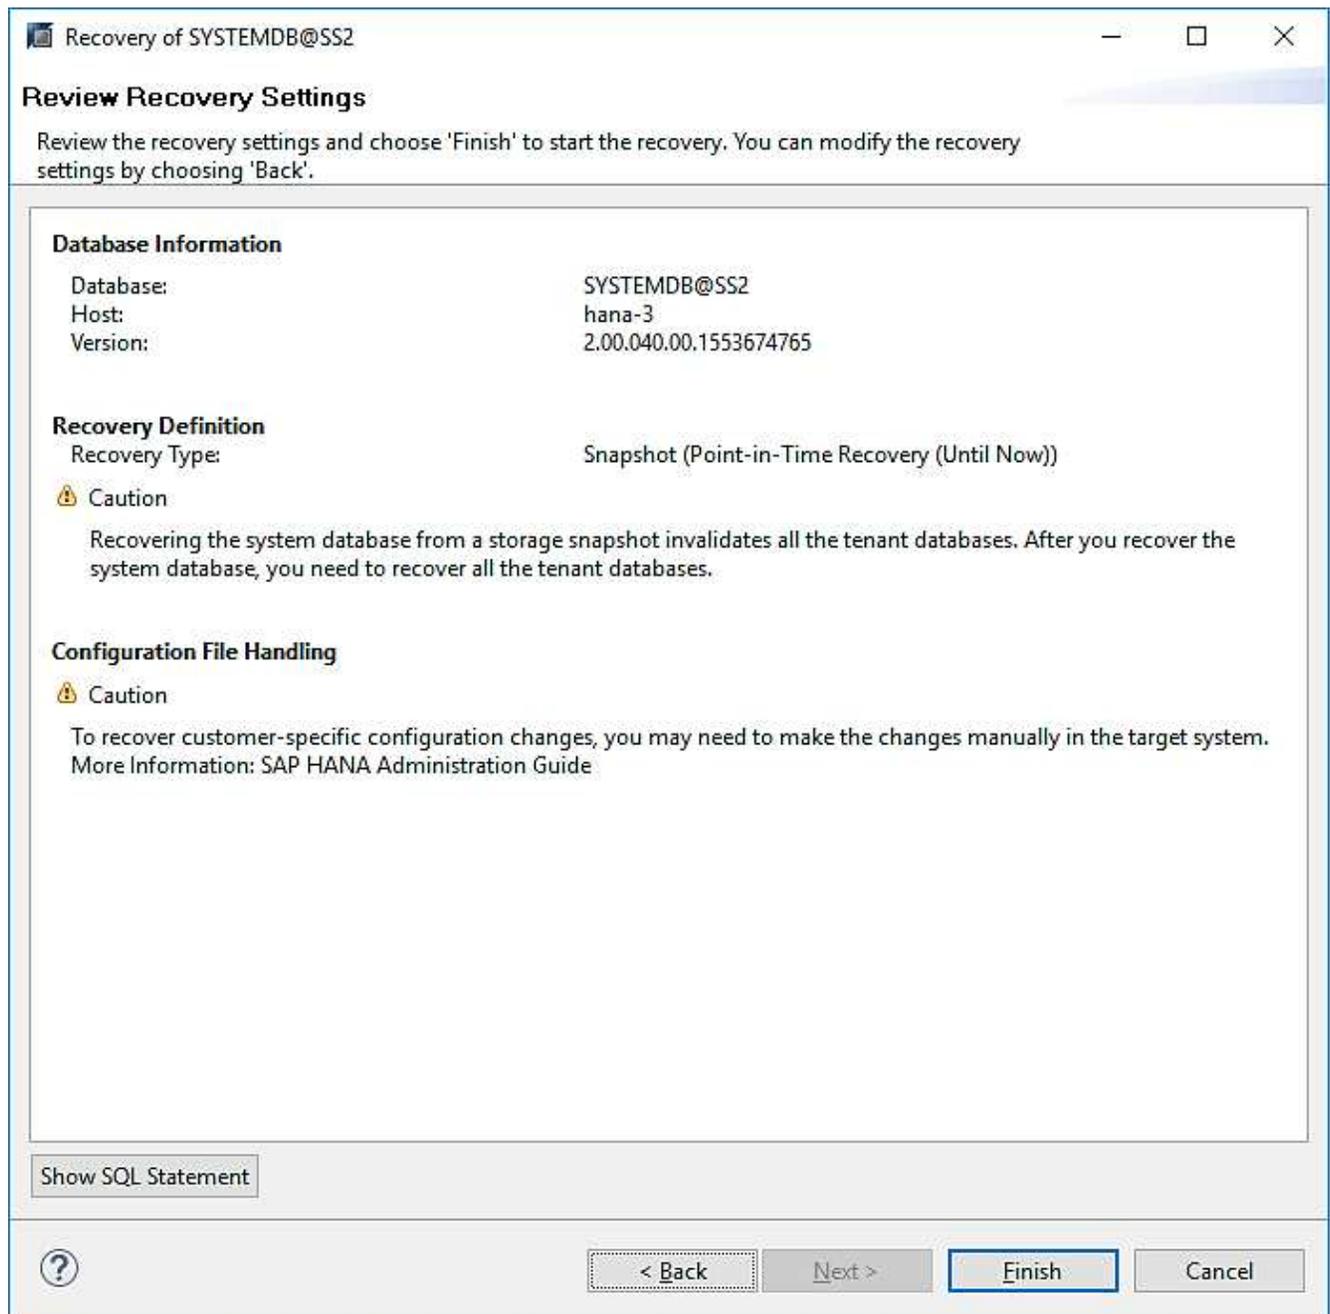
17. Provide the location of the log backups. Click Next.



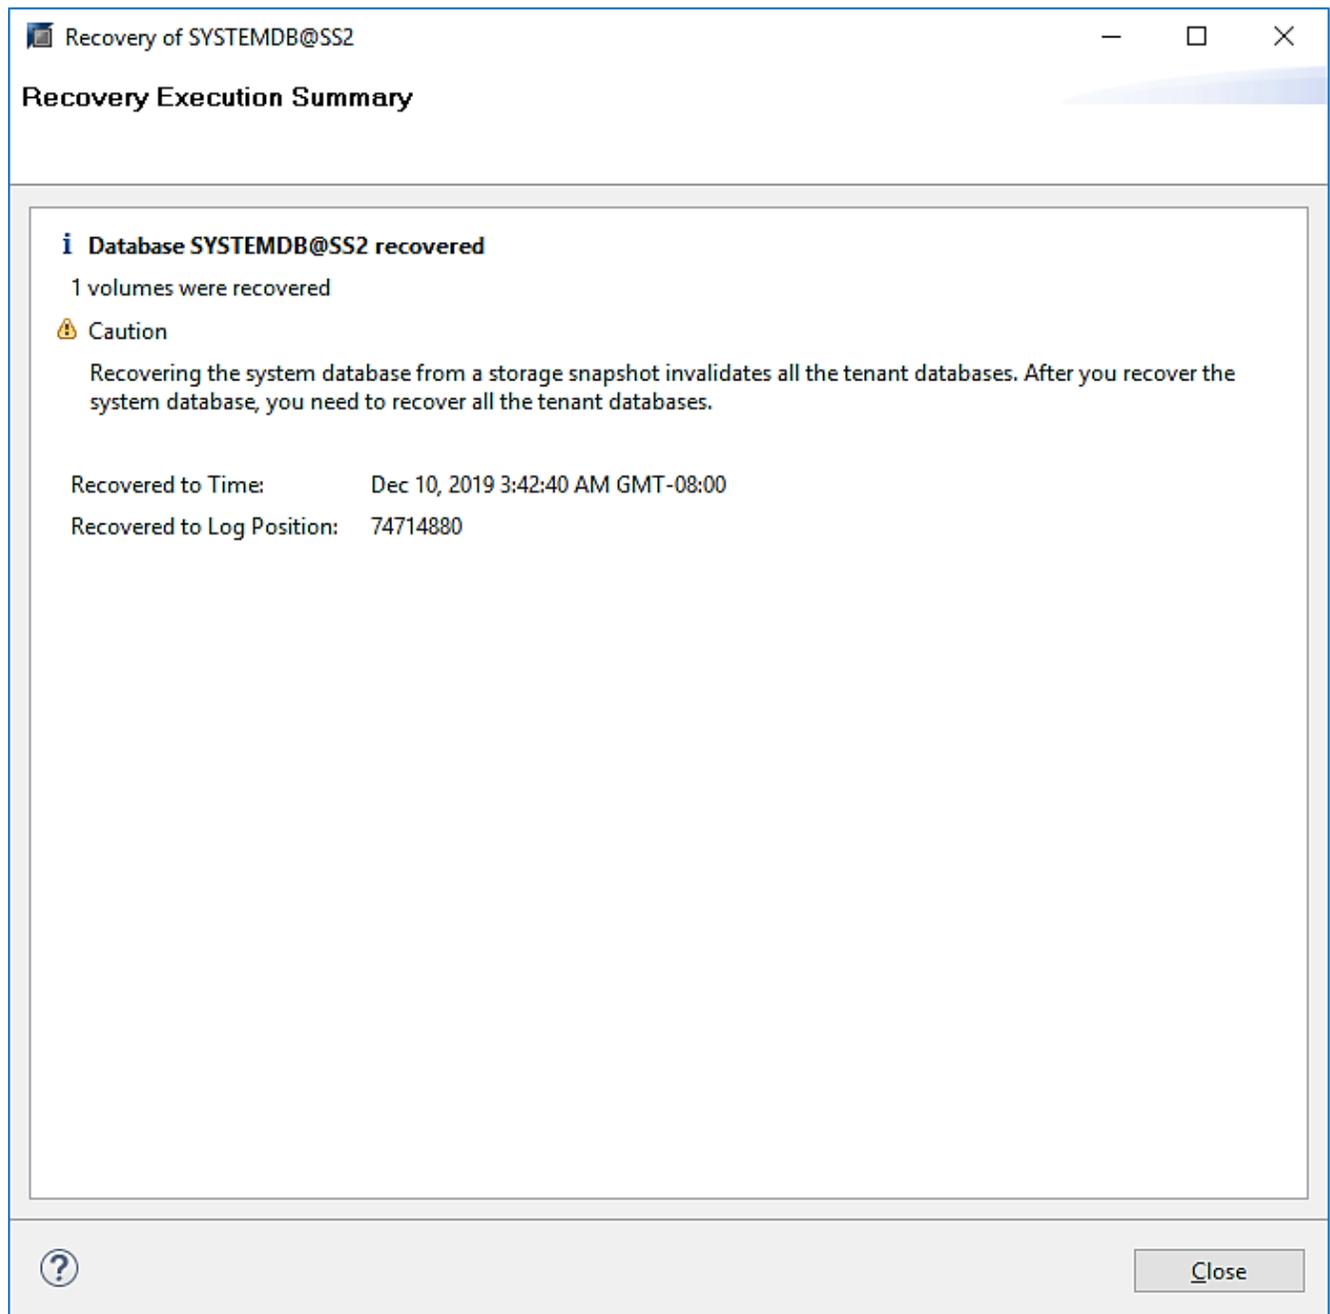
18. Select other settings as required. Make sure Use Delta Backups is not selected. Click Next.



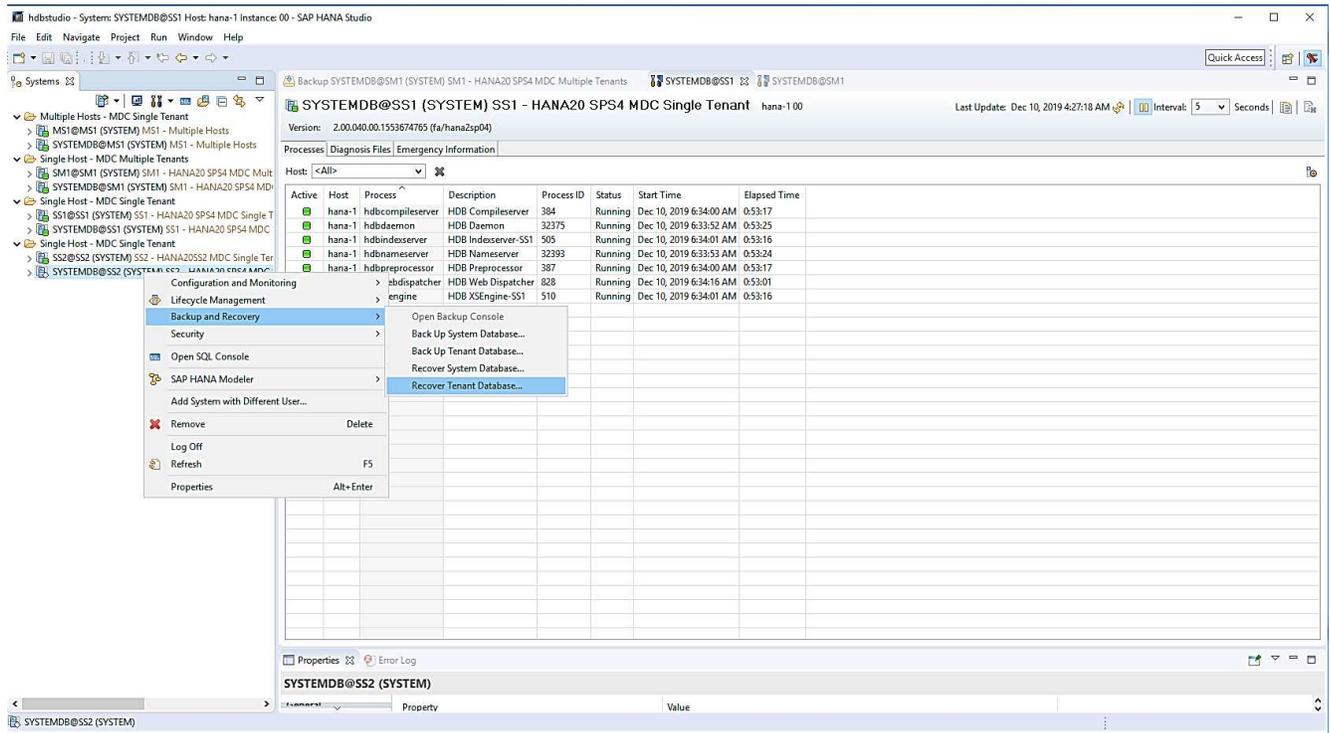
19. Review the recovery settings and click Finish.



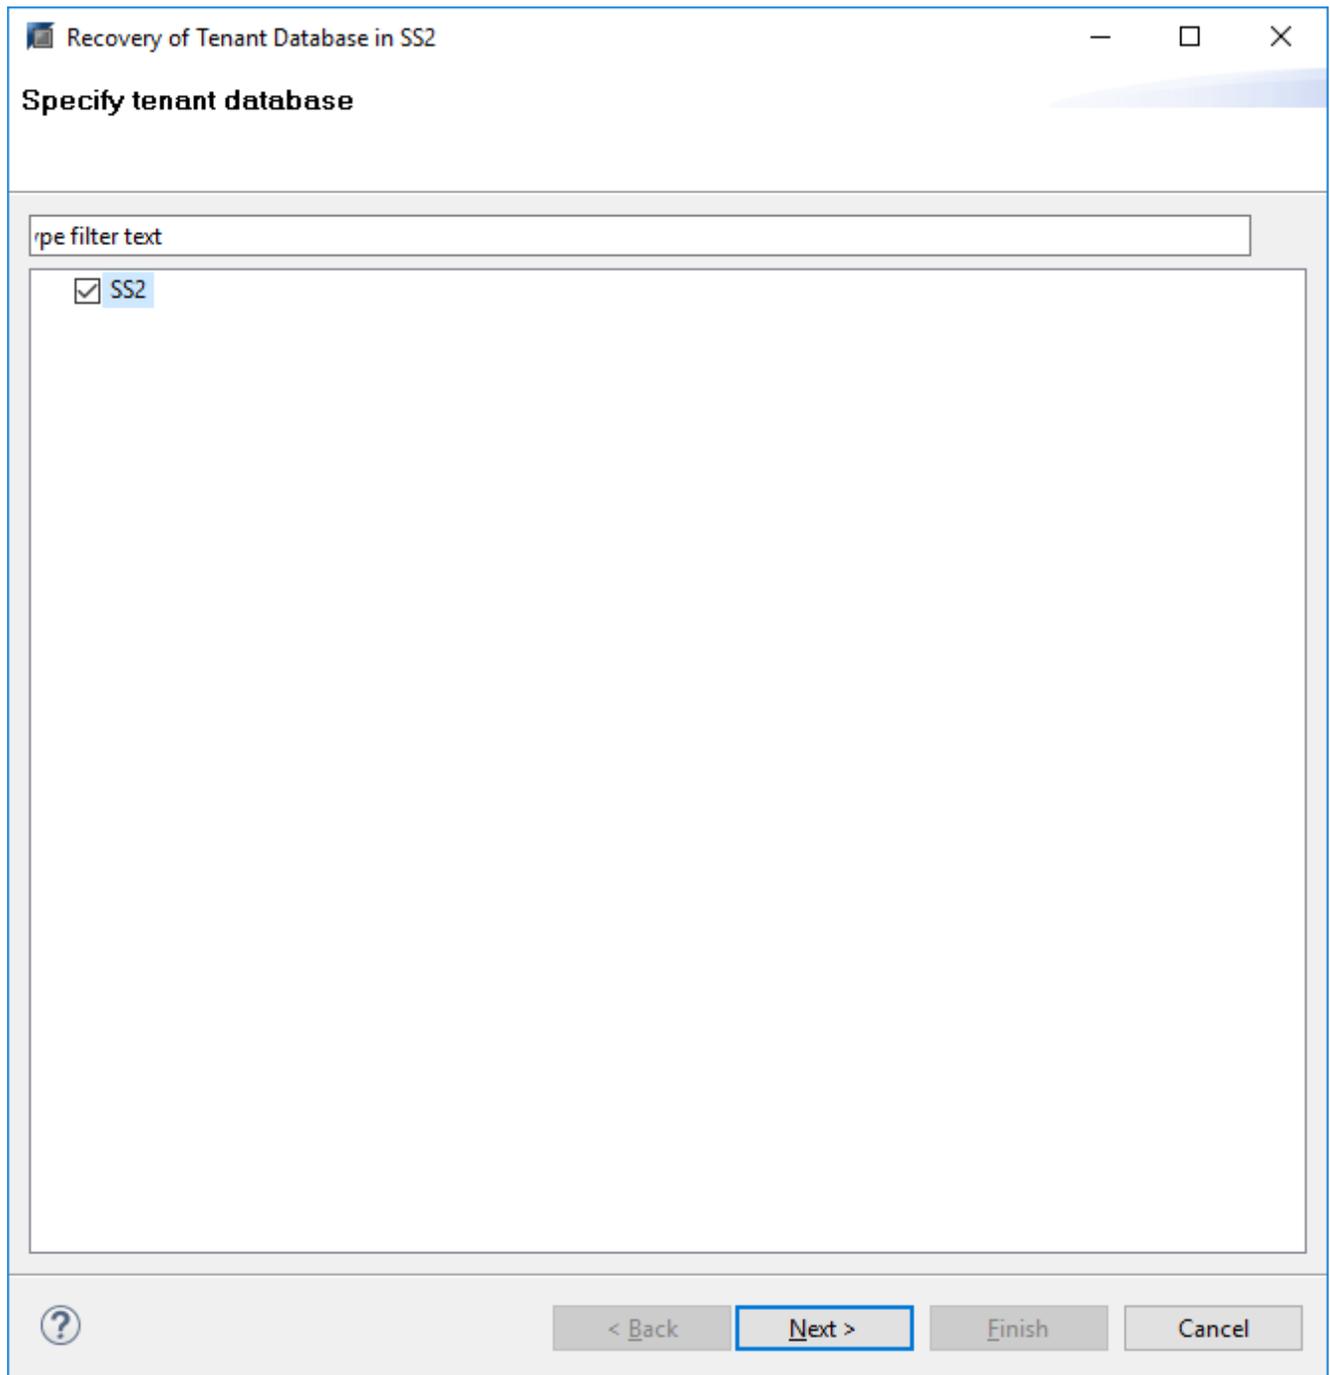
20. The recovery process starts. Wait until the recovery of the system database completes.



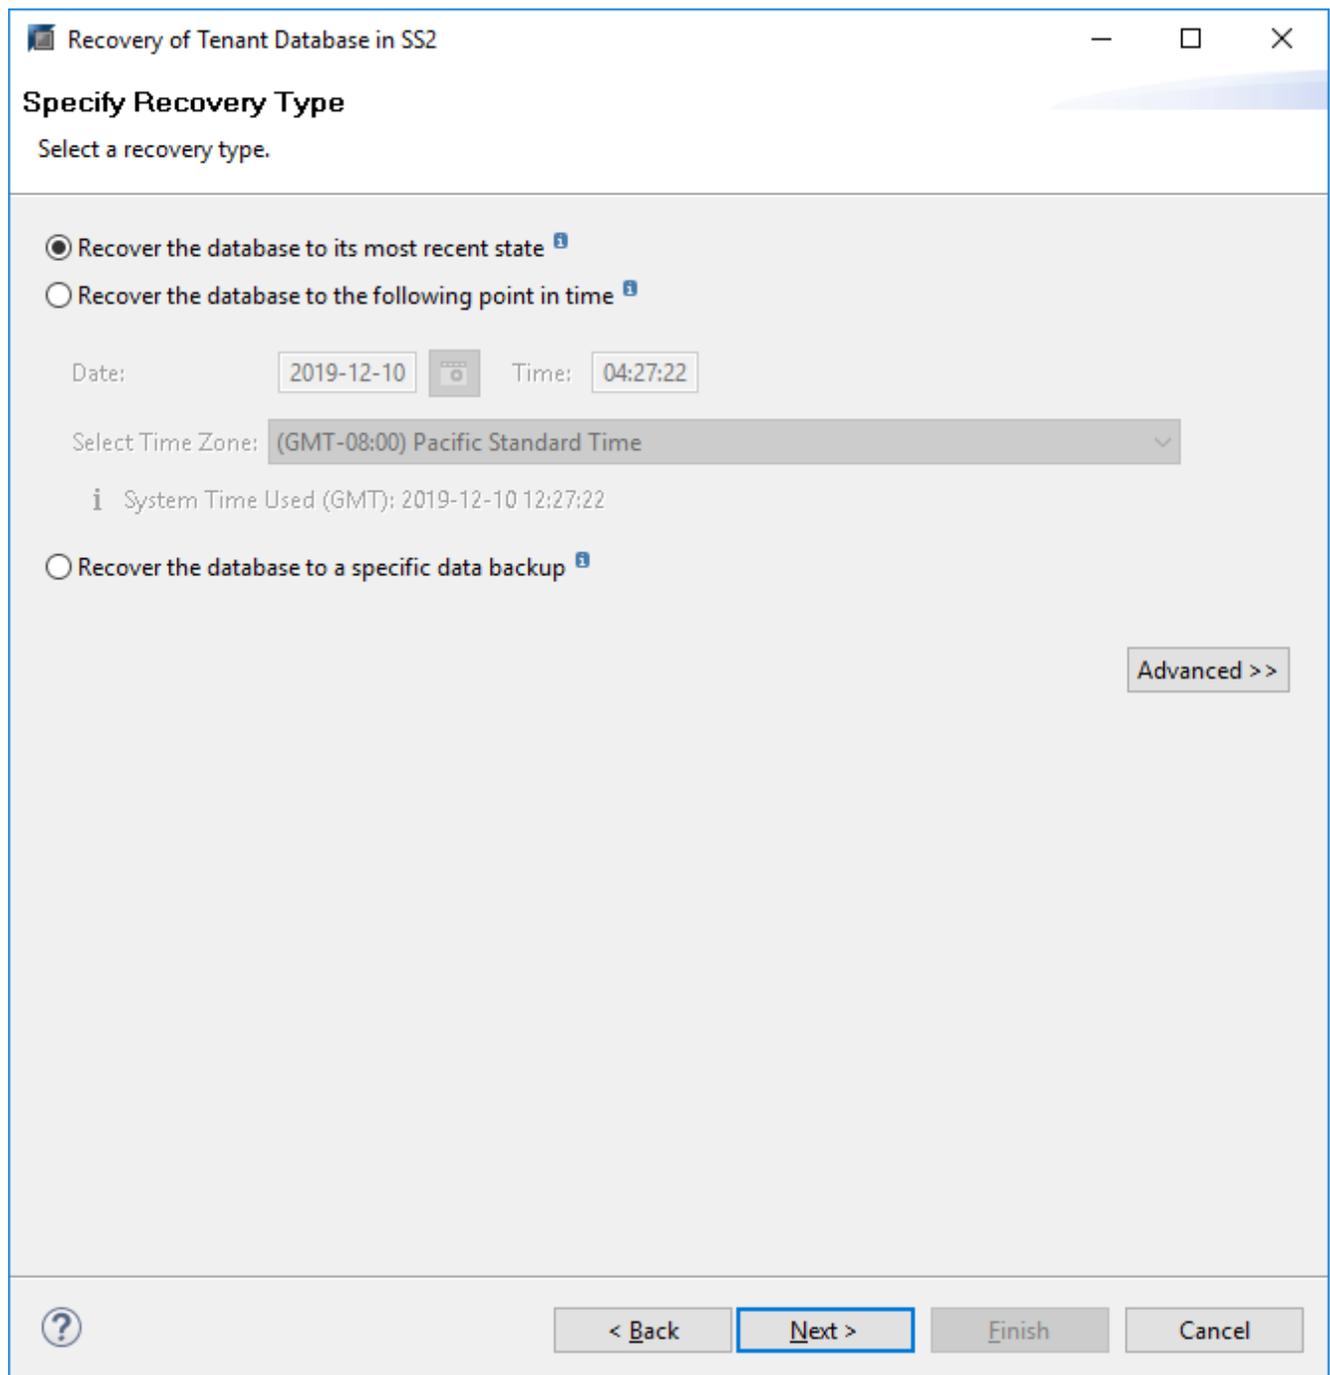
21. In SAP HANA Studio, select the entry for the system database and start Backup Recovery - Recover Tenant Database.



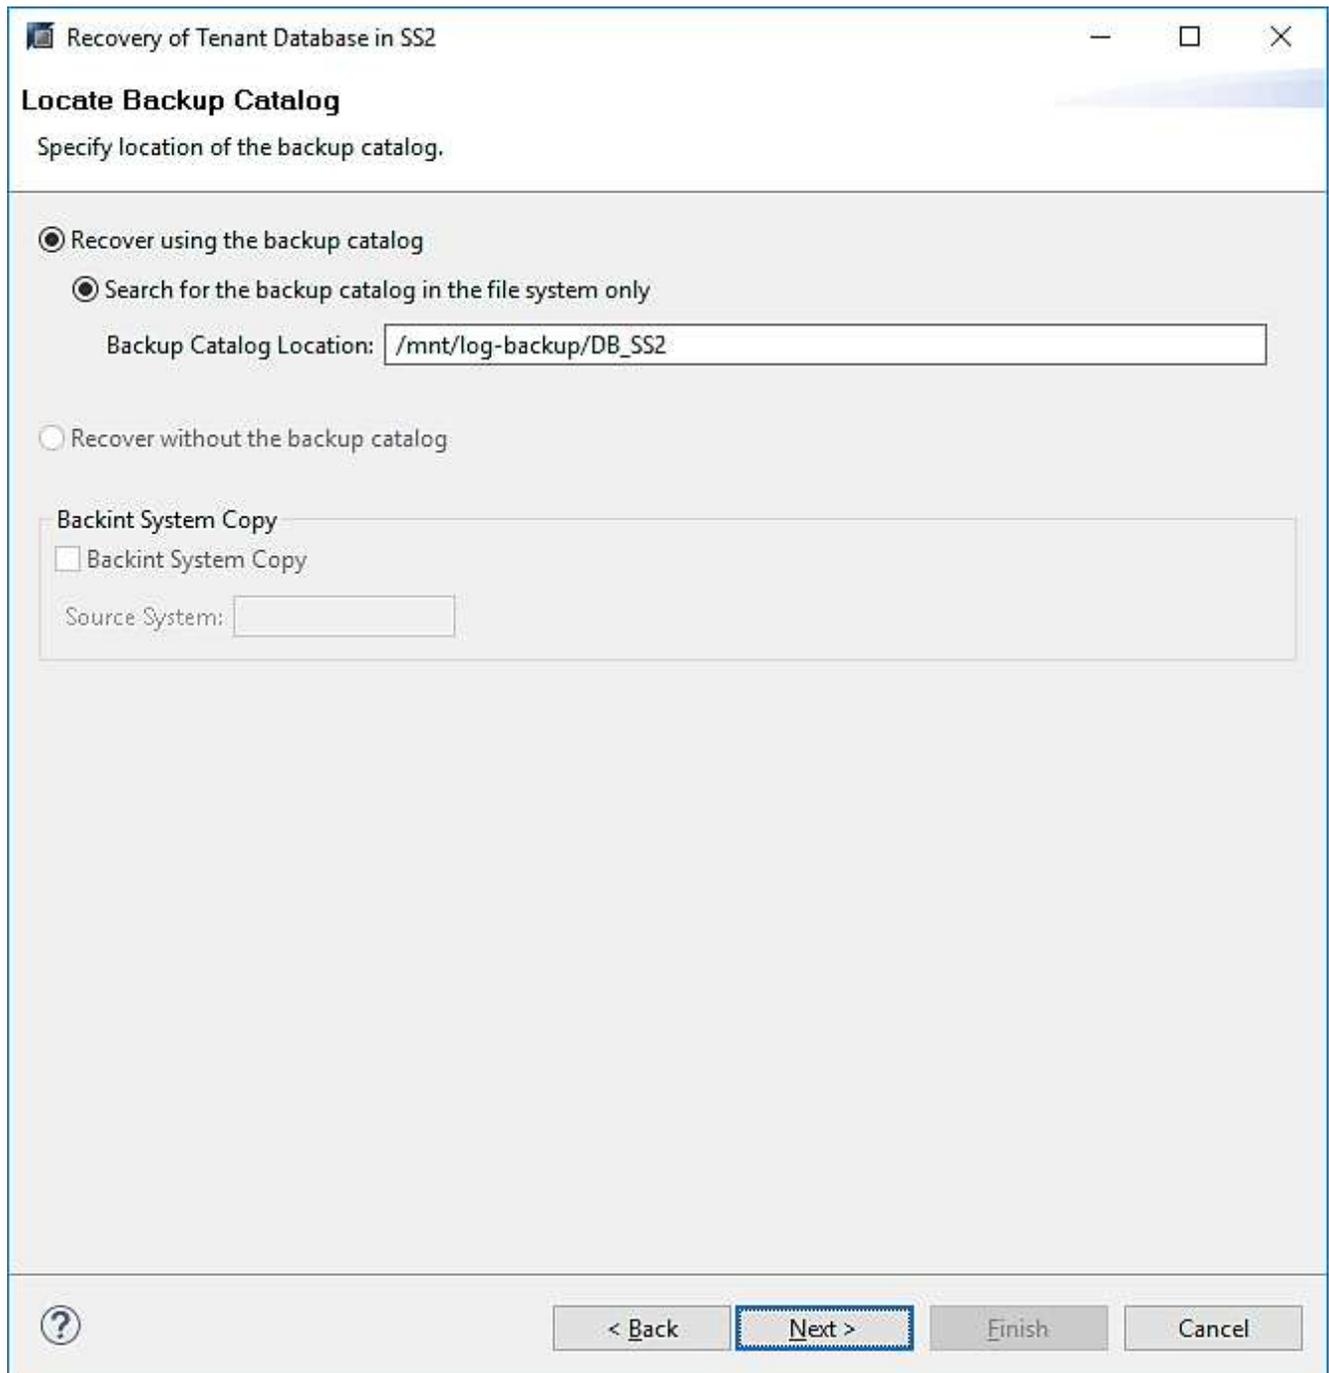
22. Select the tenant to recover and click Next.



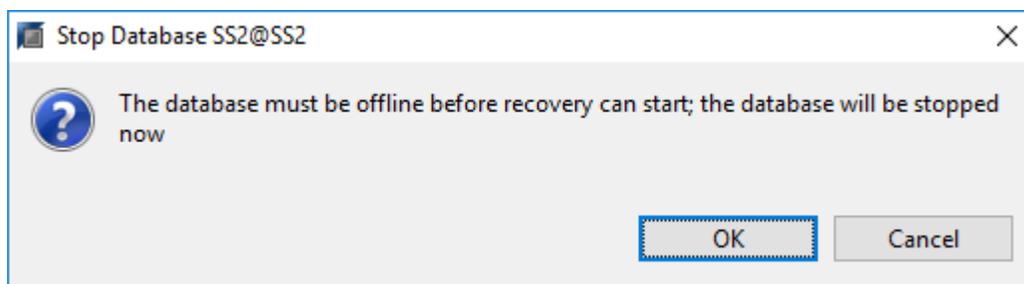
23. Specify the recovery type and click Next.



24. Confirm the backup catalog location and click Next.



25. Confirm that the tenant database is offline. Click OK to continue.



26. Because the restore of the data volume has occurred before the recovery of the system database, the tenant backup is immediately available. Select the backup highlighted in green and click Next.

Recovery of Tenant Database in SS2

### Select a Backup

Select a backup to recover the SAP HANA database

**Selected Point in Time**  
Database will be recovered to its most recent state.

**Backups**  
The overview shows backups that were recorded in the backup catalog as successful. The backup at the top is estimated to have the shortest recovery time.

Start Time	Location	Backup Prefix	Available
2019-12-10 02:05:08	/hana/data/SS2	SNAPSHOT	●
2019-12-09 22:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 18:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 14:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 10:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 06:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-09 02:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 22:05:07	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 18:05:08	/hana/data/SS2	SNAPSHOT	✘
2019-12-08 14:05:08	/hana/data/SS2	SNAPSHOT	✘

Refresh Show More

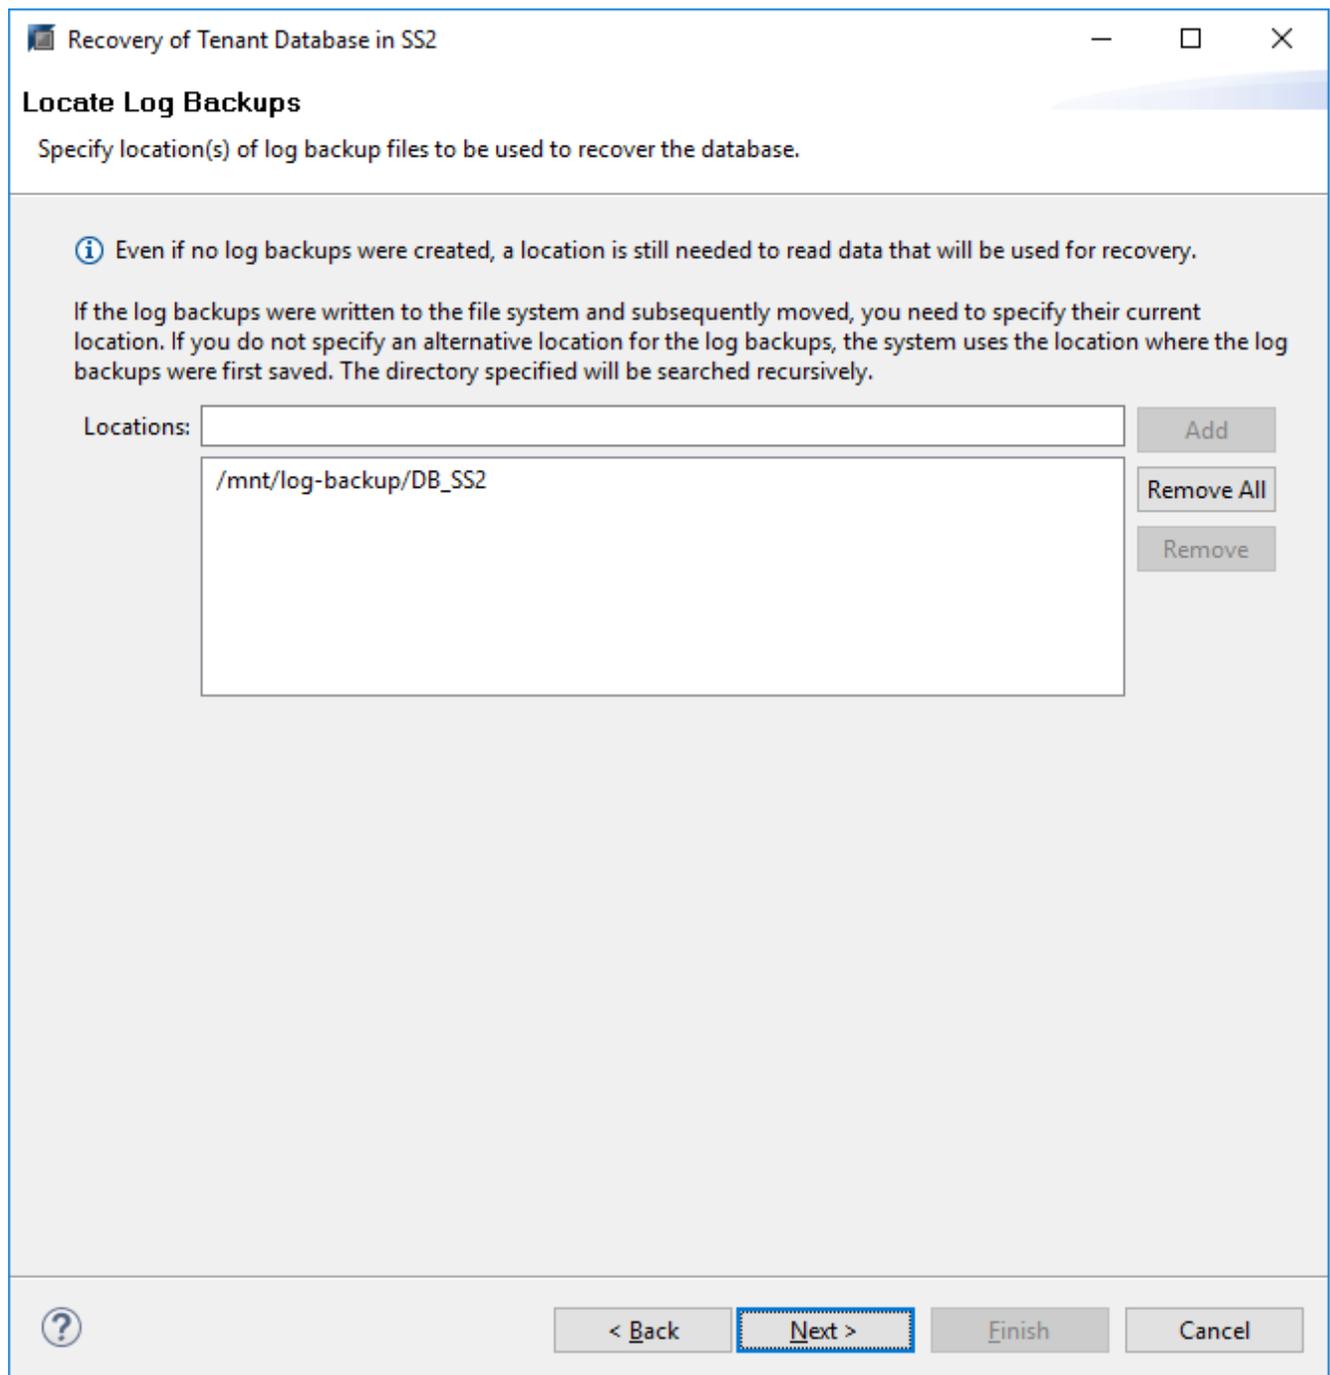
**Details of Selected Item**

Start Time: 2019-12-10 02:05:08 Destination Type: SNAPSHOT Source System: SS2@SS2  
 Size: 0 B Backup ID: 1575972308585 External Backup ID: SnapCenter\_LocalSnap\_Hourly\_12-10-2019\_02.05.01.3757  
 Backup Name: /hana/data/SS2  
 Alternative Location:

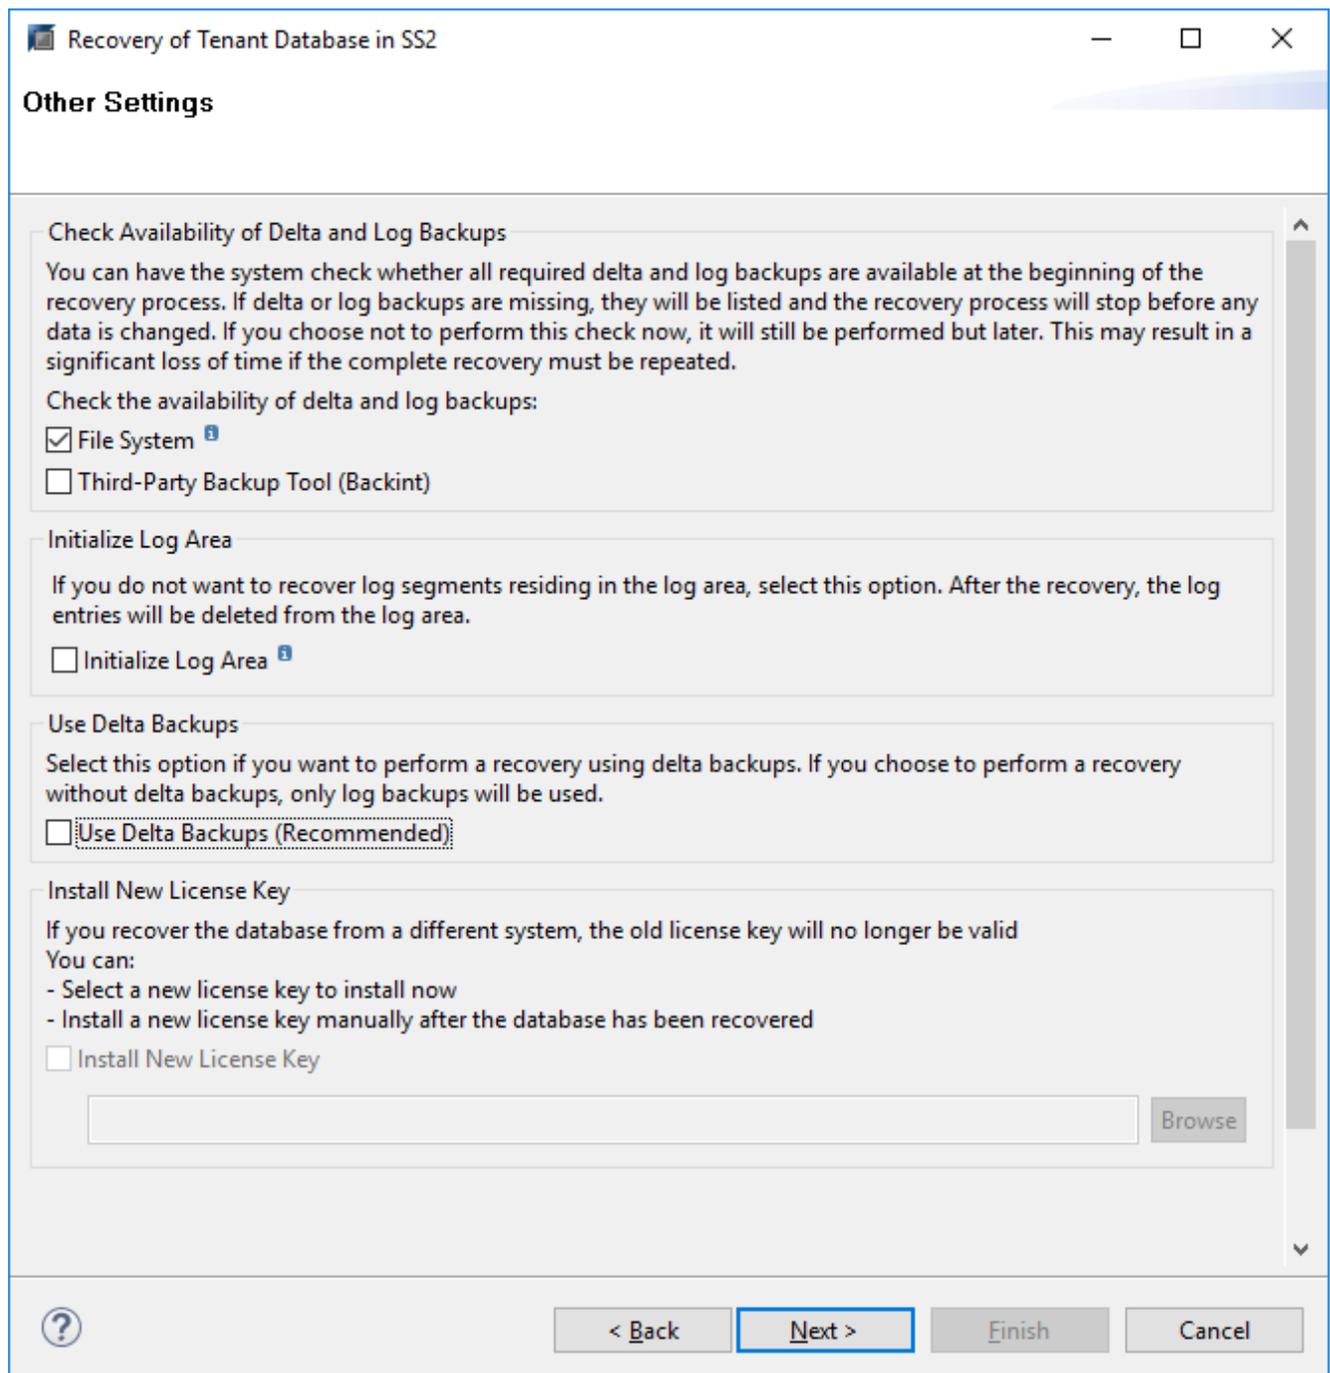
Check Availability

? < Back Next > Finish Cancel

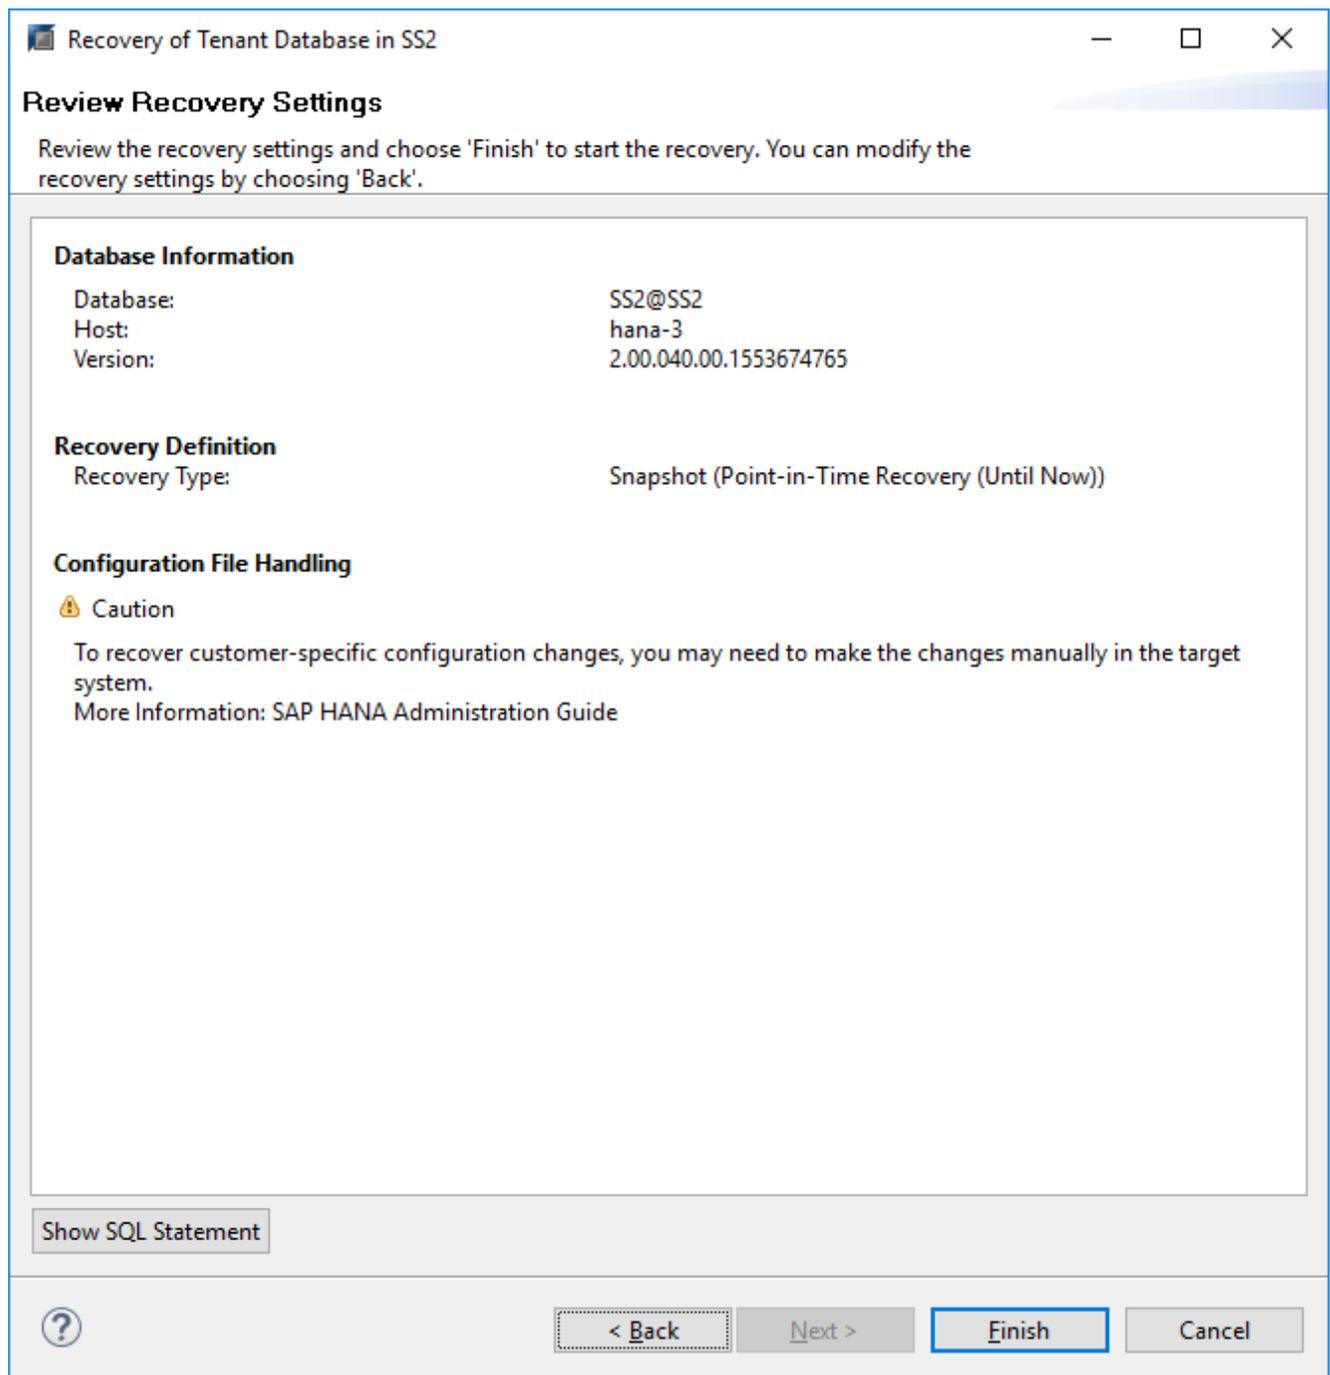
27. Confirm the log backup location and click Next.



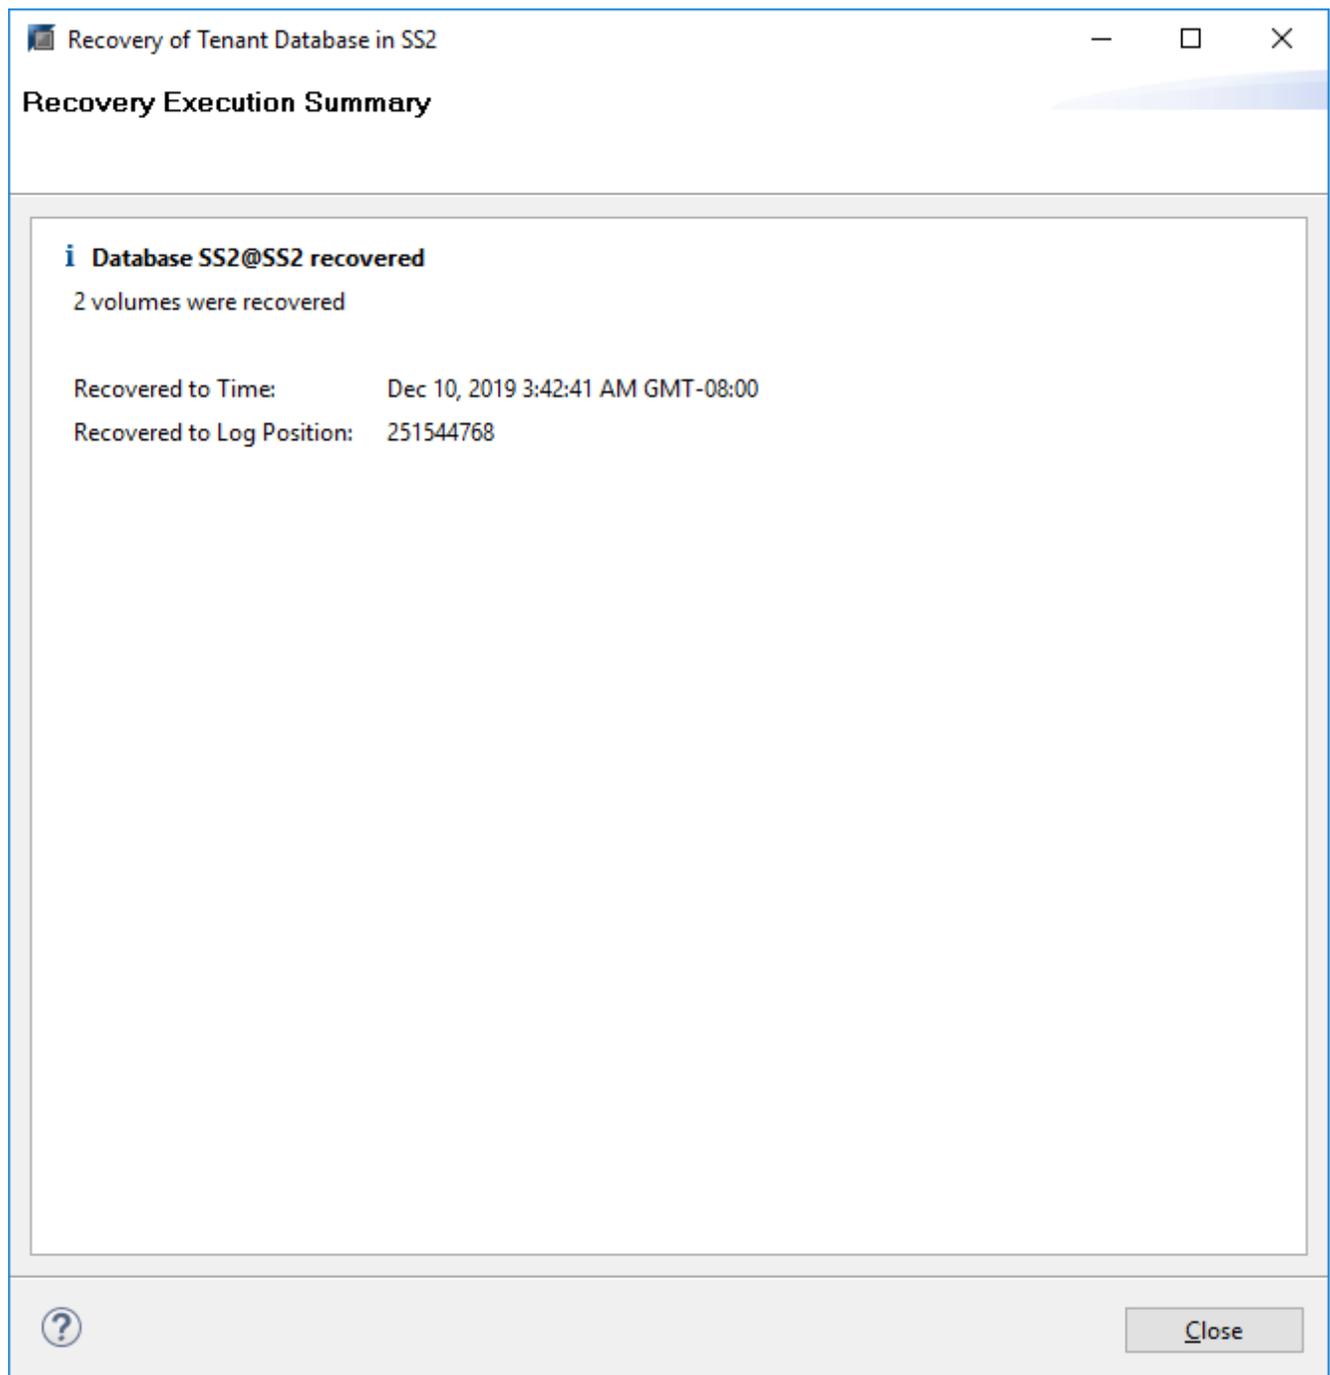
28. Select other settings as required. Make sure Use Delta Backups is not selected. Click Next.



29. Review the recovery settings and start the recovery process of the tenant database by clicking Finish.



30. Wait until the recovery has finished and the tenant database is started.



The SAP HANA system is up and running.



For an SAP HANA MDC system with multiple tenants, you must repeat steps 20–29 for each tenant.

## Advanced configuration and tuning

This section describes configuration and tuning options that customers may use to adapt the SnapCenter setup to their specific needs. Not all the settings may apply for all customer scenarios.

## Enable secure communication to HANA database

If the HANA databases are configured with secure communication, the `hdbsql` command that is executed by SnapCenter must use additional command-line options. This can be achieved by using a wrapper script which calls `hdbsql` with the required options.



There are various options to configure the SSL communication. In the following examples, the simplest client configuration is described using the command line option, where no server certificate validation is done. If certificate validation on server and/or client side is required, different `hdbsql` command line options are needed, and you must configure the PSE environment accordingly as described in the SAP HANA Security Guide.

Instead of configuring the `hdbsql` executable in the `hana.properties` files, the wrapper script is added.

For a central HANA plug-in host on the SnapCenter Windows server, you must add the following content in `C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\hana.properties`.

```
HANA_HDBSQL_CMD=C:\\Program Files\\sap\\hdbclient\\hdbsql-ssl.cmd
```

The wrapper script `hdbsql-ssl.cmd` calls `hdbsql.exe` with the required command-line options.

```
@echo off
"C:\Program Files\sap\hdbclient\hdbsql.exe" -e -ssltrustcert %*
```



The `-e -ssltrustcert hdbsql` command-line option also works for HANA systems where SSL is not enabled. This option can therefore also be used with a central HANA plug-in host, where not all HANA systems have SSL enabled or disabled.

If the HANA plug-in is deployed on individual HANA database hosts, the configuration must be done on each Linux host accordingly.

```
HANA_HDBSQL_CMD = /usr/sap/SM1/HDB12/exe/hdbsqls
```

The wrapper script `hdbsqls` calls `hdbsql` with the required command-line options.

```
#!/bin/bash
/usr/sap/SM1/HDB12/exe/hdbsql -e -ssltrustcert $*
```

## Disable auto discovery on the HANA plug-in host

To disable autodiscovery on the HANA plug-in host, complete the following steps:

1. On the SnapCenter Server, open PowerShell. Connect to the SnapCenter Server by running the `Open-SmConnection` command and specify the username and password in the opening login window.
2. To disable auto discovery, run the `Set-SmConfigSettings` command.

For a HANA host hana-2, the command is as follows:

```
PS C:\Users\administrator.SAPCC> Set-SmConfigSettings -Agent -Hostname
hana-2 -configSettings @{"DISABLE_AUTO_DISCOVERY"="true"}
Name                               Value
----                               -
DISABLE_AUTO_DISCOVERY            true
PS C:\Users\administrator.SAPCC>
```

3. Verify the configuration by running the `Get-SmConfigSettings` command.

```
PS C:\Users\administrator.SAPCC> Get-SmConfigSettings -Agent -Hostname
hana-2 -key all
Key: CUSTOMPLUGINS_OPERATION_TIMEOUT_IN_MSEC           Value: 3600000
Details: Plug-in API operation Timeout
Key: CUSTOMPLUGINS_HOSTAGENT_TO_SERVER_TIMEOUT_IN_SEC  Value: 1800
Details: Web Service API Timeout
Key: CUSTOMPLUGINS_ALLOWED_CMDS                       Value: *;
Details: Allowed Host OS Commands
Key: DISABLE_AUTO_DISCOVERY                           Value: true
Details:
Key: PORT                                               Value: 8145
Details: Port for server communication
PS C:\Users\administrator.SAPCC>
```

The configuration is written to the agent configuration file on the host and is still available after a plug-in upgrade with SnapCenter.

```
hana-2:/opt/NetApp/snapcenter/scc/etc # cat
/opt/NetApp/snapcenter/scc/etc/agent.properties | grep DISCOVERY
DISABLE_AUTO_DISCOVERY = true
hana-2:/opt/NetApp/snapcenter/scc/etc #
```

### Deactivate automated log backup housekeeping

Log backup housekeeping is enabled by default and can be disabled on the HANA plug-in host level. There are two options to change these settings.

#### Edit the hana.property file

Including the parameter `LOG_CLEANUP_DISABLE = Y` in the `hana.property` configuration file disables the log backup housekeeping for all resources using this SAP HANA plug-in host as communication host:

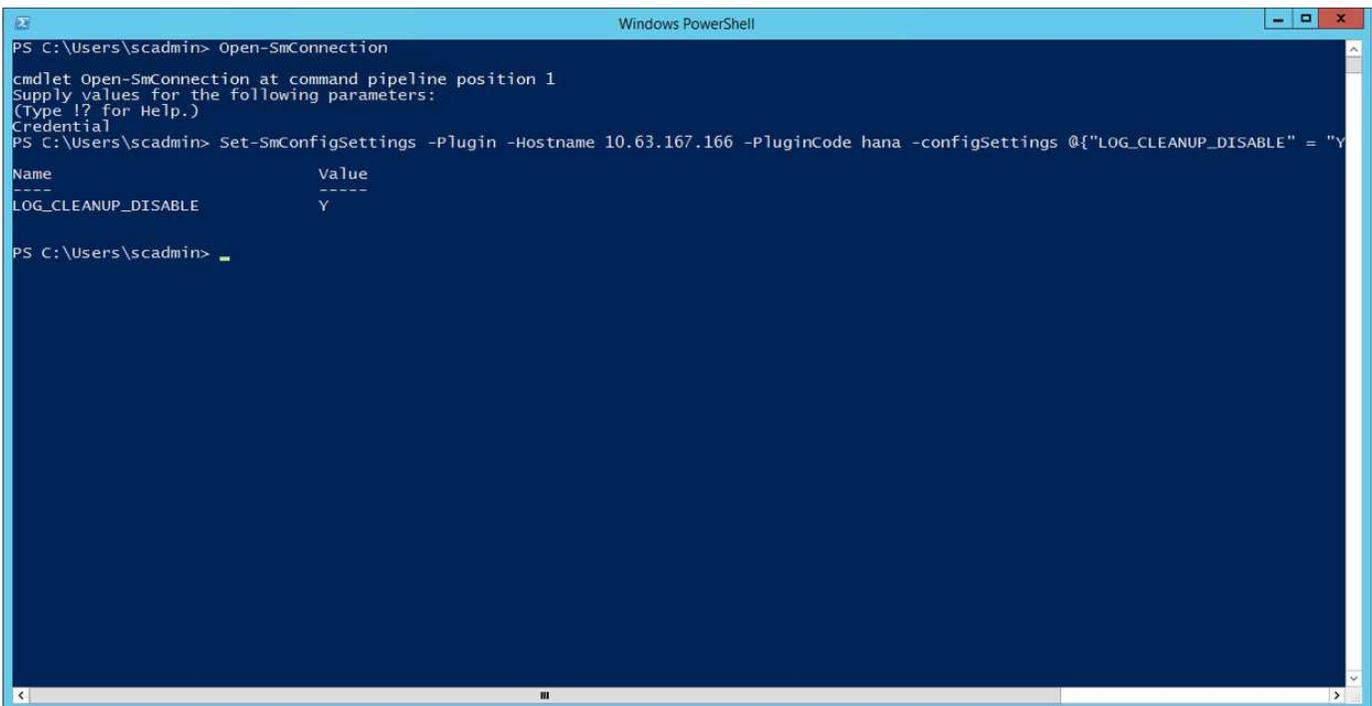
- For the Hdbsql communication host on Windows, the `hana.property` file is located at `C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc`.

- For the Hdbsql communication host on Linux, the `hana.property` file is located at `/opt/NetApp/snapcenter/scc/etc`.

### Use the PowerShell command

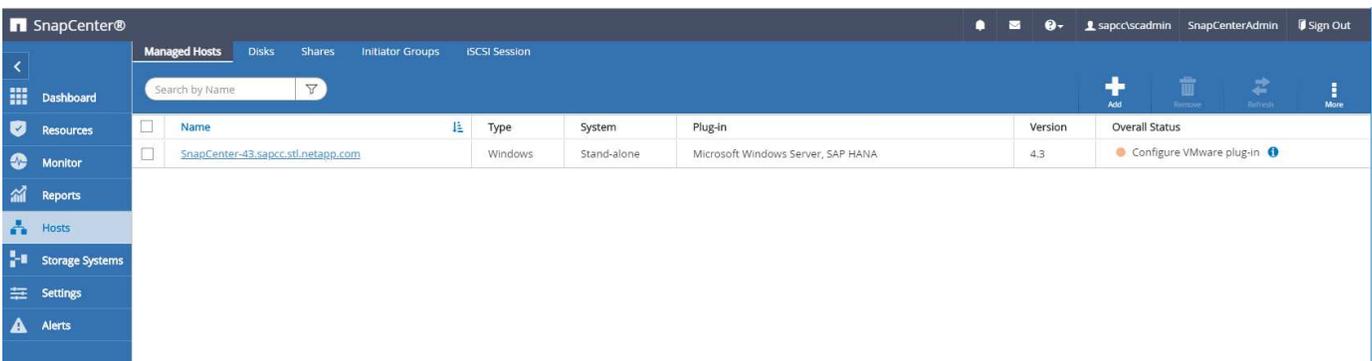
A second option to configure these settings is using a SnapCenter PowerShell command.

1. On the SnapCenter server, open a PowerShell. Connect to the SnapCenter server using the command `Open-SmConnection` and specify user name and password in the opening login window.
2. With the command `Set-SmConfigSettings -Plugin -HostName <pluginhostname> -PluginCode hana -configSettings @{"LOG_CLEANUP_DISABLE" = "Y"}`, the changes are configured for the SAP HANA plug-in host `<pluginhostname>` specified by the IP or host name (see the following figure).



### Disable warning when running SAP HANA plug-in on a virtual environment

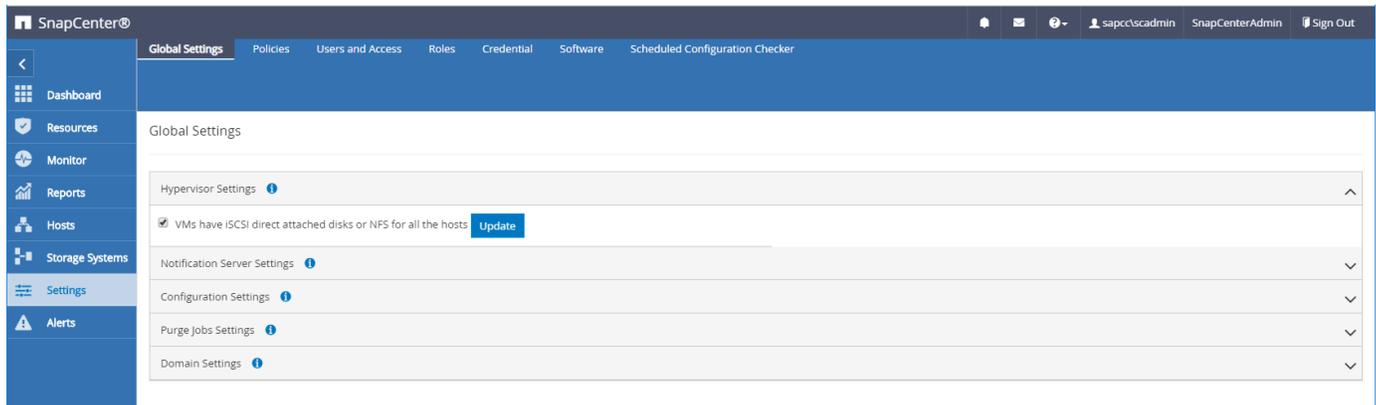
SnapCenter detects if the SAP HANA plug-in is installed on a virtualized environment. This could be a VMware environment or a SnapCenter installation at a public cloud provider. In this case, SnapCenter displays a warning to configure the hypervisor, as shown in the following figure.



It is possible to suppress this warning globally. In this case, SnapCenter is not aware of virtualized environments and, therefore, does not show these warnings.

To configure SnapCenter to suppress this warning, the following configuration must be applied:

1. From the Settings tab, select Global Settings.
2. For the hypervisor settings, select VMs Have iSCSI Direct Attached Disks or NFS For All the Hosts and update the settings.



### Change scheduling frequency of backup synchronization with off-site backup storage

As described in the section [“Retention management of backups at the secondary storage,”](#) retention management of data backups to an off-site backup storage is handled by ONTAP. SnapCenter periodically checks if ONTAP has deleted backups at the off-site backup storage by running a cleanup job with a weekly default schedule.

The SnapCenter cleanup job deletes backups in the SnapCenter repository as well as in the SAP HANA backup catalog if any deleted backups at the off-site backup storage have been identified.

The cleanup job also executes the housekeeping of SAP HANA log backups.

Until this scheduled cleanup has finished, SAP HANA and SnapCenter might still show backups that have already been deleted from the off-site backup storage.



This might result in additional log backups that are kept, even if the corresponding storage-based Snapshot backups on the off-site backup storage have already been deleted.

The following sections describe two ways to avoid this temporary discrepancy.

#### Manual refresh on resource level

In the topology view of a resource, SnapCenter displays the backups on the off-site backup storage when selecting the secondary backups, as shown in the following screenshot. SnapCenter executes a cleanup operation with the Refresh icon to synchronize the backups for this resource.

## Change the frequency of the SnapCenter cleanup job

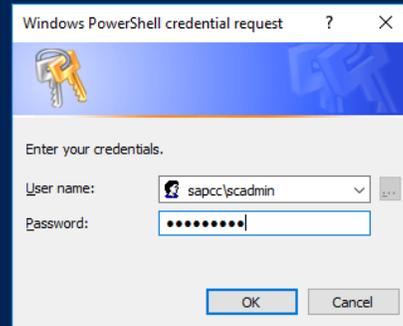
SnapCenter executes the cleanup job `SnapCenter_RemoveSecondaryBackup` by default for all resources on a weekly basis using the Windows task scheduling mechanism. This can be changed using a SnapCenter PowerShell cmdlet.

1. Start a PowerShell command window on the SnapCenter Server.
2. Open the connection to the SnapCenter Server and enter the SnapCenter administrator credentials in the login window.

```
Windows PowerShell
Copyright (C) 2016 Microsoft Corporation. All rights reserved.

PS C:\Users\scadmin> Open-SmConnection

cmdlet Open-SmConnection at command pipeline position 1
Supply values for the following parameters:
(Type !? for Help.)
Credential
```



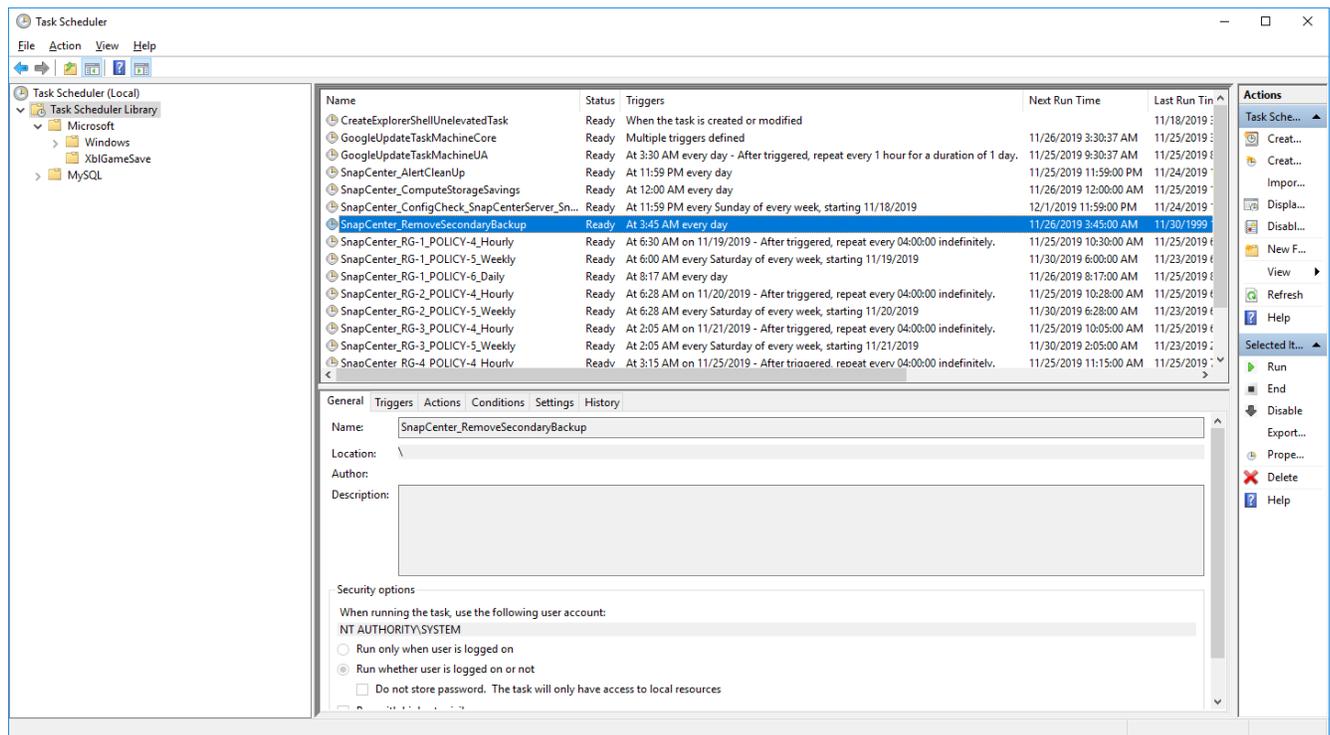
3. To change the schedule from a weekly to a daily basis, use the cmdlet `Set-SmSchedule`.

```

PS C:\Users\scadmin> Set-SmSchedule -ScheduleInformation
@{"ScheduleType"="Daily";"StartTime"="03:45 AM";"DaysInterval"=
"1"} -TaskName SnapCenter_RemoveSecondaryBackup
TaskName           : SnapCenter_RemoveSecondaryBackup
Hosts               : {}
StartTime          : 11/25/2019 3:45:00 AM
DaysOfTheMonth     :
MonthsOfTheYear    :
DaysInterval       : 1
DaysOfTheWeek      :
AllowDefaults      : False
ReplaceJobIfExists : False
UserName           :
Password           :
SchedulerType      : Daily
RepeatTask_Every_Hour :
IntervalDuration   :
EndTime            :
LocalScheduler     : False
AppType            : False
AuthMode           :
SchedulerSQLInstance : SMCoreContracts.SmObject
MonthlyFrequency   :
Hour               : 0
Minute             : 0
NodeName           :
ScheduleID         : 0
RepeatTask_Every_Mins :
CronExpression     :
CronOffsetInMinutes :
StrStartTime       :
StrEndTime         :
PS C:\Users\scadmin> Check the configuration using the Windows Task
Scheduler.

```

4. You can check the job properties in Windows task scheduler.



## Where to find additional information and version history

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

- SnapCenter Resources Page

<https://www.netapp.com/us/documentation/snapcenter-software.aspx>

- SnapCenter Software Documentation

<https://docs.netapp.com/us-en/snapcenter/index.html>

- TR-4667: Automating SAP System Copies Using the SnapCenter

[Automating SAP System Copies Using the SnapCenter](#)

- TR-4719: SAP HANA System Replication, Backup and Recovery with SnapCenter

[SAP HANA System Replication, Backup and Recovery with SnapCenter](#)

- TR-4018: Integrating NetApp ONTAP Systems with SAP Landscape Management

<https://www.netapp.com/pdf.html?item=/media/17195-tr4018pdf.pdf>

- TR-4646: SAP HANA Disaster Recovery with Storage Replication

<https://www.netapp.com/pdf.html?item=/media/8584-tr4646pdf.pdf>

## Version history

Version	Date	Document version history
Version 1.0	July 2017	<ul style="list-style-type: none"><li>• Initial release.</li></ul>
Version 1.1	September 2017	<ul style="list-style-type: none"><li>• Added the section “Advanced Configuration and Tuning.”</li><li>• Minor corrections.</li></ul>
Version 2.0	March 2018	<ul style="list-style-type: none"><li>• Updates to cover SnapCenter 4.0: New data volume resource Improved Single File SnapRestore operation</li></ul>
Version 3.0	January 2020	<ul style="list-style-type: none"><li>• Added the section “SnapCenter Concepts and Best Practices.”</li><li>• Updates to cover SnapCenter 4.3: Automatic discovery Automated restore and recovery Support of HANA MDC multiple tenants Single-tenant restore operation</li></ul>
Version 3.1	July 2020	<ul style="list-style-type: none"><li>• Minor updates and corrections: NFSv4 support with SnapCenter 4.3.1 Configuration of SSL communication Central plug-in deployment for Linux on IBM Power</li></ul>
Version 3.2	November 2020	<ul style="list-style-type: none"><li>• Added the required database user privileges for HANA 2.0 SPS5.</li></ul>
Version 3.3	May 2021	<ul style="list-style-type: none"><li>• Updated the SSL hdbsql configuration section.</li><li>• Added Linux LVM support.</li></ul>
Version 3.4	August 2021	<ul style="list-style-type: none"><li>• Added the disable auto discovery configuration description.</li></ul>

Version	Date	Document version history
Version 3.5	February 2022	<ul style="list-style-type: none"><li>Minor updates to cover SnapCenter 4.6 and auto discovery support for HANA System Replication-enabled HANA systems.</li></ul>

## SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster

### SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster

This document provides best practices for data protection with SnapCenter in a VMware environment combined with SnapMirror active sync as a high availability solution for the HANA storage resources.

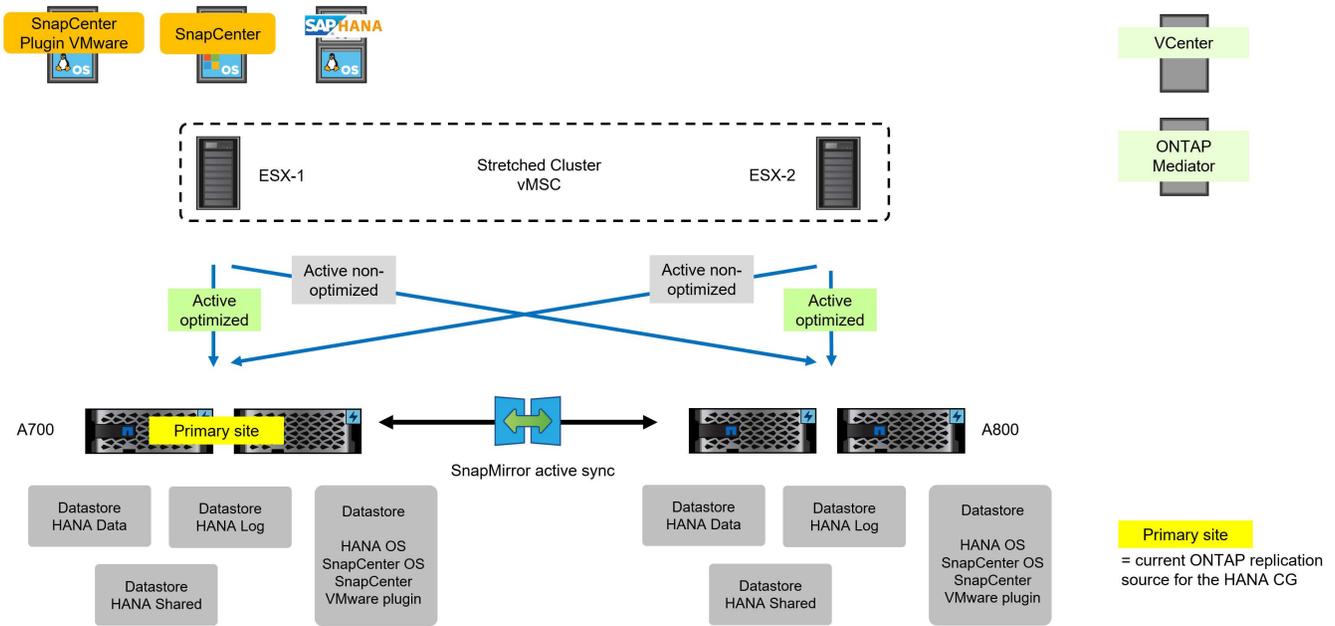
Author: Nils Bauer, NetApp

#### Scope of this document

The document is not intended to be a step-by-step description of how to setup the complete environment but will cover concepts and relevant details related to:

- Setup of SAP HANA systems with VMware VMFS
- SnapMirror active sync configuration for SAP HANA
- SnapCenter configuration for HANA on VMware with VMFS
- SnapCenter configuration for SnapMirror active sync
- SnapCenter operations with HANA on VMware and SnapMirror active sync

We will focus on a VMware Metro Storage Cluster (vMSC) configuration using a uniform access setup of SnapMirror active sync as shown in the figure below, but we will also briefly touch bare metal as well as non-uniform access configurations.



## Overview SAP HANA high availability

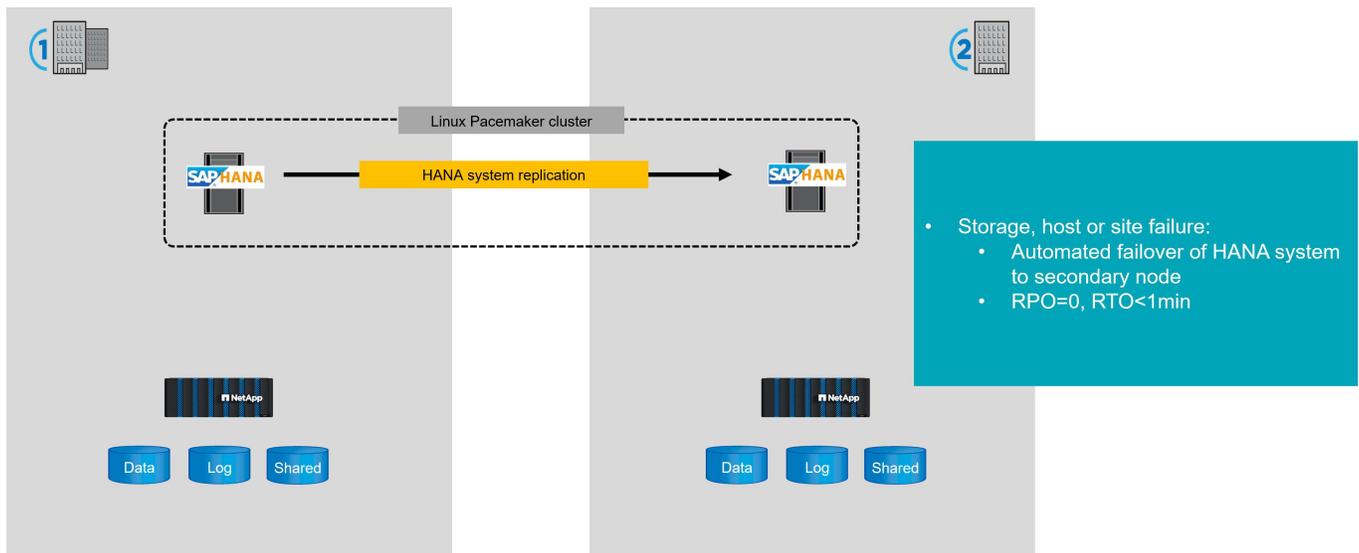
This chapter provides an overview of high availability options for SAP HANA comparing replication on application layer with storage replication.

### SAP HANA system replication (HSR)

SAP HANA system replication offers an operation mode in which the data is replicated synchronously, preloaded into memory and continuously updated at the secondary host. This mode enables very low RTO values, approximately 1 minute or less, but it also requires a dedicated server that is only used to receive the replication data from the source system. Because of the low failover time, SAP HANA system replication is also often used for near-zero-downtime maintenance operations, such as HANA software upgrades. Linux Pacemaker cluster solutions are typically used to automate failover operations.

In case of any failure at the primary site, storage, host or complete site, the HANA system automatically fails over to the secondary site controlled by the Linux Pacemaker cluster.

For a full description of all configuration options and replication scenarios, see [SAP HANA System Replication | SAP Help Portal](#).



### NetApp SnapMirror active sync

SnapMirror active sync enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. There is no manual intervention or custom scripting required to trigger a failover with SnapMirror active sync. SnapMirror active sync is supported on AFF clusters, All-Flash SAN Array (ASA) clusters, and C-Series (AFF or ASA). SnapMirror active sync protects applications with iSCSI or FCP LUNs.

Beginning with ONTAP 9.15.1, SnapMirror active sync supports a symmetric active/active capability. Symmetric active/active enable read and write I/O operations from both copies of a protected LUN with bidirectional synchronous replication so that both LUN copies can serve I/O operations locally.

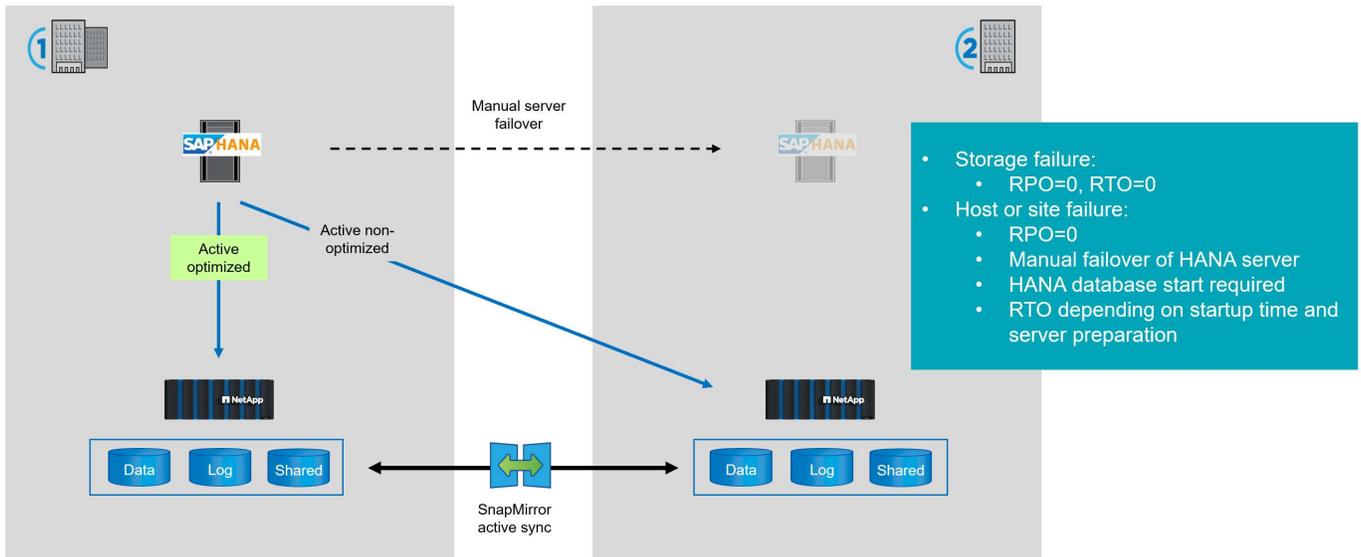
More details can be found at [SnapMirror active sync overview in ONTAP](#).

### HANA bare metal

When running SAP HANA on a bare metal server, you can use SnapMirror active sync to provide a high available storage solution. The data is replicated synchronously therefore providing an RPO=0.

In case of a storage failure, the HANA system will transparently access the mirrored copy at the secondary site using the second FCP path providing an RTO=0.

In case of a host or complete site failure, a new server at the secondary site needs to be provided to access the data from the failed host. This would typically be a test or QA system of the same size as production which will now be shut down and be used to run the production system. After the LUNs at the secondary site are connected to the new host, the HANA database needs to be started. The total RTO therefore depends on the time needed to provision the host and the startup time of the HANA database.



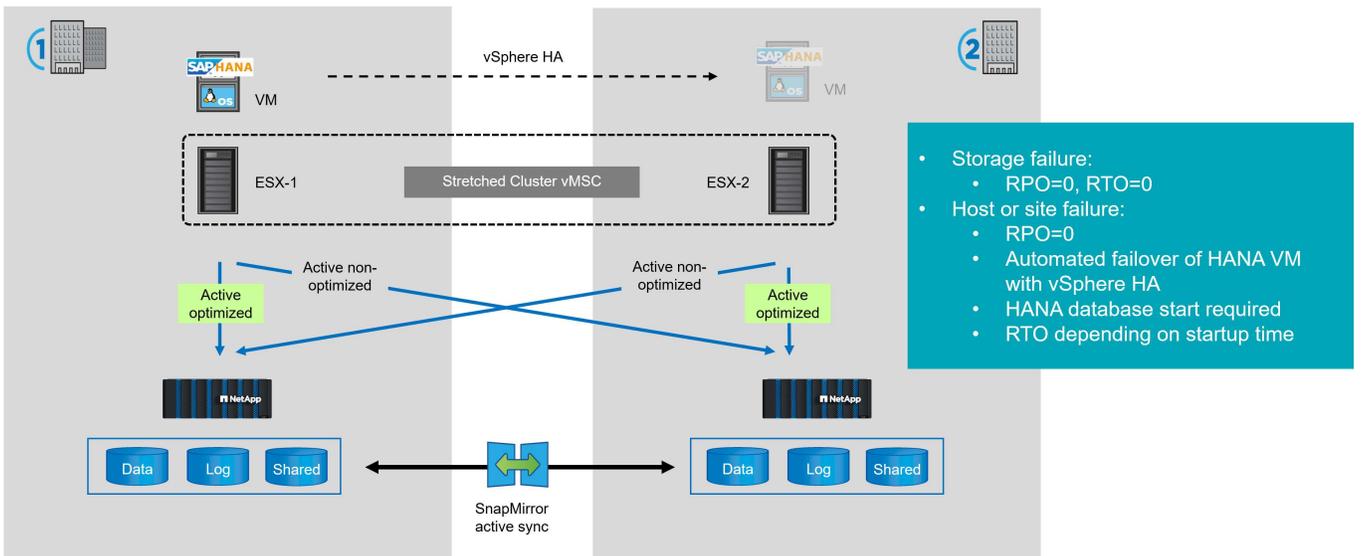
- Storage failure:
  - RPO=0, RTO=0
- Host or site failure:
  - RPO=0
  - Manual failover of HANA server
  - HANA database start required
  - RTO depending on startup time and server preparation

### vSphere Metro Storage Cluster (vMSC)

When running SAP HANA in a VMware environment using FCP attached datastores you can use SnapMirror active sync to build a VMware Metro Storage Cluster. In such a setup the datastores used by the HANA system are replicated synchronously to the secondary site.

In case of a storage failure, the ESX host will automatically access the mirrored copy at the secondary site providing an RTO=0.

In case of a host or complete site failure, vSphere HA is used to start the HANA VM at the secondary ESX host. When the HANA VM is running, the HANA database needs to be started. The total RTO therefore mainly depends on the startup time of the HANA database.



- Storage failure:
  - RPO=0, RTO=0
- Host or site failure:
  - RPO=0
  - Automated failover of HANA VM with vSphere HA
  - HANA database start required
  - RTO depending on startup time

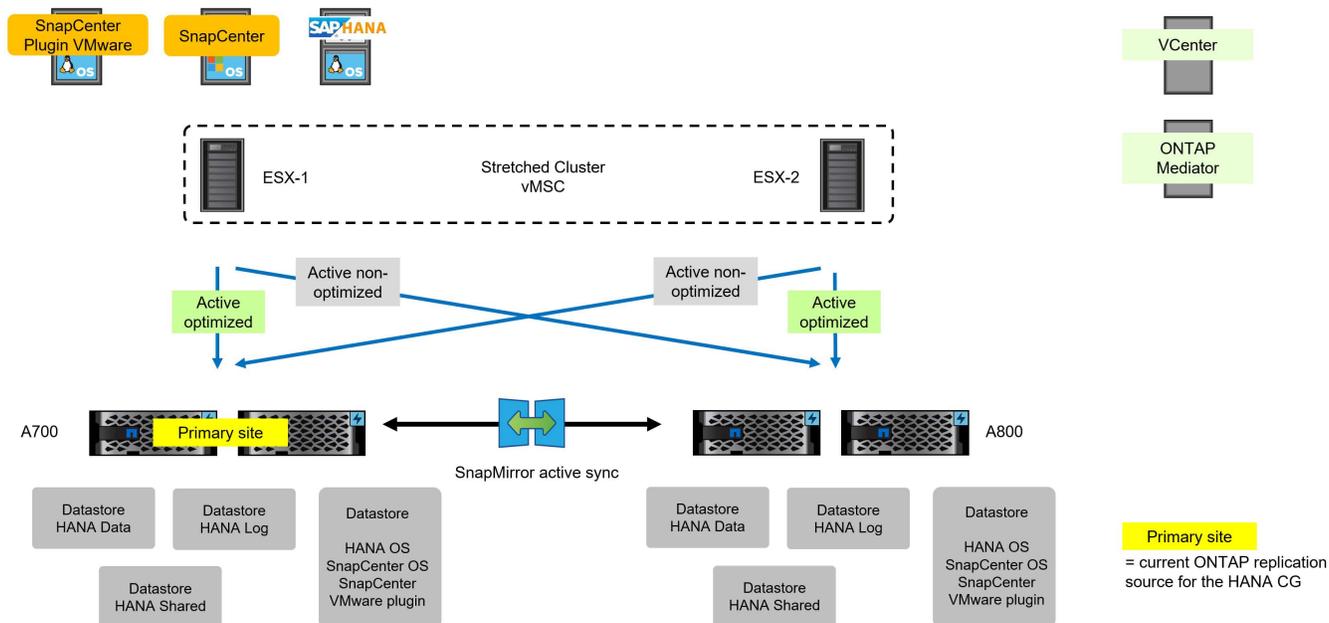
### Solution comparison

The following table provides a summary of the key characteristics of the solutions described above.

	<b>HANA System Replication</b>	<b>SnapMirror active sync – bare metal</b>	<b>SnapMirror active sync – VMware vMSC</b>
RPO with any failure	RPO=0 Synchronous replication		
RTO with storage failure	RTO < 1min	RTO=0 Transparent storage failover	
RTO with site or host failure	RTO < 1min	RTO: Depending on the time required for server preparation and HANA database startup.	RTO: Depending on the time required for HANA database startup.
Failover automation	Yes,  automated failover to secondary HSR host  controlled by pacemaker cluster.	Yes, for storage failure  No, for host or site failure  (Provisioning of host, connect storage resources, HANA database start)	Yes, for storage failure  Yes, for host or site failure  (Failover of VM to other site automated with vSphere HA, HANA database start)
Dedicated server at secondary site required	Yes,  required to preload data into memory and enable fast failover w/o database startup.	No,  server is only required in case of failover. Typically, the server used for QA would then be used for production.	No,  Resources at ESX host are only required in case of a failover. Typically, QA resources would then be used for production.

## Example configuration overview

In the lab setup, we are using a uniform access configuration, where both ESX hosts have access to both storage clusters. Within the next sections we describe the uniform access configuration but also highlight the differences for a non-uniform setup.



## Software versions

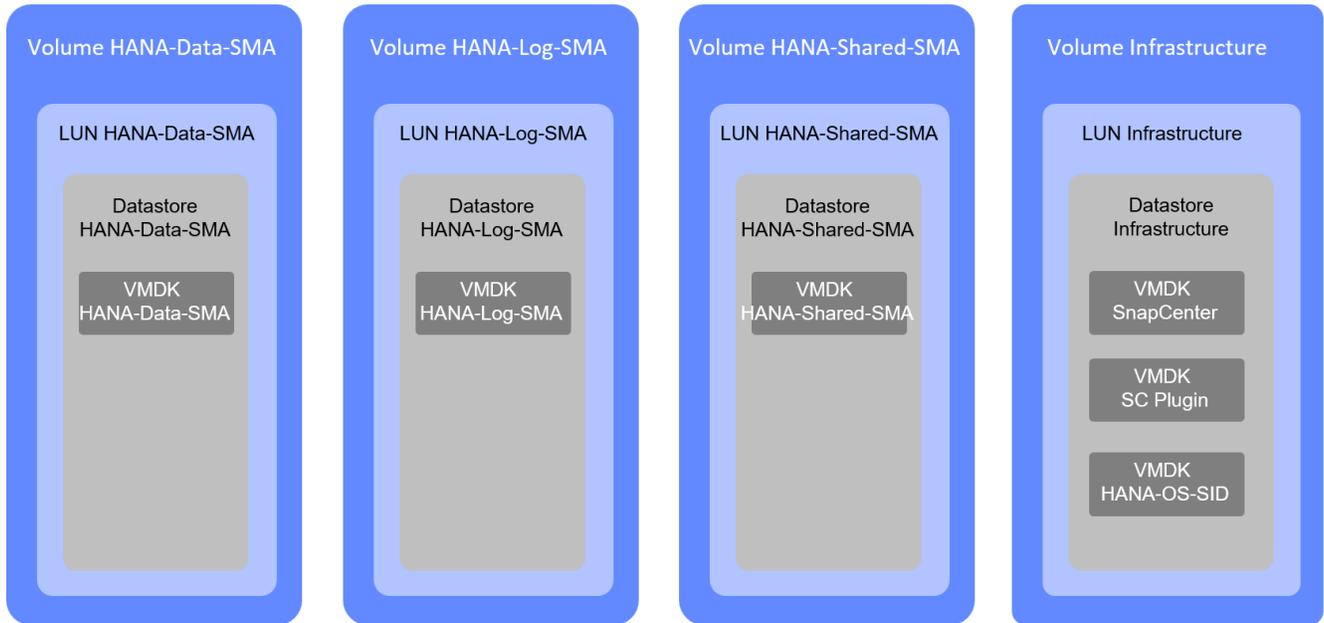
Software	Version
ONTAP	A700: 9.15.1P7, A800: 9.16.1RC1
vSphere client	8.0.3
ESXi	8.0.3
SnapCenter plugin for vSphere	6.0.1
Linux OS	SLES for SAP 15 SP5
SAP HANA	2.0 SPS8
SnapCenter	6.0.1

## HANA system provisioning and installation

This chapter describes the installation and configuration of the SAP HANA system specific to a VMware setup using VMFS. Additional generic best practices can be found at [SAP HANA on NetApp AFF Systems with Fibre Channel Protocol](#).

### Storage configuration

The figure below shows the storage and datastore configuration for the HANA system. You must configure a dedicated volume, LUN, datastore for each filesystem of the HANA system. Datastores must not be shared across multiple HANA systems or other workloads.



All three LUNs of the HANA system (hana\_data\_SMA, hana\_log\_SAM and hana\_shared\_SMA) as well as the LUN for the OS images and SnapCenter components have been provisioned at the A700 storage cluster.



All volumes of the HANA system must be provisioned in the same SVM. In the SnapMirror active sync configuration described later, we will create a consistency group across all three HANA volumes, which requires that the volumes are in the same SVM. The infrastructure volume will be in a different consistency group and could therefore be in a different SVM.

ONTAP System Manager

Search actions, objects, and pages

LUNs

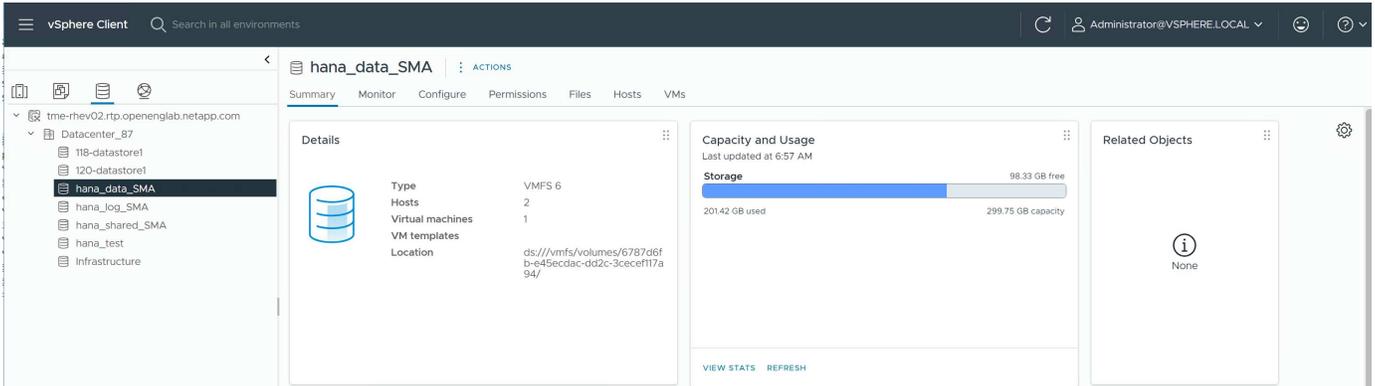
Name	Storage VM	Volume	Size	IOPS	Latency (ms)	Throughput (MB/s)
vvolPE-1724163990635	svm200_blueexpdr_a700s	vvol_FCoE_2	4 MiB	0	0	0
vvolPE-1724163990633	svm200_blueexpdr_a700s	vvol_FCoE_1	4 MiB	0	0	0
DraaS_qa_lun1	svm200_blueexpdr_a700s	DraaS_qa_lun1	200 GiB	0	0	0
DRaaS_qa_lun2	svm200_blueexpdr_a700s	DRaaS_qa_lun2	100 GiB	0	0	0
Infrastructure	svm200_blueexpdr_a700s	Infrastructure	2 TiB	50	0.31	0.58
hana_data_SMA	svm200_blueexpdr_a700s	hana_data_SMA	300 GiB	0	0.24	0
hana_log_SMA	svm200_blueexpdr_a700s	hana_log_SMA	158 GiB	0	0.24	0
hana_shared_SMA	svm200_blueexpdr_a700s	hana_shared_SMA	210 GiB	1	0.16	0.01
hana_test_lun	svm200_blueexpdr_a700s	hana_test_lun	1 TiB	0	0.39	0

Showing 1 - 9 of 9 LUNs

An initiator group must be configured, and the LUNs above must be mapped to the ESX-1 host, which is in close proximity to the A700 storage system in our lab setup.

## Datstore provisioning

We created three datastores for the HANA system using the three LUNs we have provisioned before. In addition, we created an infrastructure datastore using the infrastructure LUN.

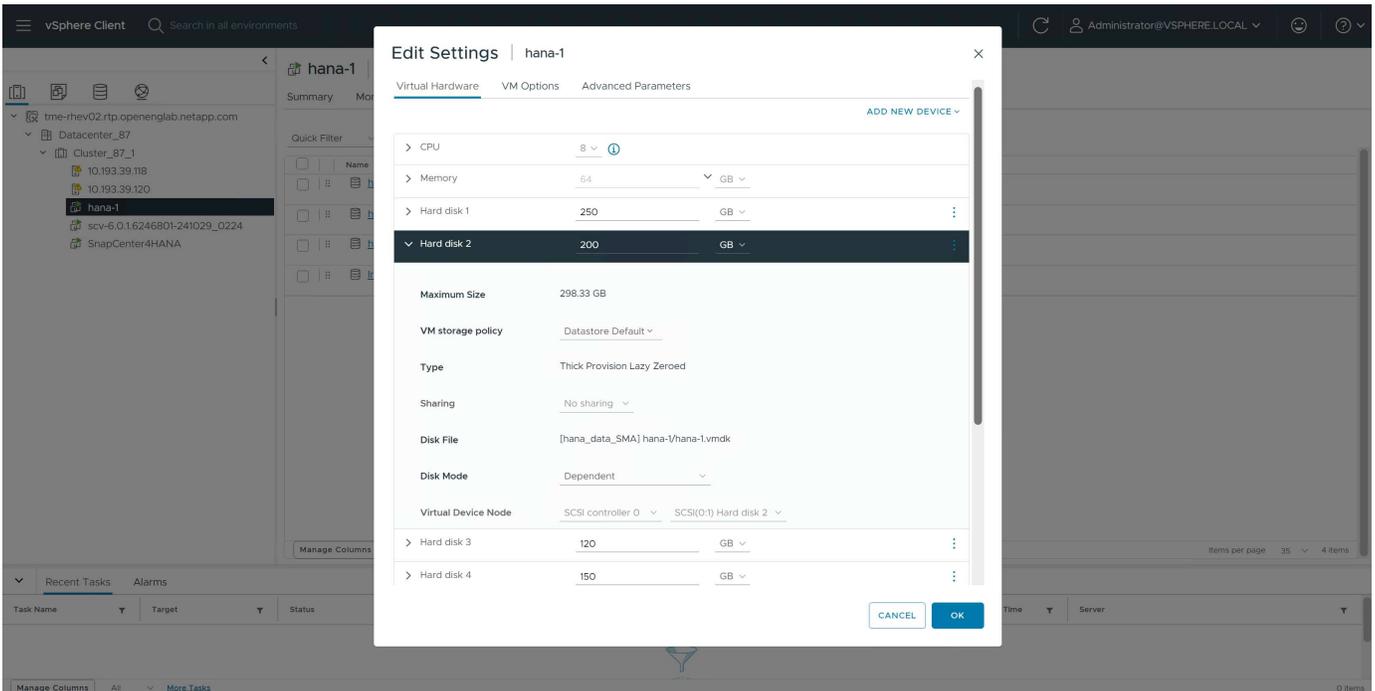


## VM provisioning and OS installation

In our lab setup we deployed a new VM and placed the VMDK for the Linux OS in the infrastructure datastore.

## VM disk configuration

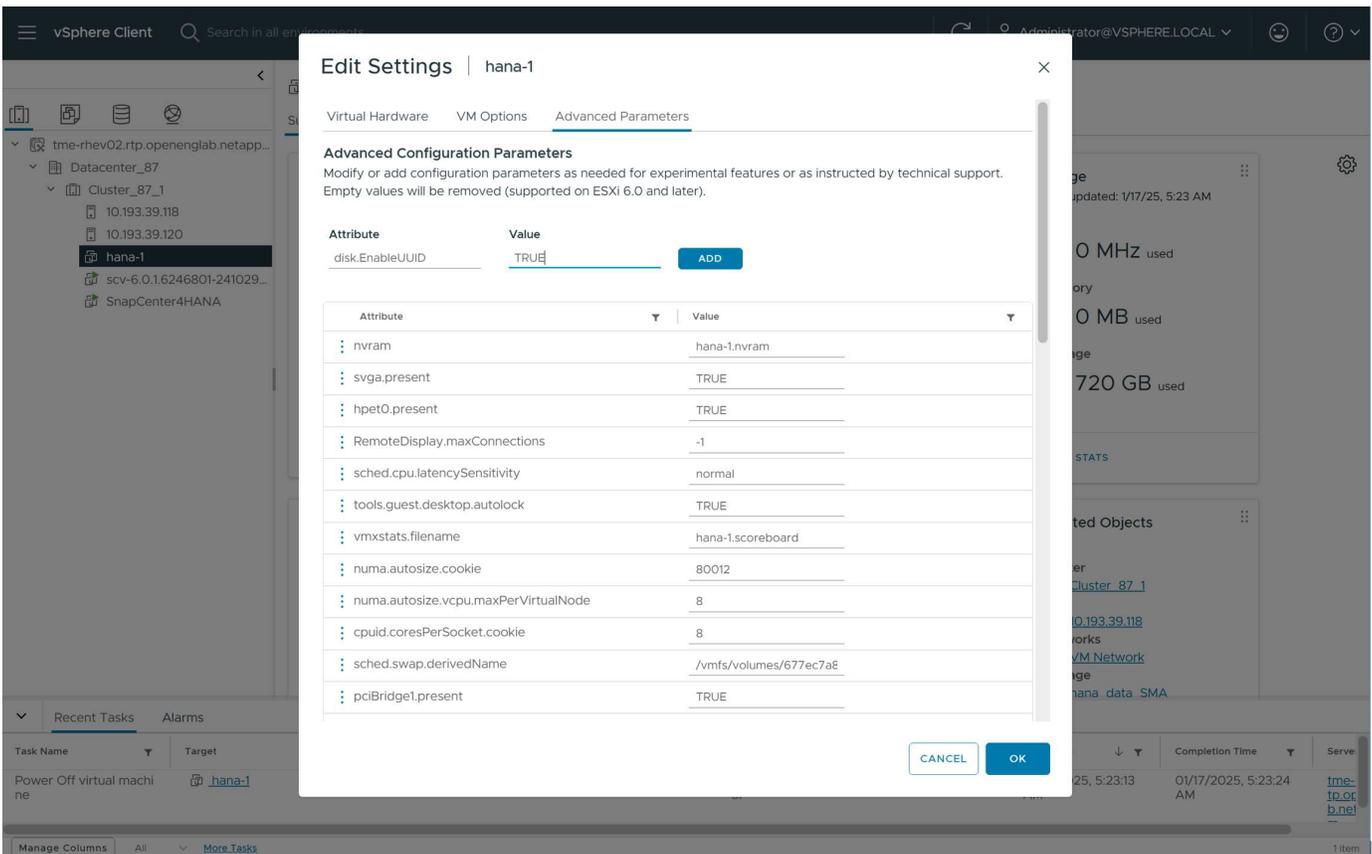
Three new disks have been added to the HANA VM, each disk within one of the datastores which have been created for the HANA system.



## VM parameter setting

The parameter `disk.EnableUUID` must be added and set to `TRUE`. The parameter is required by SnapCenter. If not set the SnapCenter "Discover virtual resource" operation will fail.

The VM must be stopped before parameter can be added.



The functionality can be checked with the command below.

```
hana-1:~ # sg_inq /dev/sdd
standard INQUIRY:
PQual=0 PDT=0 RMB=0 LU_CONG=0 hot_pluggable=0 version=0x06 [SPC-4]
[AERC=0] [TrmTsk=] NormACA=0 HiSUP=0 Resp_data_format=2
SCCS=0 ACC=0 TPGS=0 3PC=0 Protect=0 [BQue=0]
EncServ=0 MultiP=0 [MChngr=0] [ACKREQQ=0] Addr16=0
[RelAdr=0] WBus16=1 Sync=1 [Linked=0] [TranDis=0] CmdQue=1
length=36 (0x24) Peripheral device type: disk
Vendor identification: VMware
Product identification: Virtual disk
Product revision level: 2.0
Unit serial number: 6000c293fecf25ac6bc457af67fe1f54
```

## File system preparation at Linux host

### Creation of xfs filesystem on new disks

The device names of new the new disks can be checked with the command below.

```
hana-1:/install # lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINTS
sda 8:0 0 250G 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
└─sda2 8:2 0 82G 0 part
   ├─system-root 254:0 0 60G 0 lvm /root
   │ /var
   │ /usr/local
   │ /tmp
   │ /srv
   │ /opt
   │ /home
   │ /boot/grub2/x86++_++64-efi
   │ /boot/grub2/i386-pc
   │ /.snapshots
   │ /
   └─system-swap 254:1 0 2G 0 lvm SWAP
sdb 8:16 0 200G 0 disk
sdc 8:32 0 120G 0 disk
sdd 8:48 0 150G 0 disk
sr0 11:0 1 1024M 0 rom
hana-1:/install #
```

An xfs file system has been created on each of the three new disks.

```
hana-1:/install # mkfs.xfs /dev/sdb
meta-data=/dev/sdb isize=512 agcount=4, agsize=7864320 blks
sectsz=512 attr=2, projid32bit=1
crc=1 finobt=1, sparse=1, rmapbt=0
reflink=0 bigtime=0 inobtcount=0
data = bsize=4096 blocks=31457280, imaxpct=25
sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=15360, version=2
sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
```

```
hana-1:/install # mkfs.xfs /dev/sdc
meta-data=/dev/sdc isize=512 agcount=4, agsize=7864320 blks
sectsz=512 attr=2, projid32bit=1
crc=1 finobt=1, sparse=1, rmapbt=0
reflink=0 bigtime=0 inobtcount=0
data = bsize=4096 blocks=31457280, imaxpct=25
sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=15360, version=2
sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
```

```
hana-1:/install # mkfs.xfs /dev/sdd
meta-data=/dev/sdd isize=512 agcount=4, agsize=9830400 blks
sectsz=512 attr=2, projid32bit=1
crc=1 finobt=1, sparse=1, rmapbt=0
reflink=0 bigtime=0 inobtcount=0
data = bsize=4096 blocks=39321600, imaxpct=25
sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=19200, version=2
sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
hana-1:/install #
```

## Creation of mount points

```
hana-1:/ # mkdir -p /hana/data/SMA/mnt00001
hana-1:/ # mkdir -p /hana/log/SMA/mnt00001
hana-1:/ # mkdir -p /hana/shared
hana-1:/ # chmod -R 777 /hana/log/SMA
hana-1:/ # chmod -R 777 /hana/data/SMA
hana-1:/ # chmod -R 777 /hana/shared
```

### Configuration of /etc/fstab

```

hana-1:/install # cat /etc/fstab
/dev/system/root / btrfs defaults 0 0
/dev/system/root /var btrfs subvol=@/var 0 0
/dev/system/root /usr/local btrfs subvol=@/usr/local 0 0
/dev/system/root /tmp btrfs subvol=@/tmp 0 0
/dev/system/root /srv btrfs subvol=@/srv 0 0
/dev/system/root /root btrfs subvol=@/root 0 0
/dev/system/root /opt btrfs subvol=@/opt 0 0
/dev/system/root /home btrfs subvol=@/home 0 0
/dev/system/root /boot/grub2/x86_64-efi btrfs subvol=@/boot/grub2/x86_64-efi 0 0
/dev/system/root /boot/grub2/i386-pc btrfs subvol=@/boot/grub2/i386-pc 0 0
/dev/system/swap swap swap defaults 0 0
/dev/system/root /.snapshots btrfs subvol=@/.snapshots 0 0
UUID=2E8C-48E1 /boot/efi vfat utf8 0 2
/dev/sdb /hana/data/SMA/mnt00001 xfs relatime,inode64 0 0
/dev/sdc /hana/log/SMA/mnt00001 xfs relatime,inode64 0 0
/dev/sdd /hana/shared xfs defaults 0 0
hana-1:/install #

```

```

hana-1:/install # df -h
Filesystem Size Used Avail Use% Mounted on
devtmpfs 4.0M 8.0K 4.0M 1% /dev
tmpfs 49G 4.0K 49G 1% /dev/shm
tmpfs 13G 26M 13G 1% /run
tmpfs 4.0M 0 4.0M 0% /sys/fs/cgroup
/dev/mapper/system-root 60G 35G 25G 58% /
/dev/mapper/system-root 60G 35G 25G 58% /.snapshots
/dev/mapper/system-root 60G 35G 25G 58% /boot/grub2/i386-pc
/dev/mapper/system-root 60G 35G 25G 58% /boot/grub2/x86_64-efi
/dev/mapper/system-root 60G 35G 25G 58% /home
/dev/mapper/system-root 60G 35G 25G 58% /opt
/dev/mapper/system-root 60G 35G 25G 58% /srv
/dev/mapper/system-root 60G 35G 25G 58% /tmp
/dev/mapper/system-root 60G 35G 25G 58% /usr/local
/dev/mapper/system-root 60G 35G 25G 58% /var
/dev/mapper/system-root 60G 35G 25G 58% /root
/dev/sda1 253M 5.1M 247M 3% /boot/efi
tmpfs 6.3G 56K 6.3G 1% /run/user/0
/dev/sdb 200G 237M 200G 1% /hana/data/SMA/mnt00001
/dev/sdc 120G 155M 120G 1% /hana/log/SMA/mnt00001
/dev/sdd 150G 186M 150G 1% /hana/shared
hana-1:/install #

```

## HANA installation

The HANA installation can now be executed.



With the described configuration the /usr/sap/SMA directory will be on the OS VMDK. If /usr/sap/SMA should be stored in the shared VMDK, the hana shared disk could be partitioned to provide another file system for /usr/sap/SMA.

## Userstore key for SnapCenter

A user store for a system database user must be created, which should be used by SnapCenter. The HANA instance number must be set accordingly for communication port. In our setup instance number “00” is used.

A more detailed description can be found at [SnapCenter resource-specific configuration for SAP HANA database backups](#)

```
smaadm@hana-1:/usr/sap/SMA/HDB00> hdbuserstore set SMAKEY hana-1:30013
SNAPCENTER <password>
Operation succeed.
```

The connectivity can be checked with the command below.

```
smaadm@hana-1:/usr/sap/SMA/HDB00> hdbsql -U SMAKEY
Welcome to the SAP HANA Database interactive terminal.
Type: \h for help with commands
\q to quit
hdbsql SYSTEMDB=> exit
smaadm@hana-1:/usr/sap/SMA/HDB00
```

## SnapMirror active sync configuration

This article covers the configuration steps required for this solution.

### Pre-requisites

Storage clusters and relevant SVMs must be peered.

ONTAP mediator must be available and configured at both storage clusters.

**Overview**

Intercluster settings

Network interfaces

IP ADDRESS

- 192.168.100.102
- 192.168.100.101

Cluster peers

PEERED CLUSTER NAME

- tme-a800

Mediator

10.01.100.07

- tme-a800

Storage VM peers

PEERED STORAGE VMS

- 1

Protected data

Volume protection

Snapshot copies (local)

21 of the 62 volumes aren't protected.

SnapMirror (local or remote)

49 of the 62 volumes aren't protected.

Back up to cloud

66 of the 66 volumes aren't backed up to cloud.

Protect volumes

Back up volumes to cloud

Protect for business continuity

Let's you select specific volumes for protection if you don't need to protect entire storage VMs.

Let's you select which volumes you want to be backed up to a cloud destination.

Let's you protect a consistency group with a zero recovery time objective.

NetApp SnapCenter software simplifies backup, restore, and clone management for the applications hosted across ONTAP enabled platforms. [Use NetApp SnapCenter for application-consistent protection.](#)

Local policy settings

Cloud object stores

Cloud Backup Service Status: Not configured

**Storage VM peers** Protection overview

+ Peer storage VMs

Search Show/hide Filter

Storage VM	Peered cluster	Peered storage VM	Status	Applications using this peer
svm200_bluexpdr_a700s	tme-a800	svm200_bluexpdr_a800	Peered	SnapMirror

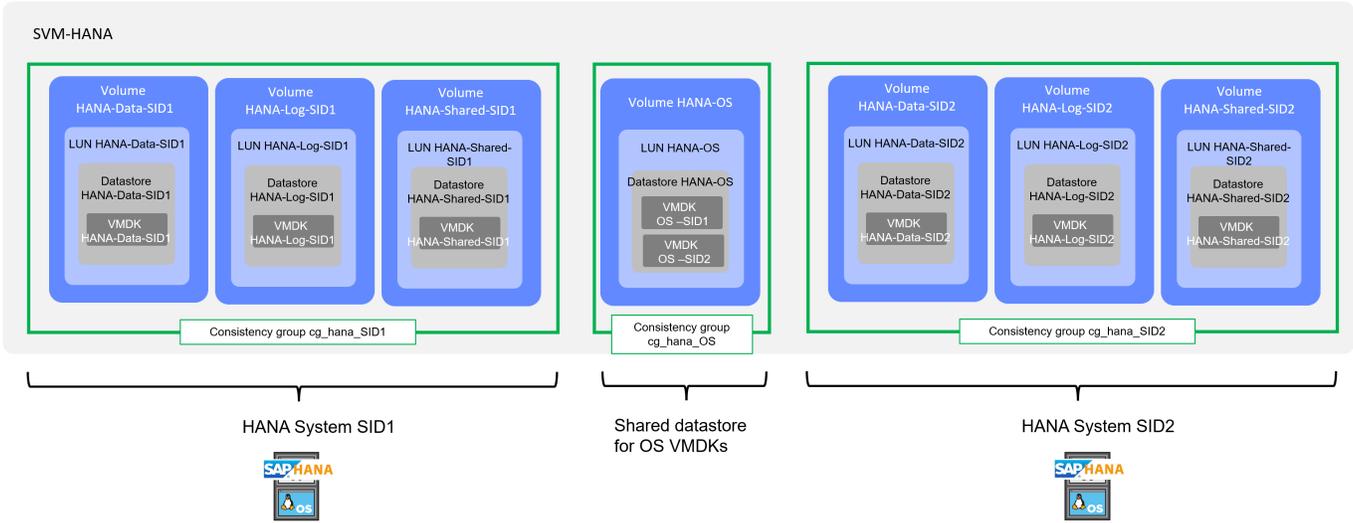
## Storage layout and consistency group configuration

In the ONTAP documentation [SnapMirror active sync overview in ONTAP](#) the concept of consistency groups with SnapMirror active sync is described as followed:

A consistency group is a collection of FlexVol volumes that provide a consistency guarantee for the application workload that must be protected for business continuity.

The purpose of a consistency group is to take simultaneous snapshot images of multiple volumes, thus ensuring crash-consistent copies of a collection of volumes at a point in time. A consistency group ensures all volumes of a dataset are quiesced and then snapped at precisely the same point in time. This provides a data-consistent restore point across volumes supporting the dataset. A consistency group thereby maintains dependent write-order consistency. If you decide to protect applications for business continuity, the group of volumes corresponding to this application must be added to a consistency group so a data protection relationship is established between a source and a destination consistency group. The source and destination consistency must contain the same number and type of volumes.

For the replication of HANA systems, the consistency group must include all volumes used by the individual HANA system (data, log and shared). Volumes which should be part of a consistency group must be stored in the same SVM. Operating system images can be stored in a separate volume with its own consistency group. The figure below illustrates a configuration example with two HANA systems.



## Initiator group configuration

In our lab setup we created an initiator group including both storage SVMs which are used for the SnapMirror active sync replication. In the SnapMirror active sync configuration described later, we will define that the initiator group will be part of the replication.

Using the proximity settings, we defined which ESX host is close to which storage cluster. In our case the A700 is close to ESX-1 and the A800 is close to ESX-2.

The screenshot shows the ONTAP System Manager interface for SAN initiator groups. The selected group is **cluster\_87\_1**. The configuration details are as follows:

- Name:** cluster\_87\_1
- Storage VM:** svm200\_bluexpdr\_a700s
- TYPE:** VMware
- PROTOCOL:** Mixed (iSCSI & FC)
- CONNECTION STATUS:** OK
- Replication:**
  - REPLICATED TO SVM: svm200\_bluexpdr\_a800
  - REPLICATION TO CLUSTER: tme-a800
  - REPLICATION STATUS: OK
- Initiators:**

Name	Description	Connection status	In proximity to
10.00.00.10.9b:17.04:69	← ESX-1	OK	svm200_bluexpdr_a700s
10.00.00.10.9b:17.04:6a		OK	svm200_bluexpdr_a700s
10.00.00.10.9b:40.b5:7f	← ESX-2	OK	svm200_bluexpdr_a800
10.00.00.10.9b:40.b9:80		OK	svm200_bluexpdr_a800



In a non-uniform access setup, the initiator group at the primary storage cluster (A700) must only include the initiators of the ESX-1 host, since there is no SAN connection to ESX-2. In addition, you need to configure another initiator group at the second storage cluster (A800) which only include the initiators of the ESX-2 host. Proximity configuration and initiator group replication is not required.

## Configure protection with ONTAP system manager

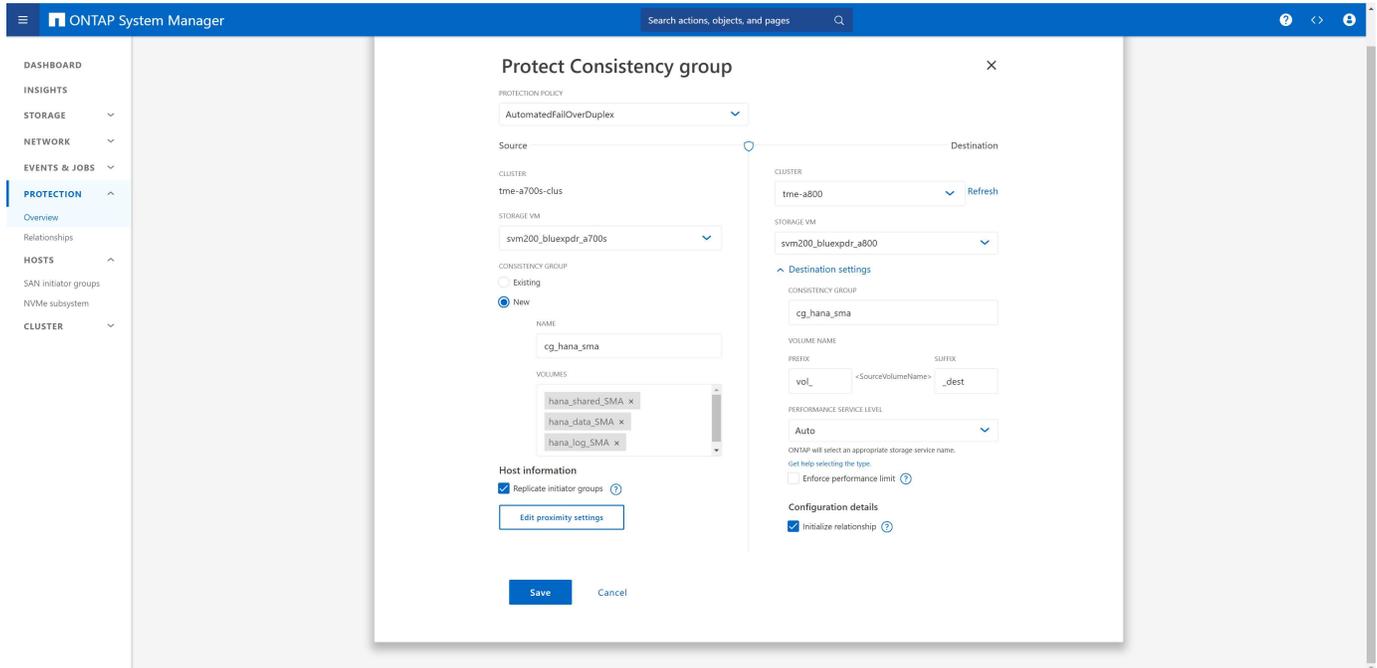
## Consistency group and initiator group replication

A new consistency group must be created, and all three LUNs of the HANA system must be added to the consistency group.

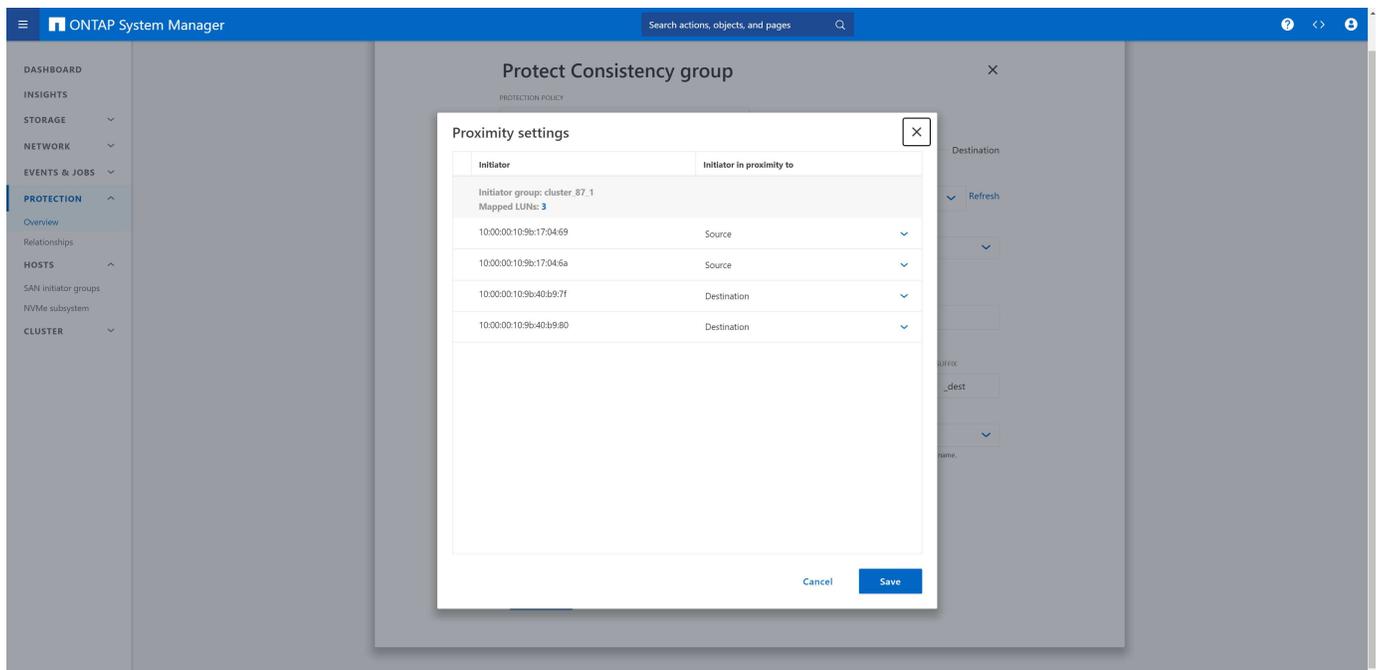
“Replicate initiator group” has been enabled. The initiator group will then stay in-sync independent where changes are made.



In a non-uniform access setup, the initiator group must not be replicated, since a separate initiator group must be configured at the second storage cluster.



By clicking on proximity settings, you can review the configuration done before in the initiator group setup.



The destination storage cluster must be configured and “initialize relationship” must be enabled.

## Synchronisation

At the A700 storage cluster (source), the new relationship is now listed.

The screenshot shows the ONTAP System Manager interface for the A700 storage cluster. The 'Relationships' page is active, showing a table of relationships between local sources and destinations. A new relationship, 'svm200\_blueexpdr\_a700c/cg/hana\_sma', is highlighted with a blue box. This relationship is synchronous and connects the source 'svm200\_blueexpdr\_a700c/cg/hana\_sma' to the destination 'svm200\_blueexpdr\_a800/cg/hana\_sma'.

Source	Destination	Policy type
svm200_blueexpdr_a700csudev1	svm200_blueexpdr_a800csudev1_dest	Asynchronous
svm200_blueexpdr_a700svvol_nvme_1	svm200_blueexpdr_a800vol_vvol_nvme_1_dest	Asynchronous
svm200_blueexpdr_a700sDraaS_qa_lun1	svm200_blueexpdr_a800DraaS_qa_lun1_dest	Asynchronous
svm200_blueexpdr_a700svvol_nvme_2	svm200_blueexpdr_a800vol_vvol_nvme_2_dest	Asynchronous
svm200_blueexpdr_a700sNVME_FC_SCV_NEW	svm200_blueexpdr_a800NVME_FC_SCV_NEW_dest_1	Asynchronous
svm200_blueexpdr_a700sVOLUME_NVME_FC_PLUGIN_PRI	svm200_blueexpdr_a800vol_VOLUME_NVME_FC_PLUGIN_PRI_dest	Asynchronous
svm200_blueexpdr_a700svvol_nvme_1	svm200_blueexpdr_a800vol_nvme_1_dest	Asynchronous
svm200_blueexpdr_a700c/cg/hana_sma	svm200_blueexpdr_a800/cg/hana_sma	Synchronous
svm200_blueexpdr_a700c/cg/hana_test	svm200_blueexpdr_a800c/cg/hana_test	Synchronous

At the A800 storage cluster (destination), the new relationship and the status of the replication is listed.

The screenshot shows the ONTAP System Manager interface for the A800 storage cluster. The 'Relationships' page is active, showing a table of relationships between local sources and destinations. A new relationship, 'svm200\_blueexpdr\_a700c/cg/hana\_sma', is highlighted with a blue box. This relationship is asynchronous and connects the source 'svm200\_blueexpdr\_a700c/cg/hana\_sma' to the destination 'svm200\_blueexpdr\_a800/cg/hana\_sma'. The relationship health is 'Healthy' and the state is 'Synchronizing'.

Source	Destination	Protection policy	Relationship health	State	Lag
svm200_blueexpdr_a700csudev1	svm200_blueexpdr_a800csudev1_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_blueexpdr_a700svvol_nvme_1	svm200_blueexpdr_a800vol_vvol_nvme_1_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_blueexpdr_a700sDraaS_qa_lun1	svm200_blueexpdr_a800DraaS_qa_lun1_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_blueexpdr_a700svvol_nvme_2	svm200_blueexpdr_a800vol_vvol_nvme_2_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_blueexpdr_a700sNVME_FC_SCV_NEW	svm200_blueexpdr_a800NVME_FC_SCV_NEW_dest_1	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_blueexpdr_a700sVOLUME_NVME_FC_PLUGIN_PRI	svm200_blueexpdr_a800vol_VOLUME_NVME_FC_PLUGIN_PRI_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_blueexpdr_a700svvol_nvme_1	svm200_blueexpdr_a800vol_nvme_1_dest	MirrorAndVault	Healthy	Mirrored	141 days, 20 hours, 5 minutes and 4 seconds
svm200_blueexpdr_a700svvol_nvme_2	svm200_blueexpdr_a800vol_nvme_2_dest	MirrorAndVault	Healthy	Mirrored	141 days, 22 hours, 12 minutes and 4 seconds
svm200_blueexpdr_a700c/cg/hana_sma	svm200_blueexpdr_a800/cg/hana_sma	AutomatedFailOverDuplex	Healthy	Synchronizing	-
svm200_blueexpdr_a700c/cg/hana_test	svm200_blueexpdr_a800c/cg/hana_test	AutomatedFailOverDuplex	Healthy	In sync	-

## Infrastructure datastore

The datastore, where the OS images of the HANA system, SnapCenter and the vSphere plugin is stored is replicated in the same way as described for the HANA database datastores.

## Primary site

SnapMirror active sync behaviour is symmetric, with one important exception - primary site configuration.

SnapMirror active sync will consider one site the "source" and the other the "destination". This implies a one-way replication relationship, but this does not apply to IO behaviour. Replication is bidirectional and symmetric and IO response times are the same on either side of the mirror.

If the replication link is lost, the LUN paths on the source copy will continue to serve data while the LUN paths on the destination copy will become unavailable until replication is reestablished and SnapMirror re-enters a synchronous state. The paths will then resume serving data.

The effect of designating one cluster as a source simply controls which cluster survives as a read-write storage system if the replication link is lost.

The primary site is detected by SnapCenter and used to execute backup, restore and cloning operations.



Keep in mind, that source and destination is not tied to the SVM or storage cluster but can be different for each replication relationship.

The screenshot shows the ONTAP System Manager interface. The main content area displays a table of replication relationships under the 'Local sources' tab. Below the table, there is a detailed view of a replication policy named 'AutomatedFailOverDuplex'.

Source	Destination	Policy type
svm200_bluexpdr_a700s:NVME_FC_SCV_NEW	svm200_bluexpdr_a800:NVME_FC_SCV_NEW_dest_1	Asynchronous
svm200_bluexpdr_a700s:VOLUME_NVME_FC_PLUGIN_PRI	svm200_bluexpdr_a800:vol_VOLUME_NVME_FC_PLUGIN_PRI_dest	Asynchronous
svm200_bluexpdr_a700s:nvme_vol_1	svm200_bluexpdr_a800:vol_nvme_1_dest	Asynchronous
svm200_bluexpdr_a700s:nvme_vol_2	svm200_bluexpdr_a800:vol_nvme_2_dest	Asynchronous
svm200_bluexpdr_a700s:cg/cg_hana_sma	svm200_bluexpdr_a800:cg/cg_hana_sma	Synchronous

**PROTECTION POLICY**  
AutomatedFailOverDuplex

STATE:  In sync

TRANSFER STATUS: Success

IS HEALTHY?:  Yes

CONTAINED LUNS (SOURCE): /vol/hana\_data\_SMA/hana\_data\_SMA, /vol/hana\_log\_SMA/hana\_log\_SMA and 1 more

FAIL OVER MODE: Planned (Completed)

The diagram below shows two consistency groups: 'tme-a700s-clus' (consistency group cg\_hana\_sma) and 'tme-a800' (consistency group cg\_hana\_sma). They are connected via a mediator with IP address 10.61.180.97.

## SnapCenter configuration

As stated at the beginning of the document, the purpose of the document is to provide best practices for a HANA environment using VMware with VMFS and SnapMirror active sync. We will only cover details and important steps relevant for this specific solution and will not explain the general SnapCenter concepts. These concepts and other additional information on SnapCenter can be found at:

[TR-4614: SAP HANA backup and recovery with SnapCenter](#)

[TR-4719: SAP HANA System Replication - Backup and Recovery with SnapCenter](#)

[TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)

## Pre-requisites

In general, SnapMirror active sync should be setup before the protected resources are added to SnapCenter. If backups have been created before the setup of SnapMirror active sync, they will only exist at the original primary storage and will not be replicated afterwards.

### SnapCenter HANA resource must be auto discovered

Resources which are configured with VMware VMFS or resources protected with SnapMirror active sync must be auto discovered by SnapCenter to allow specific operations required for these configurations.

Since HANA non-data volumes are always manual configured resources in SnapCenter, they are not supported by SnapCenter out of the box. We will discuss options and workarounds for non-data volumes later in this document.

SAP HANA multiple host systems must be configured using a central HANA plugin and are therefore manual configured resources by default. Such HANA systems are not supported by SnapCenter, when using VMware VMFS or SnapMirror active sync.

### SnapCenter for VMware vSphere plugin

The SnapCenter for VMware vSphere plugin must be deployed in the VMware environment.

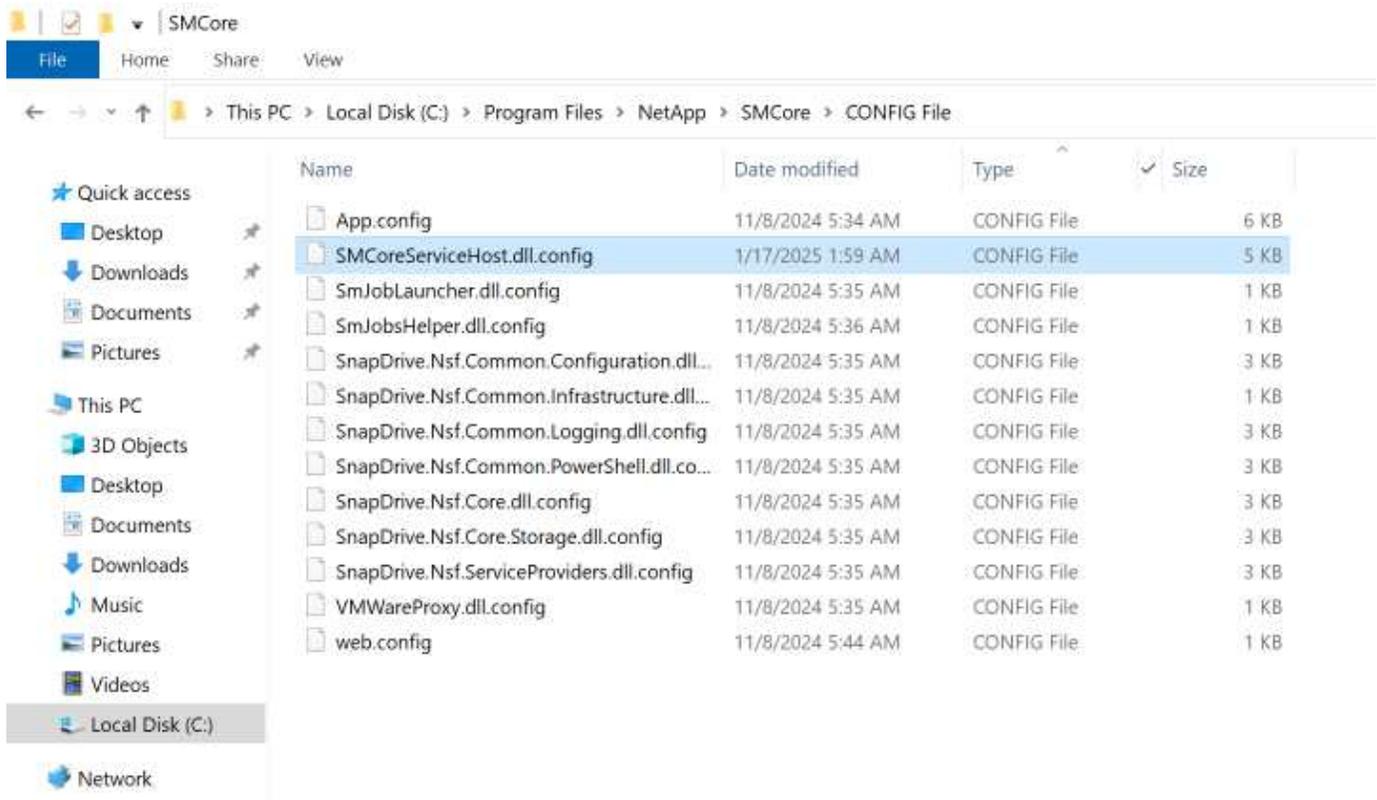
### Management IP address on SVM hosting the volumes

Even though clusters will be added to SnapCenter, the SVMs hosting the source and destination volumes must have a management IP address configured.

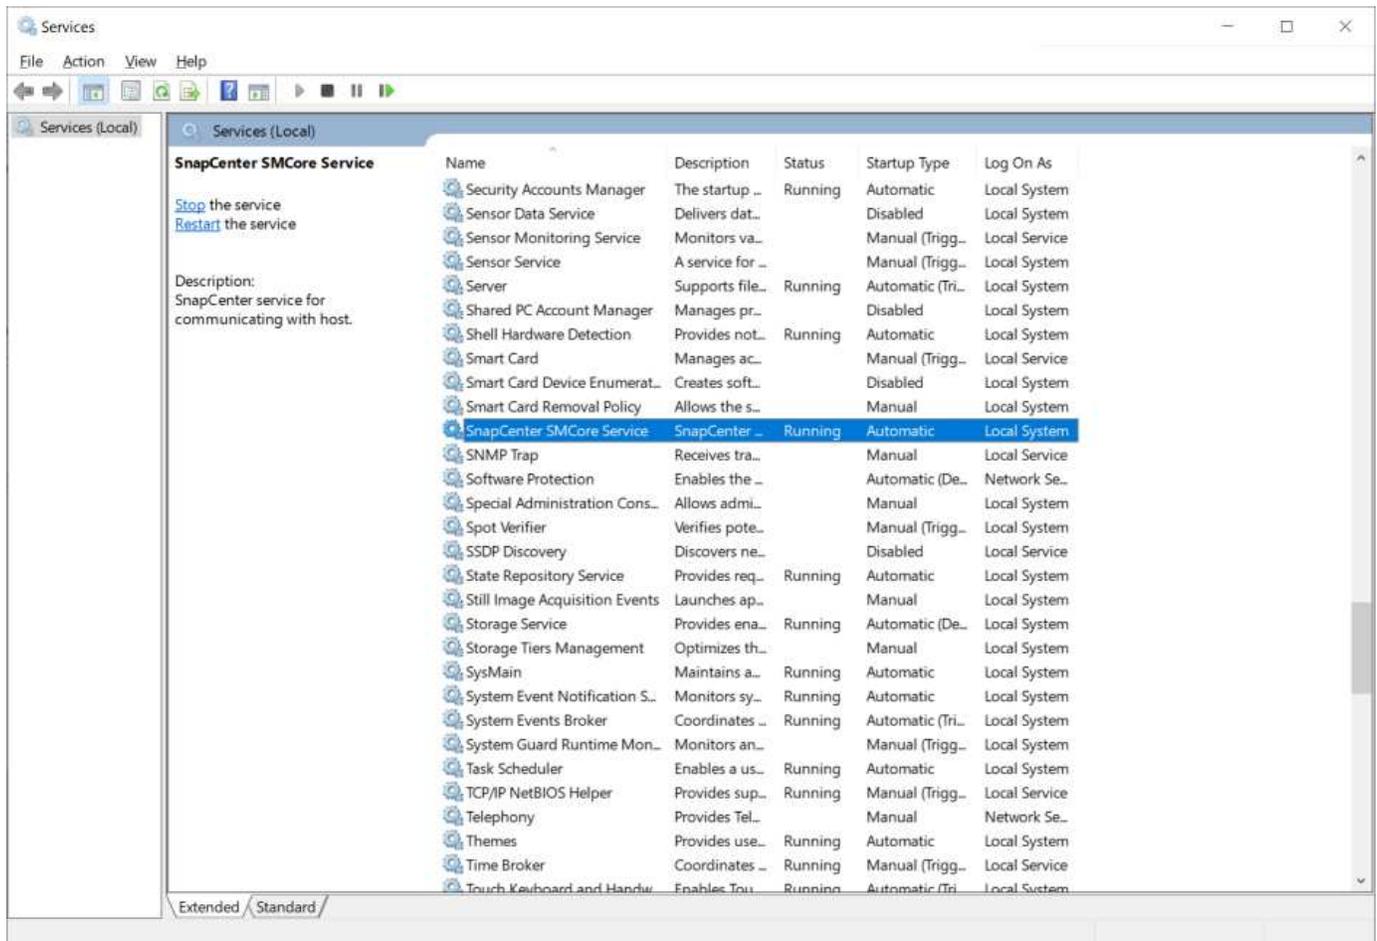
### REST APIs for storage communication

Management and monitoring of SnapMirror active sync requires REST API access. Therefore, SnapCenter must be configured to use REST APIs for storage communications. The parameter "IsRestEnabledForStorageConnection" in the configuration file C:\Program Files\NetApp\SMCore\SMCoreServiceHost.dll.config must be set to true.

```
<add key="IsRestEnabledForStorageConnection" value="true">
```

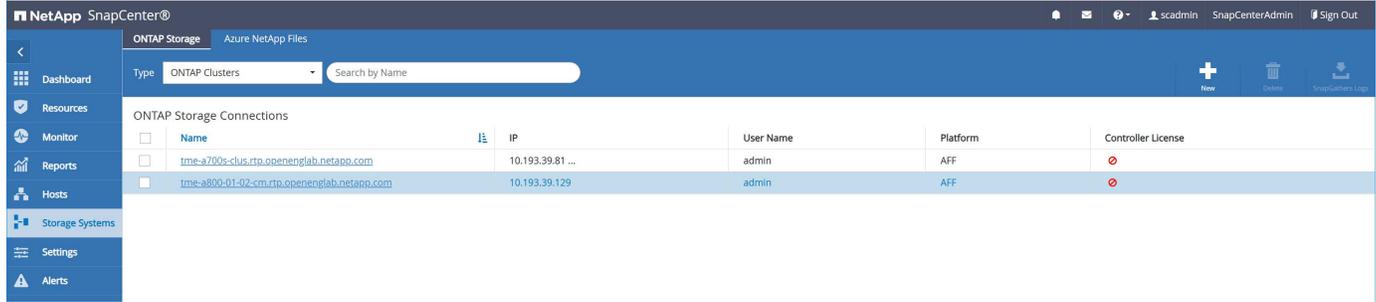


After the parameter change the SnapCenter SMCore Service must be restarted.



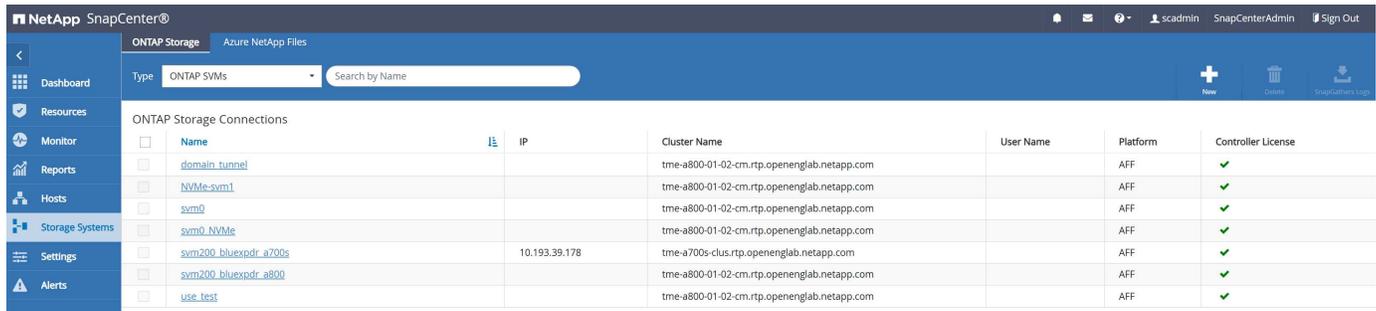
## Add storage systems

Storage systems can be added after REST API is enabled for SnapCenter. It is required to add both storage clusters, not the individual SVM's.



The screenshot shows the NetApp SnapCenter interface with the 'Storage Systems' menu selected. The 'Type' is set to 'ONTAP Clusters'. The table below lists the storage connections:

Name	IP	User Name	Platform	Controller License
tme-a700s-clus.rtp.openenglab.netapp.com	10.193.39.81 ...	admin	AFF	⊘
tme-a800-01-02-cm.rtp.openenglab.netapp.com	10.193.39.129	admin	AFF	⊘



The screenshot shows the NetApp SnapCenter interface with the 'Type' set to 'ONTAP SVMs'. The table below lists the storage connections:

Name	IP	Cluster Name	User Name	Platform	Controller License
domain_tunnel		tme-a800-01-02-cm.rtp.openenglab.netapp.com		AFF	✓
NVMe-svm1		tme-a800-01-02-cm.rtp.openenglab.netapp.com		AFF	✓
svm0		tme-a800-01-02-cm.rtp.openenglab.netapp.com		AFF	✓
svm0_NVMe		tme-a800-01-02-cm.rtp.openenglab.netapp.com		AFF	✓
svm200_bluexodr_a700s	10.193.39.178	tme-a700s-clus.rtp.openenglab.netapp.com		AFF	✓
svm200_bluexodr_a800		tme-a800-01-02-cm.rtp.openenglab.netapp.com		AFF	✓
use_test		tme-a800-01-02-cm.rtp.openenglab.netapp.com		AFF	✓

## Add host – SnapCenter for VMware vSphere plugin

If a resource in SnapCenter is running in a virtualized VMware environment, SnapCenter leverages the SnapCenter plugin for VMware vSphere to extend the SnapCenter backup, restore and cloning workflows with the required steps on the VMware layer.

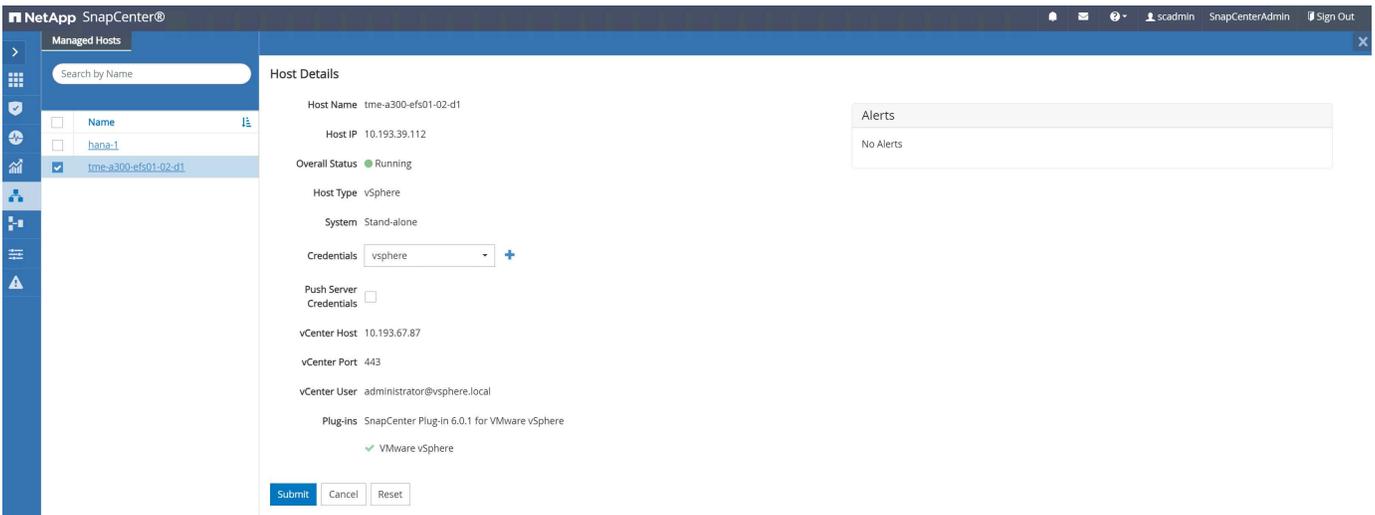
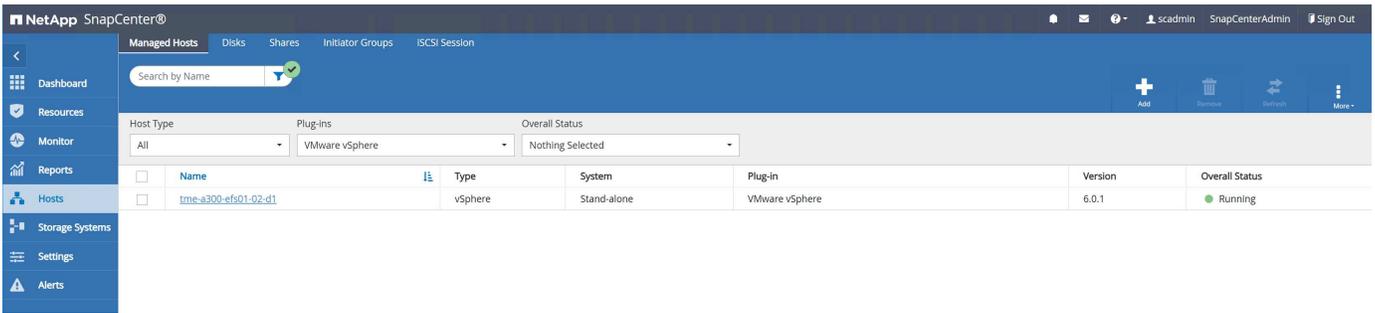
Before the host can be added in SnapCenter the SnapCenter plugin for VMware vSphere must be deployed within the VMware environment.



Credentials must be set during host add workflow, where vSphere can be selected as a host type.



The screenshot shows the 'Add Host' dialog box in the NetApp SnapCenter interface. The 'Host Type' is set to 'vSphere'. The 'Host Name' field is empty, and the 'Credentials' field is set to 'None'. There are 'Submit' and 'Cancel' buttons. A warning message at the bottom states: "Prechecks and remote installation of plug-ins cannot be performed using the credential that is set to 'None'. Plug-ins must be manually installed and plug-in services should be up and running."



No additional configuration required at the SnapCenter for vSphere plugin itself.

### Add host – HANA system

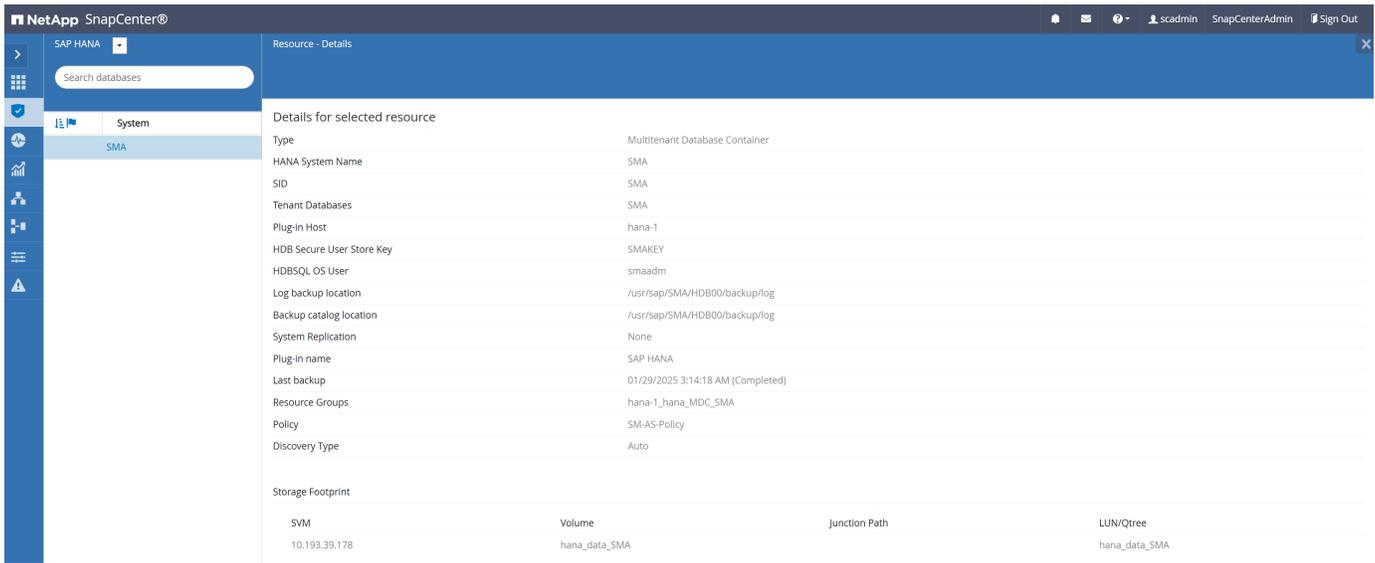


No specific requirements. Plugin deployment and auto discovery is done as usual.

With the auto discovery process SnapCenter detects that the HANA resource is running virtualized with VMFS/VMDKs. SnapCenter also detects the SnapMirror active sync setup and identifies the current primary site.

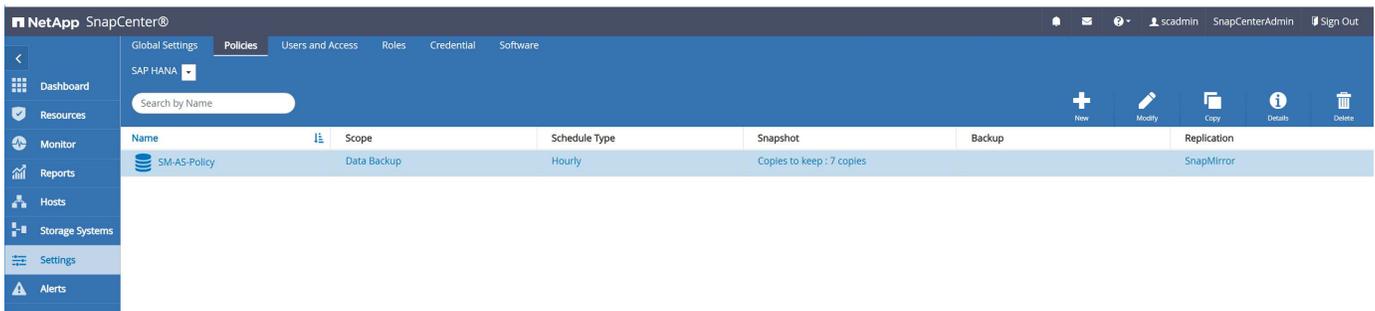
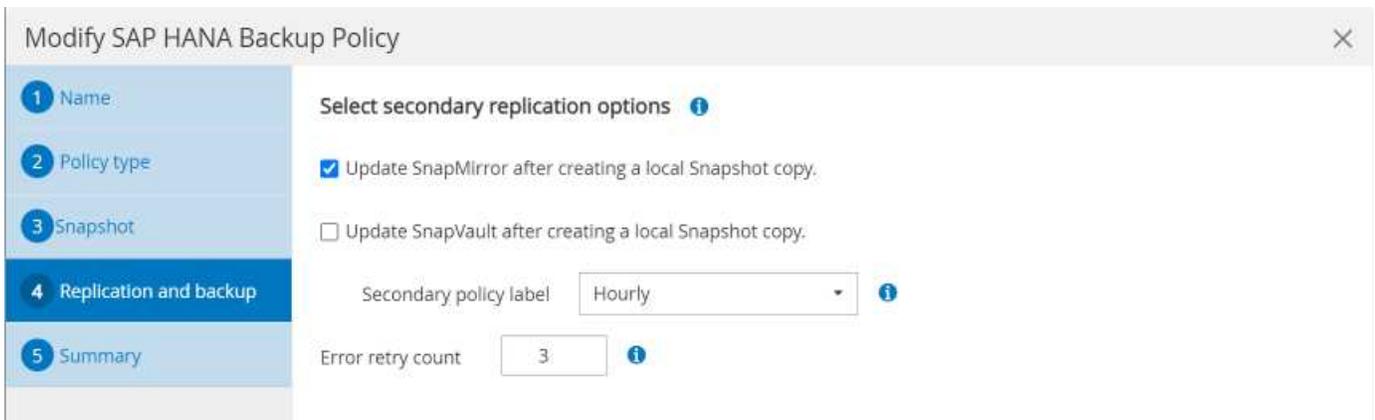
After resource auto discovery the current primary site is shown in the storage footprint section of the resource view. The detection which storage system is master is based on the output of the ONTAP command, which is used by SnapCenter.

```
volume show -vserver <vs> -volume <vol> -fields smbc-consensus,is-smbc-master
```



## Policy configuration

The policy used for the resource protected with SnapMirror active sync must be configured using SnapMirror replication even though SnapCenter does not trigger any SnapMirror update operations.



## HANA resource protection configuration

No specific requirements. Resource protection configuration is done as usual.

## SnapCenter backup operations

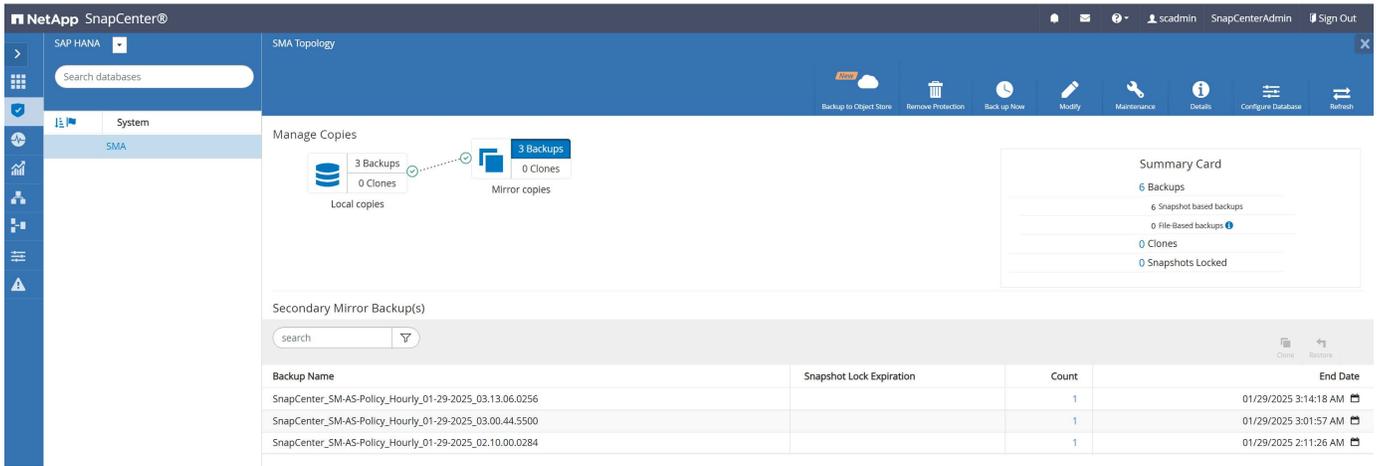
With each backup operation, SnapCenter executes the discovery on the VMware side as well as the detection of the primary site. If there is a storage failover, SnapCenter will

detect the new primary site as soon as a backup has been executed for the resource.

## Topology view

Within the topology view, SnapCenter shows the backups of both source and destination storage clusters.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_SM-AS-Policy_Hourly_01-29-2025_03.13.06.0256		1	01/29/2025 3:14:18 AM
SnapCenter_SM-AS-Policy_Hourly_01-29-2025_03.00.44.5500		1	01/29/2025 3:01:57 AM
SnapCenter_SM-AS-Policy_Hourly_01-29-2025_02.10.00.0284		1	01/29/2025 2:11:26 AM



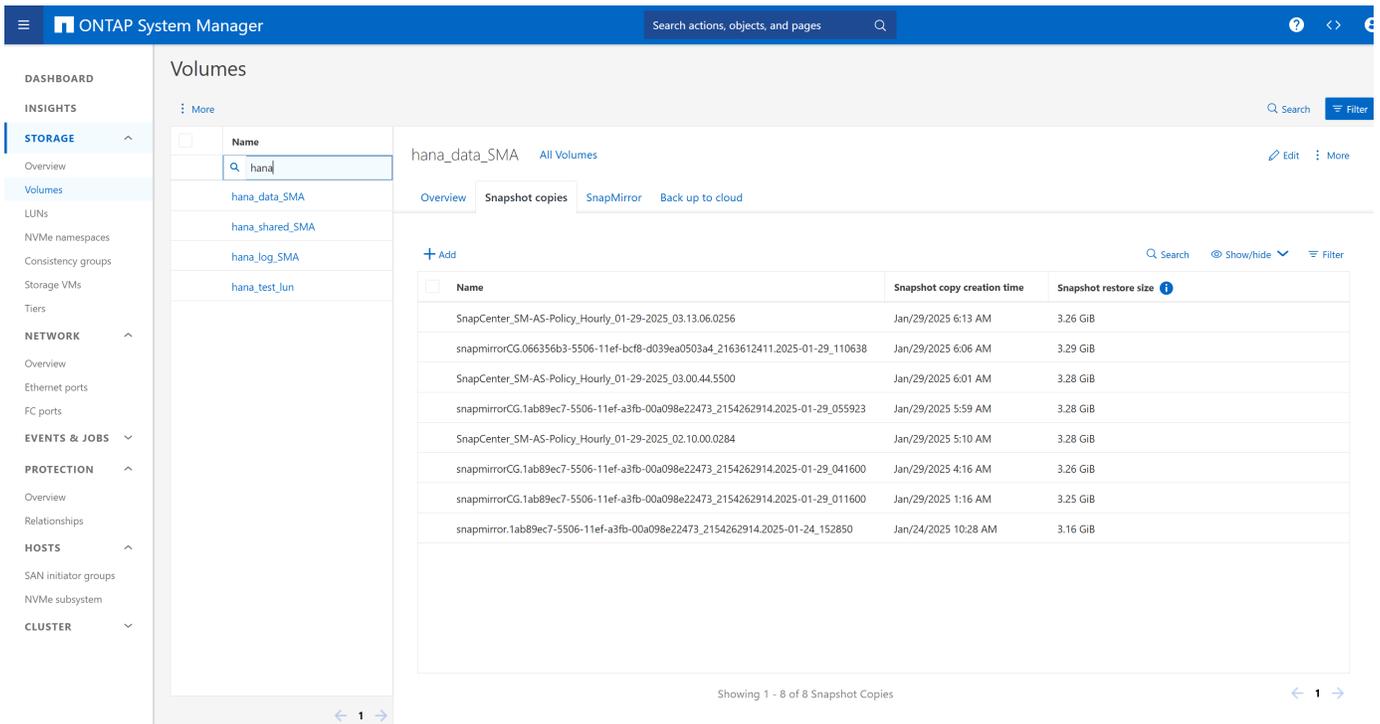
By clicking on the count number at the secondary storage, the current relationship and replication direction is shown. The source is always the current primary site. After a storage failover the primary site will change, and the display is adapted accordingly. All backups have always the same relationship dependent which storage system is currently the primary site.



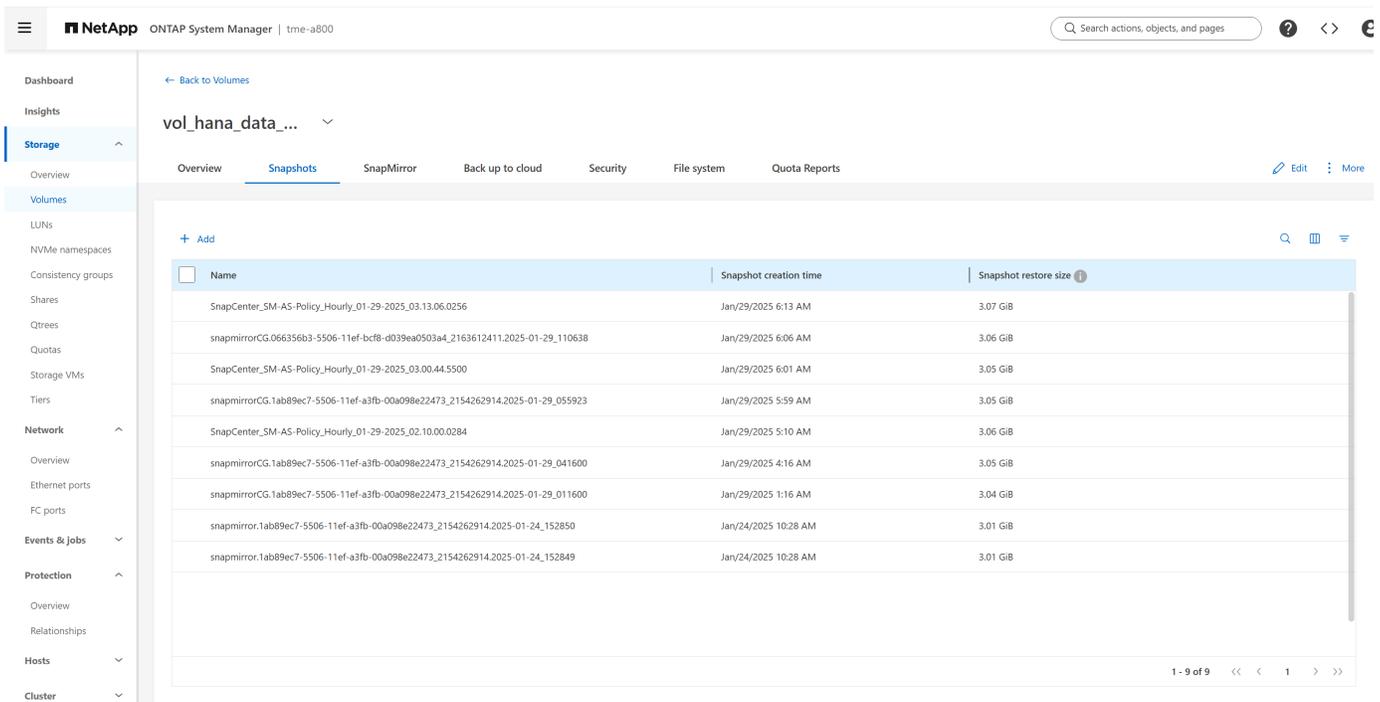
### Snapshots at storage systems

The Snapshot backups that have been created by SnapCenter are available at both HANA data volumes at both storage systems. ONTAP creates additional Snapshots on the consistency group level, which are available at all other HANA volumes as well.

The figure below shows the Snapshots of the HANA data volume at the A700 cluster.



The figure below shows the Snapshots of the HANA data volume at the A800 cluster.



## SnapCenter restore and recovery

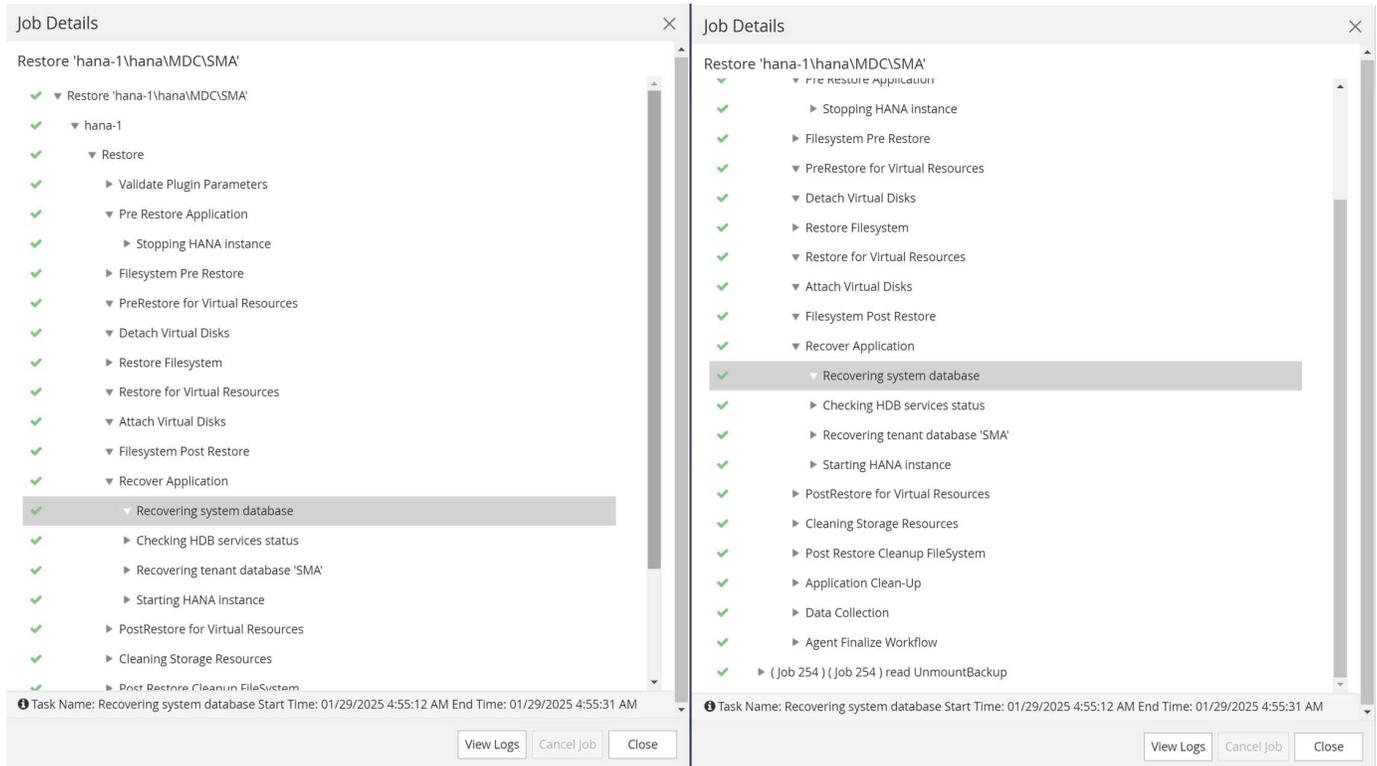
With virtual resources stored on VMFS/VMDK's a SnapCenter restore operation is always done by a clone, mount, copy operation.

1. SnapCenter creates a volume clone based on the selected Snapshot
2. SnapCenter mounts the LUN in the cloned volume as a new datastore to the ESX host

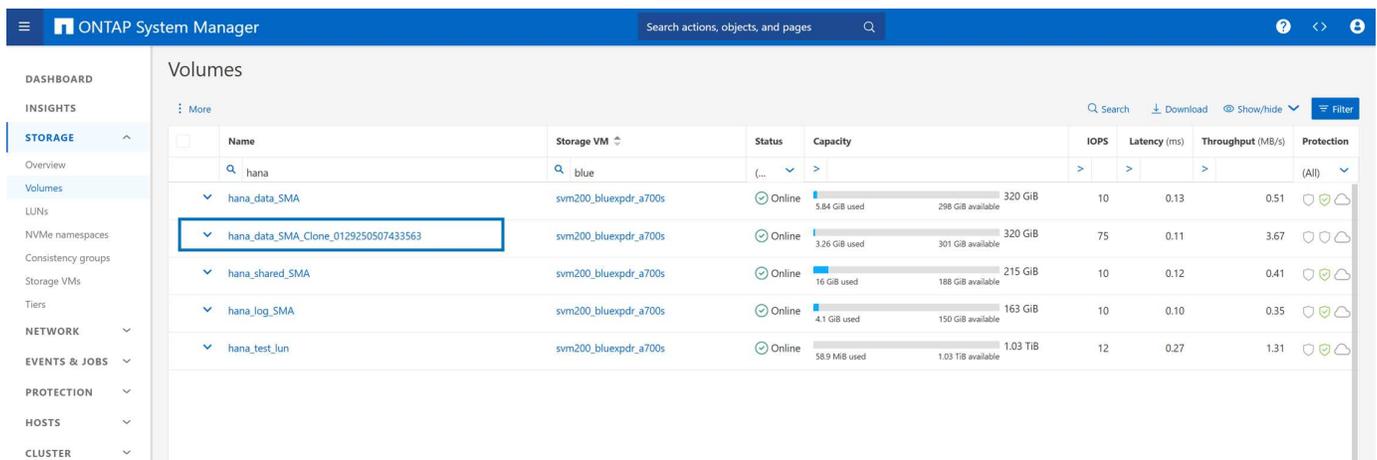
3. SnapCenter adds the VMDK within the datastore as a new disk to the HANA VM
4. SnapCenter mounts the new disk to the Linux OS
5. SnapCenter copies the data from the new disk back to the original location
6. When the copy operation is finished all above resource are removed again
7. The HANA recovery is done as usual

The overall runtime of the restore operation is therefore dependent on the database size and the throughput of the FC connection between the storage clusters and the ESX hosts.

In addition, when a resource is configured with SnapMirror active sync the SnapCenter restore operation can only be selected at the current primary site.



While the restore and recovery operation is running, you can see a new cloned volume, which has been created at the current primary site.



At the HANA Linux host, you can see a new disk, which got mounted to the host. When the restore operation is done the disk, datastore and volumes will be removed again by SnapCenter.

```
hana-1:~ # df -h
Filesystem Size Used Avail Use% Mounted on
devtmpfs 4.0M 8.0K 4.0M 1% /dev
tmpfs 49G 4.0K 49G 1% /dev/shm
tmpfs 13G 58M 13G 1% /run
tmpfs 4.0M 0 4.0M 0% /sys/fs/cgroup
/dev/mapper/system-root 60G 36G 24G 60% /
/dev/mapper/system-root 60G 36G 24G 60% /.snapshots
/dev/mapper/system-root 60G 36G 24G 60% /boot/grub2/i386-pc
/dev/mapper/system-root 60G 36G 24G 60% /home
/dev/mapper/system-root 60G 36G 24G 60% /boot/grub2/x86_64-efi
/dev/mapper/system-root 60G 36G 24G 60% /opt
/dev/mapper/system-root 60G 36G 24G 60% /srv
/dev/mapper/system-root 60G 36G 24G 60% /usr/local
/dev/mapper/system-root 60G 36G 24G 60% /tmp
/dev/mapper/system-root 60G 36G 24G 60% /root
/dev/mapper/system-root 60G 36G 24G 60% /var
/dev/sdb 200G 8.0G 192G 4% /hana/data/SMA/mnt00001
/dev/sdc 120G 7.0G 113G 6% /hana/log/SMA/mnt00001
/dev/sda1 253M 5.1M 247M 3% /boot/efi
/dev/sdd 150G 28G 123G 19% /hana/shared
tmpfs 6.3G 48K 6.3G 1% /run/user/467
tmpfs 6.3G 28K 6.3G 1% /run/user/0
/dev/sde 200G 8.0G 192G 4%
/var/opt/snapcenter/scu/clones/hana_data_SMAmnt00001_255_scu_clone_1
hana-1:~ #
```

## SAP System refresh operation

Cloning operations can be executed at the primary site or the secondary storage.

The cloned volume will not be part of the HANA consistency group and will not be replicated with SnapMirror active sync.

Detailed information on the system refresh workflows can be found at: [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)

## SnapCenter non-data volumes

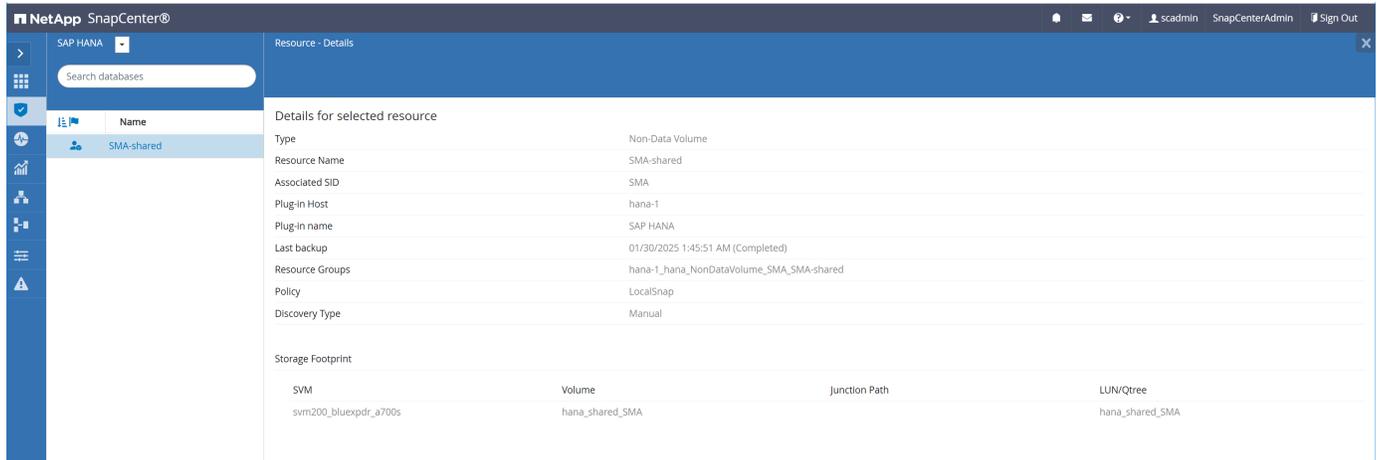
When resources are configured manually in SnapCenter and are not auto discovered, SnapCenter is not aware of VMware and SnapMirror active sync. Therefore, they are not supported natively by SnapCenter.

For non-data volumes like HANA shared, backup and restore operations could still be done using SnapCenter

when considering additional manual steps.

## Failure of the storage system configured in SnapCenter

If a failure of the storage system configured in SnapCenter occurs, SnapCenter will not automatically switch to the other storage system. The non-data volume resource must be adapted manually so that the mirrored copy of the volume is used for backup and restore operations.



The screenshot shows the NetApp SnapCenter interface. The left sidebar contains navigation icons. The main area displays 'Resource - Details' for a selected resource named 'SMA-shared'. The details include:

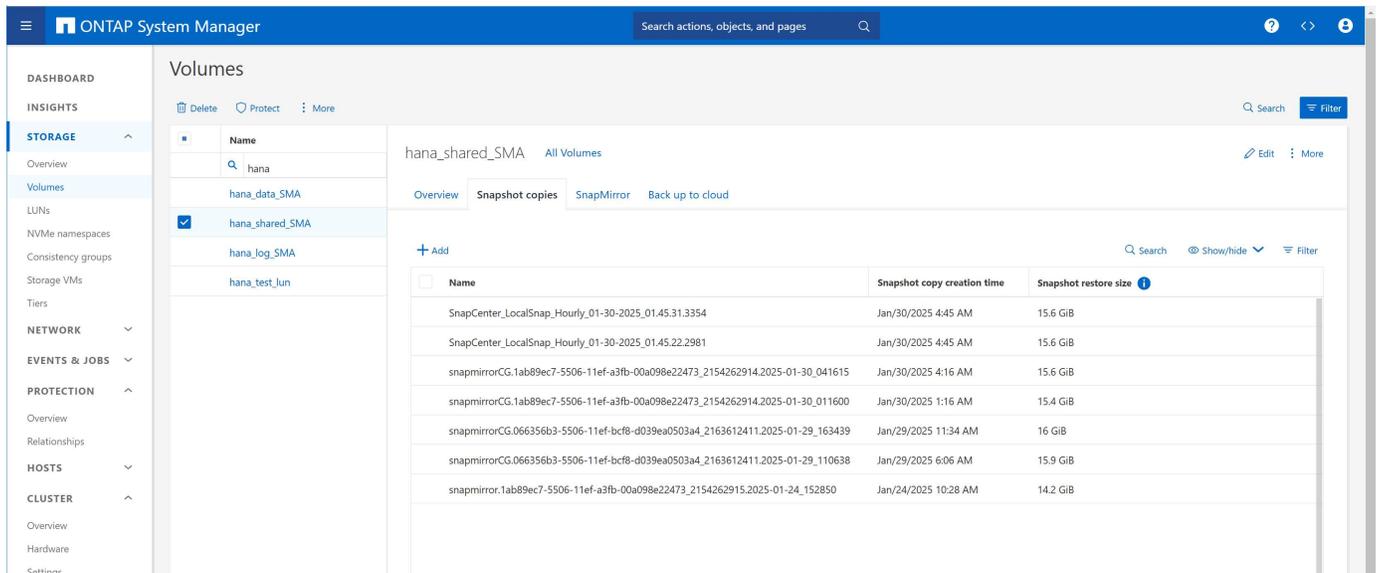
- Type: Non-Data Volume
- Resource Name: SMA-shared
- Associated SID: SMA
- Plug-in Host: hana-1
- Plug-in name: SAP HANA
- Last backup: 01/30/2025 1:45:51 AM (Completed)
- Resource Groups: hana-1\_hana\_NonDataVolume\_SMA\_SMA-shared
- Policy: LocalSnap
- Discovery Type: Manual

Below the details is a 'Storage Footprint' table:

SVM	Volume	Junction Path	LUN/Qtree
svm200_bluexpdr_a700s	hana_shared_SMA		hana_shared_SMA

## Backup operations

Even though SnapCenter is not aware of the SnapMirror active sync configuration for the HANA shared volume, Snapshots are replicated to both sites.



The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation icons. The main area displays 'Volumes' for the volume 'hana\_shared\_SMA'. The 'Snapshot copies' tab is selected, showing a list of snapshot copies:

Name	Snapshot copy creation time	Snapshot restore size
SnapCenter_LocalSnap_Hourly_01-30-2025_01.45.31.3354	Jan/30/2025 4:45 AM	15.6 GiB
SnapCenter_LocalSnap_Hourly_01-30-2025_01.45.22.2981	Jan/30/2025 4:45 AM	15.6 GiB
snapmirrorCG.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262914.2025-01-30_041615	Jan/30/2025 4:16 AM	15.6 GiB
snapmirrorCG.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262914.2025-01-30_011600	Jan/30/2025 1:16 AM	15.4 GiB
snapmirrorCG.066356b3-5506-11ef-bcf8-d039ea0503a4_2163612411.2025-01-29_163439	Jan/29/2025 11:34 AM	16 GiB
snapmirrorCG.066356b3-5506-11ef-bcf8-d039ea0503a4_2163612411.2025-01-29_110638	Jan/29/2025 6:06 AM	15.9 GiB
snapmirror.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262915.2025-01-24_152850	Jan/24/2025 10:28 AM	14.2 GiB

NetApp ONTAP System Manager | tme-a800

Search actions, objects, and pages

Dashboard

Insights

Storage

Overview

Volumes

LUNs

NVMe namespaces

Consistency groups

Shares

Qtrees

Quotas

Storage VMs

Tiers

Network

Events & Jobs

Protection

Hosts

Cluster

← Back to Volumes

vol\_hana\_share...

Overview Snapshots SnapMirror Back up to cloud Security File system Quota Reports

Edit More

+ Add

Name	Snapshot creation time	Snapshot restore size
SnapCenter_LocalSnap_Hourly_01-30-2025_01.45.31.3354	Jan/30/2025 4:45 AM	16.2 GiB
SnapCenter_LocalSnap_Hourly_01-30-2025_01.45.22.2981	Jan/30/2025 4:45 AM	16.2 GiB
snapmirrorCG.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262914.2025-01-30_041615	Jan/30/2025 4:16 AM	16.1 GiB
snapmirrorCG.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262914.2025-01-30_011600	Jan/30/2025 1:16 AM	16 GiB
snapmirrorCG.066356b3-5506-11ef-bcf8-d039ea0503a4_2163612411.2025-01-29_163439	Jan/29/2025 11:34 AM	15.7 GiB
snapmirrorCG.066356b3-5506-11ef-bcf8-d039ea0503a4_2163612411.2025-01-29_110638	Jan/29/2025 6:06 AM	15.8 GiB
snapmirror.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262915.2025-01-24_152850	Jan/24/2025 10:28 AM	14.1 GiB
snapmirror.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262915.2025-01-24_152849	Jan/24/2025 10:28 AM	14.1 GiB

## Restore operation

In case of a restore, SnapCenter would just execute a volume restore w/o any VMware specific steps. Normally you would need to unmount the HANA shared volume at the Linux host, disconnect the datastore then do the volume restore, connect the datastore again and then mount the file system at the Linux host. As a manual operation you could stop the HANA VM, restore the HANA shared volume with SnapCenter and then restart the VM again.

## Failover scenarios

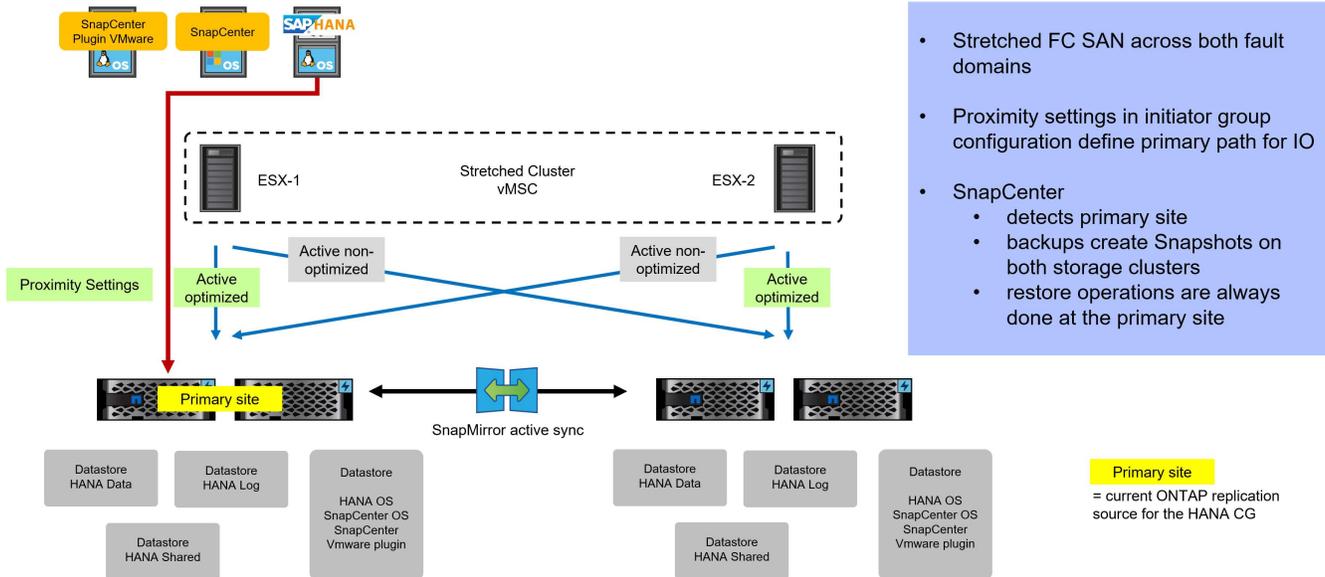
This article will highlight the failover scenarios for this solution.

## Uniform access setup

In a uniform access configuration, the fibre channel SAN is stretched across both sites. The ESX hosts at both sites could access both copies of the data sets. During normal operation, the ESX host running the HANA system is accessing the local copy of the data based on proximity settings in the initiator group configuration. Each ESX host has an active optimized path to the local copy and an active non-optimized path to the mirrored copy.

## Normal operation

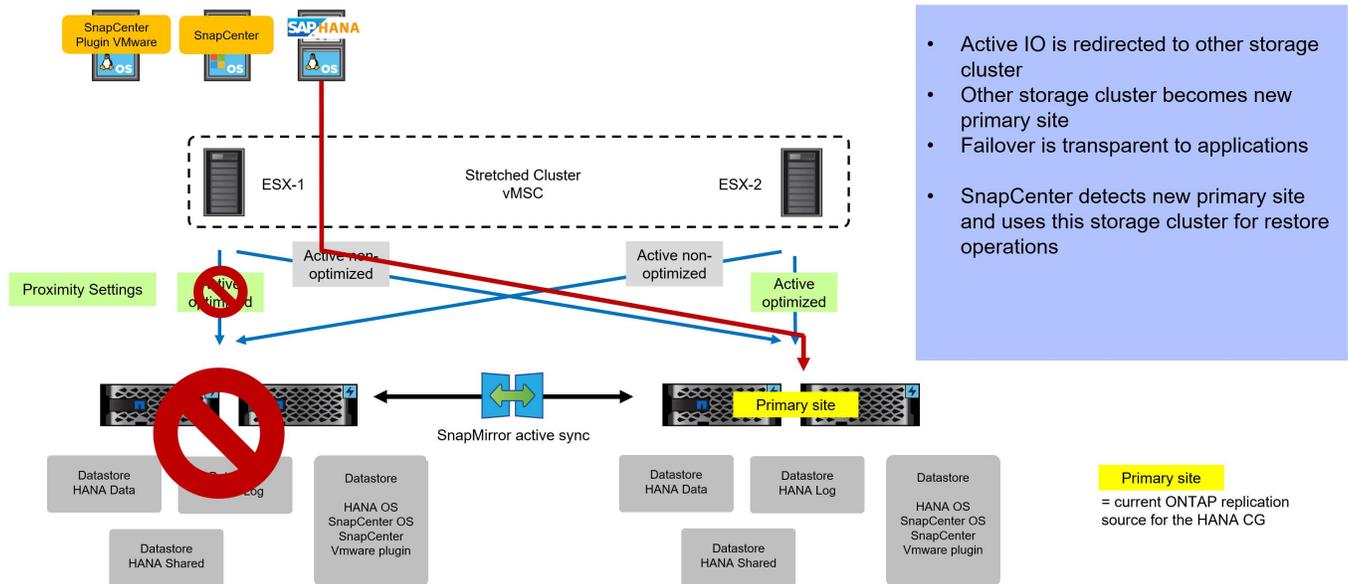
During normal operation the HANA system reads and writes from/to the local copy based on the active optimized path from ESX host ESX-1. With each backup operation, SnapCenter detects the current primary site for the replication relationship and executes the backup operations against the primary site. The Snapshots are replicated to the mirrored copy and are available at both sites. A SnapCenter restore operation would be executed at the primary site.



- Stretched FC SAN across both fault domains
- Proximity settings in initiator group configuration define primary path for IO
- SnapCenter
  - detects primary site
  - backups create Snapshots on both storage clusters
  - restore operations are always done at the primary site

### Storage failure

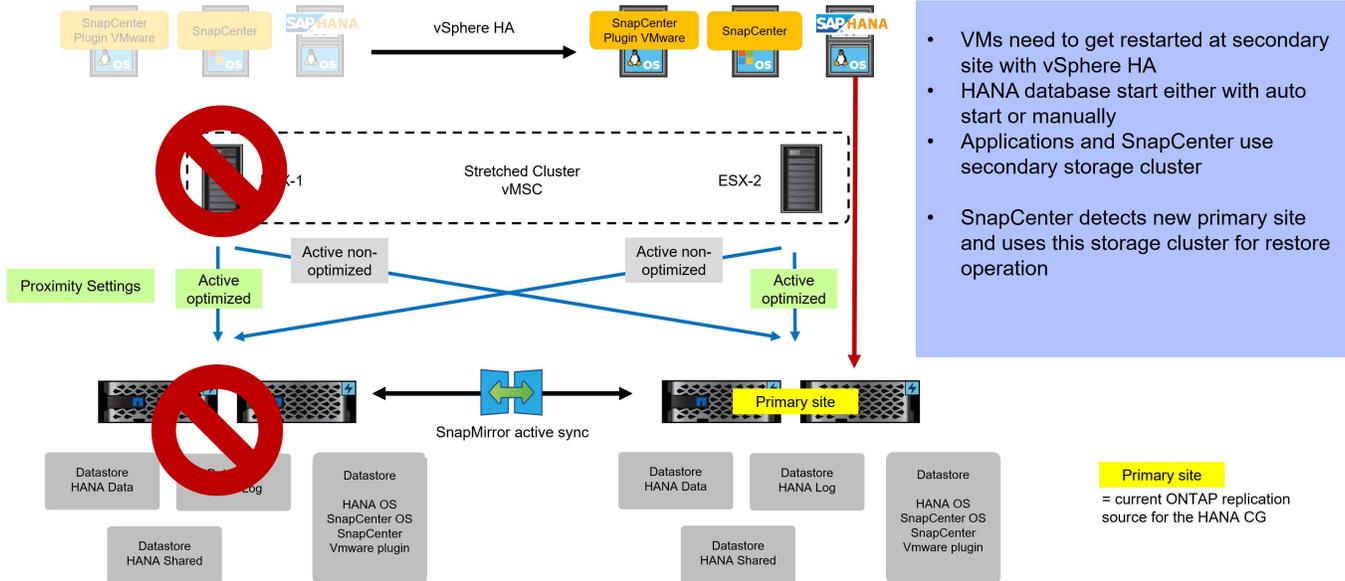
If the storage system at site 1 fails, the HANA systems access the mirrored copy at site 2 and continues operation. The primary site switches to the secondary site and SnapCenter now executes backup and restore operations at the new primary site.



- Active IO is redirected to other storage cluster
- Other storage cluster becomes new primary site
- Failover is transparent to applications
- SnapCenter detects new primary site and uses this storage cluster for restore operations

### Site failure

In case of a site failure, the HANA VM as well as SnapCenter and the SnapCenter for VMware plugin VM will fail over to the ESX host at the secondary site using vSphere HA. The HANA database needs to get started and will then access the mirrored copy at the second site. The primary site switches to the secondary site and SnapCenter now executes backup and restore operations at the new primary site.

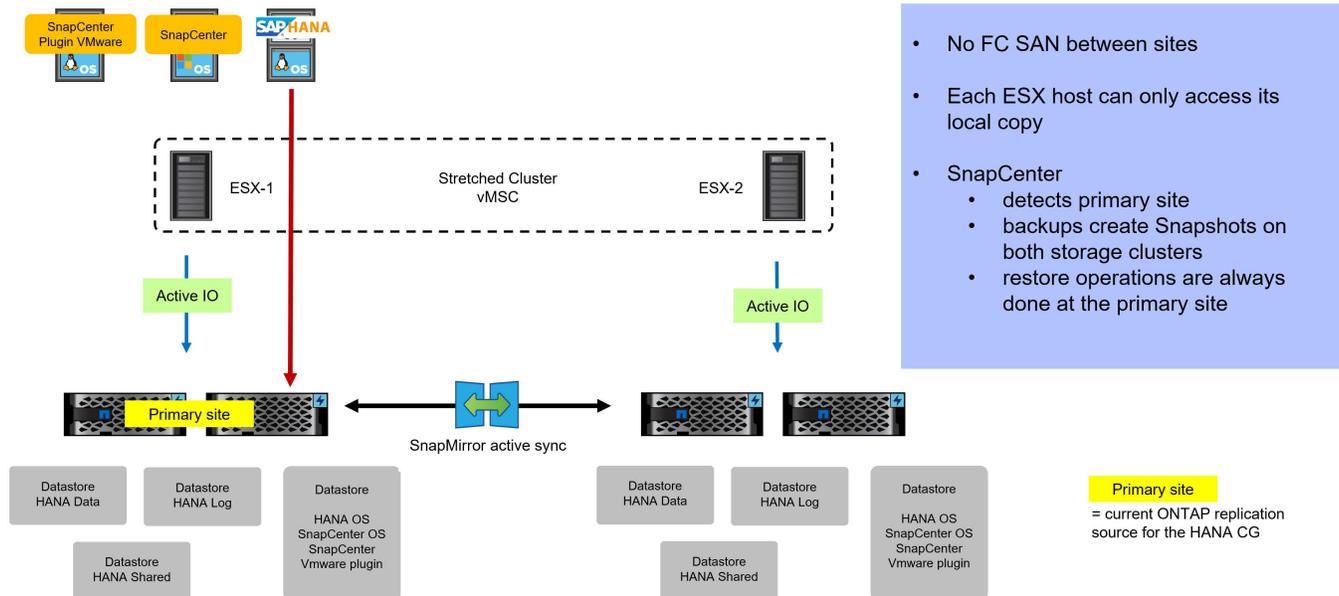


## Non-uniform access setup

In a non-uniform access configuration, the fibre channel SAN is not stretched across both sites. Each ESX host at each site can only access the local copy of the data sets.

## Normal operation

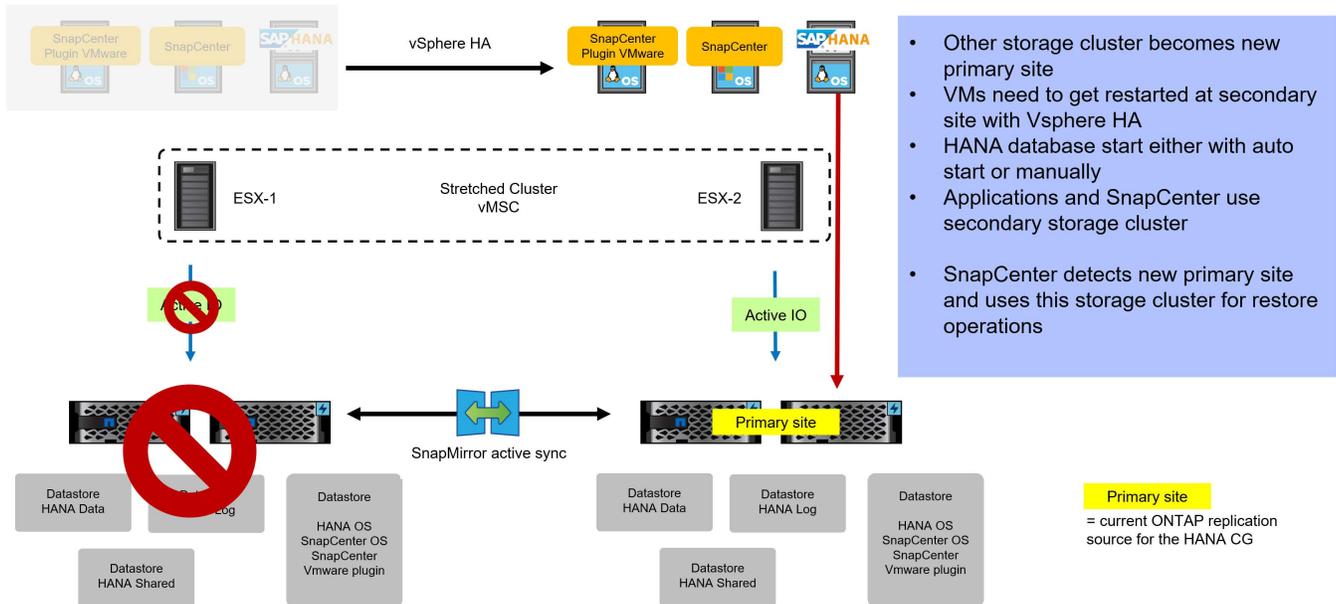
During normal operation the HANA system reads and writes from/to the local copy. With each backup operation, SnapCenter detects the current primary site for the replication relationship and executes the backup operations against the primary site. The Snapshots are replicated to the mirrored copy and are available at both sites. A SnapCenter restore operation would be executed at the primary site.



## Storage failure

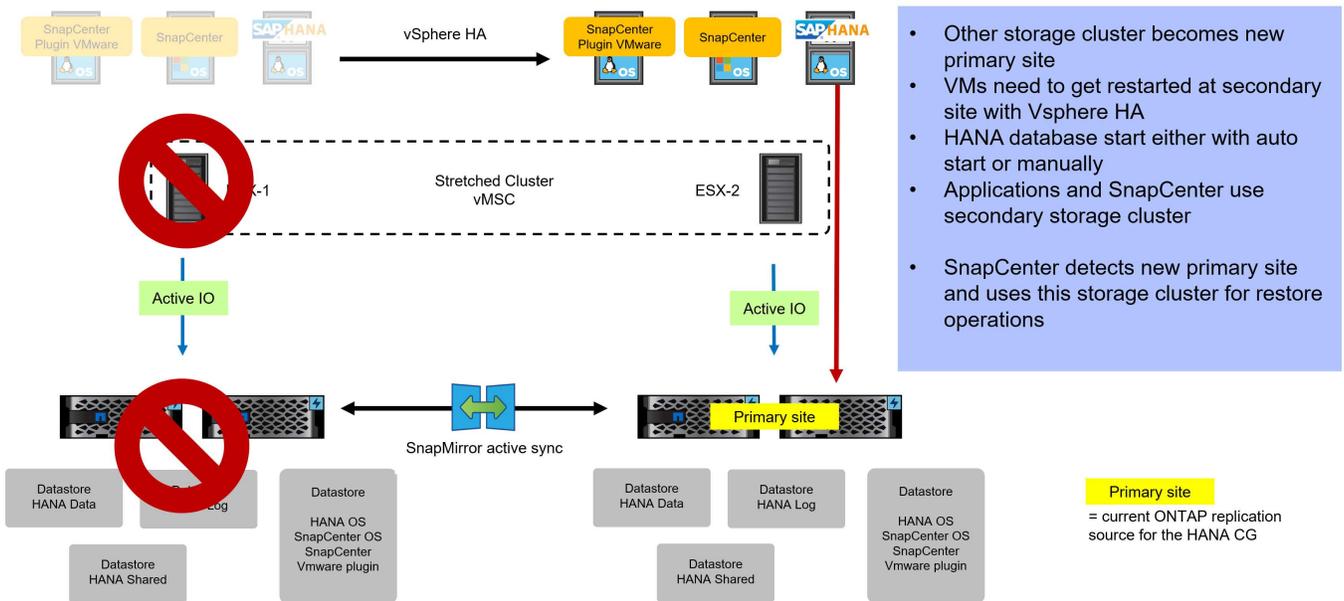
In case of a storage failure, the HANA VM as well as SnapCenter and the SnapCenter for VMware plugin VM will fail over to the ESX host at the secondary site using vSphere HA. The HANA database needs to get started and will then access the mirrored copy at the second site. The primary site switches to the secondary site and

SnapCenter now executes backup and restore operations at the new primary site.



### Site failure

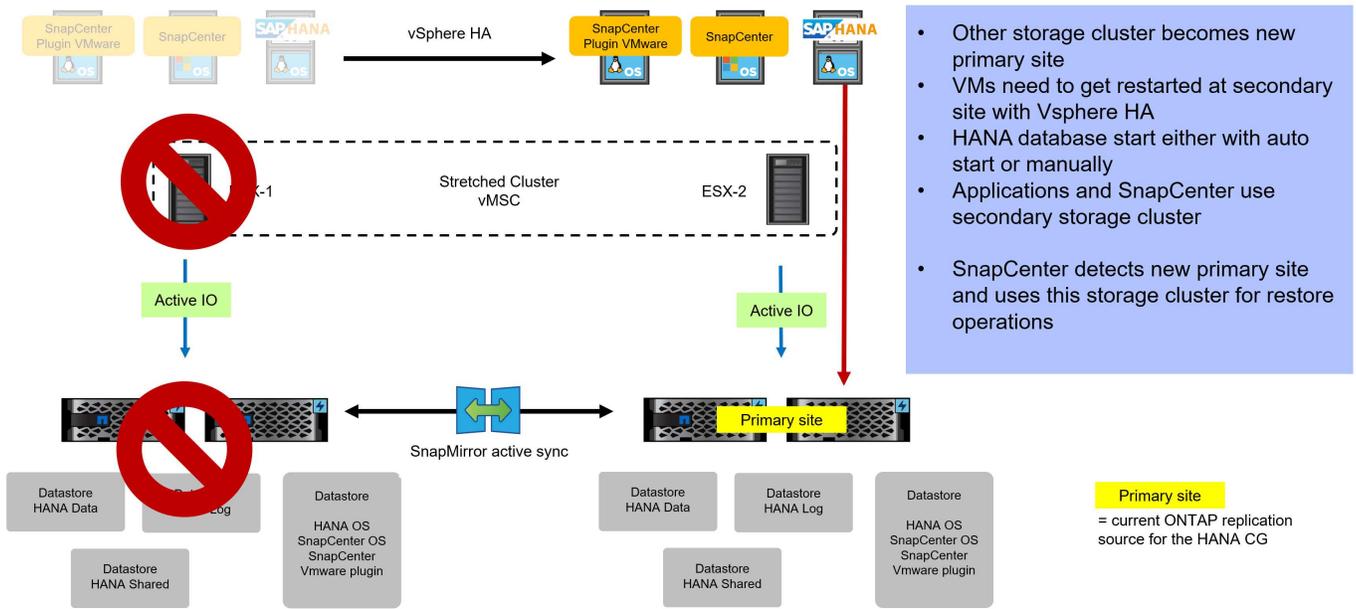
Same as storage failure.



### Relocation of HANA VM or primary site

If the HANA VM is relocated to the other ESX host and the primary site of the storage remains the same, a restore operation with SnapCenter will fail. Since SnapCenter uses the primary site to execute restore operations, the clone will be created at the left side, while the HANA VM runs on the right side. Since there is no data path between the sites, SnapCenter will not be copy the data.

As a workaround you need to make sure, that the relocation of VM and primary side is done together, or you need to failover the primary site before the restore operation with SnapCenter.



## Additional information and version history

This article provides links to additional resources relevant to this solution.

SnapCenter:

[TR-4614: SAP HANA backup and recovery with SnapCenter](#)

[TR-4719: SAP HANA System Replication - Backup and Recovery with SnapCenter](#)

[TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)

[SnapCenter Software documentation](#)

SnapMirror active sync:

[SnapMirror active sync overview in ONTAP](#)

[NetApp ONTAP with NetApp SnapMirror active sync with VMware vSphere Metro Storage Cluster \(vMSC\).](#)

[VMware vSphere Metro Storage Cluster with SnapMirror active sync](#)

[VMware vSphere Metro Storage Cluster \(vMSC\)](#)

Version history:

Version	Date	Comment
Version 1.0	March 2025	Initial version

# SAP HANA data protection with SnapCenter with VMware VMFS and NetApp ASA systems

# SAP HANA data protection with SnapCenter with VMware VMFS and NetApp ASA systems

This document outlines the best practices for data protection using SnapCenter for HANA systems running on VMware with datastores using VMFS and LUNs stored on NetApp ASA systems.

Author: Nils Bauer, NetApp

## Scope of this document

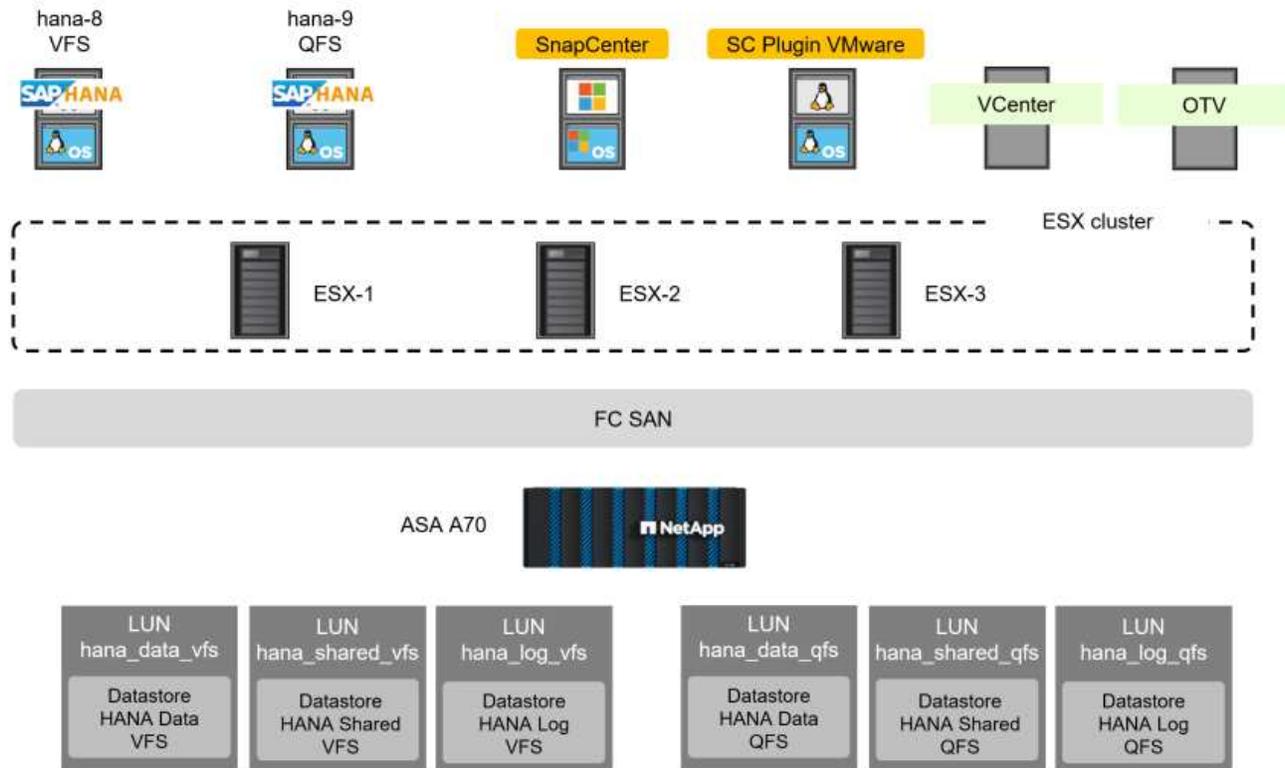
It does not serve as a step-by-step guide for configuring the entire environment but focuses on details specific to SnapCenter and HANA on VMFS, including:

- Setting up SAP HANA systems with VMware VMFS
- Specific SnapCenter configurations for HANA on VMware with VMFS
- SnapCenter backup, restore, and recovery operations for HANA on VMware with VMFS
- SnapCenter SAP System Refresh operations for HANA on VMware with VMFS

For further information and detailed configuration instructions, refer to the documents listed in the [“Additional Information”](#) chapter.

## Lab setup used for this document

The figure below presents a high-level overview of the lab setup utilized. Two single-host HANA MDC systems are used to demonstrate the various operations. The HANA system VFS is designated for executing backup, restore, and recovery operations, while the HANA system QFS serves as the target system for SAP System Refresh operations. The SnapCenter plug-in for VMware is essential for enabling SnapCenter to manage HANA resources configured with VMware VMFS. Although ONTAP tools for VMware were used to provision the storage units for the HANA systems, they are not a mandatory component.



## Software versions

Software	Version
ONTAP	ASA A70 ONTAP 9.16.1
vSphere client	8.0.3
ESXi	8.0.3
SnapCenter plugin for vSphere	6.1.0
ONTAP tools for VMware vSphere	10.4
Linux OS	SLES for SAP 15 SP6
SAP HANA	2.0 SPS8
SnapCenter	6.1P1

## HANA system provisioning and installation

This chapter describes the installation and configuration of the SAP HANA system specific to a VMware setup using VMFS. Additional generic best practices can be found at [SAP HANA on NetApp ASA Systems with Fibre Channel Protocol](#).

### Storage configuration

To meet the storage performance KPIs defined by SAP for production HANA systems, dedicated LUNs and datastores must be configured for the data and log filesystems of the HANA system. Datastores must not be shared among multiple HANA systems or other workloads.

ONTAP tools for VMware (OTV) has been used to provision the three datastores for the HANA system VFS.

- hana\_data\_VFS
- hana\_log\_VFS
- hana\_shared\_VFS



The datastore for the HANA shared filesystem can also be shared across multiple HANA systems.

The screenshot shows the vSphere Client interface for a SAPCC environment. The left sidebar displays a tree view of the environment, including datastores like 'hana\_data\_VFS', 'hana\_log\_VFS', and 'hana\_shared\_VFS'. The main panel shows 'Datacenter Details' with 3 hosts, 38 virtual machines, 1 cluster, 8 networks, and 10 datastores. The 'Capacity and Usage' section provides a breakdown of resources: CPU (152.74 GHz free, 9.26 GHz used), Memory (2.63 TB free, 375.09 GB used), and Storage (7.87 TB capacity, 1.86 TB used). A 'Tags' section indicates no tags are assigned. Below this, a 'Recent Tasks' table shows the completion of VMFS datastore updates and VMFS creation tasks.

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Process VMFS datastore updates	10.63.167.6	Completed		System	4 ms	05/19/2025, 9:20:23 AM	05/19/2025, 9:20:23 AM	vcenter@sapcc-stf.netapp.com
Process VMFS datastore updates	10.63.167.4	Completed		System	5 ms	05/19/2025, 9:20:23 AM	05/19/2025, 9:20:23 AM	vcenter@sapcc-stf.netapp.com
Create VMFS datastore	10.63.167.14	Completed		SAPCC.VCENTER.Administrat	10 ms	05/19/2025, 9:20:22	05/19/2025, 9:20:23	vcenter@sapcc-stf.netapp.com

At the storage system three LUNs have been created by OTV.

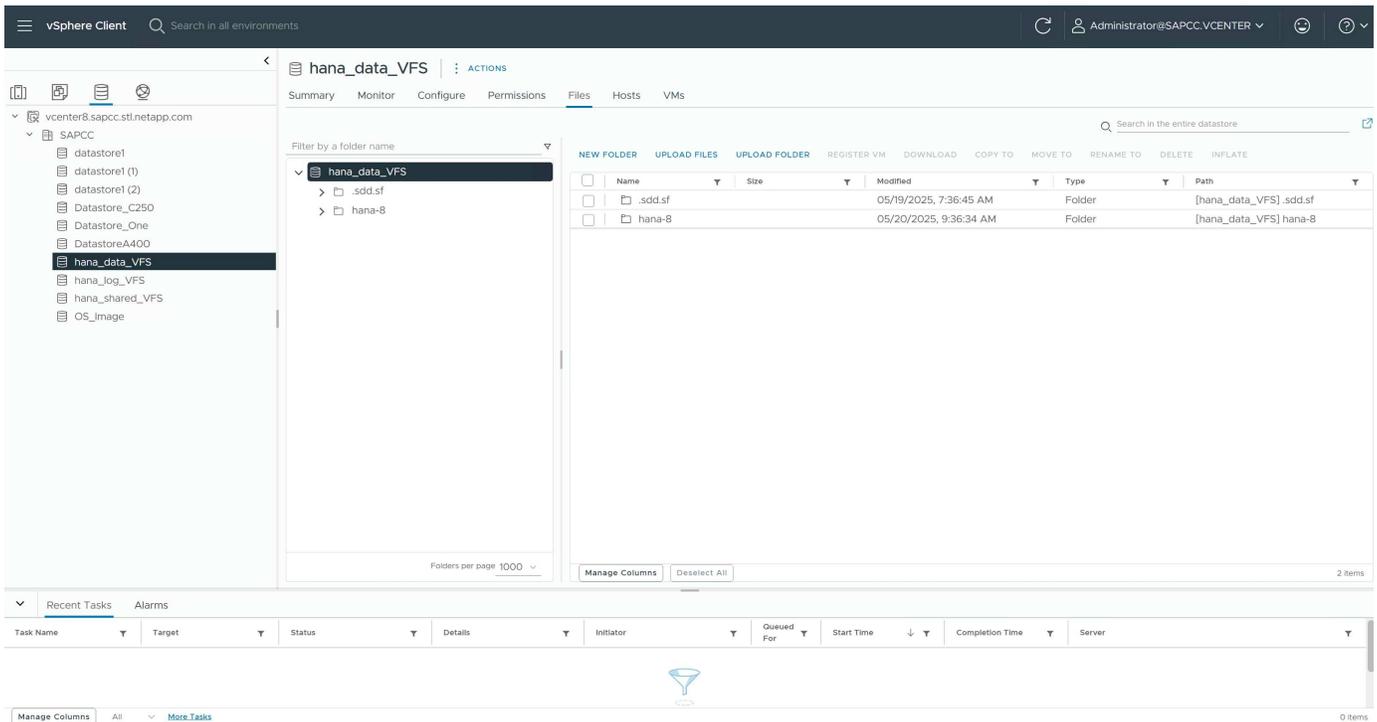
The screenshot shows the NetApp ONTAP System Manager interface for an A70-SAPCC system. The 'Storage' section displays 19 storage units, 68.6 TiB available, and 19 online units. A table below lists three LUNs created by OTV:

Name	Consistency group	Capacity	Data reduction	Host mapping	IOPS	Latency (ms)	Throughput (MB/s)
hana_data_VFS	-	100 GiB	8.75 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	0	0	0
hana_log_VFS	-	100 GiB	8.69 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	0	0	0
hana_shared_VFS	-	100 GiB	3.13 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	0	0	0

## VM disk configuration

Three new disks (VMDK) must be added to the HANA VM. Each disk within one of the datastores which have been created before as illustrated in the picture below.





When the three disk have been added to the VM, they can be listed at the OS level.

```

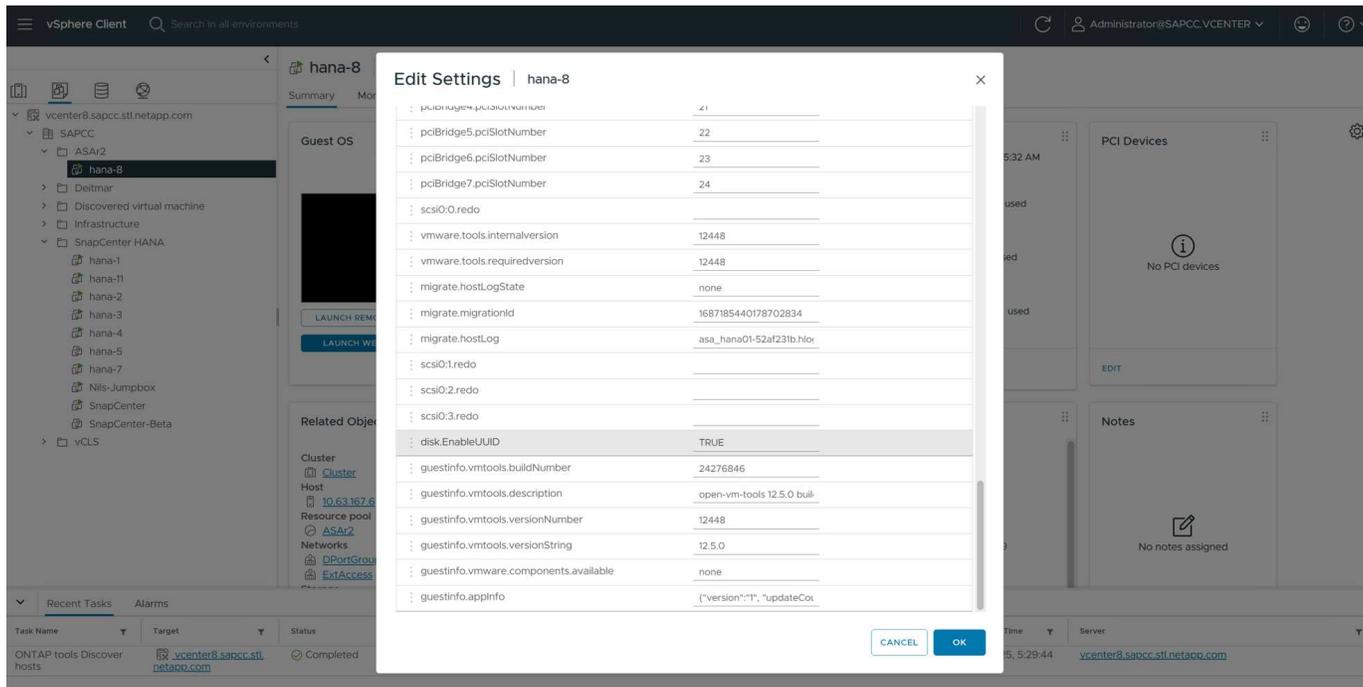
hana-8:~ # lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINTS
sda 8:0 0 100G 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
└─sda2 8:2 0 82G 0 part
   └─system-root 254:0 0 60G 0 lvm /root
      /var
      /usr/local
      /tmp
      /srv
      /opt
      /home
      /boot/grub2/x86++_++64-efi
      /boot/grub2/i386-pc
      /.snapshots
      /
   └─system-swap 254:1 0 2G 0 lvm [SWAP]
sdb 8:16 0 95G 0 disk
sdc 8:32 0 95G 0 disk
sdd 8:48 0 95G 0 disk
sr0 11:0 1 17.1G 0 rom

```

### VM parameter disk.EnableUUID

This parameter must be set accordingly, otherwise SnapCenter database auto discovery will fail.

1. Shutdown VM
2. Add new parameter “disk.EnableUUID” and set to “TRUE”
3. Start VM



## File system preparation at Linux host

### Creation of xfs filesystem on new disks

An xfs file system has been created on each of the three new disks.

```
hana-8:~ # mkfs.xfs /dev/sdb
meta-data=/dev/sdb isize=512 agcount=4, agsize=6225920 blks
= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=1, sparse=1, rmapbt=1
= reflink=1 bigtime=1 inobtcount=0 nnext64=0
data = bsize=4096 blocks=24903680, imaxpct=25
= sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=16384, version=2
= sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
Discarding blocks...Done.
```

```
hana-8:~ # mkfs.xfs /dev/sdc
meta-data=/dev/sdc isize=512 agcount=4, agsize=6225920 blks
= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=1, sparse=1, rmapbt=1
= reflink=1 bigtime=1 inobtcount=0 nnext64=0
data = bsize=4096 blocks=24903680, imaxpct=25
= sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=16384, version=2
= sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
Discarding blocks...Done.
```

```
hana-8:~ # mkfs.xfs /dev/sdd
meta-data=/dev/sdd isize=512 agcount=4, agsize=6225920 blks
= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=1, sparse=1, rmapbt=1
= reflink=1 bigtime=1 inobtcount=0 nnext64=0
data = bsize=4096 blocks=24903680, imaxpct=25
= sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=16384, version=2
= sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
Discarding blocks...Done.
```

```
hana-8:~ #
```

## Creation of mount points

```
hana-8:/ # mkdir -p /hana/data/VFS/mnt00001
hana-8:/ # mkdir -p /hana/log/VFS/mnt00001
hana-8:/ # mkdir -p /hana/shared
hana-8:/ # chmod -R 777 /hana/log/SMA
hana-8:/ # chmod -R 777 /hana/data/SMA
hana-8:/ # chmod -R 777 /hana/shared
```

### Configuration of /etc/fstab

```
hana-8:/ # cat /etc/fstab

/dev/system/root / btrfs defaults 0 0
/dev/system/root /var btrfs subvol=@/var 0 0
/dev/system/root /usr/local btrfs subvol=@/usr/local 0 0
/dev/system/root /tmp btrfs subvol=@/tmp 0 0
/dev/system/root /srv btrfs subvol=@/srv 0 0
/dev/system/root /root btrfs subvol=@/root 0 0
/dev/system/root /opt btrfs subvol=@/opt 0 0
/dev/system/root /home btrfs subvol=@/home 0 0
/dev/system/root /boot/grub2/x86+_+_64-efi btrfs
subvol=@/boot/grub2/x86+_+_64-efi 0 0
/dev/system/root /boot/grub2/i386-pc btrfs subvol=@/boot/grub2/i386-pc 0
0
/dev/system/swap swap swap defaults 0 0
/dev/system/root /.snapshots btrfs subvol=@/.snapshots 0 0
UUID=FB79-24DC /boot/efi vfat utf8 0 2
### SAPCC_share
192.168.175.86:/sapcc_share /mnt/sapcc-share nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0
0
/dev/sdb /hana/data/VFS/mnt00001 xfs relatime,inode64 0 0
/dev/sdc /hana/log/VFS/mnt00001 xfs relatime,inode64 0 0
/dev/sdd /hana/shared xfs defaults 0 0
hana-8:/ #

hana-8:/ # df -h
Filesystem Size Used Avail Use% Mounted on
/dev/mapper/system-root 60G 4.4G 54G 8% /
devtmpfs 4.0M 0 4.0M 0% /dev
tmpfs 49G 0 49G 0% /dev/shm
efivarfs 256K 57K 195K 23% /sys/firmware/efi/efivars
tmpfs 13G 18M 13G 1% /run
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup-dev-
early.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-sysctl.service
```

```

tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup-dev.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-vconsole-setup.service
/dev/mapper/system-root 60G 4.4G 54G 8% /.snapshots
/dev/mapper/system-root 60G 4.4G 54G 8% /boot/grub2/i386-pc
/dev/mapper/system-root 60G 4.4G 54G 8% /boot/grub2/x86++_++64-efi
/dev/mapper/system-root 60G 4.4G 54G 8% /home
/dev/mapper/system-root 60G 4.4G 54G 8% /opt
/dev/mapper/system-root 60G 4.4G 54G 8% /srv
/dev/mapper/system-root 60G 4.4G 54G 8% /tmp
/dev/mapper/system-root 60G 4.4G 54G 8% /usr/local
/dev/mapper/system-root 60G 4.4G 54G 8% /var
/dev/sda1 253M 5.9M 247M 3% /boot/efi
/dev/mapper/system-root 60G 4.4G 54G 8% /root
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup.service
tmpfs 6.3G 72K 6.3G 1% /run/user/464
tmpfs 1.0M 0 1.0M 0% /run/credentials/getty@tty1.service
tmpfs 6.3G 52K 6.3G 1% /run/user/0
192.168.175.86:/sapcc_share 1.4T 840G 586G 59% /mnt/sapcc-share
/dev/sdb 95G 1.9G 94G 2% /hana/data/VFS/mnt00001
/dev/sdc 95G 1.9G 94G 2% /hana/log/VFS/mnt00001
/dev/sdd 95G 1.9G 94G 2% /hana/shared

hana-8:/ #

```

## HANA installation

The HANA installation can now be executed.

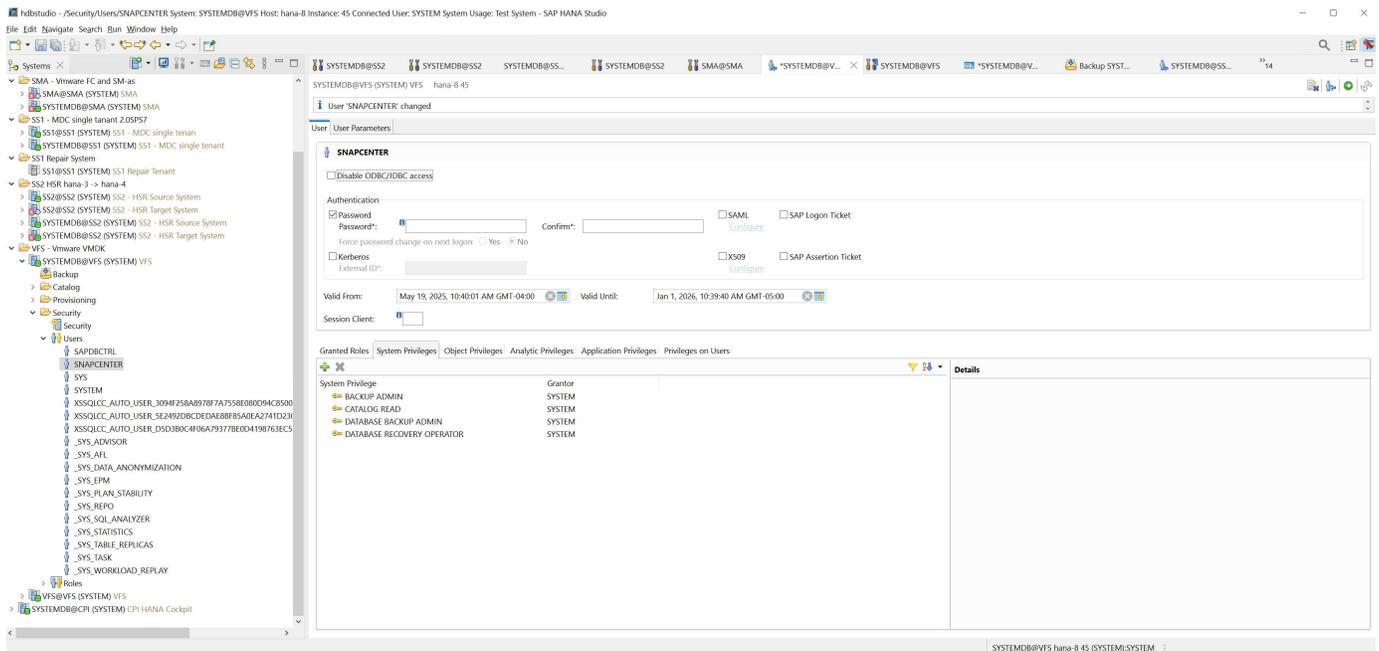


With the described configuration the `/usr/sap/VFS` directory will be on the OS VMDK. If `/usr/sap/VFS` should be stored in the shared VMDK, the hana shared disk could be partitioned to provide another file system for `/usr/sap/VFS`.

## HANA configuration

### Configure SnapCenter database user

A user store for a system database user must be created, which should be used by SnapCenter.



## Configure hdb userstore key

A user store key must be created for the user vfsadm. The HANA instance number must be set accordingly for communication the port. In our setup instance number “45” is used.

```
vfsadm@hana-8:/usr/sap/VFS/HDB45> hdbuserstore SET VFSKEY hana-8:34513
SNAPCENTER <password>
```

Retroactive report: Operation succeed.

Check access with:

```
vfsadm@hana-8:/usr/sap/VFS/HDB45> hdbsql -U VFSKEY

Welcome to the SAP HANA Database interactive terminal.
Type: \h for help with commands
\q to quit
hdbsql SYSTEMDB=> exit

vfsadm@hana-8:/usr/sap/VFS/HDB45>
```

## SnapCenter configuration

### Pre-requisites

#### SnapCenter HANA resource must be auto discovered

Resources configured with VMware VMFS must be auto discovered by SnapCenter to enable specific operations required for these configurations.

Since HANA non-data volumes are always manually configured resources in SnapCenter, they are not supported by SnapCenter with VMFS.

SAP HANA multiple host systems must be configured using a central HANA plugin and are therefore manually configured by default. Such systems are also not supported by SnapCenter when using VMware VMFS.

#### **SnapCenter for VMWare vSphere plugin**

The SnapCenter for VMware vSphere plugin must be deployed in the VMware environment.

#### **Storage SVM management IP**

Storage SVMs hosting the LUN's must have a management interface configured, otherwise the SVMs will not be listed in SnapCenter when adding storage with the "add cluster" option and auto discovery operation will fail.

## Job Details



Discover resources for host 'hana-8.sapcc.stl.netapp.com'

✘ ▼ Discover resources for host 'hana-8.sapcc.stl.netapp.com'

✘ ▼ hana-8.sapcc.stl.netapp.com

✘ ▼ Discover

✔ ▶ Complete Application Discovery

✔ ▶ Discover Filesystem Resources

✘ ▶ Discover Virtual Resources

✔ ▶ Discover\_OnFailure

✘ Failure in virtual resources discovery: [Failed to resolve the storage associated with the VMware virtual disks 6000c2964ec4375910dc9953d9f870ca]

View Logs

Cancel Job

Close

Name	IP	Cluster Name	User Name	Platform	Controller License
svm1	10.63.167.55	10.63.167.54		ASA	✓
hana		10.63.150.245		AFF	✓
hana-backup	10.63.150.246	10.63.150.245		AFF	✓
hana-cloud-dr		10.1.2.175		FSx	Not applicable
hana-dr	10.63.150.247	10.63.150.245		AFF	✓
hana-primary	10.63.150.248 ...	10.63.150.245		AFF	✓

### VM disk parameter

The parameter must be set as described in chapter “VM parameter disk.EnableUUID”, otherwise SnapCenter database auto discovery will fail.

**Configure Database**

Plug-in host: hana-8.sapcc.stl.netapp.com

HDBSQL OS User: vfsadm

HDB Secure User Store Key: VFSKEY

**Failure in getting storage details: [Failed to retrieve the unit serial number for the device '/dev/sdb', Reason: 'SCSI inquiry failed. Check if the disk.EnableUUID parameter is set to TRUE in the VM configuration file.']**

Buttons: Cancel, OK

### Configure SnapCenter to use REST APIs for storage communication

SnapCenter must be configured to use REST APIs for storage communications. Otherwise, the create Snapshot operation will fail with the error message shown below.

Job Details ✕

Backup of Resource Group 'hana-8\_sapcc\_stl\_netapp\_com\_hana\_MDC\_VFS' with policy 'LocalSnap'

- ✕ ▾ Backup of Resource Group 'hana-8\_sapcc\_stl\_netapp\_com\_hana\_MDC\_VFS' with policy 'LocalSnap'
- ✕ ▾ hana-8.sapcc.stl.netapp.com
- ✕ ▾ Backup
  - ✓ ▶ Validate Dataset Parameters
  - ✓ ▶ Validate Plugin Parameters
  - ✓ ▶ Complete Application Discovery
  - ✓ ▶ Initialize Filesystem Plugin
  - ✓ ▶ Discover Filesystem Resources
  - ✓ ▶ Discover Virtual Resources
  - ✓ ▶ Populate storage details
  - ✓ ▶ Validate Retention Settings
  - ✓ ▶ Quiesce Application
  - ✓ ▶ Quiesce Filesystem
  - ✕ ▾ Create Snapshot
  - ⚠ ▶ Backup\_OnFailure

✖ SCC-STORAGE-02002: Creating Snapshot copy [SnapCenter\_hana-8\_LocalSnap\_Hourly\_05-20-2025\_10.33.58.2195] on storage resource [svm1:hana\_data\_VFS] failed with error [Snapshot operation failed. [400]: POST, DELETE, and PATCH requests on the snapshot session endpoint are not supported on this platform.]

The parameter "IsRestEnabledForStorageConnection" in the configuration file C:\Program Files\NetApp\SMCore\SMCoreServiceHost.dll.config must be set to "true".

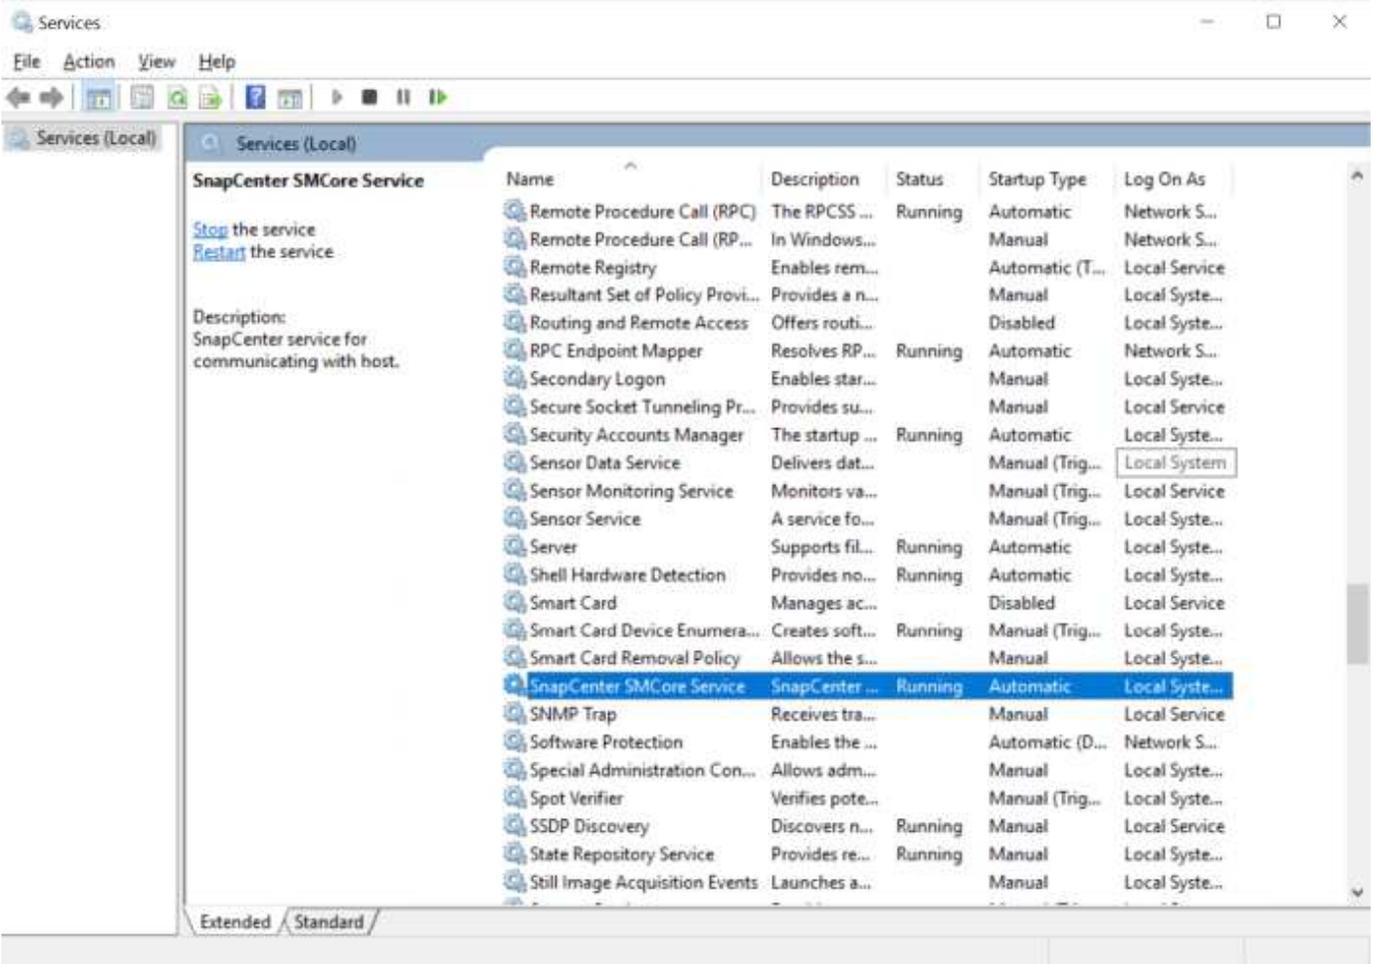
```
<add key="IsRestEnabledForStorageConnection" value="true" />
```

```

SMCoreServiceHost.dll.config - Notepad
File Edit Format View Help
<add key="EnableCancelJob" value="true" />
<add key="PSErrorString" value="Internal network error,API invoke failed,No such file or directory" />
<add key="CommandErrorDuringMccFailure" value="timed out,Unknown internal error,API invoke failed,metrocluster" />
<add key="VolumeEnumerationOptimized" value="true" />
<add key="CloneSplitStatusCheckPollTime" value="300000" />
<add key="ConfigCheckerJobStatusTimeout" value="20" />
<add key="ConfigCheckerJobStatusRetry" value="30" />
<add key="AzureEnvironment" value="AzureGlobalCloud" />
<add key="AzureLongRunningOperationRetryTimeoutInSec" value="20" />
<add key="AzureClientType" value="sdk" />
<add key="AzureThreadSleepTime" value="10000" />
<add key="AzureRestVersion" value="2019-11-01" />
<add key="GetStorageIDBeforeCacheInitialize" value="true" />
<add key="SccCloneSuffix" value="Clone" />
<add key="SourceComponent" value="smcore" />
<add key="WmiTimeoutIntervalMinutes" value="30" />
<add key="IsWmiTimeoutSet" value="true" />
<add key="OracleAlmActivityParallelExecution" value="true" />
<add key="OracleAlmActivityParallelMountInterval" value="20" />
<add key="OracleAlmActivityParallelUnmountInterval" value="10" />
<add key="SkipOracleAlmBackupsCatalogAndUncatalog" value="false" />
<add key="UseVolumeFilterInGetSnapshot" value="true" />
<add key="EnablePredefinedWindowsScriptsDirectory" value="true" />
<add key="PredefinedWindowsScriptsDirectory" value="C:\Program Files\NetApp\SMCore\Scripts" />
<add key="IsRestEnabledForStorageConnection" value="true" />
<add key="ExecutePredefinedWindowsScriptsCommands" value="Add-NetLunMap" />
<add key="MinOntapVersionToUseREST" value="9.13.1" />
<add key="IS_COLO_SNAPCENTER_AGENT" value="true" />
<add key="IS_SCN_PLUGIN_SERVICE_PRESENT" value="false" />
<add key="SMCORE_IMAGE_PATH" value="C:\Program Files\NetApp\SMCore\" />
<add key="REPOSITORY_PATH" value="C:\ProgramData\NetApp\SnapCenter" />
<add key="SNAPGATHERS_PATH" value="C:\Program Files\NetApp\SnapGathers\" />
<add key="SNAPGATHERS_PATH_WINDOWS" value="C:\Program Files\NetApp\SnapCenter\SnapGathers\" />
<add key="smcoreprotocol" value="https" />
<add key="SERVICE_CERTIFICATE_PATH" value="/var/opt/snapcenter/certs/snapcenter.pfx" />
<add key="SERVICE_CERTIFICATE_PASSWORD" value="" />
<add key="ForceSHA256EncryptionKey" value="false" />
<add key="WINRM_PROTOCOL" value="http" />
<add key="WINRM_PORT" value="5985" />
<add key="WINRM_AUTH_TYPE" value="ntlm" />
<add key="DoNotSaveOracleBlob" value="false" />
<add key="IsRestEnabledForLowerONTAP" value="false" />
</appSettings>
</configuration>

```

After the change has been made, SnapCenter SMCore Service must be stopped and started.

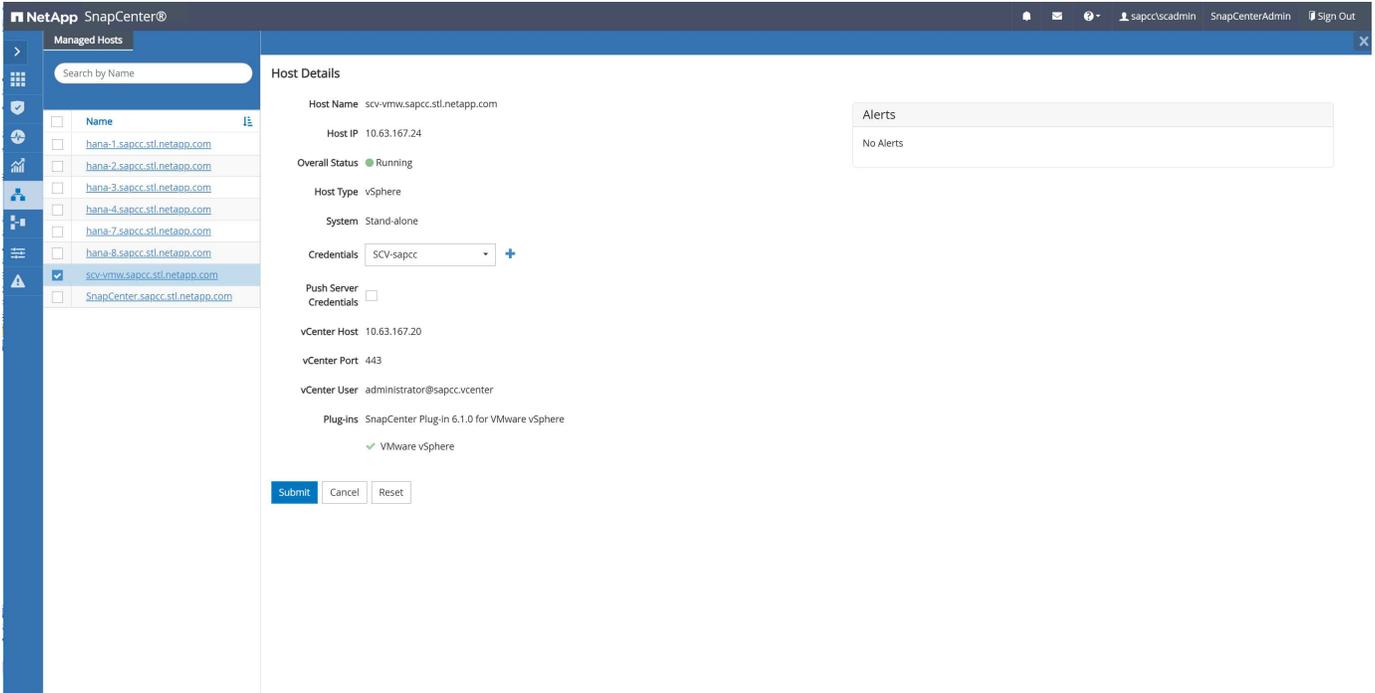


## Add VMware Plugin to SnapCenter

Before the host can be added in SnapCenter the SnapCenter plugin for VMware vSphere must be deployed within the VMware environment. See also [Deploy SnapCenter Plug-in for VMware vSphere](#).



Credentials must be set during host add workflow, where vSphere can be selected as a host type.

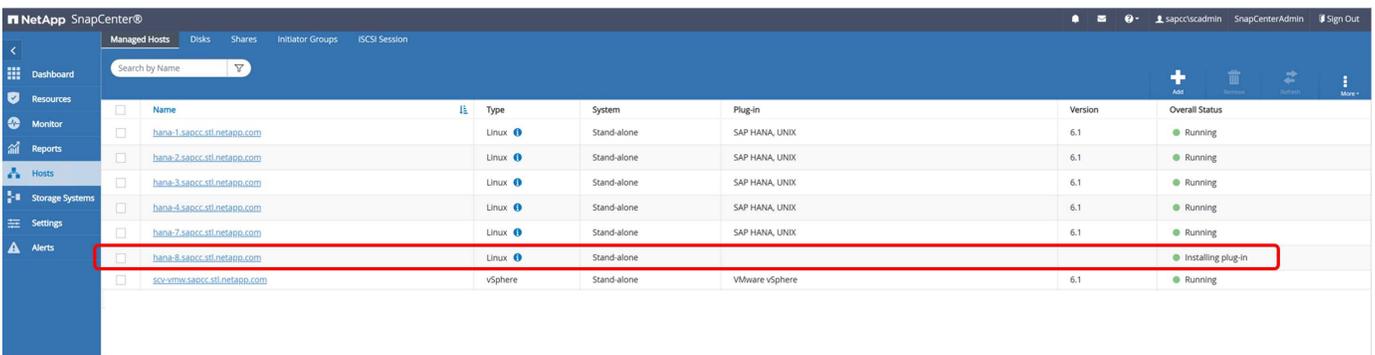


The screenshot shows the NetApp SnapCenter interface. On the left, a sidebar lists 'Managed Hosts' with a search bar and a list of hosts. The main area displays 'Host Details' for the host 'scv-vmw.sapcc.stl.netapp.com'. The details include: Host Name, Host IP (10.63.167.24), Overall Status (Running), Host Type (vSphere), System (Stand-alone), Credentials (SCV-sapcc), Push Server Credentials (unchecked), vCenter Host (10.63.167.20), vCenter Port (443), vCenter User (administrator@sapcc.vcenter), and Plug-ins (SnapCenter Plug-in 6.1.0 for VMware vSphere, with VMware vSphere checked). At the bottom are 'Submit', 'Cancel', and 'Reset' buttons. An 'Alerts' box on the right shows 'No Alerts'.

## Add HANA host



No specific requirements. Plugin deployment and auto discovery is done as usual.



The screenshot shows the NetApp SnapCenter interface displaying a list of managed hosts. The table has columns for Name, Type, System, Plug-in, Version, and Overall Status. A red box highlights the row for 'hana-8.sapcc.stl.netapp.com', which is in the 'Installing plug-in' state. The other hosts are in the 'Running' state.

Name	Type	System	Plug-in	Version	Overall Status
hana-1.sapcc.stl.netapp.com	Linux	Stand-alone	SAP HANA, UNIX	6.1	Running
hana-2.sapcc.stl.netapp.com	Linux	Stand-alone	SAP HANA, UNIX	6.1	Running
hana-3.sapcc.stl.netapp.com	Linux	Stand-alone	SAP HANA, UNIX	6.1	Running
hana-4.sapcc.stl.netapp.com	Linux	Stand-alone	SAP HANA, UNIX	6.1	Running
hana-7.sapcc.stl.netapp.com	Linux	Stand-alone	SAP HANA, UNIX	6.1	Running
hana-8.sapcc.stl.netapp.com	Linux	Stand-alone	SAP HANA, UNIX	6.1	Installing plug-in
scv-vmw.sapcc.stl.netapp.com	vSphere	Stand-alone	VMware vSphere	6.1	Running

With the auto discovery process SnapCenter detects that the HANA resource is running virtualized with VMFS.

The screenshot displays the NetApp SnapCenter interface for a SAP HANA resource. The left sidebar shows a list of systems: QS1, SM1, SS1, SS2, SS2, and VFS. The main panel shows details for the selected resource, VFS.

**Resource - Details**

Details for selected resource

Type	Multitenant Database Container
HANA System Name	VFS
SID	VFS
Tenant Databases	VFS
Plug-In Host	hana-8.sapcc.stl.netapp.com
HDB Secure User Store Key	VFSKEY
HDBSQL OS User	vfsadm
Log backup location	/usr/sap/VFS/HDB45/backup/log
Backup catalog location	/usr/sap/VFS/HDB45/backup/log
System Replication	None
Plug-in name	SAP HANA
Last backup	None
Resource Groups	None
Policy	None
Discovery Type	Auto

Storage Footprint

SVM	Volume	Junction Path	LUN/Qtree
svm1			hana_data_VFS

Activity: The 5 most recent jobs are displayed. 5 Completed, 0 Warnings, 0 Failed, 0 Canceled, 0 Running, 0 Queued.

## Policy and resource protection configuration

Nothing specific to VMware with VMFS.

## Backup operations

Nothing specific to VMware with VMFS.

## Job Details



Backup of Resource Group 'hana-8\_sapcc\_stl\_ne.....na\_MDC\_VFS' with policy 'LocalSnap'

✓ Backup of Resource Group 'hana-8\_sapcc\_stl\_netapp\_com\_hana\_MDC\_VFS' with policy 'LocalSnap'

✓ ▼ hana-8.sapcc.stl.netapp.com

✓ ▼ Backup

- ✓ ▶ Validate Dataset Parameters
- ✓ ▶ Validate Plugin Parameters
- ✓ ▶ Complete Application Discovery
- ✓ ▶ Initialize Filesystem Plugin
- ✓ ▶ Discover Filesystem Resources
- ✓ ▶ Discover Virtual Resources
- ✓ ▶ Populate storage details
- ✓ ▶ Validate Retention Settings
- ✓ ▶ Quiesce Application
- ✓ ▶ Quiesce Filesystem
- ✓ ▶ Create Snapshot
- ✓ ▶ UnQuiesce Filesystem
- ✓ ▶ UnQuiesce Application
- ✓ ▶ Get Snapshot Details
- ✓ ▶ Get Filesystem Metadata
- ✓ ▶ Get Virtualization Metadata
- ✓ ▶ Finalize Filesystem Plugin
- ✓ ▶ Collect Autosupport data
- ✓ ▶ Register Backup and Apply Retention
- ✓ ▶ Register Snapshot attributes
- ✓ ▶ Application Clean-Up
- ✓ ▶ Data Collection
- ✓ ▶ Agent Finalize Workflow

**i** Task Name: Backup Start Time: 05/21/2025 10:29:05 PM End Time: 05/21/2025 10:30:38 PM

View Logs

Cancel Job

Close

The screenshot shows the NetApp SnapCenter interface for a SAP HANA system. The left sidebar contains navigation options like System, Q51, SM1, SS1, SS2, and VFS. The main area is titled 'VFS Topology' and includes a 'Manage Copies' section with '12 Backups' and '0 Clones'. A 'Summary Card' on the right provides a high-level overview: 12 Backups, 12 Snapshot based backups, 0 File-Based backups, 0 Clones, and 0 Snapshots Locked. Below this is a 'Primary Backup(s)' table with columns for Backup Name, Snapshot Lock Expiration, Count, and End Date. The table lists 12 backup entries with their respective IDs and timestamps. At the bottom, an 'Activity' bar shows 5 Completed jobs, 0 Warnings, 0 Failed, 0 Canceled, 0 Running, and 0 Queued.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-8_LocalSnap_Hourly_05-22-2025_06.29.00.3706		1	05/22/2025 6:30:14 AM
SnapCenter_hana-8_LocalSnap_Hourly_05-22-2025_02.29.00.3541		1	05/22/2025 2:30:12 AM
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_22.29.03.2699		1	05/21/2025 10:30:19 PM
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_18.29.00.3956		1	05/21/2025 6:30:12 PM
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_14.29.00.3696		1	05/21/2025 6:30:12 PM
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_10.29.00.3581		1	05/21/2025 2:30:12 AM
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_06.29.00.3960		1	05/21/2025 6:30:12 AM
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_02.29.00.3515		1	05/21/2025 2:30:12 AM
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_22.29.00.3896		1	05/20/2025 10:30:12 PM
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_18.29.00.3611		1	05/20/2025 6:30:12 PM
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_14.29.00.3840		1	05/20/2025 2:30:12 PM
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_11.03.44.3420		1	05/20/2025 11:05:03 AM

SnapCenter creates a consistency group (CG) and adds the storage unit hana\_data\_VFS to the CG. Snapshots are created at CG level.

The screenshot shows the NetApp ONTAP System Manager 'Storage' page. It displays 19 Storage units, with 68.5 TiB available. A summary shows 19 Online units and 0 Offline units. Below this is a table listing storage units with columns for Name, Consistency group, Capacity, Data reduction, Host mapping, IOPS, Latency (ms), and Throughput (MB/s).

Name	Consistency group	Capacity	Data reduction	Host mapping	IOPS	Latency (ms)	Throughput (MB/s)
hana_data_VFS	sc20250520_110422_689	100 GiB	1 to 1	otv_host-44_e3d7e9d4-46f3-4fd1	1	0.07	0
hana_log_VFS	-	100 GiB	1.19 to 1	otv_host-44_e3d7e9d4-46f3-4fd1	4	0.23	0.41
hana_shared_VFS	-	100 GiB	2.8 to 1	otv_host-44_e3d7e9d4-46f3-4fd1	6	0.23	0.43

NetApp ONTAP System Manager | A70-SAPCC

Search actions, objects, and pages

Dashboard

Insights

Storage

Hosts

Network

Events & Jobs

Protection

Consistency groups

Policies

Replication

Cluster

← Back to consistency groups

sc20250520\_11...

Overview Snapshots Replication

Storage VM svm1

Storage units 1

Application type VMware

Protection Show uninitialized

Snapshots None

Replication None

Storage units

Delete Remove from consistency group

Name	Capacity	Host mapping
hana_data_VFS	100 GiB	otv_host-44_e3d7e9d4-46f3-4fda-aba3-00c1be4c0fcf +2

NetApp ONTAP System Manager | A70-SAPCC

Search actions, objects, and pages

Dashboard

Insights

Storage

Hosts

Network

Events & Jobs

Protection

Consistency groups

Policies

Replication

Cluster

← Back to consistency groups

sc20250520\_110422...

Overview Snapshots Replication

Policy: -

Name	Created	SnapMirror label
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_11.03.44.3420	May/20/2025 11:10 AM	
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_14.29.00.3840	May/20/2025 2:36 PM	
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_18.29.00.3611	May/20/2025 6:36 PM	
SnapCenter_hana-8_LocalSnap_Hourly_05-20-2025_22.29.00.3896	May/20/2025 10:36 PM	
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_02.29.00.3515	May/21/2025 2:36 AM	
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_06.29.00.3960	May/21/2025 6:36 AM	
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_10.29.00.3581	May/21/2025 10:36 AM	
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_14.29.00.3696	May/21/2025 2:36 PM	
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_18.29.00.3956	May/21/2025 6:36 PM	
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_22.29.03.2699	May/21/2025 10:36 PM	
SnapCenter_hana-8_LocalSnap_Hourly_05-22-2025_02.29.00.3541	May/22/2025 2:36 AM	
SnapCenter_hana-8_LocalSnap_Hourly_05-22-2025_06.29.00.3706	May/22/2025 6:36 AM	

## Restore and recovery operations

With virtual resources stored on VMFS/VMDK's SnapCenter restore operations are always done by a clone, mount, copy operation.

1. SnapCenter creates a storage clone based on the selected Snapshot
2. SnapCenter mounts the LUN as a new datastore to the ESX host
3. SnapCenter adds the VMDK within the datastore as a new disk to the HANA VM
4. SnapCenter mounts the new disk to the Linux OS
5. SnapCenter copies the data from the new disk back to the original location

6. When the copy operation is finished all above resource are removed again
7. SnapCenter executes recovery of the HANA system database
8. SnapCenter executes recovery of the HANA tenant database

The overall runtime of the restore operation is dependent on the database size and the throughput of the FC connection between the storage clusters and the ESX hosts. In our lab setup with an initial HANA installation the runtime has been around 12 minutes.

Restore from SnapCenter\_hana-8\_LocalSnap\_Hourly\_05-22-2025\_06.29.00.3706
✕

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

**Select the restore types**

Complete Resource i

Tenant Database

Restore from SnapCenter\_hana-8\_LocalSnap\_Hourly\_05-22-2025\_06.29.00.3706
✕

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

**Recover database files using**

Recover to most recent state i

Recover to point in time i

Recover to specified data backup i

No recovery i

**Specify log backup locations** i

[Add](#)

/usr/sap/VFS/HDB45/backup/log

**Specify backup catalog location** i

/usr/sap/VFS/HDB45/backup/log

While the restore and recovery operation is running, you can see a new cloned storage unit.

The screenshot shows the NetApp ONTAP System Manager interface. The 'Storage' section displays 20 storage units, 68.6 TiB available, and 20 online units. A table below lists storage configurations with a red box highlighting a new LUN based on a cloned storage unit.

Name	Consistency group	Capacity	Data reduction	Host mapping	IOPS	Latency (ms)	Throughput (MB/s)
hana_data_VFS	sc20250520_110422_689	100 GiB	1.01 to 1	otv_host-44_e3d7e9d4-46f3-4f6a	0	0	0
hana_data_VFS_Clone_0522250947396031	-	100 GiB	1 to 1	otv_host-57_e3d7e9d4-46f3-4f6a	-	-	-
hana_log_VFS	-	100 GiB	1.19 to 1	otv_host-44_e3d7e9d4-46f3-4f6a	0	0	0
hana_shared_VFS	-	100 GiB	2.33 to 1	otv_host-44_e3d7e9d4-46f3-4f6a	0	0	0

The new LUN (datastore) based on the cloned storage unit gets attached to the ESX cluster.

The screenshot shows the vSphere Client interface. The left pane displays a tree view of storage resources, with 'hana\_data\_VFS(sc-20250522094807386)' selected. The main pane shows the contents of this datastore, including folders 'sdd.sf' and 'hana-8'. Below the datastore view, a 'Recent Tasks' table shows the completion of several tasks related to the datastore configuration.

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Reconfigure virtual machine	hana-8	Completed		SAPCC.VCENTER\Administrat or	7 ms	05/22/2025, 9:48:25 AM	05/22/2025, 9:48:26 AM	ycenter8.sapcc.stf.netapp.com
Rename datastore	snip-5781colf2-han a_data_VFS	Completed		SAPCC.VCENTER\Administrat or	5 ms	05/22/2025, 9:48:15 AM	05/22/2025, 9:48:21 AM	ycenter8.sapcc.stf.netapp.com
Resignature unprovisioned		Completed		SAPCC.VCENTER\Administrat or	4 ms	05/22/2025, 9:48:05 AM	05/22/2025, 9:48:06 AM	ycenter8.sapcc.stf.netapp.com

The VMDK within the datastore gets mapped to the target HANA VM and mounted to the HANA system.

```
hana-8:~ # df -h
```

```
Filesystem Size Used Avail Use% Mounted on
/dev/mapper/system-root 60G 5.3G 54G 9% /
devtmpfs 4.0M 8.0K 4.0M 1% /dev
tmpfs 49G 0 49G 0% /dev/shm
efivarfs 256K 57K 195K 23% /sys/firmware/efi/efivars
tmpfs 13G 26M 13G 1% /run
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup-dev-
early.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-sysctl.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-sysusers.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup-dev.service
/dev/mapper/system-root 60G 5.3G 54G 9% /.snapshots
/dev/mapper/system-root 60G 5.3G 54G 9% /boot/grub2/i386-pc
/dev/mapper/system-root 60G 5.3G 54G 9% /boot/grub2/x86++_++64-efi
/dev/mapper/system-root 60G 5.3G 54G 9% /home
/dev/mapper/system-root 60G 5.3G 54G 9% /opt
/dev/mapper/system-root 60G 5.3G 54G 9% /root
/dev/mapper/system-root 60G 5.3G 54G 9% /srv
/dev/mapper/system-root 60G 5.3G 54G 9% /usr/local
/dev/mapper/system-root 60G 5.3G 54G 9% /tmp
/dev/mapper/system-root 60G 5.3G 54G 9% /var
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-vconsole-setup.service
/dev/sdc 95G 8.9G 87G 10% /hana/log/VFS/mnt00001
/dev/sdb 95G 7.6G 88G 8% /hana/data/VFS/mnt00001
/dev/sdd 95G 15G 81G 16% /hana/shared
/dev/sda1 253M 5.9M 247M 3% /boot/efi
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup.service
192.168.175.86:/sapcc_share 1.4T 858G 568G 61% /mnt/sapcc-share
tmpfs 6.3G 72K 6.3G 1% /run/user/464
tmpfs 1.0M 0 1.0M 0% /run/credentials/getty@tty1.service
tmpfs 6.3G 52K 6.3G 1% /run/user/0
/dev/sde 95G 9.2G 86G 10%
/var/opt/snapcenter/scu/clones/hana_data_VFS_mnt00001_142592_scu_clone_1
```

```
hana-8:~ #
```

## Job Details



Restore 'hana-8.sapcc.stl.netapp.com\hana\MDC\VFS'

- ✓ ▾ Restore 'hana-8.sapcc.stl.netapp.com\hana\MDC\VFS'
- ✓ ▾ hana-8.sapcc.stl.netapp.com
  - ✓ ▾ Restore
    - ✓ ▶ Validate Plugin Parameters
    - ✓ ▾ Pre Restore Application
      - ✓ ▾ Stopping HANA Instance
    - ✓ ▾ Filesystem Pre Restore
    - ✓ ▾ PreRestore for Virtual Resources
    - ✓ ▾ Detach Virtual Disks
    - ✓ ▶ Restore Filesystem
    - ✓ ▶ Restore for Virtual Resources
    - ✓ ▶ Attach Virtual Disks
    - ✓ ▶ Filesystem Post Restore
    - ✓ ▶ Recover Application
    - ✓ ▶ PostRestore for Virtual Resources
    - ✓ ▶ Cleaning Storage Resources
    - ✓ ▶ Post Restore Cleanup FileSystem
    - ✓ ▶ Application Clean-Up
    - ✓ ▶ Data Collection
    - ✓ ▶ Agent Finalize Workflow
  - ✓ ▶ ( Job 142596 ) ( Job 142596 ) read UnmountBackup

**i** Task Name: Recover Application Start Time: 05/22/2025 9:56:13 AM End Time: 05/22/2025 9:58:15 AM

View Logs

Cancel Job

Close

## SAP System Refresh

Detailed information on SAP System Refresh operations using SnapCenter can be found at [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#).

The second HANA system QFS has been provisioned in the same way as described in chapter [“HANA system provisioning and installation”](#).

### Prerequisites

The current version of SnapCenter (6.1P1) has some limitations which are planned to get fixed with next releases.

1. It is required to restart the spl process after each “clone create” and “clone delete” workflows using the command “systemctl restart spl” at the target HANA host.
2. The HANA VMs used as source and target of the SAP system refresh operation must run on the same ESX host.

### Workflow summary

Before the first SAP System Refresh operation can be executed, the target HANA system must be installed, and the host must be added to SnapCenter. Then the HANA system must be shut down and the HANA data disk must be unmounted from host.

### SnapCenter clone create workflow

1. Create storage clone
2. Configure host mapping for storage clone
3. Attach storage clone (datastore) to ESX host
4. Add new disk from datastore to target HANA VM
5. Mount disk to HANA VM OS
6. Recover HANA system using post-script

Runtime: 12 minutes



Compared to the restore operation, the runtime of the clone operation is independent from the size of the HANA database. The runtime of step 1 – 5 will be similar also for very large databases. Recovery will of course take longer for larger HANA systems.

### SnapCenter clone delete workflow

1. Shutdown HANA system using pre-script
2. Unmount disk from HANA VM OS
3. Remove disk from HANA VM
4. Remove datastore from ESX host
5. Delete storage clone

Runtime: 11 minutes

### SnapCenter clone create workflow

The clone create workflow is started by selecting the desired Snapshot and by clicking on the clone button.

The screenshot shows the NetApp SnapCenter interface. On the left, a sidebar lists systems: QFS, OS1, SM1, SS1, SS2, SS2, and VFS. The main area is titled 'VFS Topology' and contains a 'Manage Copies' section with '12 Backups' and '0 Clones'. Below this is a 'Primary Backup(s)' table with columns for Backup Name, Snapshot Lock Expiration, Count, and End Date. The table lists 12 backup entries with their respective IDs and dates. A 'Summary Card' on the right shows: 13 Backups, 12 Snapshot based backups, 1 File-based backup, 0 Clones, and 0 Snapshots Locked. At the bottom, an activity bar shows 3 Completed, 0 Warnings, 0 Failed, 0 Canceled, 2 Running, and 0 Queued jobs.

The target host and SID must be provided.

This screenshot shows the 'Clone From Backup' dialog box, specifically the 'Location' step. The dialog has a sidebar with steps: 1 Location, 2 Settings, 3 Scripts, 4 Notification, and 5 Summary. The main content area is titled 'Select the host to create the clone'. It contains two fields: 'Plug-in host' with a dropdown menu showing 'hana-9.sapcc.stl.netapp.com' and 'Target Clone SID' with a text input field containing 'QFS'. Information icons are visible next to both fields.

This screenshot shows the 'Clone From Backup' dialog box, specifically the 'Settings' step. The sidebar shows '2 Settings' is selected. The main content area is titled 'LUN Map Settings' and contains a dropdown menu for 'igroup protocol' with 'FCP' selected. A list of options is shown below the dropdown: Select, Mixed, FCP (highlighted), and iSCSI.

In our example we are using a post-script to execute the recovery at the target host.

# Clone From Backup



1 Location

The following commands will run on the Plug-in Host: **hana-9.sapcc.stl.netapp.com**

2 Settings

Enter optional commands to run before performing a clone operation **i**

3 Scripts

Pre clone command

4 Notification

Enter optional commands to run after performing a clone operation **i**

5 Summary

Post clone command  
`/mnt/sapcc-share/SAP-System-Refresh/sc-system-refresh.sh  
recover`

When the workflow is started SnapCenter creates a cloned storage unit based on the selected Snapshot.

The screenshot shows the NetApp ONTAP System Manager interface. The 'Storage' section displays 22 storage units, 68.5 TiB available, and 22 online units. A table lists storage units with columns for Name, Consistency group, Capacity, Data reduction, Host mapping, IOPS, Latency (ms), and Throughput (MB/s).

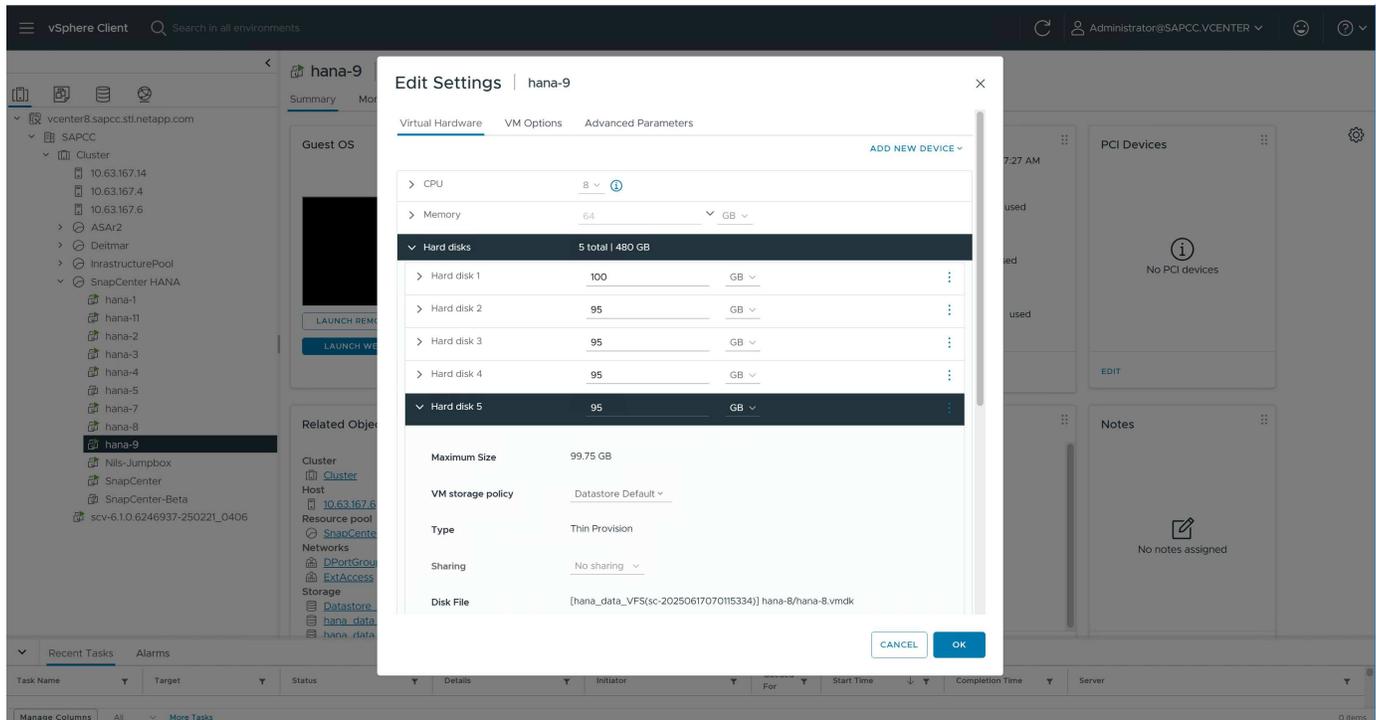
Name	Consistency group	Capacity	Data reduction	Host mapping	IOPS	Latency (ms)	Throughput (MB/s)
hana_data_QFS	-	100 GiB	5.46 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	4	0.11	0.39
hana_data_VFS	vc20250520_110422_689	100 GiB	1 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	5	0.12	0.39
hana_data_VFS_Clone_06172507005937511	-	100 GiB	1 to 1	otv_host-57_e3d7e9d4-46f3-4f5d	23	0.11	1.24
hana_log_QFS	-	100 GiB	4.1 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	5	0.10	0.39
hana_log_VFS	-	100 GiB	1.22 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	8	0.12	0.40
hana_shared_QFS	-	100 GiB	2.81 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	5	0.11	0.39
hana_shared_VFS	-	100 GiB	1.69 to 1	otv_host-44_e3d7e9d4-46f3-4f5d	5	0.13	0.39

SnapCenter then attaches the LUN (datastore) to the ESX host, on which the target HANA VM is running.

The screenshot shows the vSphere Client interface displaying the 'Datastores' page for host 10.63.167.6. A table lists datastores with columns for Name, Status, Type, Datastore Cluster, Capacity, and Free space.

Name	Status	Type	Datastore Cluster	Capacity	Free
datastore1[2]	✓ Norm	VMFS 6		766 GB	764.58 GB
Datastore_C250	✓ Norm	NFS 3		1.95 TB	1.95 TB
Datastore_One	✓ Norm	NFS 3		2.85 TB	1.22 TB
Datastore-A40Q	✓ Norm	NFS 3		500 GB	271.24 GB
hana_data_QFS	✓ Norm	VMFS 6		99.75 GB	87.26 GB
hana_data_VFS	✓ Norm	VMFS 6		99.75 GB	90.94 GB
hana_data_VFS/sc-2025061707018334	✓ Norm	VMFS 6		99.75 GB	90.94 GB
hana_log_QFS	✓ Norm	VMFS 6		99.75 GB	91.31 GB
hana_log_VFS	✓ Norm	VMFS 6		99.75 GB	91.3 GB
hana_shared_QFS	✓ Norm	VMFS 6		99.75 GB	87 GB
hana_shared_VFS	✓ Norm	VMFS 6		99.75 GB	80.55 GB
OS_Image	✓ Norm	NFS 3		142.5 GB	55.39 GB

The VMDK within the new datastore is then added to the HANA VM.



SnapCenter then configures and mounts the new disk at the HANA Linux system.

```
hana-9:/mnt/sapcc-share/SAP-System-Refresh # df -h

Filesystem Size Used Avail Use% Mounted on
/dev/mapper/system-root 60G 5.2G 52G 10% /
devtmpfs 4.0M 4.0K 4.0M 1% /dev
tmpfs 49G 0 49G 0% /dev/shm
efivarfs 256K 57K 195K 23% /sys/firmware/efi/efivars
tmpfs 13G 26M 13G 1% /run
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup-dev-early.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-sysctl.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-sysusers.service
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup-dev.service
/dev/mapper/system-root 60G 5.2G 52G 10% /.snapshots
/dev/mapper/system-root 60G 5.2G 52G 10% /boot/grub2/i386-pc
/dev/mapper/system-root 60G 5.2G 52G 10% /boot/grub2/x86++_++64-efi
/dev/mapper/system-root 60G 5.2G 52G 10% /home
/dev/mapper/system-root 60G 5.2G 52G 10% /opt
/dev/mapper/system-root 60G 5.2G 52G 10% /srv
/dev/mapper/system-root 60G 5.2G 52G 10% /root
/dev/mapper/system-root 60G 5.2G 52G 10% /tmp
/dev/mapper/system-root 60G 5.2G 52G 10% /usr/local
/dev/mapper/system-root 60G 5.2G 52G 10% /var
```

```

tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-vconsole-setup.service
/dev/sdc 95G 8.9G 87G 10% /hana/log/QFS/mnt00001
/dev/sdd 95G 14G 82G 14% /hana/shared
/dev/sda1 253M 5.9M 247M 3% /boot/efi
tmpfs 1.0M 0 1.0M 0% /run/credentials/systemd-tmpfiles-setup.service
192.168.175.86:/sapcc+_++share 1.4T 858G 568G 61% /mnt/sapcc-share
tmpfs 6.3G 72K 6.3G 1% /run/user/464
tmpfs 1.0M 0 1.0M 0% /run/credentials/getty@tty1.service
tmpfs 6.3G 52K 6.3G 1% /run/user/0
/dev/sde 95G 9.2G 86G 10% /hana/data/QFS/mnt00001
tmpfs 6.3G 56K 6.3G 1% /run/user/1001
hana-9:/mnt/sapcc-share/SAP-System-Refresh #

hana-9:/mnt/sapcc-share/SAP-System-Refresh # cat /etc/fstab
/dev/system/root / btrfs defaults 0 0
/dev/system/root /var btrfs subvol=@/var 0 0
/dev/system/root /usr/local btrfs subvol=@/usr/local 0 0
/dev/system/root /tmp btrfs subvol=@/tmp 0 0
/dev/system/root /srv btrfs subvol=@/srv 0 0
/dev/system/root /root btrfs subvol=@/root 0 0
/dev/system/root /opt btrfs subvol=@/opt 0 0
/dev/system/root /home btrfs subvol=@/home 0 0
/dev/system/root /boot/grub2/x86+_++64-efi btrfs
subvol=@/boot/grub2/x86+_++64-efi 0 0
/dev/system/root /boot/grub2/i386-pc btrfs subvol=@/boot/grub2/i386-pc 0
0
/dev/system/swap swap swap defaults 0 0
/dev/system/root /.snapshots btrfs subvol=@/.snapshots 0 0
UUID=FB79-24DC /boot/efi vfat utf8 0 2
192.168.175.86:/sapcc+_++share /mnt/sapcc-share nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0
0
#/dev/sdb /hana/data/QFS/mnt00001 xfs relatime,inode64 0 0
/dev/sdc /hana/log/QFS/mnt00001 xfs relatime,inode64 0 0
/dev/sdd /hana/shared xfs defaults 0 0
# The following entry has been added by NetApp (SnapCenter Plug-in for
UNIX)
/dev/sde /hana/data/QFS/mnt00001 xfs
rw,relatime,attr2,inode64,logbufs=8,logbsize=32k,noquota 0 0
hana-9:/mnt/sapcc-share/SAP-System-Refresh #

```

The following screenshot shows the job steps executed by SnapCenter.

Job Details
✕

Clone from backup 'SnapCenter\_hana-8\_LocalSnap\_Hourly\_06-17-2025\_10.29.00.4260'

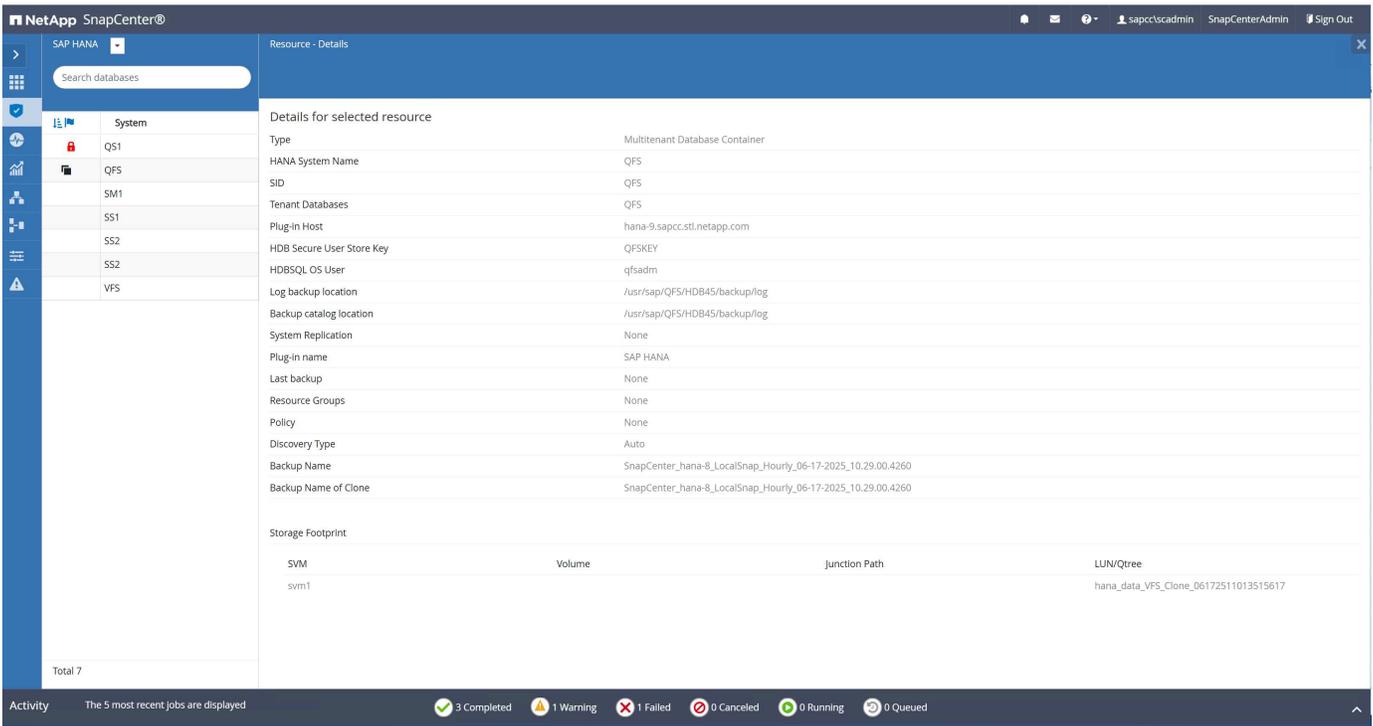
- ✓ ▾ Clone from backup 'SnapCenter\_hana-8\_LocalSnap\_Hourly\_06-17-2025\_10.29.00.4260'
- ✓ ▾ hana-9.sapcc.stl.netapp.com
  - ✓ ▾ Clone
    - ✓ ▶ Application Pre Clone
    - ✓ ▶ Storage Clone
    - ✓ ▶ Can Execute Clone Virtual or RDM disks
    - ✓ ▶ Clone Virtual or RDM disks
    - ✓ ▶ Unmount Filesystem
    - ✓ ▾ Mount Filesystem
      - ✓ ▶ Performing rescan of devices
      - ✓ ▶ Building clone for data file systems and associated entities
    - ✓ ▾ Application Post Clone
    - ✓ ▾ Register Clone Metadata
    - ✓ ▾ Clean-up Snapshot entries on Server
    - ✓ ▾ Application Clean-Up
      - ✓ ▶ Data Collection
      - ✓ ▶ Agent Finalize Workflow

**i** Task Name: Mount Filesystem Start Time: 06/17/2025 11:02:42 AM End Time: 06/17/2025 11:10:17 AM

View Logs
Cancel Job
Close

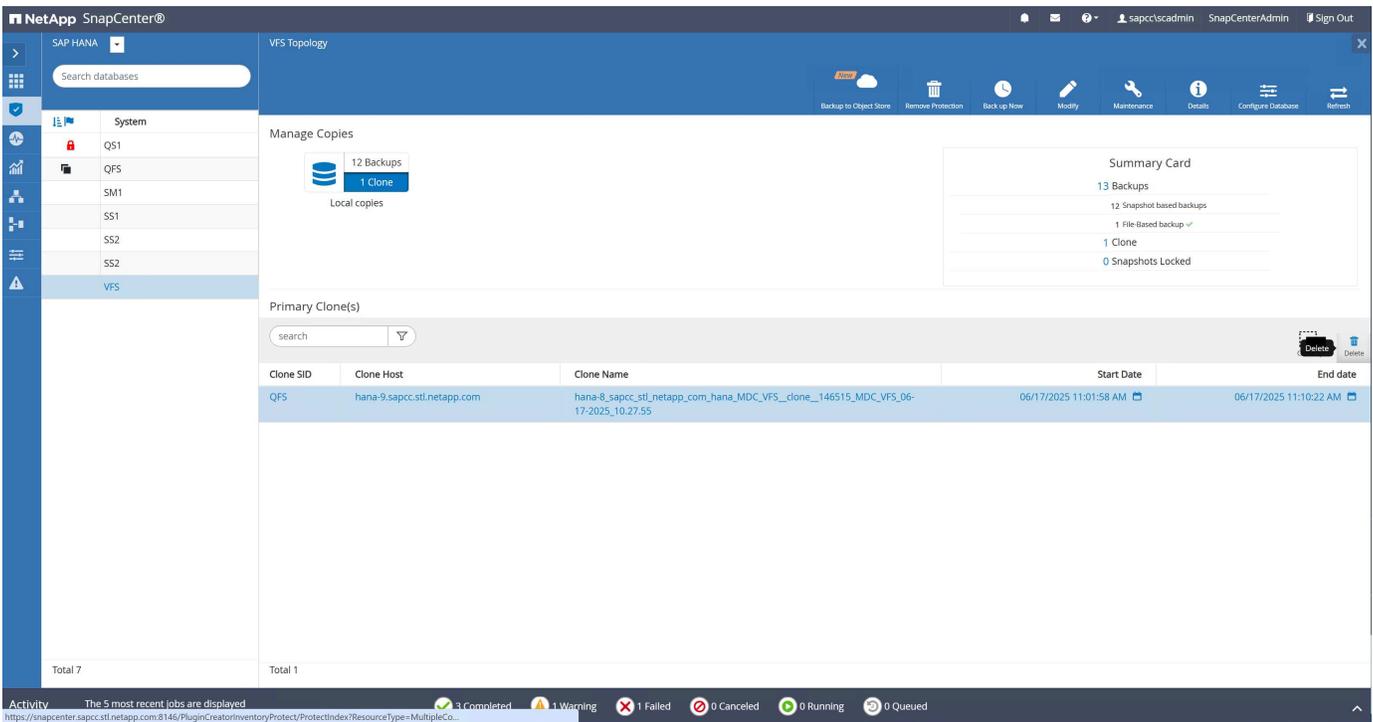
As mentioned in the “Pre-requisites” section, the SnapCenter spl service at the HANA host must be restarted using the command “systemctl restart spl” to initiate proper cleanup. This must be done when the job has finished.

When the clone workflow is finished, the auto discovery can be started by clicking on the resource QFS. When the auto discovery process is finished the new storage footprint is listed in the details view of the resource.



## SnapCenter clone delete workflow

The clone delete workflow is started by selecting the clone at the source HANA resource and by clicking on the delete button.



In our example we are using a pre-script to shutdown the target HANA database.

## Delete Clone



**i** 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

The following screenshot shows the job steps executed by SnapCenter.

Job Details
✕

Deleting clone 'hana-8\_sapcc\_stl\_netapp\_com\_h.....S\_clone\_\_146534\_MDC\_VFS\_06-17-2025\_10.27.55'

- ✓ ▼ Deleting clone 'hana-8\_sapcc\_stl\_netapp\_com\_hana\_MDC\_VFS\_clone\_\_146534\_MDC\_VFS\_06-17-2025\_10.27.55'
- ✓ ▼ hana-9.sapcc.stl.netapp.com
  - ✓ ▼ Delete Clone
    - ▶ Validate Plugin Parameters
    - ▶ Application Clone Delete
    - ▶ Delete Pre Clone Commands
    - ▼ Unmount Filesystem
      - ▶ Deporting cloned file systems and associated entities
      - ▶ Performing rescan of devices
    - ▶ Deleting Virtual Resources
    - ▼ Delete Storage Clone
    - ▼ Unregister Clone Metadata
    - ▼ Filesystem Clone Metadata Cleanup
      - ▶ Performing rescan of devices
    - ▶ Agent Finalize Workflow

**i** Task Name: Application Clone Delete Start Time: 06/17/2025 1:36:24 PM End Time: 06/17/2025 1:37:02 PM

View Logs
Cancel Job
Close

As mentioned in the “Pre-requisites” section, the SnapCenter spl service at the HANA host must be restarted using the command “systemctl restart spl” to initiate proper cleanup.

## Additional information and version history

HANA best practices:

- [SAP HANA on NetApp ASA Systems with Fibre Channel Protocol.](#)

SnapCenter:

- [TR-4614: SAP HANA backup and recovery with SnapCenter](#)
- [TR-4719: SAP HANA System Replication - Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- [SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster](#)
- [SnapCenter Software documentation](#)

Version history:

Version	Date	Comment
Version 1.0	07/2025	Initial version

## BlueXP Backup and Recovery for SAP HANA - Cloud Object storage as backup destination

### BlueXP Backup and Recovery for SAP HANA - Cloud Object storage as backup destination

This technical report provides best practices for SAP HANA data protection using NetApp BlueXP Backup and Recovery for Application. This document covers concepts, configuration recommendations, and operation workflows, including configuration, backup operations, and restore operations.

#### Overview

This document describes how to setup and configure SAP HANA for data protection from on-premises to cloud based object stores with NetApp BlueXP. It covers the BlueXP backup and recovery part of the solution. This solution is an enhancement of the on-premises SAP HANA backup solution using NetApp Snap Center and provides a cost-efficient way for long-term archiving of SAP HANA backups to cloud based object storage and offers optional tiering of object storage to archival storage like AWS Glacier/Deep Glacier, Microsoft Azure Blob Archive, and GCP Archive Storage.

The setup and configuration of the on-premises SAP HANA backup and recovery solution is described in [TR-4614: SAP HANA backup and recovery with SnapCenter \(netapp.com\)](#).

This TR only describes how to enhance the on-premises SnapCenter based SAP HANA backup and recovery solution with BlueXP backup and recovery for SAP HANA using AWS S3 object storage as example. The setup and configuration using Microsoft Azure and GCP object storage instead of AWS S3 is similar, but is not described within this document.

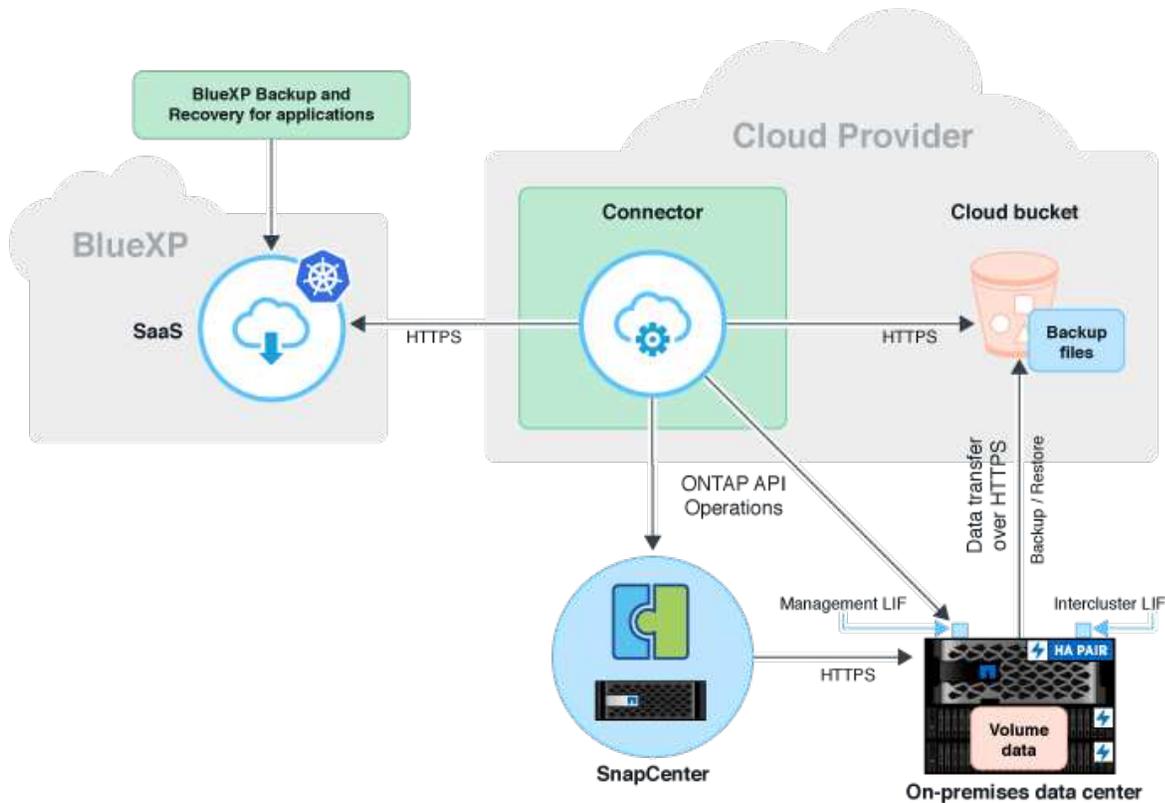
#### BlueXP Backup and Recovery architecture

BlueXP backup and recovery is a SaaS solution that provides data protection capabilities for applications running on NetApp on-premises Storage to the cloud. It offers efficient, application consistent, policy-based protection of SAP HANA using NetApp storage. In addition, BlueXP backup and recovery provides centralized control and oversight, while delegating the ability for users to manage application-specific backup and restore operations.

BlueXP backup and recovery runs as SaaS within NetApp BlueXP and leverages the framework and UI. The BlueXP working environment framework is used to configure and manage the credentials for NetApp ONTAP based on-premises storage and the NetApp SnapCenter Server.

A BlueXP connector needs to be deployed within the customer virtual network. A connection between the on-premises environment and the cloud environment is required such as a site to site VPN connection. The communication between the NetApp SaaS components and the customer environment is exclusively done via the connector. The connector is executing the storage operations by using the ONTAP and SnapCenter management APIs.

The data transfer between the on-premises storage and the cloud bucket is end-to-end protected with AES 256-bit encryption at rest, TLS/HTTPS encryption in flight, and customer-managed key (CMK) support. The backed-up data can be stored in an immutable and indelible WORM state. The only way to access the data from the object storage is to restore it to NetApp ONTAP based storage including NetApp CVO.



## Overview of installation and configuration steps

The required installation and configuration steps can be split in three areas.

Prerequisite is that the SAP HANA backup configuration has been configured at NetApp Snap Center. For setting up Snap Center for SAP HANA in the first place refer to [SnapCenter configuration \(netapp.com\)](https://www.netapp.com/learn/topics/snapcenter-configuration).

1. Installation and configuration of NetApp BlueXP components.

Needs to be done once during the initial setup of the data protection solution.

2. Preparation steps at NetApp SnapCenter.

Needs to be done for each SAP HANA database, which should be protected.

### 3. Configuration steps in BlueXP backup and recovery.

Needs to be done for each SAP HANA database, which should be protected.

## Installation and configuration of NetApp BlueXP Hybrid Application Backup

The installation and configuration of the NetApp BlueXP components are described in [Protect your on-premises applications data | NetApp Documentation](#).

1. Sign-up to BlueXP and setup NetApp account at <https://bluexp.netapp.com/>.
2. Deploy BlueXP connector in your environment. Description is available at [Learn about Connectors | NetApp Documentation](#).
3. Add/buy a Cloud Backup license at BlueXP: <https://docs.netapp.com/us-en/cloud-manager-backup-restore/task-licensing-cloud-backup.html>.
4. Create working environment for NetApp on-premises environment and your cloud destination in BlueXP by adding your on-premises storage.
5. Create a new object store relationship for the on-premises storage into an AWS S3 bucket.
6. Configure SAP HANA system resource at SnapCenter.
7. Add Snap Center to your working environment.
8. Create a policy for your environment.
9. Protect you SAP HANA System.

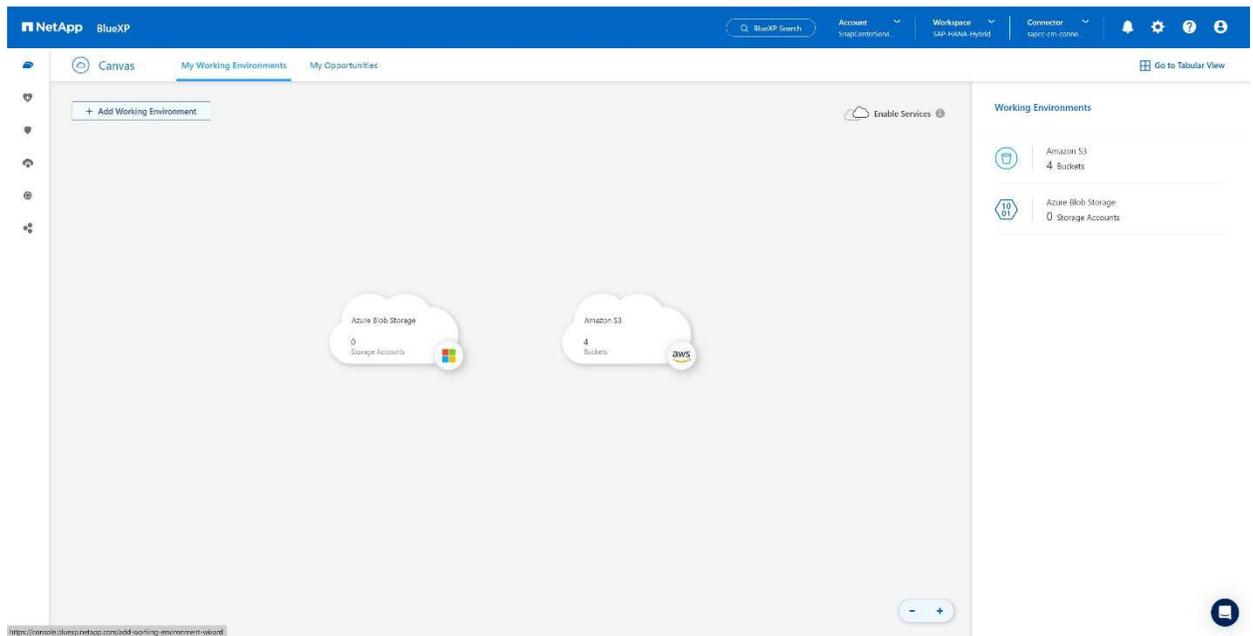
## Configuring BlueXP Backup and Recovery for SAP HANA

This section describes how to setup the working environment, how to configure SnapCenter, and how to configure and activate SAP HANA backup within BlueXP.

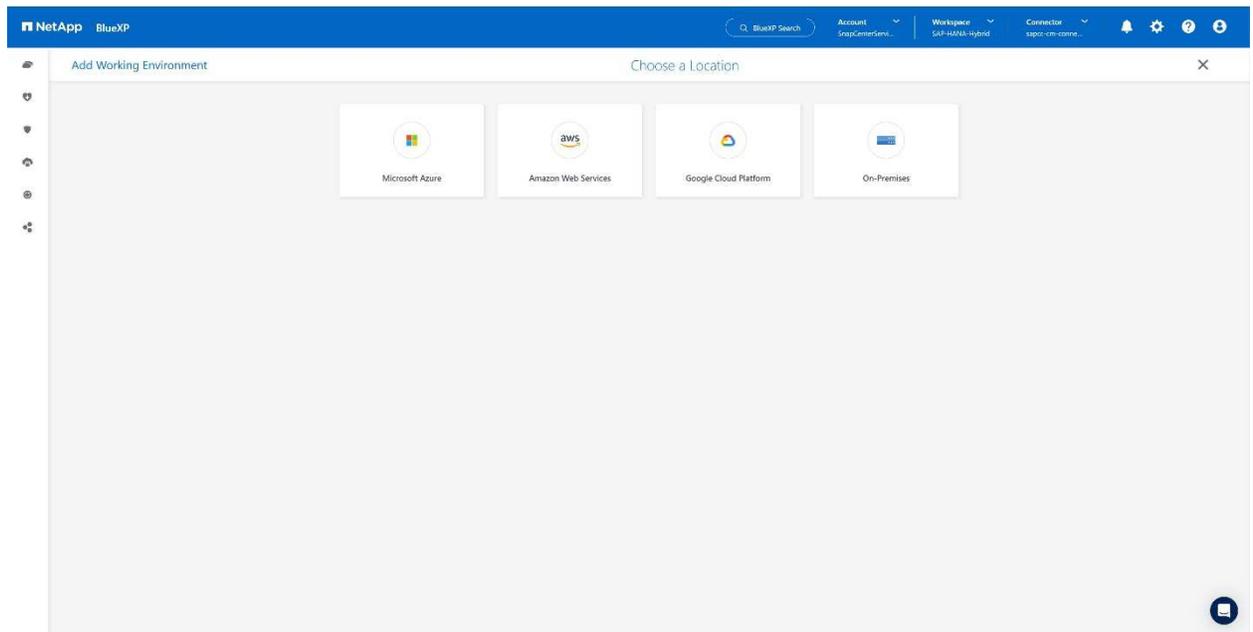
### Create working environment for BlueXP

Add the on-premises storage system to you work environment.

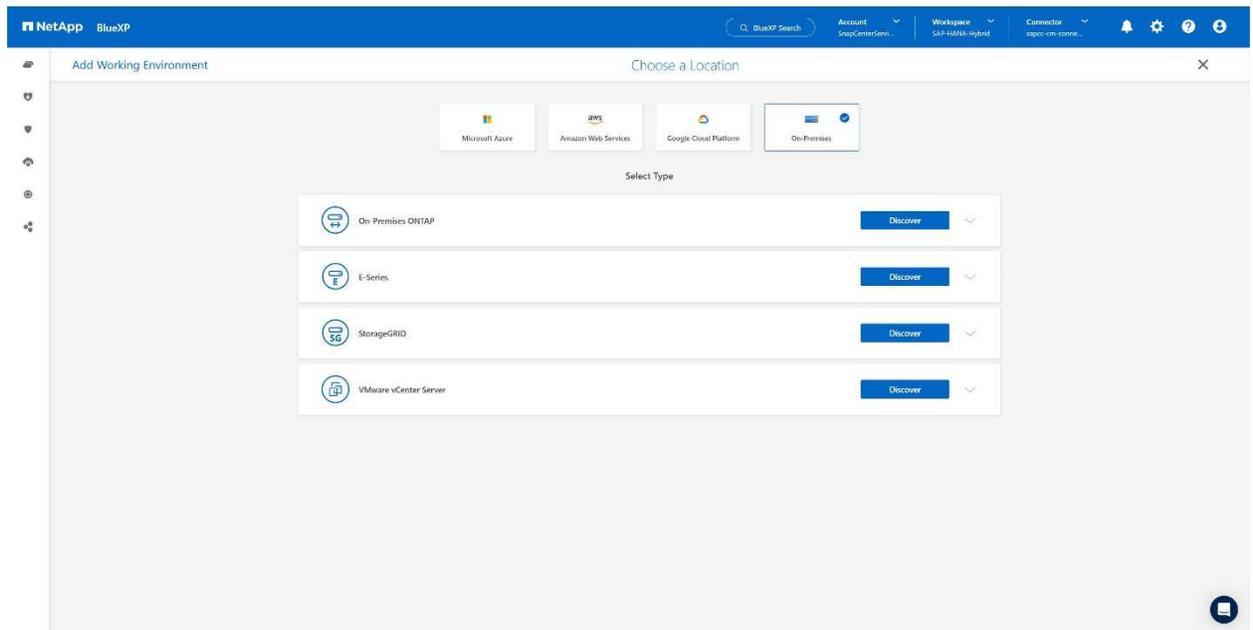
1. At the left menu choose **Storage** → **Canvas** → **My Working** Environment.
2. Press **+ Add Working Environment**.



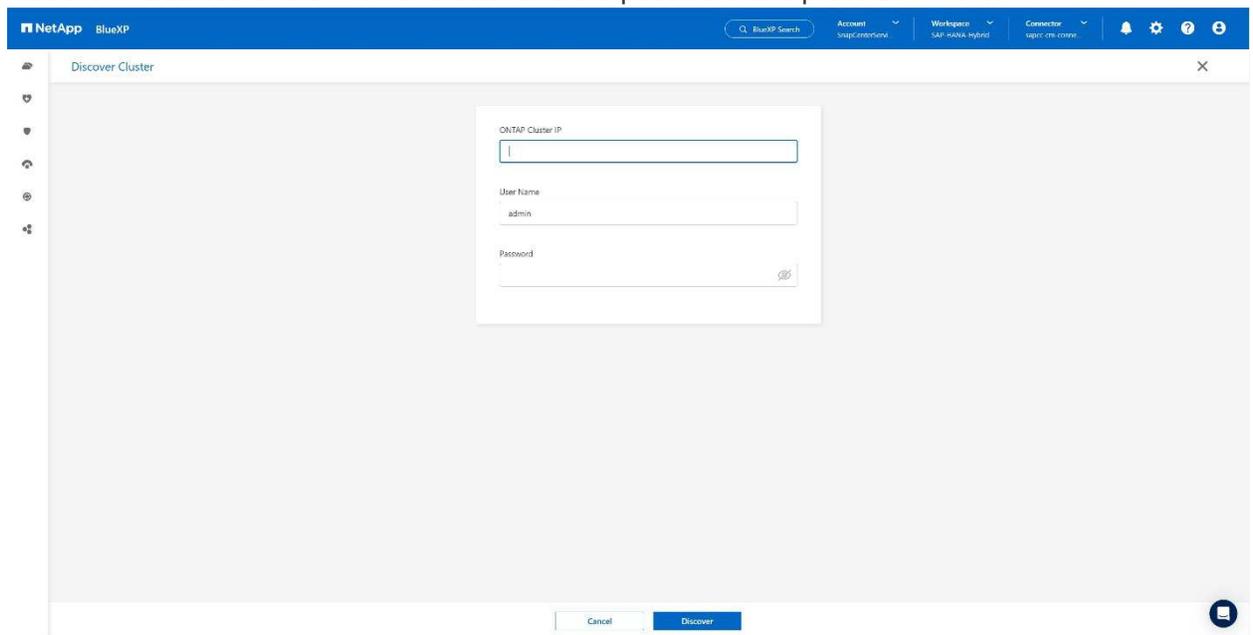
3. Choose **On-Premises**.



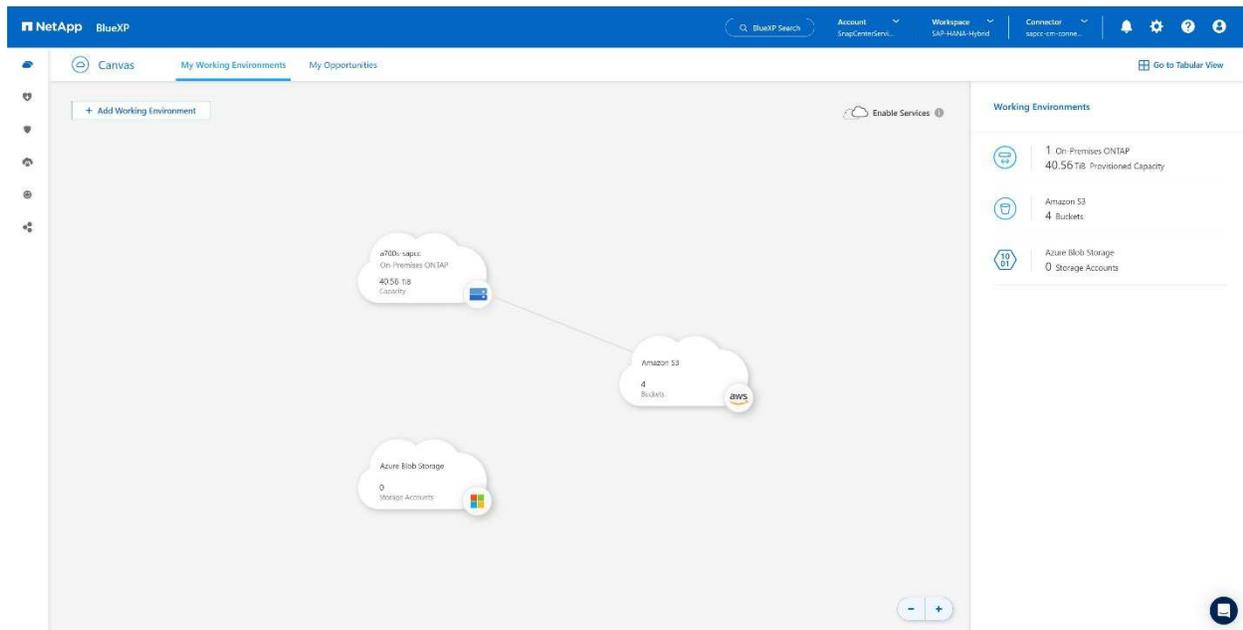
4. Choose **Discover On-Premises ONTAP**.



5. Add the IP address of the ONTAP cluster and the password and press **Discover**.



6. The ONTAP cluster is now available.



## Create a relationship between the on-premises storage system and an object storage bucket

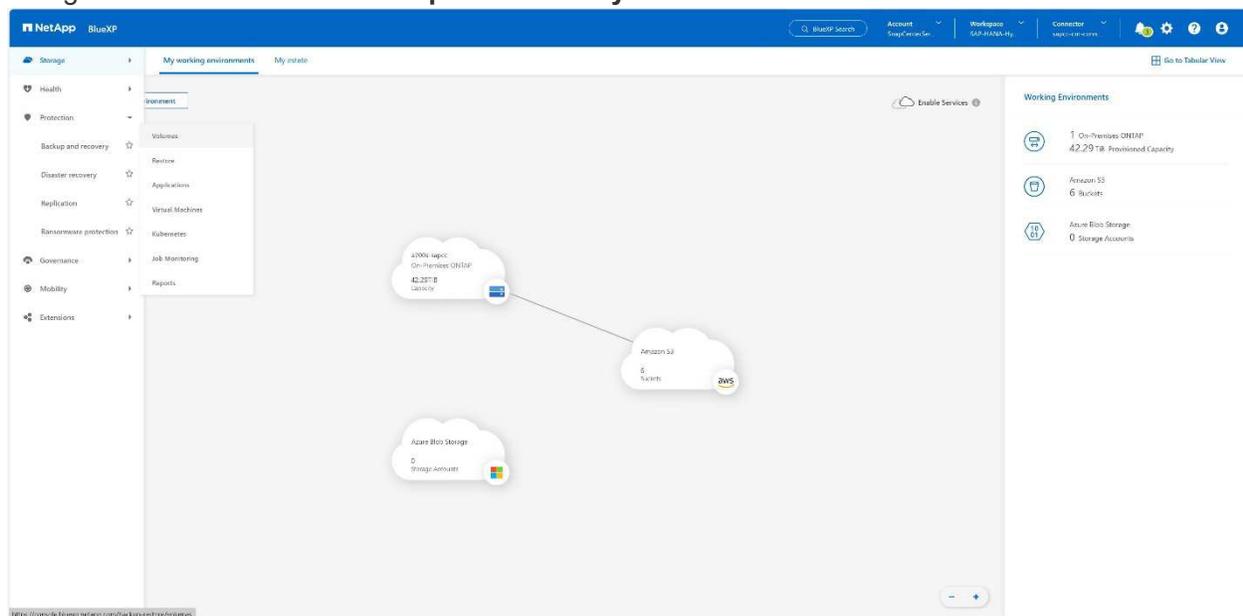
The relationship between the on-premises storage and the S3 bucket is done by creating a backup for a volume or by activating a backup of an application. If an existing site-to-site VPN shall be used for transferring the data from on-premises to S3, a volume backup needs to be used for creating the relationship between the on-premises storage and S3 bucket as VPC endpoints need to be used.

At creation of this documentation the application backup workflow doesn't offer to choose VPC endpoints to access S3 buckets.

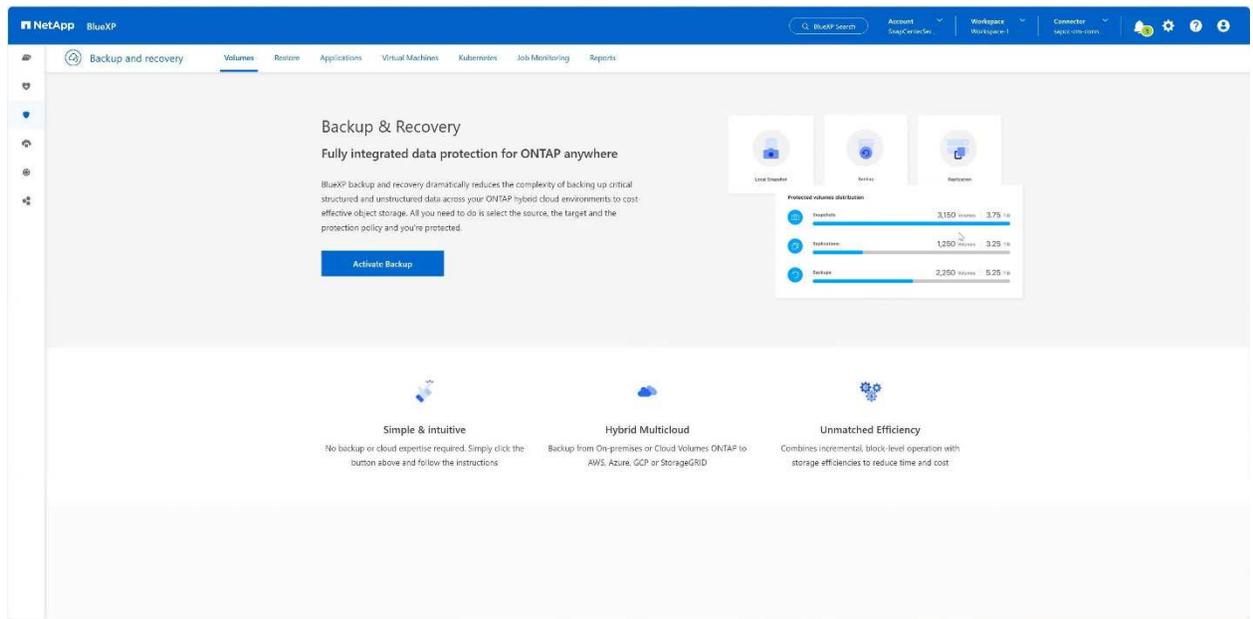
Refer to [Gateway endpoints for Amazon S3 - Amazon Virtual Private Cloud](#) how to setup VPC endpoints for S3 within your VPC.

To create a first volume backup, perform the following steps:

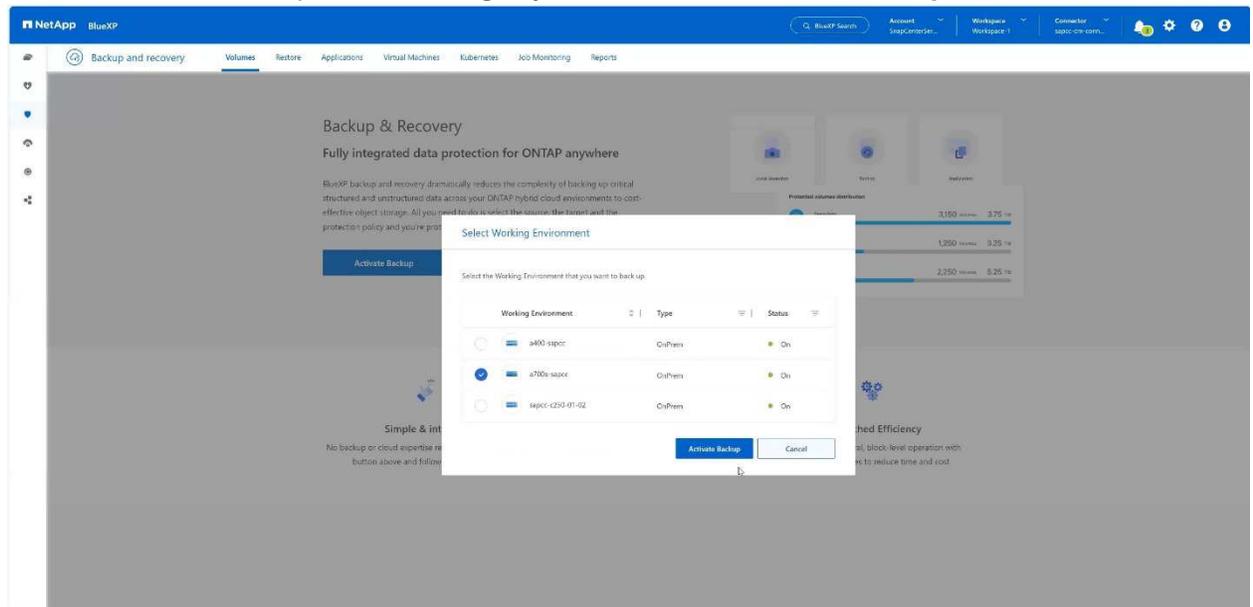
1. Navigate via **Protection** to **Backup and recovery** and choose **Volumes**.



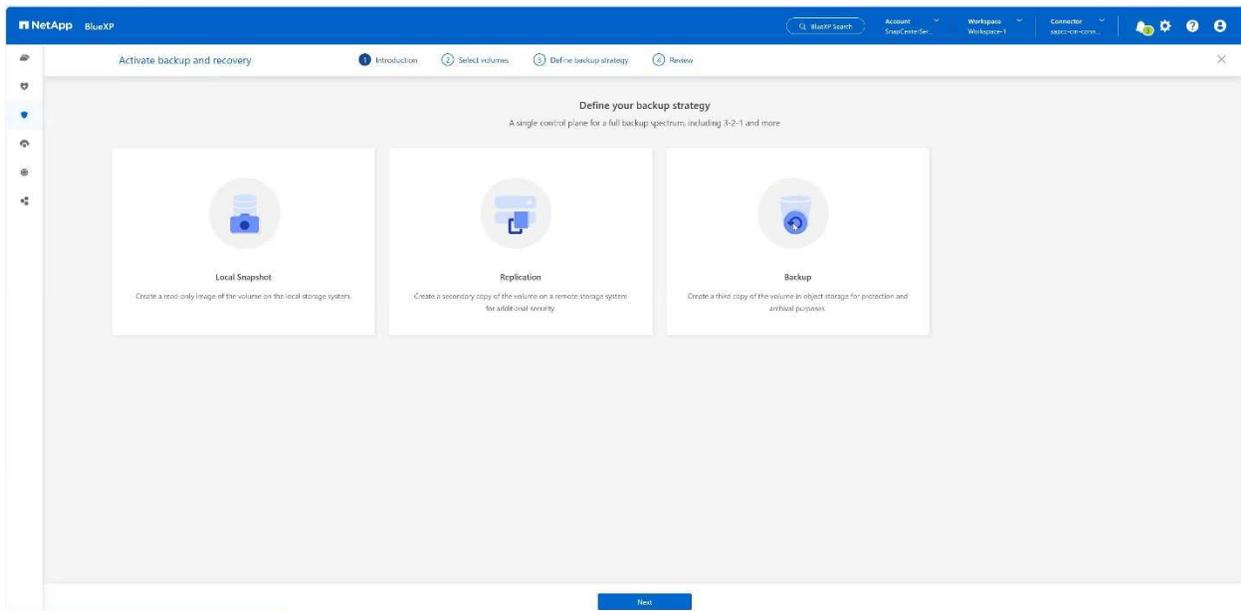
2. Press the **Activate Backup** button.



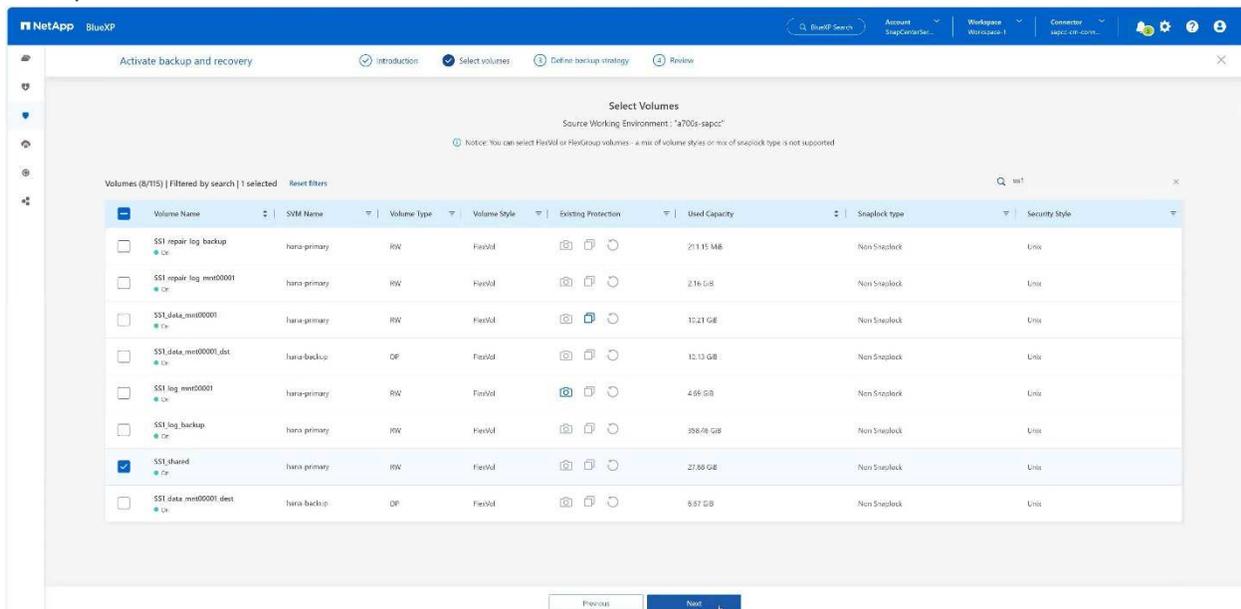
3. Choose the desired on-premises storage system and click **Activate Backup**.



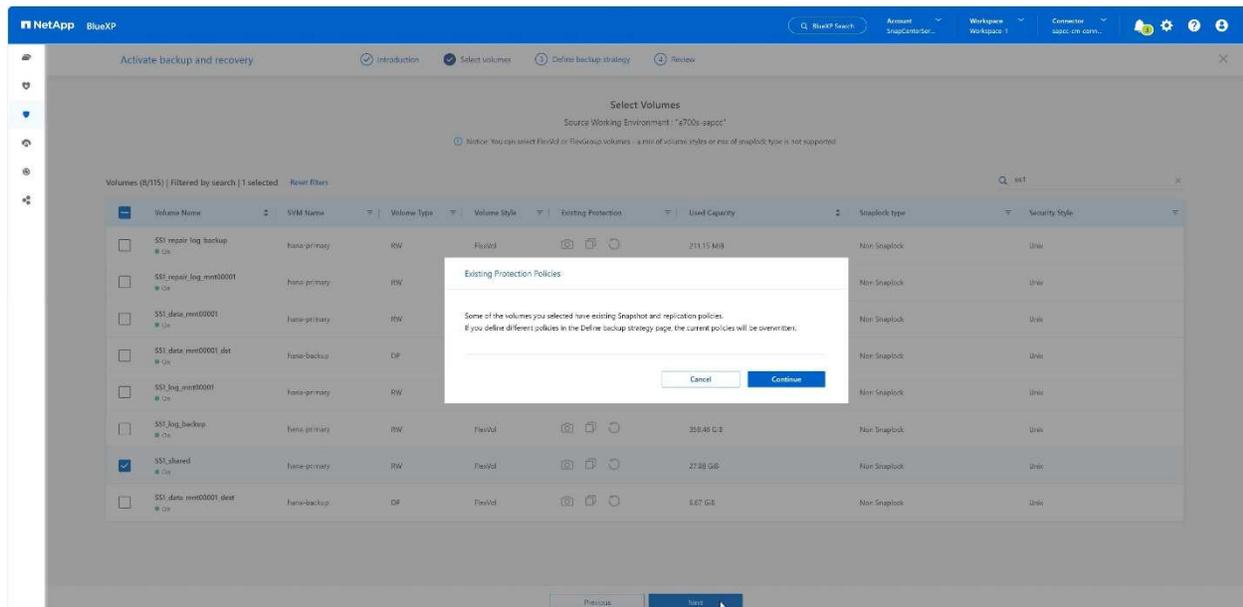
4. Choose **Backup**.



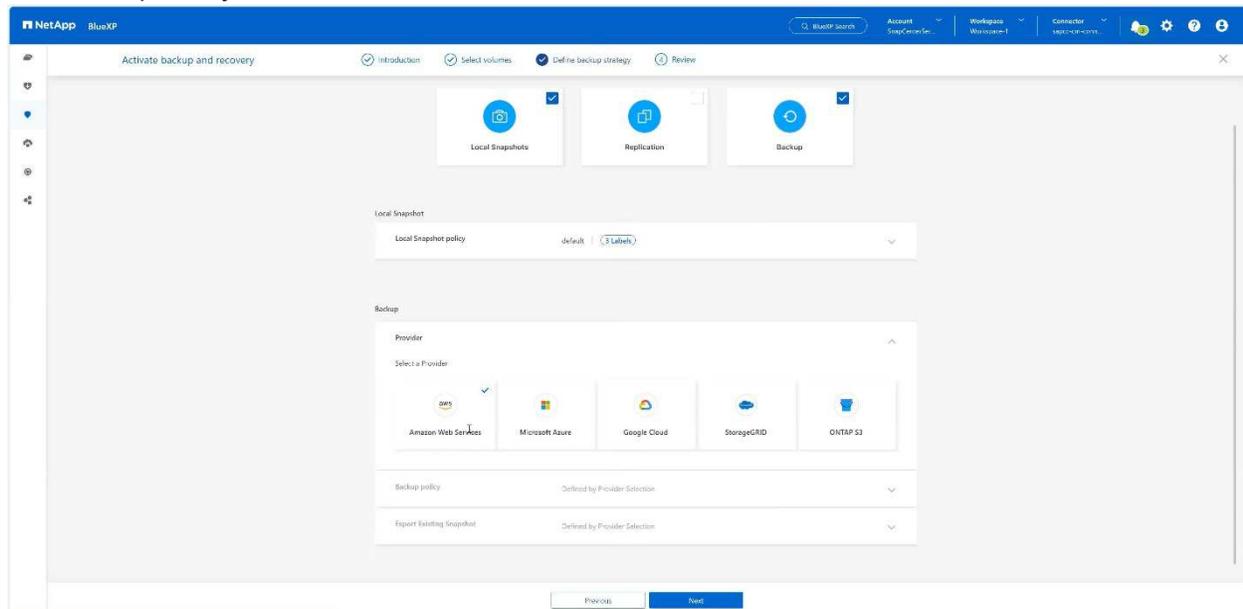
5. Choose a volume which is stored at the same SVM as your SAP HANA data files and press **Next**. In this example the volume for /hana/shared has been chosen.



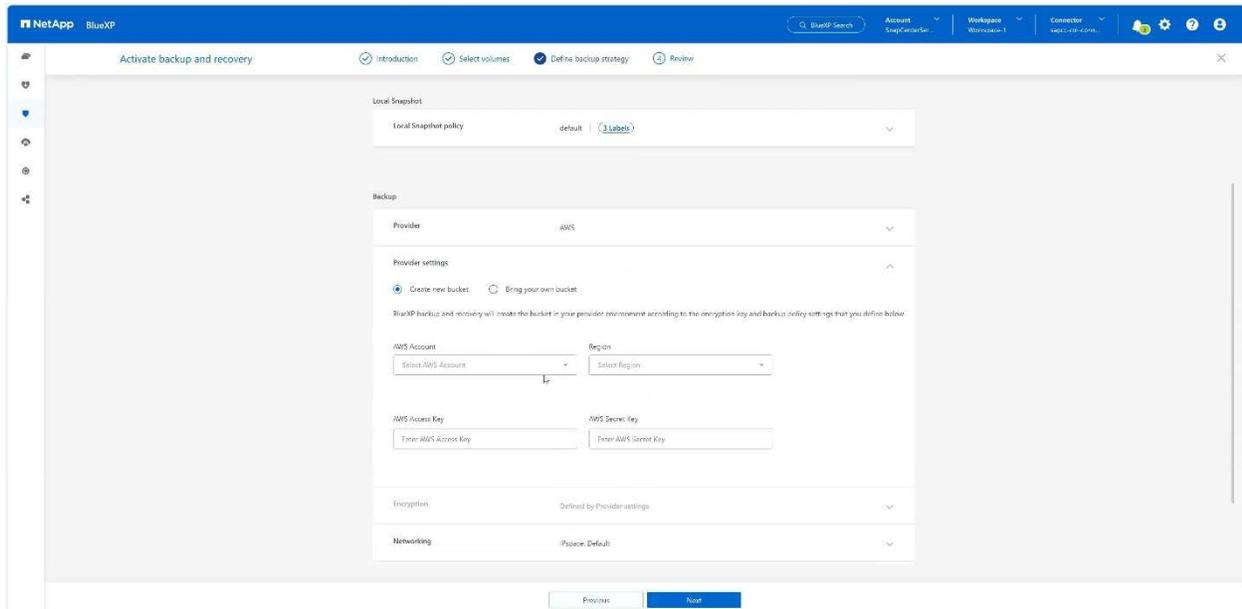
6. **Continue**, if an existing policy exists.



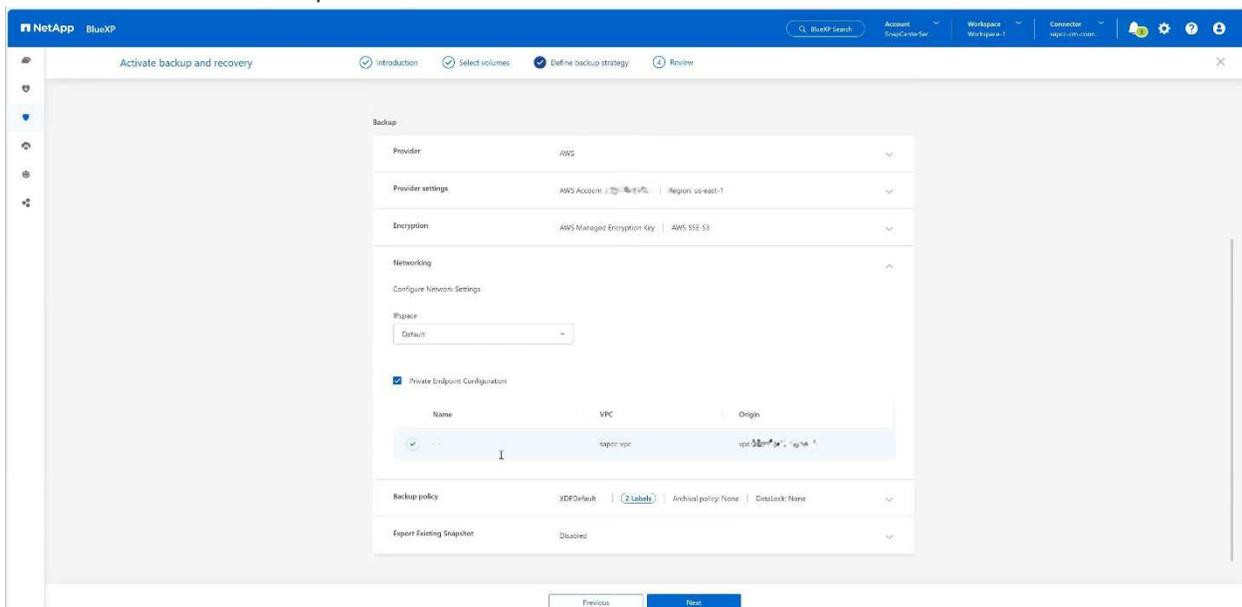
7. Check the **Backup Option** and choose your desired Backup Provider. In this example AWS. Keep the option checked for already existing policies. Uncheck options you do not want to use.



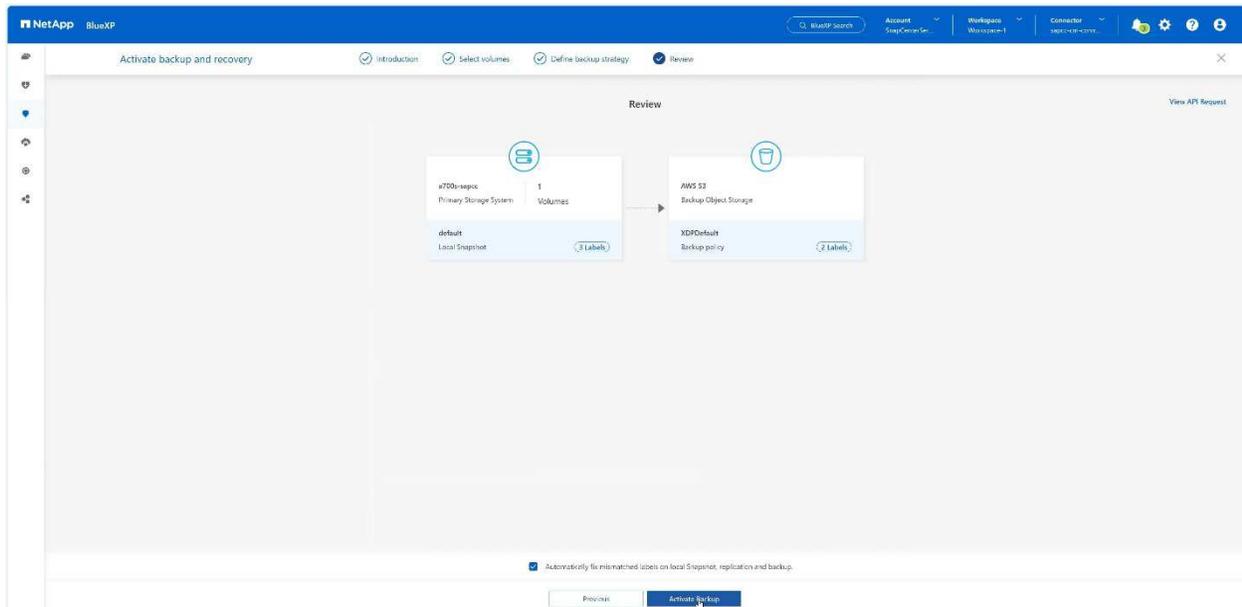
8. Create a new bucket or choose an existing one. Provide your AWS account settings, the region, your access key, and the secret key. Press **Next**.



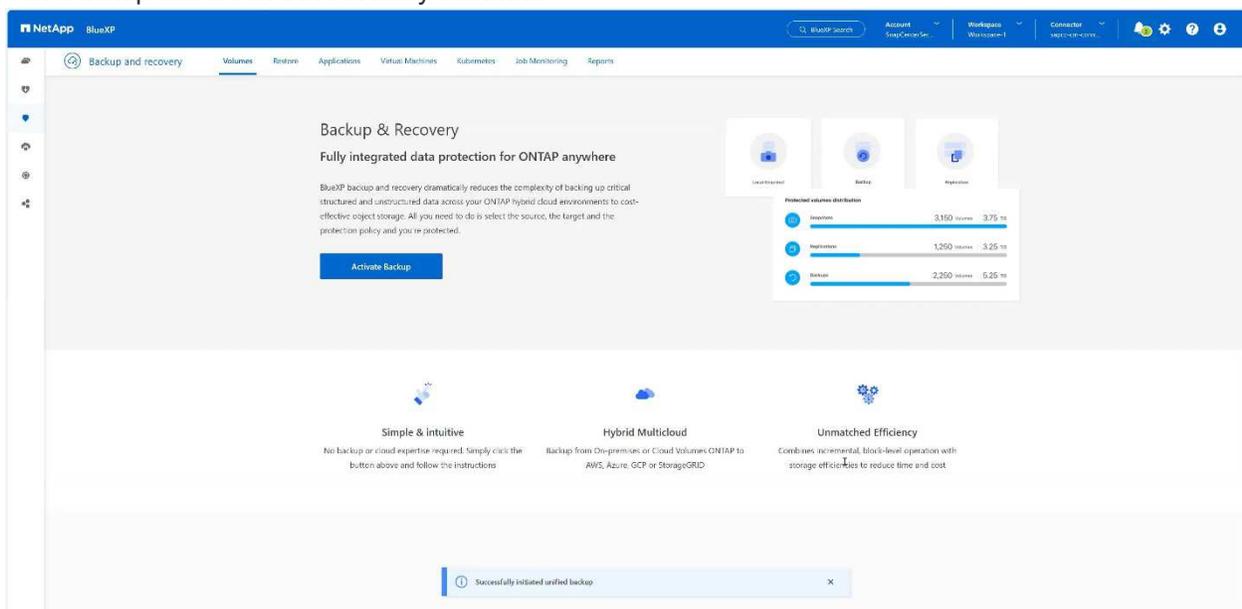
9. Choose the correct IPspace of your on-premises storage system, select **Privat Endpoint Configuration** and choose the VPC endpoint for the S3. Press **Next**.



10. Review your configuration and press **Activate Backup**.

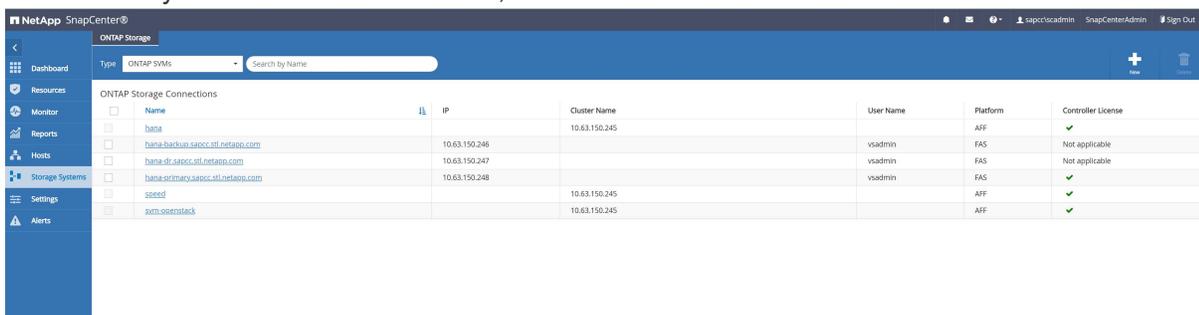


11. The backup has been successfully initiated.

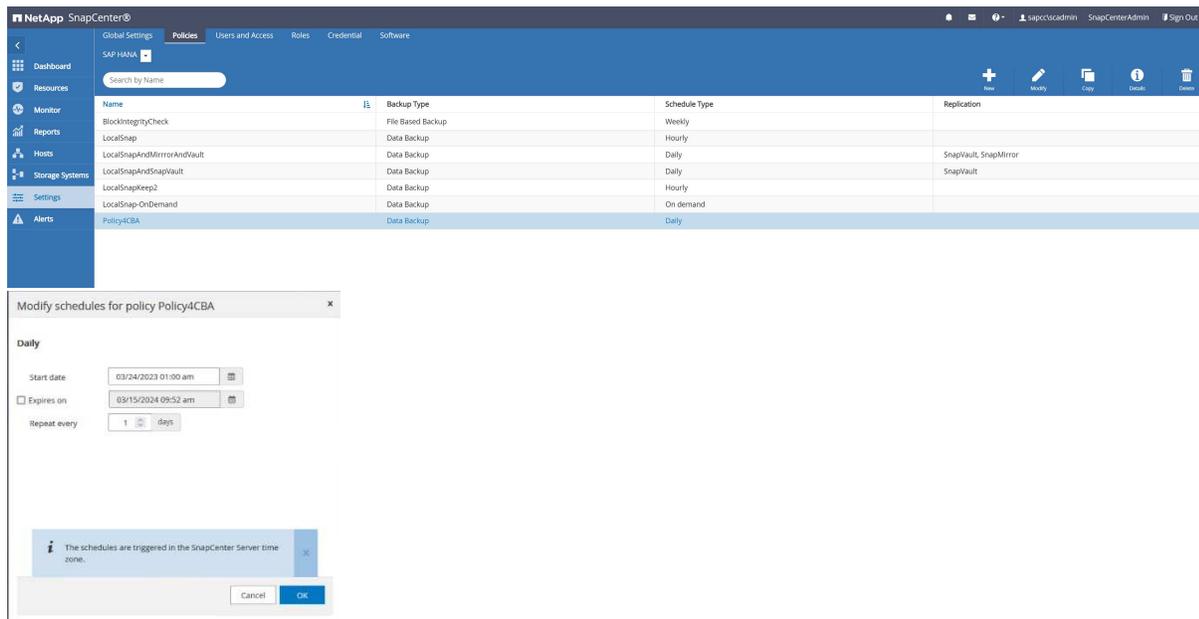


## Configure the SAP HANA system resource at SnapCenter

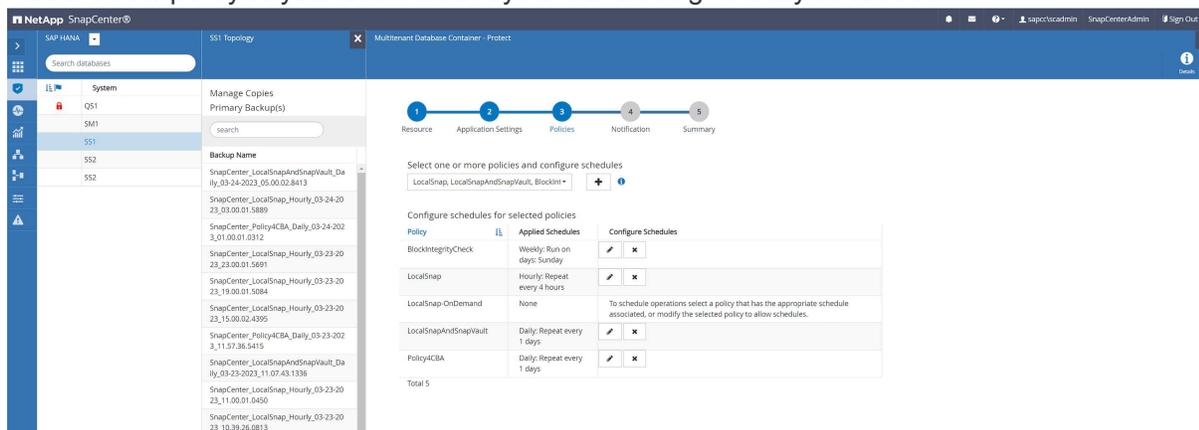
1. Check, if the SVM (hana in this example) where your SAP HANA system is stored has been added via the cluster. If only the SVM has been added, add the cluster.



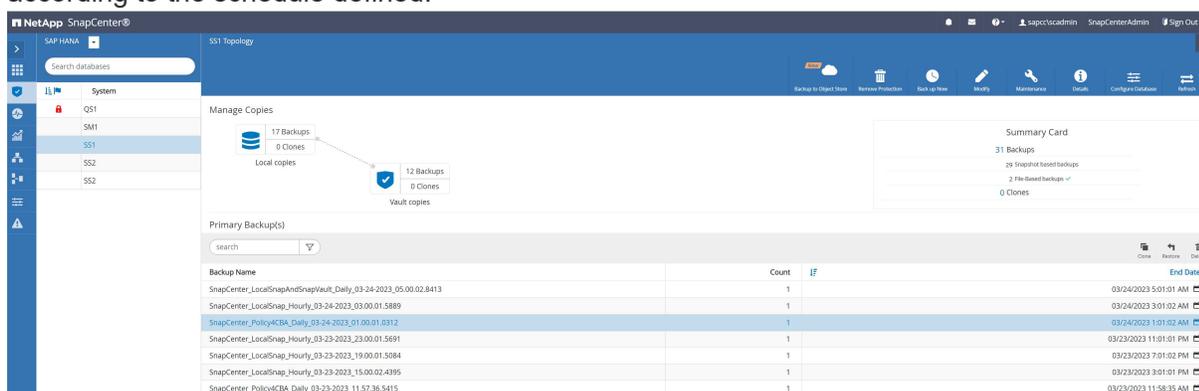
2. Define a schedule policy with either daily, weekly, or monthly schedule type.



3. Add the new policy to your SAP HANA system and assign a daily schedule.

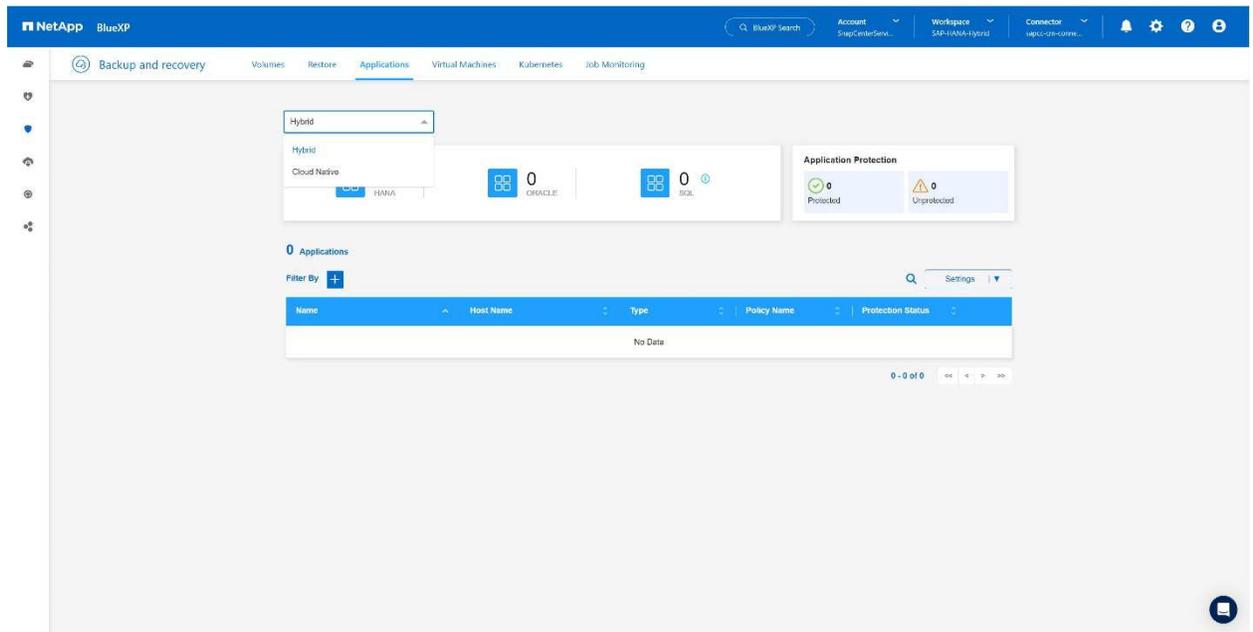


4. Once configured new backups with this policy will be available after the policy has been executed according to the schedule defined.

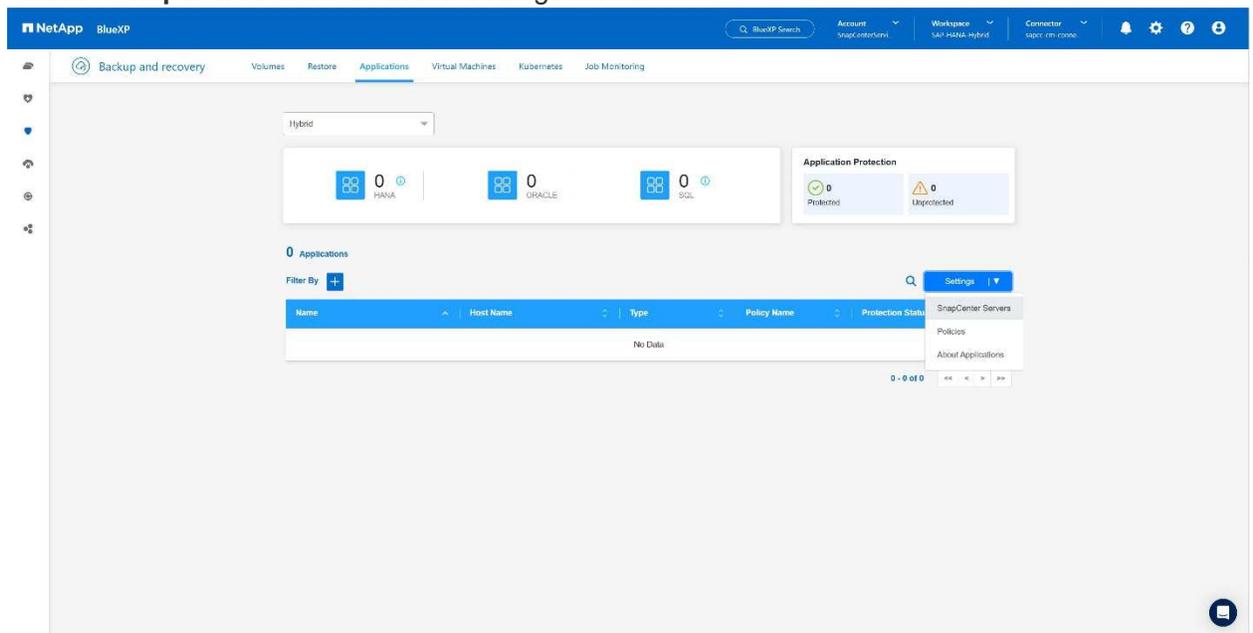


## Adding SnapCenter to the BlueXP Working Environment

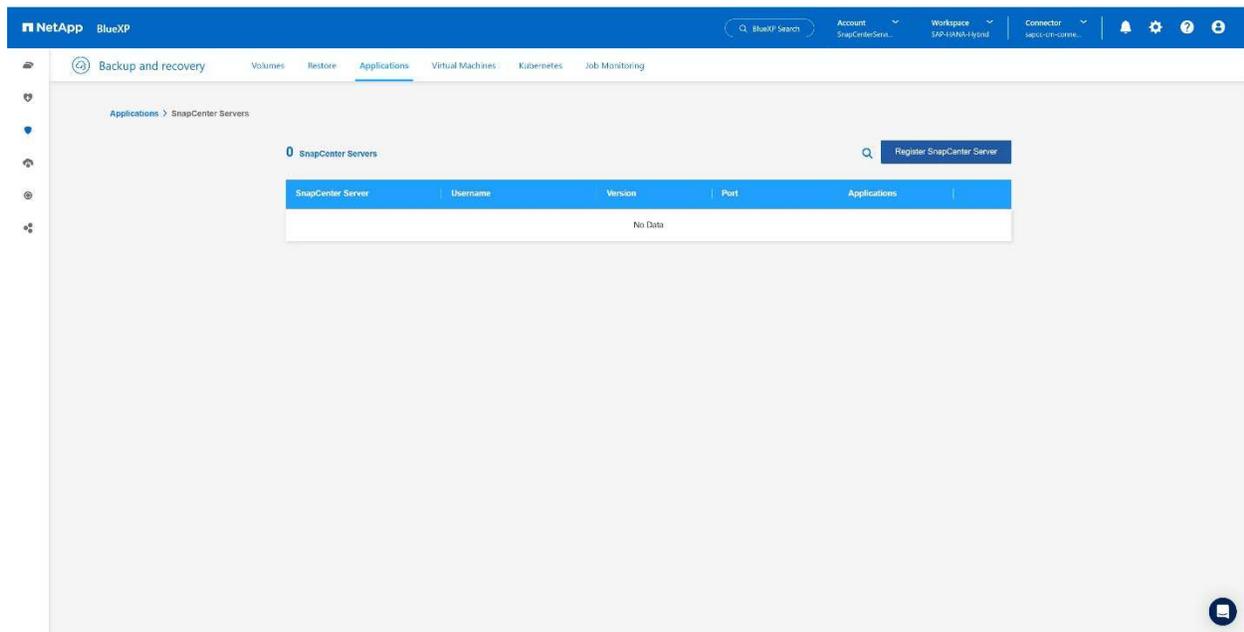
1. At the left menu choose **Protection** → **Backup and recovery** → **Applications**.
2. Choose **Hybrid** from the pulldown menu.



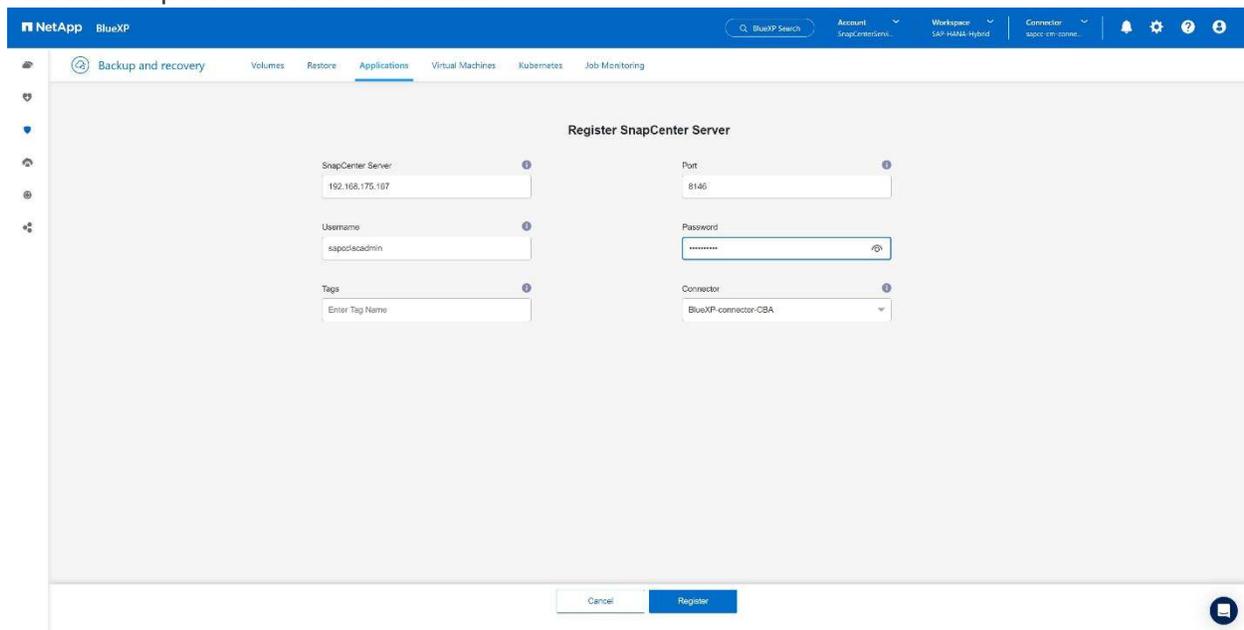
3. Choose **SnapCenter Servers** at the Settings menu.



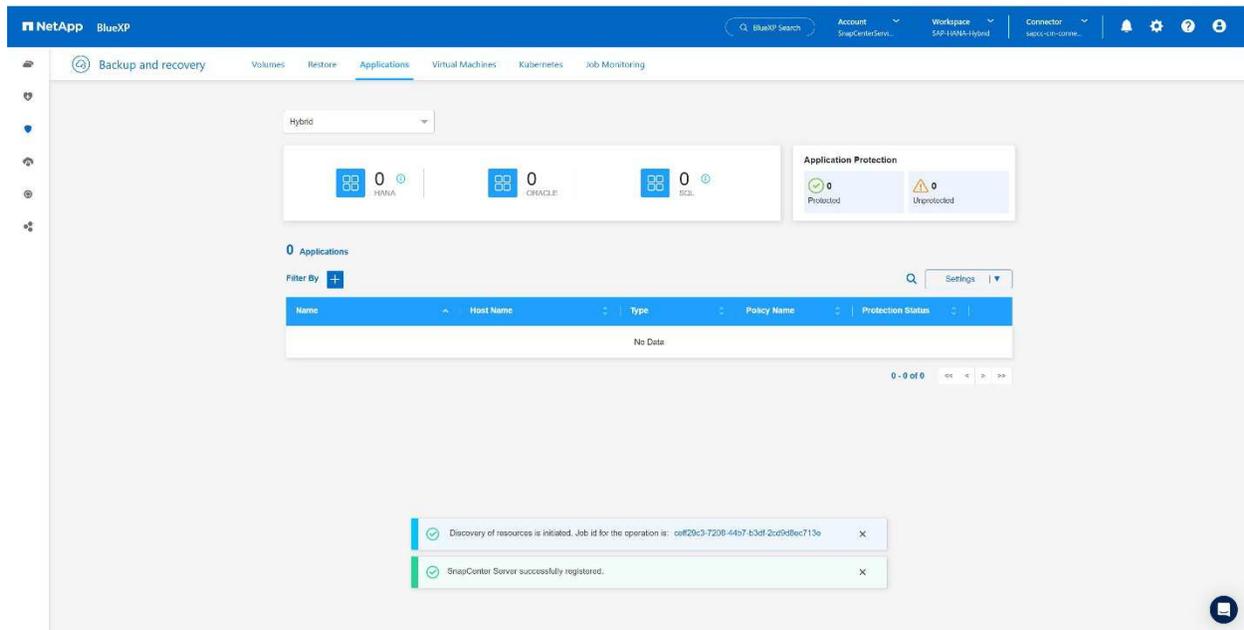
4. Register the SnapCenter Server.



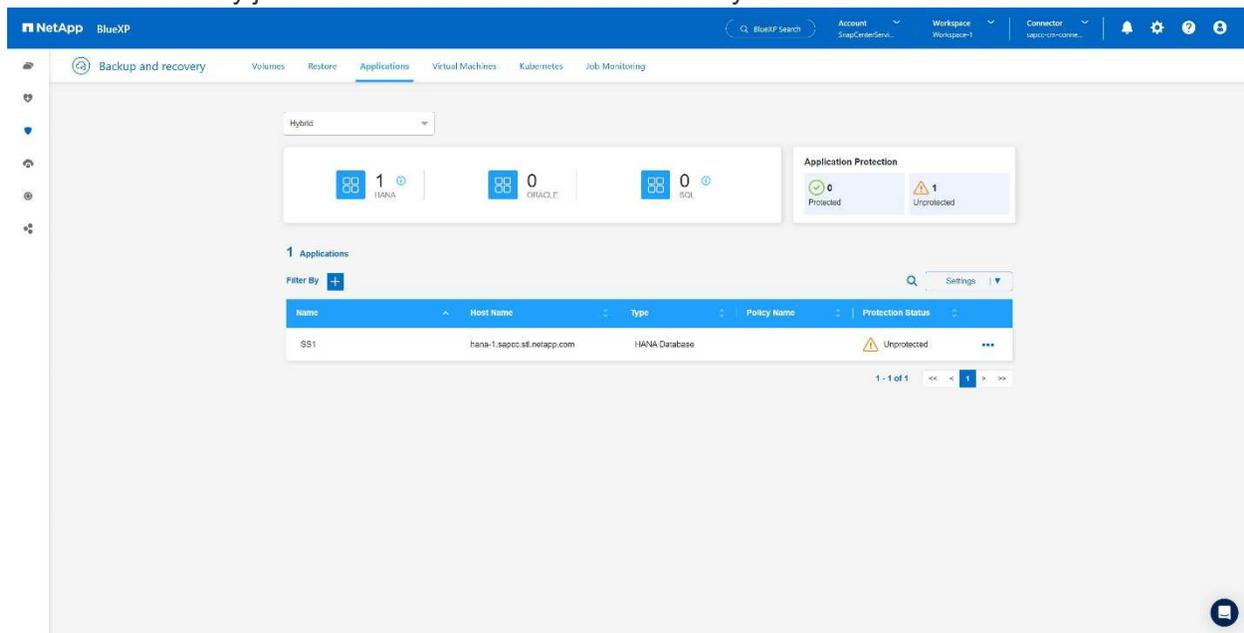
5. Add the SnapCenter Server credentials.



6. The SnapCenter Servers has been added and data will be discovered.

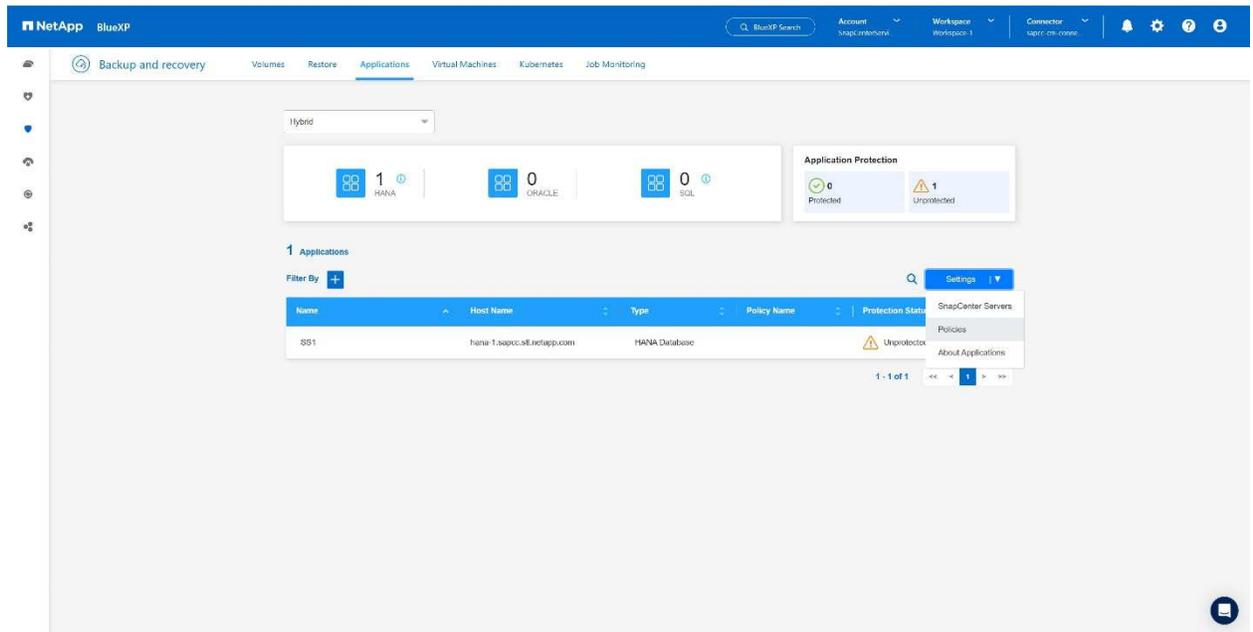


7. Once the discovery job has been finished the SAP HANA system will be available.

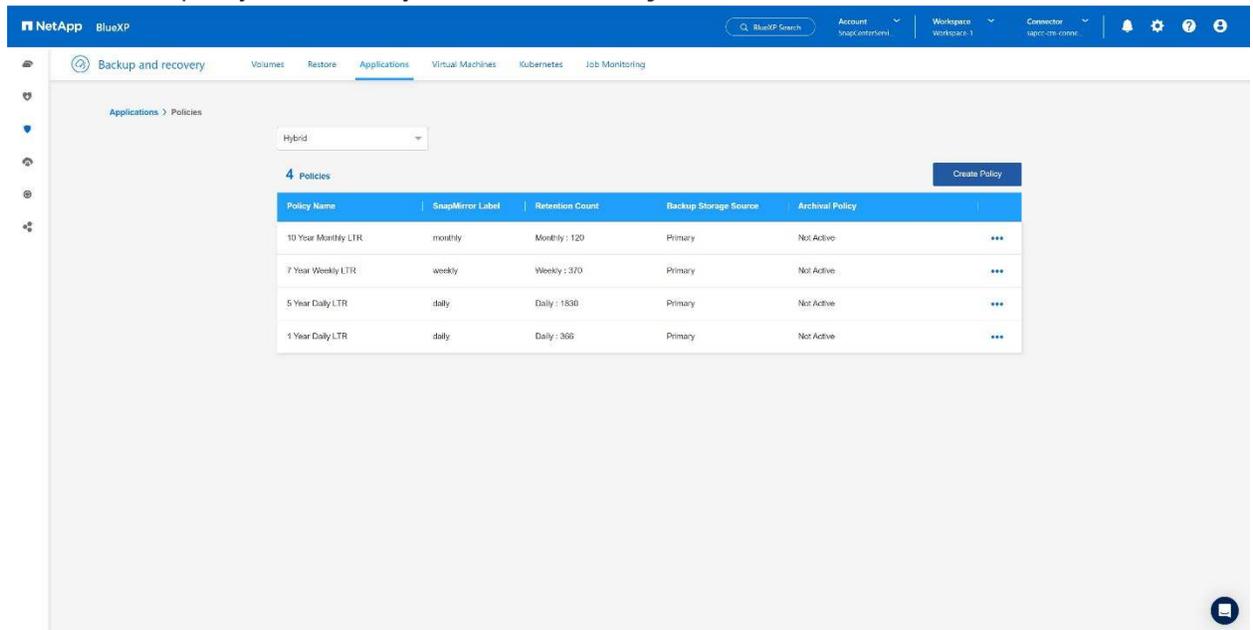


## Creating a Backup Policy for Application Backup

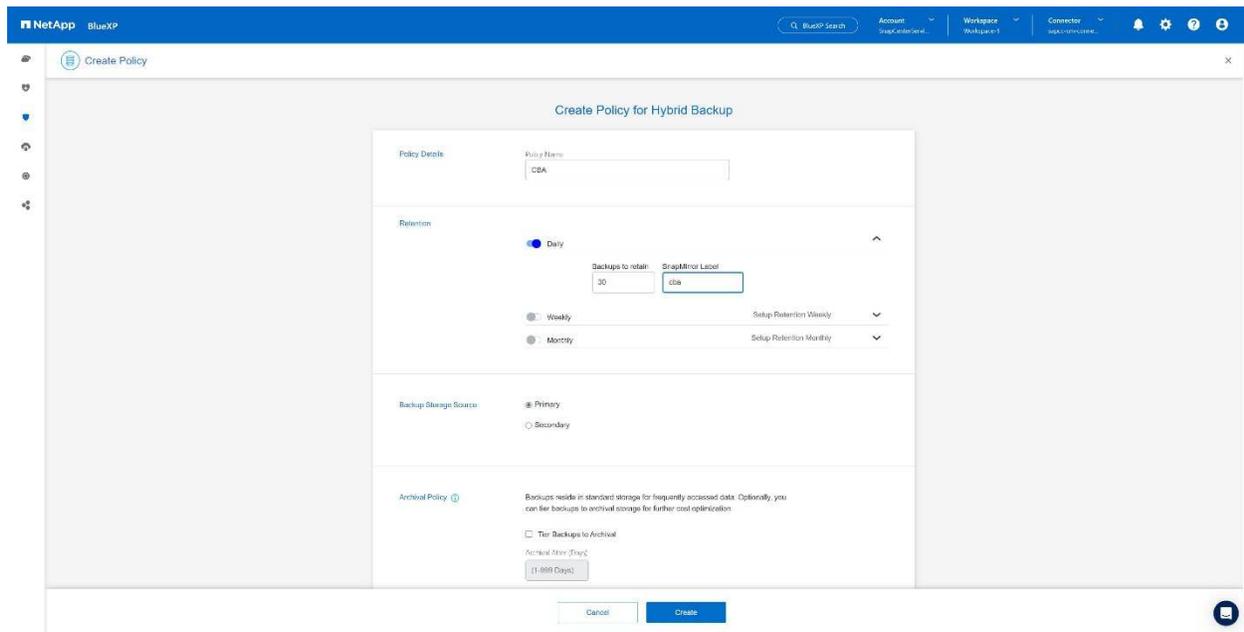
1. Choose **Policies** within the settings menu.



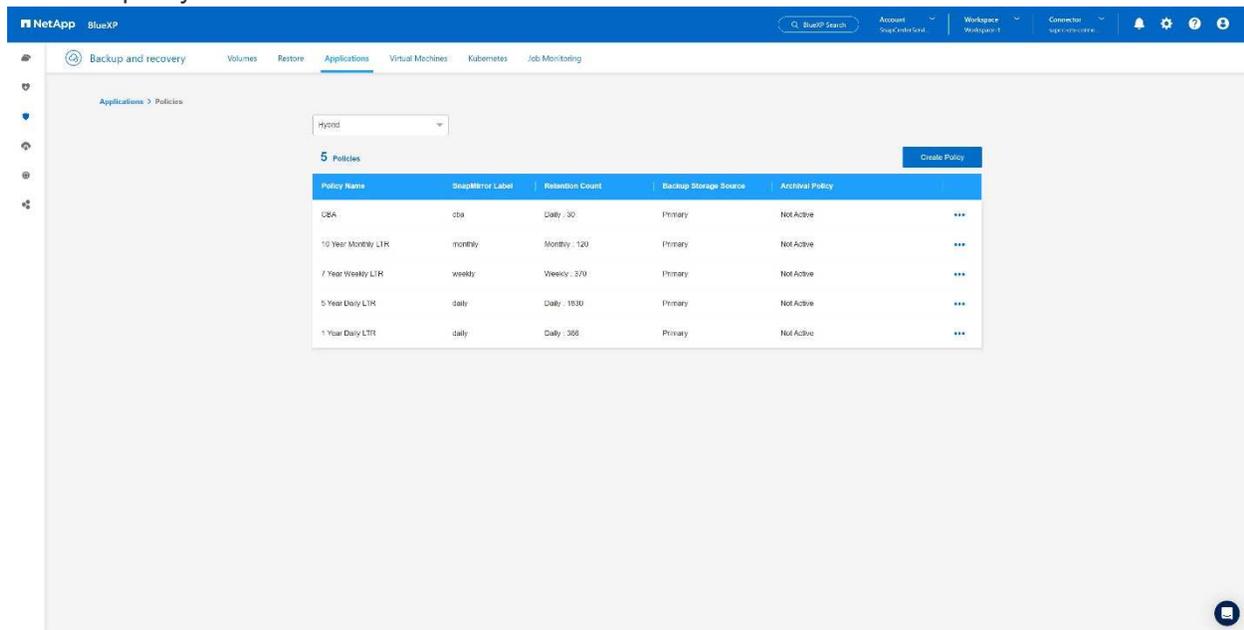
2. Create a new policy, if desired by click **Create Policy**.



3. Provide the policy name, desired SnapMirror label, choose your desired options, and press **Create**.

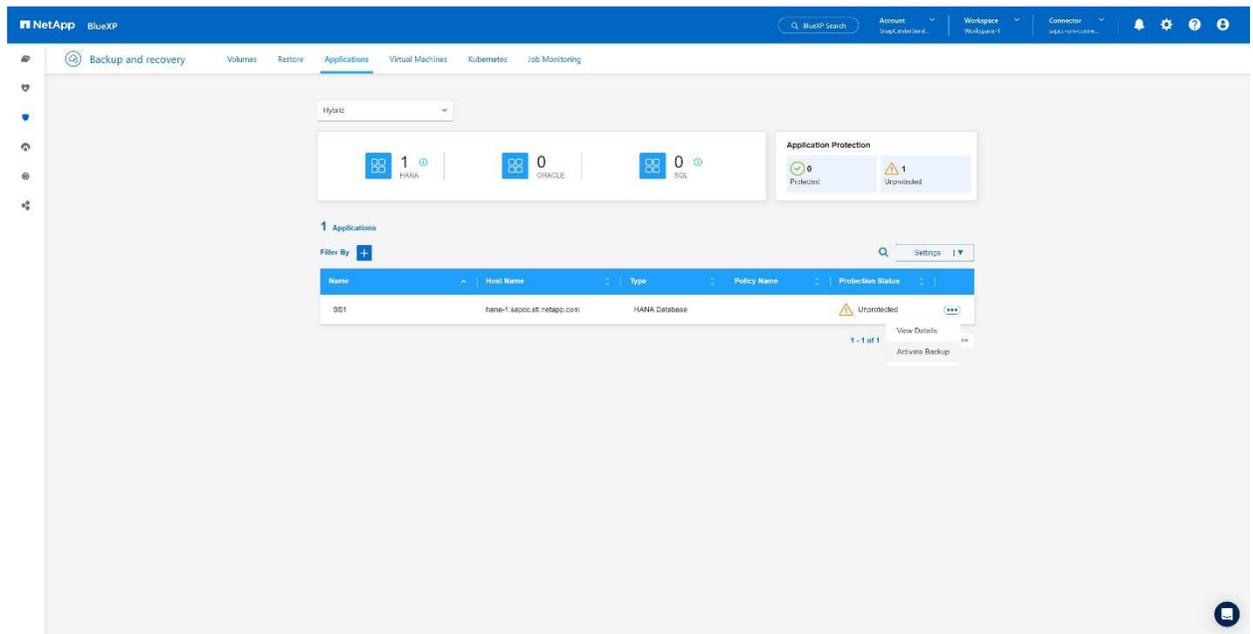


4. The new policy is available.

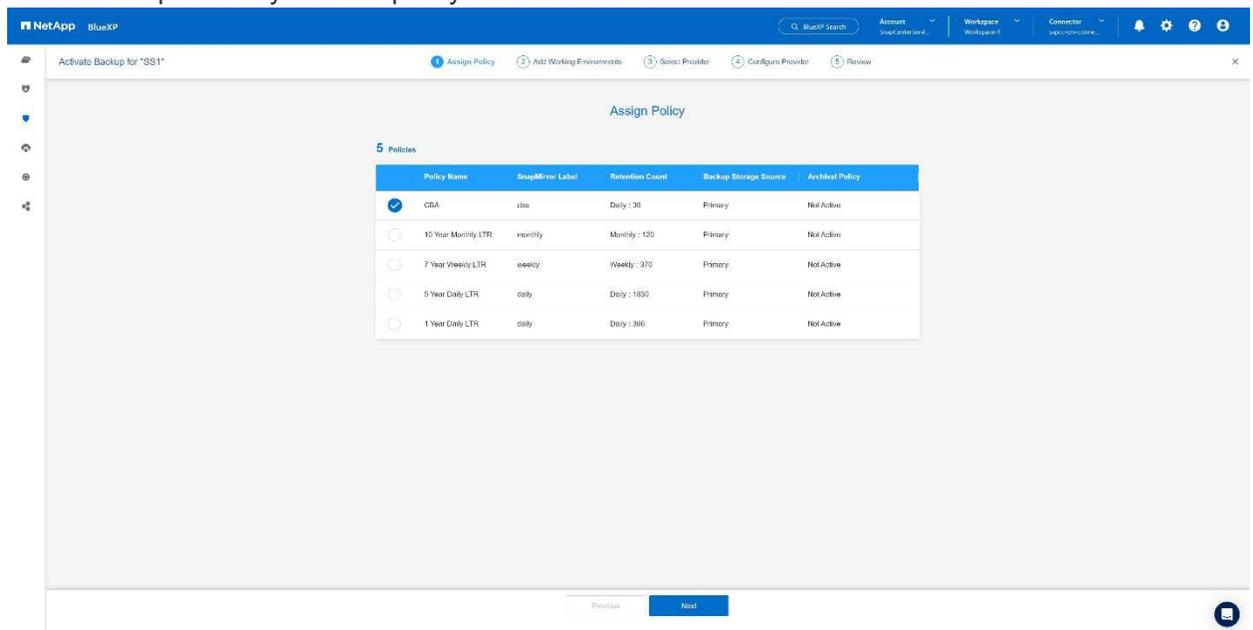


## Protecting the SAP HANA database with Cloud Backup for Applications

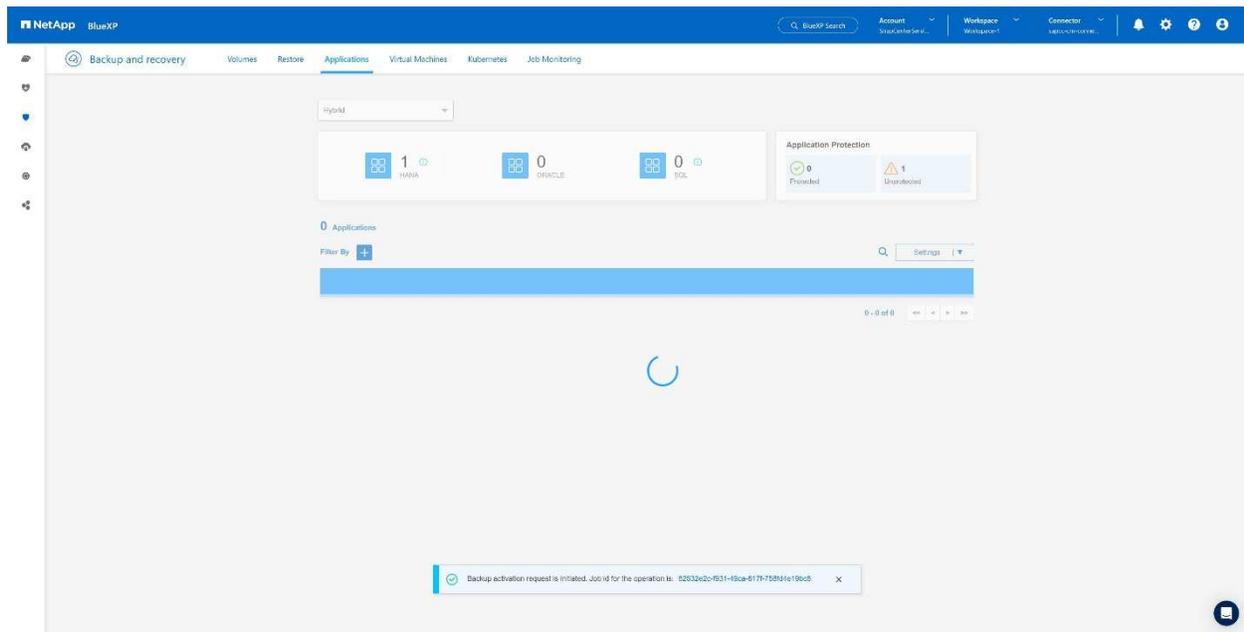
1. Choose **Activate Backup** for the SAP HANA system.



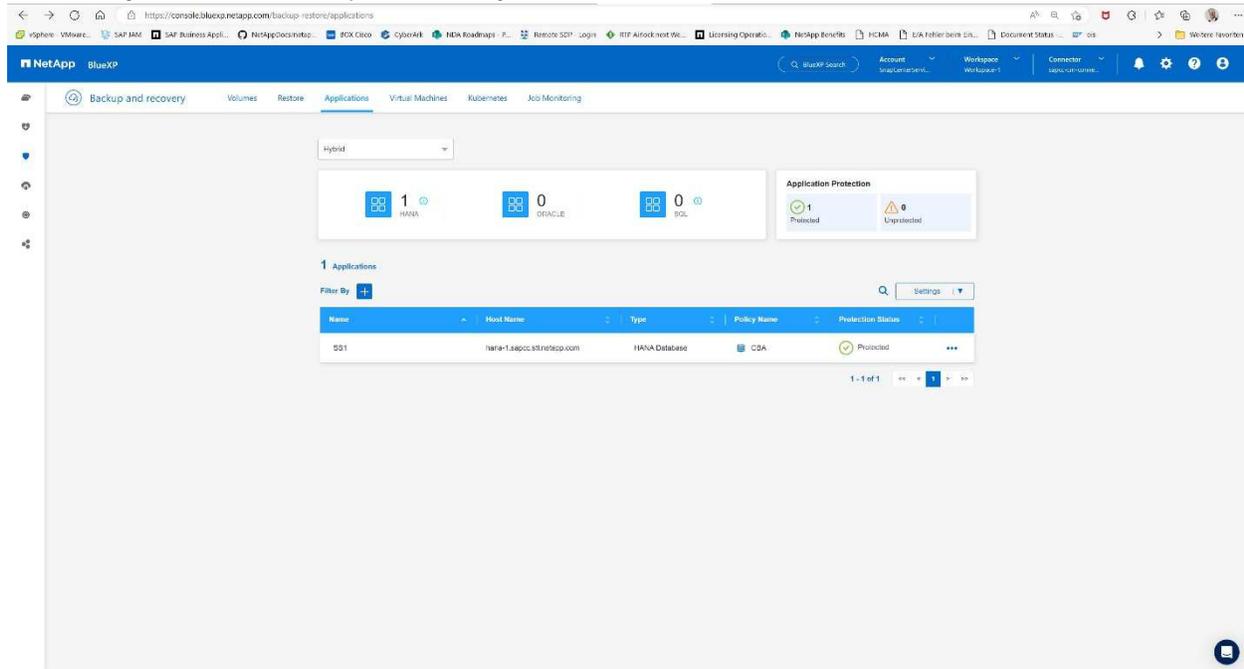
2. Choose the previously created policy and click **Next**.



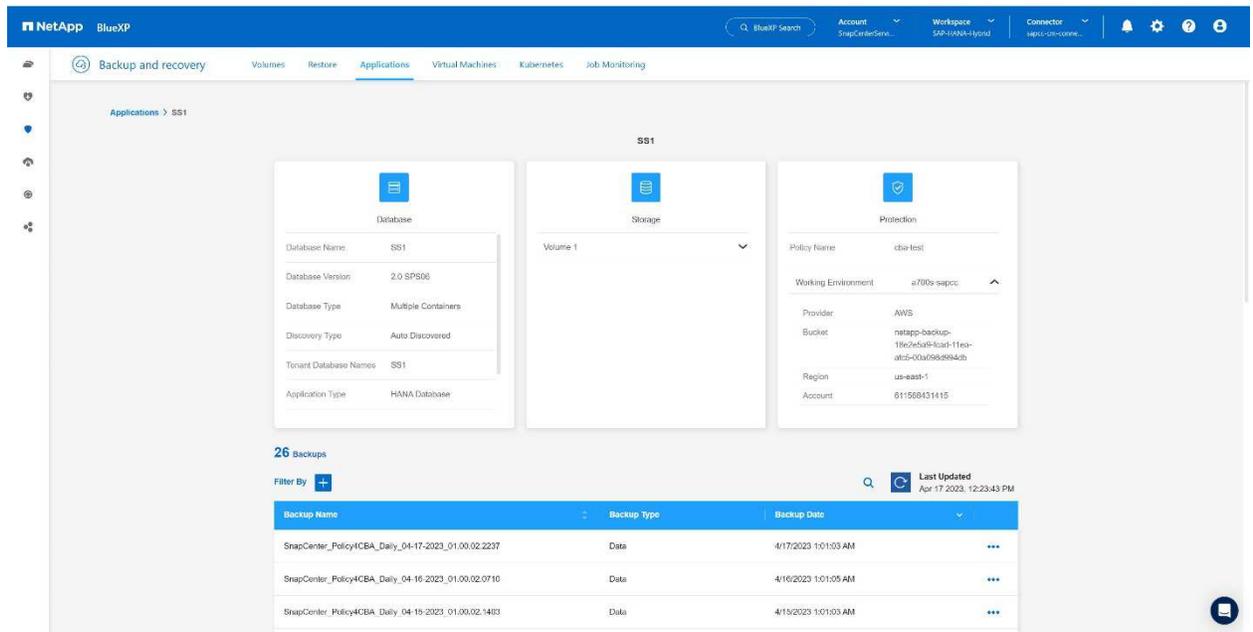
3. As the storage system and the connector have configured upfront the backup will be activated.



4. Once the job has been completed the System will be listed.



5. After some time the backups will be listed at the detail view of the SAP HANA System. A daily backup will be listed the next day.

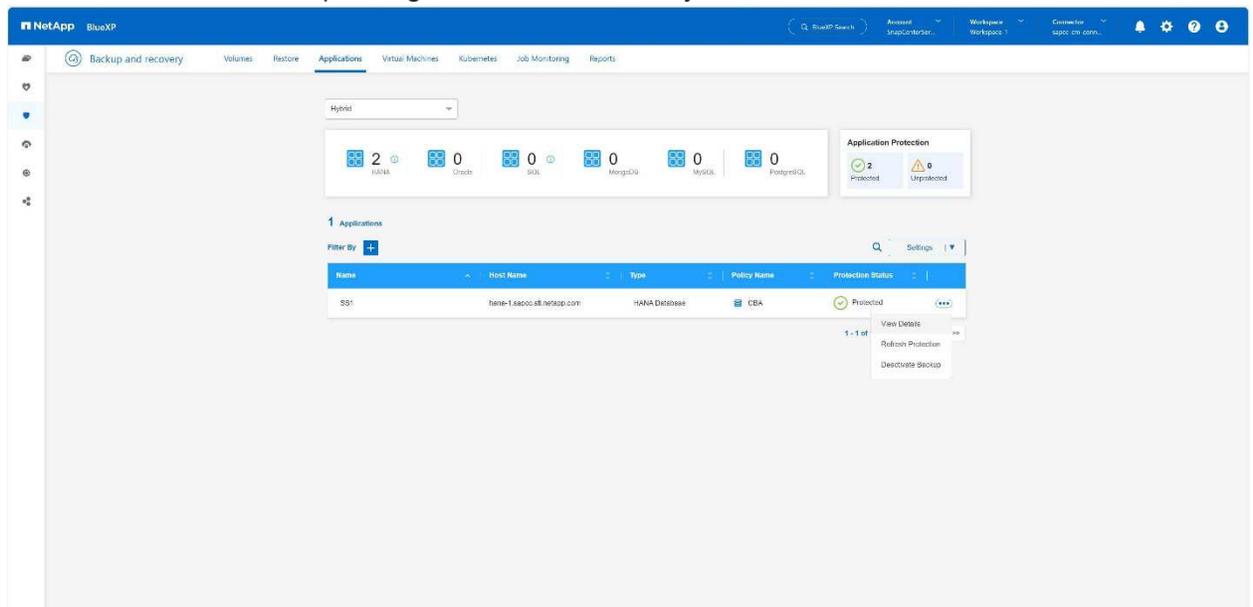


In some environments it might be necessary to remove any existing schedule settings of the snapmirror source. To do so execute the following command at the source ONTAP system: `snapmirror modify -destination -path <hana-cloud-svm>:<SID_data_mnt00001>_copy -schedule ""`.

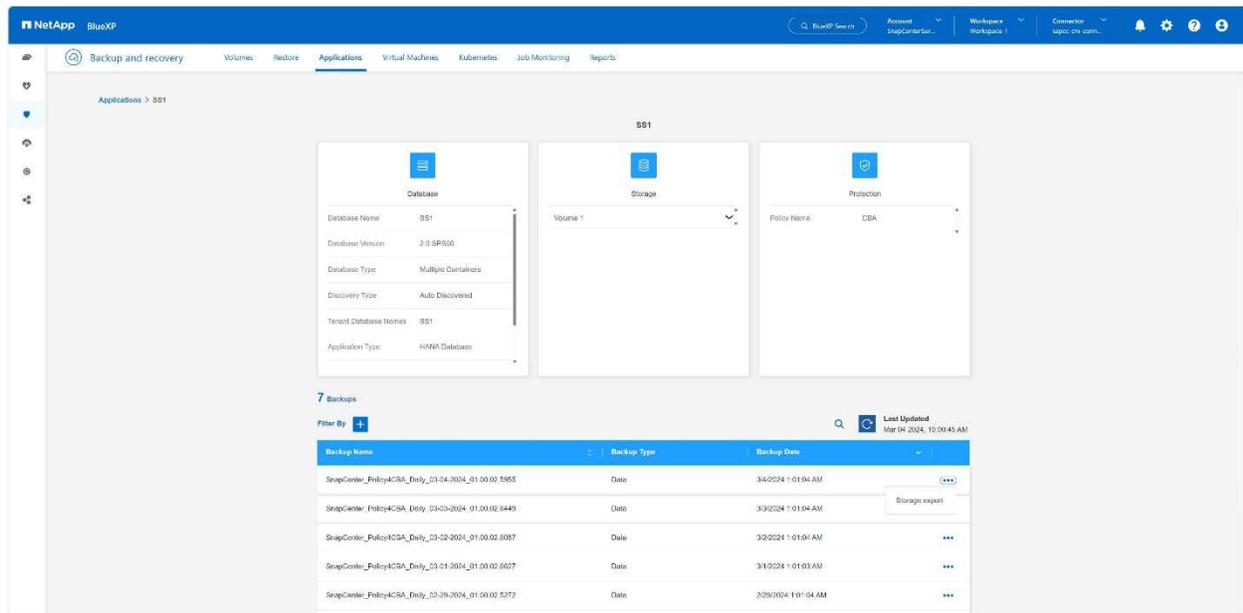
## Restoring SAP HANA BlueXP Backup

A restore from of the backup can only be done to an on-premises NetApp ONTAP based storage system or NetApp CVO within the cloud. A restore can be done by doing the following steps:

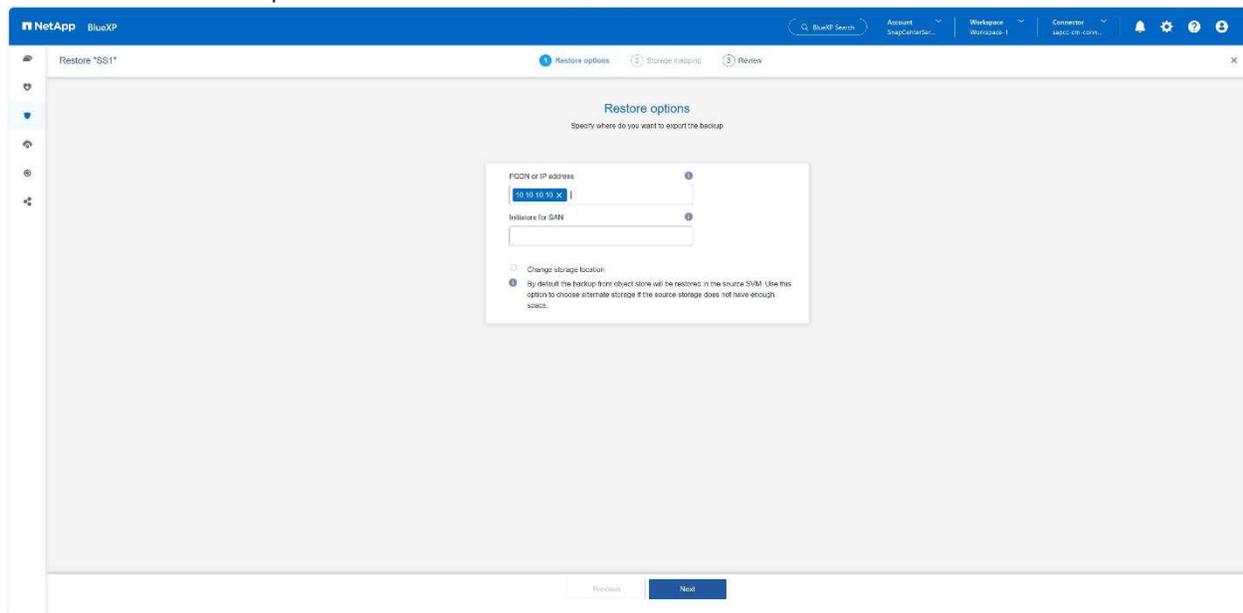
1. In BlueXP UI, click **Protection > Backup and recovery > Applications** and choose Hybrid.
2. In the **Filter By** field, select the filter **Type** and from the drop-down select **HANA**.
3. Click **View Details** corresponding to the database that you want to restore.



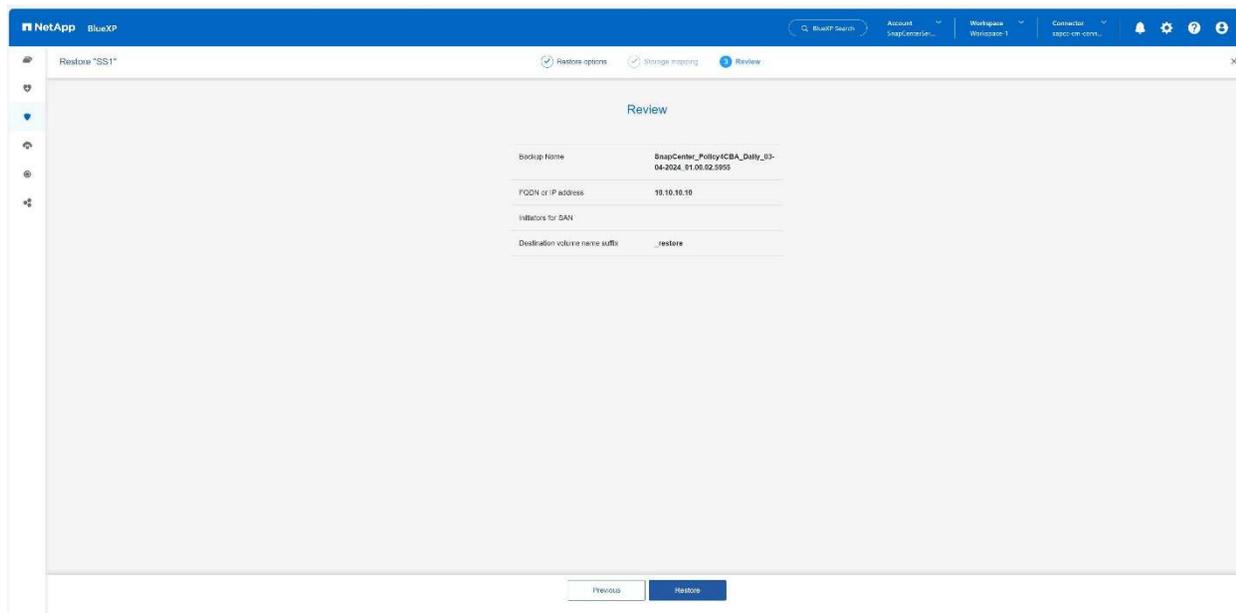
4. Select the desired backup and choose Storage Export.



5. Provide the desired options:



- a. For NAS environment, specify the FQDN or IP address of the host to which the volumes restored from object store are to be exported.
  - b. For SAN environment, specify the initiators of the host to which LUNs of the volumes restored from object store are to be mapped.
6. If the snapshot is in archival storage, select the priority to restore your data from the archival storage.
  7. If there is not enough space on the source storage or the source storage is down, select **Change storage location**.
  8. If you select **Change storage location**, you can append a suffix to the destination volume. If you have not selected the checkbox, then by default **\_restore** is appended to the destination volume. Click **Next**.
  9. If you selected Change Storage Location, specify the alternate storage location details where the data restored from the object store will be stored in the Storage mapping page and click **Next**.
  10. Review the details and click **Restore**.



This operation does only the storage export of the restored backup for the given host. You must manually mount the filesystem at the host and bring up the database. After utilizing the volume, the storage Administrator can delete the volume from the ONTAP cluster.

## Additional Information and Version History

This section lists where to find additional information and shows the version history.

### Where to find additional information

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

- NetApp BlueXP backup and recovery Product Documentation  
[Protect your on-premises applications data | NetApp Documentation](#)
- [SAP HANA backup and recovery with SnapCenter](#)

### Version history

Version	Date	Document version history
Version 1.0	March 2024	Initial version

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

## SAP HANA System Replication Backup and Recovery with SnapCenter

## TR-4719: SAP HANA System Replication - Backup and Recovery with SnapCenter

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

Author: Nils Bauer, NetApp

There are two primary use cases that can be combined:

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

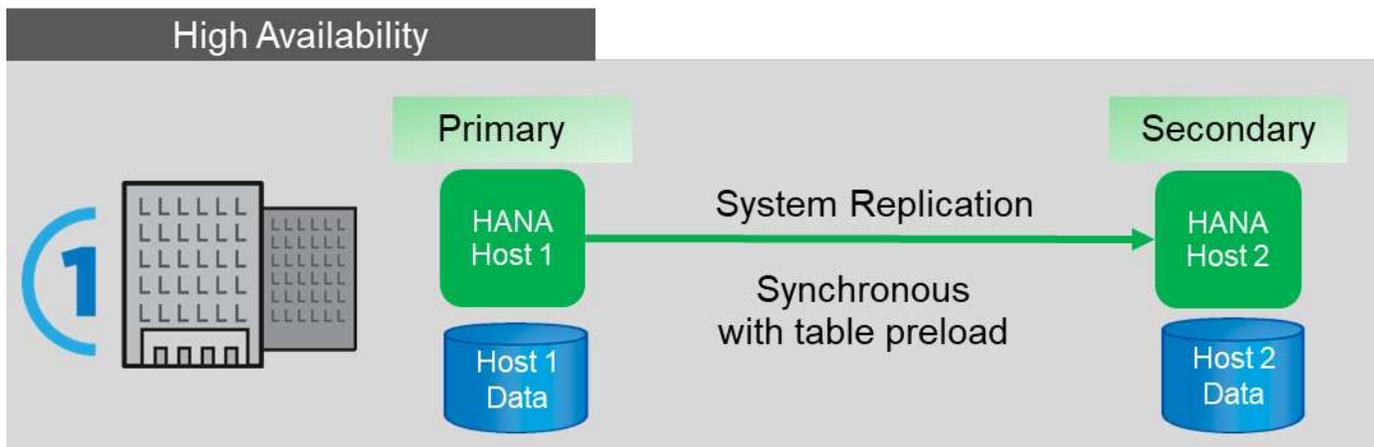
### High availability with an RPO of zero and a minimal RTO

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

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

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

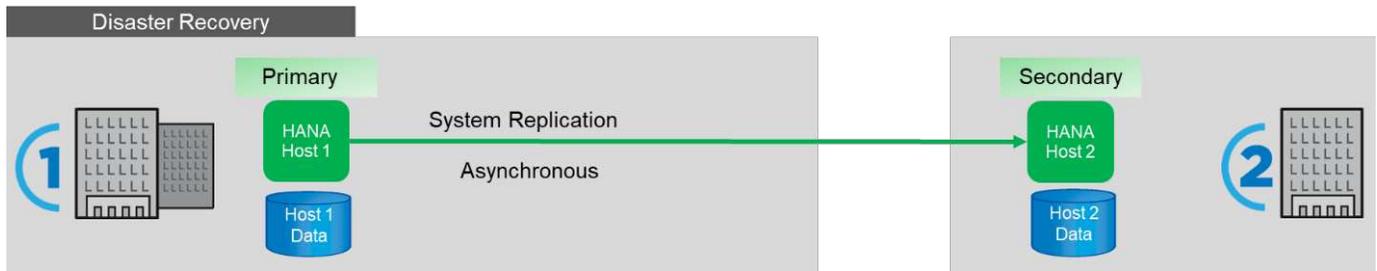
The rest of this document focuses on backup operations with SAP System Replication configured as a high-availability solution.



### Disaster recovery over a large distance

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

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



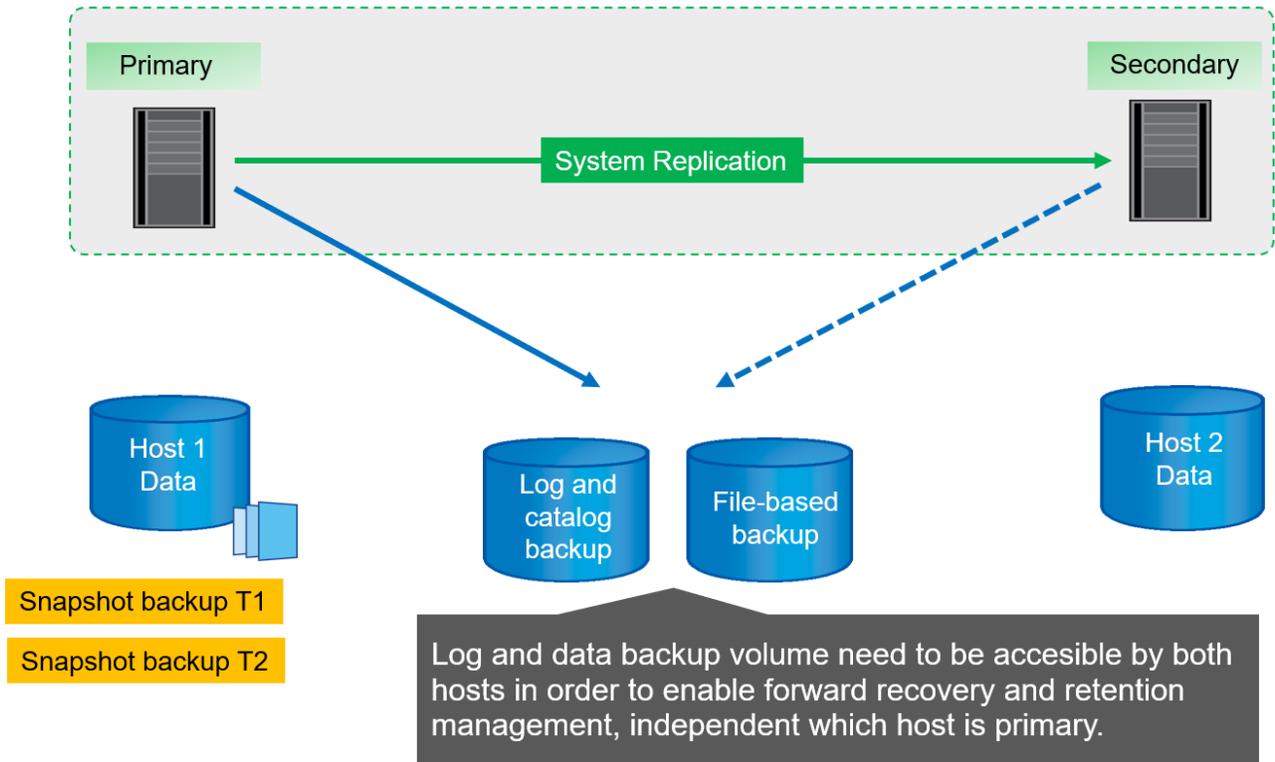
## Storage Snapshot backups and SAP System Replication

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

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

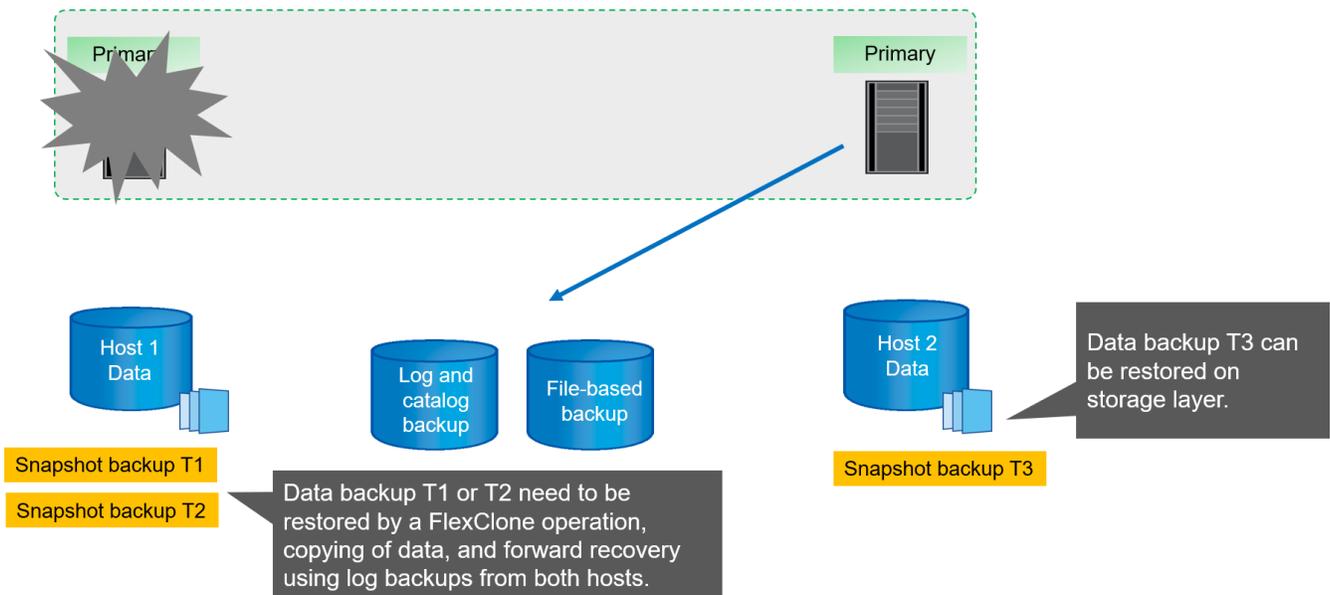
The ability to use any backup for restore and to do forward recovery using log backups from both hosts requires a shared log backup location that is accessible from both hosts. NetApp recommends that you use a shared storage volume. However, you should also separate the log backup destination into subdirectories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

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



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

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

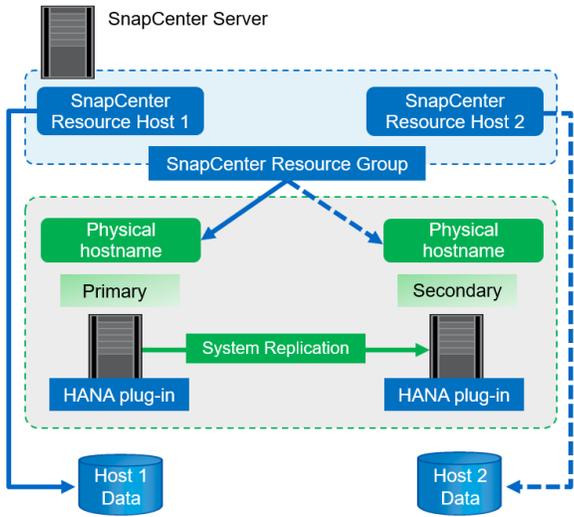


## SnapCenter configuration options for SAP System Replication

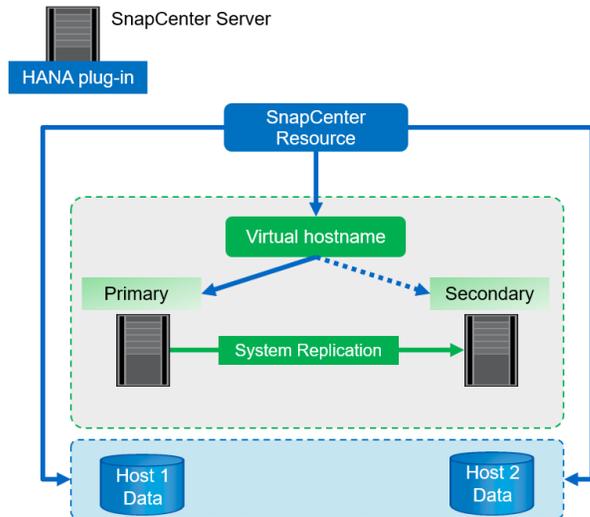
There are two options for configuring data protection with NetApp SnapCenter software in an SAP HANA System Replication environment:

- A SnapCenter resource group including both SAP HANA hosts and auto discovery with SnapCenter version 4.6 or higher.
- A single SnapCenter resource for both SAP HANA hosts using a virtual IP address.

**Option 1: SnapCenter 4.6 auto discovery of HANA System Replication**



**Option 2: SnapCenter manual resource configuration with central HANA plug-in**



Starting with SnapCenter 4.6, SnapCenter supports auto-discovery of HANA systems configured in a HANA System Replication relationship. Each host is configured using its physical IP address (host name) and its individual data volume on the storage layer. The two Snapcenter resources are combined in a resource group, and SnapCenter automatically identifies which host is primary or secondary and executes the required backup operations accordingly. Retention management for Snapshot and file-based backups created by SnapCenter is performed across both hosts to ensure that old backups also get deleted at the current secondary host.

With a single-resource configuration for both SAP HANA hosts, the single SnapCenter resource is configured using the virtual IP address of the SAP HANA System Replication hosts. Both data volumes of the SAP HANA hosts are included in the SnapCenter resource. Because it is a single SnapCenter resource, retention management for Snapshot and file-based backups created by SnapCenter works independent of which host is currently primary or secondary. This options is possible with all SnapCenter releases.

The following table summarizes the key differences of the two configuration options.

	<b>Resource group with SnapCenter 4.6</b>	<b>Single SnapCenter resource and virtual IP address</b>
Backup operation (Snapshot and file-based)	Automatic identification of primary host in resource group	Automatically use virtual IP address
Retention management (Snapshot and file-based)	Automatically executed across both hosts	Automatically use single resource
Backup capacity requirements	Backups are only created at primary host volume	Backups are always created at both hosts volumes. The backup of the second host is only crash consistent and cannot be used to do a roll forward.

	<b>Resource group with SnapCenter 4.6</b>	<b>Single SnapCenter resource and virtual IP address</b>
Restore operation	Backups from current active host are available for restore operation	Pre-backup script required to identify which backups are valid and can be used for restore
Recovery operation	All recovery options available, same as for any auto-discovered resource	Manual recovery required



In general, NetApp recommends using the resource group configuration option with SnapCenter 4.6 to protect HANA systems with enabled HANA System Replication. Using a single SnapCenter resource configuration is only required if the SnapCenter operation approach is based on a central plug-in host and the HANA plug-in is not deployed on the HANA database hosts.

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

## SnapCenter 4.6 configuration using a resource group

SnapCenter 4.6 supports auto discovery for HANA systems configured with HANA System Replication. SnapCenter 4.6 includes the logic to identify primary and secondary HANA hosts during backup operations and also handles retention management across both HANA hosts. In addition, automated restore and recovery is now also available for HANA System Replication environments.

### SnapCenter 4.6 configuration of HANA System Replication environments

The following figure shows the lab setup used for this chapter. Two HANA hosts, hana-3 and hana-4, were configured with HANA System Replication.

A database user “SnapCenter” was created for the HANA system database with the required privileges to execute backup and recovery operations (see [SAP HANA Backup and Recovery with SnapCenter](#)). A HANA user store key must be configured at both hosts using the above database user.

```
ss2adm@hana- 3: / > hdbuserstore set SS2KEY hana- 3:33313 SNAPCENTER
<password>
```

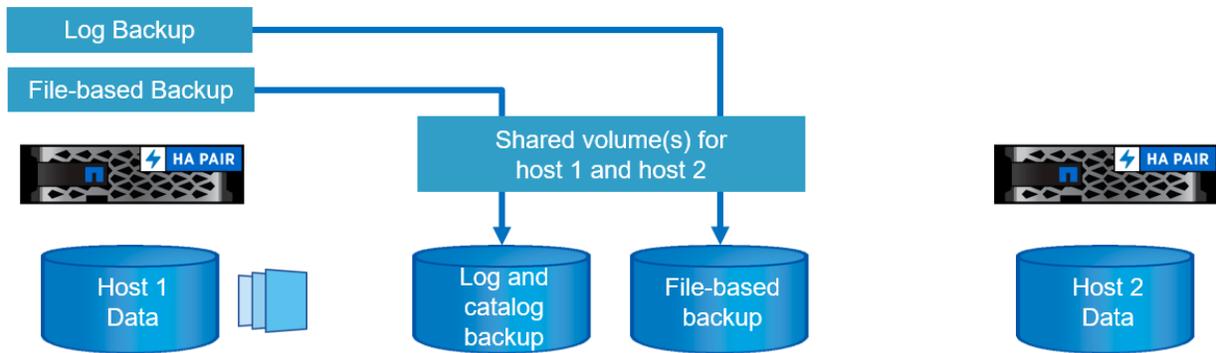
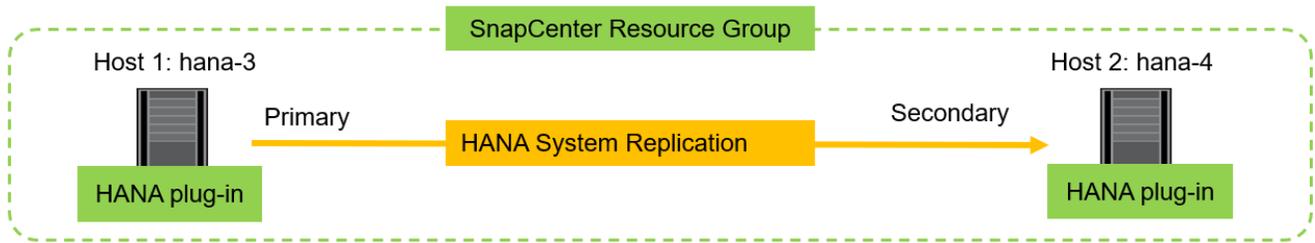
```
ss2adm@hana- 4:/ > hdbuserstore set SS2KEY hana-4:33313 SNAPCENTER
<password>
```

From a high-level perspective, you must perform the following steps to set up HANA System Replication within SnapCenter.

1. Install the HANA plugin on the primary and secondary host. Autodiscovery is executed and the HANA System Replication status is detected for each primary or secondary host.
2. Execute `SnapCenter configure database` and provide the `hdbuserstore` key. Further autodiscovery

operations are executed.

3. Create a resource group, including both hosts and configure protection.



After you have installed the SnapCenter HANA plug-in on both HANA hosts, the HANA systems are shown in the SnapCenter resource view in the same way as other autodiscovered resources. Starting with SnapCenter 4.6, an additional column is displayed that shows the status of HANA system replication (enabled/disabled, primary/secondary).

System	System ID (SID)	Tenant Databases	Replication	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
SS2	SS2	SS2	Enabled (Primary)	hana-3.sapcc.stl.netapp.com				Not protected
SS2	SS2	SS2	Enabled (Secondary)	hana-4.sapcc.stl.netapp.com				Not protected

By clicking the resource, SnapCenter requests the HANA user store key for the HANA system.

### Configure Database ✕

Plug-in host: hana-3.sapcc.stl.netapp.com

HDBSQL OS User: ss2adm

HDB Secure User Store Key:  i

Cancel
OK

Additional autodiscovery steps are executed, and SnapCenter show the resource details. With SnapCenter 4.6, the system replication status and the secondary server are listed in this view.

NetApp SnapCenter - Resource - Details

Details for selected resource

Type	Multitenant Database Container
HANA System Name	SS2
SID	SS2
Tenant Databases	SS2
Plug-in Host	hana-3.sapcc.stl.netapp.com
HDB Secure User Store Key	SS2KEY
HDBSQL OS User	ss2adm
Log backup location	/mnt/backup/SS2
Backup catalog location	/mnt/backup/SS2
System Replication	Enabled (Primary)
Secondary Servers	hana-4
plug-in name	SAP HANA
Last backup	None
Resource Groups	None
Policy	None
Discovery Type	Auto

Storage Footprint

SVM	Volume	Junction Path	LUN/Qtree
hana-primary.sapcc.stl.netapp.com	SS2_data_mnt00001	/SS2_data_mnt00001	

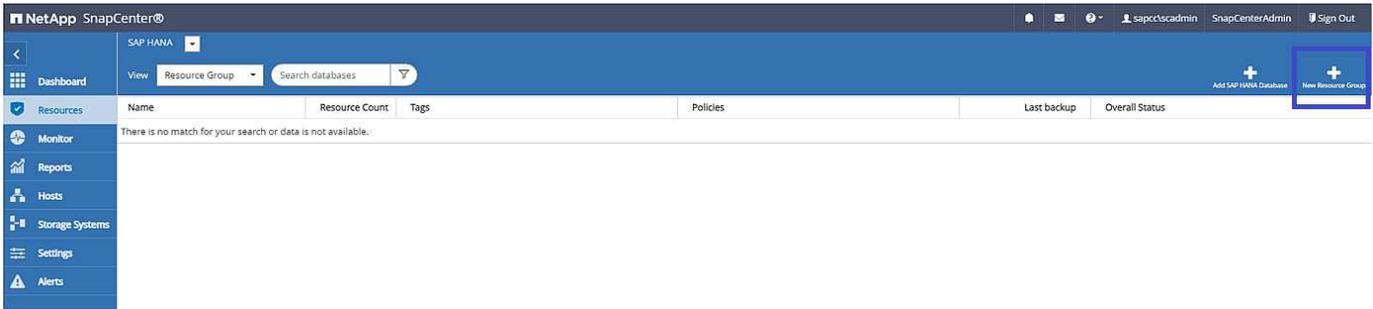
Activity: The 5 most recent jobs are displayed. 0 Completed, 0 Warnings, 0 Failed, 0 Canceled, 0 Running, 0 Queued.

After performing the same steps for the second HANA resource, the autodiscovery process is complete and both HANA resources are configured in SnapCenter.

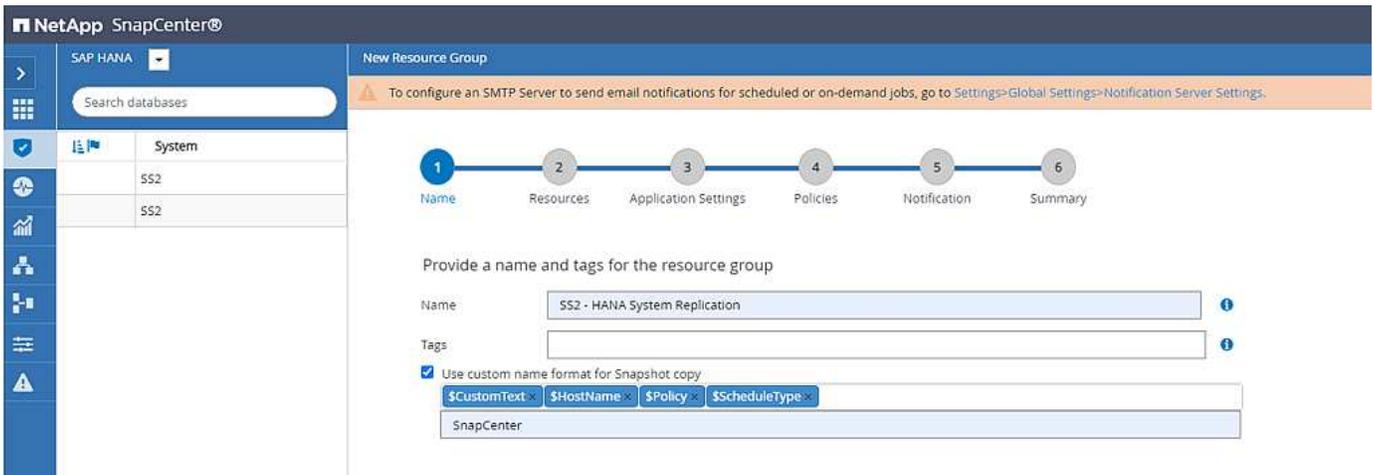
NetApp SnapCenter - Resources

System	System ID (SID)	Tenant Databases	Replication	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
SS2	SS2	SS2	Enabled (Primary)	hana-3.sapcc.stl.netapp.com				Not protected
SS2	SS2	SS2	Enabled (Secondary)	hana-4.sapcc.stl.netapp.com				Not protected

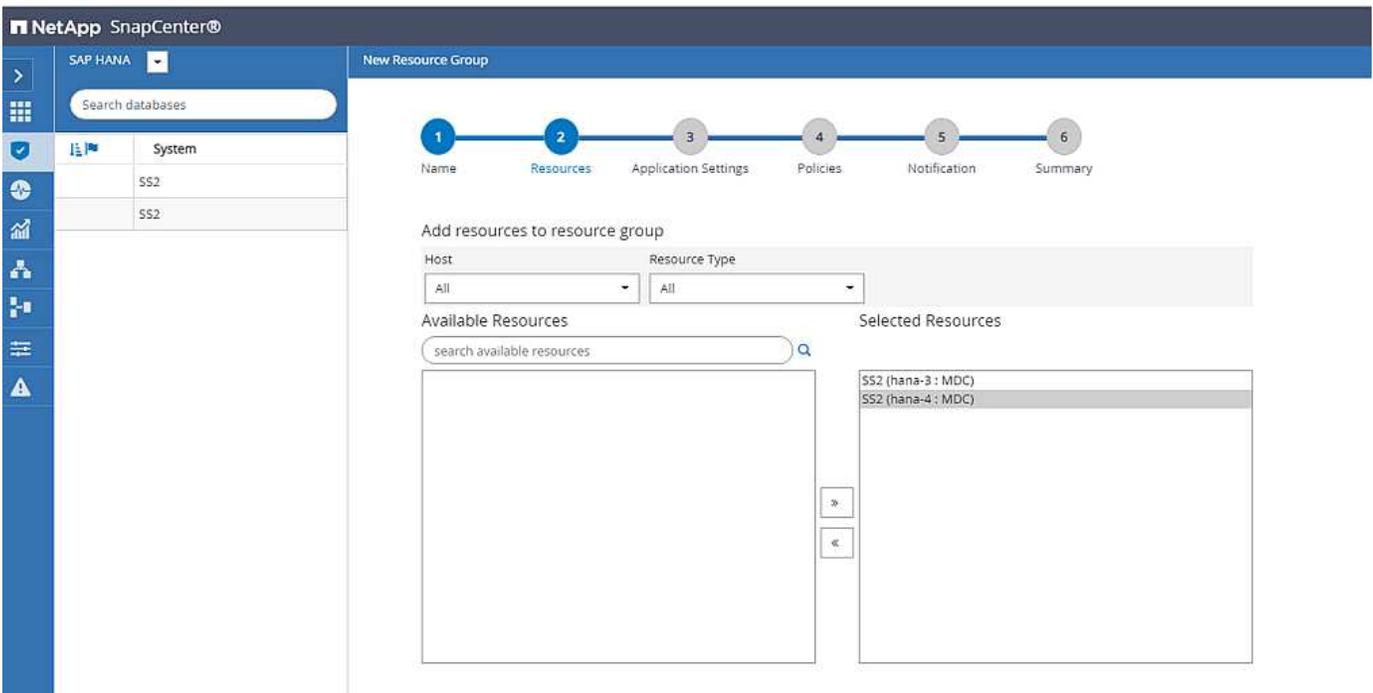
For HANA System Replication- enabled systems, you must configure a SnapCenter resource group, including both HANA resources.



NetApp recommends using a custom name format for the Snapshot name, which should include the hostname, the policy, and the schedule.



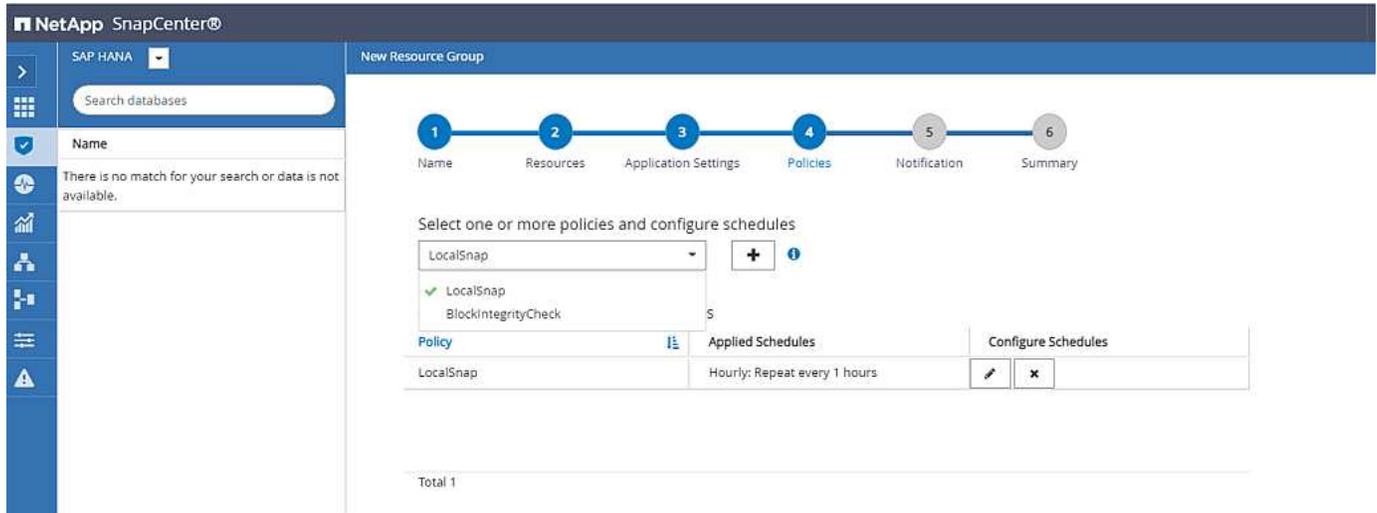
You must add both HANA hosts to the resource group.



Policies and schedules are configured for the resource group.



The retention defined in the policy is used across both HANA hosts. If, for example, a retention of 10 is defined in the policy, the sum of backups of both hosts is used as a criteria for backup deletion. SnapCenter deletes the oldest backup independently if it has been created at the current primary or secondary host.

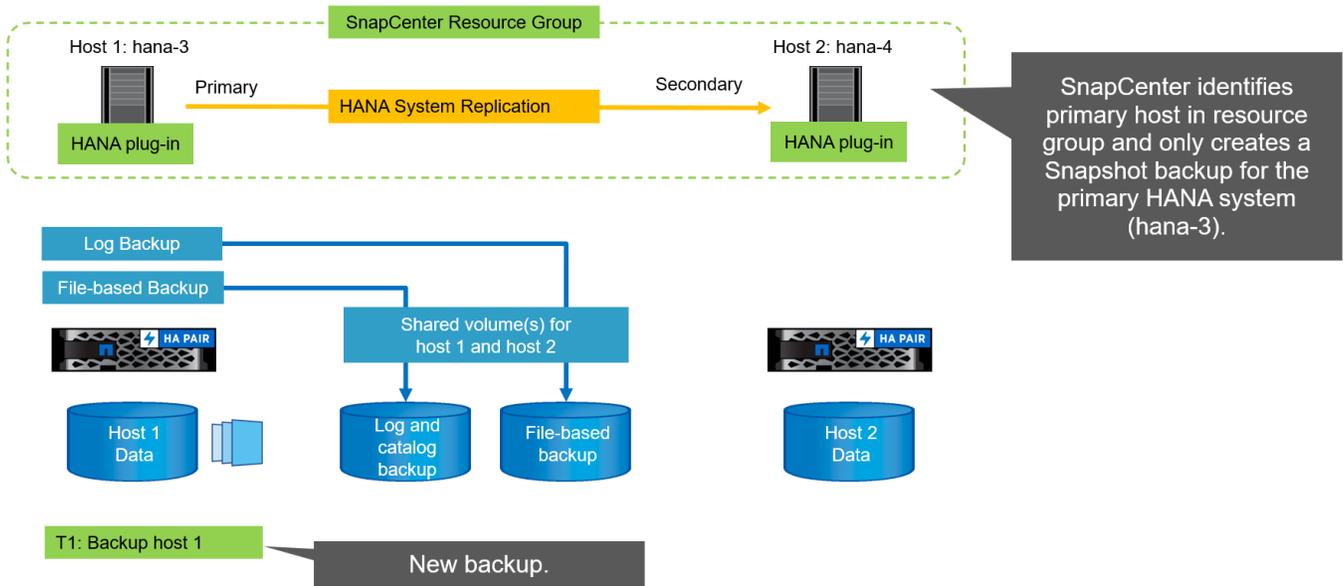


The resource group configuration is now finished and backups can be executed.

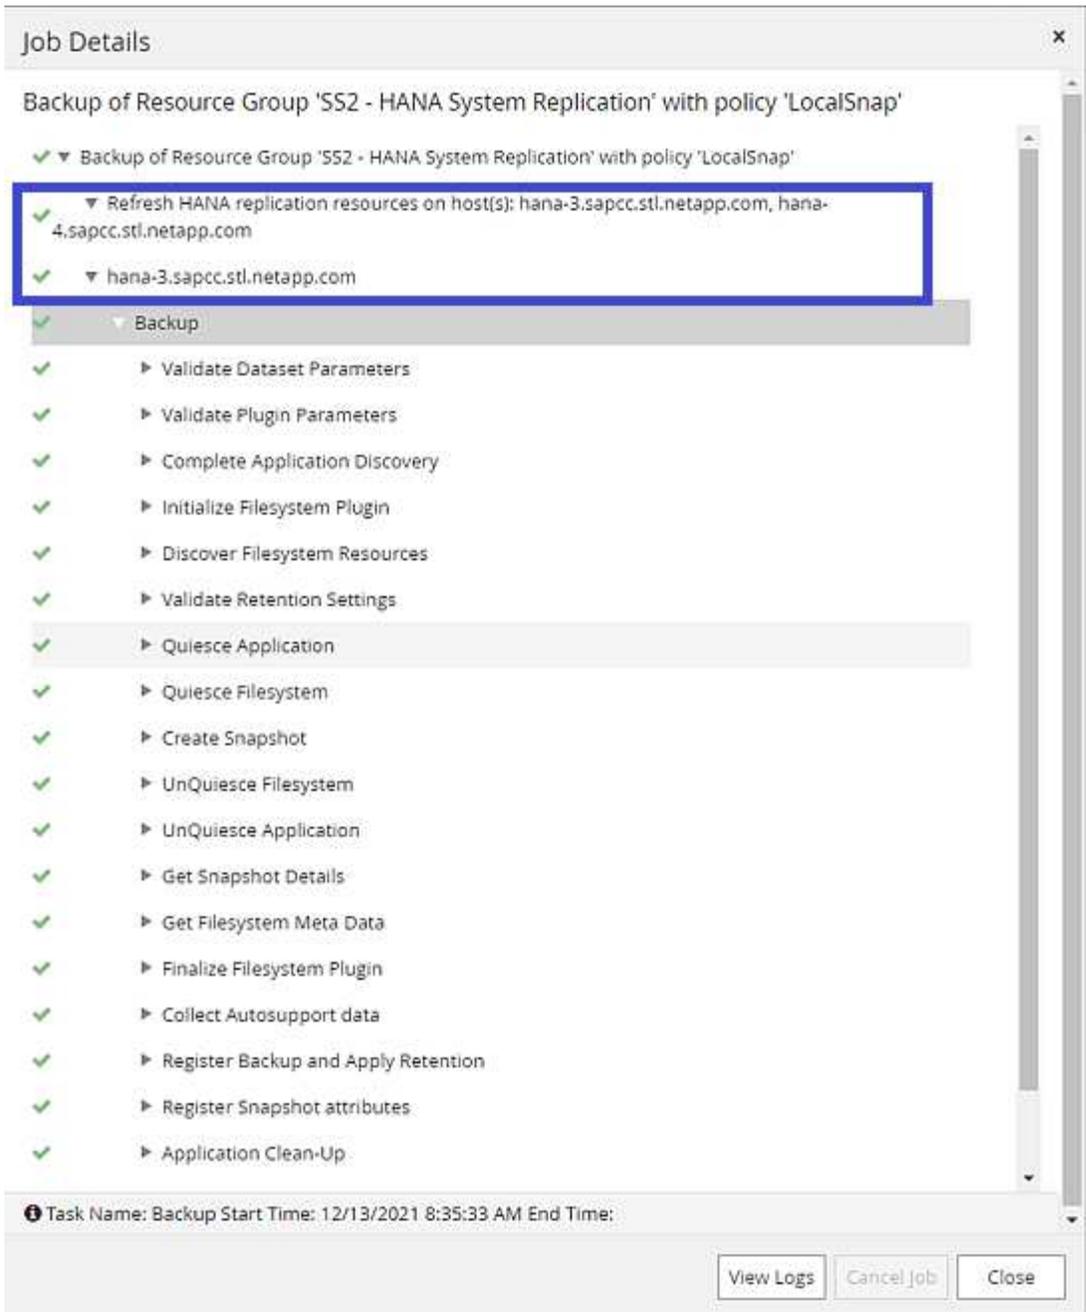


## Snapshot backup operations

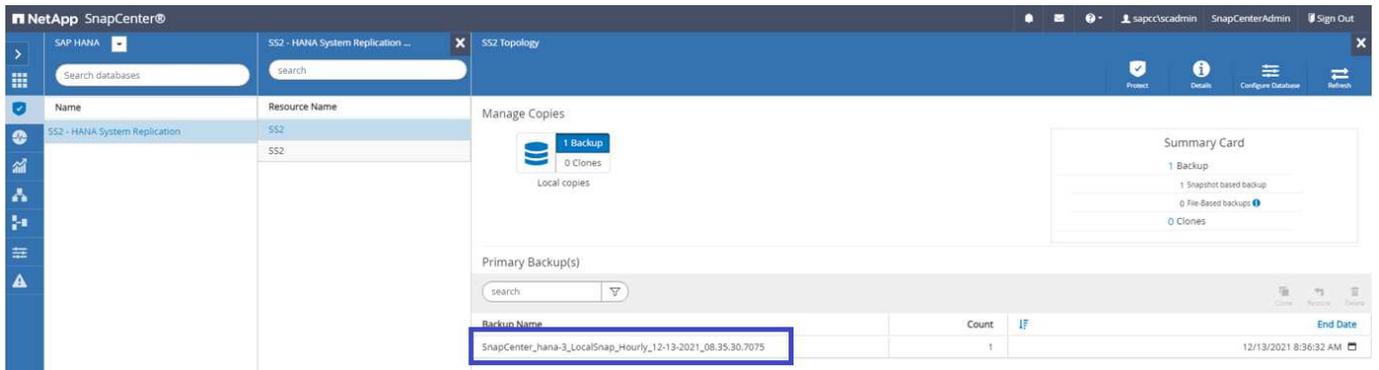
When a backup operation of the resource group is executed, SnapCenter identifies which host is primary and only triggers a backup at the primary host. This means, only the data volume of the primary host will be snapshotted. In our example, hana-3 is the current primary host and a backup is executed at this host.



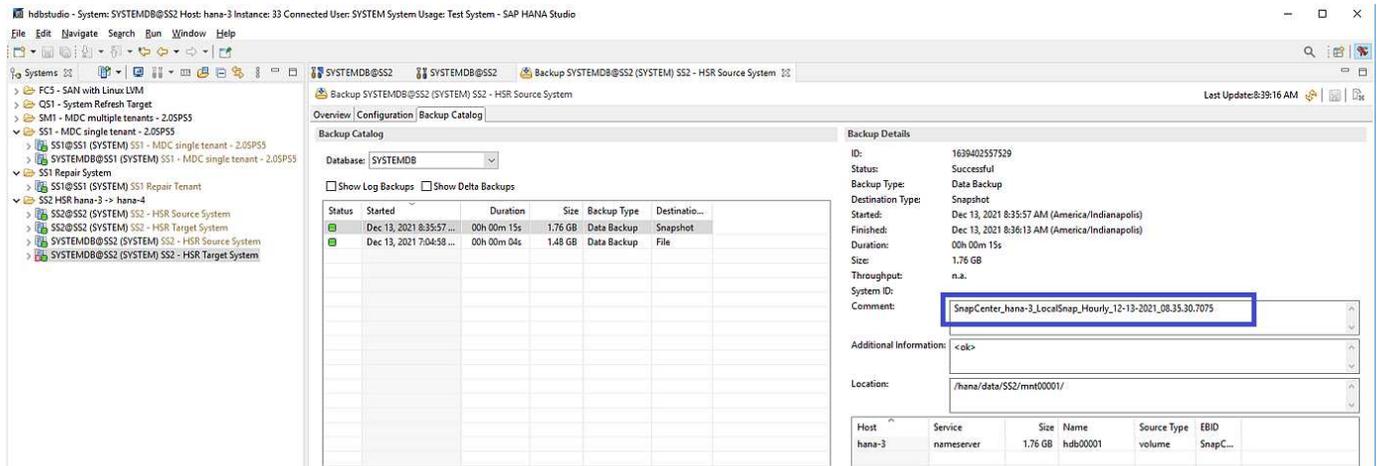
The SnapCenter job log shows the identification operation and the execution of the backup at the current primary host hana-3.



A Snapshot backup has now been created at the primary HANA resource. The hostname included in the backup name shows hana-3.



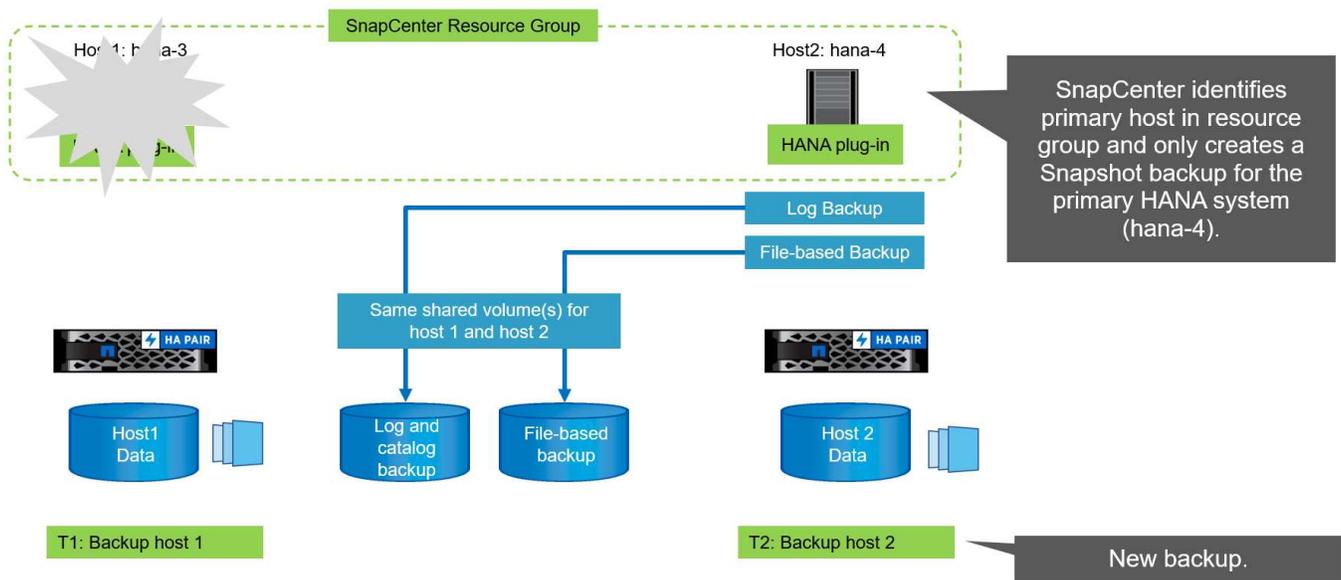
The same Snapshot backup is also visible in the HANA backup catalog.



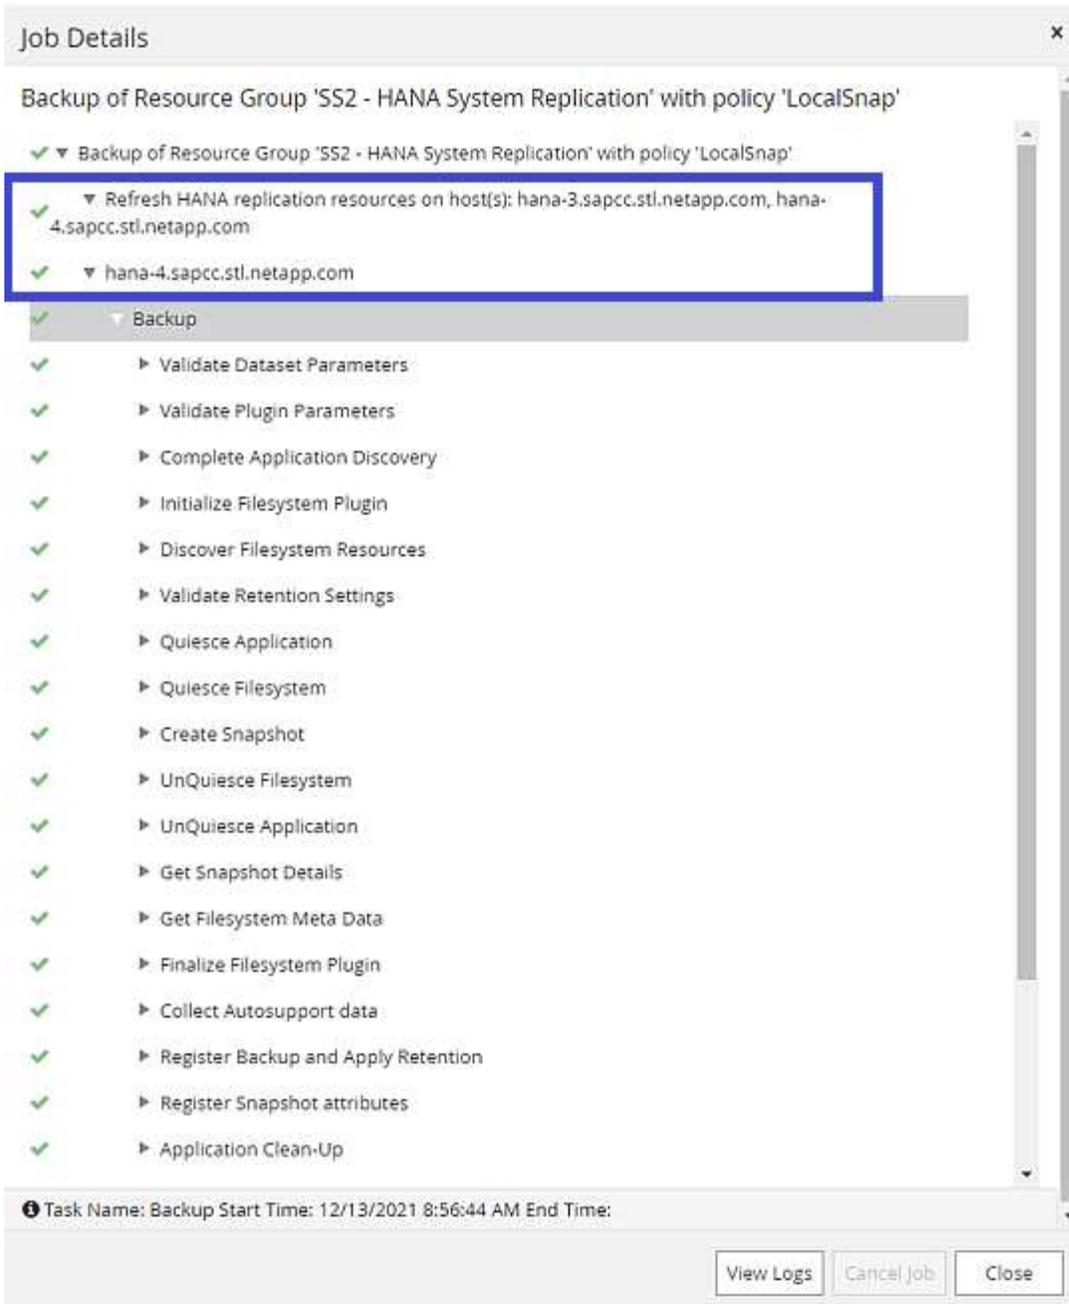
If a takeover operation is executed, further SnapCenter backups now identify the former secondary host (hana-4) as primary, and the backup operation is executed at hana-4. Again, only the data volume of the new primary host (hana-4) is snapshotted.



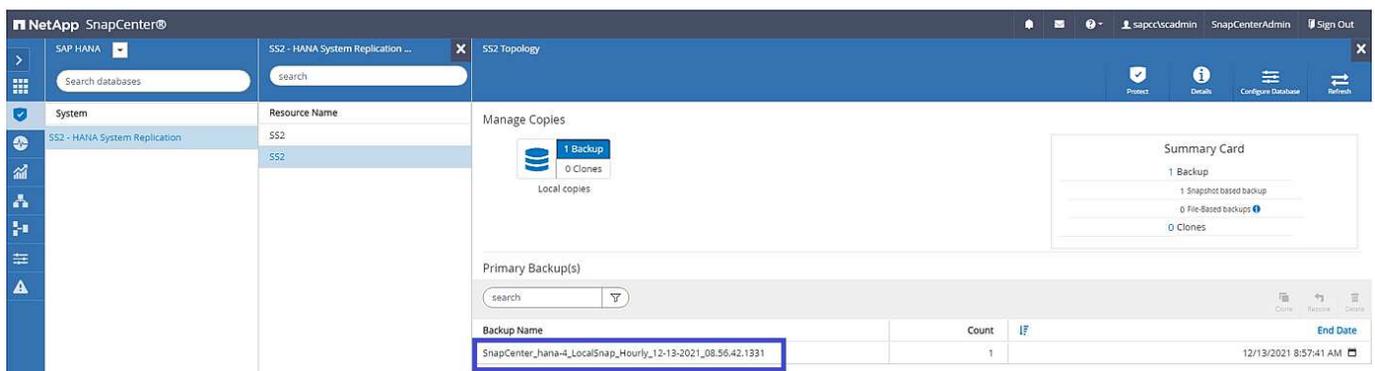
The SnapCenter identification logic only covers scenarios in which the HANA hosts are in a primary-secondary relation or when one of the HANA hosts is offline.



The SnapCenter job log shows the identification operation and the execution of the backup at the current primary host hana-4.



A Snapshot backup has now been created at the primary HANA resource. The hostname included in the backup name shows hana-4.



The same Snapshot backup is also visible in the HANA backup catalog.

The screenshot shows the SAP HANA Studio interface. The 'Backup Catalog' tab is active, displaying a table of backup operations for the 'SYSTEMDB' database. The table has columns for Status, Started, Duration, Size, Backup Type, and Destination. A 'Snapshot' backup is highlighted in blue. To the right, the 'Backup Details' pane shows information for the selected backup, including ID, Status (Successful), Backup Type (Data Backup), Destination Type (Snapshot), and Start/Finish times. The 'Comment' field contains the text 'SnapCenter\_hana-4\_LocalSnap\_Hourly\_12-13-2021\_08:56:42.1331'.

Status	Started	Duration	Size	Backup Type	Destination...
Success	Dec 13, 2021 8:57:07...	00h 00m 15s	1.69 GB	Data Backup	Snapshot
Success	Dec 13, 2021 8:50:40...	00h 00m 14s	1.76 GB	Data Backup	Snapshot
Success	Dec 13, 2021 8:43:45...	00h 00m 04s	1.48 GB	Data Backup	File
Success	Dec 13, 2021 7:04:58...	00h 00m 04s	1.48 GB	Data Backup	File

### Block-integrity check operations with file-based backups

SnapCenter 4.6 uses the same logic as described for Snapshot backup operations for block-integrity check operations with file-based backups. SnapCenter identifies the current primary HANA host and executes the file-based backup for this host. Retention management is also performed across both hosts, so the oldest backup is deleted regardless of which host is currently the primary.

### SnapVault replication

To allow transparent backup operations without manual interaction in case of a takeover and independent of which HANA host is currently the primary host, you must configure a SnapVault relationship for the data volumes of both hosts. SnapCenter executes a SnapVault update operation for the current primary host with each backup run.

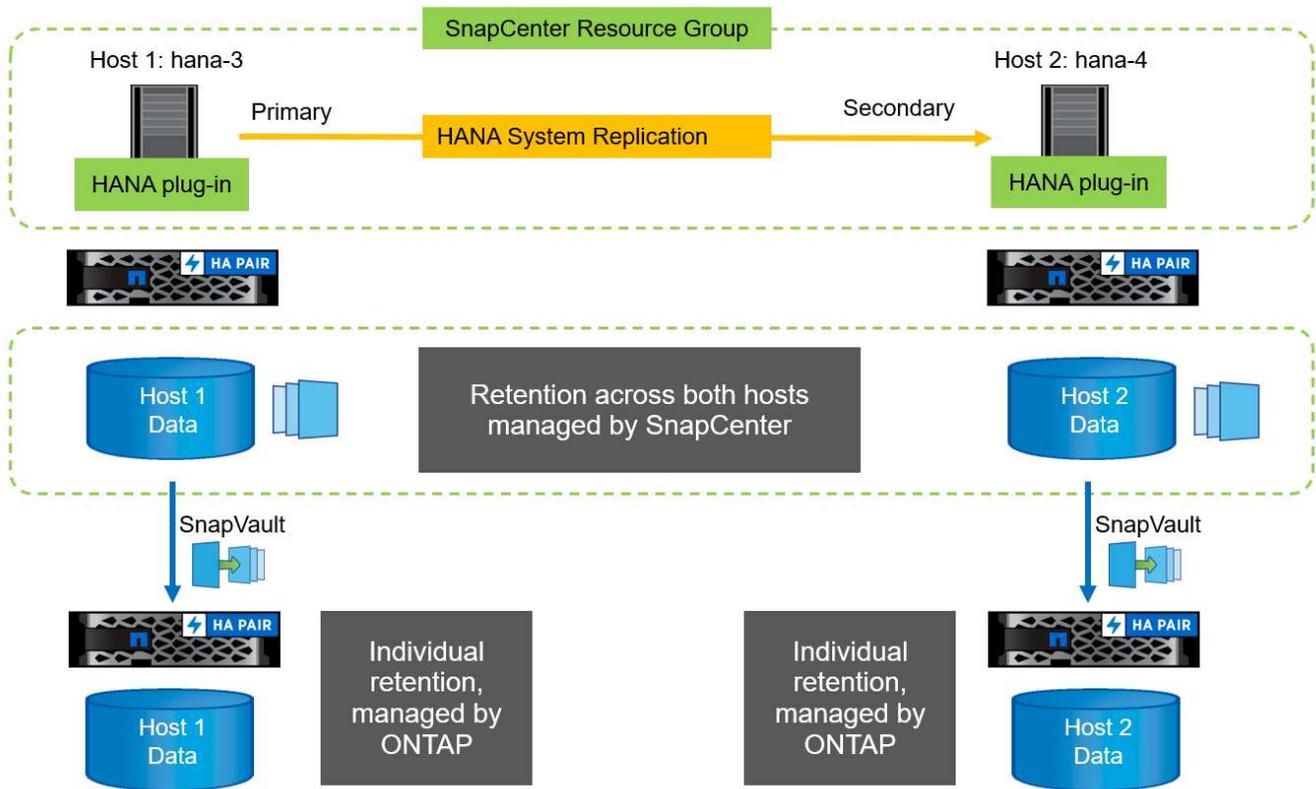


If a takeover to the secondary host is not performed for a long time, the number of changed blocks for the first SnapVault update at the secondary host will be high.

Since the retention management at the SnapVault target is managed outside of SnapCenter by ONTAP, the retention can't be handled across both HANA hosts. Therefore backups that have been created before a takeover are not deleted with backup operations at the former secondary. These backups remain until the former primary becomes primary again. So that these backups do not block the retention management of log backups, they must be deleted manually either at the SnapVault target or within the HANA backup catalog.



A cleanup of all SnapVault Snapshot copies is not possible, because one Snapshot copy is blocked as a synchronization point. If the latest Snapshot copy needs to be deleted as well, the SnapVault replication relationship must be deleted. In this case, NetApp recommends deleting the backups in the HANA backup catalog to unblock log backup retention management.



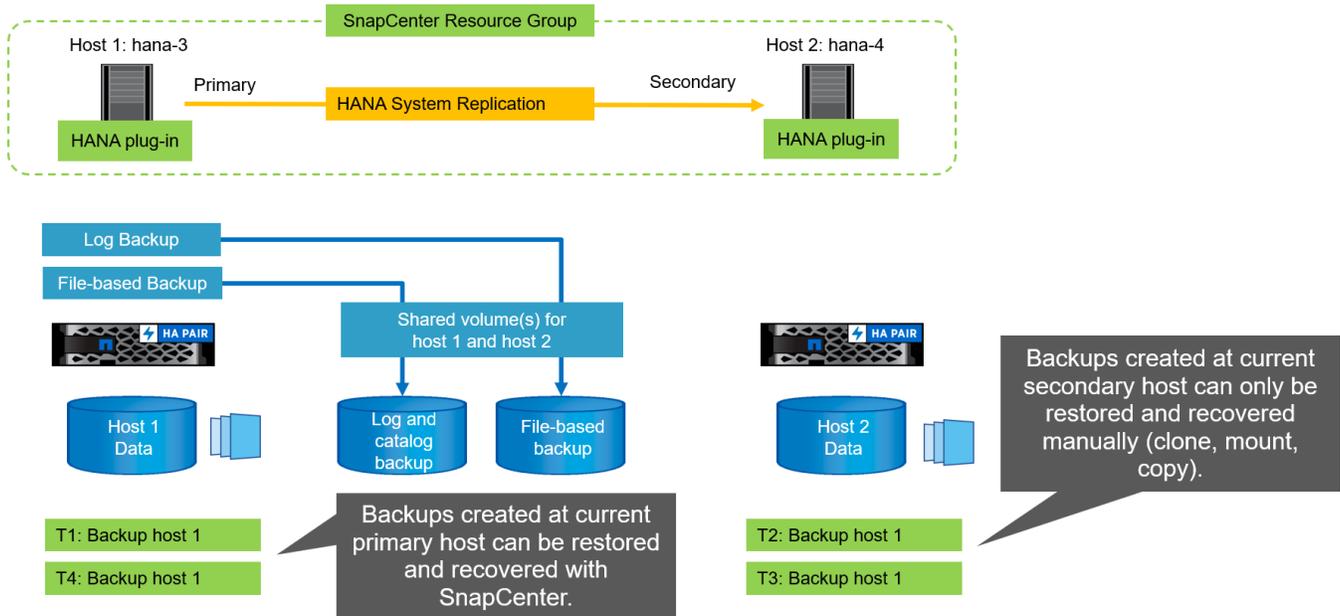
## Retention management

SnapCenter 4.6 manages retention for Snapshot backups, block-integrity check operations, HANA backup catalog entries, and log backups (if not disabled) across both HANA hosts, so it doesn't matter which host is currently primary or secondary. Backups (data and log) and entries in the HANA catalog are deleted based on the defined retention, regardless of whether a delete operation is necessary on the current primary or secondary host. In other words, no manual interaction is required if a takeover operation is performed and/or the replication is configured in the other direction.

If SnapVault replication is part of the data protection strategy, manual interaction is required for specific scenarios, as described in the section [SnapVault Replication](#)

## Restore and recovery

The following figure depicts a scenario in which multiple takeovers have been executed and Snapshot backups have been created at both sites. With the current status, the host hana-3 is the primary host and the latest backup is T4, which has been created at host hana-3. If you need to perform a restore and recovery operation, the backups T1 and T4 are available for restore and recovery in SnapCenter. The backups, which have been created at host hana-4 (T2, T3), can't be restored using SnapCenter. These backups must be copied manually to the data volume of hana-3 for recovery.



Restore and recovery operations for a SnapCenter 4.6 resource group configuration are identical to an autodiscovered non-System Replication setup. All options for restore and automated recovery are available. For further details, see the technical report [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

A restore operation from a backup that was created at the other host is described in the section [Restore and Recovery from a Backup Created at the Other Host](#).

## SnapCenter configuration with a single resource

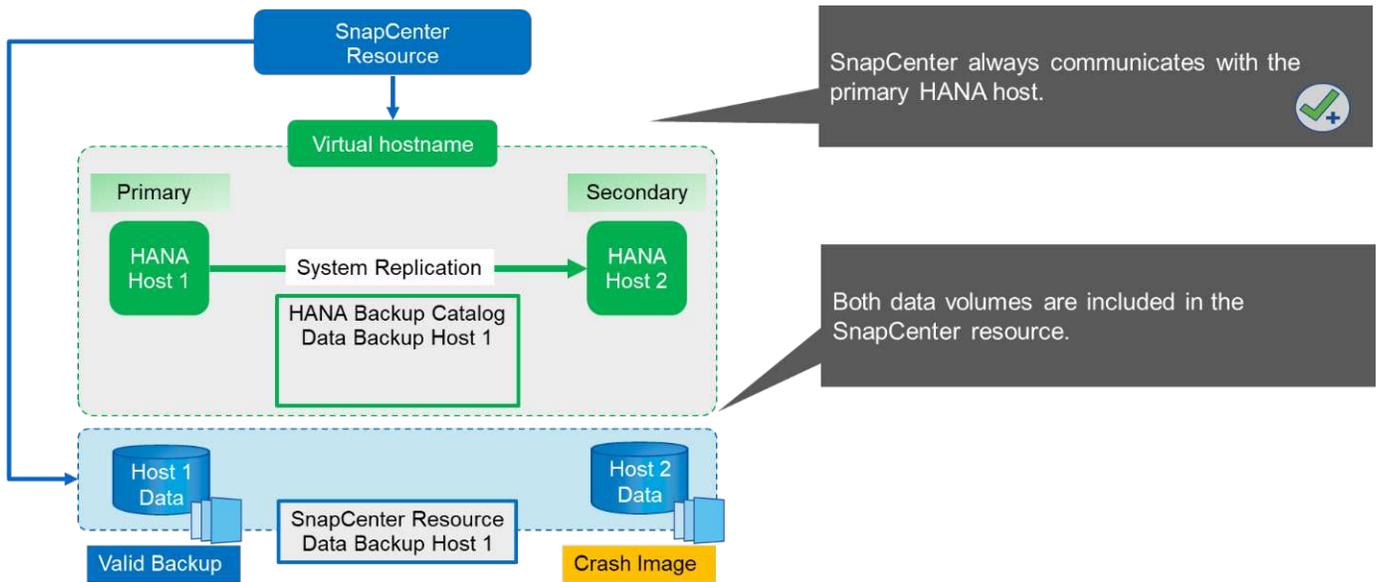
A SnapCenter resource is configured with the virtual IP address (host name) of the HANA System Replication environment. With this approach, SnapCenter always communicates with the primary host, regardless of whether host 1 or host 2 is primary. The data volumes of both SAP HANA hosts are included in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



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

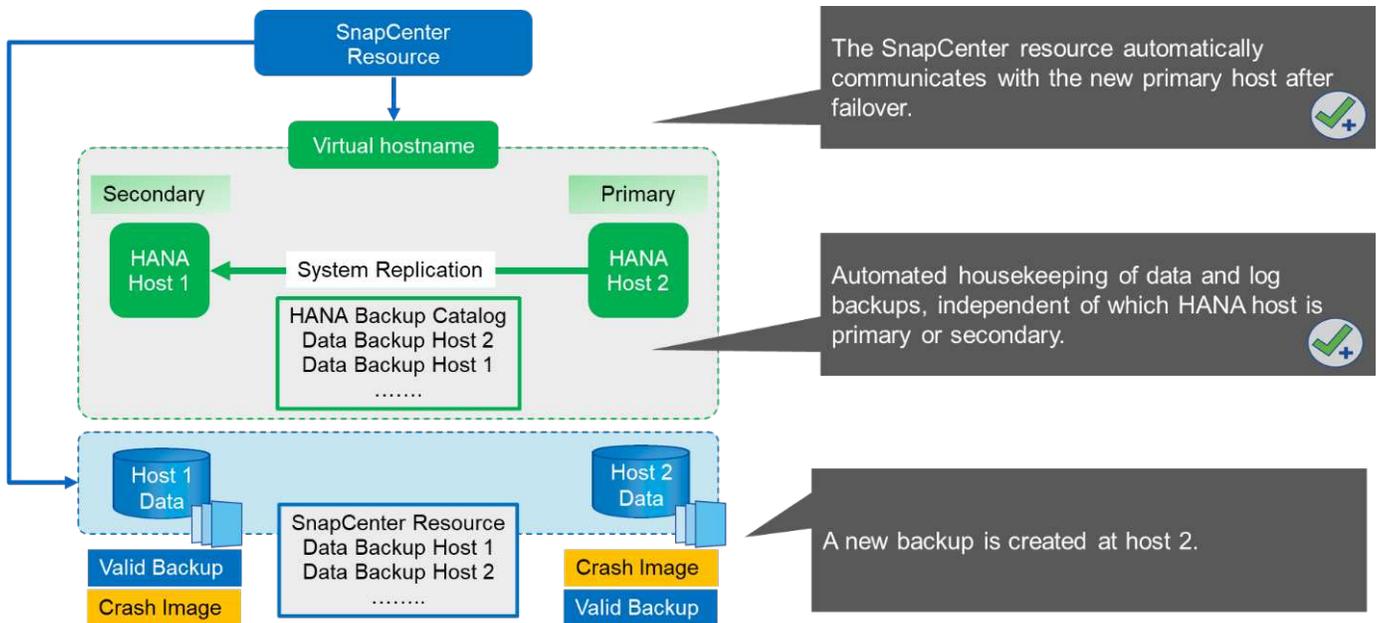
When a backup is executed with host 1 as the primary host, a database-consistent Snapshot backup is created at the data volume of host 1. Because the data volume of host 2 is part of the SnapCenter resource, another Snapshot copy is created for this volume. This Snapshot copy is not database consistent;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

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



The following figure shows the backup operation after failover to host 2 and replication from host 2 to host 1. SnapCenter automatically communicates with host 2 by using the virtual IP address configured in the SnapCenter resource. Backups are now created at host 2. Two Snapshot copies are created by SnapCenter: a database-consistent backup at the data volume at host 2 and a crash image Snapshot copy at the data volume at host 1. The SAP HANA backup catalog and the SnapCenter resource now include the backup created at host 1 and the backup created at host 2.

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

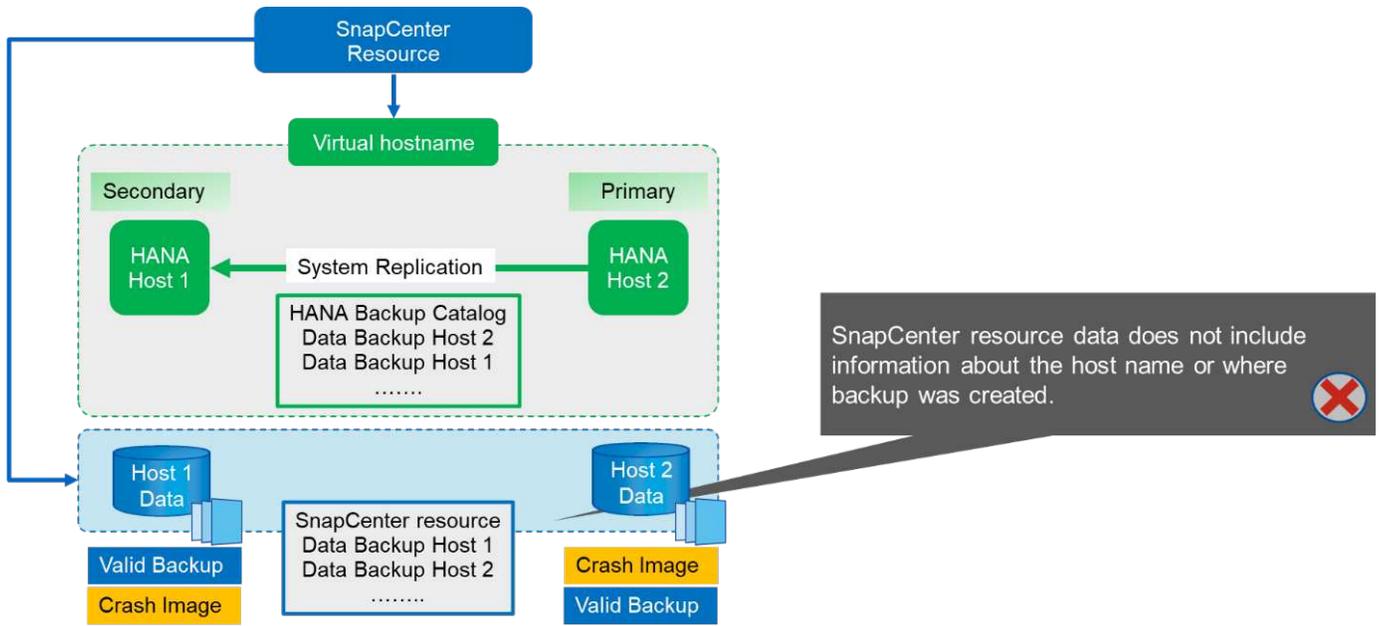


As discussed in the section [Storage Snapshot Backups and SAP System Replication](#), a restore operation with storage-based Snapshot backups is different, depending on which backup must be restored. It is important to identify which host the backup was created at to determine if the restore can be performed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

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

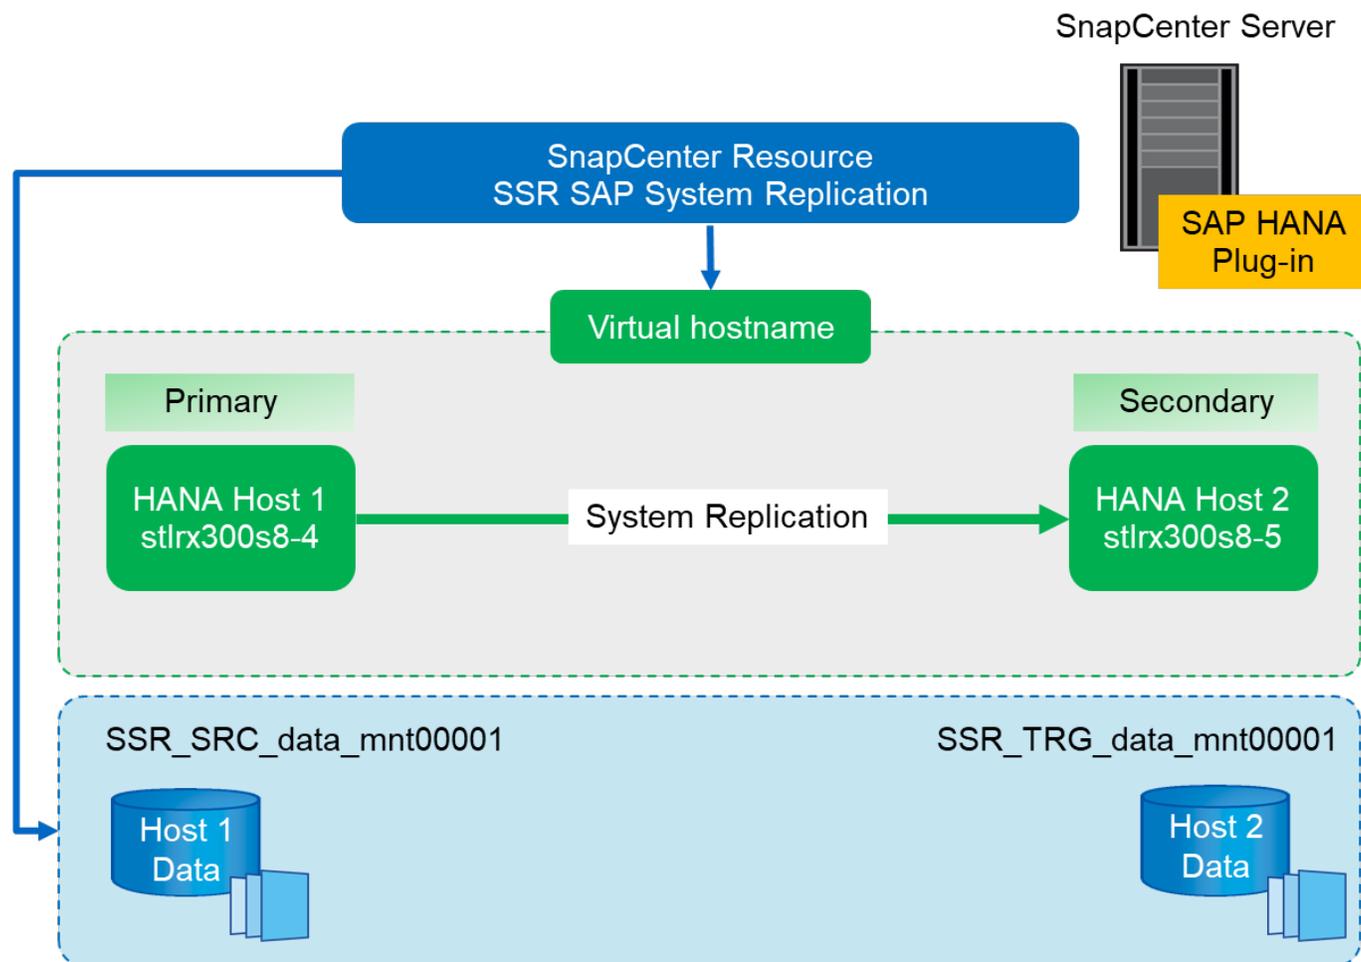
identify which host is currently the primary SAP HANA host.

The following figure depicts identification of the backup host.



### SnapCenter configuration

The following figure shows the lab setup and an overview of the required SnapCenter configuration.



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

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

```
hdbuserstore.exe -u SYSTEM set SSRKEY ssr-vip:31013 SNAPCENTER <password>
```

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

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

Add SAP HANA Database
✕

1 Name

2 Storage Footprint

3 Summary

### Provide Resource Details

Resource Type

Single Container

Multitenant Database Container (MDC) - Single Tenant

Non-data Volumes

HANA System Name

SID

Tenant Database

HDBSQL Client Host

HDB Secure User Store Keys

HDBSQL OS User

Previous
Next

The data volumes of both SAP HANA hosts are included in the storage footprint configuration, as the following figure shows.

x
Add SAP HANA Database

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

### Provide Storage Footprint Details

Storage Systems for storage footprint hana

Modify hana
x

Select one or more volumes and if required their associated Qtrees and LUNs

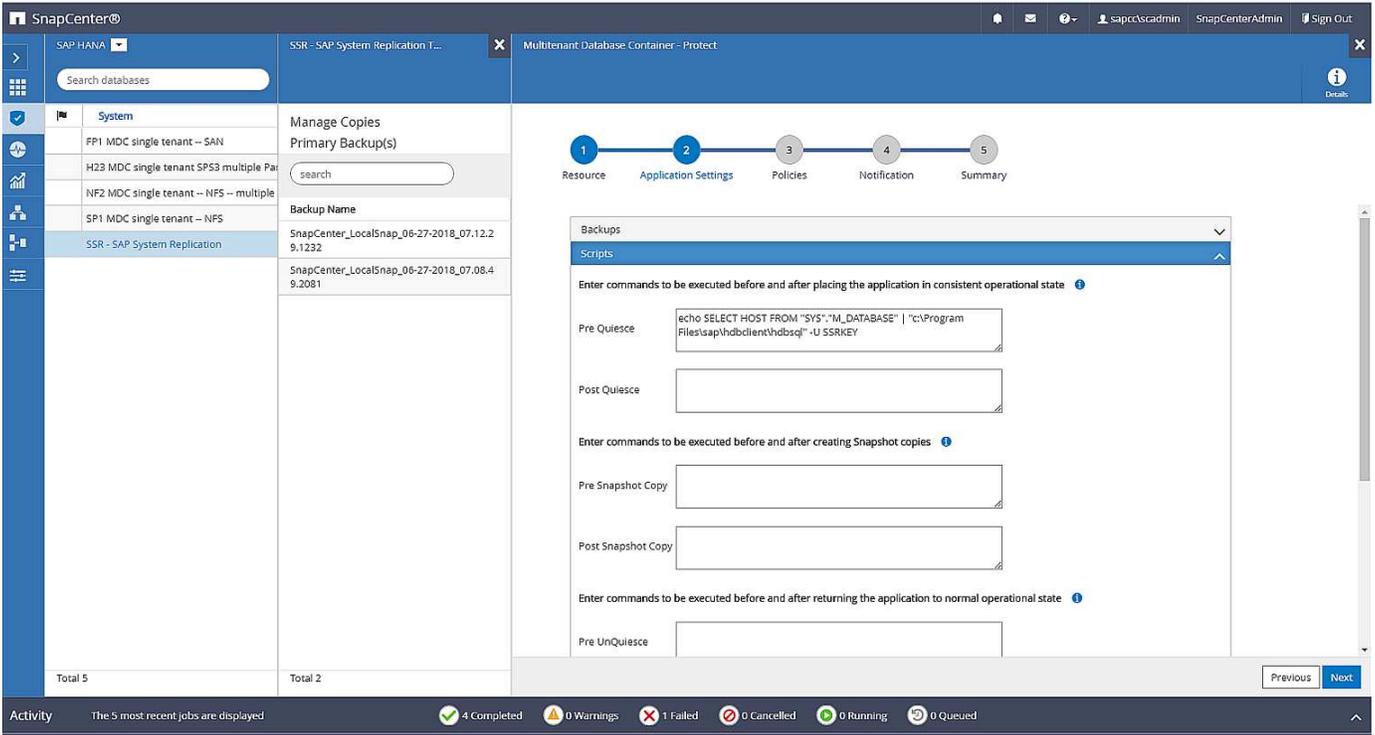
Volume Name	LUNs or Qtrees
<span style="border: 1px solid #ccc; padding: 2px;">SSR_TRG_data_mnt00001</span> ▼	<span style="border: 1px solid #ccc; padding: 2px;">Default is 'None' or type to find</span>
<span style="border: 1px solid #ccc; padding: 2px;">SSR_SRC_data_mnt00001</span> ▼	<span style="border: 1px solid #ccc; padding: 2px;">Default is 'None' or type to find</span>

Save

Previous
Next

As discussed before, SnapCenter is not aware of where the backup was created. NetApp therefore recommends that you add a pre- backup script in the SnapCenter backup workflow to identify which host is currently the primary SAP HANA host. You can perform this identification using a SQL statement that is added to the backup workflow, as the following figure shows.

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

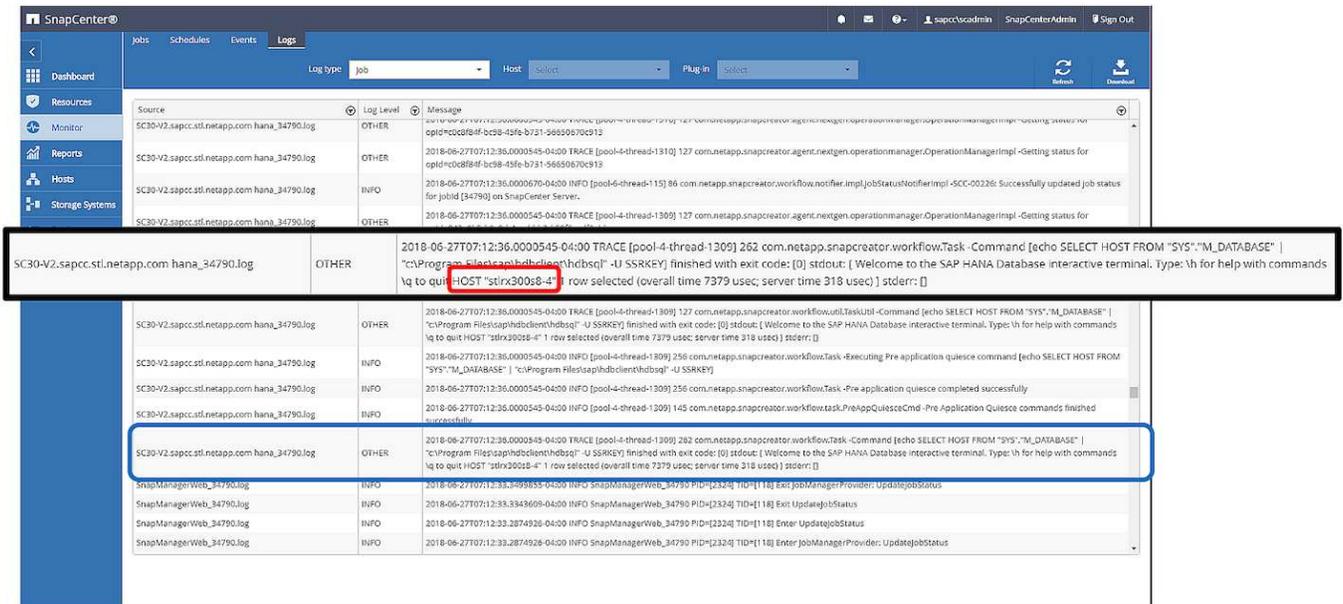


## SnapCenter backup operation

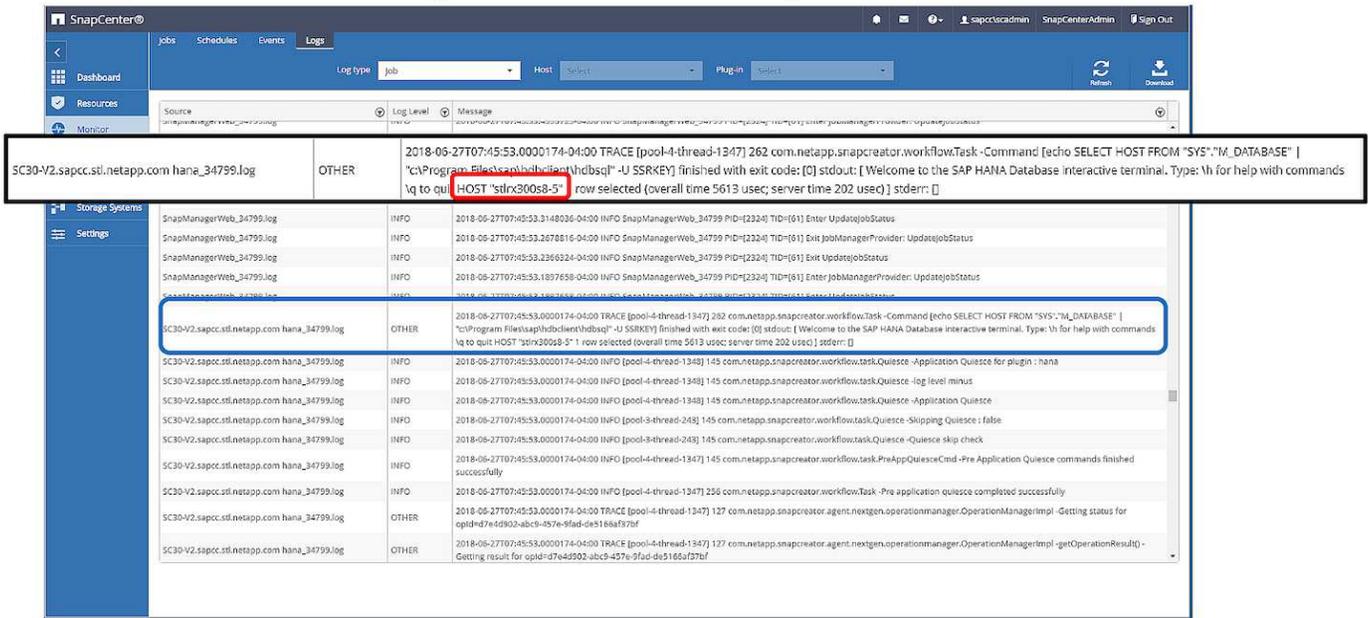
Backup operations are now executed as usual. Housekeeping of data and log backups is performed independent of which SAP HANA host is primary or secondary.

The backup job logs include the output of the SQL statement, which allows you to identify the SAP HANA host where the backup was created.

The following figure shows the backup job log with host 1 as the primary host.



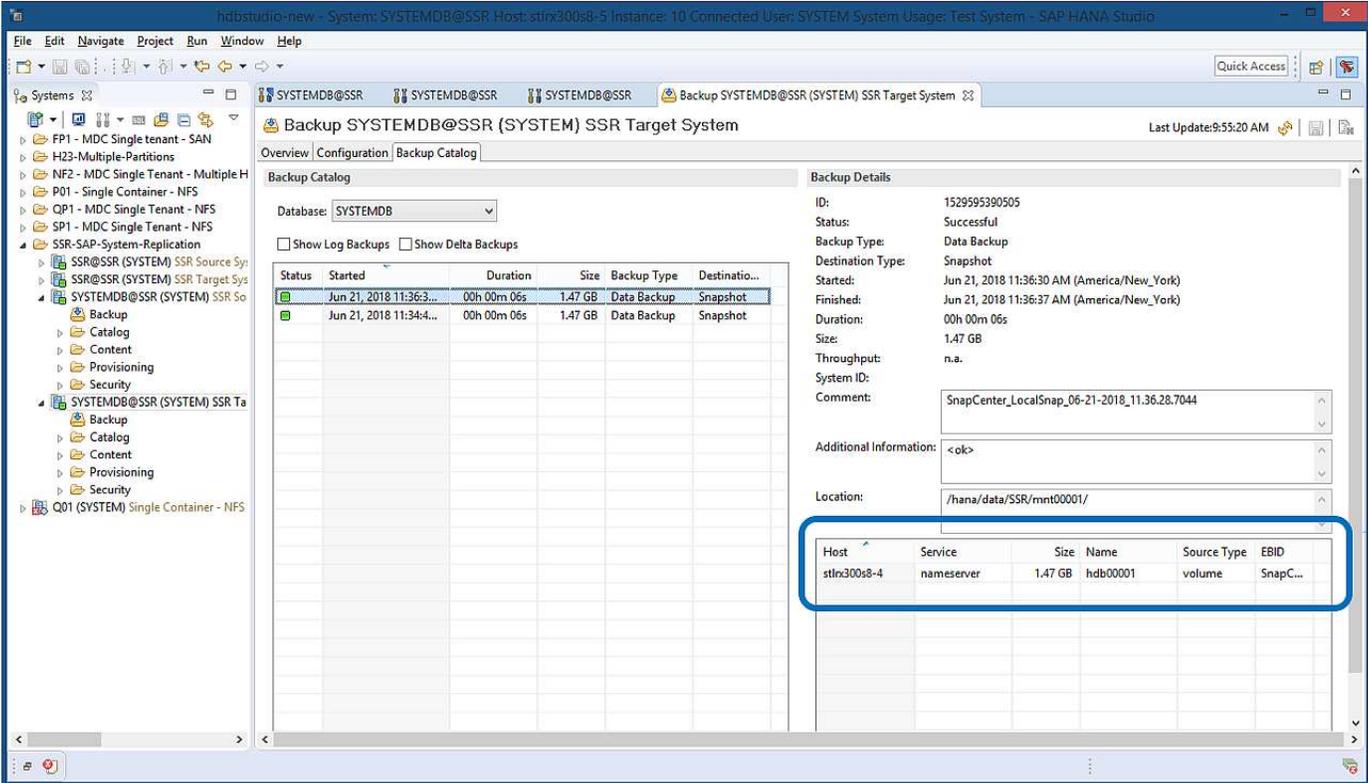
This figure shows the backup job log with host 2 as the primary host.



The following figure shows the SAP HANA backup catalog in SAP HANA Studio. When the SAP HANA database is online, the SAP HANA host where the backup was created is visible in SAP HANA Studio.



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



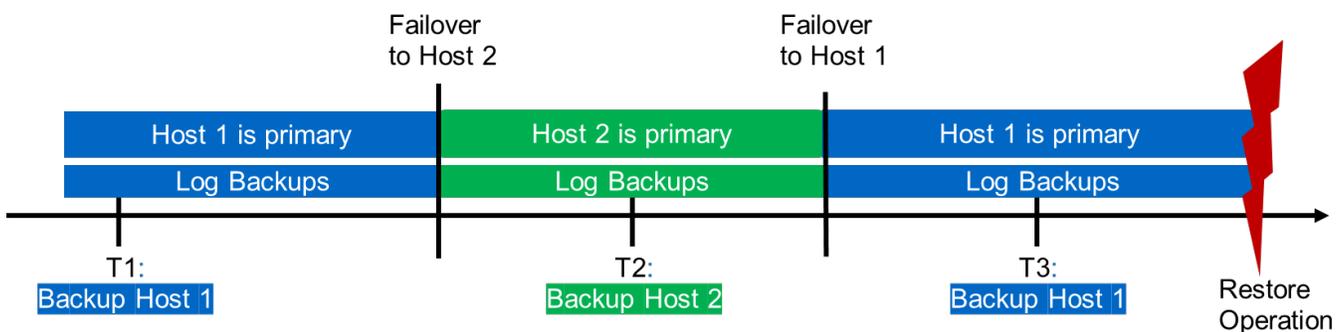
## Restore and recovery

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

The following figure illustrates the different restore operations depending on the selected backup.

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

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



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

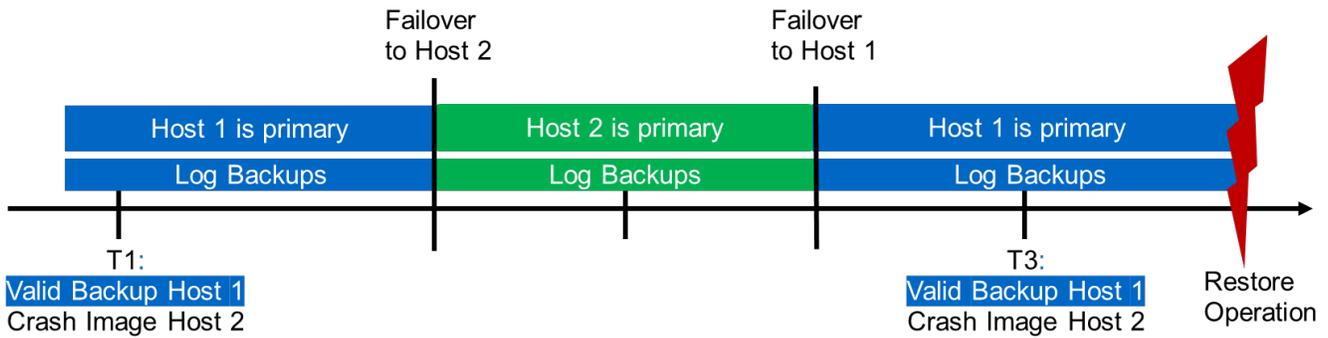
With a single SnapCenter resource configuration, Snapshot copies are created at both storage volumes of both SAP HANA System Replication hosts. Only the Snapshot backup that is created at the storage volume of the primary SAP HANA host is valid to use for forward recovery. The Snapshot copy created at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

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

- Restore only the valid backup
  - Restore the complete resource, including the valid backup and the crash image
- The following sections discuss the two different restore operations in more detail.

A restore operation from a backup that was created at the other host is described in the section [Restore and Recovery from a Backup Created at the Other Host](#).

The following figure depicts restore operations with a single SnapCenter resource configuration.

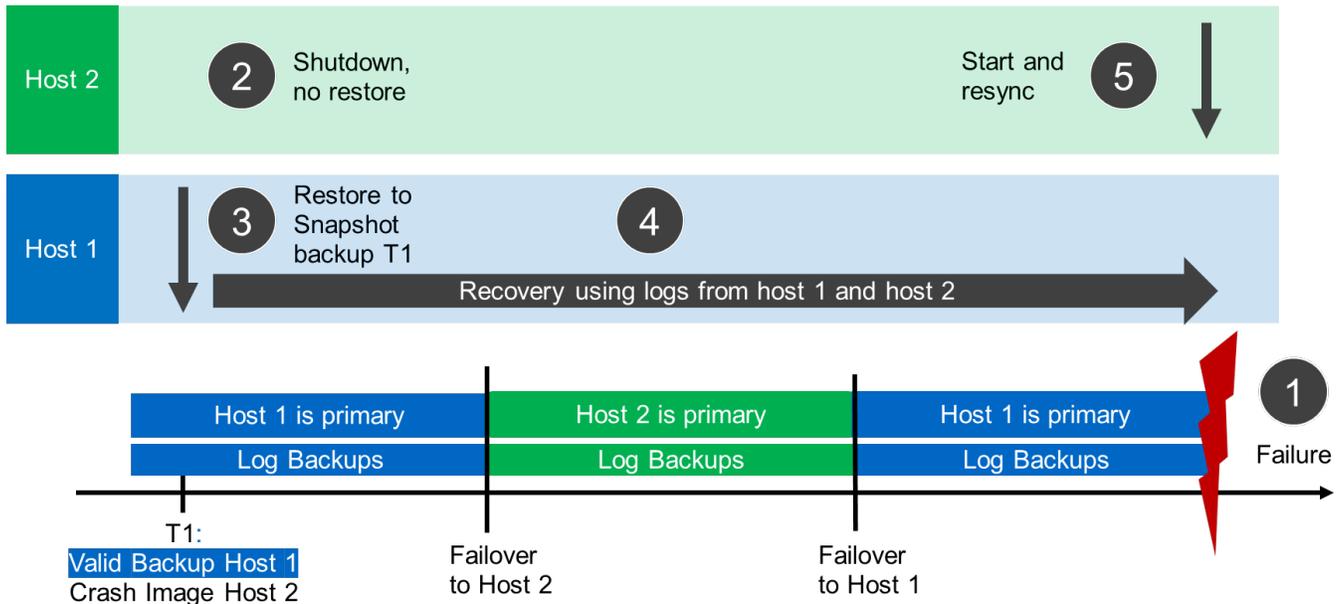


### SnapCenter restore of the valid backup only

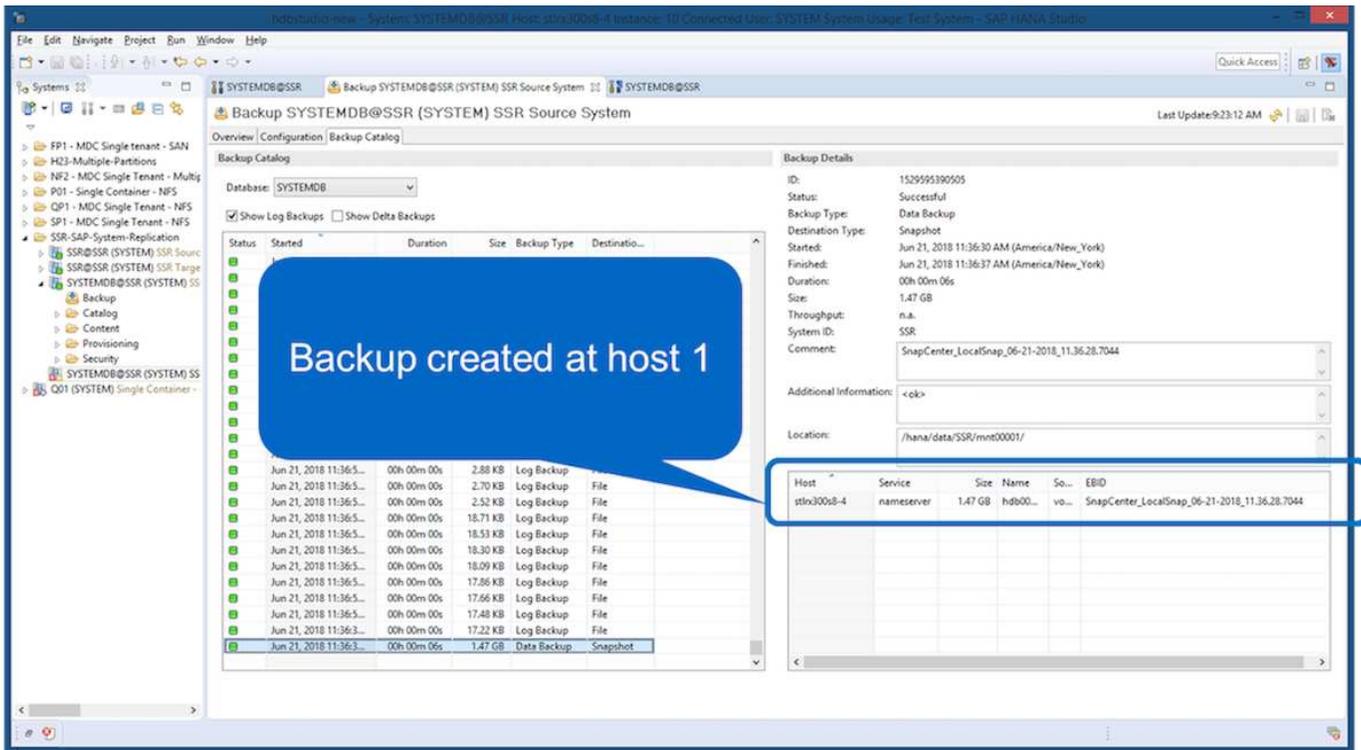
The following figure shows an overview of the restore and recovery scenario described in this section.

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

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



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

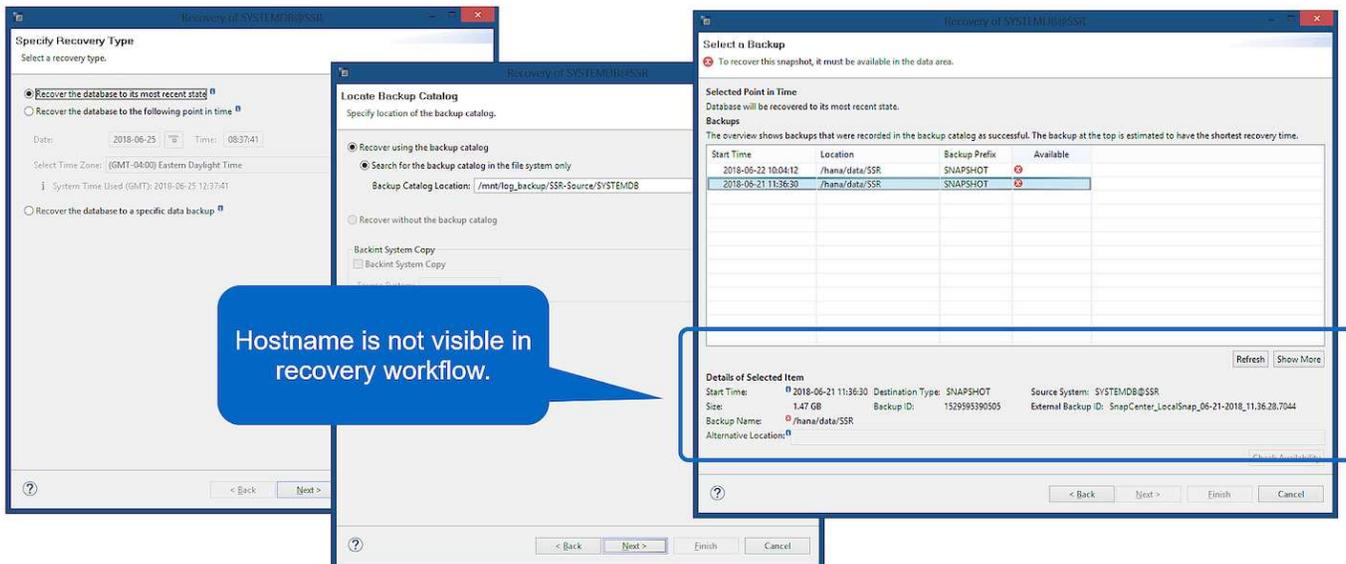


25

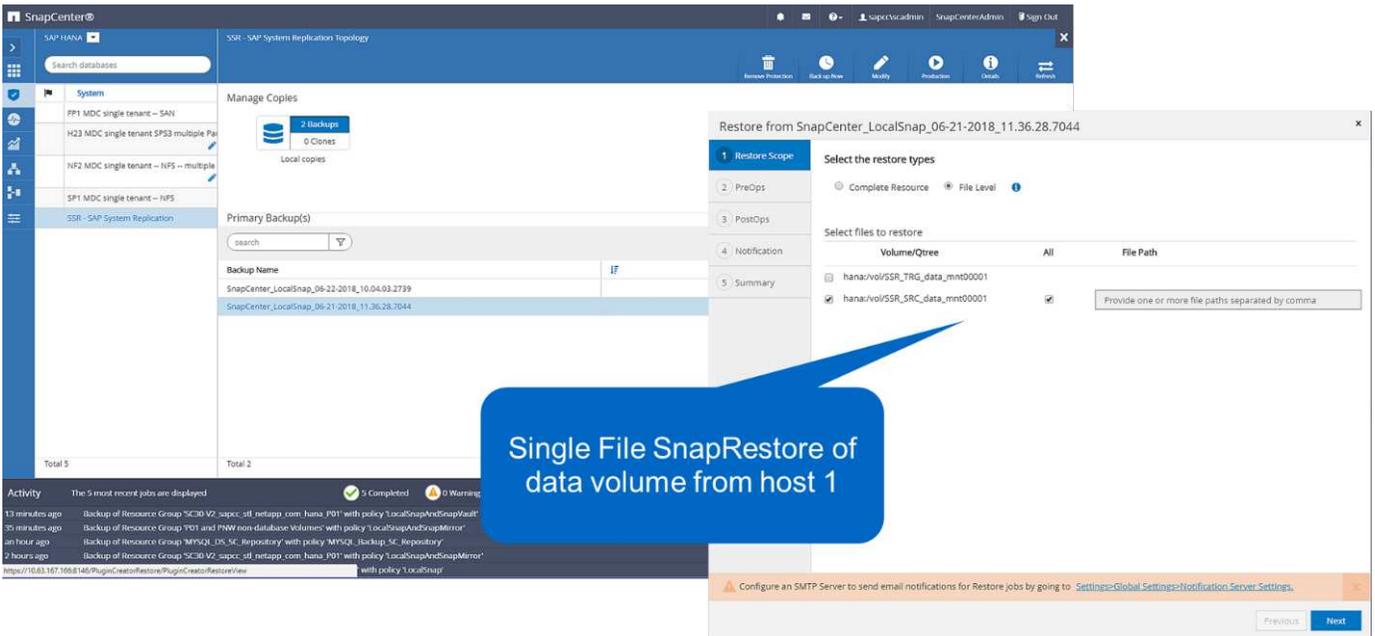
A restore and recovery operation is started in SAP HANA Studio. As the following figure shows, the name of the host where the backup was created is not visible in the restore and recovery workflow.



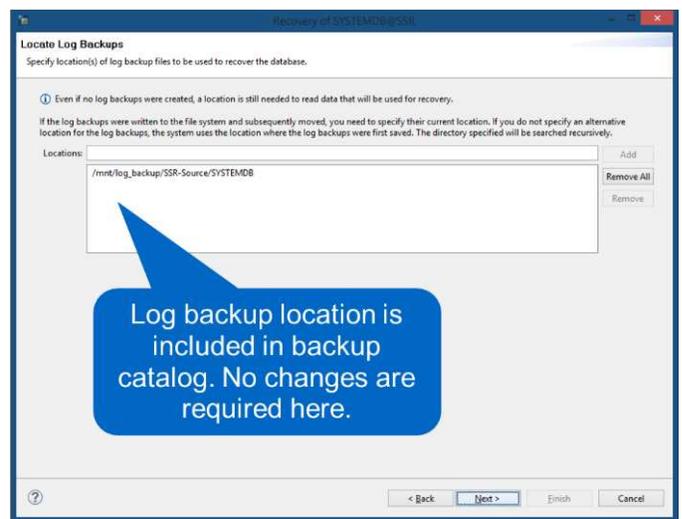
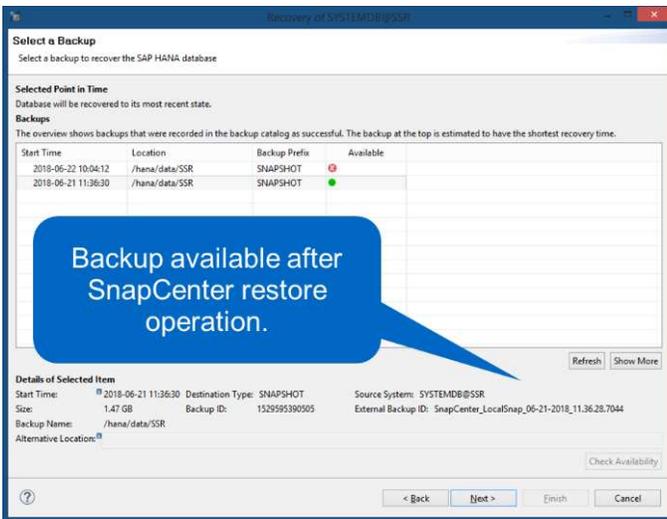
In our test scenario, we were able to identify the correct backup (the backup created at host 1) in SAP HANA Studio when the database was still online. If the database is not available, you must check the SnapCenter backup job log to identify the right backup.



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



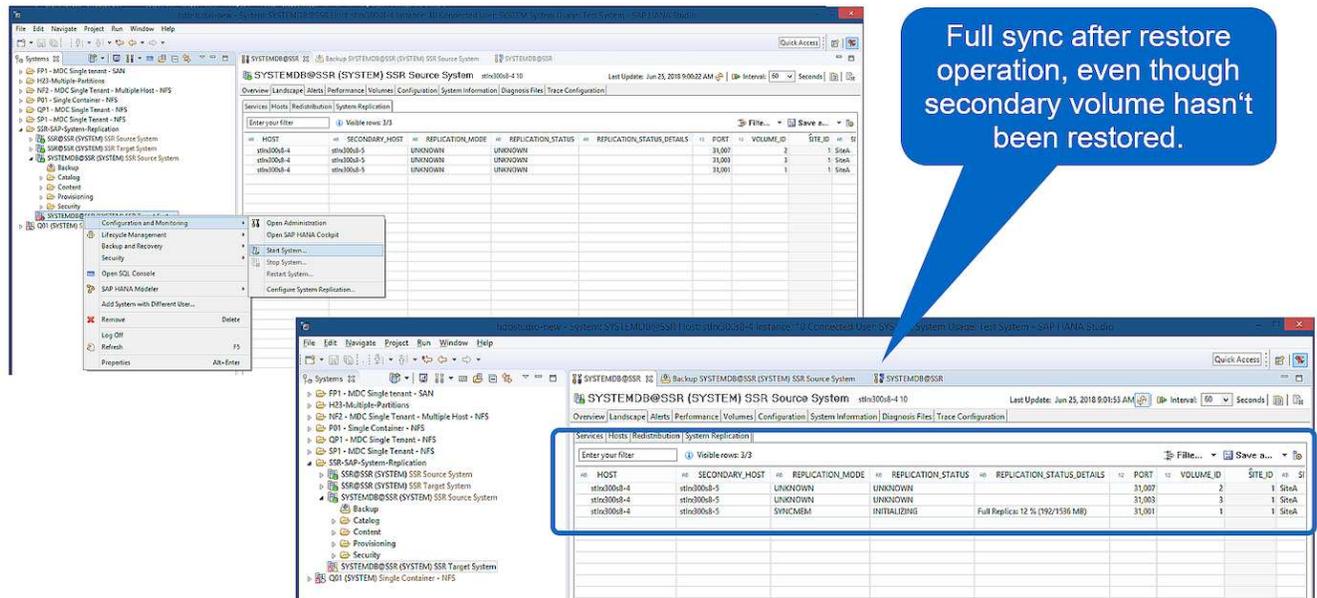
After the restore operation, the backup is highlighted in green in SAP HANA Studio. You don't have to enter an additional log backup location, because the file path of log backups of host 1 and host 2 are included in the backup catalog.



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



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

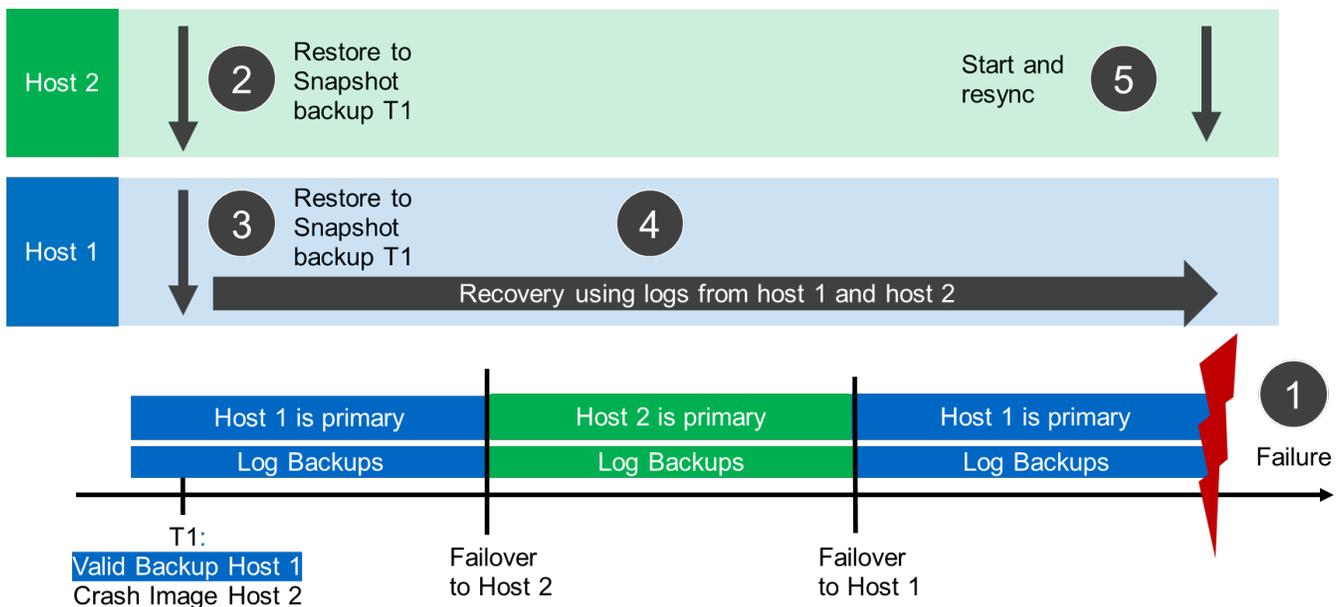


### SnapCenter restore of valid backup and crash image

The following figure shows an overview of the restore and recovery scenario described in this section.

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

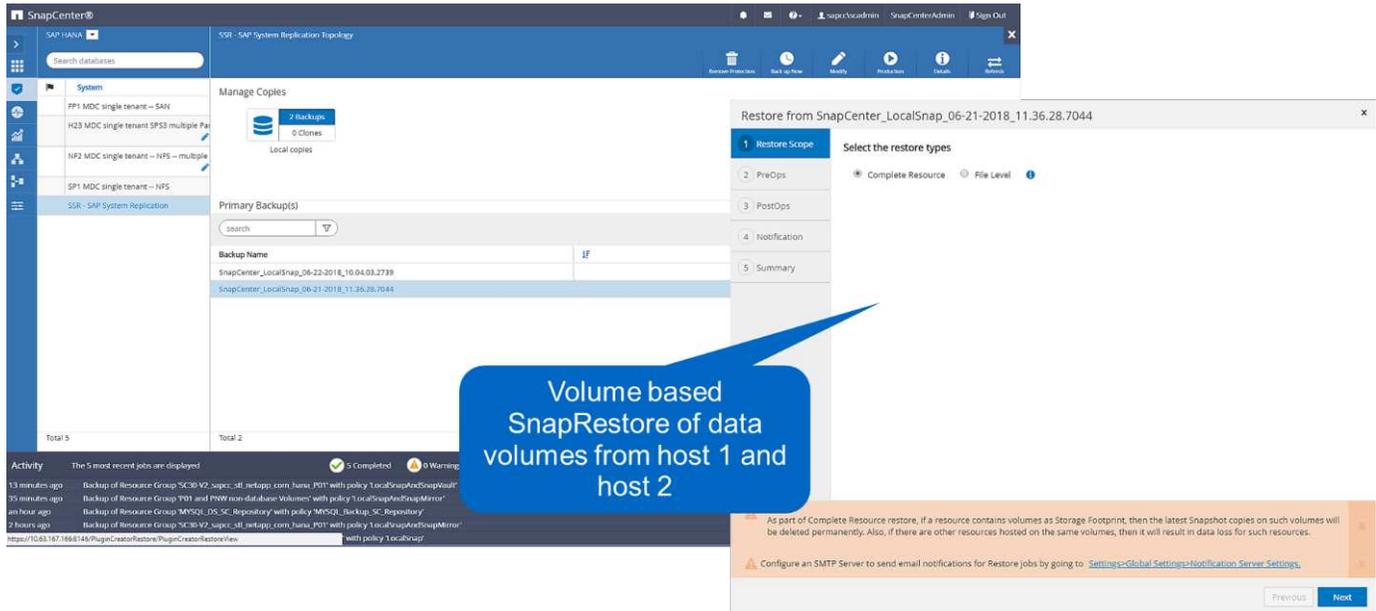
1. A failure occurred and you must restore the backup created at T1 at host 1.
2. The secondary host (host 2) is shut down and the T1 crash image is restored.
3. The storage volume of host 1 is restored to the backup created at T1.
4. A forward recovery is performed with logs from host 1 and host 2.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



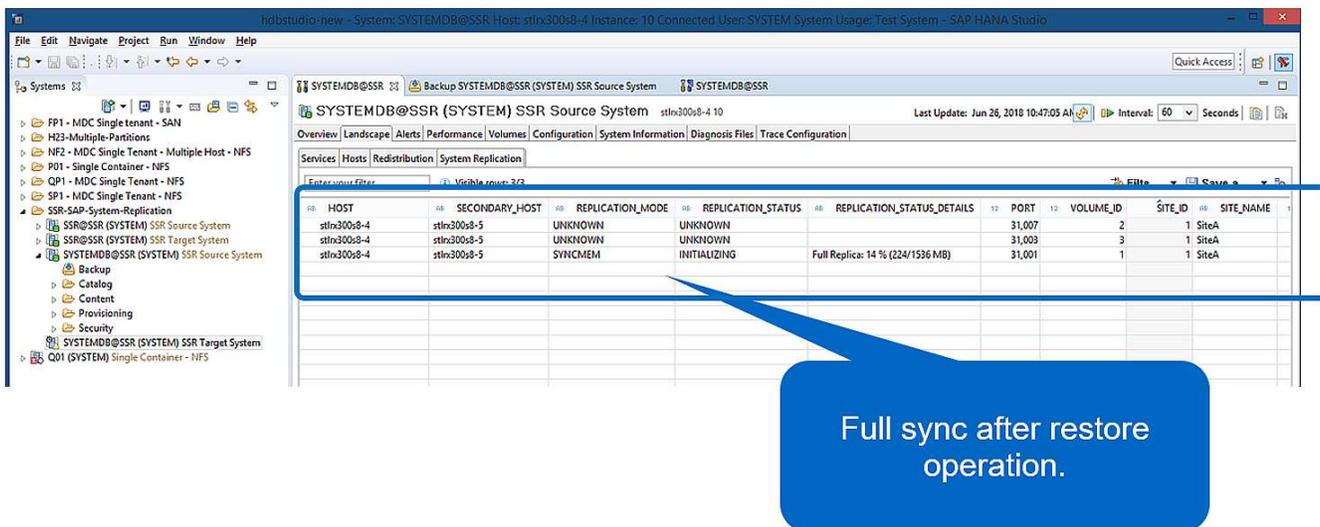
The restore and recovery operation with SAP HANA Studio is identical to the steps described in the section

## SnapCenter restore of the valid backup only.

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



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



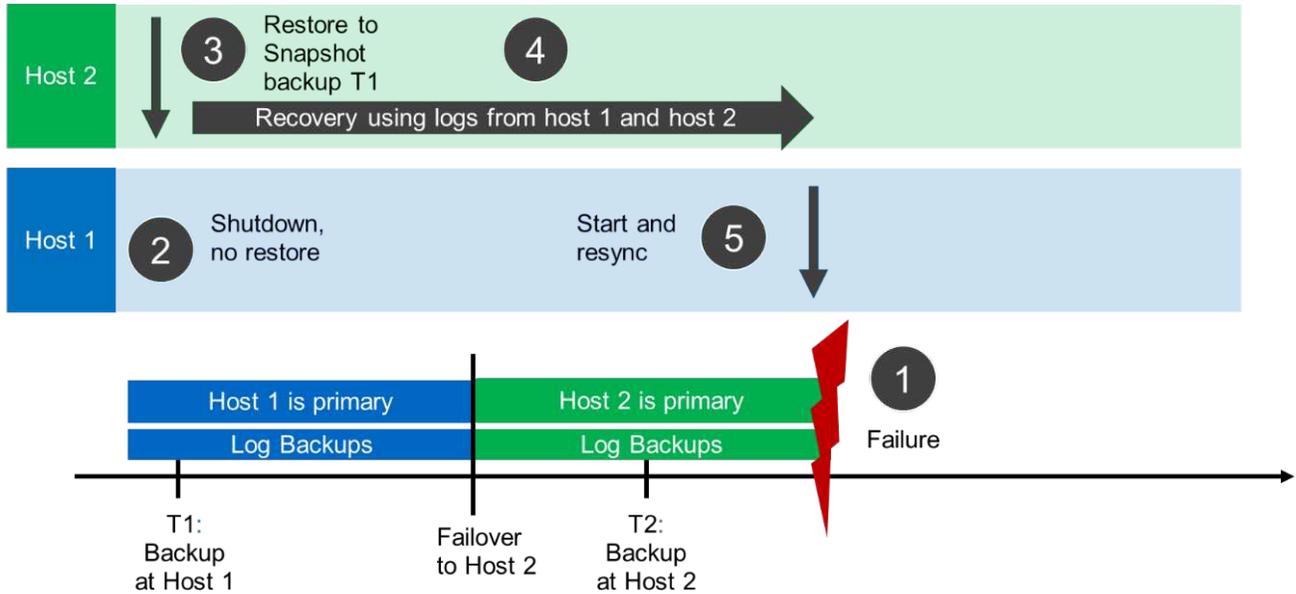
## Restore and recovery from a backup created at the other host

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

The following figure shows an overview of the restore and recovery scenario described in this section.

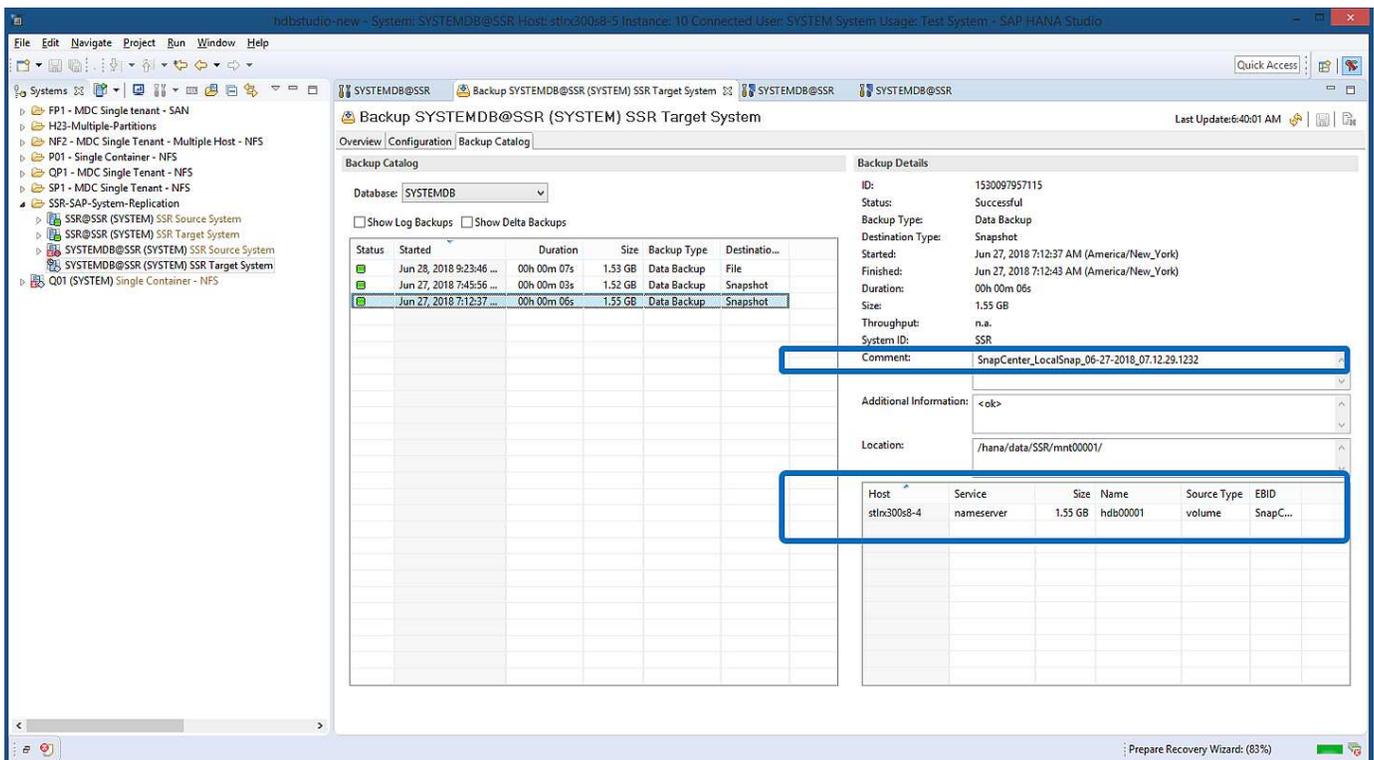
A backup has been created at T1 at host 1. A failover has been performed to host 2. At the current point in time, host 2 is the primary host.

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



31

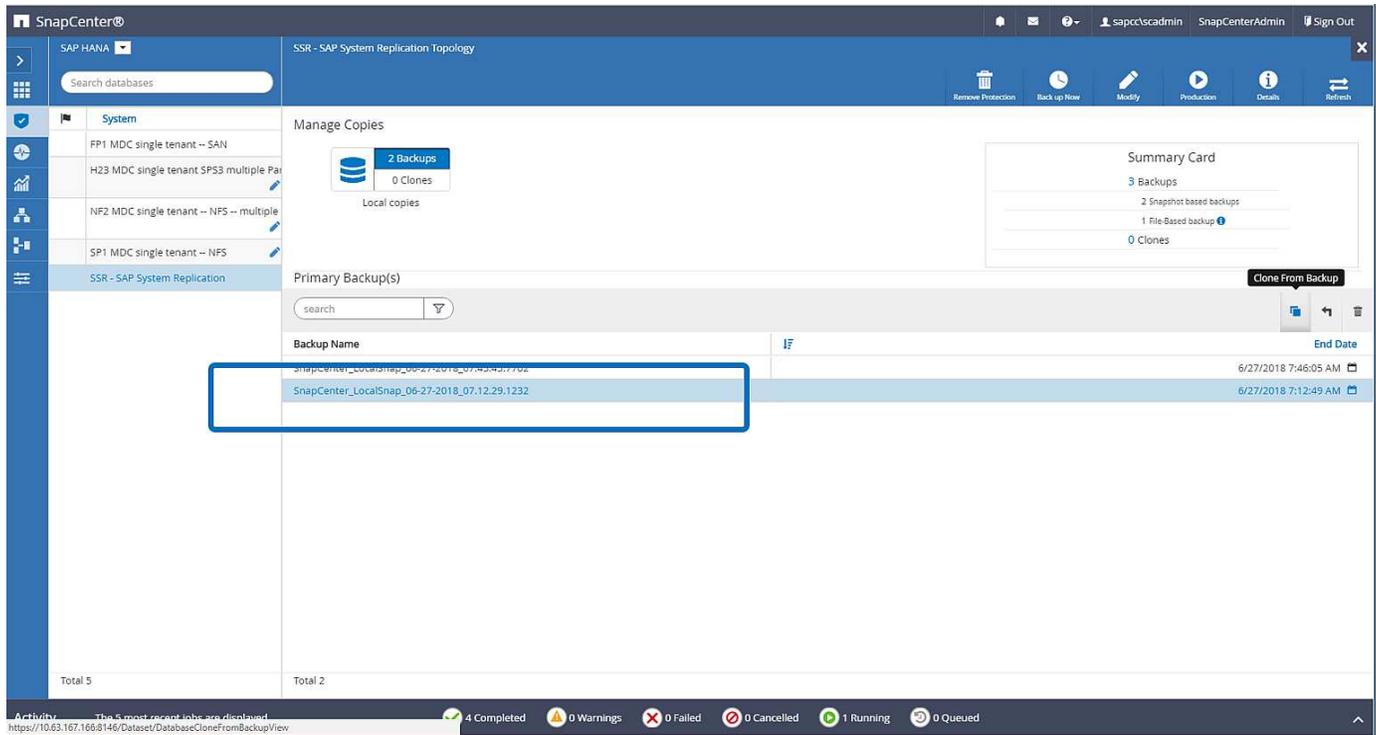
The following figure shows the SAP HANA backup catalog and highlights the backup, created at host 1, that was used for the restore and recovery operation.



The restore operation involves the following steps:

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

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



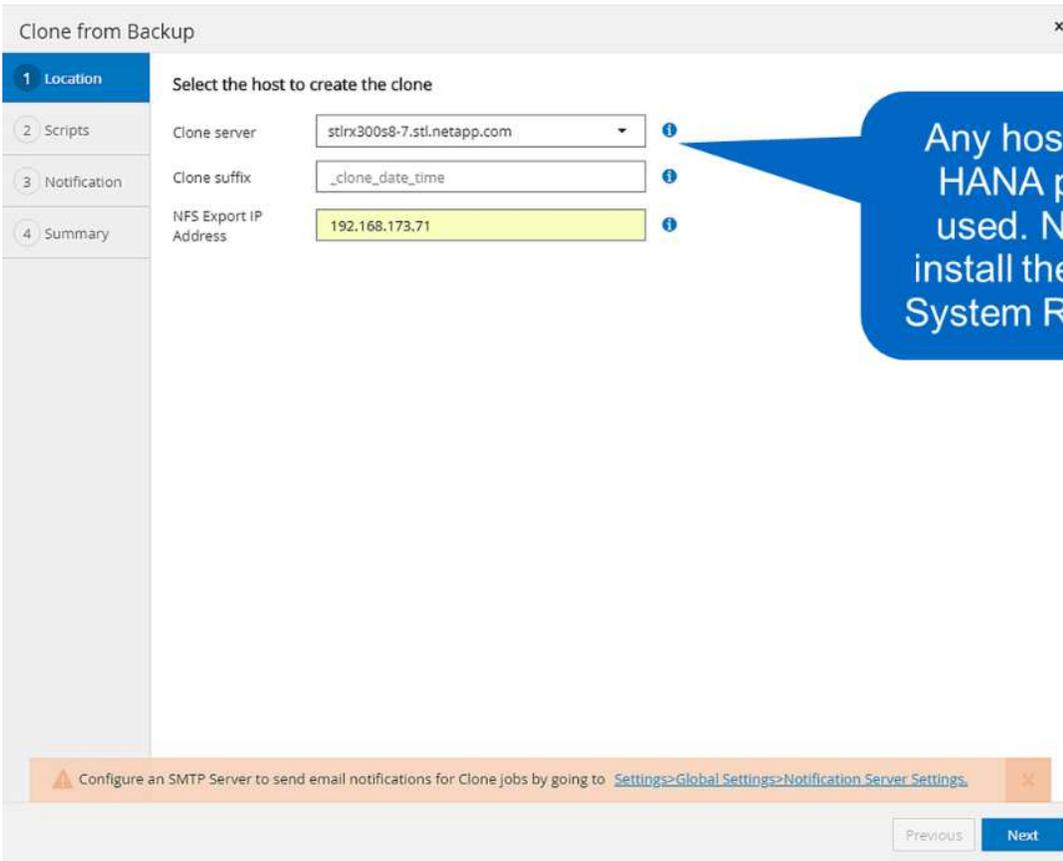
You must provide the clone server and the NFS export IP address.



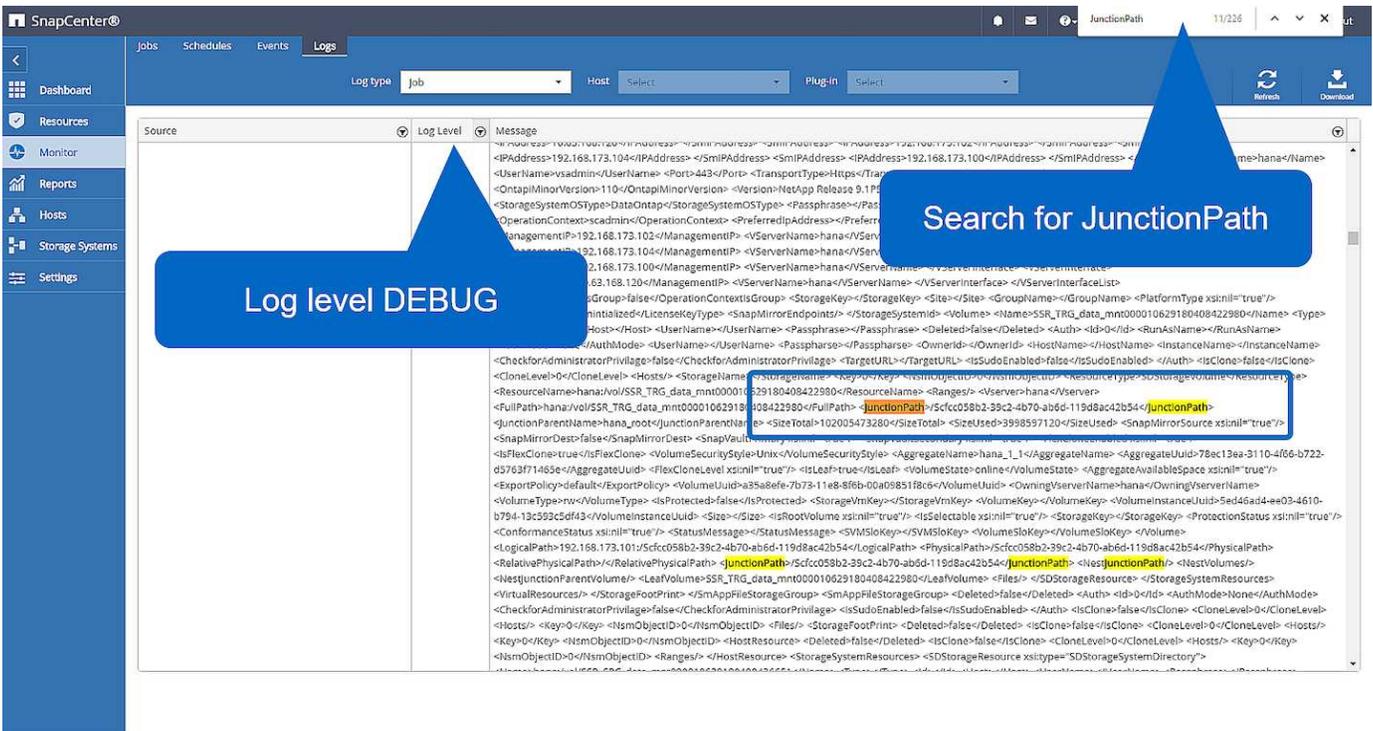
In a SnapCenter single-resource configuration, the SAP HANA plug-in is not installed at the database host. To execute the SnapCenter clone workflow, any host with an installed HANA plug-in can be used as a clone server.

+

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



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



The cloned volume can now be mounted.

```
stlrx300s8-5:/mnt/tmp # mount 192.168.173.101:/Scc373da37-00ff-4694-b1e1-8153dbd46caf /mnt/tmp
```

The cloned volume contains the data of the HANA database.

```
stlrx300s8-5:/mnt/tmp/# ls -al
drwxr-x--x 2 ssradm sapsys 4096 Jun 27 11:12 hdb00001
drwx----- 2 ssradm sapsys 4096 Jun 21 09:38 hdb00002.00003
drwx----- 2 ssradm sapsys 4096 Jun 27 11:12 hdb00003.00003
-rw-r--r-- 1 ssradm sapsys  22 Jun 27 11:12 nameserver.lck
```

The data is copied to the original location.

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

The recovery with SAP HANA Studio is performed as described in the section [SnapCenter restore of the valid backup only](#).

## Where to find additional information

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

- [SAP HANA Backup and Recovery with SnapCenter](#)
- [Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- [SAP HANA Disaster Recovery with Storage Replication](#)

<https://www.netapp.com/us/media/tr-4646.pdf>

## Version history

Version History:

Version	Date	Document Version History
Version 1.0	October 2018	Initial version
Version 2.0	January 2022	Update to cover SnapCenter 4.6 HANA System Replication support

# SAP HANA Disaster Recovery with Azure NetApp Files

## TR-4891: SAP HANA disaster recovery with Azure NetApp Files

Studies have shown that business application downtime has a significant negative impact on the business of enterprises.

Authors:

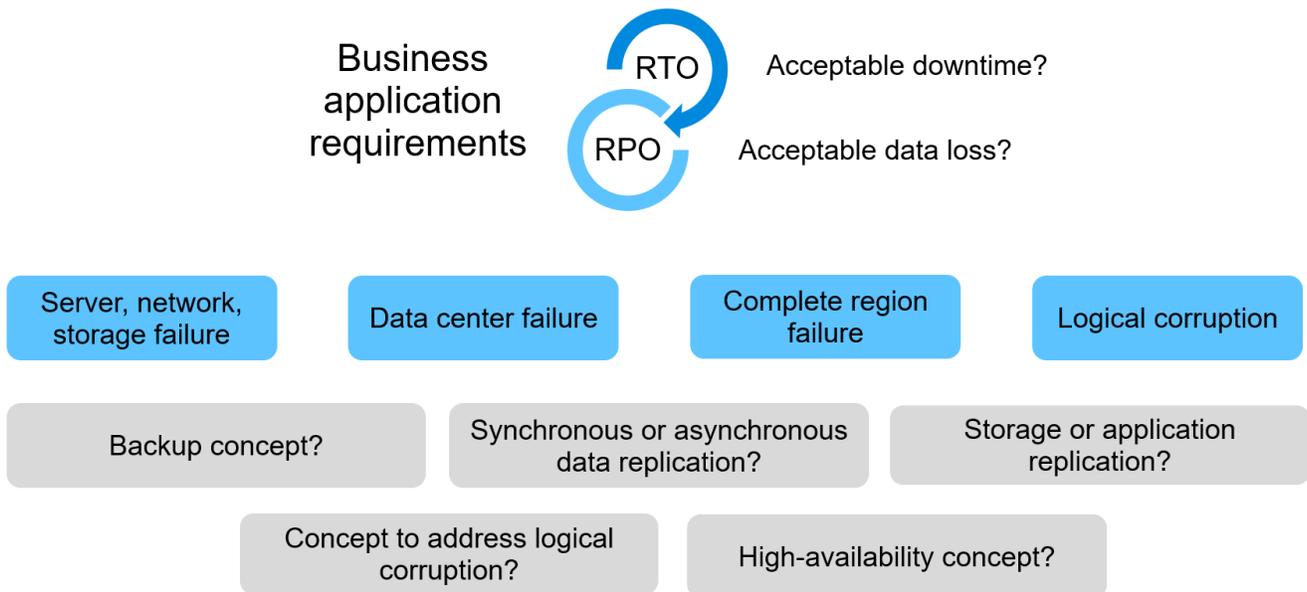
Nils Bauer, NetApp

Ralf Klahr, Microsoft

In addition to the financial impact, downtime can also damage the company's reputation, staff morale, and customer loyalty. Surprisingly, not all companies have a comprehensive disaster recovery policy.

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

To develop a comprehensive disaster recovery policy, customers must understand the business application requirements and technical capabilities they need for data protection and disaster recovery. The following figure provides an overview of data protection.



### Business application requirements

There are two key indicators for business applications:

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

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

## High availability

The infrastructure for SAP HANA, such as virtual machines, network, and storage, must have redundant components to make sure that there is no single point of failure. MS Azure provides redundancy for the different infrastructure components.

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

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

## Logical corruption

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

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

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



The required steps to setup a repair system are identical to a disaster recovery testing scenario described in this document. The described disaster recovery solution can therefore easily be extended to address logical corruption as well.

## Backups

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

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

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

In contrast, a backup solution must address logical corruption, because data replication solutions cannot cover all kinds of logical corruption.

### **Synchronous or asynchronous data replication**

The RPO primarily determines which data replication method you should use. If the RPO must be zero, even when the primary and backup storage is lost, the data must be replicated synchronously. However, there are technical limitations for synchronous replication, such as the distance between two Azure regions. In most cases, synchronous replication is not appropriate for distances greater than 100km due to latency, and therefore this is not an option for data replication between Azure regions.

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

### **HANA system replication with or without data preload**

The startup time for an SAP HANA database is much longer than that of traditional databases because a large amount of data must be loaded into memory before the database can provide the expected performance. Therefore, a significant part of the RTO is the time needed to start the database. With any storage-based replication as well as with HANA System Replication without data preload, the SAP HANA database must be started in case of failover to the disaster recovery site.

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

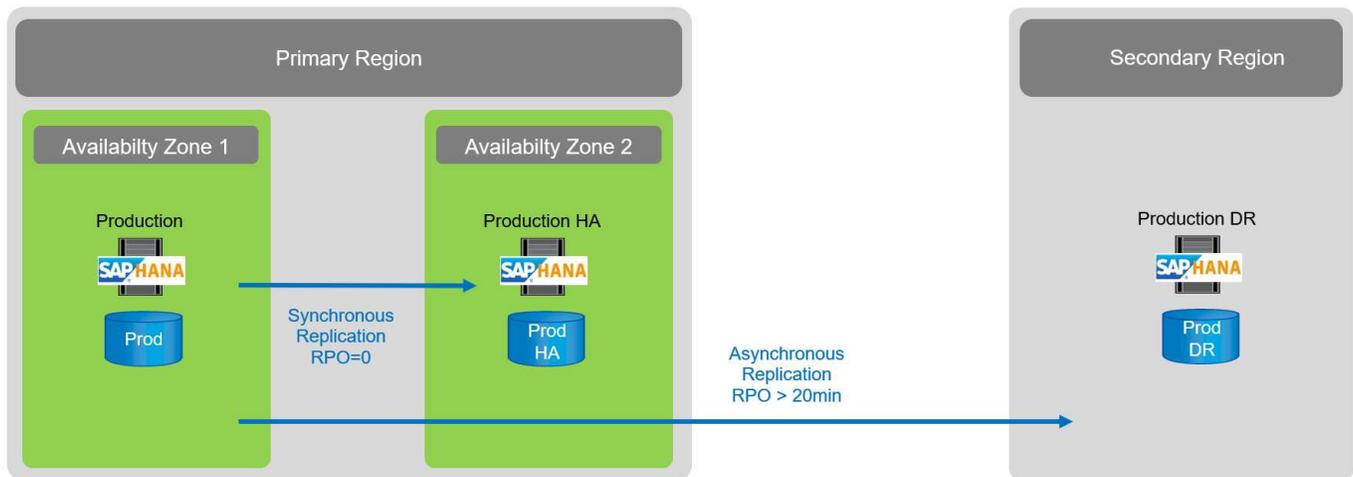
### **Disaster recovery solution comparison**

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

The following figure shows a typical configuration in which the data is replicated synchronously within the same Azure region into a second availability zone. The short distance allows you to replicate the data synchronously to achieve an RPO of zero (typically used to provide HA).

In addition, data is also replicated asynchronously to a secondary region to be protected from disasters, when the primary region is affected. The minimum achievable RPO depends on the data replication frequency, which is limited by the available bandwidth between the primary and the secondary region. A typical minimal RPO is in the range of 20 minutes to multiple hours.

This document discusses different implementation options of a two- region disaster recovery solution.



## SAP HANA System Replication

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

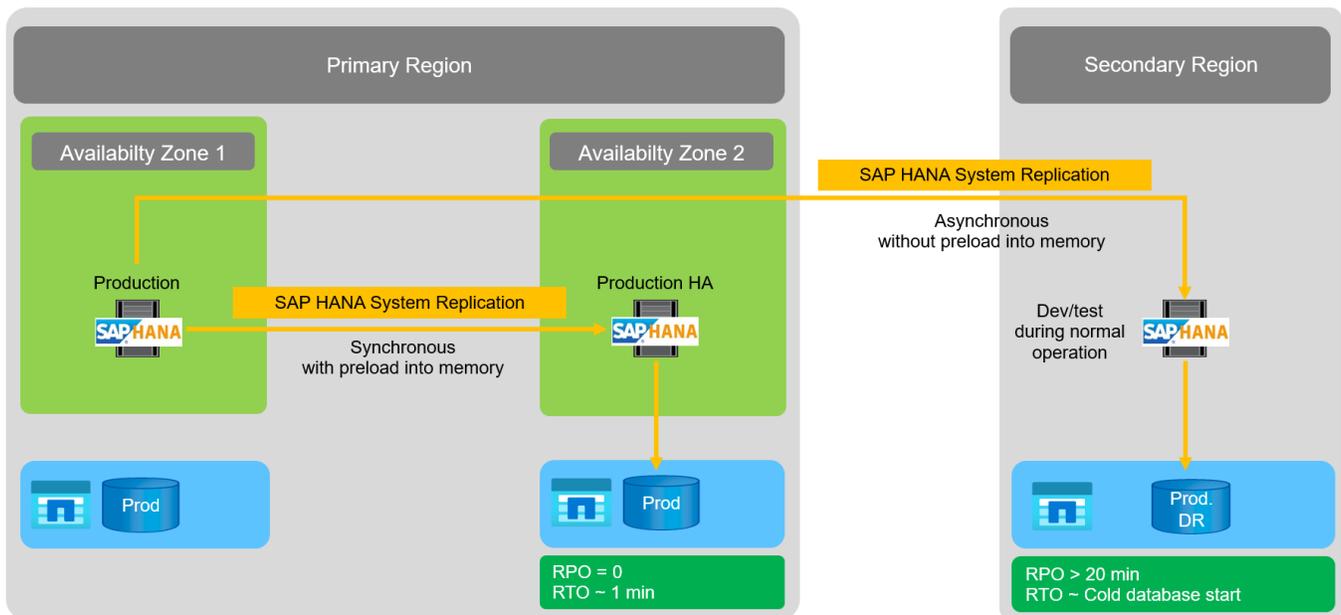
SAP HANA System Replication can be operated in one of two modes:

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

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

The following figure shows the setup of a two-region disaster recovery solution with SAP HANA System Replication. Synchronous replication with data preloaded into memory is used for local HA in the same Azure region, but in different availability zones. Asynchronous replication without data preloaded is configured for the remote disaster recovery region.

The following figure depicts SAP HANA System Replication.



### SAP HANA System Replication with data preloaded into memory

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

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

### SAP System Replication without data preloaded into memory

For less stringent RTO requirements, you can use SAP HANA System Replication without data preloaded. In this operational mode, the data at the disaster recovery region is not loaded into memory. The server at the DR region is still used to process SAP HANA System Replication running all the required SAP HANA processes. However, most of the server's memory is available to run other services, such as SAP HANA dev/test systems.

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

### SAP HANA disaster recovery with ANF Cross-Region Replication

ANF Cross-Region Replication is built into ANF as a disaster recovery solution using asynchronous data replication. ANF Cross-Region Replication is configured through a data protection relationship between two ANF volumes on a primary and a secondary Azure region. ANF Cross-Region Replication updates the secondary volume by using efficient block delta replications. Update schedules can be defined during the replication configuration.

The following figure shows a two-region disaster recovery solution example, using ANF Cross-Region Replication. In this example the HANA system is protected with HANA System Replication within the primary region as discussed in the previous chapter. The replication to a secondary region is performed using ANF

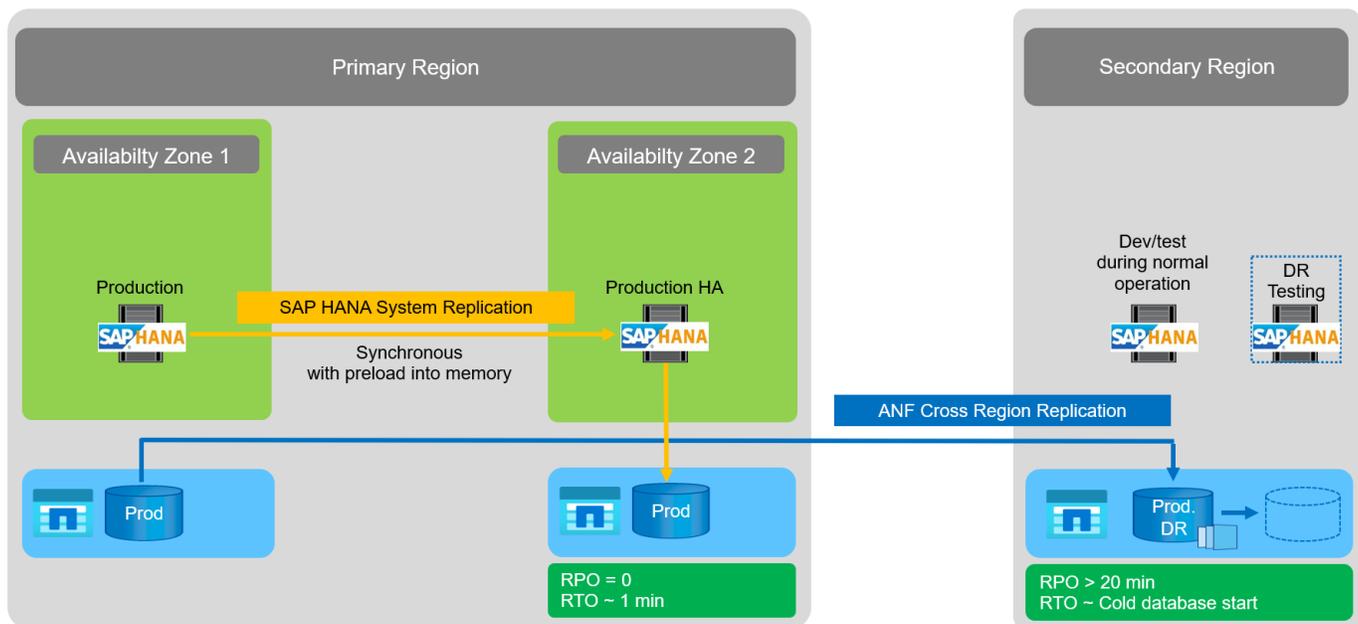
cross region replication. The RPO is defined by the replication schedule and replication options.

The RTO depends mainly on the time needed to start the HANA database at the disaster recovery site and to load the data into memory. With the assumption that the data is read with a throughput of 1000MB/s, loading 1TB of data would take approximately 18 minutes. Depending on the replication configuration, forward recovery is required as well and will add to the total RTO value.

More details on the different configuration options are provided in chapter [Configuration options for cross region replication with SAP HANA](#).

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

ANF Cross-Region Replication allows you to test the DR workflow without impacting the RPO and RTO. This is accomplished by creating volume clones and attaching them to the DR testing server.



## Summary of disaster recovery solutions

The following table compares the disaster recovery solutions discussed in this section and highlights the most important indicators.

The key findings are as follows:

- If a very low RTO is required, SAP HANA System Replication with preload into memory is the only option.
  - A dedicated server is required at the DR site to receive the replicated data and load the data into memory.
- In addition, storage replication is needed for the data that resides outside of the database (for example shared files, interfaces, and so on).
- If RTO/RPO requirements are less strict, ANF Cross-Region Replication can also be used to:
  - Combine database and nondatabase data replication.
  - Cover additional use cases such as disaster recovery testing and dev/test refresh.
  - With storage replication the server at the DR site can be used as a QA or test system during normal operation.

- A combination of SAP HANA System Replication as an HA solution with RPO=0 with storage replication for long distance makes sense to address the different requirements.

The following table provides a comparison of disaster recovery solutions.

	Storage replication	SAP HANA system replication	
	Cross-region replication	With data preload	Without data preload
RTO	Low to medium, depending on database startup time and forward recovery	Very low	Low to medium, depending on database startup time
RPO	RPO > 20min asynchronous replication	RPO > 20min asynchronous replication RPO=0 synchronous replication	RPO > 20min asynchronous replication RPO=0 synchronous replication
Servers at DR site can be used for dev/test	Yes	No	Yes
Replication of nondatabase data	Yes	No	No
DR data can be used for refresh of dev/test systems	Yes	No	No
DR testing without affecting RTO and RPO	Yes	No	No

## ANF Cross-Region Replication with SAP HANA

### ANF Cross-Region Replication with SAP HANA

Application agnostic information on Cross-Region Replication can be found at the following location.

[Azure NetApp Files documentation | Microsoft Docs](#) in the concepts and how- to guide sections.

### Configuration options for Cross-Region Replication with SAP HANA

The following figure shows the volume replication relationships for an SAP HANA system using ANF Cross-Region Replication. With ANF Cross-Region Replication, the HANA data and the HANA shared volume must be replicated. If only the HANA data volume is replicated, typical RPO values are in the range of one day. If lower RPO values are required, the HANA log backups must be also replicated for forward recovery.



The term “log backup” used in this document includes the log backup and the HANA backup catalog backup. The HANA backup catalog is required to execute forward recovery operations.

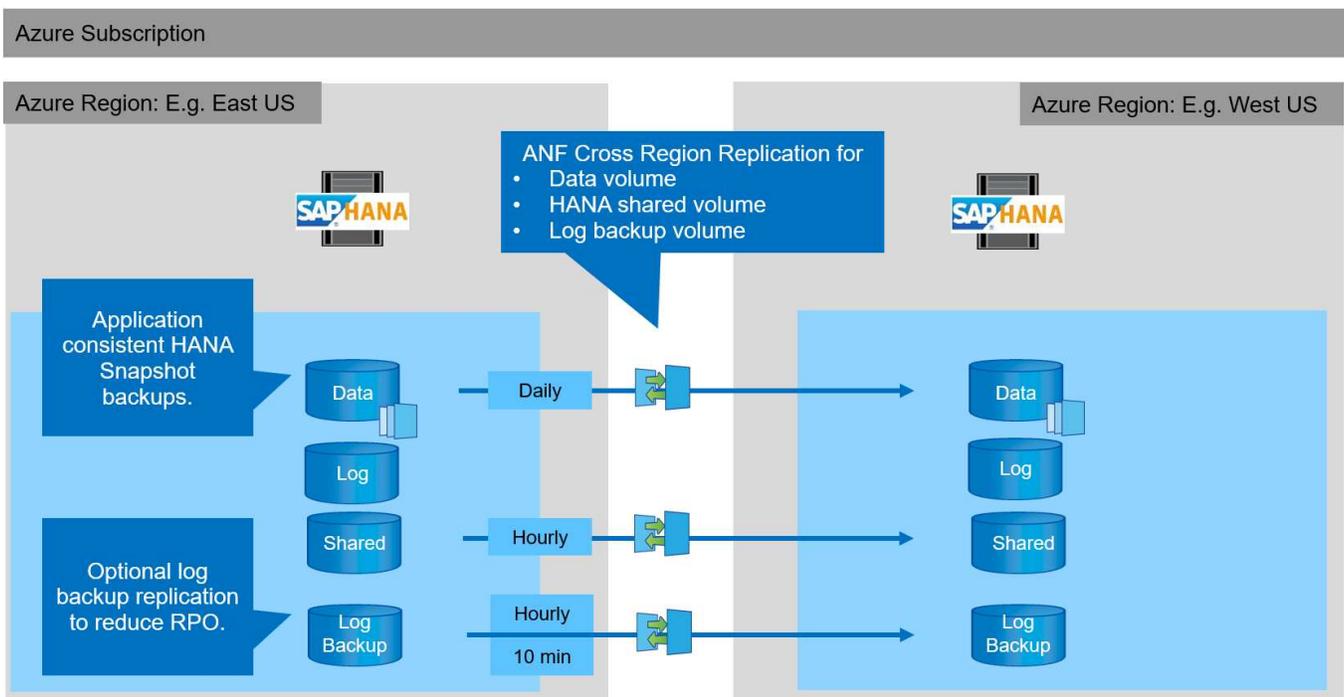


The following description and the lab setup focus on the HANA database. Other shared files, for example the SAP transport directory would be protected and replicated in the same way as the HANA shared volume.

To enable HANA save-point recovery or forward recovery using the log backups, application-consistent data Snapshot backups must be created at the primary site for the HANA data volume. This can be done for example with the ANF backup tool AzAcSnap (see also [What is Azure Application Consistent Snapshot tool for Azure NetApp Files | Microsoft Docs](#)). The Snapshot backups created at the primary site are then replicated to the DR site.

In the case of a disaster failover, the replication relationship must be broken, the volumes must be mounted to the DR production server, and the HANA database must be recovered, either to the last HANA save point or with forward recovery using the replicated log backups. The chapter [Disaster recovery failover](#), describes the required steps.

The following figure depicts the HANA configuration options for cross-region replication.



With the current version of Cross-Region Replication, only fixed schedules can be selected, and the actual replication update time cannot be defined by the user. Available schedules are daily, hourly and every 10 minutes. Using these schedule options, two different configurations make sense depending on the RPO requirements: data volume replication without log backup replication and log backup replication with different schedules, either hourly or every 10 minutes. The lowest achievable RPO is around 20 minutes. The following table summarizes the configuration options and the resulting RPO and RTO values.

	Data volume replication	Data and log backup volume replication	Data and log backup volume replication
CRR schedule data volume	Daily	Daily	Daily
CRR schedule log backup volume	n/a	Hourly	10 min

	Data volume replication	Data and log backup volume replication	Data and log backup volume replication
Max RPO	24 hours + Snapshot schedule (e.g., 6 hours)	1 hour	2 x 10 min
Max RTO	Primarily defined by HANA startup time	HANA startup time + recovery time	HANA startup time + recovery time
Forward recovery	NA	Logs for the last 24 hours + Snapshot schedule (e.g., 6 hours)	Logs for the last 24 hours + Snapshot schedule (e.g., 6 hours)

## Requirements and best practices

Microsoft Azure does not guarantee the availability of a specific virtual machine (VM) type upon creation or when starting a deallocated VM. Specifically, in case of a region failure, many clients might require additional VMs at the disaster recovery region. It is therefore recommended to actively use a VM with the required size for disaster failover as a test or QA system at the disaster recovery region to have the required VM type allocated.

For cost optimization it makes sense to use an ANF capacity pool with a lower performance tier during normal operation. The data replication does not require high performance and could therefore use a capacity pool with a standard performance tier. For disaster recovery testing, or if a disaster failover is required, the volumes must be moved to a capacity pool with a high-performance tier.

If a second capacity pool is not an option, the replication target volumes should be configured based on capacity requirements and not on performance requirements during normal operations. The quota or the throughput (for manual QoS) can then be adapted for disaster recovery testing in the case of disaster failover.

Further information can be found at [Requirements and considerations for using Azure NetApp Files volume cross-region replication | Microsoft Docs](#).

## Lab setup

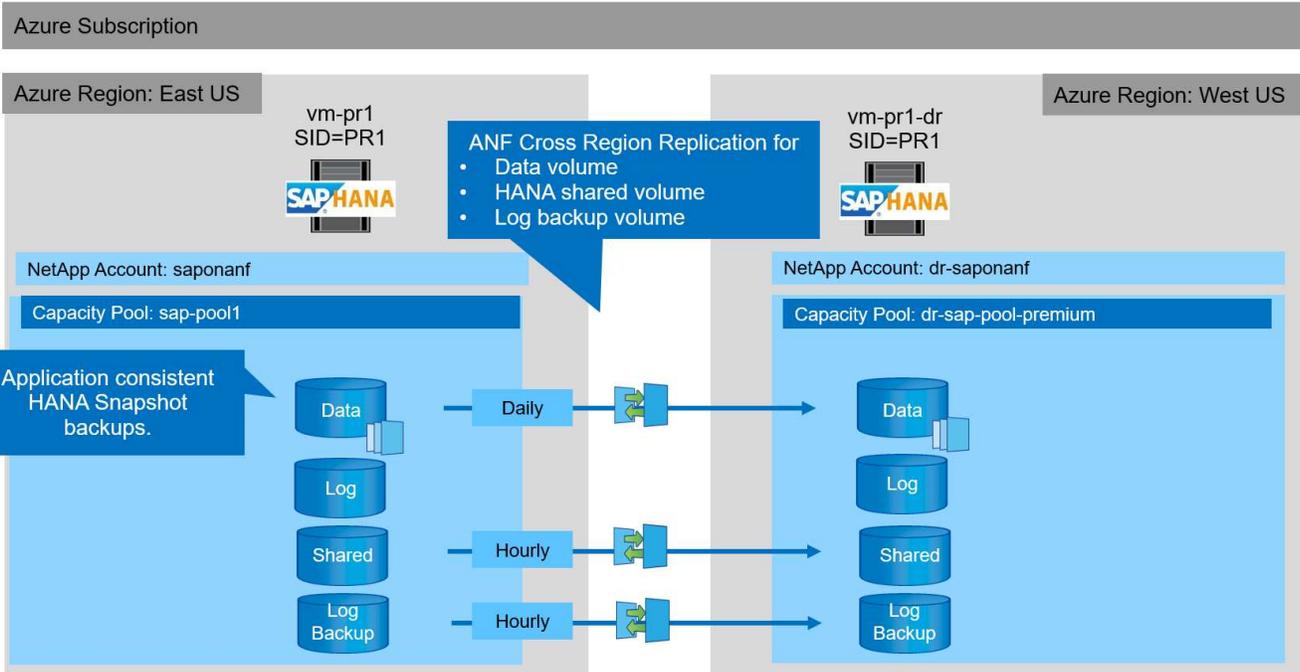
Solution validation has been performed with an SAP HANA single-host system. The Microsoft AzAcSnap Snapshot backup tool for ANF has been used to configure HANA application-consistent Snapshot backups. A daily data volume, hourly log backup, and shared volume replication were all configured. Disaster recover testing and failover was validated with a save point as well as with forward recovery operations.

The following software versions have been used in the lab setup:

- Single host SAP HANA 2.0 SPS5 system with a single tenant
- SUSE SLES for SAP 15 SP1
- AzAcSnap 5.0

A single capacity pool with manual QoS has been configured at the DR site.

The following figure depicts the lab setup.



### Snapshot backup configuration with AzAcSnap

At the primary site, AzAcSnap was configured to create application-consistent Snapshot backups of the HANA system PR1. These Snapshot backups are available at the ANF data volume of the PR1 HANA system, and they are also registered in the SAP HANA backup catalog, as shown in the following two figures. Snapshot backups were scheduled for every 4 hours.

With the replication of the data volume using ANF Cross-Region Replication, these Snapshot backups are replicated to the disaster recovery site and can be used to recover the HANA database.

The following figure shows the Snapshot backups of the HANA data volume.

1-data-mnt00001

PR1-data-mnt00001 (saponanf/sap-pool1/PR1-data-mnt00001) | Snapshots

Search (Ctrl+/)    + Add snapshot    Refresh

Overview  
Activity log  
Access control (IAM)  
Tags  
Settings  
Properties  
Locks  
Storage service  
Mount instructions  
Export policy  
Snapshots  
Replication  
Monitoring  
Metrics

Name	Location	Created
azacsnap__2021-02-12T145015-1799555Z	East US	02/12/2021, 03:49:48 PM
azacsnap__2021-02-12T145227-1245630Z	East US	02/12/2021, 03:51:24 PM
azacsnap__2021-02-12T145828-3863442Z	East US	02/12/2021, 03:58:01 PM
azacsnap__2021-02-16T134021-9431230Z	East US	02/16/2021, 02:39:18 PM
azacsnap__2021-02-16T134917-6284160Z	East US	02/16/2021, 02:48:55 PM
azacsnap__2021-02-16T135737-3778546Z	East US	02/16/2021, 02:56:32 PM
azacsnap__2021-02-16T160002-1354654Z	East US	02/16/2021, 04:59:40 PM
azacsnap__2021-02-16T200002-0790339Z	East US	02/16/2021, 08:59:42 PM
azacsnap__2021-02-17T000002-1753859Z	East US	02/17/2021, 12:59:32 AM
azacsnap__2021-02-17T040001-5454808Z	East US	02/17/2021, 04:59:31 AM
azacsnap__2021-02-17T080002-2933611Z	East US	02/17/2021, 08:59:40 AM

The following figure shows the SAP HANA backup catalog.

The screenshot shows the SAP HANA backup catalog interface. The main window displays a list of backups for the SYSTEMDB database. The backup details panel on the right shows information for a specific backup, including its ID, status, type, and location.

Status	Started	Duration	Size	Backup Type	Destinatio...
Success	Feb 17, 2021 8:00:02 ...	00h 00m 42s	3.13 GB	Data Backup	Snapshot
Success	Feb 17, 2021 4:00:01 ...	00h 00m 35s	3.13 GB	Data Backup	Snapshot
Success	Feb 17, 2021 12:00:00 ...	00h 00m 36s	3.13 GB	Data Backup	Snapshot
Success	Feb 16, 2021 8:00:02 ...	00h 00m 34s	3.13 GB	Data Backup	Snapshot
Success	Feb 16, 2021 4:00:02 ...	00h 00m 38s	3.13 GB	Data Backup	Snapshot
Success	Feb 16, 2021 1:57:37 ...	00h 00m 32s	3.13 GB	Data Backup	Snapshot
Success	Feb 16, 2021 1:49:17 ...	00h 00m 32s	3.13 GB	Data Backup	Snapshot
Success	Feb 16, 2021 1:40:22 ...	00h 00m 34s	3.13 GB	Data Backup	Snapshot
Success	Feb 12, 2021 2:58:28 ...	00h 00m 32s	3.13 GB	Data Backup	Snapshot
Success	Feb 12, 2021 2:52:27 ...	00h 00m 32s	3.13 GB	Data Backup	Snapshot
Success	Feb 12, 2021 2:50:15 ...	00h 00m 32s	3.13 GB	Data Backup	Snapshot

**Backup Details:**

- ID: 1613141415533
- Status: Successful
- Backup Type: Data Backup
- Destination Type: Snapshot
- Started: Feb 12, 2021 2:50:15 PM (UTC)
- Finished: Feb 12, 2021 2:50:48 PM (UTC)
- Duration: 00h 00m 32s
- Size: 3.13 GB
- Throughput: n.a.
- System ID:
- Comment: Snapshot prefix: azacsnap  
Tools version: 5.0 Preview (20201214.65524)
- Additional Information: <ok>
- Location: /hana/data/PR1/mnt00001/

Host	Service	Size	Name	Source ...	EBID
vm-pr1	nameserver	3.13 GB	hdb00001	volume	azacsnap_2021-02-12T14501...

### Configuration steps for ANF Cross-Region Replication

A few preparation steps must be performed at the disaster recovery site before volume replication can be configured.

- A NetApp account must be available and configured with the same Azure subscription as the source.
- A capacity pool must be available and configured using the above NetApp account.
- A virtual network must be available and configured.
- Within the virtual network, a delegated subnet must be available and configured for use with ANF.

Protection volumes can now be created for the HANA data, the HANA shared and the HANA log backup volume. The following table shows the configured destination volumes in our lab setup.



To achieve the best latency, the volumes must be placed close to the VMs that run the SAP HANA in case of a disaster failover. Therefore, the same pinning process is required for the DR volumes as for any other SAP HANA production system.

HANA volume	Source	Destination	Replication schedule
HANA data volume	PR1-data-mnt00001	PR1-data-mnt00001-sm-dest	Daily
HANA shared volume	PR1-shared	PR1-shared-sm-dest	Hourly
HANA log/catalog backup volume	hanabackup	hanabackup-sm-dest	Hourly

For each volume, the following steps must be performed:

1. Create a new protection volume at the DR site:
  - a. Provide the volume name, capacity pool, quota, and network information.

- b. Provide the protocol and volume access information.
  - c. Provide the source volume ID and a replication schedule.
  - d. Create a target volume.
2. Authorize replication at the source volume.
    - Provide the target volume ID.

The following screenshots show the configuration steps in detail.

At the disaster recovery site, a new protection volume is created by selecting volumes and clicking Add Data Replication. Within the Basics tab, you must provide the volume name, capacity pool and network information.



The quota of the volume can be set based on capacity requirements, because volume performance does not have an effect on the replication process. In the case of a disaster recovery failover, the quota must be adjusted to fulfill the real performance requirements.



If the capacity pool has been configured with manual QoS, you can configure the throughput in addition to the capacity requirements. Same as above, you can configure the throughput with a low value during normal operation and increase it in case of a disaster recovery failover.

# Create a new protection volume

[Basics](#) [Protocol](#) [Replication](#) [Tags](#) [Review + create](#)

This page will help you create an Azure NetApp Files volume in your subscription and enable you to access the volume from within your virtual network. [Learn more about Azure NetApp Files](#)

## Volume details

Volume name *	<input type="text" value="PR1-data-mnt00001-sm-dest"/> ✓
Capacity pool * ⓘ	<input type="text" value="dr-sap-pool1"/> ▼
Available quota (GiB) ⓘ	<input type="text" value="4096"/> 4 TiB
Quota (GiB) * ⓘ	<input type="text" value="500"/> ✓ 500 GiB
Virtual network * ⓘ	<input type="text" value="dr-vnet (10.2.0.0/16,10.0.2.0/24)"/> ▼ <a href="#">Create new</a>
Delegated subnet * ⓘ	<input type="text" value="default (10.0.2.0/28)"/> ▼ <a href="#">Create new</a>
Show advanced section	<input type="checkbox"/>

[Review + create](#)

[< Previous](#)

[Next : Protocol >](#)

In the Protocol tab, you must provide the network protocol, the network path, and the export policy.



The protocol must be the same as the protocol used for the source volume.

# Create a new protection volume

Basics Protocol Replication Tags Review + create

Configure access to your volume.

## Access

Protocol type  NFS  SMB  Dual-protocol (NFSv3 and SMB)

## Configuration

File path \* ⓘ

Versions \*  ▼

Kerberos  Enabled  Disabled

## Export policy

Configure the volume's export policy. This can be edited later. [Learn more](#)

↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom 🗑 Delete

<input checked="" type="checkbox"/>	Index	Allowed clients	Access	Root Access	
<input checked="" type="checkbox"/>	1	<input type="text" value="0.0.0.0/0"/>	<input type="text" value="Read &amp; Write"/> ▼	<input type="text" value="On"/> ▼	⋮
		<input type="text"/>	<input type="text"/> ▼	<input type="text"/> ▼	

**Review + create**

< Previous

Next : Replication >

Within the Replication tab, you must configure the source volume ID and the replication schedule. For data volume replication, we configured a daily replication schedule for our lab setup.



The source volume ID can be copied from the Properties screen of the source volume.

# Create a new protection volume

Basics Protocol **Replication** Tags Review + create

Source volume ID ⓘ

Replication schedule ⓘ 

Daily ^

Every 10 minutes

Hourly

Daily

**Review + create**

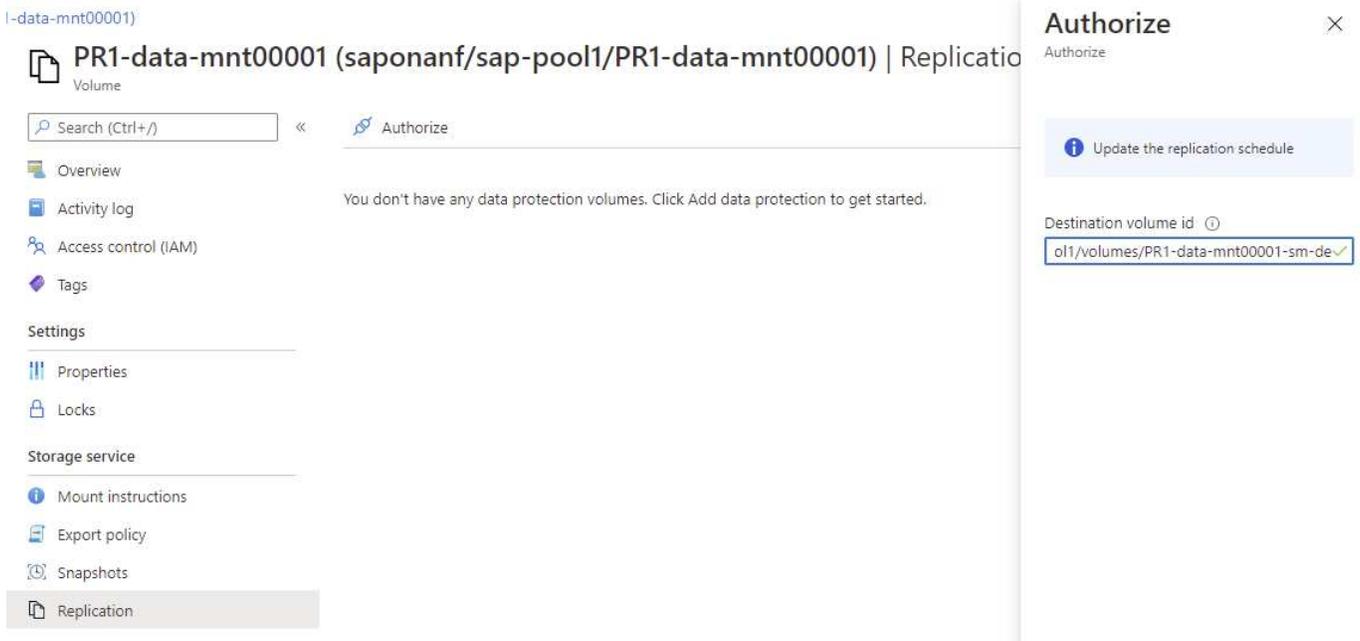
< Previous

Next : Tags >

As a final step, you must authorize replication at the source volume by providing the ID of the target volume.



You can copy the destination volume ID from the Properties screen of the destination volume.



The same steps must be performed for the HANA shared and the log backup volume.

### Monitoring ANF Cross-Region Replication

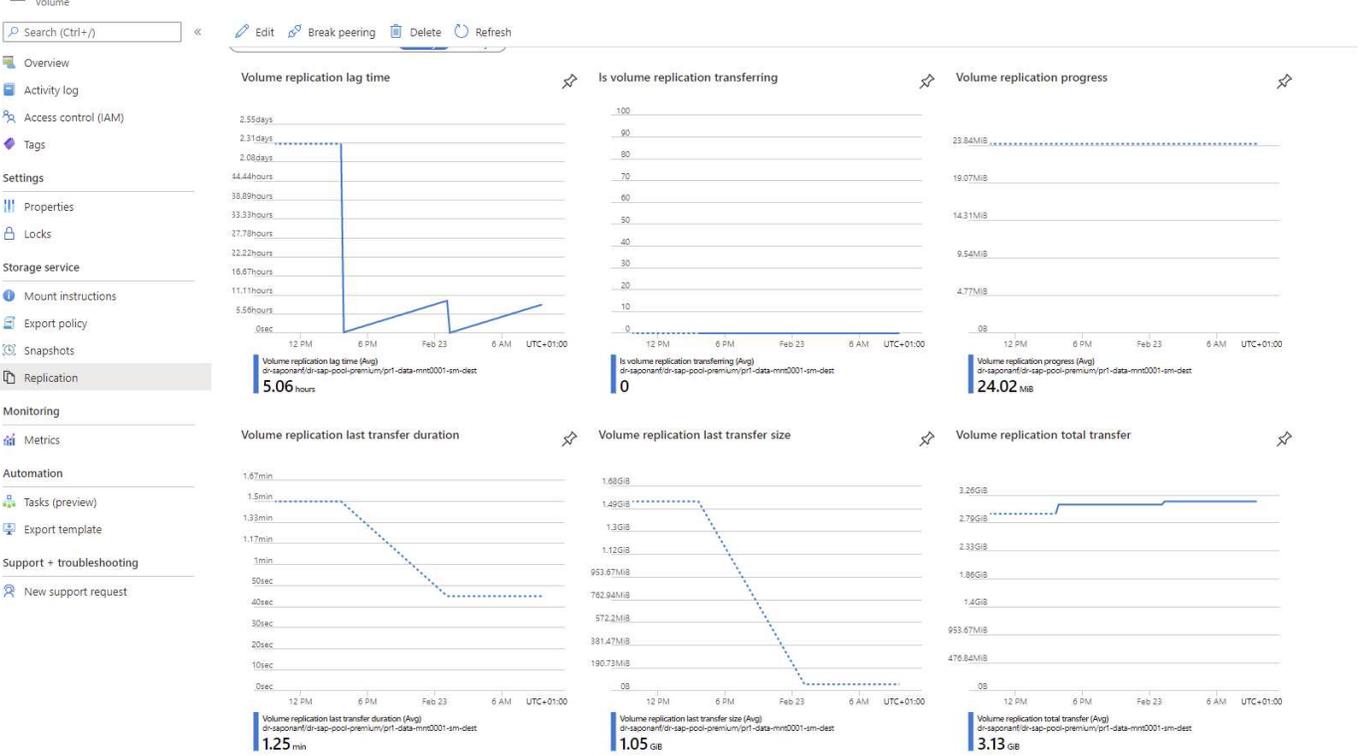
The following three screenshots show the replication status for the data, log backup, and shared volumes.

The volume replication lag time is a useful value to understand RPO expectations. For example, the log backup volume replication shows a maximum lag time of 58 minutes, which means that the maximum RPO has the same value.

The transfer duration and transfer size provide valuable information on bandwidth requirements and change the rate of the replicated volume.

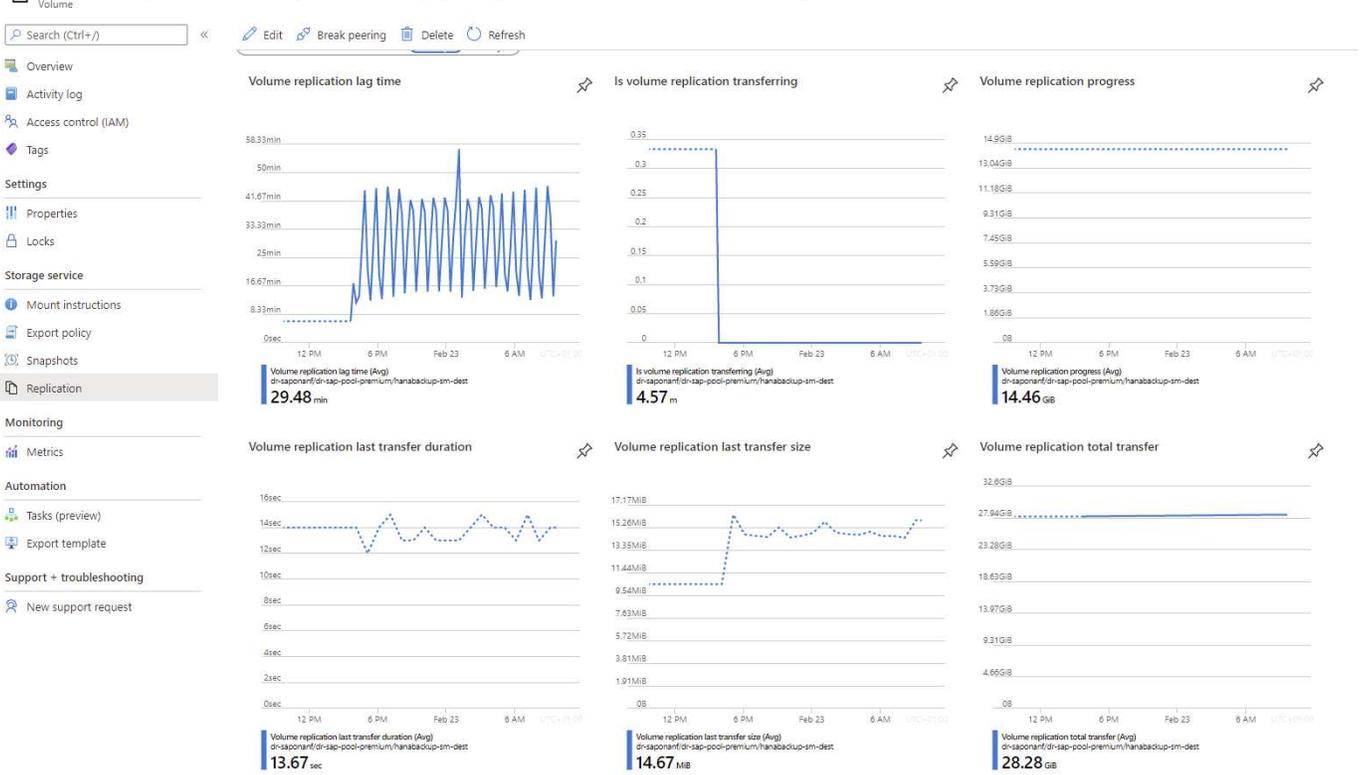
The following screenshot shows the replication status of HANA data volume.

### PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest) | Replication



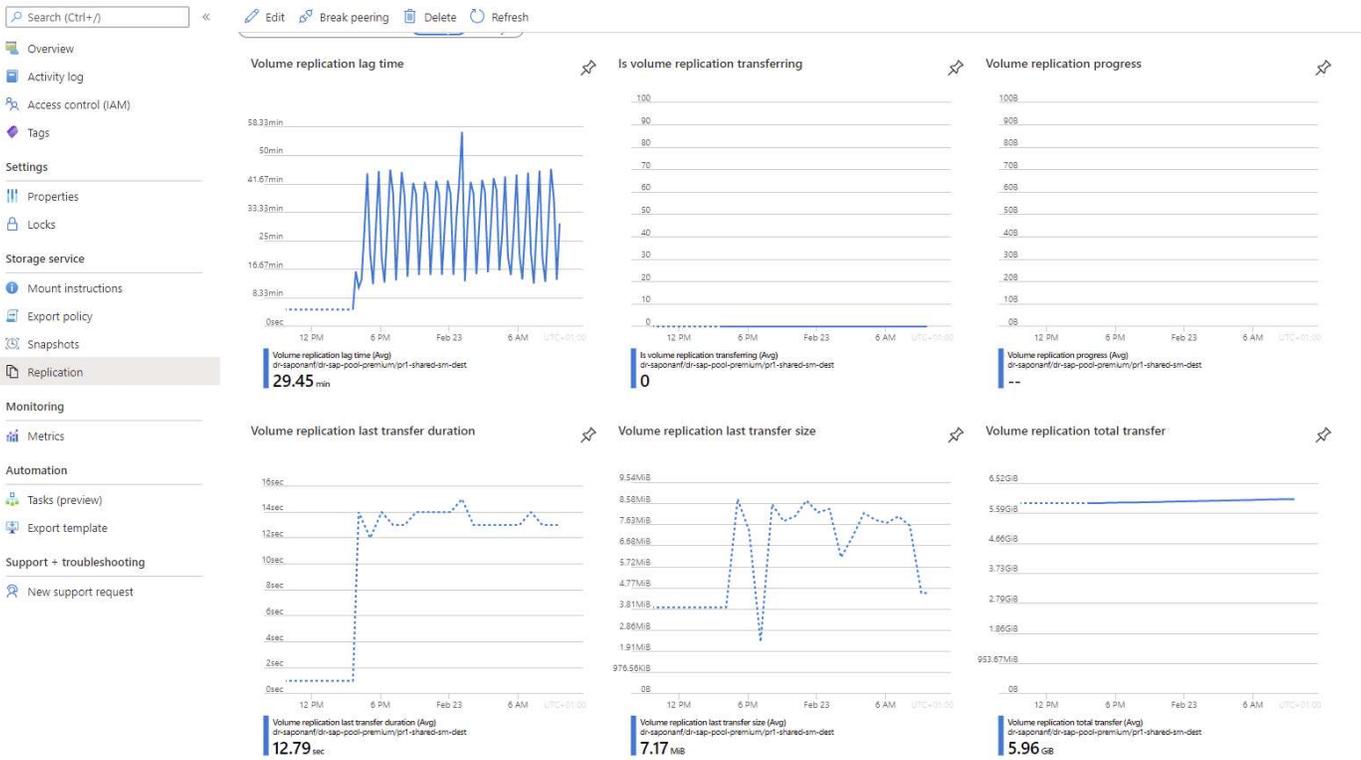
The following screenshot shows the replication status of HANA log backup volume.

### hanabackup-sm-dest (dr-saponanf/dr-sap-pool-premium/hanabackup-sm-dest) | Replication



The following screenshot shows the replication status of HANA shared volume.

PR1-shared-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-shared-sm-dest) | Replication



### Replicated snapshot backups

With each replication update from the source to the target volume, all block changes that happened between the last and the current update are replicated to the target volume. This also includes the snapshots, which have been created at the source volume. The following screenshot shows the snapshots available at the target volume. As already discussed, each of the snapshots created by the AzAcSnap tool are application-consistent images of the HANA database that can be used to execute either a savepoint or a forward recovery.



Within the source and the target volume, SnapMirror Snapshot copies are created as well, which are used for resync and replication update operations. These Snapshot copies are not application consistent from the HANA database perspective; only the application-consistent snapshots created via AzaCSnap can be used for HANA recovery operations.

PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest)   Snapshots		
Volume		
Search (Ctrl+/) << + Add snapshot Refresh		
Overview	Search snapshots	
Activity log		
Access control (iAM)		
Tags		
Settings		
Properties		
Locks		
Storage service		
Mount instructions		
Export policy		
<b>Snapshots</b>		
Replication		
Monitoring		
Metrics		
Automation		
Tasks (preview)		
Export template		
Support + troubleshooting		
New support request		
Name	Location	Created
azacsnap__2021-02-18T120002-2150721Z	West US	02/18/2021, 01:00:05 PM
azacsnap__2021-02-18T160002-1442691Z	West US	02/18/2021, 05:00:49 PM
azacsnap__2021-02-18T200002-0756687Z	West US	02/18/2021, 09:00:05 PM
azacsnap__2021-02-19T000002-0039686Z	West US	02/19/2021, 01:00:05 AM
azacsnap__2021-02-19T040001-8773746Z	West US	02/19/2021, 05:00:05 AM
azacsnap__2021-02-19T080001-5198653Z	West US	02/19/2021, 09:00:05 AM
azacsnap__2021-02-19T120002-1495322Z	West US	02/19/2021, 01:00:05 PM
azacsnap__2021-02-19T160002-3698678Z	West US	02/19/2021, 05:00:05 PM
azacsnap__2021-02-22T120002-3145396Z	West US	02/22/2021, 01:00:05 PM
snapmirrorb1e8e48d-7114-11eb-b147-d039ea1e211e_2155791247.2021-02-23_143159	West US	02/23/2021, 03:32:00 PM
azacsnap__2021-02-22T160002-0144647Z	West US	02/22/2021, 05:00:05 PM
azacsnap__2021-02-22T200002-0649581Z	West US	02/22/2021, 09:00:05 PM
azacsnap__2021-02-23T000002-0311379Z	West US	02/23/2021, 01:00:05 AM
snapmirrorb1e8e48d-7114-11eb-b147-d039ea1e211e_2155791247.2021-02-23_001000	West US	02/23/2021, 01:10:00 AM

## Disaster recovery testing

### Disaster Recovery Testing

To implement an effective disaster recovery strategy, you must test the required workflow. Testing demonstrates whether the strategy works and whether the internal documentation is sufficient, and it also allows administrators to train on the required procedures.

ANF Cross-Region Replication enables disaster recovery testing without putting RTO and RPO at risk. Disaster recovery testing can be done without interrupting data replication.

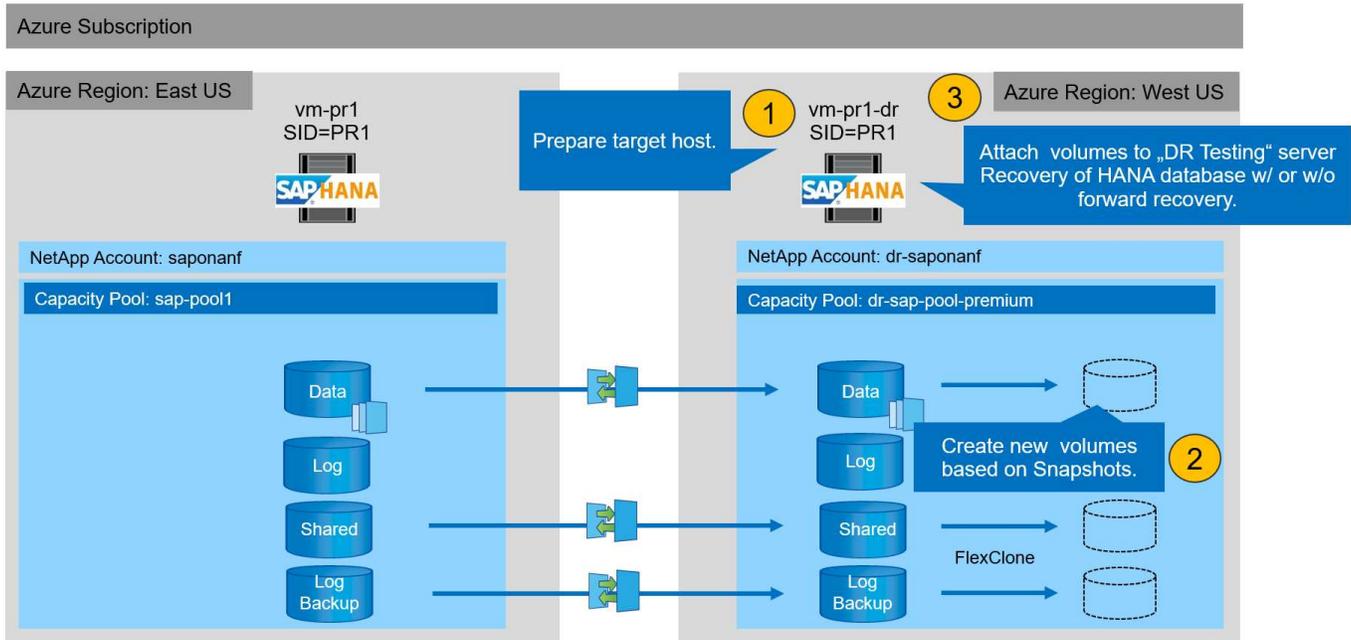
The disaster recovery testing workflow leverages the ANF feature set to create new volumes based on existing Snapshot backups at the disaster recovery target. See [How Azure NetApp Files snapshots work | Microsoft Docs](#).

Depending on whether log backup replication is part of the disaster recovery setup or not, the steps for disaster recovery are slightly different. This section describes the disaster recovery testing for data-backup-only replication as well as for data volume replication combined with log backup volume replication.

To perform disaster recovery testing, complete the following steps:

1. Prepare the target host.
2. Create new volumes based on Snapshot backups at the disaster recovery site.
3. Mount the new volumes at the target host.
4. Recover the HANA database.
  - Data volume recovery only.
  - Forward recovery using replicated log backups.

The following subsections describe these steps in detail.



## Prepare the target host

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

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

The target host preparation also includes shutting down the HANA QA or test system as well as stopping all services using `systemctl stop sapinit`.

### 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, resulting in logically corrupted data.

### Install required software

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



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

## Configure users, ports, and SAP services

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

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

```
vm-pr1:~ # cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/PR1/HDB01/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/PR1/HDB01/exe/sapstartsrv
pf=/usr/sap/PR1/SYS/profile/PR1_HDB01_vm-pr1 -D -u pr1adm
limit.descriptors=1048576
```

## Prepare HANA log volume

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

```
vm-pr1:~ # ls -al /hana/log/PR1/mnt00001/
total 16
drwxrwxrwx 5 root    root    4096 Feb 19 16:20 .
drwxr-xr-x 3 root    root     22 Feb 18 13:38 ..
drwxr-xr-- 2 pr1adm sapsys 4096 Feb 22 10:25 hdb00001
drwxr-xr-- 2 pr1adm sapsys 4096 Feb 22 10:25 hdb00002.00003
drwxr-xr-- 2 pr1adm sapsys 4096 Feb 22 10:25 hdb00003.00003
vm-pr1:~ #
```

## Prepare log backup volume

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

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

## Prepare file system mounts

The following table shows the naming conventions used in the lab setup. The volume names at the disaster recovery site are included in `/etc/fstab`.

HANA PR1 volumes	Volume and subdirectories at disaster recovery site	Mount point at target host
Data volume	PR1-data-mnt00001-sm-dest	/hana/data/PR1/mnt00001
Shared volume	PR1-shared-sm-dest/shared PR1-shared-sm-dest/usr-sap-PR1	/hana/shared /usr/sap/PR1
Log backup volume	hanabackup-sm-dest	/hanabackup



The mount points from this table must be created at the target host.

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

```
vm-pr1:~ # cat /etc/fstab
# HANA ANF DB Mounts
10.0.2.4:/PR1-data-mnt00001-sm-dest /hana/data/PR1/mnt00001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-log-mnt00001-dr /hana/log/PR1/mnt00001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
# HANA ANF Shared Mounts
10.0.2.4:/PR1-shared-sm-dest/hana-shared /hana/shared nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-shared-sm-dest/usr-sap-PR1 /usr/sap/PR1 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
# HANA file and log backup destination
10.0.2.4:/hanabackup-sm-dest /hanabackup nfs
rw,vers=3,hard,timeo=600,rsize=262144,wsiz=262144,nconnect=8,bg,noatime,n
olock 0 0
```

### Create new volumes based on snapshot backups at the disaster recovery site

Depending on the disaster recovery setup (with or without log backup replication), two or three new volumes based on snapshot backups must be created. In both cases, a new volume of the data and the HANA shared volume must be created.

A new volume of the log backup volume must be created if the log backup data is also replicated. In our example, data and the log backup volume have been replicated to the disaster recovery site. The following steps use the Azure Portal.

1. One of the application-consistent snapshot backups is selected as a source for the new volume of the HANA data volume. Restore to New Volume is selected to create a new volume based on the snapshot backup.

PR1-data-mnt00001-sm-dest (dr-saponanf/dr-sap-pool1/PR1-data-mnt00001-sm-dest) | Snapshots

Volume

Search (Ctrl+/) << + Add snapshot Refresh

- Overview
- Activity log
- Access control (IAM)
- Tags
- Settings
  - Properties
  - Locks
  - Storage service
    - Mount instructions
    - Export policy
    - Snapshots**
    - Replication
  - Monitoring
    - Metrics
  - Automation
    - Tasks (preview)
    - Export template
  - Support + troubleshooting
    - New support request

Search snapshots

Name	Location	Created	
azacsnap_2021-02-16T134021-9431230Z	West US	02/16/2021, 02:40:27 PM	...
azacsnap_2021-02-16T134917-6284160Z	West US	02/16/2021, 02:49:20 PM	...
azacsnap_2021-02-16T135737-3778546Z	West US	02/16/2021, 02:57:41 PM	...
azacsnap_2021-02-16T160002-1354654Z	West US	02/16/2021, 05:00:05 PM	...
azacsnap_2021-02-16T200002-0790339Z	West US	02/16/2021, 09:00:08 PM	...
azacsnap_2021-02-17T000002-1753859Z	West US	02/17/2021, 01:00:06 AM	...
azacsnap_2021-02-17T040001-5454808Z	West US	02/17/2021, 05:00:05 AM	...
azacsnap_2021-02-17T080002-2933611Z	West US	02/17/2021, 09:00:18 AM	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/17/2021, 12:46:22 PM	...
azacsnap_2021-02-17T120001-9196266Z	West US	02/17/2021, 01:00:08 PM	...
azacsnap_2021-02-17T160002-2801612Z	West US	02/17/2021, 05:00:06 PM	...
azacsnap_2021-02-17T200001-9149055Z	West US	02/17/2021, 09:00:05 PM	...
azacsnap_2021-02-18T000001-7955243Z	West US	02/18/2021, 01:00:07 AM	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 01:10:00 AM	...

Restore to new volume  
Revert volume  
Delete

2. The new volume name and quota must be provided in the user interface.

## Create a volume

Basics Protocol Tags Review + create

This page will help you create an Azure NetApp Files volume in your subscription and enable you to access the volume from within your virtual network. [Learn more about Azure NetApp Files](#)

### Volume details

Volume name \*  ✓

Restoring from snapshot ⓘ

Available quota (GiB) ⓘ  2.05 TiB

Quota (GiB) \* ⓘ  500 GiB

Virtual network ⓘ  ▼

Delegated subnet ⓘ  ▼

Show advanced section

3. Within the protocol tab, the file path and export policy are configured.

Home > Azure NetApp Files > dr-saponanf > dr-sap-pool1 (dr-saponanf/dr-sap-pool1) > PR1-data-mnt00001-sm-dest (d

## Create a volume

Basics Protocol Tags Review + create

Configure access to your volume.

### Access

Protocol type  NFS  SMB  Dual-protocol (NFSv3 and SMB)

### Configuration

File path \* ⓘ

Versions  ▼

Kerberos  Enabled  Disabled

### Export policy

Configure the volume's export policy. This can be edited later. [Learn more](#)

↑ Move up ↓ Move down ↕ Move to top ⬇ Move to bottom 🗑 Delete

<input checked="" type="checkbox"/>	Index	Allowed clients	Access	Root Access	
<input checked="" type="checkbox"/>	1	<input type="text" value="0.0.0.0/0"/>	Read & Write ▼	On ▼	⋮
		<input type="text"/>	▼	▼	

4. The Create and Review screen summarizes the configuration.

# Create a volume

Validation passed

Basics Protocol Tags Review + create

## Basics

Subscription Pay-As-You-Go  
 Resource group dr-rg-sap  
 Region West US  
 Volume name PR1-data-mnt00001-sm-dest-clone  
 Capacity pool dr-sap-pool1  
 Service level Standard  
 Quota 500 GiB

## Networking

Virtual network dr-vnet (10.2.0.0/16,10.0.2.0/24)  
 Delegated subnet default (10.0.2.0/28)

## Protocol

Protocol NFSv4.1  
 File path PR1-data-mnt00001-sm-dest-clone

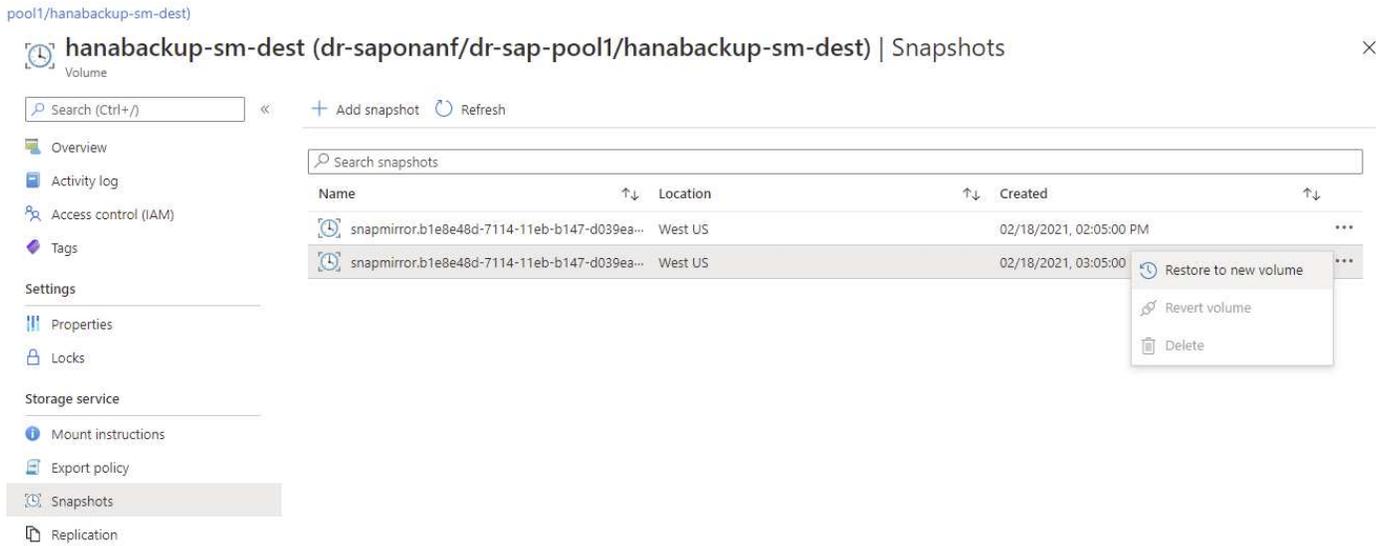
5. A new volume has now been created based on the HANA snapshot backup.

dr-saponanf | Volumes

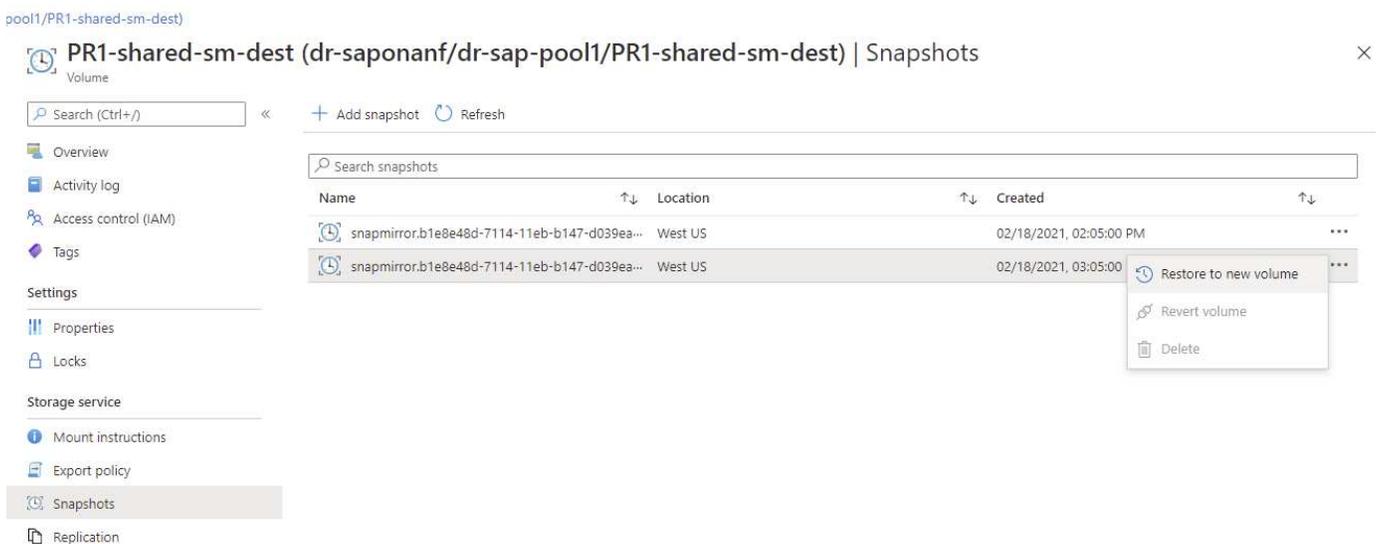
Search (Ctrl+/) << + Add volume + Add data replication Refresh

Name	Quota	Protocol type	Mount path	Service level	Capacity pool
hanabackup-sm-dest	1000 GiB	NFSv3	10.0.2.4/hanabackup-sm-dest	Standard	dr-sap-pool1
PR1-data-mnt00001-sm-dest	500 GiB	NFSv4.1	10.0.2.4/PR1-data-mnt00001-s	Standard	dr-sap-pool1
<b>PR1-data-mnt00001-sm-dest-clone</b>	500 GiB	NFSv4.1	10.0.2.4/PR1-data-mnt00001-s	Standard	dr-sap-pool1
PR1-log-mnt00001-dr	250 GiB	NFSv4.1	10.0.2.4/PR1-log-mnt00001-dr	Standard	dr-sap-pool1
PR1-shared-sm-dest	250 GiB	NFSv4.1	10.0.2.4/PR1-shared-sm-dest	Standard	dr-sap-pool1

The same steps must now be performed for the HANA shared and the log backup volume as shown in the following two screenshots. Since no additional snapshots have been created for the HANA shared and log backup volume, the newest SnapMirror Snapshot copy must be selected as the source for the new volume. This is unstructured data, and the SnapMirror Snapshot copy can be used for this use case.



The following screenshot shows the HANA shared volume restored to new volume.



If a capacity pool with a low performance tier has been used, the volumes must now be moved to a capacity pool that provides the required performance.

All three new volumes are now available and can be mounted at the target host.

### Mount the new volumes at the target host

The new volumes can now be mounted at the target host, based on the `/etc/fstab` file created before.

```
vm-pr1:~ # mount -a
```

The following output shows the required file systems.

```
vm-pr1:/hana/data/PR1/mnt00001/hdb00001 # df
Filesystem                                1K-blocks      Used
Available Use% Mounted on
devtmpfs                                  8190344         8
8190336   1% /dev
tmpfs                                      12313116         0
12313116   0% /dev/shm
tmpfs                                      8208744        17292
8191452   1% /run
tmpfs                                      8208744         0
8208744   0% /sys/fs/cgroup
/dev/sda4                                  29866736    2438052
27428684   9% /
/dev/sda3                                  1038336       101520
936816   10% /boot
/dev/sda2                                  524008         1072
522936   1% /boot/efi
/dev/sdb1                                  32894736       49176
31151560   1% /mnt
tmpfs                                      1641748         0
1641748   0% /run/user/0
10.0.2.4:/PR1-log-mnt00001-dr              107374182400      256
107374182144   1% /hana/log/PR1/mnt00001
10.0.2.4:/PR1-data-mnt00001-sm-dest-clone  107377026560    6672640
107370353920   1% /hana/data/PR1/mnt00001
10.0.2.4:/PR1-shared-sm-dest-clone/hana-shared 107377048320 11204096
107365844224   1% /hana/shared
10.0.2.4:/PR1-shared-sm-dest-clone/usr-sap-PR1 107377048320 11204096
107365844224   1% /usr/sap/PR1
10.0.2.4:/hanabackup-sm-dest-clone         107379429120 35293440
107344135680   1% /hanabackup
```

## HANA database recovery

The following shows the steps for HANA database recovery

Start the required SAP services.

```
vm-pr1:~ # systemctl start sapinit
```

The following output shows the required processes.

```
vm-pr1:/ # ps -ef | grep sap
root      23101      1  0 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/saphostexec pf=/usr/sap/hostctrl/exe/host_profile
pr1adm    23191      1  3 11:29 ?          00:00:00
/usr/sap/PR1/HDB01/exe/sapstartsrv
pf=/usr/sap/PR1/SYS/profile/PR1_HDB01_vm-pr1 -D -u pr1adm
sapadm    23202      1  5 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/hostctrl/exe/host_profile -D
root      23292      1  0 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/saposcol -l -w60
pf=/usr/sap/hostctrl/exe/host_profile
root      23359    2597  0 11:29 pts/1      00:00:00 grep --color=auto sap
```

The following subsections describe the recovery process with and without forward recovery using the replicated log backups. The recovery is executed using the HANA recovery script for the system database and hdbsql commands for the tenant database.

#### Recovery to latest HANA data volume backup savepoint

The recovery to the latest backup savepoint is executed with the following commands as user pr1adm:

- System database

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

- Tenant database

```
Within hdbsql: RECOVER DATA FOR PR1 USING SNAPSHOT CLEAR LOG
```

You can also use HANA Studio or Cockpit to execute the recovery of the system and the tenant database.

The following command output show the recovery execution.

#### System database recovery

```

pr1adm@vm-pr1:/usr/sap/PR1/HDB01> HDBSettings.sh recoverSys.py
--command="RECOVER DATA USING SNAPSHOT CLEAR LOG"
[139702869464896, 0.008] >> starting recoverSys (at Fri Feb 19 14:32:16
2021)
[139702869464896, 0.008] args: ()
[139702869464896, 0.009] keys: {'command': 'RECOVER DATA USING SNAPSHOT
CLEAR LOG'}
using logfile /usr/sap/PR1/HDB01/vm-pr1/trace/backup.log
recoverSys started: =====2021-02-19 14:32:16 =====
testing master: vm-pr1
vm-pr1 is master
shutdown database, timeout is 120
stop system
stop system on: vm-pr1
stopping system: 2021-02-19 14:32:16
stopped system: 2021-02-19 14:32:16
creating file recoverInstance.sql
restart database
restart master nameserver: 2021-02-19 14:32:21
start system: vm-pr1
sapcontrol parameter: ['-function', 'Start']
sapcontrol returned successfully:
2021-02-19T14:32:56+00:00 P0027646      177bab4d610 INFO      RECOVERY
RECOVER DATA finished successfully
recoverSys finished successfully: 2021-02-19 14:32:58
[139702869464896, 42.017] 0
[139702869464896, 42.017] << ending recoverSys, rc = 0 (RC_TEST_OK), after
42.009 secs
pr1adm@vm-pr1:/usr/sap/PR1/HDB01>

```

## Tenant database recovery

If a user store key has not been created for the pr1adm user at the source system, a key must be created at the target system. The database user configured in the key must have privileges to execute tenant recovery operations.

```

pr1adm@vm-pr1:/usr/sap/PR1/HDB01> hdbuserstore set PR1KEY vm-pr1:30113
<backup-user> <password>

```

The tenant recovery is now executed with hdbsql.

```
pr1adm@vm-pr1:/usr/sap/PR1/HDB01> hdbsql -U PR1KEY
Welcome to the SAP HANA Database interactive terminal.
Type:  \h for help with commands
       \q to quit
hdbsql SYSTEMDB=> RECOVER DATA FOR PR1 USING SNAPSHOT CLEAR LOG
0 rows affected (overall time 66.973089 sec; server time 66.970736 sec)
hdbsql SYSTEMDB=>
```

The HANA database is now up and running, and the disaster recovery workflow for the HANA database has been tested.

#### Recovery with forward recovery using log/catalog backups

Log backups and the HANA backup catalog are being replicated from the source system.

The recovery using all available log backups is executed with the following commands as user pr1adm:

- System database

```
recoverSys.py --command "RECOVER DATABASE UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT"
```

- Tenant database

```
Within hdbsql: RECOVER DATABASE FOR PR1 UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT
```



To recover using all available logs, you can just use any time in the future as the timestamp in the recovery statement.

You can also use HANA Studio or Cockpit to execute the recovery of the system and the tenant database.

The following command output show the recovery execution.

#### System database recovery

```

prladm@vm-pr1:/usr/sap/PR1/HDB01> HDBSettings.sh recoverSys.py --command
"RECOVER DATABASE UNTIL TIMESTAMP '2021-02-20 00:00:00' CLEAR LOG USING
SNAPSHOT"
[140404915394368, 0.008] >> starting recoverSys (at Fri Feb 19 16:06:40
2021)
[140404915394368, 0.008] args: ()
[140404915394368, 0.008] keys: {'command': "RECOVER DATABASE UNTIL
TIMESTAMP '2021-02-20 00:00:00' CLEAR LOG USING SNAPSHOT"}
using logfile /usr/sap/PR1/HDB01/vm-pr1/trace/backup.log
recoverSys started: =====2021-02-19 16:06:40 =====
testing master: vm-pr1
vm-pr1 is master
shutdown database, timeout is 120
stop system
stop system on: vm-pr1
stopping system: 2021-02-19 16:06:40
stopped system: 2021-02-19 16:06:41
creating file recoverInstance.sql
restart database
restart master nameserver: 2021-02-19 16:06:46
start system: vm-pr1
sapcontrol parameter: ['-function', 'Start']
sapcontrol returned successfully:
2021-02-19T16:07:19+00:00 P0009897      177bb0b4416 INFO      RECOVERY
RECOVER DATA finished successfully, reached timestamp 2021-02-
19T15:17:33+00:00, reached log position 38272960
recoverSys finished successfully: 2021-02-19 16:07:20
[140404915394368, 39.757] 0
[140404915394368, 39.758] << ending recoverSys, rc = 0 (RC_TEST_OK), after
39.749 secs

```

## Tenant database recovery

```

prladm@vm-pr1:/usr/sap/PR1/HDB01> hdbsql -U PR1KEY
Welcome to the SAP HANA Database interactive terminal.
Type:  \h for help with commands
       \q to quit

hdbsql SYSTEMDB=> RECOVER DATABASE FOR PR1 UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT
0 rows affected (overall time 63.791121 sec; server time 63.788754 sec)

hdbsql SYSTEMDB=>

```

The HANA database is now up and running, and the disaster recovery workflow for the HANA database has been tested.

### Check consistency of latest log backups

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

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

```
pr1adm@hana-10: > hdbbackupcheck
/hanabackup/PR1/log/SYSTEMDB/log_backup_0_0_0_0.1589289811148
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivercache'
Backup '/mnt/log-backup/SYSTEMDB/log_backup_0_0_0_0.1589289811148'
successfully checked.
```

The check must be executed for the latest log backup files of the system and the tenant database.

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

## Disaster recovery failover

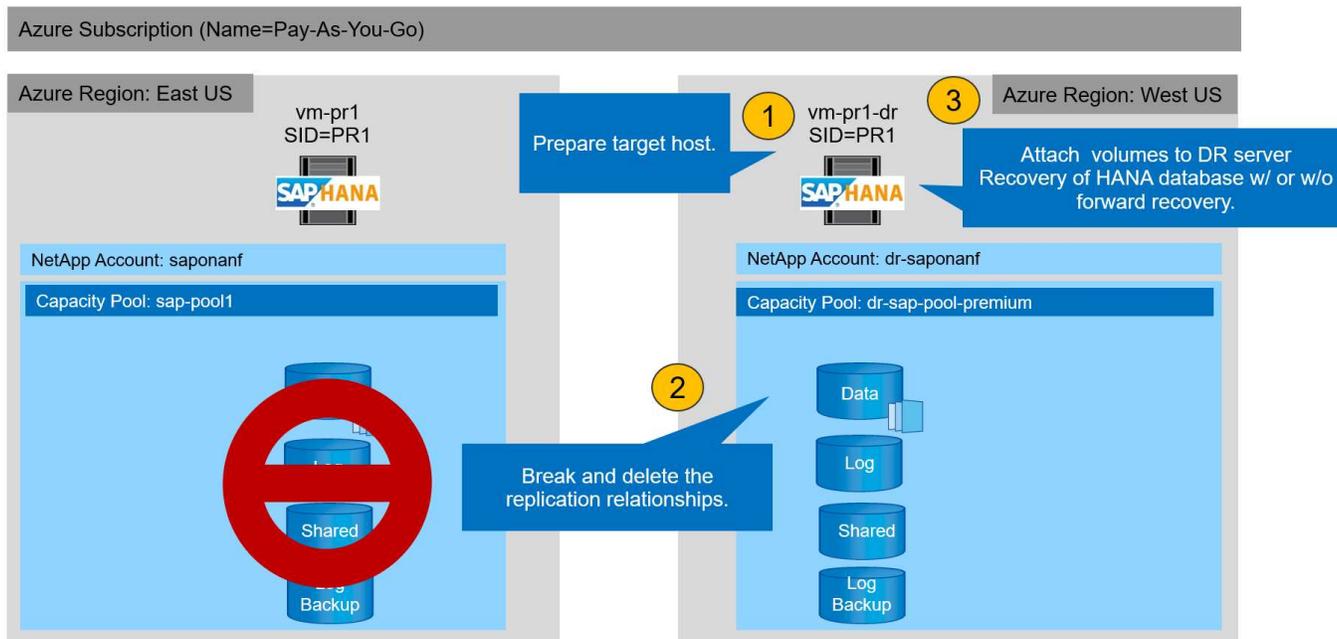
### Disaster recovery failover

Depending on whether the log backup replication is part of the disaster recovery setup, the steps for disaster recovery are slightly different. This section describes the disaster recovery failover for data-backup-only replication as well as for data volume replication combined with log backup volume replication.

To execute disaster recovery failover, complete these steps:

1. Prepare the target host.
2. Break and delete the replication relationships.
3. Restore the data volume to the latest application- consistent snapshot backup.
4. Mount the volumes at the target host.
5. Recover the HANA database.
  - Data volume recovery only.
  - Forward recovery using replicated log backups.

The following subsections describe these steps in detail, and the following figure depicts disaster failover testing.



## Prepare the target host

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

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

The target host preparation also includes shutting down the HANA QA or test system as well as stopping all services using `systemctl stop sapinit`.

### 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, resulting in logically corrupted data.

### Install required software

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



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

## Configure users, ports, and SAP services

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

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

```
vm-pr1:~ # cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/PR1/HDB01/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/PR1/HDB01/exe/sapstartsrv
pf=/usr/sap/PR1/SYS/profile/PR1_HDB01_vm-pr1 -D -u pr1adm
limit.descriptors=1048576
```

## Prepare HANA log volume

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

```
vm-pr1:~ # ls -al /hana/log/PR1/mnt00001/
total 16
drwxrwxrwx 5 root root 4096 Feb 19 16:20 .
drwxr-xr-x 3 root root 22 Feb 18 13:38 ..
drwxr-xr-- 2 pr1adm sapsys 4096 Feb 22 10:25 hdb00001
drwxr-xr-- 2 pr1adm sapsys 4096 Feb 22 10:25 hdb00002.00003
drwxr-xr-- 2 pr1adm sapsys 4096 Feb 22 10:25 hdb00003.00003
vm-pr1:~ #
```

## Prepare log backup volume

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

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

## Prepare file system mounts

The following table shows the naming conventions used in the lab setup. The volume names at the disaster recovery site are included in `/etc/fstab`.

HANA PR1 volumes	Volume and subdirectories at disaster recovery site	Mount point at target host
Data volume	PR1-data-mnt00001-sm-dest	/hana/data/PR1/mnt00001
Shared volume	PR1-shared-sm-dest/shared PR1-shared-sm-dest/usr-sap-PR1	/hana/shared /usr/sap/PR1
Log backup volume	hanabackup-sm-dest	/hanabackup



The mount points from this table must be created at the target host.

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

```
vm-pr1:~ # cat /etc/fstab
# HANA ANF DB Mounts
10.0.2.4:/PR1-data-mnt00001-sm-dest /hana/data/PR1/mnt00001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-log-mnt00001-dr /hana/log/PR1/mnt00001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
# HANA ANF Shared Mounts
10.0.2.4:/PR1-shared-sm-dest/hana-shared /hana/shared nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-shared-sm-dest/usr-sap-PR1 /usr/sap/PR1 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsiz=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
# HANA file and log backup destination
10.0.2.4:/hanabackup-sm-dest /hanabackup nfs
rw,vers=3,hard,timeo=600,rsize=262144,wsiz=262144,nconnect=8,bg,noatime,n
olock 0 0
```

### Break and delete replication peering

In case of a disaster failover, the target volumes must be broken off so that the target host can mount the volumes for read and write operations.



For the HANA data volume, you must restore the volume to the latest HANA snapshot backup created with AzAcSnap. This volume revert operation is not possible if the latest replication snapshot is marked as busy due to the replication peering. Therefore, you must also delete the replication peering.

The next two screenshots show the break and delete peering operation for the HANA data volume. The same operations must be performed for the log backup and the HANA shared volume as well.

ir-sap-pool-premium/PR1-data-mnt0001-sm-dest

### PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest)

Volume

Search (Ctrl+/) Edit Break peering Delete Refresh

Overview Activity log Access control (IAM) Tags

Settings Properties Locks

Storage service Mount instructions Export policy Snapshots Replication

Essentials

End point type : Destination Source  
 Healthy : Healthy Relationship status :  
 Mirror state : Mirrored Replication schedule : Total progress :

Show data for last: 1 hour 6 hours 12 hours 1 day 7 days

Volume replication lag time Is volume replication transfer

Volume replication lag time	Is volume replication transfer
9.72hours	100
8.33hours	90
6.94hours	80
5.56hours	70
	60
	50

### Break replication peering

Break replication peering

Warning! This action will stop data replication between the volumes and might result in loss of data.

Type 'yes' to proceed

yes

ir-sap-pool-premium/PR1-data-mnt0001-sm-dest

### PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest)

Volume

Search (Ctrl+/) Resync Delete Refresh

Overview Activity log Access control (IAM) Tags

Settings Properties Locks

Storage service Mount instructions Export policy Snapshots Replication

Essentials

End point type : Destination Source  
 Healthy : Healthy Relationship status :  
 Mirror state : Broken Replication schedule : Total progress :

Show data for last: 1 hour 6 hours 12 hours 1 day 7 days

Volume replication lag time Is volume replication transfer

Volume replication lag time	Is volume replication transfer
1.67min	100
1.5min	90
1.33min	80
1.17min	70
1min	60
50sec	50

### Delete replication

Delete replication object

Warning this operation will delete the connection between PR1-data-mnt00001 and PR1-data-mnt0001-sm-dest

This will delete the replication object of PR1-data-mnt00001, type 'yes' to proceed

yes

Since replication peering was deleted, it is possible to revert the volume to the latest HANA snapshot backup. If peering is not deleted, the selection of revert volume is grayed out and is not selectable. The following two screenshots show the volume revert operation.

PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest) | Snapshots

Search (Ctrl+/) << + Add snapshot Refresh

- Overview
- Activity log
- Access control (IAM)
- Tags
- Settings
  - Properties
  - Locks
- Storage service
  - Mount instructions
  - Export policy
  - Snapshots**
  - Replication
- Monitoring
  - Metrics
- Automation
  - Tasks (preview)
  - Export template
- Support + troubleshooting
  - New support request

Search snapshots

Name	Location	Created	
azacsnap__2021-02-18T120002-2150721Z	West US	02/18/2021, 01:00:05 PM	...
azacsnap__2021-02-18T160002-1442691Z	West US	02/18/2021, 05:00:49 PM	...
azacsnap__2021-02-18T200002-0758687Z	West US	02/18/2021, 09:00:05 PM	...
azacsnap__2021-02-19T000002-0039686Z	West US	02/19/2021, 01:00:05 AM	...
azacsnap__2021-02-19T040001-8773748Z	West US	02/19/2021, 05:00:06 AM	...
azacsnap__2021-02-19T080001-5198653Z	West US	02/19/2021, 09:00:05 AM	...
azacsnap__2021-02-19T120002-1495322Z	West US	02/19/2021, 01:00:06 PM	...
azacsnap__2021-02-19T160002-3698678Z	West US	02/19/2021, 05:00:05 PM	...
azacsnap__2021-02-22T120002-3145398Z	West US	02/22/2021, 01:00:06 PM	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/22/2021, 03:32:00 PM	...
azacsnap__2021-02-22T160002-0144647Z	West US	02/22/2021, 05:00:05 PM	...
azacsnap__2021-02-22T200002-0649581Z	West US	02/22/2021, 09:00:05 PM	...
azacsnap__2021-02-23T000002-0311379Z	West US	02/23/2021, 01:00:05 PM	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/23/2021, 01:10:00 PM	...

Restore to new volume  
Revert volume  
Delete

PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest) | Snapshots

Search (Ctrl+/) << + Add snapshot Refresh

- Overview
- Activity log
- Access control (IAM)
- Tags
- Settings
  - Properties
  - Locks
- Storage service
  - Mount instructions
  - Export policy
  - Snapshots**
  - Replication
- Monitoring
  - Metrics
- Automation
  - Tasks (preview)
  - Export template
- Support + troubleshooting
  - New support request

Search snapshots

Name	Location
azacsnap__2021-02-18T120002-2150721Z	West US
azacsnap__2021-02-18T160002-1442691Z	West US
azacsnap__2021-02-18T200002-0758687Z	West US
azacsnap__2021-02-19T000002-0039686Z	West US
azacsnap__2021-02-19T040001-8773748Z	West US
azacsnap__2021-02-19T080001-5198653Z	West US
azacsnap__2021-02-19T120002-1495322Z	West US
azacsnap__2021-02-19T160002-3698678Z	West US
azacsnap__2021-02-22T120002-3145398Z	West US
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US
azacsnap__2021-02-22T160002-0144647Z	West US
azacsnap__2021-02-22T200002-0649581Z	West US
azacsnap__2021-02-23T000002-0311379Z	West US
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US

### Revert volume to snapshot

Revert volume PR1-data-mnt0001-sm-dest to snapshot azacsnap\_\_2021-02-23T000002-0311379Z?

**Warning:** This action is irreversible and it will delete all the volumes snapshots that are newer than azacsnap\_\_2021-02-23T000002-0311379Z. Please type 'PR1-data-mnt0001-sm-dest' to confirm.

Are you sure you want to revert 'PR1-data-mnt0001-sm-dest' to state of 'azacsnap\_\_2021-02-23T000002-0311379Z'?

PR1-data-mnt0001-sm-dest ✓

After the volume revert operation, the data volume is based on the consistent HANA snapshot backup and can now be used to execute forward recovery operations.



If a capacity pool with a low performance tier has been used, the volumes must now be moved to a capacity pool that can provide the required performance.

## Mount the volumes at the target host

The volumes can now be mounted at the target host, based on the `/etc/fstab` file created before.

```
vm-pr1:~ # mount -a
```

The following output shows the required file systems.

```
vm-pr1:~ # df
Filesystem                1K-blocks    Used
Available Use% Mounted on
devtmpfs                  8201112         0
8201112    0% /dev
tmpfs                     12313116         0
12313116   0% /dev/shm
tmpfs                     8208744         9096
8199648    1% /run
tmpfs                     8208744         0
8208744    0% /sys/fs/cgroup
/dev/sda4                 29866736 2543948
27322788   9% /
/dev/sda3                 1038336       79984
958352     8% /boot
/dev/sda2                 524008        1072
522936    1% /boot/efi
/dev/sdb1                 32894736     49180
31151556   1% /mnt
10.0.2.4:/PR1-log-mnt00001-dr 107374182400   6400
107374176000 1% /hana/log/PR1/mnt00001
tmpfs                   1641748         0
1641748    0% /run/user/0
10.0.2.4:/PR1-shared-sm-dest/hana-shared 107377178368 11317248
107365861120 1% /hana/shared
10.0.2.4:/PR1-shared-sm-dest/usr-sap-PR1 107377178368 11317248
107365861120 1% /usr/sap/PR1
10.0.2.4:/hanabackup-sm-dest 107379678976 35249408
107344429568 1% /hanabackup
10.0.2.4:/PR1-data-mnt0001-sm-dest 107376511232 6696960
107369814272 1% /hana/data/PR1/mnt00001
vm-pr1:~ #
```

## HANA database recovery

The following shows the steps for HANA database recovery

Start the required SAP services.

```
vm-pr1:~ # systemctl start sapinit
```

The following output shows the required processes.

```
vm-pr1:/ # ps -ef | grep sap
root      23101      1  0 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/saphostexec pf=/usr/sap/hostctrl/exe/host_profile
pr1adm    23191      1  3 11:29 ?          00:00:00
/usr/sap/PR1/HDB01/exe/sapstartsrv
pf=/usr/sap/PR1/SYS/profile/PR1_HDB01_vm-pr1 -D -u pr1adm
sapadm    23202      1  5 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/hostctrl/exe/host_profile -D
root      23292      1  0 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/saposcol -l -w60
pf=/usr/sap/hostctrl/exe/host_profile
root      23359    2597  0 11:29 pts/1      00:00:00 grep --color=auto sap
```

The following subsections describe the recovery process with and without forward recovery using the replicated log backups. The recovery is executed using the HANA recovery script for the system database and hdbsql commands for the tenant database.

### Recovery to latest HANA data volume backup savepoint

The recovery to the latest backup savepoint is executed with the following commands as user pr1adm:

- System database

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

- Tenant database

```
Within hdbsql: RECOVER DATA FOR PR1 USING SNAPSHOT CLEAR LOG
```

You can also use HANA Studio or Cockpit to execute the recovery of the system and the tenant database.

The following command output show the recovery execution.

## System database recovery

```
pr1adm@vm-pr1:/usr/sap/PR1/HDB01> HDBSettings.sh recoverSys.py
--command="RECOVER DATA USING SNAPSHOT CLEAR LOG"
[139702869464896, 0.008] >> starting recoverSys (at Fri Feb 19 14:32:16
2021)
[139702869464896, 0.008] args: ()
[139702869464896, 0.009] keys: {'command': 'RECOVER DATA USING SNAPSHOT
CLEAR LOG'}
using logfile /usr/sap/PR1/HDB01/vm-pr1/trace/backup.log
recoverSys started: =====2021-02-19 14:32:16 =====
testing master: vm-pr1
vm-pr1 is master
shutdown database, timeout is 120
stop system
stop system on: vm-pr1
stopping system: 2021-02-19 14:32:16
stopped system: 2021-02-19 14:32:16
creating file recoverInstance.sql
restart database
restart master nameserver: 2021-02-19 14:32:21
start system: vm-pr1
sapcontrol parameter: ['-function', 'Start']
sapcontrol returned successfully:
2021-02-19T14:32:56+00:00 P0027646      177bab4d610 INFO      RECOVERY
RECOVER DATA finished successfully
recoverSys finished successfully: 2021-02-19 14:32:58
[139702869464896, 42.017] 0
[139702869464896, 42.017] << ending recoverSys, rc = 0 (RC_TEST_OK), after
42.009 secs
pr1adm@vm-pr1:/usr/sap/PR1/HDB01>
```

## Tenant database recovery

If a user store key has not been created for the pr1adm user at the source system, a key must be created at the target system. The database user configured in the key must have privileges to execute tenant recovery operations.

```
pr1adm@vm-pr1:/usr/sap/PR1/HDB01> hdbuserstore set PR1KEY vm-pr1:30113
<backup-user> <password>
```

The tenant recovery is now executed with hdbsql.

```
pr1adm@vm-pr1:/usr/sap/PR1/HDB01> hdbsql -U PR1KEY
Welcome to the SAP HANA Database interactive terminal.
Type:  \h for help with commands
       \q to quit
hdbsql SYSTEMDB=> RECOVER DATA FOR PR1 USING SNAPSHOT CLEAR LOG
0 rows affected (overall time 66.973089 sec; server time 66.970736 sec)
hdbsql SYSTEMDB=>
```

The HANA database is now up and running, and the disaster recovery workflow for the HANA database has been tested.

#### Recovery with forward recovery using log/catalog backups

Log backups and the HANA backup catalog are being replicated from the source system.

The recovery using all available log backups is executed with the following commands as user pr1adm:

- System database

```
recoverSys.py --command "RECOVER DATABASE UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT"
```

- Tenant database

```
Within hdbsql: RECOVER DATABASE FOR PR1 UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT
```



To recover using all available logs, you can just use any time in the future as the timestamp in the recovery statement.

You can also use HANA Studio or Cockpit to execute the recovery of the system and the tenant database.

The following command output show the recovery execution.

#### System database recovery

```

prladm@vm-pr1:/usr/sap/PR1/HDB01> HDBSettings.sh recoverSys.py --command
"RECOVER DATABASE UNTIL TIMESTAMP '2021-02-20 00:00:00' CLEAR LOG USING
SNAPSHOT"
[140404915394368, 0.008] >> starting recoverSys (at Fri Feb 19 16:06:40
2021)
[140404915394368, 0.008] args: ()
[140404915394368, 0.008] keys: {'command': "RECOVER DATABASE UNTIL
TIMESTAMP '2021-02-20 00:00:00' CLEAR LOG USING SNAPSHOT"}
using logfile /usr/sap/PR1/HDB01/vm-pr1/trace/backup.log
recoverSys started: =====2021-02-19 16:06:40 =====
testing master: vm-pr1
vm-pr1 is master
shutdown database, timeout is 120
stop system
stop system on: vm-pr1
stopping system: 2021-02-19 16:06:40
stopped system: 2021-02-19 16:06:41
creating file recoverInstance.sql
restart database
restart master nameserver: 2021-02-19 16:06:46
start system: vm-pr1
sapcontrol parameter: ['-function', 'Start']
sapcontrol returned successfully:
2021-02-19T16:07:19+00:00 P0009897      177bb0b4416 INFO      RECOVERY
RECOVER DATA finished successfully, reached timestamp 2021-02-
19T15:17:33+00:00, reached log position 38272960
recoverSys finished successfully: 2021-02-19 16:07:20
[140404915394368, 39.757] 0
[140404915394368, 39.758] << ending recoverSys, rc = 0 (RC_TEST_OK), after
39.749 secs

```

## Tenant database recovery

```

prladm@vm-pr1:/usr/sap/PR1/HDB01> hdbsql -U PR1KEY
Welcome to the SAP HANA Database interactive terminal.
Type:  \h for help with commands
       \q to quit

hdbsql SYSTEMDB=> RECOVER DATABASE FOR PR1 UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT
0 rows affected (overall time 63.791121 sec; server time 63.788754 sec)

hdbsql SYSTEMDB=>

```

The HANA database is now up and running, and the disaster recovery workflow for the HANA database has been tested.

### Check consistency of latest log backups

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

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

```
pr1adm@hana-10: > hdbbackupcheck
/hanabackup/PR1/log/SYSTEMDB/log_backup_0_0_0_0.1589289811148
Loaded library 'libhdbcsaccessor'
Loaded library 'libhdblivercache'
Backup '/mnt/log-backup/SYSTEMDB/log_backup_0_0_0_0.1589289811148'
successfully checked.
```

The check must be executed for the latest log backup files of the system and the tenant database.

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

## Update history

The following technical changes have been made to this solution since its original publication.

Version	Date	Update summary
Version 1.0	April 2021	Initial version

## TR-4646: SAP HANA Disaster Recovery with Storage Replication

TR-4646 is an overview of the options for disaster recovery protection for SAP HANA. It includes detailed setup information and a use case description of a three-site disaster recovery solution based on synchronous and asynchronous NetApp SnapMirror Storage replication. The described solution uses NetApp SnapCenter with the SAP HANA plug-in to manage database consistency.

Author: Nils Bauer, NetApp

<https://www.netapp.com/pdf.html?item=/media/8584-tr4646pdf.pdf>

# TR-4711: SAP HANA Backup and Recovery Using NetApp Storage Systems and Commvault Software

TR-4711 describes the design of a NetApp and Commvault solution for SAP HANA, which includes Commvault IntelliSnap snapshot management technology and NetApp Snapshot technology. The solution is based on NetApp storage and the Commvault data protection suite.

Authors: Marco Schoen, NetApp; Dr. Tristan Daude, Commvault Systems

<https://www.netapp.com/pdf.html?item=/media/17050-tr4711pdf.pdf>

## SnapCenter Integration for SAP ASE Database

This document describes the SnapCenter integration specifics for SAP ASE Database used in an SAP environment.

### Introduction

The document is not intended to be a step-by-step description of how to setup the complete environment but will cover concepts and relevant details related to:

- Example configuration overview
- Sample Layout
- Protect SAP ASE Instance
- Restore and Recover SAP ASE Instance

Author: Michael Schlosser, NetApp

### Example configuration overview

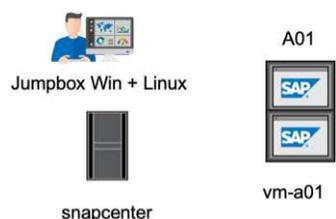
Example Implementation for SnapCenter ASE Plugin for an SAP System running on the Azure Platform.



This implementation describes the minimal required volume configuration. Data Dump Backups and Log Dump Backups are configured according to SAP Note 1588316.

Alternatively, the volume structure described in this [MS Technical Community Blog](#) could be used.

### Demo Environment



## Software versions

Software	Version
Linux OS	SLES for SAP 15 SP5
SAP	SAP NetWeaver 7.5
SAP ASE	16.0 SP04 PL06 HF1
SnapCenter	6.1

## ASE Volume Design

Following least volume Layout must be used to enable backup / recovery and clone use-cases for the SAP ASE database. The example configuration use <SID>: A01.

Volume Name	Directory (qtree) on Volume	Mount Point on Server	Comment
<SID>-sapase	sybase	/sybase	Parent directory for ASE related files
		/sybase/<SID>/backups	Data Dump Backups (might be placed on a different volume)
		/sybase/<SID>/log_archives	Log Dump Backups (might be placed on a different volume)
	<sid>adm	/home/<sid>adm	Home directory of user <sid>adm
	usrsaptrans	/usr/sap/trans	Transport directory
	usrsap<SID>	/usr/sap/<SID>	Usr sap
	sapmnt<SID>	/sapmnt/<SID>	SAP GlobalHost Dir
<SID>-datalog	sapdata_1	/sybase/<SID>/sapdata_1	DB Data (SID)

Volume Name	Directory (qtree) on Volume	Mount Point on Server	Comment
	saplog_1	/sybase/<SID>/saplog_1	DB Log (SID)
	saptemp	/sybase/<SID>/saptemp	PSAPTEMP
	sybsecurity	/sybase/<SID>/sybsecurity	Sybase security DB
	sybssystem	/sybase/<SID>/sybssystem	Sybase system DB
	sybtemp	/sybase/<SID>/sybtemp	Sybase system DB - Temp
	sapdiag	/sybase/<SID>/sapdiag	'saptools' database

### Steps to Protect Database A01

- Check File distribution, according to the sample Layout
- Check Prerequisites for the Host (vm-a01)
- Check Prerequisites for the Database (A01)
- Deploy / Install SnapCenter Agent on Host (vm-a01)
- Create SnapCenter Instance Resource Configuration

### Prerequisites on Host

More current information might be available [here](#).

Before you add a host and install the plug-ins package for Linux, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You can either use the password-based authentication for the root or non-root user or SSH key based authentication.
- SnapCenter Plug-in for Unix File Systems can be installed by a non-root user. However, you should configure the sudo privileges for the non-root user to install and start the plug-in process. After installing the plug-in, the processes will be running as an effective non-root user.
- Create credentials with authentication mode as Linux for the install user.
- You must have installed Java 11 on your Linux host.
- Ensure that you have installed only the certified edition of JAVA 11 on the Linux host
- For information to download JAVA, see: Java Downloads for All Operating Systems
- You should have bash as the default shell for plug-in installation.

### Prerequisites for the Database – Enable Logging and Backups

- Create Directories for backups and log\_archives (/sybase/A01/backups, /sybase/A01/log\_archives)
- Connect to database A01 (as OS-user syba01)
  - isql -S A01 -U sapsa -X -w 1024
- Create Dump configuration for DATA (A01DB) according to SAP Note 1588316
  - use master
  - go

- exec sp\_config\_dump @config\_name='A01DB', @stripe\_dir = '/sybase/A01/backups', @compression = '101', @verify = 'header'
- go
- Create Dump configuration for LOG (A01LOG) according to SAP Note 1588316
  - use master
  - go
  - sp\_config\_dump @config\_name='A01LOG', @stripe\_dir = '/sybase/A01/log\_archives', @compression = '101', @verify = 'header'
  - go
- Enable full logging for Database A01
  - sp\_dboption A01, 'trunc log on chkpt', false
  - go
  - sp\_dboption A01, 'full logging for all', 'true'
  - go
  - sp\_dboption A01, 'enforce dump tran sequence', 'true'
  - go
- Database DUMP Backup to enable Log DUMP Backup
  - dump database A01 using config ='A01DB'
  - go
  - Log Dump
  - dump transaction A01 using config = 'A01LOG'
  - go
- Ensure, that regular Log Backups are configured, according to SAP Note 1588316

### Optional – create dedicated database user

For SAP Environments user sapsa could be used.

- Connect to database A01 (as OS-user syba01)
  - isql -S A01 -U sapsa -X -w 1024
- create user
  - create login backup with password <password>
  - go
- assign permissions / roles to the user
  - grant role sa\_role,sso\_role,oper\_role,sybase\_ts\_role to backup
  - go

### Deploy SnapCenter Agent to Host vm-a01

Further information could be found in the [SnapCenter documentation](#).

Select SAP ASE and Unix File Systems Plugins.

## Add Host

Host Type	Linux	
Host Name	vm-a01	
Credentials	snapcenter-linux	+ ⓘ

## Select Plug-ins to Install SnapCenter Plug-ins Package 6.1 for Linux

- |   |  |
|---|--|
| <input type="checkbox"/> IBM DB2                      | <input type="checkbox"/> MongoDB               |
| <input type="checkbox"/> MySQL                        | <input type="checkbox"/> Oracle Applications ⓘ |
| <input type="checkbox"/> Oracle Database              | <input checked="" type="checkbox"/> SAP ASE    |
| <input type="checkbox"/> PostgreSQL                   | <input type="checkbox"/> SAP MaxDB             |
| <input type="checkbox"/> SAP HANA                     | <input type="checkbox"/> Storage ⓘ             |
| <input checked="" type="checkbox"/> Unix File Systems |  |

 [More Options](#): Port, Install Path, Custom Plug-Ins...

Submit

Cancel

## Create SnapCenter Instance Resource Configuration for Database A01

Resources → SAP ASE → Add Resources

Add SAP ASE Resource
✕

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

### Provide Resource Details

Name  ⓘ

Host Name  ▾

Type  ▾

Credential Name  ▾ + ⓘ

Add information for the credential

Credential Name

Username

Password



If Password contains Special Characters, they must be masked with a backslash.  
 E.g. Test!123! → Test\!123\!

Add SAP ASE Resource
✕

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

### Provide Resource Details

Name  ⓘ

Host Name  ▾

Type  ▾

Credential Name  ▾ + ⓘ

- 1 Name
- 2 Storage Footprint
- 3 Resource Settings
- 4 Summary

### Provide Storage Footprint Details

Storage Type  ONTAP  Azure NetApp Files

Storage Systems for storage footprint

SAP-EastUS	A01-datalog	✎	✕
------------	-------------	---	---

**Modify SAP-EastUS** ✕

Select one or more Capacity pools and their associated Volumes

Capacity pool	Volume
sap-premium-mqos	A01-datalog ✕



If you are using the volume design out of the [MS Technical Community Blog](#).

Volumes /vol<SID>sybase, /vol<SID>data, /vol<SID>log has to be configured as Storage Footprint

Following Resource Settings Custom key-value pairs must be made (at least).

Add SAP ASE Resource
✕

- 1 Name
- 2 Storage Footprint
- 3 Resource Settings
- 4 Summary

**Resource Settings** ⓘ

Custom key-value pairs for SAP ASE plug-in

Name	Value	
<input type="text" value="SYBASE_ISQL_CMD"/>	<input type="text" value="isql -X"/>	✕
<input type="text" value="SYBASE_USER"/>	<input type="text" value="syba01"/>	✕
<input type="text" value="SYBASE_SERVER"/>	<input type="text" value="A01"/>	✕
<input type="text" value="SYBASE_EXCLUDE_TEMPDB"/>	<input type="text" value="Y"/>	✕
<input type="text" value="SYBASE_DATABASES_EXCLUDE"/>	<input type="text" value="saptempdb"/>	+ ✕

Previous

Next

The following table lists the Sybase plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
SYBASE_ISQL_CMD	Example: /opt/sybase/OCS-15__0/bin/isql -X	Defines the path to the isql command. Available Options: <a href="https://infocenter.sybase.com/help/index.jsp?topic=/com.sybase.infocenter.dc34237.1500/html/mvsinst/CIHHFDGC.htm">https://infocenter.sybase.com/help/index.jsp?topic=/com.sybase.infocenter.dc34237.1500/html/mvsinst/CIHHFDGC.htm</a>
SYBASE_USER	user_name	Specifies the operating system user who can run the isql command. This parameter is required for UNIX. This parameter is required if the user running the Snap Creator Agentstart and stop commands (usually the root user) and the user running the isql command are different.
SYBASE_SERVER	data_server_name	Specifies the Sybase data server name (-S option on isql command).For example: A01

Parameter	Setting	Description
SYBASE_DATABASES	db_name:user_name/password	<p>Lists the databases within the instance to back up. The master database is added; for example: DBAtest2:sa/53616c7404351e.If a database named +ALL is used, then database automatic discovery is used, and the sybsyntax, sybssystemdb, sybssystemprocs, and tempdb databases are excluded.</p> <p>For example: +ALL:sa/53616c71a6351e</p> <p>Encrypted passwords are supported if the NTAP_PWD_PROTECTION parameter is set.</p>
SYBASE_DATABASES_EXCLUDE	db_name	Allows databases to be excluded if the +ALL construct is used. You can specify multiple databases by using a semicolon-separated list.For example, pubs2;test_db1
SYBASE_TRAN_DUMP	db_name:directory_path	<p>Enables you to perform a Sybase transaction dump after creating a Snapshot copy.For example: pubs2:/sybasedumps/pubs2</p> <p>You must specify each database that requires a transaction dump.</p>
SYBASE_TRAN_DUMP_FORMAT	%S_%D_%T.cmn	<p>Enables you to specify the dump naming convention. The following keys can be specified:</p> <p>%S = instance name from SYBASE_SERVER</p> <p>%D = database from SYBASE_DATABASES</p> <p>%T = unique timestamp</p> <p>Here is an example: %S_%D_%T.log</p>
SYBASE_TRAN_DUMP_COMPRESS	(Y / N)	Enables or disables native Sybase transaction dump compression.
SYBASE	Example: /Sybase	Specifies the location of the Sybase installation.
SYBASE_MANIFEST	Example: A01:/sybase/A01/sapdiag	Specifies the databases for which the manifest file should be created, along with the location where the manifest file should be placed.
SYBASE_MANIFEST_FORMAT	%S_%D_.manifest Example: %S_%D_.manifest	<p>Enables you to specify the manifest file naming convention. The following keys can be specified:</p> <p>%S = Instance name from SYBASE_SERVER</p> <p>%D = database from SYBASE_DATABASES</p>

Parameter	Setting	Description
SYBASE_MANIFEST_DELETE	(Y / N)	Allows the manifest to be deleted after the Snapshot copy has been created. The manifest file should be captured in the Snapshot copy so that it is always available with the backup.
SYBASE_EXCLUDE_TEMPDB	(Y / N)	Enables automatic exclusion of user-created temporary databases.

### Sequence to Recover System A01

1. stop SAP System A01 (including database), stop sapinit
2. umount Filesystems
3. restore Volumes A01-datalog (using SnapCenter)
4. mount Filesystems
5. start Database A01 (with option -q, to avoid automatic online and keep database forward recoverable – according to SAP Note 1887068)
6. start BackupServer A01
7. online database saptools, sybsecurity , sybmgmtdb
8. recover Database A01 (using isql)
9. online database A01
10. start sapinit, SAP System A01

### Recover Instance A01

- Stop SAP System + DB A01 on host vm-a01
  - User a01adm: stopsap
  - User root: /etc/init.d/sapinit stop
  - User root: umount -a -t nfs
- Restore Backup
  - SnapCenter GUI: Select required Backup for Restore

The screenshot shows the SnapCenter GUI interface for managing copies. At the top, there are navigation tabs: Remove Protection, Back up Now, Modify, Maintenance, Details, and Refresh. Below this, the 'Manage Copies' section shows a diagram with 'Local copies' (6 Backups, 0 Clones) and 'Backups' (3 Backups). A 'Summary Card' on the right displays '9 Backups' and '0 Clones'. The 'Primary Backup(s)' section includes a search bar and a table of backup entries.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_sybase_ondemand_02-07-2025_13_23_21_3633		1	02/07/2025 1:23:58 PM
SnapCenter_sybase_daily_02-07-2025_11_08_28_9176		1	02/07/2025 11:09:07 AM
SnapCenter_sybase_ondemand_02-07-2025_09_31_42_2639		1	02/07/2025 9:32:23 AM
SnapCenter_sybase_daily_02-06-2025_16_35_19_5734		1	02/06/2025 4:36:32 PM
SnapCenter_sybase_ondemand_02-06-2025_16_34_01_6115		1	02/06/2025 4:34:36 PM
SnapCenter_sybase_ondemand_02-06-2025_15_41_33_6630		1	02/06/2025 3:42:21 PM

- For ANF Deployment – only Complete Resource is available



Selecting Complete Resource will trigger a Volume Based Snap Restore (VBSR). Within Azure it is called [volume revert](#).

**Important**

Active filesystem data and snapshots that were taken after the selected snapshot will be lost. The snapshot revert operation will replace *all* the data in the targeted volume with the data in the selected snapshot. You should pay attention to the snapshot contents and creation date when you select a snapshot. You cannot undo the snapshot revert operation.



For other deployment Types (e.g. On-Prem ANF) a Single File Snap Restore (SFSR) operation could be orchestrated. Select File Level and the according Volume and Checkmark "All" – see following screenshot.

Restore from SnapCenter\_sybase\_ondemand\_02-10-2025\_18.16.17.1615

1 Restore scope

2 PreOps

3 PostOps

4 Notification

5 Summary

Select the restore types

Complete Resource ⓘ

File Level ⓘ

Select files to restore

Volume/Qtree	All	File Path
<input checked="" type="checkbox"/> svm-sap01.muccbc.hq.netapp.com:/vol/A0...	<input checked="" type="checkbox"/>	Provide one or more file paths separated by comma
<input type="checkbox"/> svm-sap01.muccbc.hq.netapp.com:/vol/A0...		

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

Previous Next

Summary would be displayed and with Finish the actual restore is started.

Restore from SnapCenter\_sybase\_ondemand\_02-07-2025\_13\_23\_21\_3633
✕

- 1 Restore scope
- 2 PreOps
- 3 PostOps
- 4 Notification
- 5 Summary

### Summary

Backup Name	SnapCenter_sybase_ondemand_02-07-2025_13_23_21_3633
Backup date	02/07/2025 1:23:58 PM
Restore scope	Complete Resource
Pre restore command	
Unmount command	
Mount command	
Post restore command	
Send email	No

⚠ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous
Finish

- Mount Filesystems (vm-a01)
  - User root: mount -a -t nfs
- Start Database A01 + BackupServer
  - Modify RUN\_A01 and add -q \ (according to SAP Note 1887068)
  - User syba01: RUN\_A01 &
  - User syba01: RUN\_A01\_BS&
- Online databases saptools, sybsecurity , sybmgmtdb
  - User syba01: isql -S A01 -U sapsa -X -w 1024
  - online database saptools
  - go
  - online database sybsecurity
  - go
  - online database sybmgmtdb
  - go



### External Documentation

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

- [SAP Installation Azure on ANF](#)
- [SnapCenter Prerequisites for Plugins](#)
- [SnapCenter Install Plugins](#)
- [Sybase Infocenter - isql](#)
- [Sybase Infocenter - load transaction log dumps](#)
- SAP Notes (login required)
  - 1887068 - SYB: Using external backup and restore with SAP ASE: <https://me.sap.com/notes/1887068/E>
  - 1618817 - SYB: How to restore an SAP ASE database server (UNIX): <https://me.sap.com/notes/1618817/E>
  - 1585981 - SYB: Ensuring Recoverability for SAP ASE: <https://me.sap.com/notes/1585981/E>
  - 1588316 - SYB: Configure automatic database and log backups: <https://me.sap.com/notes/1588316/E>
  - NetApp Product Documentation: <https://www.netapp.com/support-and-training/documentation/>
  - [NetApp SAP Solutions – Informations about Use-Cases, Best-Practices and Benefits](#)

### Version history

Version	Date	Document version history
Version 1.0	April 2025	Initial version – backup / recovery ASE database

## SnapCenter Integration for IBM DB2 Database

This document describes the SnapCenter integration specifics for IBM DB2 Database used in an SAP environment.

### Introduction

The document is not intended to be a step-by-step description of how to setup the complete environment but will cover concepts and relevant details related to:

- Example configuration overview
- Sample Layout
- Protect DB2 database
- Restore and Recover DB2 database

Author: Michael Schlosser, NetApp

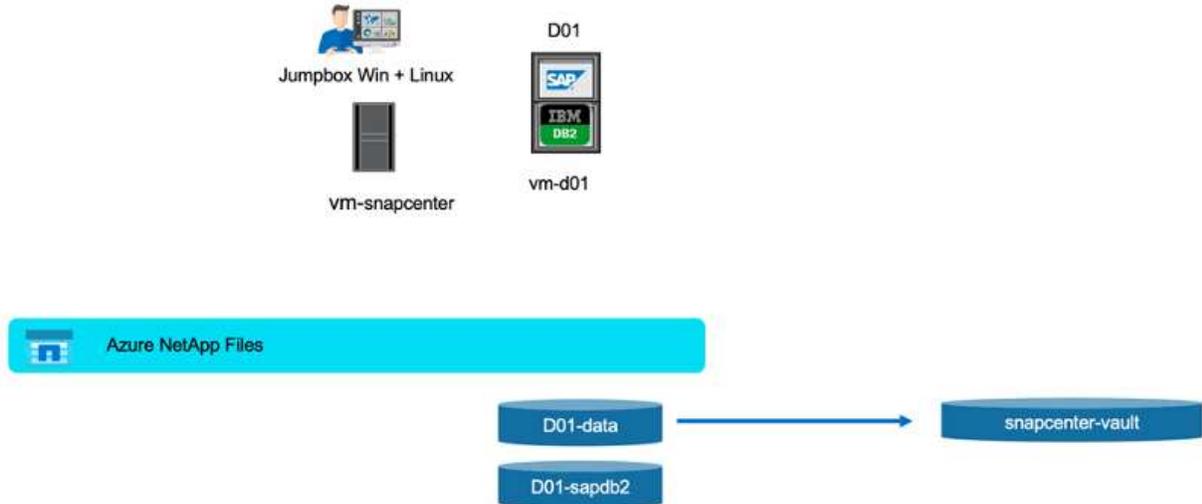
## Example configuration overview

Example Implementation for SnapCenter DB2 Plugin for an SAP System running on the Azure Platform.



This implementation describes the minimal required volume configuration.

Alternatively, the volume structure described in this [MS Technical Community blog](#) could be used.



## Demo Environment

### Software versions

Software	Version
Linux OS	SLES for SAP 15 SP5
SAP	SAP NetWeaver 7.5
DB2	10.5.0.7
SnapCenter	6.1

### DB2 Volume Design

Following least volume Layout must be used to enable backup / recovery and clone use-cases for the DB2 database. The example configuration use <SID>: D01.

Volume Name	Directory (qtree) on Volume	Mount Point on Server	Comment
<SID>-sapdb2	db2	/db2	
		/db2/<SID>	Parent directory for DB2 related files

Volume Name	Directory (qtree) on Volume	Mount Point on Server	Comment
		/db2/db2<sid>	Home directory of user db2<sid> and DB2 Software
		/db2/<SID>/db2dump	DB2 diagnostic log and dump files
		/db2/<SID>/backup	Backup dlocation (might be placed on a different volume)
		/db2/<SID>/log_arch	Offline Redo Logs (might be placed on a different volume – snapshot will be triggered)
		/db2/<SID>/log_dir	Online Redo Logs (might be placed on a different volume – snapshot will be triggered)
	<sid>adm	/home/<sid>adm	Home directory of user <sid>adm
	sap<sid>	/home/sap<sid>	Home directory of user sap<sid>
	usrsaptrans	/usr/sap/trans	Transport directory
	usrsap<SID>	/usr/sap/<SID>	Usr sap
	sapmnt<SID>	/sapmnt/<SID>	SAP GlobalHost Dir
<SID>-data	sapdata1	/db2/<SID>/sapdata1	DB Data
	sapdata2	/db2/<SID>/sapdata2	DB Data
	sapdata3	/db2/<SID>/sapdata3	DB Data
	sapdata4	/db2/<SID>/sapdata4	DB Data
	saptmp1	/db2/<SID>/saptmp1	DB Temp Files
	saptmp2	/db2/<SID>/saptmp2	DB Temp Files
	saptmp3	/db2/<SID>/saptmp3	DB Temp Files
	saptmp4	/db2/<SID>/saptmp4	DB Temp Files
	db2<sid>	/db2/<SID>/db2<sid>	Instance Files

Because auto-discovery is enabled by default for the DB2 plug-in, a snapshot is created for volumes that match the following file paths.

Database StoragePath	/db2/D01/saptmp4/, /db2/D01/saptmp3/, /db2/D01/saptmp2/, /db2/D01/saptmp1/, /db2/D01/sapdata4/, /db2/D01/sapdata3/, /db2/D01/sapdata2/, /db2/D01/sapdata1/
Database LogPath	/db2/D01/log_dir/NODE0000/LOGSTREAM0000/
Database Archive Path (Primary)	DISK:/db2/D01/log_arch/

## Steps to Protect Database D01

- Check File distribution, according to the sample Layout
- Check Prerequisites for the Host (vm-d01)
- Check Prerequisites for the Database (D01)
- Deploy / Install SnapCenter Agent on Host (vm-d01)
- Create SnapCenter Instance Resource Configuration

## Prerequisites on Host

More current information might be available here:

- [https://docs.netapp.com/us-en/snapcenter/protect-scu/reference\\_prerequisites\\_for\\_adding\\_hosts\\_and\\_installing\\_snapcenter\\_plug\\_ins\\_package\\_for\\_linux.html](https://docs.netapp.com/us-en/snapcenter/protect-scu/reference_prerequisites_for_adding_hosts_and_installing_snapcenter_plug_ins_package_for_linux.html)
- <https://docs.netapp.com/us-en/snapcenter/protect-db2/prerequisites-for-using-snapcenter-plug-in-for-ibm-db2.html>

Before you add a host and install the plug-ins package for Linux, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You can either use the password-based authentication for the root or non-root user or SSH key based authentication.
- SnapCenter Plug-in for Unix File Systems can be installed by a non-root user. However, you should configure the sudo privileges for the non-root user to install and start the plug-in process. After installing the plug-in, the processes will be running as an effective non-root user.
- Create credentials with authentication mode as Linux for the install user.
- You must have installed Java 11 on your Linux host.
- Ensure that you have installed only the certified edition of JAVA 11 on the Linux host
- For information to download JAVA, see: Java Downloads for All Operating Systems
- You should have bash as the default shell for plug-in installation.

## Prerequisites for the Database – Enable Logging and Backups



to enable offline logs a offline full backup of the database is required. Typically it is already enabled for productive systems.

- Create Directories for backup and log\_arch (/db2/D01/backup, /sybase/D01/log\_arch)
- Enable logarchmeth1 (as OS-user db2d01)
  - db2 update db cfg for D01 using logarchmeth1 DISK:/db2/D01/log\_arch/
- Create offline backup (as OS-user db2d01)
  - db2stop force
  - db2start admin mode restricted access
  - db2 backup db D01 to /db2/D01/backup
  - db2 activate db D01

## Deploy SnapCenter Agent to Host vm-d01

Further information could be found in the [SnapCenter documentation](#).

Select IBM DB2 and Unix File Systems Plugins.

### Add Host

Host Type:

Host Name:

Credentials:   

#### Select Plug-ins to Install SnapCenter Plug-ins Package 6.1 for Linux

<input checked="" type="checkbox"/> IBM DB2	<input type="checkbox"/> MongoDB
<input type="checkbox"/> MySQL	<input type="checkbox"/> Oracle Applications 
<input type="checkbox"/> Oracle Database	<input type="checkbox"/> SAP ASE
<input type="checkbox"/> PostgreSQL	<input type="checkbox"/> SAP MaxDB
<input type="checkbox"/> SAP HANA	<input type="checkbox"/> Storage 
<input checked="" type="checkbox"/> Unix File Systems	

 [More Options](#) : Port, Install Path, Custom Plug-Ins...



After the installation a discovery of the Databases on the host is triggered.

ID	Status	Name
189	✓	Discover resources for host 'vm-d01.1h05kdpkcgauj4qsseqldygg.bx.internal.cloudapp.net'
188	✓	Discover resources for host 'vm-d01.1h05kdpkcgauj4qsseqldygg.bx.internal.cloudapp.net'
187	⚠	Package Installation on host 'vm-d01.1h05kdpkcgauj4qsseqldygg.bx.internal.cloudapp.net'
186	✓	Add Host 'vm-d01.1h05kdpkcgauj4qsseqldygg.bx.internal.cloudapp.net'
185	✓	Validate Host 'vm-d01.1h05kdpkcgauj4qsseqldygg.bx.internal.cloudapp.net'

## Create Resource Configuration for Database D01

Select discovered Resource D01

Name	Type	Instance	Host	Resource Groups	Policies	Last backup	Overall Status
D01	Database	db2d01	vm-d01.1h05kdpkcgauj4qsseqldygg.bx.internal.cloudapp.net				Not protected

Configure Snapshot Name

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

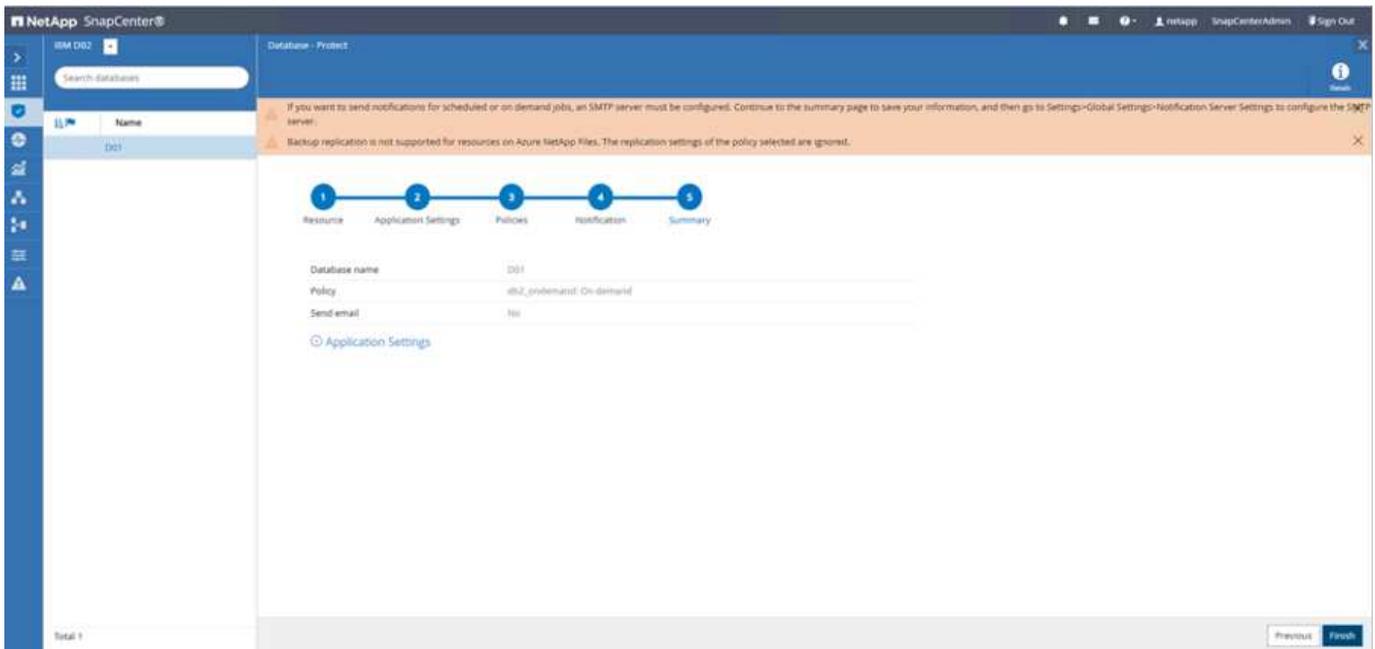
Use custom name format for Snapshot copy

Format: `&CustomText; &Policy`

Example: SnapCenter

Previous Next

No specific application settings required, configure policy and notification settings as required.



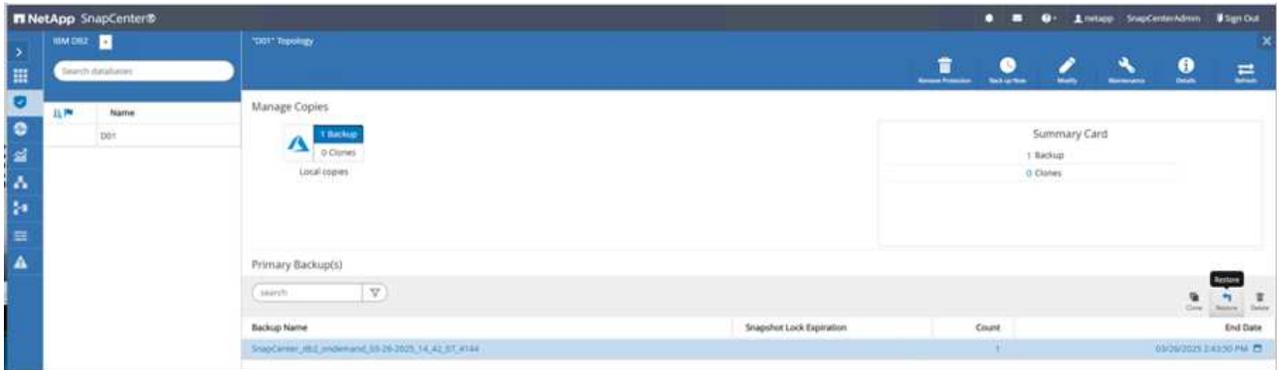
And finish the configuration.

### Sequence to Recover System D01

1. Stop SAP System D01 (including database)
2. Restore SnapCenter Backup (Volume D01-data)
  - a. Unmount Filesystems
  - b. Restore Volume
  - c. Mount Filesystems
  - d. Init database as mirror db
3. Recover Database D01 (using db2 rollforward)
4. Start SAP System D01

### Recover Database D01

- Stop SAP System + DB D01 on host vm-d01
  - User d01adm: stopsap
- Restore Backup
  - SnapCenter GUI: Select required Backup for Restore



- For ANF Deployment – only Complete Resource is available



Summary would be displayed and with Finish the actual restore is started.

Restore from SnapCenter\_db2\_ondemand\_03-26-2025\_14\_42\_07\_4144
✕

- 1 Restore scope
- 2 PreOps
- 3 PostOps
- 4 Notification
- 5 Summary

### Summary

Backup Name	SnapCenter_db2_ondemand_03-26-2025_14_42_07_4144
Backup date	03/26/2025 2:43:50 PM
Restore scope	Complete Resource without Volume Revert
Pre restore command	
Post restore command	
Send email	No

⚠ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous
Finish



“db2inidb D01 as mirror” is done as part of SnapCenter Restore Workflow.

- Check recover status Database D01 (as user db2d01)
  - db2 rollforward db D01 query status
- Recover database as needed – here an losless recovery is initiated (as user db2d01)
  - db2 rollforward db D01 to end of logs
- Stop database recovery and online database D01 (as user db2d01)
  - db2 rollforward db D01 stop
- Start SAP System (as user d01adm)
  - startsap

## Additional information and version history

Following recoded Demos are available to support the documentation.

[Installation and Configuration DB2 Plugin, Backup of DB2 database](#)

## Restore and Recovery of DB2 database

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

- [SAP on DB2 Installation Azure on ANF](#)
- [SnapCenter Prerequisites for Plugins](#)
- [SnapCenter Install Plugins](#)
- [SnapCenter DB2 Plugin Documentation](#)
- SAP Notes (login required)
  - 83000 - DB2/390: Backup and Recovery Options: <https://me.sap.com/notes/83000>
  - 594301 - DB6: Admin Tools and Split Mirror: <https://me.sap.com/notes/594301>
- NetApp Product Documentation: <https://www.netapp.com/support-and-training/documentation/>
- [NetApp SAP Solutions – Information about Use-Cases, Best-Practices and Benefits](#)

### Version history

Version	Date	Document version history
Version 1.0	April 2025	Initial version – backup / recovery DB2 database

## SnapCenter Integration for SAP MaxDB Database

This document describes the SnapCenter integration specifics for SAP MaxDB Database used in an SAP environment.

### Introduction

The document is not intended to be a step-by-step description of how to setup the complete environment but will cover concepts and relevant details related to:

- Example configuration overview
- Sample Layout
- Protect SAP MaxDB Instance
- Restore and Recover SAP MaxDB Instance

### Example configuration overview

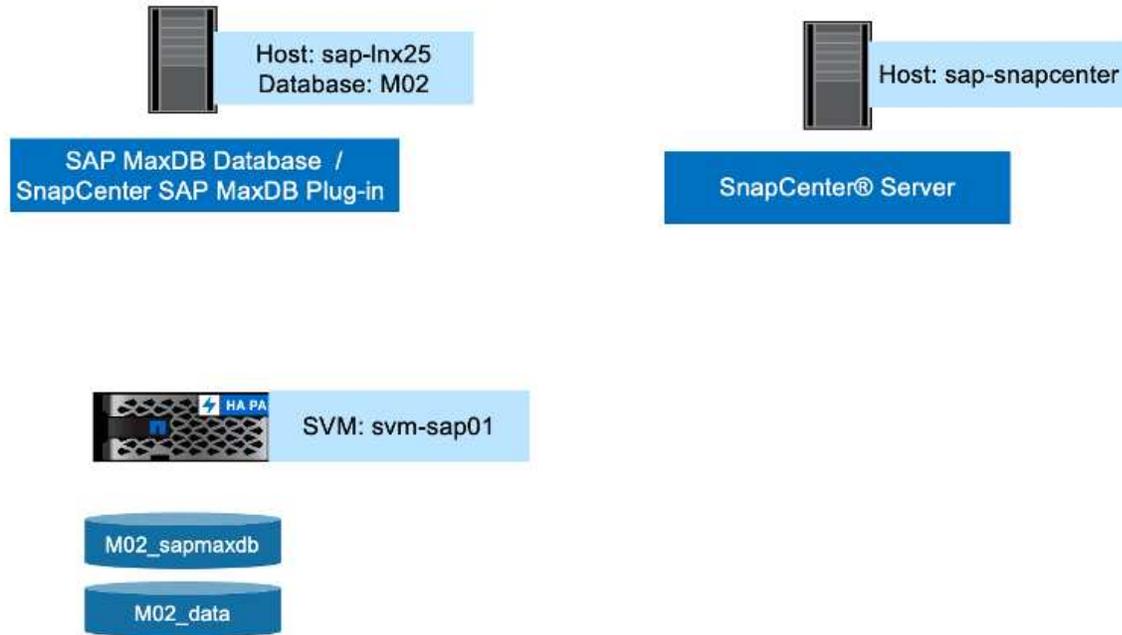
Example Implementation for SnapCenter MaxDB Plugin for an SAP System running in our Demo Center.



This implementation describes the minimal required volume configuration. Data Dump Backups and Log Dump Backups, Backup Template, etc. are configured according to SAP Note “1928060 - Data backup and recovery with file system backup” and referenced Notes from there.

Alternatively, the volume structure described in [MS Techcommunity Blog](#) could be used.

## Demo Environment



## Software versions

Software	Version
Linux OS	SLES for SAP 15 SP5
SAP	SAP NetWeaver 7.5
SAP MaxDB	DBMServer 7.9.10 Build 004-123-265-969
SnapCenter	6.1

## MaxDB Volume Design

Following least volume Layout must be used to enable backup / recovery and clone use-cases for the SAP MaxDB database. The example configuration use <SID>: M02.

Volume Name	Directory (qtree) on Volume	Mount Point on Server	Comment
<SID>_sapmaxdb	sapdb	/sapdb	Parent directory for MaxDB related files
		/sapdb/<SID>/saplog	Redo Logs (might be placed on a different volume)
		/sapdb/<SID>/backup	Dump Backups (Data + Log) (might be placed on a different volume)
	<sid>adm	/home/<sid>adm	Home directory of user <sid>adm
	sdb	/home/sdb	Home directory of User sdb
	sqd<sid>	/home/sqd<sid>	Home directory of User sqd<sid>

Volume Name	Directory (qtree) on Volume	Mount Point on Server	Comment
	usrsaptrans	/usr/sap/trans	Transport directory
	usrsap<SID>	/usr/sap/<SID>	Usr sap
	sapmnt<SID>	/sapmnt/<SID>	SAP GlobalHost Dir
<SID>_data	sapdata	/sapdb/<SID>/sapdata	DB Data Files (SID)

## Steps to Protect Database M02

- Check File distribution, according to the sample Layout
- Check Prerequisites for the Host (sap-Inx25)
- Check Prerequisites for the Database (M02)
- Deploy / Install SnapCenter Agent on Host (sap-Inx25)
- Create SnapCenter Instance Resource Configuration

## Prerequisites on Host

More current information might be available [here](#).

Before you add a host and install the plug-ins package for Linux, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You can either use the password-based authentication for the root or non-root user or SSH key based authentication.
- SnapCenter Plug-in for Unix File Systems can be installed by a non-root user. However, you should configure the sudo privileges for the non-root user to install and start the plug-in process. After installing the plug-in, the processes will be running as an effective non-root user.
- Create credentials with authentication mode as Linux for the install user.
- You must have installed Java 11 on your Linux host.
- Ensure that you have installed only the certified edition of JAVA 11 on the Linux host
- For information to download JAVA, see: Java Downloads for All Operating Systems
- You should have bash as the default shell for plug-in installation.

## Prerequisites for the Database – Create Backup Templates, Enable Logbackup

- Create Directories for data and log backups (/sapdb/M02/backup/data, /sapdb/M02/backup/log – owner sdb:sdba – Permissions 755)
- Connect to database M02 (as OS-user sqdm02)
  - dbmcli -d M02 -u CONTROL,<password>
- Create Data File Backup Template (M02\_DATA) according to SAP Note 1928060
  - backup\_template\_create M02\_DATA to FILE /sapdb/M02/backup/data/M02\_DATA content DATA
- Create Data Backup Template (M02\_LOG) according to SAP Note 1928060

- backup\_template\_create M02\_LOG to FILE /sapdb/M02/backup/log/M02\_LOG content LOG
- Create Data Snapshot Backup Template (M02\_SNAP) according to SAP Note 1928060
  - backup\_template\_create M02\_SNAP to EXTERNAL SNAPSHOT
- Create Fake-Backup to enable LOG Backup
  - util\_connect
  - backup\_start M02\_SNAP
  - backup\_finish M02\_SNAP ExternalBackupID first\_full\_fake\_backup
- Switch Database Logging Mode
  - autolog\_off
  - autolog\_on M02\_LOG INTERVAL 300
  - autolog\_show

## Deploy SnapCenter Agent to Host sap-Inx25

Further Information could be found in the [SnapCenter documentation](#).

Select SAP MaxDB and Unix File Systems Plugins.

## Add Host

Host Type

Host Name

Credentials   

### Select Plug-ins to Install SnapCenter Plug-ins Package 6.1 for Linux

- |   |  |
|---|--|
| <input type="checkbox"/> IBM DB2                      | <input type="checkbox"/> MongoDB   |
| <input type="checkbox"/> MySQL                        | <input type="checkbox"/> Oracle Applications  |
| <input type="checkbox"/> Oracle Database              | <input type="checkbox"/> SAP ASE   |
| <input type="checkbox"/> PostgreSQL                   | <input checked="" type="checkbox"/> SAP MaxDB  |
| <input type="checkbox"/> SAP HANA                     | <input type="checkbox"/> Storage              |
| <input checked="" type="checkbox"/> Unix File Systems |  |

 [More Options](#): Port, Install Path, Custom Plug-Ins...

## Create SnapCenter Resource Configuration for Database M02

Resources → SAP MaxDB → Add Resources

Add SAP MaxDB Resource

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

Provide Resource Details

Name: M02

Host Name: sap-lnx25.muccbc.hq.netapp.com

Type: Database

Credential Name: None

Add information for the credential

Credential Name: control-M02

Username: control

Password: .....

Add

Previous Next



If Password contains Special Characters, they must be masked with a backslash (e.g. Test!123! → Test!\!123!\!).

Add SAP MaxDB Resource

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

Provide Resource Details

Name: M02

Host Name: sap-lnx25.muccbc.hq.netapp.com

Type: Database

Credential Name: control-M02

Add SAP MaxDB Resource

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

Provide Storage Footprint Details

Storage Type  ONTAP  Azure NetApp Files

Add Storage Footprint

Storage System

Select one or more volumes and if required their associated Qtrees and LUNs

Volume name	LUNs or Qtrees
<input type="text" value="M02_data"/>	<input type="text" value="Default is 'None' or type to find"/>
<input type="text" value="M02_sapmaxdb"/>	<input type="text" value="Default is 'None' or type to find"/>

+ x

Save

Following Resource Settings Custom key-value pairs must be made (at least).

Add SAP MaxDB Resource
✕

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

Resource Settings i

Custom key-value pairs for SAP MaxDB plug-in ^

Name	Value	
DBMCLICMD	/sapdb/M02/db/bin/dbmcli	✕
SQLCLICMD	/sapdb/M02/db/bin/sqlcli	✕
MAXDB_UPDATE_HIST_LOG	Y	✕
MAXDB_BACKUP_TEMPLATES	M02:M02_\$SNAP	+ ✕

Previous
Next

The following table lists the MaxDB plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
HANDLE_LOGWRITER	(Y / N)	Executes suspend logwriter (N) or resume logwriter (Y) operations.
DBMCLICMD	path_to_dbmcli_cmd	Specifies the path to the MaxDB dbmcli command. If not set, dbmcli on the search path is used.
SQLCLICMD	path_to_sqlcli_cmd	Specifies the path for the MaxDB sqlcli command. If not set, sqlcli is used on the search path.
MAXDB_UPDATE_HIST_LOG	(Y / N)	Instructs the MaxDB backup program whether or not to update the MaxDB history log.

Parameter	Setting	Description
MAXDB_BACKUP_TEMP LATES	template_name (e.g. M02_SNAP)	Specifies a backup template for each database. The template must already exist and be an external type of backup template.  To enable Snapshot copy integration for MaxDB 7.8 and later, you must have MaxDB background server functionality and already configured MaxDB backup template.
MAXDB_BG_SERVER_PREFIX	bg_server_prefix (e.g. na_bg)	Specifies the prefix for the background server name. If the MAXDB_BACKUP_TEMPLATES parameter is set, you must also set the MAXDB_BG_SERVER_PREFIX parameter. If you do not set the prefix, the default value na_bg_DATABASE is used.

Add SAP MaxDB Resource
✕

- 1 Name
- 2 Storage Footprint
- 3 Resource Settings
- 4 Summary

### Summary

Name	M02
Type	Database
Host	sap-lnx25.muccbc.hq.netapp.com
Credential Name	control-M02

Storage Footprint

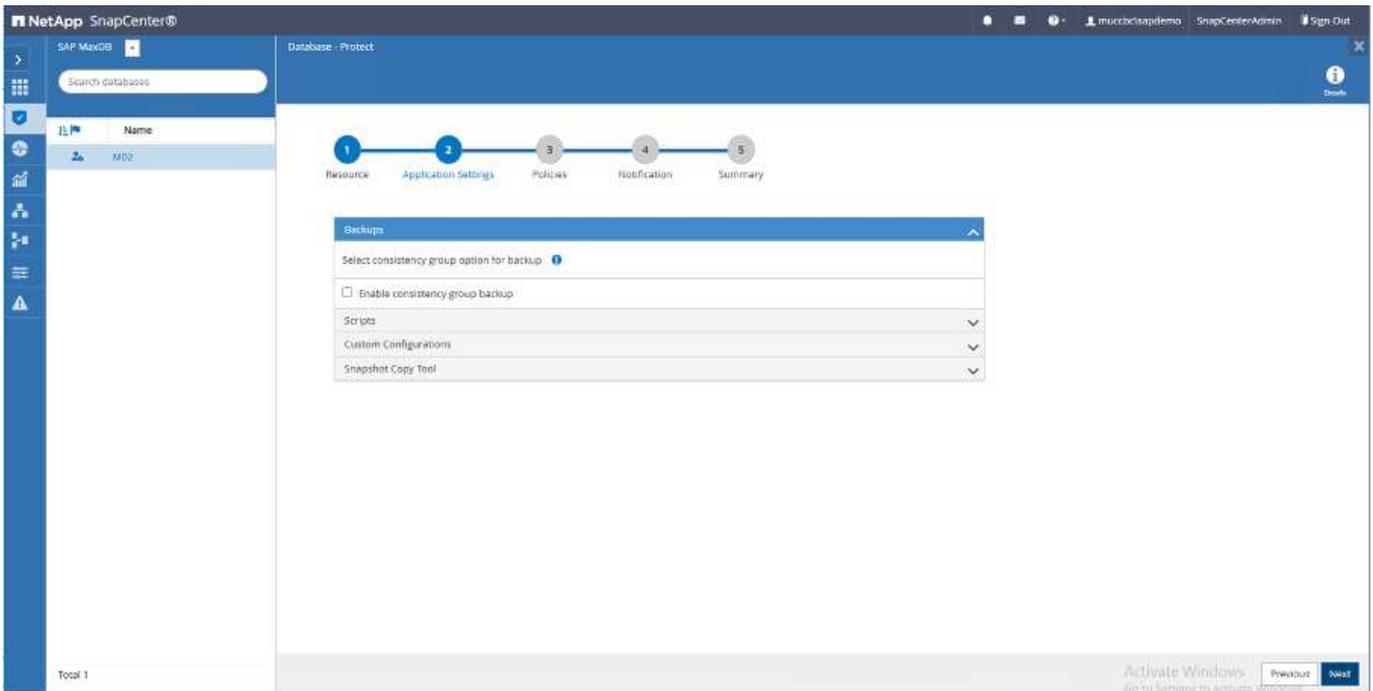
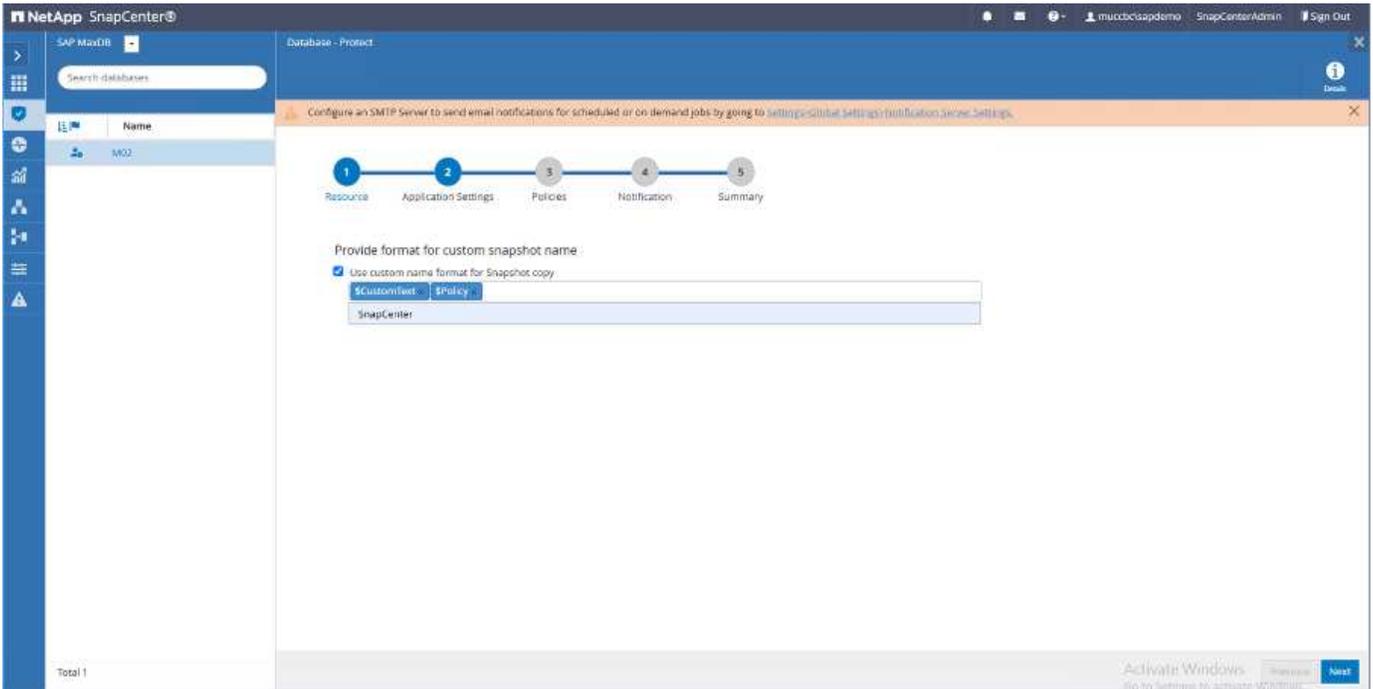
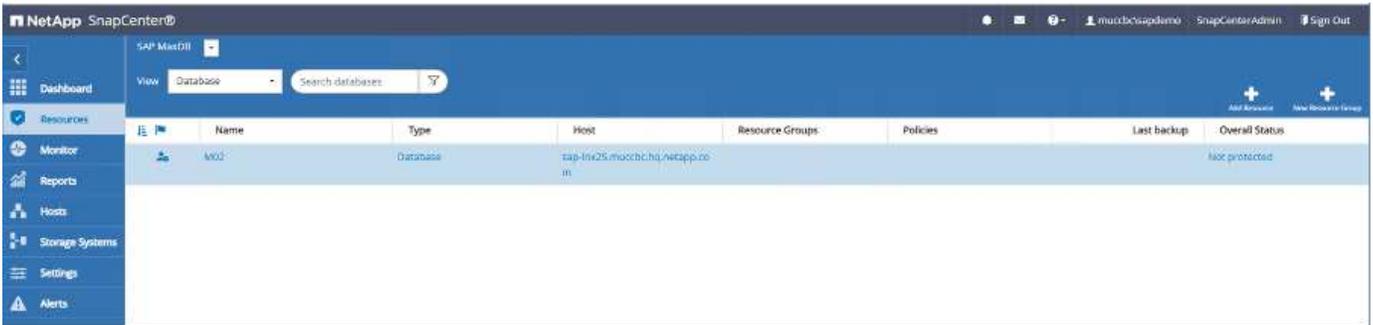
Storage System	Volume	LUN/Qtree
svm-sap01.muccbc.hq.netapp.com	M02_data	
	M02_sapmaxdb	

Custom Resource Parameters

Key	Value
DBMCLICMD	/sapdb/M02/db/bin/dbmcli
SQLCLICMD	/sapdb/M02/db/bin/sqlcli
MAXDB_UPDATE_HIST_LOG	Y
MAXDB_BACKUP_TEMPLATES	M02:M02_SNAP

Previous
Finish

Now the configuration could be finished and Backup scheduled according to the overall protection concept.



NetApp SnapCenter Database - Protect

SAP MaxDB

Search databases

Name

M02

1 Resource 2 Application Settings 3 Policies 4 Notification 5 Summary

Select one or more policies and configure schedules

maxdb\_ondemand +

Configure schedules for selected policies

Policy	Applied Schedules	Configure Schedules	Secondary Protection
maxdb_ondemand	None	To schedule operations select a policy that has the appropriate schedule associated, or modify the selected policy to allow schedules.	No

Total 1

Activate Windows  
Go to Settings to activate Windows.

Previous Next

NetApp SnapCenter Database - Protect

SAP MaxDB

Search databases

Name

M02

1 Resource 2 Application Settings 3 Policies 4 Notification 5 Summary

If you want to send notifications for scheduled or on-demand jobs, an SMTP server must be configured. Continue to the summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Provide email settings

Select the service accounts or people to notify regarding protection issues.

Email preference: Never

From: From email

To: Email to

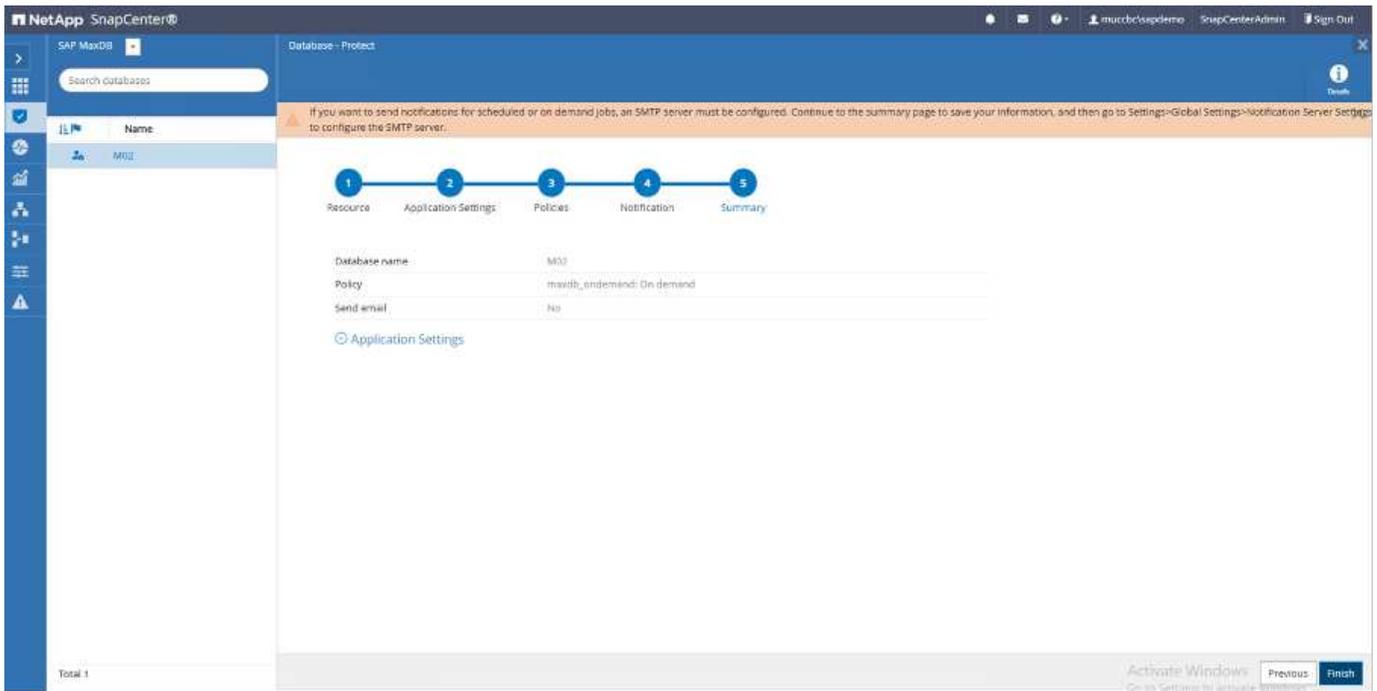
Subject: Notification

Attach job report

Total 1

Activate Windows  
Go to Settings to activate Windows.

Previous Next

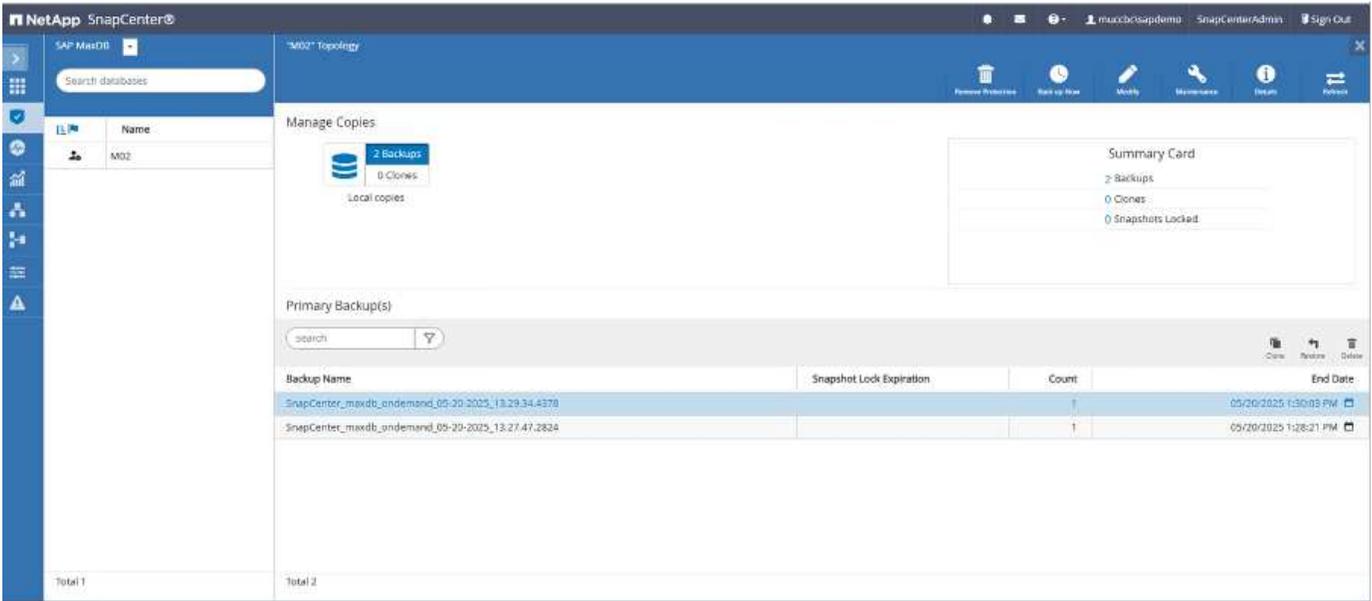


## Sequence to Recover System M02

1. stop SAP System M02 (including database), stop sapinit
2. umount Filesystem /sapdb/M02/sapdata
3. restore Volumes M02\_data (using SnapCenter)
4. mount Filesystem /sapdb/M02/sapdata
5. start Database M02 and connect (admin mode)
6. Gather Backup Information
7. recover database data backup
8. recover database log backups
9. stop database
10. start sapinit, SAP System M02

## Recover Instance M02

- Stop SAP System + DB M02 on host sap-lnx25
  - User m02adm: stopsap
  - Optional – if database has not been stopped successfully – User: sqdm02
  - dbmcli -d M02 -u CONTROL,<password>
    - db\_offline
  - User root: /etc/init.d/sapinit stop
  - User root: umount /sapdb/M02/sapdata
- Restore Backup
  - SnapCenter GUI: Select required Backup for Restore



Selecting Complete Resource will trigger a Volume Based Snap Restore (VBSR). Within Azure it is called [volume revert](#). For ANF Deployment **only Complete Resource is available**.

#### Important

Active filesystem data and snapshots that were taken after the selected snapshot will be lost. The snapshot revert operation will replace *all* the data in the targeted volume with the data in the selected snapshot. You should pay attention to the snapshot contents and creation date when you select a snapshot. You cannot undo the snapshot revert operation.



For other deployment Types (e.g. On-Prem ANF) a Single File Snap Restore (SFSR) Operation could be orchestrated. Select File Level and the according Volume and Checkmark "All" – see following screenshot.

Restore from SnapCenter\_maxdb\_ondemand\_05-20-2025\_13.29.34.4378

**1 Restore scope**

2 PreOps

3 PostOps

4 Notification

5 Summary

### Select the restore types

Complete Resource ⓘ

File Level ⓘ

### Select files to restore

Volume/Qtree	All	File Path
<input checked="" type="checkbox"/> svm-sap01.muccbc.hq.netapp.com:/vol/M...	<input checked="" type="checkbox"/>	Provide one or more file paths separated by comma
<input type="checkbox"/> svm-sap01.muccbc.hq.netapp.com:/vol/M...	<input type="checkbox"/>	

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

Previous
Next

Summary would be displayed and with Finish the actual restore is started.

Restore from SnapCenter\_maxdb\_ondemand\_05-20-2025\_13.29.34.4378
✕

- 1 Restore scope
- 2 PreOps
- 3 PostOps
- 4 Notification
- 5 Summary

### Summary

Backup Name	SnapCenter_maxdb_ondemand_05-20-2025_13.29.34.4378
Backup date	05/20/2025 1:30:03 PM
Restore scope	File Level
Pre restore command	
Unmount command	
Mount command	
Post restore command	
Send email	No

⚠ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous
Finish

- Mount Filesystems (sap-Inx25)
  - User root: mount /sapdb/M02/sapdata
- Start Database M02 in admin mode an connect
  - User: sqdm02: dbmcli -d M02 -u CONTROL,<password>
    - db\_admin
    - db\_connect
- Gather Backup Information
  - backup\_history\_open
  - backup\_history\_list -c label,action,pages,stop,media -r last

```
[dbmcli on M02>backup_history_list -c label,action,pages,stop,media -r last
OK
END
DAT_000000008|SAVE WARM|          0|2025-05-20 13:29:50|M02_SNAP
---
```

- Recover Database

- Recover Data Backup
  - recover\_start M02\_SNAP data ExternalBackupID DAT\_000000008

```
[dbmcli on M02>recover_start M02_SNAP data ExternalBackupID DAT_000000008
OK
Returncode                0
Date                      20250520
Time                      00151550
Server                    sap-lnx25
Database                  M02
Kernel Version            Kernel    7.9.10    Build 004-123-265-969
Pages Transferred         0
Pages Left
Volumes
Medianame                 M02_SNAP
Location
ErrorText
Label                    DAT_000000008
Is Consistent             true
First LOG Page           512226
Last LOG Page
DB Stamp 1 Date          20250520
DB Stamp 1 Time          00132933
DB Stamp 2 Date
DB Stamp 2 Time
Page Count
Devices Used              0
Database ID              sap-lnx25:M02_20241203_104036
Max Used Data Page       3187892
Converter Page Count
---
```

- Recover Log Backup as necessary
  - e.g. recover\_start M02\_LOG LOG 147

```

[dbmcli on M02>recover_start M02_LOG LOG 147
OK
Returncode           0
Date                 20250521
Time                 00112001
Server                sap-lnx25
Database              M02
Kernel Version       Kernel    7.9.10   Build 004-123-265-969
Pages Transferred    24
Pages Left           0
Volumes              1
Medianame            M02_LOG
Location              /sapdb/M02/backup/log/M02_LOG.147
Errortext
Label                 LOG_000000147
Is Consistent
First LOG Page       514072
Last LOG Page        514075
DB Stamp 1 Date      20250520
DB Stamp 1 Time      00180238
DB Stamp 2 Date      20250520
DB Stamp 2 Time      00180539
Page Count           4
Devices Used         1
Database ID          sap-lnx25:M02_20241203_104036
Max Used Data Page
Converter Page Count

```

- Optional Information – autorecover to a specific time stamp (without need to specify dedicated data / log backup
  - e.g. autorecover until 20250520 200000

```

---
[dbmcli on M02>autorecover until 20250520 200000
OK
Returncode           0
Date                 20250521
Time                 00131559
Server                sap-lnx25
Database              M02
Kernel Version       Kernel    7.9.10   Build 004-123-265-969
Pages Transferred    10096
Pages Left           0
Volumes              1
Medianame            M02_LOG
Location              /sapdb/M02/backup/log/M02_LOG.102
Errortext
Label                 LOG_000000102
Is Consistent
First LOG Page       256227
Last LOG Page        341559
DB Stamp 1 Date      20241203
DB Stamp 1 Time      00190348
DB Stamp 2 Date      20241226
DB Stamp 2 Time      00193615
Page Count           85333
Devices Used         1
Database ID          sap-lnx25:M02_20241203_104036
Max Used Data Page
Converter Page Count

```

- End Recovery and stop Database

- db\_offline



Further information about Recovery is available in the [MaxDB Documentation](#)

- start SAP System
  - User root: /etc/init.d/sapinit start
  - User m02adm: startsap

## Additional information and version history

### Recorded Demos

Following recoded Demos are available to support the documentation.

[Installation MaxDB Plugin, Configuration MaxDB Plugin, Backup of MaxDB database](#)

[Restore and Recovery of MaxDB database](#)

### External Documentation

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

- [SAP Installation Azure on ANF](#)
- [SnapCenter Prerequisites for Plugins](#)
- [SnapCenter Install Plugins](#)
- [MaxDB Recovery Documentation](#)
- SAP Notes (login required)
  - [1928060 - Data backup and recovery with file system backup](#)
  - [2282054 - Background DBM server](#)
  - [616814 - Suspend log writer for split mirror or snapshot](#)
- [HowTo - SAP MaxDB Backup with Database Manager CLI](#)
- [HowTo - SAP MaxDB Recovery with Database Manager CLI](#)
- [NetApp Product Documentation](#)
- [NetApp SAP Solutions – Informations about Use-Cases, Best-Practices and Benefits](#)

### Version history

Version	Date	Document version history
Version 1.0	May 2025	Initial version – backup / recovery MaxDB database

# Lifecycle Management

## NetApp SAP Landscape Management Integration using Ansible

### TR-4953: NetApp SAP Landscape Management Integration using Ansible

SAP Landscape Management (LaMa) enables SAP system administrators to automate SAP system operations, including end-to-end SAP system clone, copy, and refresh operations.

Authors: Michael Schlosser, Nils Bauer, NetApp

NetApp offers a rich set of Ansible modules that allows SAP LaMa to access technologies such as NetApp Snapshot and FlexClone through SAP LaMa Automation Studio. These technologies help to simplify and accelerate SAP system clone, copy, and refresh operations.

The integration can be used by customers who run NetApp storage solutions on-premises or by customers using NetApp storage services at public cloud providers such as Amazon Web Services, Microsoft Azure, or Google Cloud Platform.

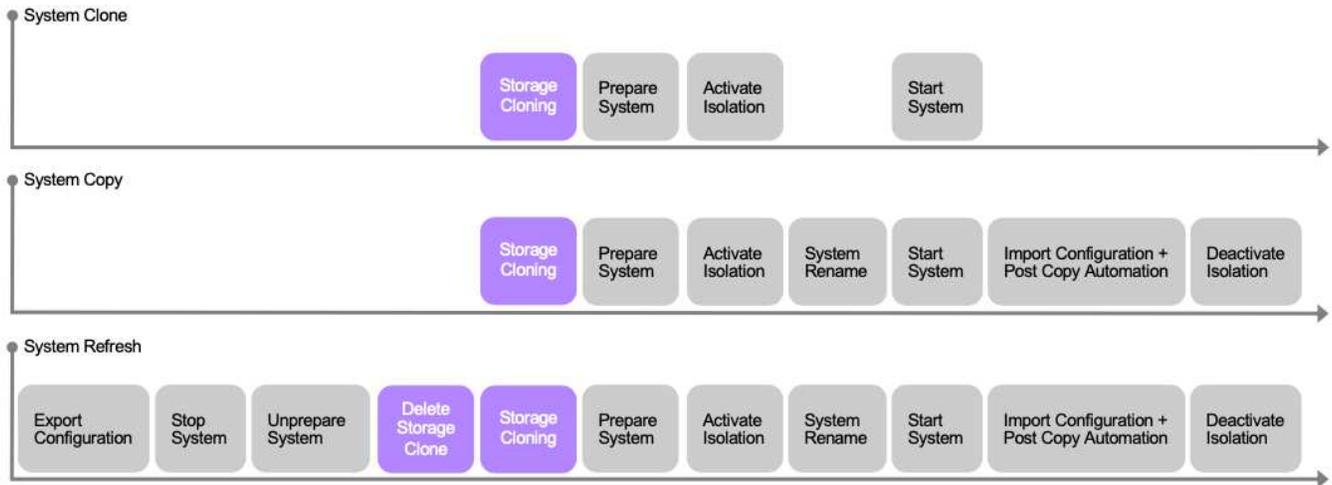
This document describes the configuration of SAP LaMa with NetApp storage features for SAP system copy, clone, and refresh operations using Ansible automation.

### SAP system clone, copy, and refresh scenarios

The term SAP system copy is often used as a synonym for three different processes: SAP system clone, SAP system copy, or SAP system refresh. It is important to distinguish between the different operations because the workflows and use cases differ for each one.

- **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.
- **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 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.

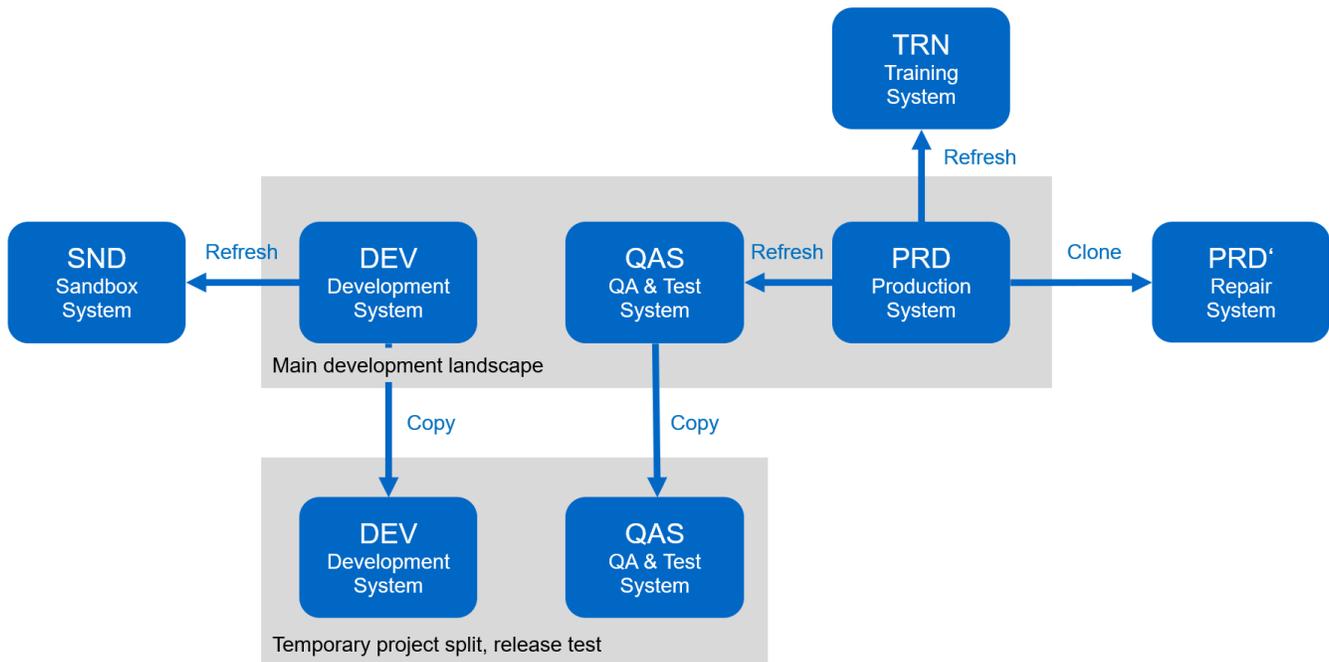
The following figure illustrates the main steps that must be performed during a system clone, system copy, or system refresh operation. The purple boxes indicate steps where NetApp storage features can be integrated. All three operations can be fully automated by using SAP LaMa.



### Use cases for system refresh, copy, and cloning

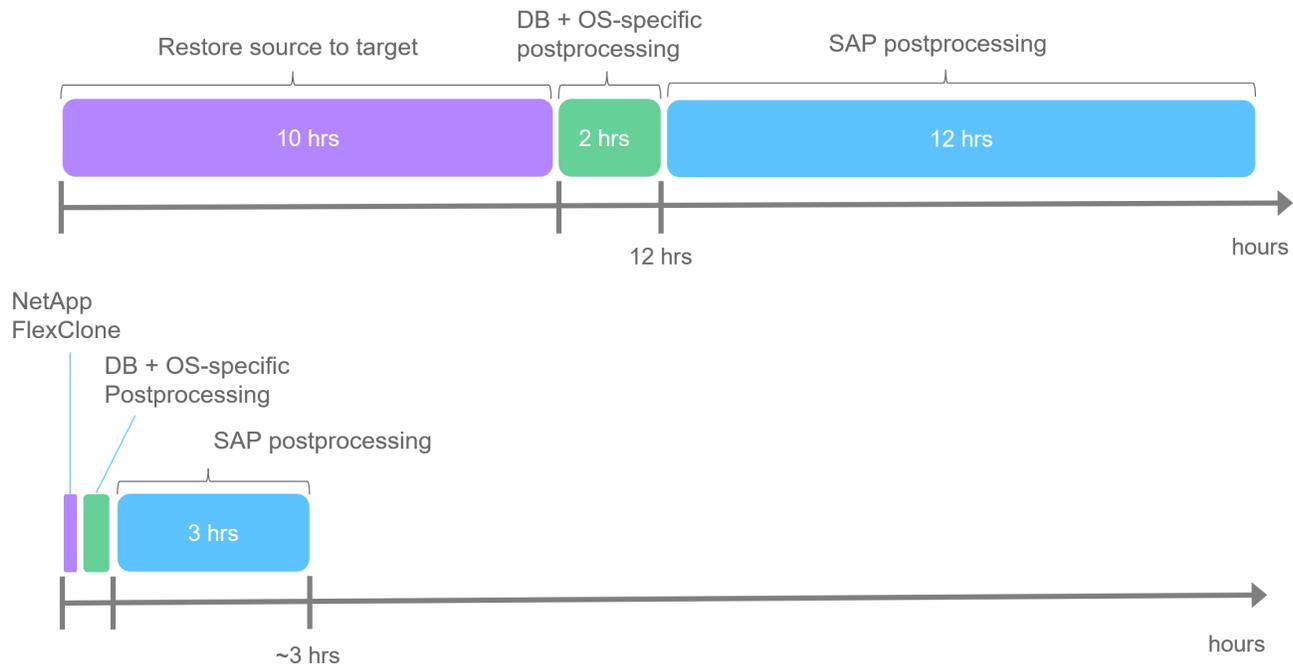
There are multiple scenarios in which data from a source system must be made available to a target system for testing or training purposes. These test and training systems must be updated with data from the source system on a regular basis to make sure that testing and training is performed with the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



SAP LaMa and NetApp 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,

SAP LaMa uses NetApp Snapshot copy and NetApp FlexClone technology so that required tasks up to a started HANA database can be performed in minutes instead of hours as shown in the following figure. 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. Further reduction of the runtime is accomplished by automating tasks on the operating system and database layer as well as on the SAP post processing side.



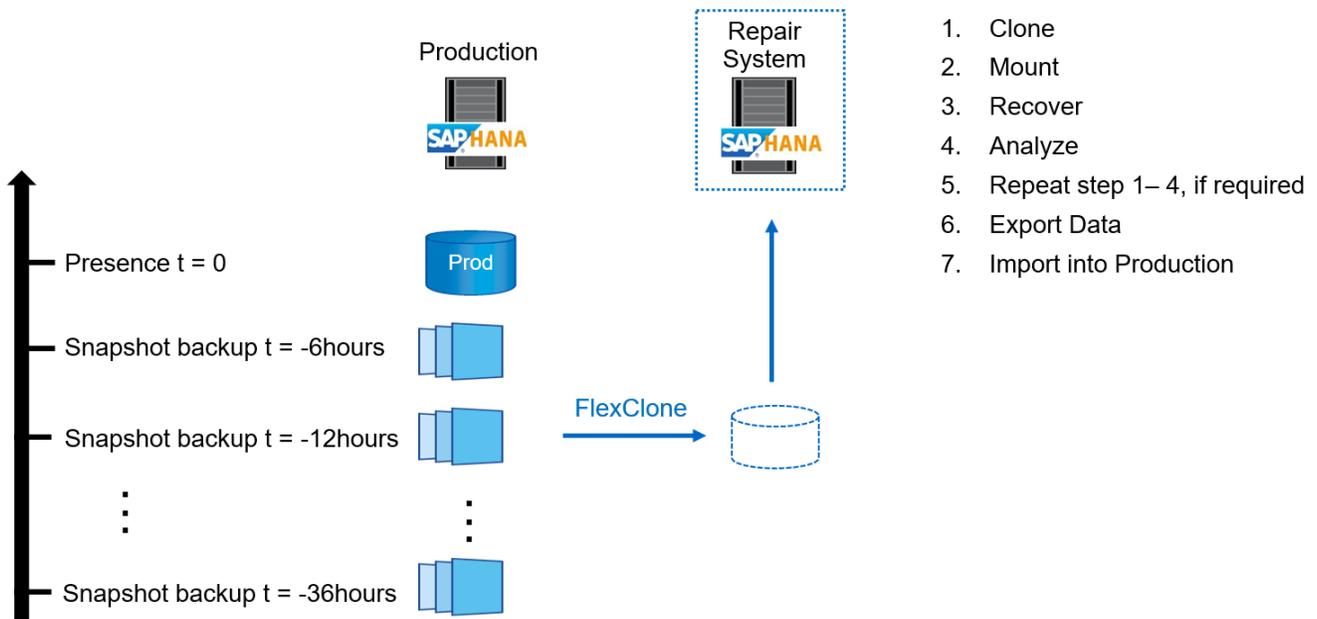
### Address logical corruption

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 acceptable 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 into 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 speed are crucial. With NetApp storage-based Snapshot backups, multiple consistent database images are available to create a clone of the production system by using NetApp FlexClone technology. FlexClone volumes can be created in a matter of seconds rather than multiple hours if a redirected restore from a file-based backup is used to set up the repair system.

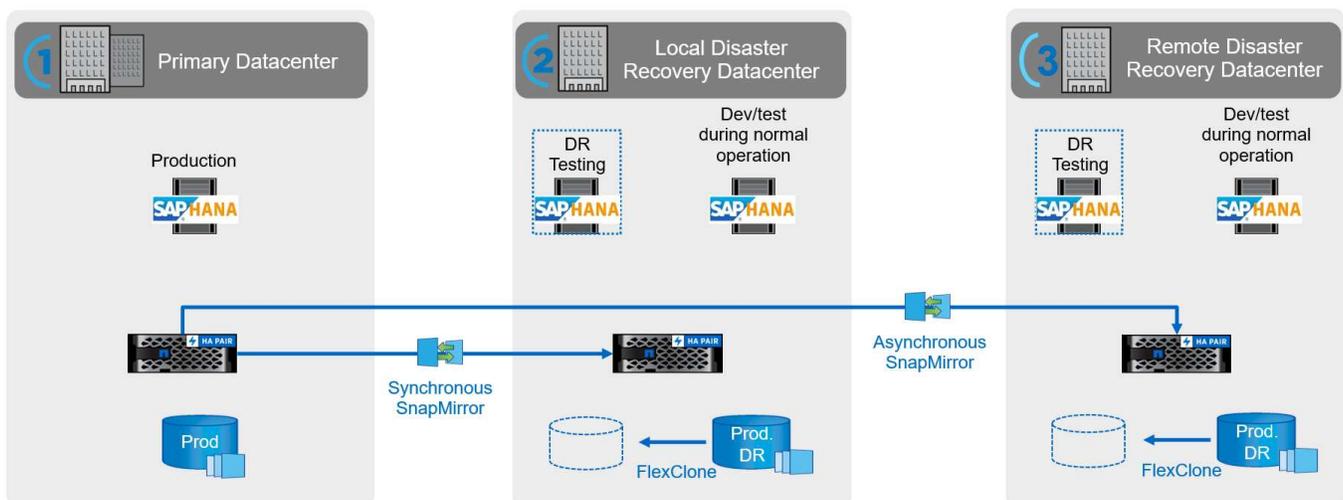


## Disaster recovery testing

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

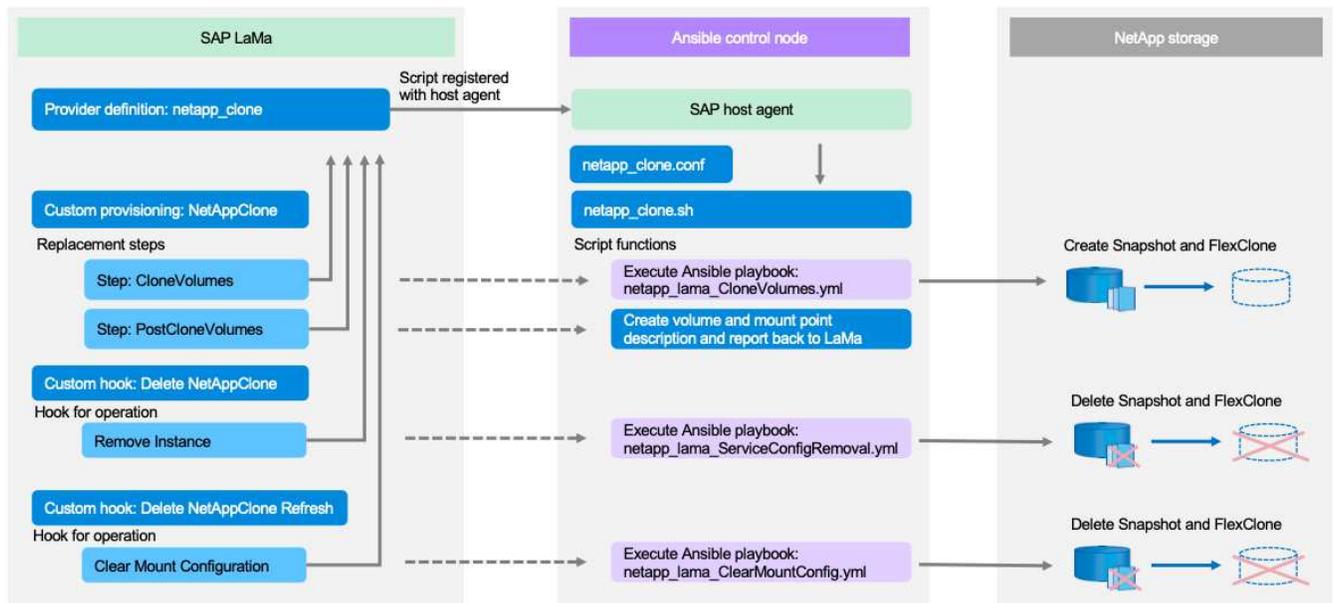
Storage replication with SnapMirror makes it possible to execute disaster recovery testing without putting RTO and RPO at risk. Disaster recovery testing can be performed without interrupting data replication. Disaster recovery testing for both asynchronous and synchronous SnapMirror uses Snapshot backups and FlexClone volumes at the disaster recovery target.

SAP LaMa can be used to orchestrate the entire testing procedure, and it also takes care of network fencing, target host maintenance, and so on.



## NetApp SAP LaMa integration using Ansible

The integration approach uses SAP LaMa custom provisioning and operation hooks combined with Ansible playbooks for NetApp storage management. The following figure shows a high-level overview of the configuration on the LaMa side as well as the corresponding components of the example implementation.



A central host acting as an Ansible control node is used to execute the requests from SAP LaMa and to trigger the NetApp storage operations using Ansible playbooks. The SAP host agent components must be installed on this host so that the host can be used as a communication gateway to SAP LaMa.

Within LaMa Automation Studio, a provider is defined that is registered at the Ansible host's SAP host agent. A host agent configuration file points to a shell script that is called by SAP LaMa with a set of command line parameters, depending on the requested operation.

Within LaMa Automation Studio, custom provisioning and a custom hook is defined to execute storage cloning operations during provisioning and also during clean-up operations when the system is deprovisioned. The shell script on the Ansible control node then executes the corresponding Ansible playbooks, which trigger the Snapshot and FlexClone operations as well as the deletion of the clones with the deprovisioning workflow.

More information on NetApp Ansible modules and the LaMa provider definitions can be found at:

- [NetApp Ansible modules](#)
- [SAP LaMa documentation – provider definitions](#)

### Example implementation

Due to the large number of options available for system and storage setups, the example implementation should be used as a template your individual system setup and configuration requirements.



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

## Validated configurations and limitations

The following principles were applied to the example implementation and might need to be adapted to meet customer needs:

- Managed SAP systems used NFS to access NetApp storage volumes and were set up based on the adaptive design principle.
- You can use all ONTAP releases supported by NetApp Ansible modules (ZAPI and REST API).
- Credentials for a single NetApp cluster and SVM were hard coded as variables in the provider script.
- Storage cloning was performed on the same storage system that was used by the source SAP system.
- Storage volumes for the target SAP system had the same names as the source with an appendix.
- No cloning at secondary storage (SV/SM) was implemented.
- FlexClone split was not implemented.
- Instance numbers were identical for the source and target SAP systems.

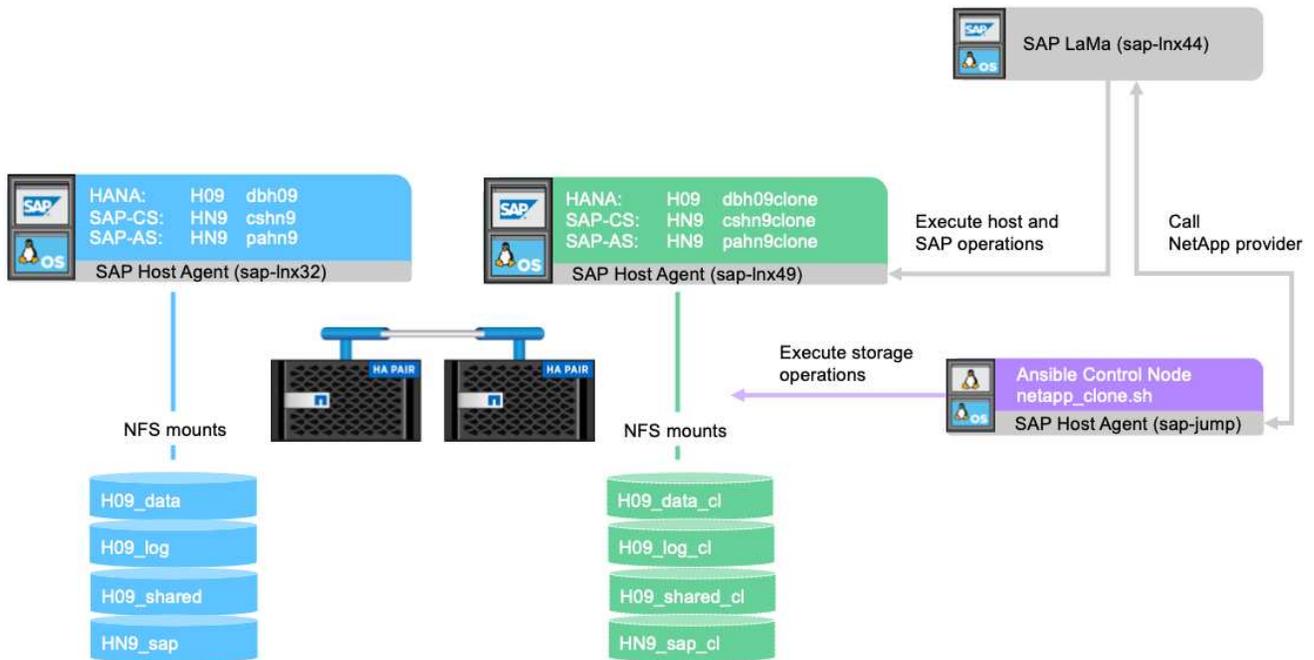
## Lab setup

The following figure shows the lab setup we used. The source SAP system HN9 used for the system clone operation consisted of the database H09, the SAP CS, and the SAP AS services running on the same host (sap-lnx32) with installed [adaptive design](#) enabled. An Ansible control node was prepared according to the [Ansible Playbooks for NetApp ONTAP](#) documentation.

The SAP host agent was installed on this host as well. The NetApp provider script as well as the Ansible playbooks were configured on the Ansible control node as described in the “[Appendix: Provider Script Configuration](#).”

The host `sap-lnx49` was used as the target for the SAP LaMa cloning operations, and the isolation-ready feature was configured there.

Different SAP systems (HNA as source and HN2 as target) were used for system copy and refresh operations, because Post Copy Automation (PCA) was enabled there.



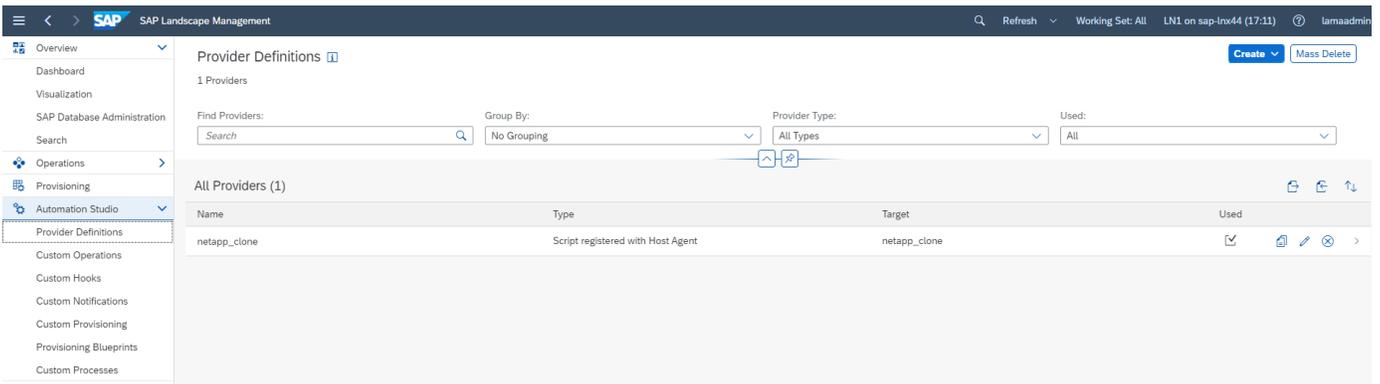
The following software releases were used in the lab setup:

- SAP LaMa Enterprise Edition 3.00 SP23\_2
- SAP HANA 2.00.052.00.1599235305
- SAP 7.77 Patch 27 (S/4 HANA 1909)
- SAP Host Agent 7.22 Patch 56
- SAPACEXT 7.22 Patch 69
- Linux SLES 15 SP2
- Ansible 2.13.7
- NetApp ONTAP 9.8P8

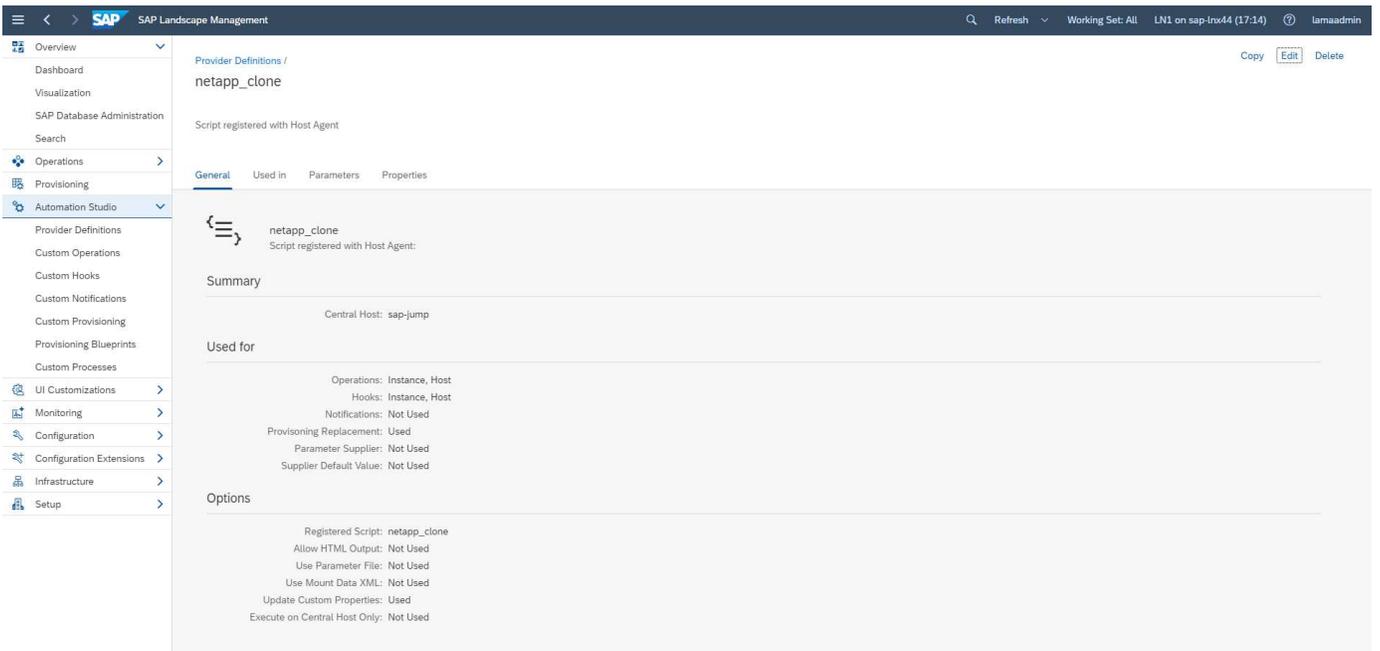
## SAP LaMa configuration

### SAP LaMa provider definition

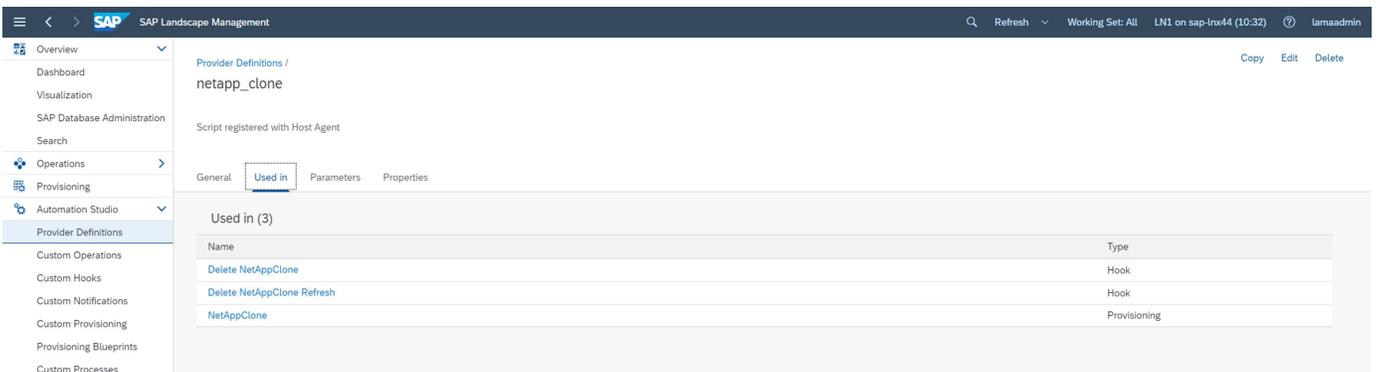
The provider definition is performed within Automation Studio of SAP LaMa as shown in the following screenshot. The example implementation uses a single provider definition that is used for different custom provisioning steps and operation hooks as explained before.



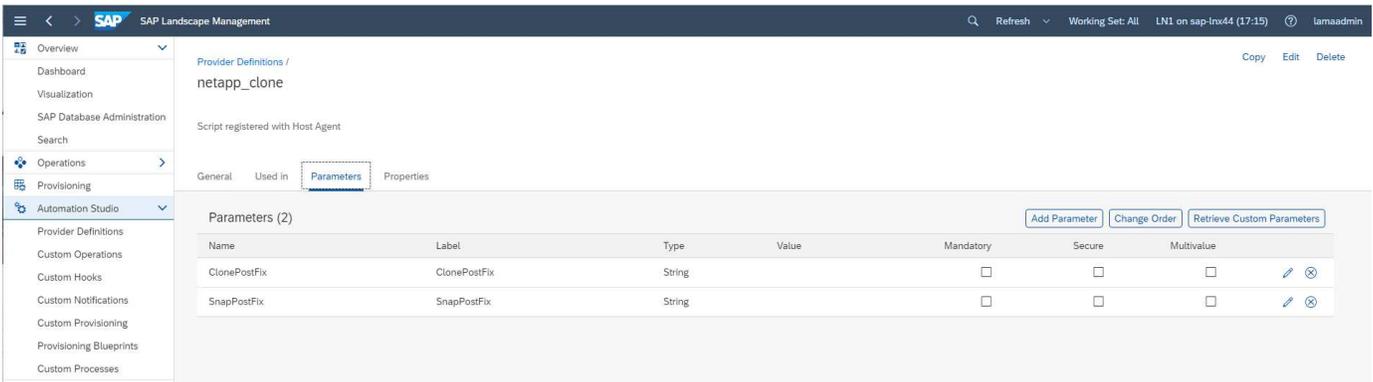
The provider `netapp_clone` is defined as the script `netapp_clone.sh` registered at the SAP host agent. The SAP host agent runs on the central host `sap-jump`, which also acts as the Ansible control node.



The **Used in** tab shows which custom operations the provider is used for. The configuration for the custom provisioning **NetAppClone** and the custom hooks **Delete NetAppClone** and **Delete NetAppClone Refresh** are shown in the next chapters.

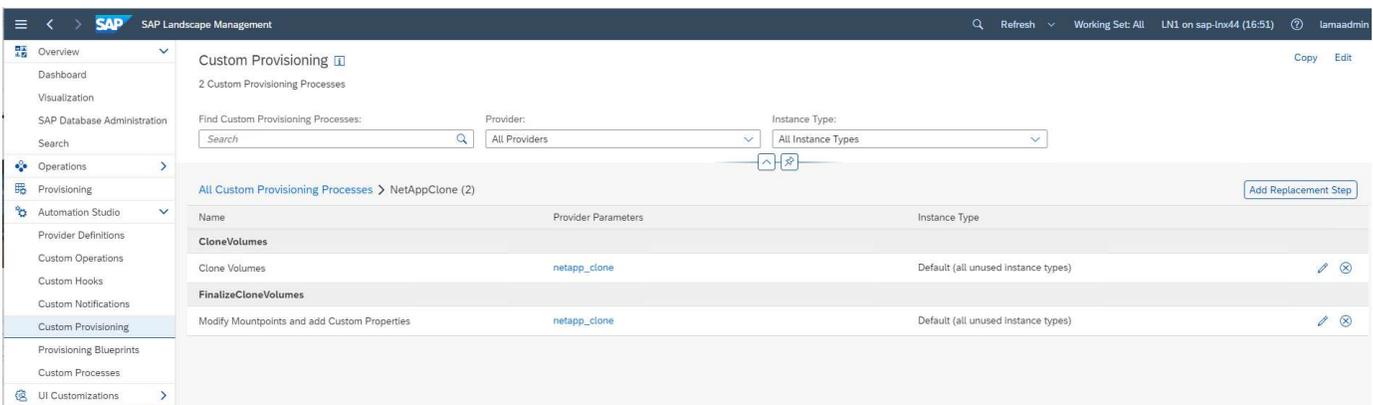


The parameters **ClonePostFix** and **SnapPostFix** are requested during the execution of the provisioning workflow and are used for the Snapshot and FlexClone volume names.



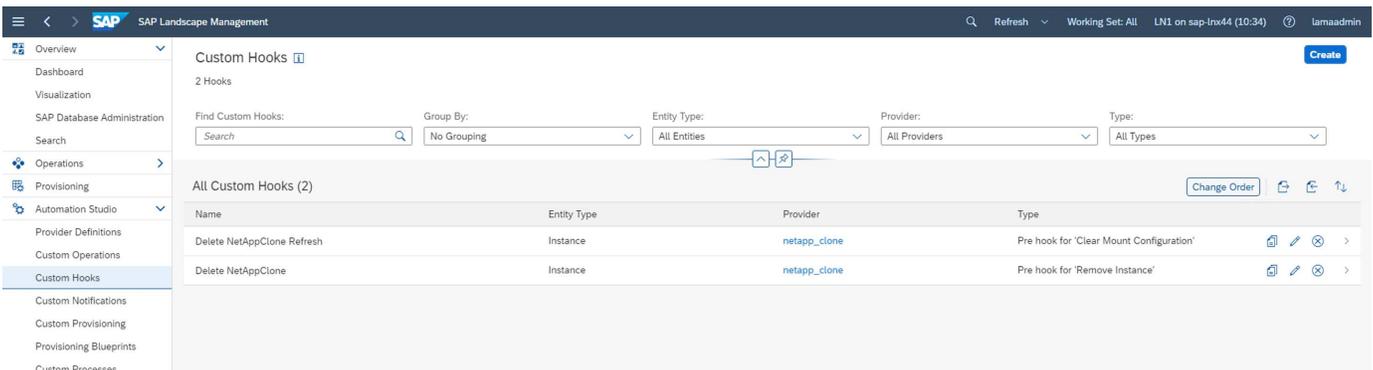
## SAP LaMa custom provisioning

In the SAP LaMa custom provisioning configuration, the customer provider described before is used to replace the provisioning workflow steps **Clone Volumes** and **PostCloneVolumes**.



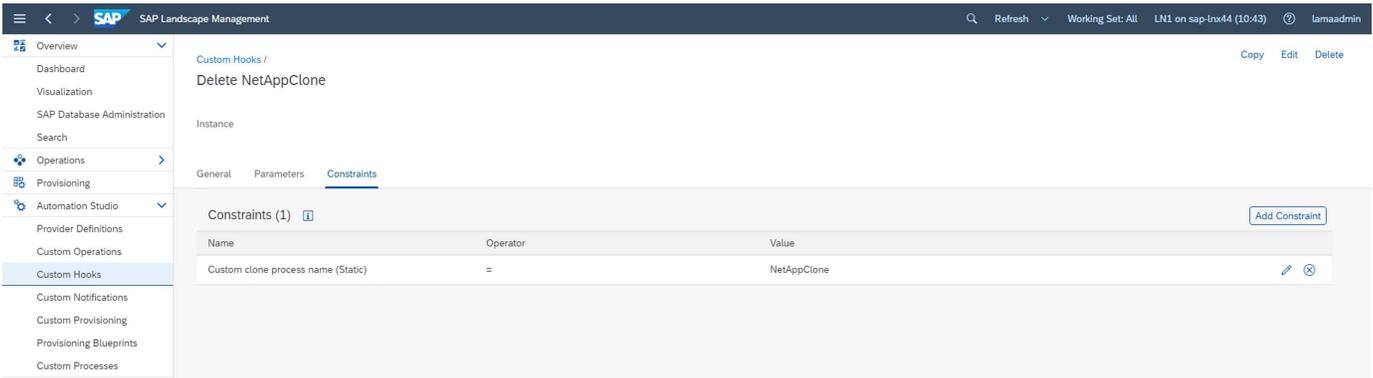
## SAP LaMa custom hook

If a system is deleted with the system destroy workflow, the hook **Delete NetAppClone** is used to call the provider definition `netapp_clone`. The **Delete NetApp Clone Refresh** hook is used during the system refresh workflow because the instance is preserved during the execution.



It is important to configure **Use Mount Data XML** for the custom hook, so that SAP LaMa provides the information of the mount point configuration to the provider.

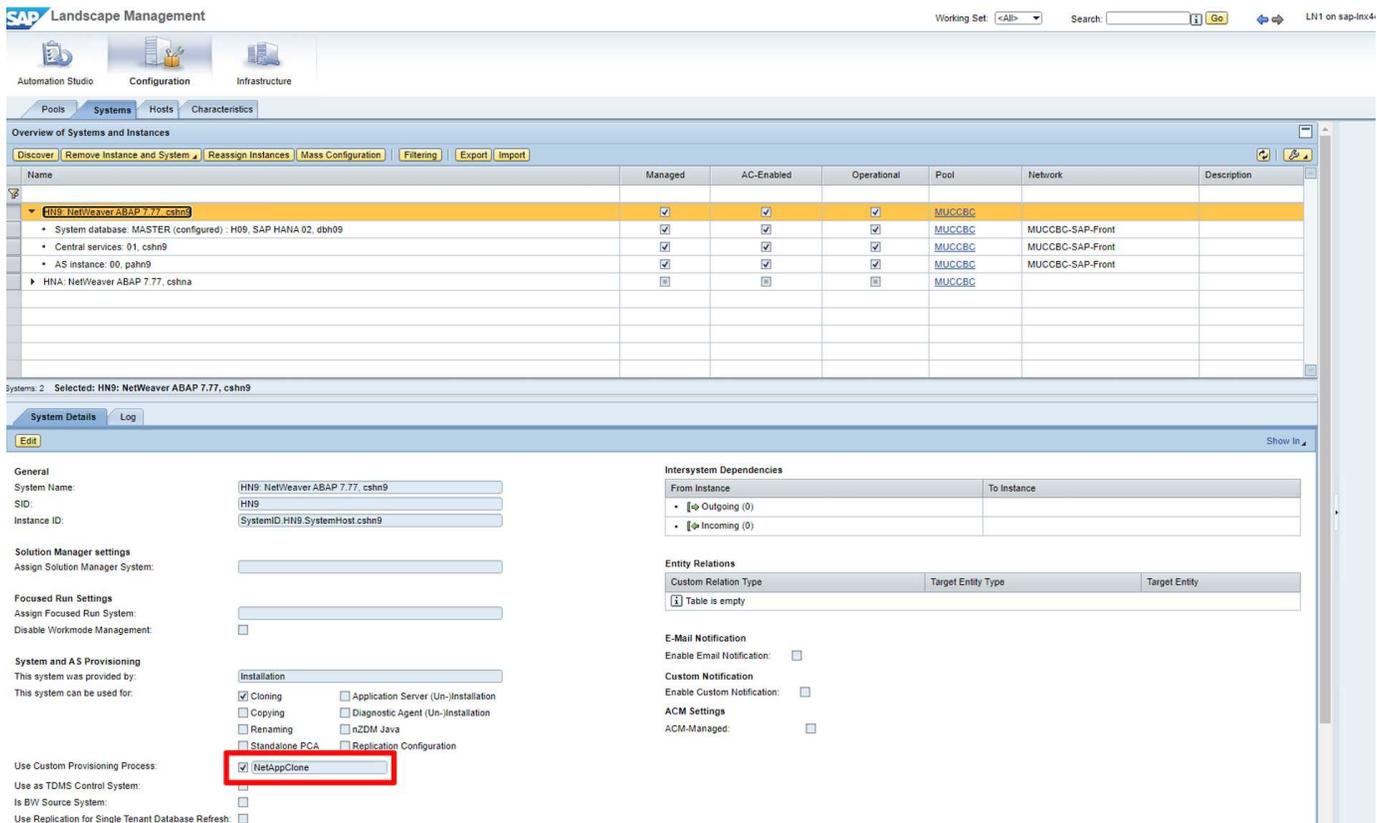
To ensure that the custom hook is only used and executed when the system was created with a custom provisioning workflow, the following constraint is added to it.



More information about the use of custom hooks can be found in the [SAP LaMa Documentation](#).

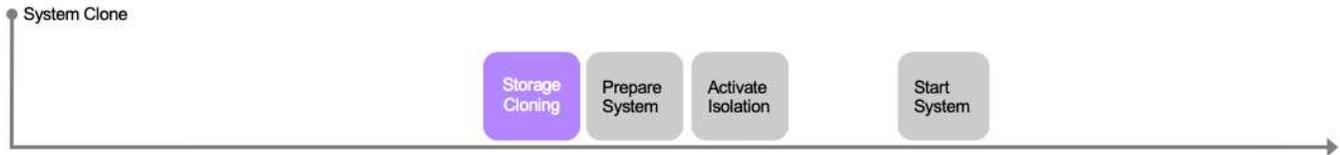
### Enable custom provisioning workflow for SAP source system

To enable the custom Hooks provisioning workflow for the source system, it must be adapted in the configuration. The **Use Custom Provisioning Process** checkbox with the corresponding custom provisioning definition must be selected.

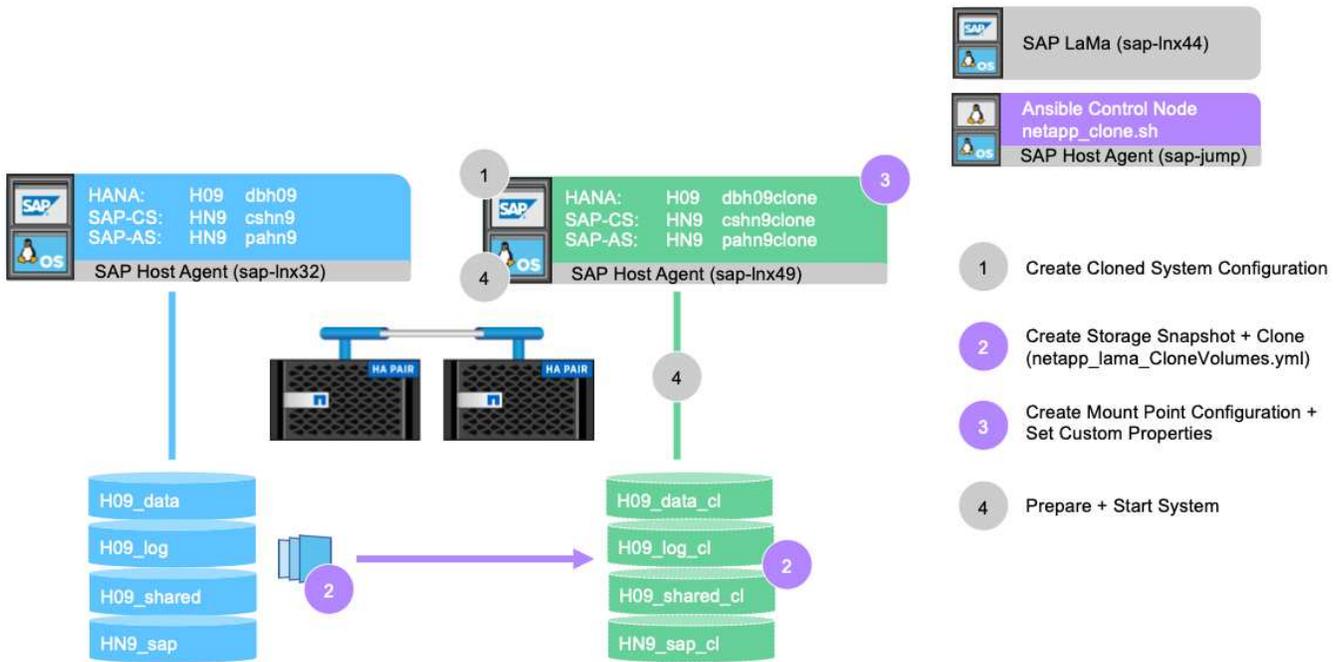


### SAP LaMa provisioning workflow - clone system

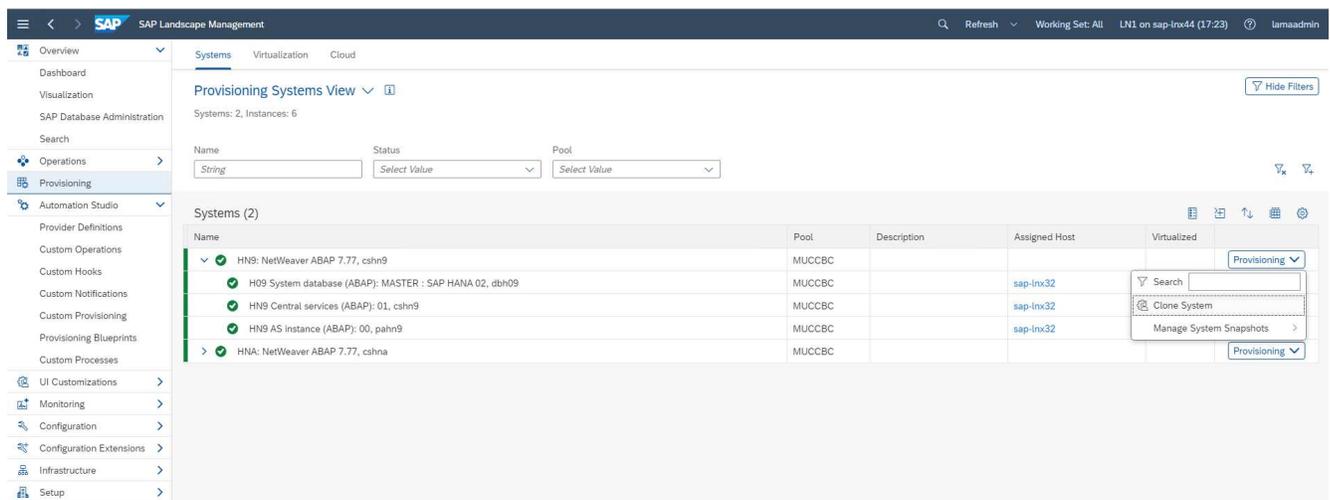
The following figure highlights the main steps executed with the system clone workflow.



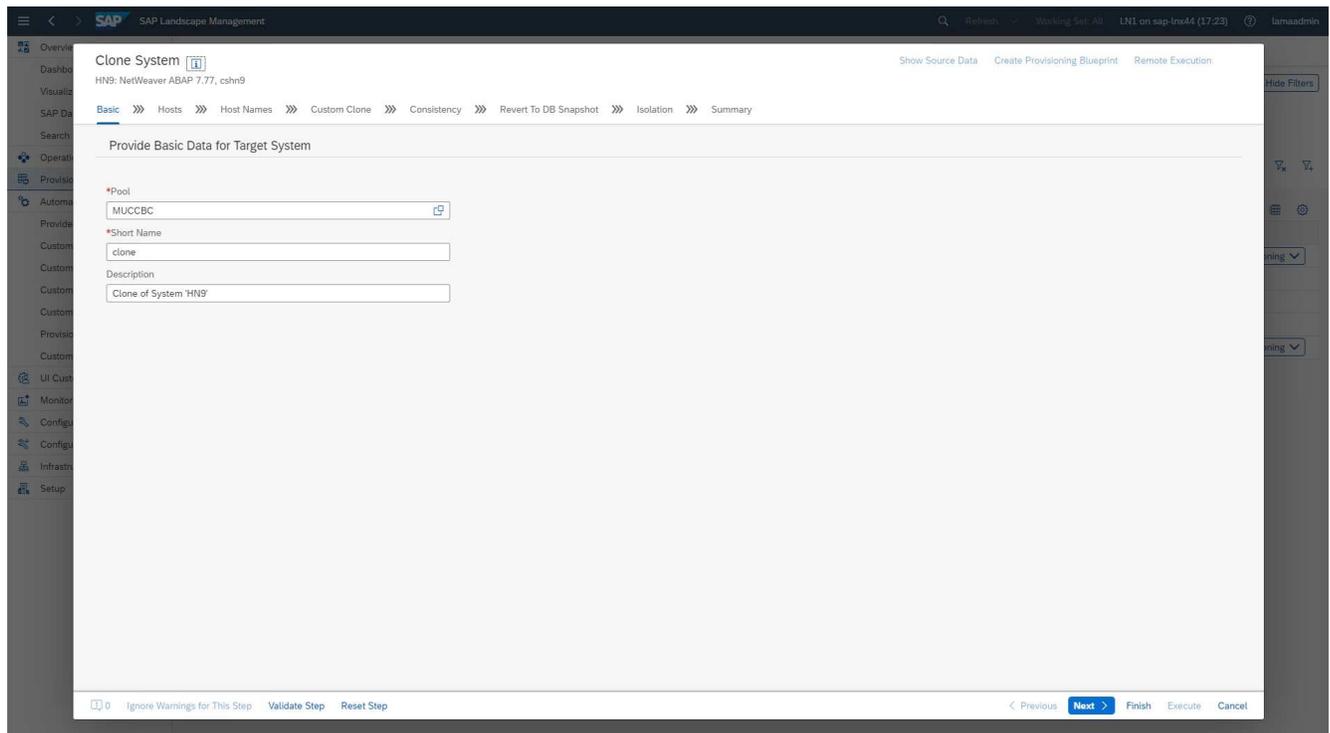
In this section, we go through the complete SAP LaMa system cloning workflow based on the source SAP system HN9 with HANA database H09. The following picture gives an overview of the steps executed during the workflow.



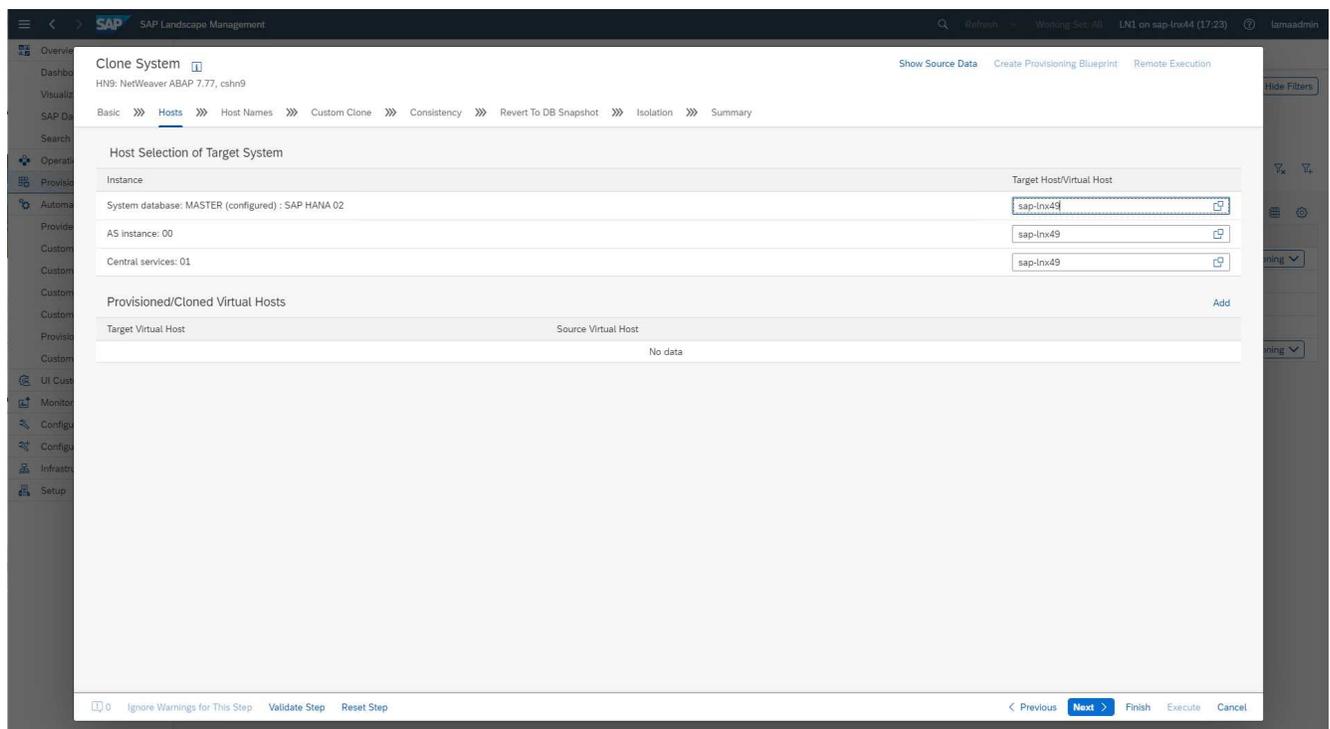
1. To start the cloning workflow, open **Provisioning** in the menu tree and select the source system (in our example HN9). Then start the **Clone System** wizard.



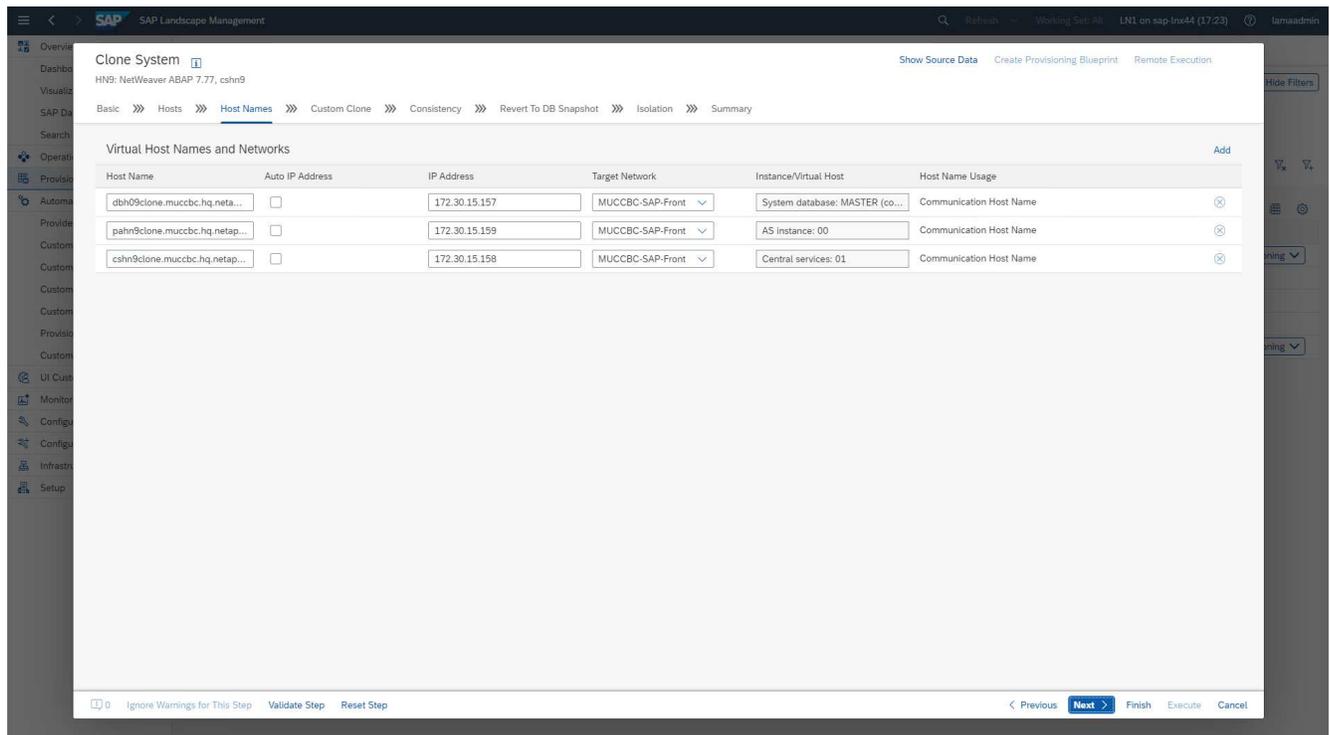
2. Enter the requested values. Screen 1 of the wizard asks for the pool name for the cloned system. This step specifies the instances (virtual or physical) on which the cloned system will be started. The default is to clone the system into the same pool as the target system.



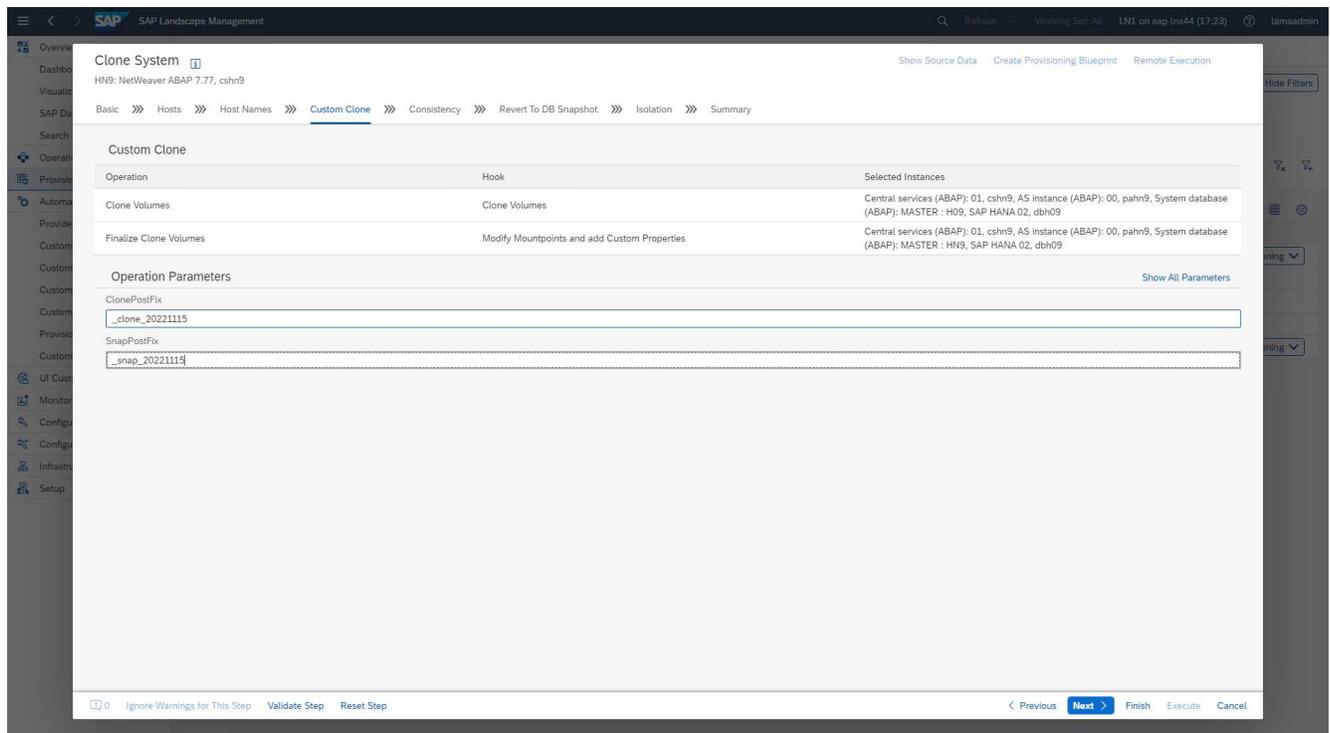
3. Screen 2 of the wizard asks for the target hosts that the new SAP instances are started on. The target hosts for this instance(s) can be selected out of the host pool specified in the previous screen. Each instance or service can be started on a different host. In our example, all three services run on the same host.



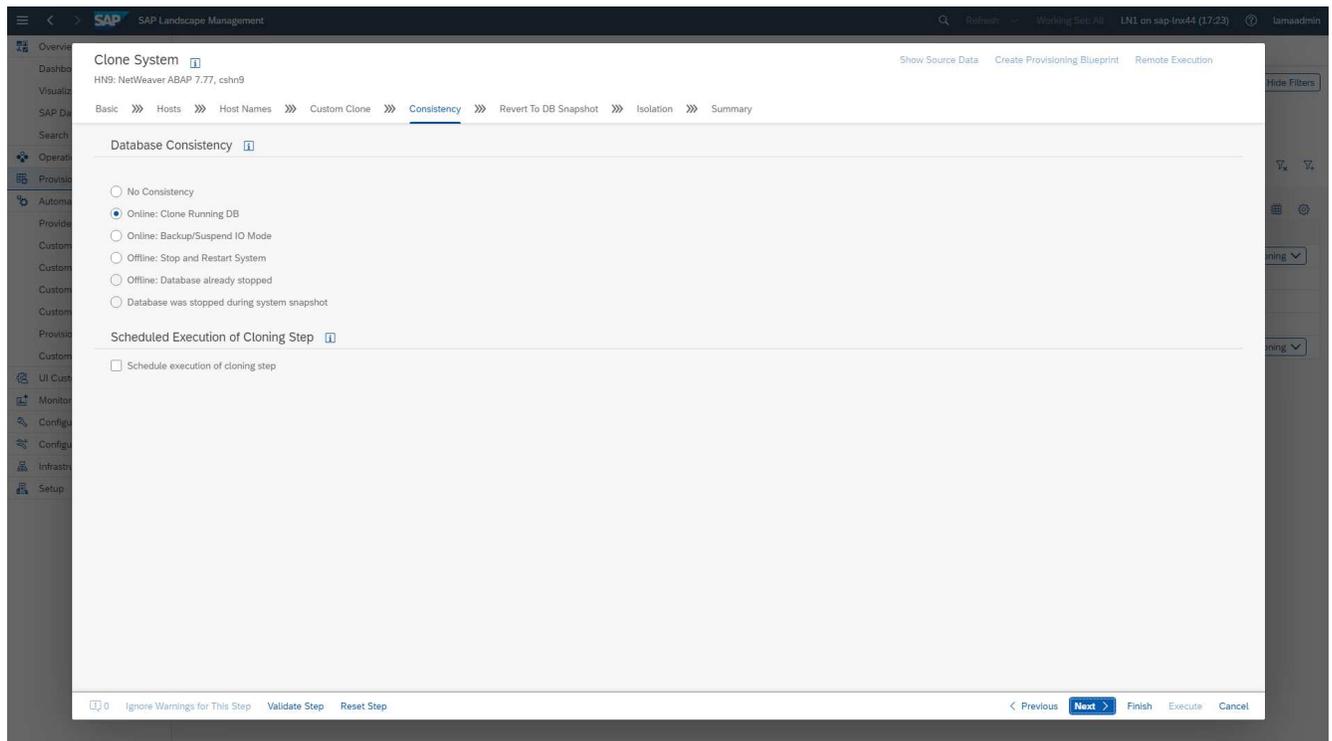
4. Provide the information requested in screen 3, which asks for virtual host names and networks. Typically, the host names are maintained in DNS, so the IP addresses are prepopulated accordingly.



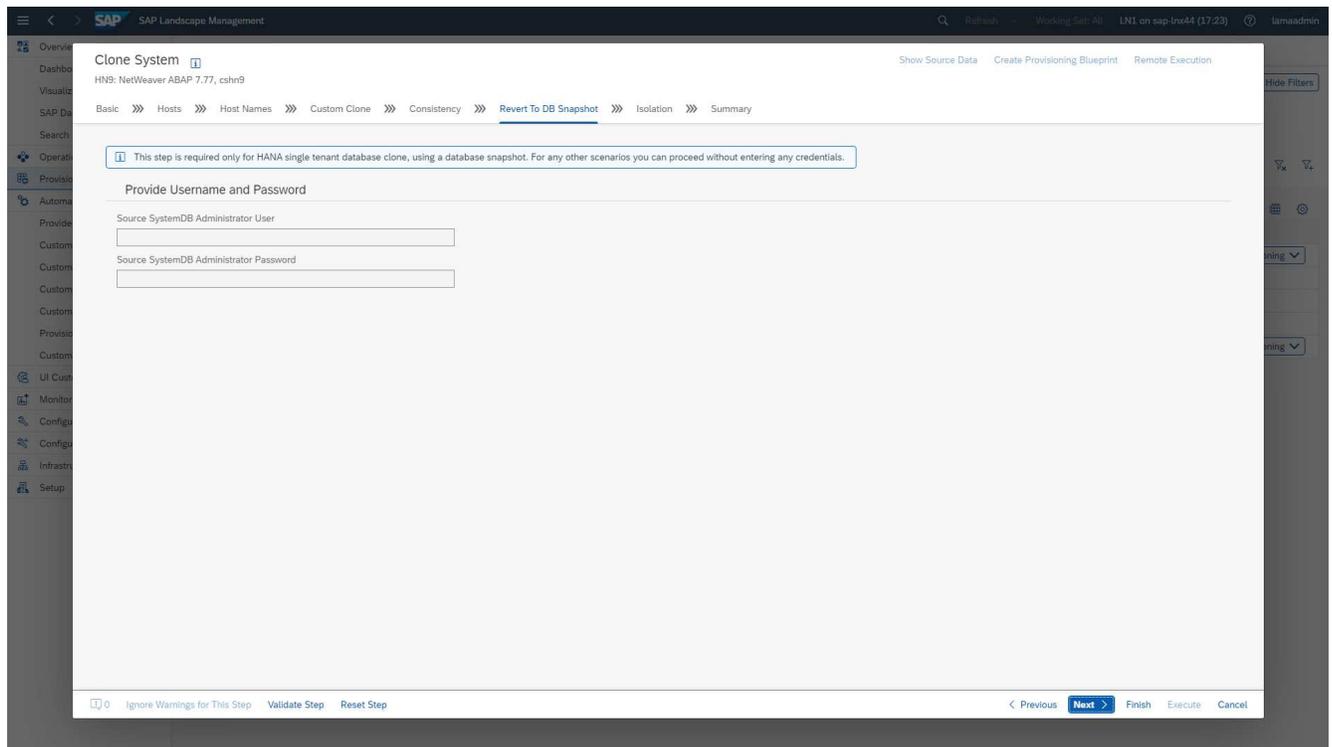
5. In screen 4, the custom clone operations are listed. A clone and a **SnapPostfix** name are provided, which are used during the storage clone operation for the FlexClone volume and Snapshot name, respectively. If you leave these fields empty, the default value configured in the variable section of the provider script `netapp_clone.sh` is used.



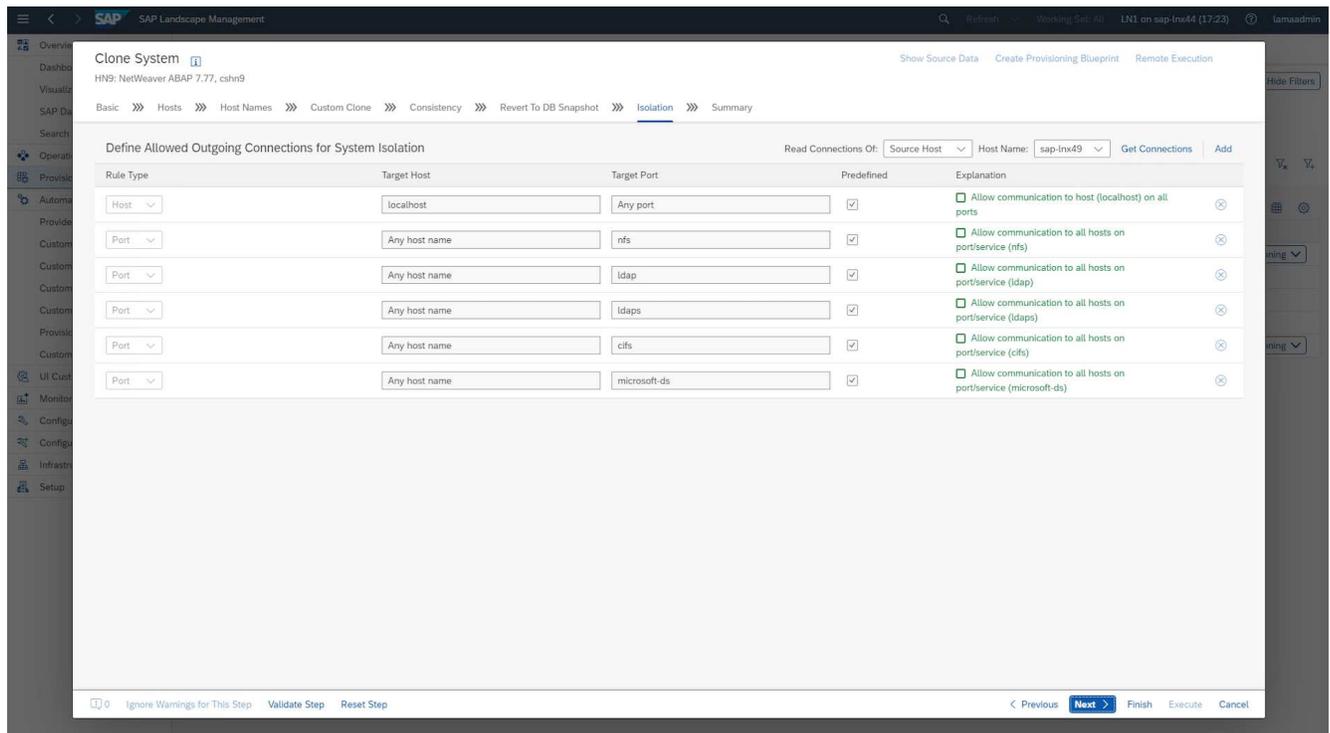
6. In screen 5, the database consistency option is selected. In our example, we selected **Online: Clone running DB**.



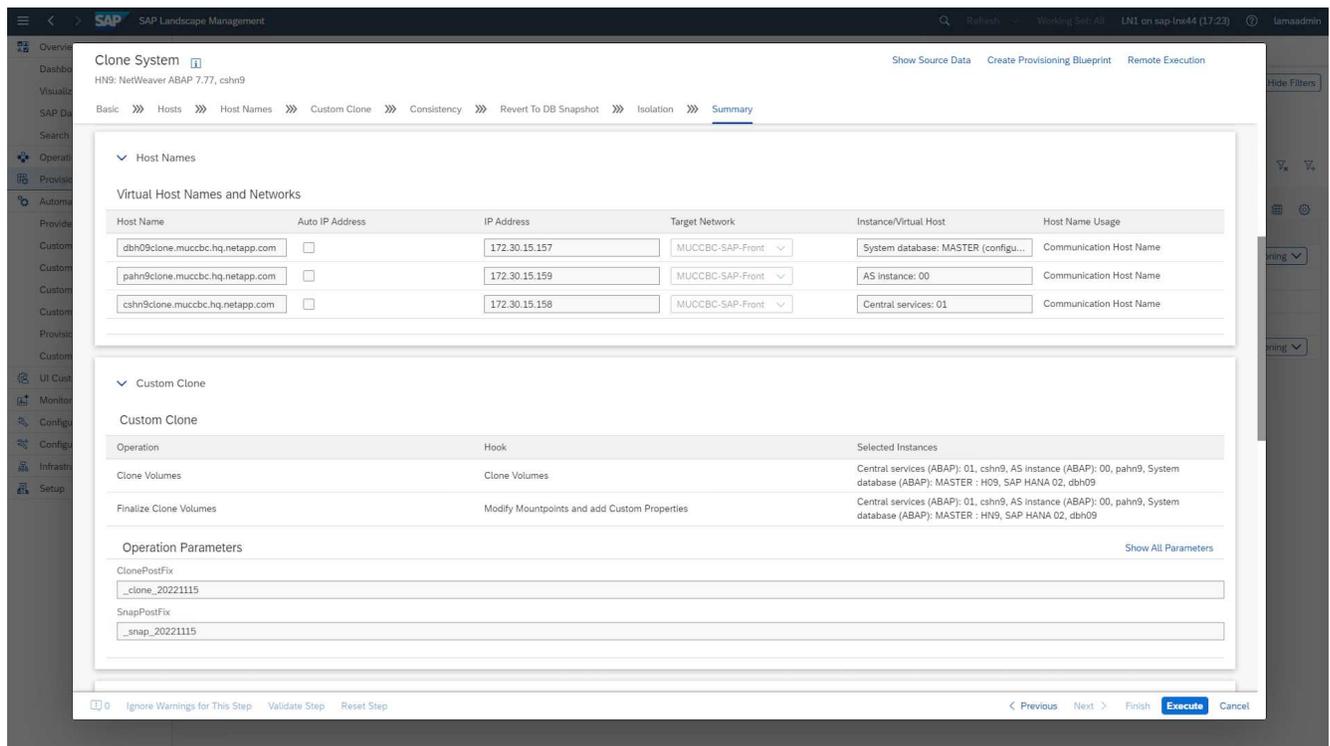
7. In screen 6, input is only required if you perform a tenant clone.



8. In screen 7, system isolation can be configured.

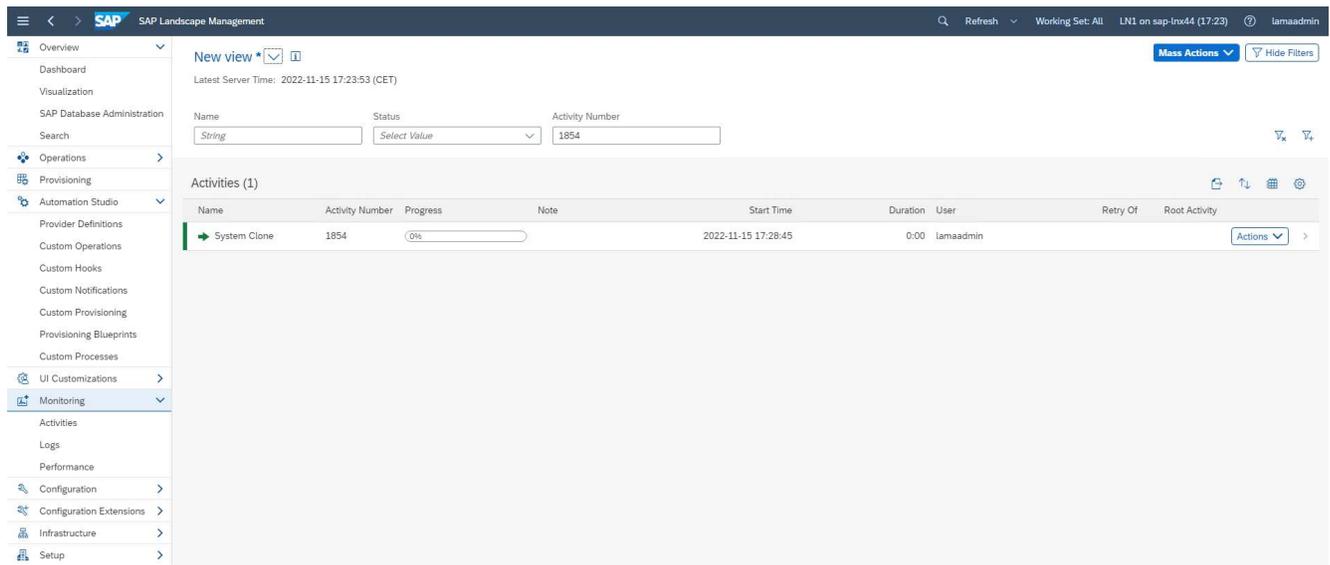


9. In screen 8, a summary page contains all the settings for final confirmation before the workflow is started. Click **Execute** to start the workflow.

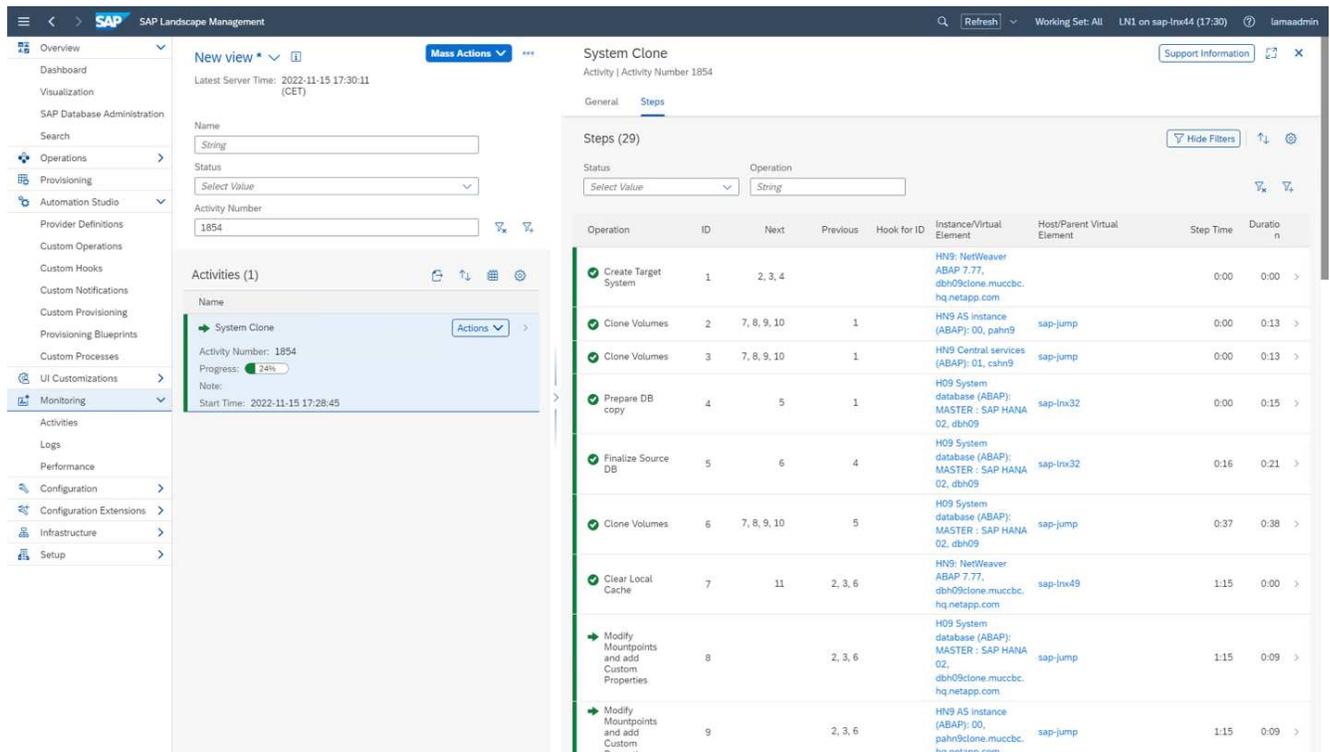


SAP LaMa now performs all the actions indicated in the configuration. These actions include creating the storage volume clones and exports, mounting them to the target host, adding the firewall rules for isolation, and starting the HANA database and SAP services.

10. You can monitor the progress of the clone workflow under the **Monitoring** menu.



Within the detailed log, the operations **Clone Volume** and **Modify Mountpoints and add Custom Properties** are executed at the Ansible node, the `sap-jump` host. These steps are executed for each service, the HANA database, the SAP central services, and the SAP AS service.



11. By selecting the **Clone Volumes** task the detailed log for that step is displayed and the execution of the Ansible Playbook is shown here. You can see, that the Ansible playbook `netapp_lama_CloneVolumes.yml` is executed for each HANA database volume, data, log, and shared.

The screenshot displays the SAP Landscape Management interface for a 'System Clone' activity. The 'Messages' pane on the right is highlighted with a red box, showing several debug messages related to cloning volumes and NetApp clones. The messages include:

- DEBUG | ID: 59 | Message Code: OSP-0200
- DEBUG | ID: 58 | Message Code: TMP-1001
- DEBUG | ID: 57 | Message Code: FWD-0003
- DEBUG | ID: 56 | Message Code: HALog
- DEBUG | ID: 55 | Message Code: LVM
- DEBUG | ID: 39 | Message Code: NetApp Clone for Custom Provis
- DEBUG | ID: 31 | Message Code: NetApp Clone for Custom Provis
- DEBUG | ID: 23 | Message Code: NetApp Clone for Custom Provis
- DEBUG | ID: 22 | Message Code: NetApp Clone for Custom Provis
- DEBUG | ID: 21 | Message Code: NetApp Clone for Custom Provis
- DEBUG | ID: 20 | Message Code: NetApp Clone for Custom Provis

12. In the details view of the step **Modify Mountpoints and add Custom Properties**, you can find information about the mount points and the custom properties handed over by the execution script.

The screenshot displays the SAP Landscape Management interface for the 'Modify Mountpoints and add Custom Properties' activity. The 'Messages' pane on the right is highlighted with a red box, showing a result message with custom properties. The message includes:

- RESULT | ID: 24 | Message Code: NetApp Clone for Custom Provis
- Got new property SnapPostFix: \_snap\_20221115
- RESULT | ID: 23 | Message Code: NetApp Clone for Custom Provis
- Got new property ClonePostFix: \_clone\_20221115

Below the result message, there is a detailed log of the script execution, including the following parameters:

```

netapp_clone.sh --HookOperationName=FinalizeCloneVolumes --SAPSYSTEMNAME=HN9 --SAPSYSTEM=01 --
MOUNT_XML_PATH=--PARAM_ClonePostFix=_clone_20221115 --PARAM_SnapPostFix=_snap_20221115 --
PROP_ClonePostFix=--PROP_SnapPostFix=--SAP_LVM_SRC_SID=HN9 --SAP_LVM_TARGET_SID=HN9
  
```

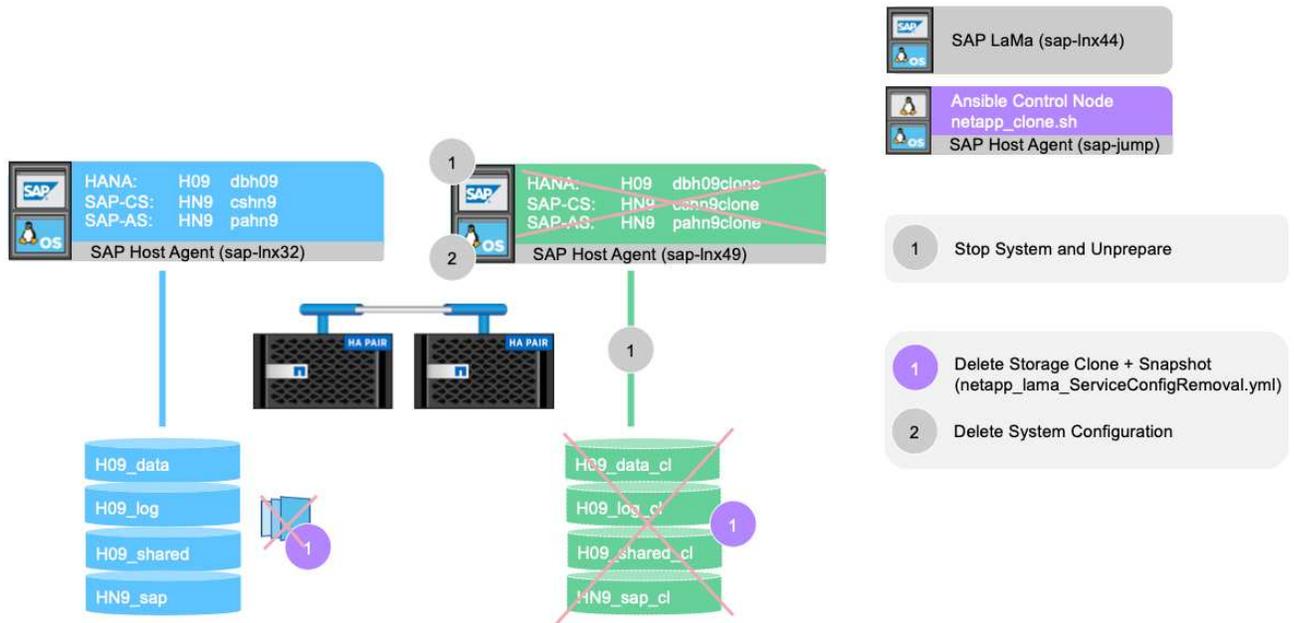
After the workflow has been completed, the cloned SAP system is prepared, started, and ready for use.

## SAP LaMa deprovisioning workflow - system destroy

The following figure highlights the main steps executed with the system destroy workflow.



1. To decommission a cloned system, it must be stopped and prepared in advance. Afterwards the system destroy workflow can be started.



2. In this example, we run the system destroy workflow for the system created before. We select the system in the **System View** screen and start the system destroy workflow under **Destroy Processes**.

3. All the mount points maintained during the provisioning phase are shown here and are deleted during the system destroy workflow process.

SAP Landscape Management

Destroy System **HN9: NetWeaver ABAP 7.77, dbh09clone.mucccbc.hq.netapp.com**

Delete Storage Volumes **>>>** Delete Host Names **>>>** Summary

### Storage Volumes

Delete	Volume	Storage Manager	Storage System	Storage Pool	Volume Group	Latest Monitoring Time
No data						

### Mount Data Without Corresponding Storage Volume

Instance	Storage Type	Export Path	Mount Point	Mount Options
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/hn9a...	/home/hn9adm	rw,noatime,vers=3,rsize=65536,wsiz...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/sapmnt	/sapmnt/HN9	rw,noatime,vers=3,rsize=65536,wsiz...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/HN9	/usr/sap/HN9	rw,noatime,vers=3,rsize=65536,wsiz...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/ccms	/usr/sap/ccms/HN9_00	rw,noatime,vers=3,rsize=65536,wsiz...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/saptr...	/usr/sap/trans	rw,noatime,vers=3,rsize=65536,wsiz...
System database: MASTER : H09, SAP HANA 02	NETFS	192.168.10.14:/H09_data_clone_20221115/data	/hana/data/H09	rw,noatime,vers=3,rsize=65536,wsiz...
System database: MASTER : H09, SAP HANA 02	NETFS	192.168.10.14:/H09_log_clone_20221115/log	/hana/log/H09	rw,noatime,vers=3,rsize=65536,wsiz...
System database: MASTER : H09, SAP HANA 02	NETFS	192.168.10.14:/H09_shared_clone_20221115/sh...	/hana/shared/H09	rw,noatime,vers=3,rsize=65536,wsiz...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/hn9a...	/home/hn9adm	rw,noatime,vers=3,rsize=65536,wsiz...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/sapmnt	/sapmnt/HN9	rw,noatime,vers=3,rsize=65536,wsiz...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/HN9	/usr/sap/HN9	rw,noatime,vers=3,rsize=65536,wsiz...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/ccms	/usr/sap/ccms/HN9_00	rw,noatime,vers=3,rsize=65536,wsiz...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/saptr...	/usr/sap/trans	rw,noatime,vers=3,rsize=65536,wsiz...

Monitoring Time:  [Monitoring Data](#)

Ignore Warnings for This Step [Validate Step](#) [Reset Step](#) [Previous](#) [Next](#) [Finish](#) [Execute](#) [Cancel](#)

No virtual hostnames are deleted because they are maintained through DNS and have been assigned automatically.

SAP Landscape Management

Destroy System **HN9: NetWeaver ABAP 7.77, dbh09clone.mucccbc.hq.netapp.com**

Delete Storage Volumes **>>>** Delete Host Names **>>>** Summary

### Host Names

Delete	DNS Server	Host Name	IP Address
No data			

Ignore Warnings for This Step [Validate Step](#) [Reset Step](#) [Previous](#) [Next](#) [Finish](#) [Execute](#) [Cancel](#)

4. The operation is started by clicking the execute button.

Destroy System Show Source Data Create Provisioning Blueprint Remote Execution

HN9: NetWeaver ABAP 7.77, dbh09clone.muccbc.hq.netapp.com

Delete Storage Volumes >>> Delete Host Names >>> **Summary**

SAP advises that it is the customer's responsibility to ensure that no data is lost when the selected volumes/virtual hosts are deleted by SAP Landscape Management.

**Delete Storage Volumes**

Storage Volumes

Delete	Volume	Storage Manager	Storage System	Storage Pool	Volume Group	Latest Monitoring Time
No data						

Mount Data Without Corresponding Storage Volume

Instance	Storage Type	Export Path	Mount Point	Mount Options
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/hn9...	/home/hn9adm	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/sap...	/sapmnt/HN9	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/HN9	/usr/sap/HN9	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/ccms	/usr/sap/ccms/HN9_00	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
AS instance: 00	NETFS	192.168.10.14:/HN9_sap_clone_20221115/sapt...	/usr/sap/trans	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
System database: MASTER : H09, SAP HANA 02	NETFS	192.168.10.14:/H09_data_clone_20221115/data	/hana/data/H09	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
System database: MASTER : H09, SAP HANA 02	NETFS	192.168.10.14:/H09_log_clone_20221115/log	/hana/log/H09	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
System database: MASTER : H09, SAP HANA 02	NETFS	192.168.10.14:/H09_shared_clone_20221115/s...	/hana/shared/H09	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/hn9...	/home/hn9adm	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/sap...	/sapmnt/HN9	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/HN9	/usr/sap/HN9	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/ccms	/usr/sap/ccms/HN9_00	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...
Central services: 01	NETFS	192.168.10.14:/HN9_sap_clone_20221115/sapt...	/usr/sap/trans	rw,noatime,vers=3,rsize=65536,wsiz=65536,n...

Monitoring Time:  Monitoring Data

0 Ignore Warnings for This Step Validate Step Reset Step

< Previous Next > Finish Execute Cancel

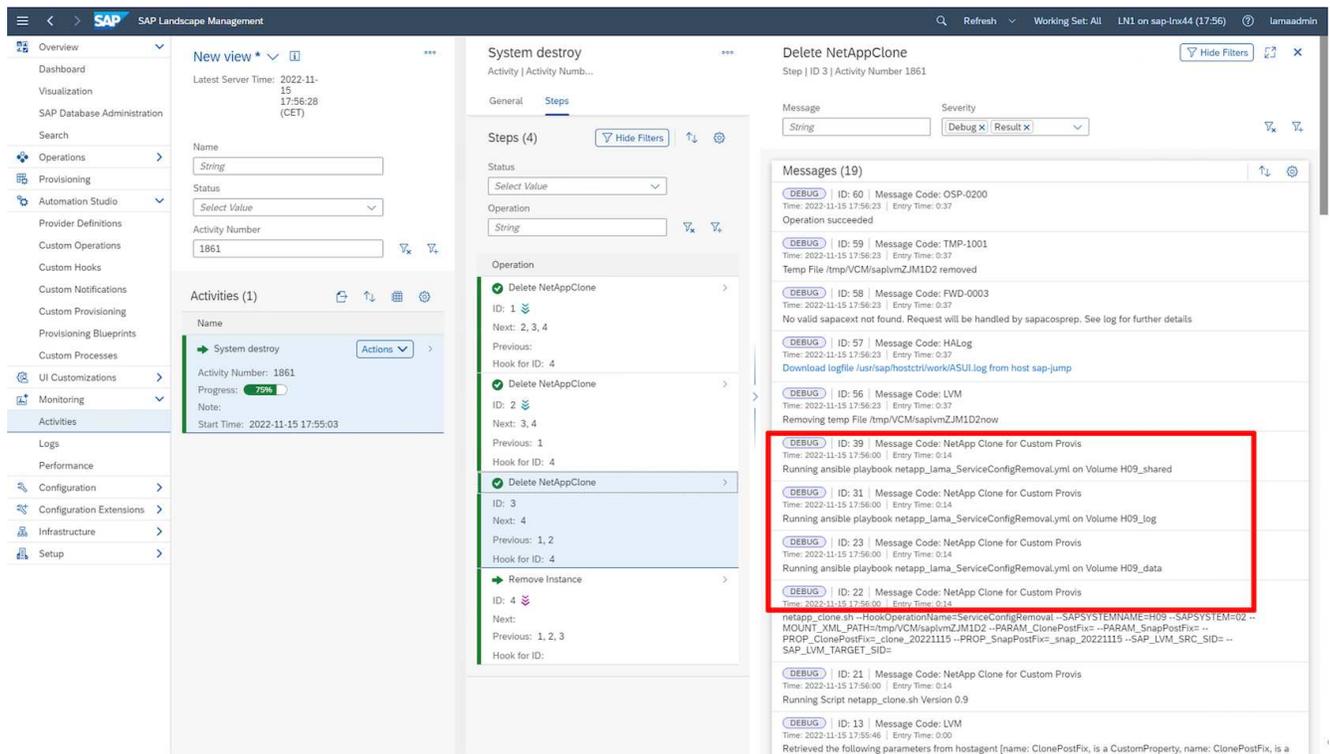
SAP LaMa now performs the deletion of the volume clones and deletes the configuration of the cloned system.

5. You can monitor the progress of the clone workflow under the **Monitoring** menu.

The screenshot shows the SAP Landscape Management interface. On the left is a navigation menu with 'Monitoring' selected. The main area displays a 'System destroy' activity (Activity Number 1861) with a progress bar at 0%. Below the activity, a table shows the following steps:

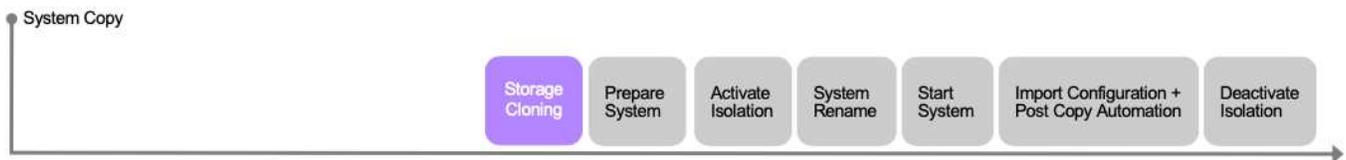
Operation	ID	Next	Previous	Hook for ID	Instance/Virtual Element	Host/Parent Virtual Element	Step Time	Duration
Delete NetAppClone	1	2, 3, 4		4	HN9 Central services (ABAP): 01, csnr9clone.muccbc.hq.netapp.com	sap-jump	0:00	0:11
Delete NetAppClone	2	3, 4	1	4	HN9 AS Instance (ABAP): 00, pah9clone.muccbc.hq.netapp.com	sap-jump		
Delete NetAppClone	3	4	1, 2	4	H09 System database (ABAP): MASTER : SAP HANA 02, dbh09clone.muccbc.hq.netapp.com	sap-jump		
Remove Instance	4		1, 2, 3		HN9: NetWeaver ABAP 7.77, dbh09clone.muccbc.hq.netapp.com			

6. By selecting the **Delete NetAppClone** task, the detailed log for that step is displayed. The execution of the Ansible Playbook is shown here. As you can see, the Ansible playbook `netapp_lama_ServiceConfigRemoval.yml` is executed for each HANA database volume, data, log, and shared.

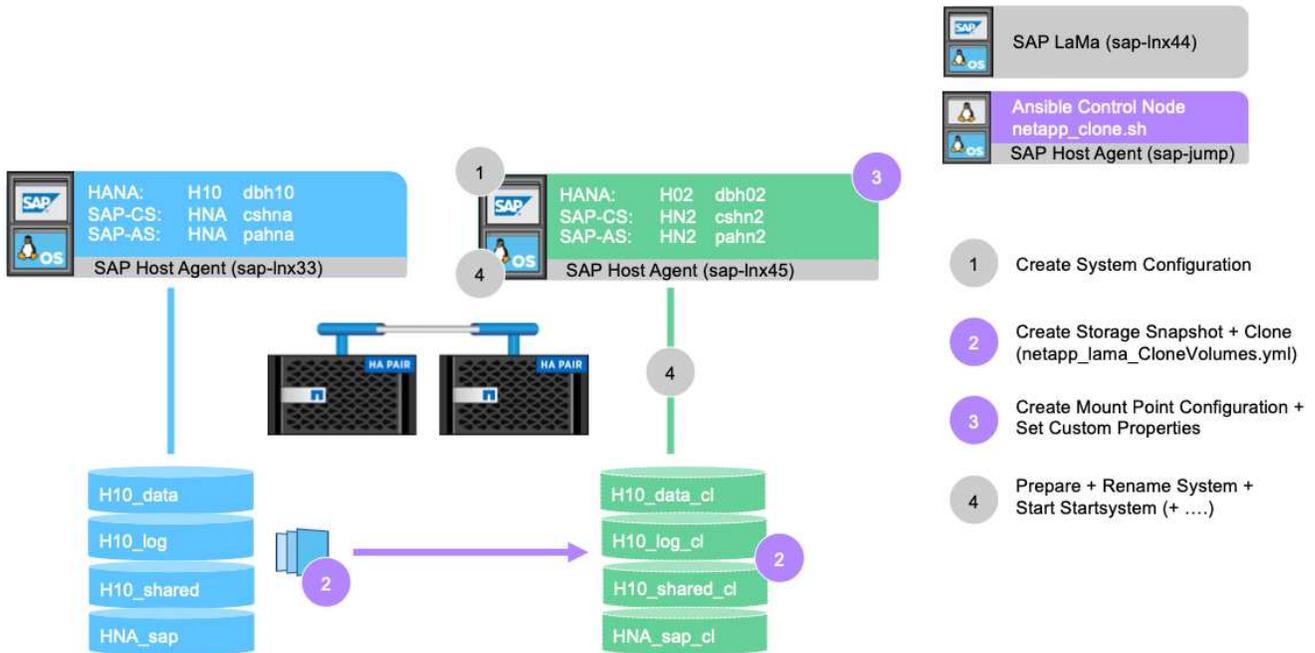


## SAP LaMa provisioning workflow - copy system

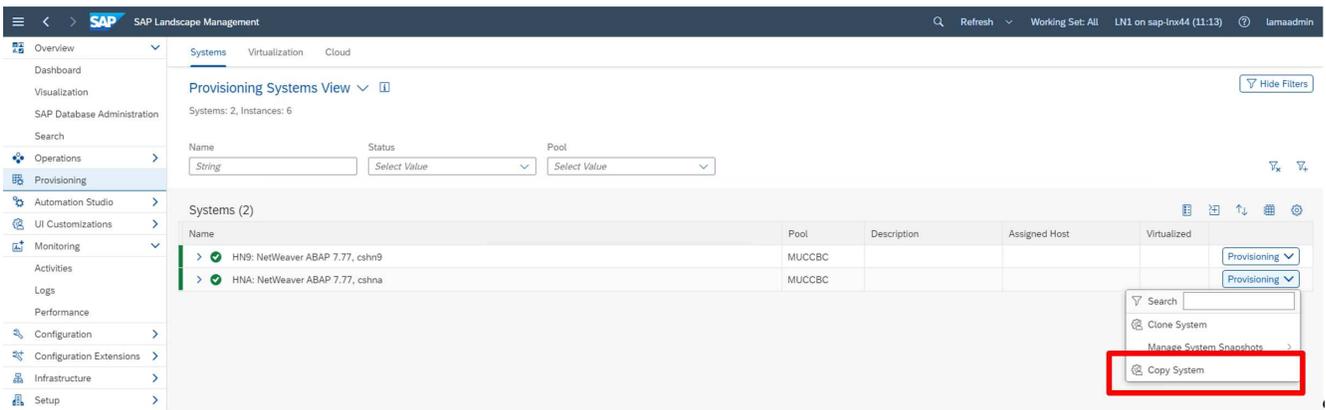
The following figure highlights the primary steps executed with the system copy workflow.



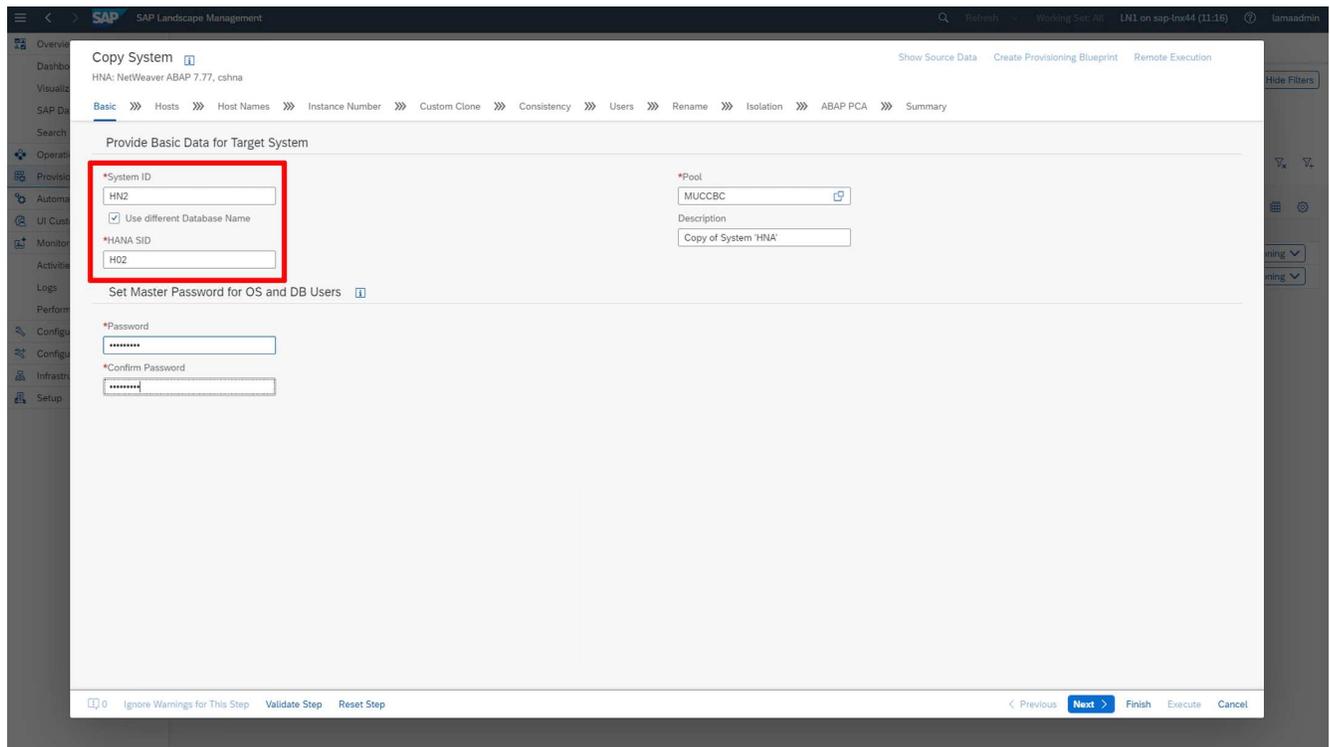
In this chapter, we briefly discuss the differences for the system clone workflow and input screens. As you can see in the following image, nothing changes in the storage workflow.



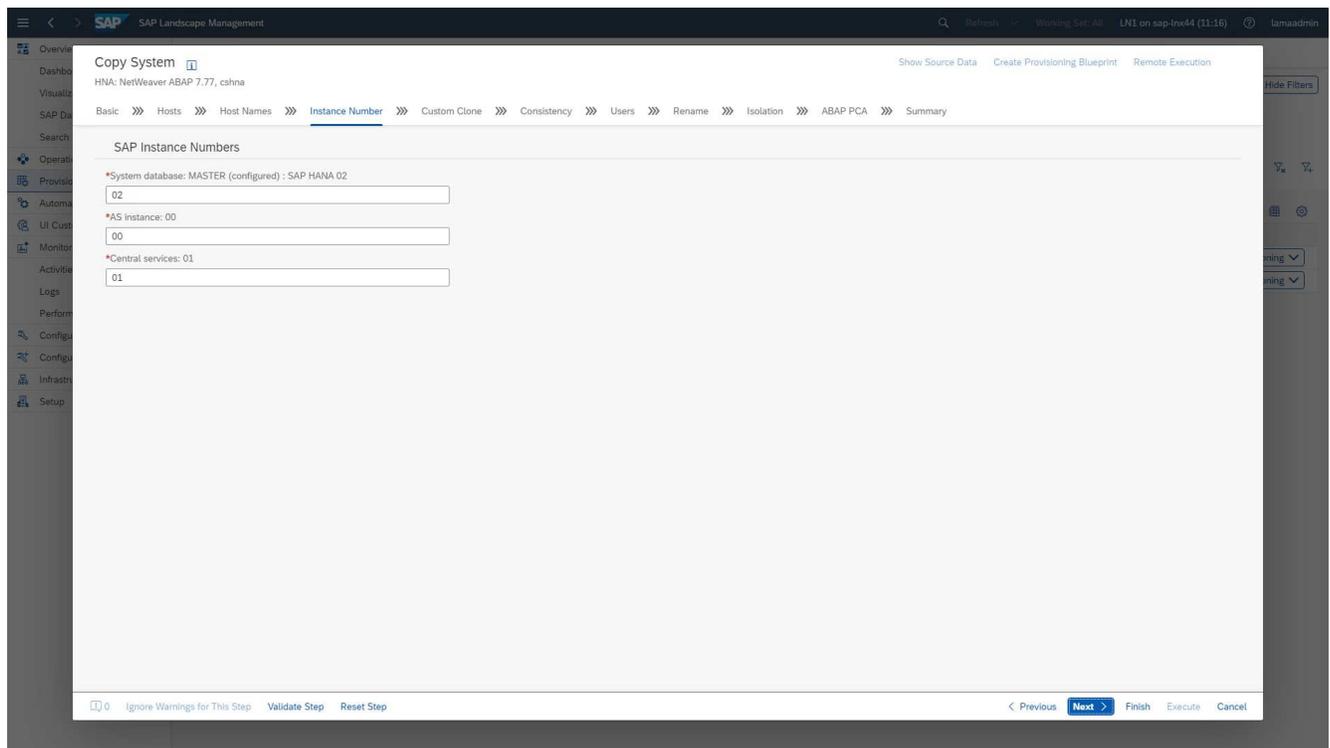
1. The system copy workflow can be started when the system is prepared accordingly. This is not a specific task for this configuration, and we do not explain it in detail. If you need further information, review the SAP LaMa documentation.



2. During the copy workflow, the system is renamed, as must be specified in the first screen.

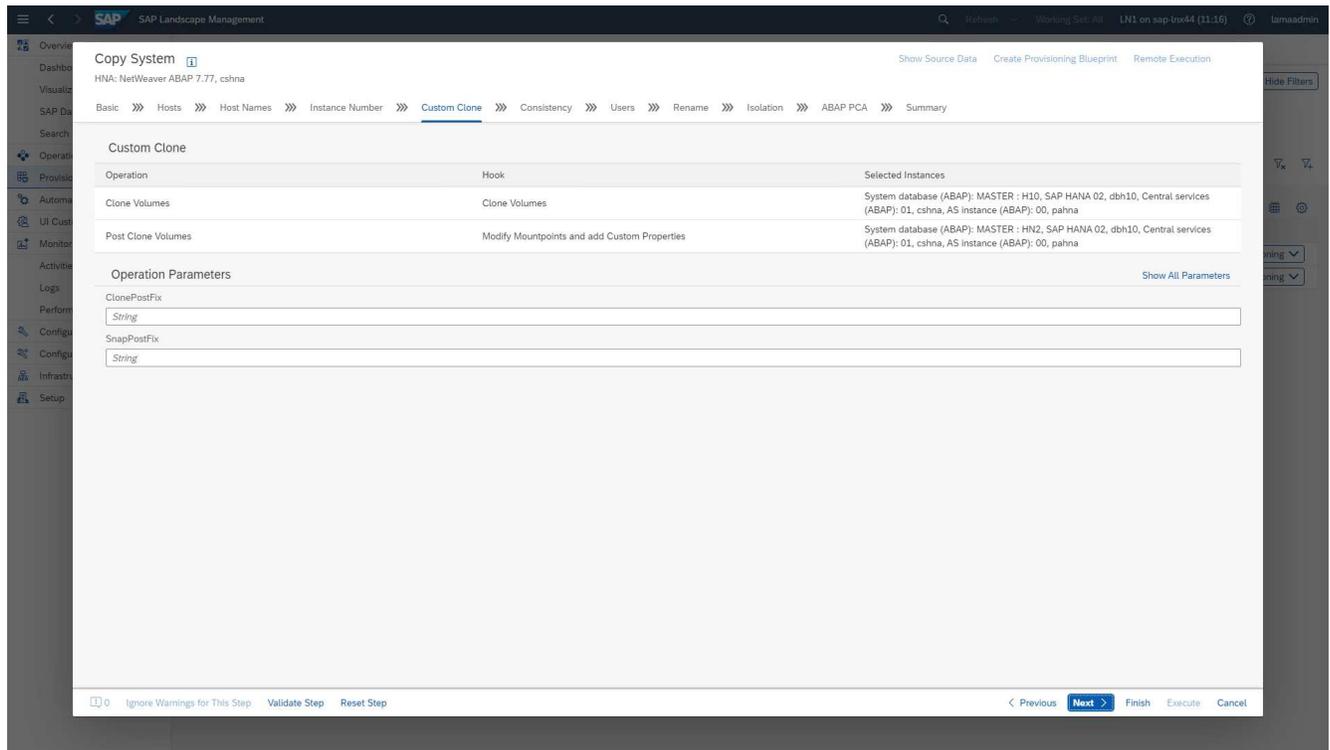


3. During the workflow, you can change the instance numbers.

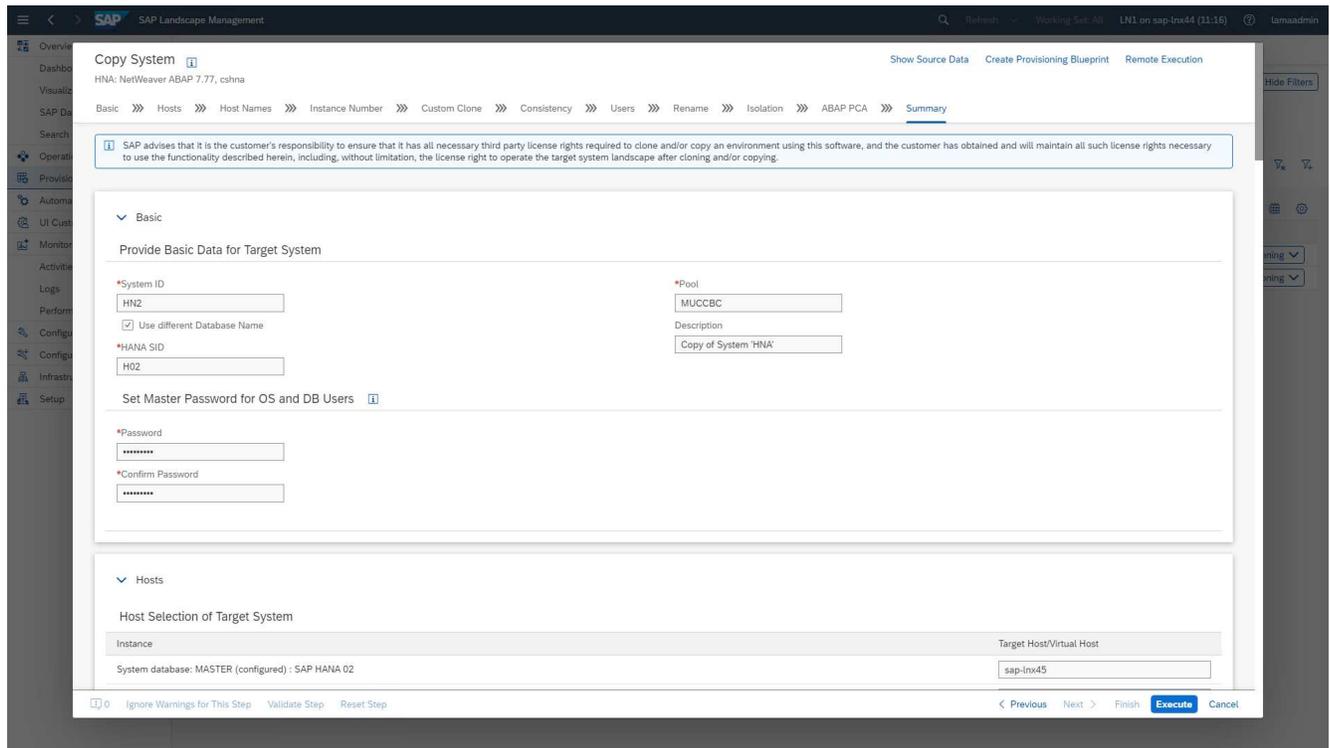


Changing instance numbers has not been tested and might require changes in the provider script.

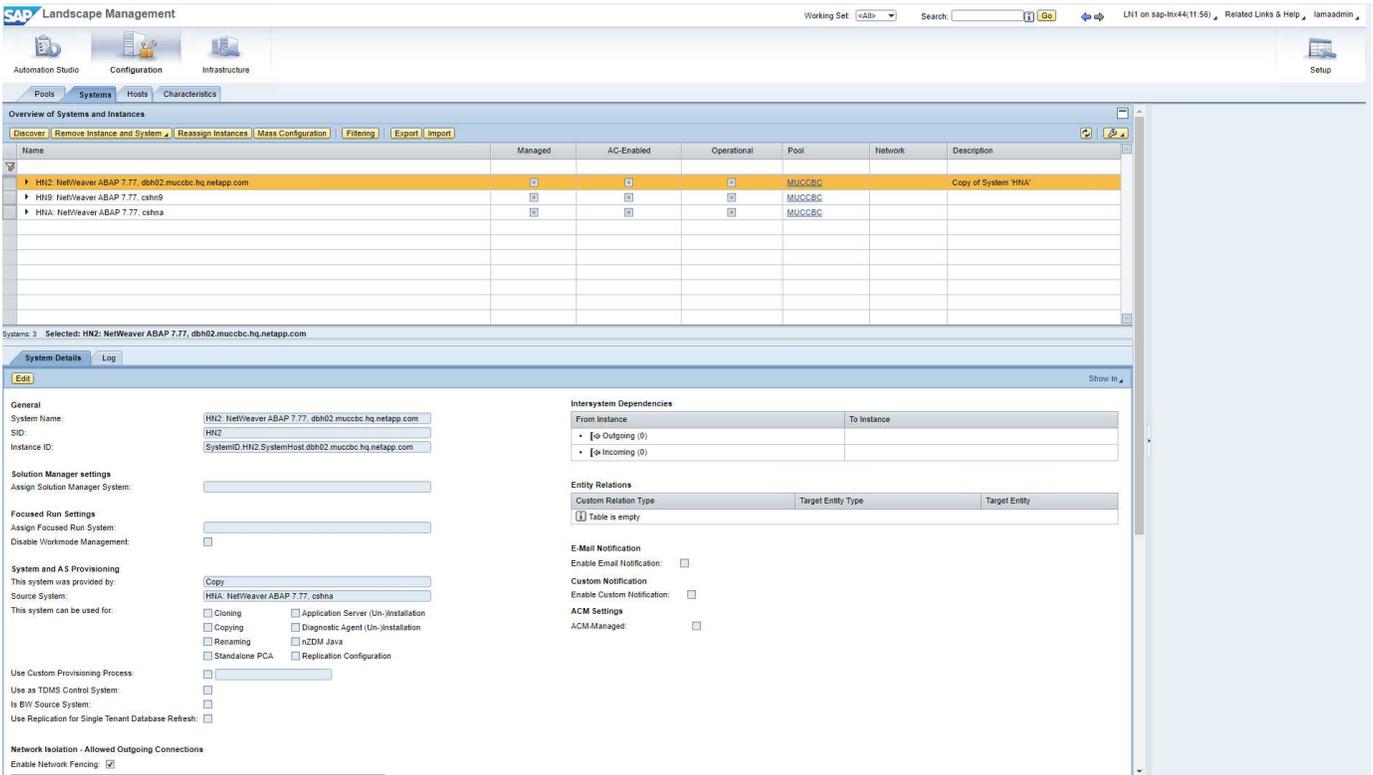
4. As described, the **Custom Clone** screen does not differ from the cloning workflow, as is shown here.



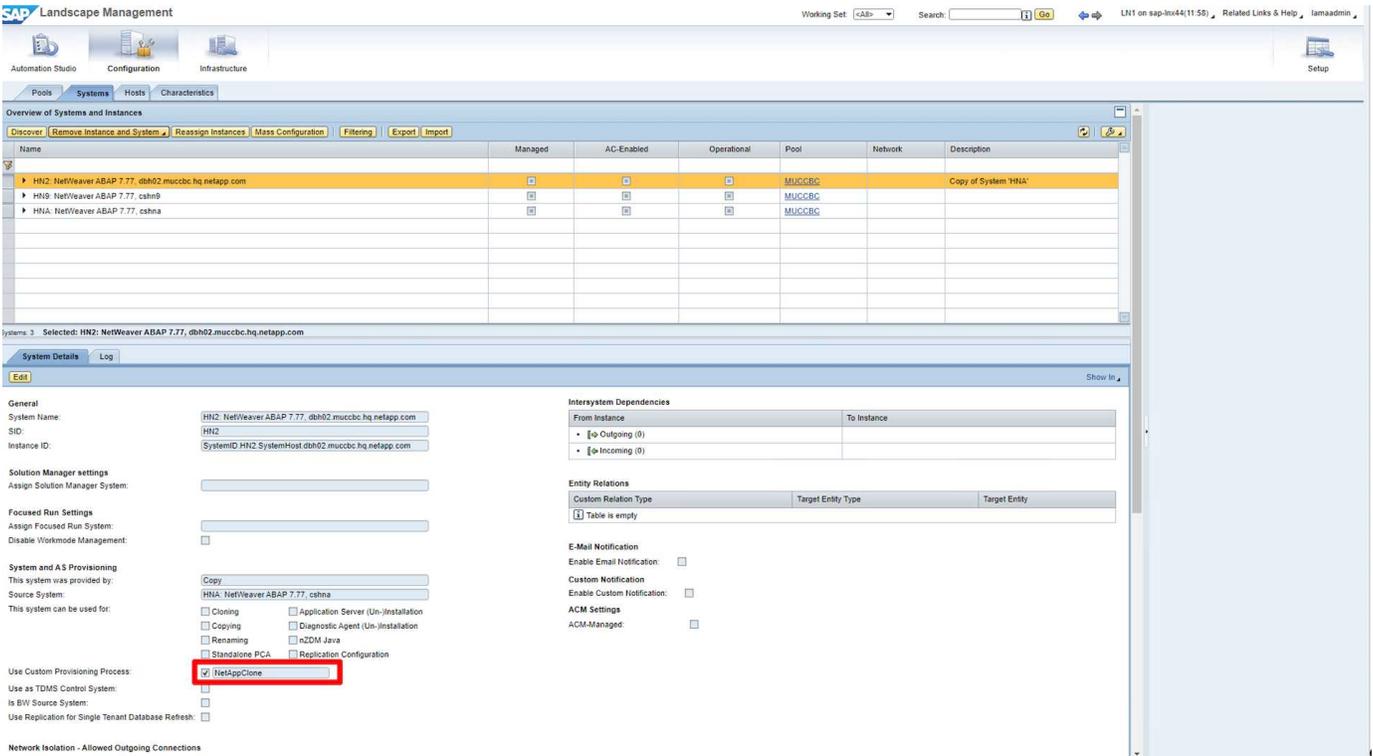
5. As we already described, the remaining input masks do not deviate from the standard, and we do not go into them any further here. The final screen shows a summary, and execution can now be started.



After the copy process, the target instance is not enabled for the custom cloning process.



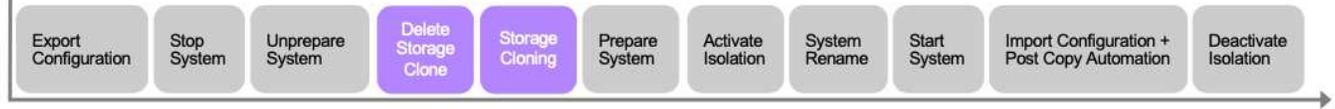
It must be adopted manually to run the pre-hook step during the system destroy process because a constraint is set and would prevent execution.



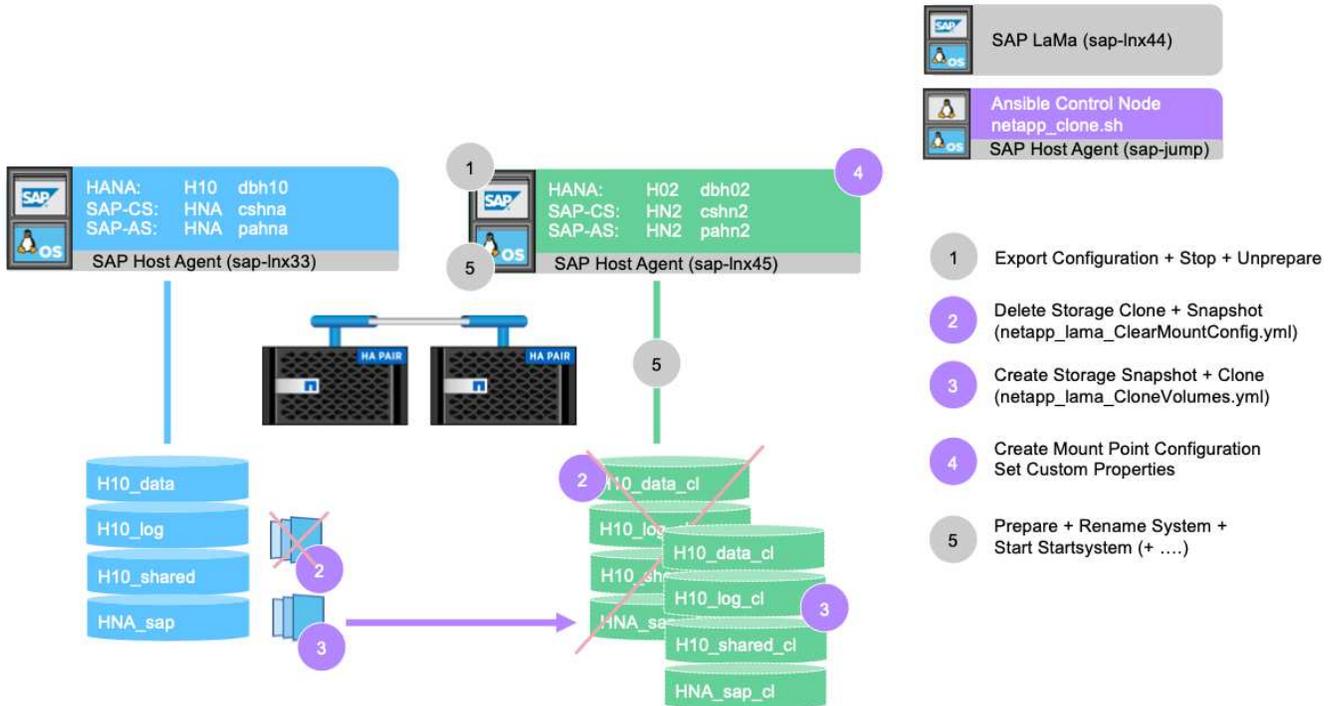
## SAP LaMa provisioning workflow - system refresh

The following figure highlights the main steps executed with the system refresh workflow.

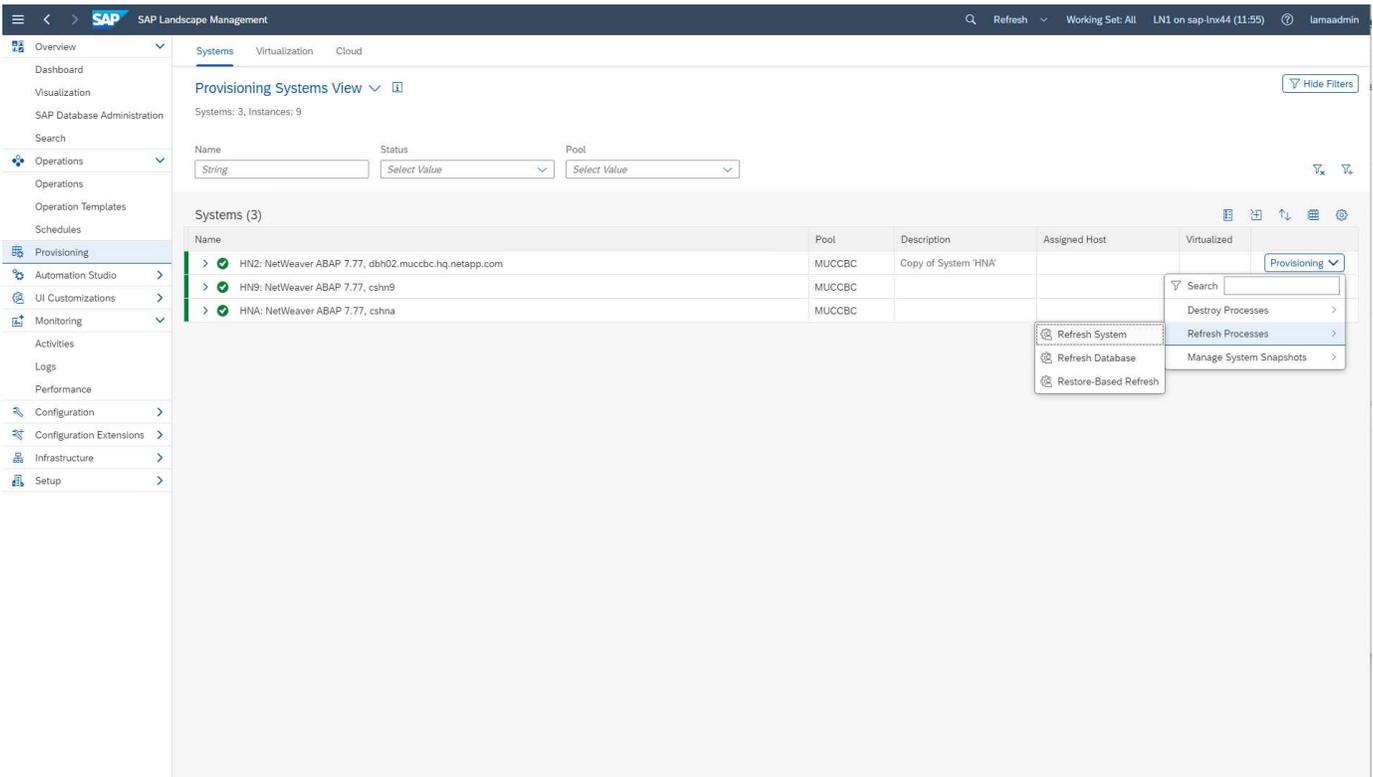
## System Refresh



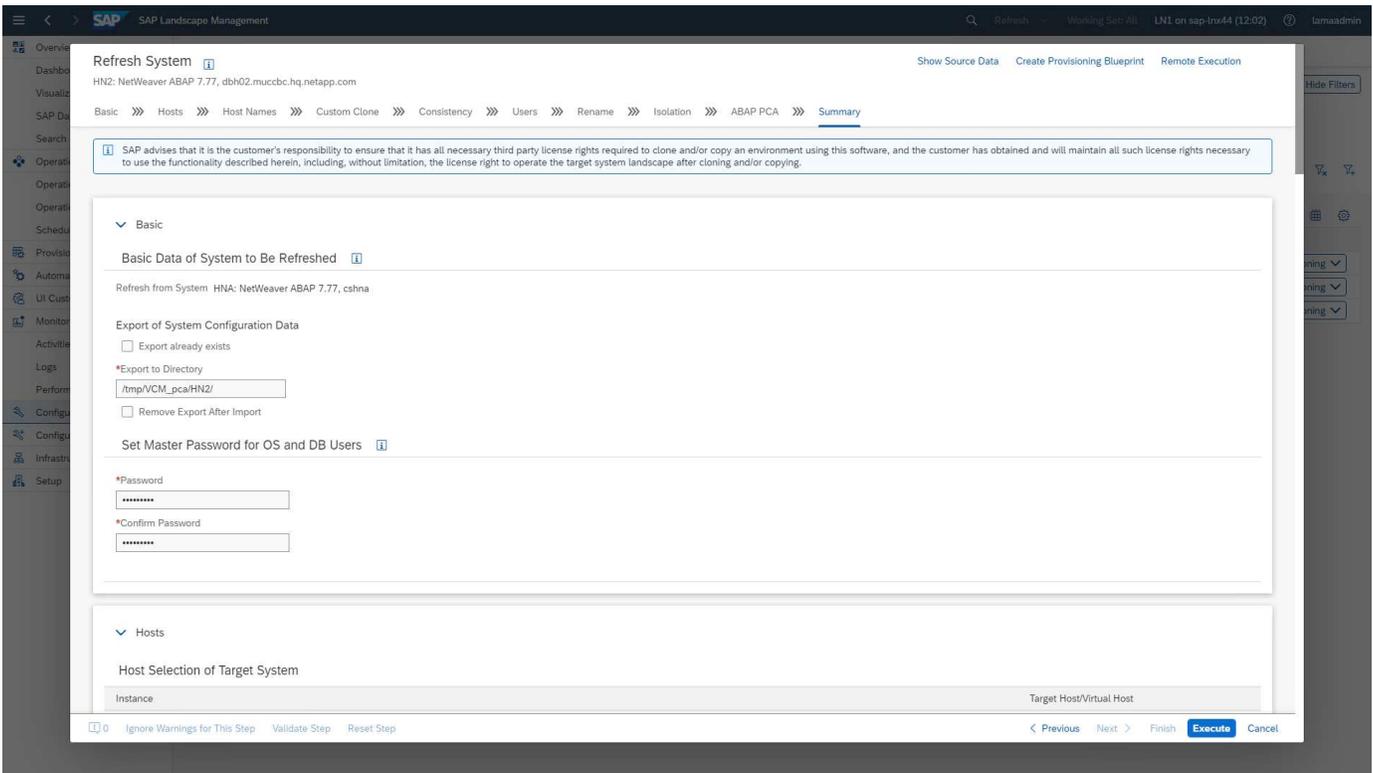
During the refresh workflow, the storage clone must be deleted. You can use the same Ansible playbook as for the system destroy workflow. However, the custom hook is defined to a different step, so the playbook is named accordingly. The process step for the clone doesn't differ.



The refresh workflow can be triggered through the provisioning screen for a copied system.



Again, nothing differs in the input screens from the standard, and the workflow execution can be started from the summary screen.



## Provider script configuration and Ansible playbooks

The following provider configuration file, execution script, and Ansible playbooks are used

during the sample deployment and workflow execution in this documentation.



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

### Provider configuration file `netapp_clone.conf`

The configuration file is created as described in the [SAP LaMa Documentation - Configuring SAP Host Agent Registered Scripts](#). This configuration file must be located on the Ansible control node where the SAP host agent is installed.

The configured os-user `sapuser` must have the appropriate permissions to execute the script and the called Ansible playbooks. You can place the script in a common script directory. SAP LaMa can provide multiple parameters when calling the script.

In addition to the custom parameters, `PARAM_ClonePostFix`, `PROP_ClonePostFix`, `PARAM_ClonePostFix`, and `PROP_ClonePostFix`, many others can be handed over, as is shown in the [SAP LaMa Documentation](#).

```
root@sap-jump:~# cat /usr/sap/hostctrl/exe/operations.d/netapp_clone.conf
Name: netapp_clone
Username: sapuser
Description: NetApp Clone for Custom Provisioning
Command: /usr/sap/scripts/netapp_clone.sh
--HookOperationName=${HookOperationName} --SAPSYSTEMNAME=${SAPSYSTEMNAME}
--SAPSYSTEM=${SAPSYSTEM} --MOUNT_XML_PATH=${MOUNT_XML_PATH}
--PARAM_ClonePostFix=${PARAM_ClonePostFix} --PARAM_SnapPostFix=${PARAM
-SnapPostFix} --PROP_ClonePostFix=${PROP_ClonePostFix}
--PROP_SnapPostFix=${PROP_SnapPostFix}
--SAP_LVM_SRC_SID=${SAP_LVM_SRC_SID}
--SAP_LVM_TARGET_SID=${SAP_LVM_TARGET_SID}
ResulConverter: hook
Platform: Unix
```

### Provider script `netapp_clone.sh`

The provider script must be stored in `/usr/sap/scripts` as configured in the provider configuration file.

#### Variables

The following variables are hard coded in the script and must be adapted accordingly.

- `PRIMARY_CLUSTER=<hostname of netapp cluster>`
- `PRIMARY_SVM=<SVM name where source system volumes are stored>`

The certificate files `PRIMARY_KEYFILE=/usr/sap/scripts/ansible/certs/ontap.key` and `PRIMARY_CERTFILE=/usr/sap/scripts/ansible/certs/ontap.pem` must be provided as described in [NetApp Ansible modules - Prepare ONTAP](#).



If different clusters or SVMs are required for different SAP systems, these variables can be added as parameters in the SAP LaMa provider definition.

### Function: create inventory file

To make Ansible playbook execution more dynamic, an `inventory.yml` file is created on the fly. Some static values are configured in the variable section and some are dynamically created during execution.

### Function: run Ansible playbook

This function is used to execute the Ansible playbook together with the dynamically created `inventory.yml` file. The naming convention for the playbooks is `netapp_lama_{{HookOperationName}}.yml`. The values for `{{HookOperationName}}` is dependent on the LaMa operation and handed over by LaMa as a command line parameter.

### Section Main

This section contains the main execution plan. The variable `{{HookOperationName}}` contains the name of the LaMa replacement step and is provided by LaMa when the script is called.

- Values with the system clone and system copy provisioning workflow:
  - `CloneVolumes`
  - `PostCloneVolumes`
- Value with the system destroy workflow:
  - `ServiceConfigRemoval`
- Value with the system refresh workflow:
  - `ClearMountConfig`

### HookOperationName = CloneVolumes

With this step, the Ansible playbook is executed, which triggers the Snapshot copy and cloning operation. The volume names and mount configuration are handed over by SAP LaMa through an XML file defined in the variable `$MOUNT_XML_PATH`. This file is saved because it is used later in the step `FinalizeCloneVolumes` to create the new mount-point configuration. The volume names are extracted from the XML file and the Ansible cloning playbook is executed for each volume.



In this example, the AS instance and the central services share the same volume. Therefore, volume cloning is only executed when the SAP instance number (`$SAPSYSTEM`) is not 01. This might differ in other environments and must be changed accordingly.

### HookOperationName = PostCloneVolumes

During this step, the custom properties `ClonePostFix` and `SnapPostFix` and the mount point configuration for the target system are maintained.

The custom properties are used later as input when the system is decommissioned during the `ServiceConfigRemoval` or `ClearMountConfig` phase. The system is designed to preserve the settings of the custom parameters that were specified during the system provisioning workflow.

The values used in this example are `ClonePostFix=_clone_20221115` and

SnapPostFix=\_snap\_20221115.

For the volume HN9\_sap, the dynamically created Ansible file includes the following values: datavolumename: HN9\_sap, snapshotpostfix: \_snap\_20221115, and clonepostfix: \_clone\_20221115.

Which leads into the snapshot name on the volume HN9\_sap HN9\_sap\_snap\_20221115 and the created volume clone name HN9\_sap\_clone\_20221115.



Custom properties could be used in any way to preserve parameters used during the provisioning process.

The mount point configuration is extracted from the XML file that has been handed over by LaMa in the CloneVolume step. The ClonePostFix is added to the volume names and send back to LaMa through the default script output. The functionality is described in [SAP Note 1889590](#).



In this example, qtrees on the storage system are used as a common way to place different data on a single volume. For example, HN9\_sap holds the mount points for /usr/sap/HN9, /sapmnt/HN9, and /home/hn9adm. Subdirectories work in the same way. This might differ in other environments and must be changed accordingly.

#### **HookOperationName = ServiceConfigRemoval**

In this step, the Ansible playbook that is responsible for the deletion of the volume clones is running.

The volume names are handed over by SAP LaMa through the mount configuration file, and the custom properties ClonePostFix and SnapPostFix are used to hand over the values of the parameters originally specified during the system provisioning workflow (see the note at HookOperationName = PostCloneVolumes).

The volume names are extracted from the xml file, and the Ansible cloning playbook is executed for each volume.



In this example, the AS instance and the central services share the same volume. Therefore, the volume deletion is only executed when the SAP instance number (\$SAPSYSTEM) is not 01. This might differ in other environments and must be changed accordingly.

#### **HookOperationName = ClearMountConfig**

In this step, the Ansible playbook that is responsible for the deletion of the volume clones during a system refresh workflow is running.

The volume names are handed over by SAP LaMa through the mount configuration file, and the custom properties ClonePostFix and SnapPostFix are used to hand over the values of the parameters originally specified during the system provisioning workflow.

The volume names are extracted from the XML file and the Ansible cloning playbook is executed for each volume.



In this example, the AS instance and the central services share the same volume. Therefore, volume deletion is only executed when the SAP instance number (\$SAPSYSTEM) is not 01. This might differ in other environments and must be changed accordingly.

```

root@sap-jump:~# cat /usr/sap/scripts/netapp_clone.sh
#!/bin/bash
#Section - Variables
#####
VERSION="Version 0.9"
#Path for ansible play-books
ANSIBLE_PATH=/usr/sap/scripts/ansible
#Values for Ansible Inventory File
PRIMARY_CLUSTER=grenada
PRIMARY_SVM=svm-sap01
PRIMARY_KEYFILE=/usr/sap/scripts/ansible/certs/ontap.key
PRIMARY_CERTFILE=/usr/sap/scripts/ansible/certs/ontap.pem
#Default Variable if PARAM ClonePostFix / SnapPostFix is not maintained in
LaMa
DefaultPostFix=_clone_1
#TMP Files - used during execution
YAML_TMP=/tmp/inventory_ansible_clone_tmp_$$$.yml
TMPFILE=/tmp/tmpfile.$$
MY_NAME="`basename $0`"
BASE_SCRIPT_DIR="`dirname $0`"
#Sendig Script Version and run options to LaMa Log
echo "[DEBUG]: Running Script $MY_NAME $VERSION"
echo "[DEBUG]: $MY_NAME $@"
#Command declared in the netapp_clone.conf Provider definition
#Command: /usr/sap/scripts/netapp_clone.sh
--HookOperationName=${HookOperationName} --SAPSYSTEMNAME=${SAPSYSTEMNAME}
--SAPSYSTEM=${SAPSYSTEM} --MOUNT_XML_PATH=${MOUNT_XML_PATH}
--PARAM_ClonePostFix=${PARAM-ClonePostFix} --PARAM_SnapPostFix=${PARAM
-SnapPostFix} --PROP_ClonePostFix=${PROP-ClonePostFix}
--PROP_SnapPostFix=${PROP-SnapPostFix}
--SAP_LVM_SRC_SID=${SAP_LVM_SRC_SID}
--SAP_LVM_TARGET_SID=${SAP_LVM_TARGET_SID}
#Reading Input Variables hand over by LaMa
for i in "$@"
do
case $i in
--HookOperationName=*)
HookOperationName="${i#*=}";shift;;
--SAPSYSTEMNAME=*)
SAPSYSTEMNAME="${i#*=}";shift;;
--SAPSYSTEM=*)
SAPSYSTEM="${i#*=}";shift;;
--MOUNT_XML_PATH=*)
MOUNT_XML_PATH="${i#*=}";shift;;
--PARAM_ClonePostFix=*)

```

```

PARAM_ClonePostFix="${i#*=}";shift;;
--PARAM_SnapPostFix=*)
PARAM_SnapPostFix="${i#*=}";shift;;
--PROP_ClonePostFix=*)
PROP_ClonePostFix="${i#*=}";shift;;
--PROP_SnapPostFix=*)
PROP_SnapPostFix="${i#*=}";shift;;
--SAP_LVM_SRC_SID=*)
SAP_LVM_SRC_SID="${i#*=}";shift;;
--SAP_LVM_TARGET_SID=*)
SAP_LVM_TARGET_SID="${i#*=}";shift;;
*)
# unknown option
;;
esac
done
#If Parameters not provided by the User - defaulting to DefaultPostFix
if [ -z $PARAM_ClonePostFix ]; then PARAM_ClonePostFix=$DefaultPostFix;fi
if [ -z $PARAM_SnapPostFix ]; then PARAM_SnapPostFix=$DefaultPostFix;fi
#Section - Functions
#####
#Function Create (Inventory) YML File
#####
create_yml_file()
{
echo "ontapservers:">$YAML_TMP
echo " hosts:">>$YAML_TMP
echo "   ${PRIMARY_CLUSTER}:">>$YAML_TMP
echo "   ansible_host: "'"$PRIMARY_CLUSTER'">>$YAML_TMP
echo "   keyfile: "'"$PRIMARY_KEYFILE'">>$YAML_TMP
echo "   certfile: "'"$PRIMARY_CERTFILE'">>$YAML_TMP
echo "   svmname: "'"$PRIMARY_SVM'">>$YAML_TMP
echo "   datavolumename: "'"$datavolumename'">>$YAML_TMP
echo "   snapshotpostfix: "'"$snapshotpostfix'">>$YAML_TMP
echo "   clonepostfix: "'"$clonepostfix'">>$YAML_TMP
}
#Function run ansible-playbook
#####
run_ansible_playbook()
{
echo "[DEBUG]: Running ansible playbook
netapp_lama_${HookOperationName}.yml on Volume $datavolumename"
ansible-playbook -i $YAML_TMP
$ANSIBLE_PATH/netapp_lama_${HookOperationName}.yml
}
#Section - Main

```

```
#####
#HookOperationName - CloneVolumes
#####
if [ $HookOperationName = CloneVolumes ] ;then
#save mount xml for later usage - used in Section FinalizeCloneVolumes to
generate the mountpoints
echo "[DEBUG]: saving mount config..."
cp $MOUNT_XML_PATH /tmp/mount_config_${SAPSYSTEMNAME}_${SAPSYSTEM}.xml
#Instance 00 + 01 share the same volumes - clone needs to be done once
if [ $SAPSYSTEM != 01 ]; then
#generating Volume List - assuming usage of qtrees - "IP-
Adress:/VolumeName/qtrees"
xmlFile=/tmp/mount_config_${SAPSYSTEMNAME}_${SAPSYSTEM}.xml
if [ -e $TMPFILE ];then rm $TMPFILE;fi
numMounts=`xml_grep --count "/mountconfig/mount" $xmlFile | grep "total: "
| awk '{ print $2 }'`
i=1
while [ $i -le $numMounts ]; do
    xmllint --xpath "/mountconfig/mount[$i]/exportpath/text()" $xmlFile
|awk -F"/" '{print $2}' >>$TMPFILE
i=$((i + 1))
done
DATAVOLUMES=`cat $TMPFILE |sort -u`
#Create yml file and rund playbook for each volume
for I in $DATAVOLUMES; do
datavolumename="$I"
snapshotpostfix="$PARAM_SnapPostFix"
clonepostfix="$PARAM_ClonePostFix"
create_yml_file
run_ansible_playbook
done
else
echo "[DEBUG]: Doing nothing .... Volume cloned in different Task"
fi
fi
#HookOperationName - PostCloneVolumes
#####
if [ $HookOperationName = PostCloneVolumes ] ;then
#Reporting Properties back to LaMa Config for Cloned System
echo "[RESULT]:Property:ClonePostFix=$PARAM_ClonePostFix"
echo "[RESULT]:Property:SnapPostFix=$PARAM_SnapPostFix"
#Create MountPoint Config for Cloned Instances and report back to LaMa
according to SAP Note: https://launchpad.support.sap.com/#/notes/1889590
echo "MountDataBegin"
echo '<?xml version="1.0" encoding="UTF-8"?>'
echo "<mountconfig>"
```

```

xmlFile=/tmp/mount_config_${SAPSYSTEMNAME}_${SAPSYSTEM}.xml
numMounts=`xml_grep --count "/mountconfig/mount" $xmlFile | grep "total: "
| awk '{ print $2 }'`
i=1
while [ $i -le $numMounts ]; do
MOUNTPOINT=`xmllint --xpath "/mountconfig/mount[$i]/mountpoint/text()"
$xmlFile`;
    EXPORTPATH=`xmllint --xpath
"/mountconfig/mount[$i]/exportpath/text()" $xmlFile`;
    OPTIONS=`xmllint --xpath "/mountconfig/mount[$i]/options/text()"
$xmlFile`;
#Adopt Exportpath and add Clonepostfix - assuming usage of qtrees - "IP-
Adress:/VolumeName/qtrees"
TMPFIELD1=`echo $EXPORTPATH|awk -F"/" '{print $1}'`
TMPFIELD2=`echo $EXPORTPATH|awk -F"/" '{print $2}'`
TMPFIELD3=`echo $EXPORTPATH|awk -F"/" '{print $3}'`
EXPORTPATH=$TMPFIELD1":/"${TMPFIELD2}$PARAM_ClonePostFix"/"${TMPFIELD3
echo -e "\t<mount fstype="nfs" storagetype="NETFS">"
echo -e "\t\t<mountpoint>${MOUNTPOINT}</mountpoint>"
echo -e "\t\t<exportpath>${EXPORTPATH}</exportpath>"
echo -e "\t\t<options>${OPTIONS}</options>"
echo -e "\t</mount>"
i=$((i + 1))
done
echo "</mountconfig>"
echo "MountDataEnd"
#Finished MountPoint Config
#Cleanup Temporary Files
rm $xmlFile
fi
#HookOperationName - ServiceConfigRemoval
#####
if [ $HookOperationName = ServiceConfigRemoval ] ;then
#Assure that Properties ClonePostFix and SnapPostfix has been configured
through the provisioning process
if [ -z $PROP_ClonePostFix ]; then echo "[ERROR]: Propertiy ClonePostFix
is not handed over - please investigate";exit 5;fi
if [ -z $PROP_SnapPostFix ]; then echo "[ERROR]: Propertiy SnapPostFix is
not handed over - please investigate";exit 5;fi
#Instance 00 + 01 share the same volumes - clone delete needs to be done
once
if [ $SAPSYSTEM != 01 ]; then
#generating Volume List - assuming usage of qtrees - "IP-
Adress:/VolumeName/qtrees"
xmlFile=$MOUNT_XML_PATH
if [ -e $TMPFILE ];then rm $TMPFILE;fi

```

```

numMounts=`xml_grep --count "/mountconfig/mount" $xmlFile | grep "total: "
| awk '{ print $2 }'`
i=1
while [ $i -le $numMounts ]; do
    xmllint --xpath "/mountconfig/mount[$i]/exportpath/text()" $xmlFile
|awk -F"/" '{print $2}' >>$TMPFILE
i=$((i + 1))
done
DATAVOLUMES=`cat $TMPFILE |sort -u| awk -F $PROP_ClonePostFix '{ print $1
}'`
#Create yml file and rund playbook for each volume
for I in $DATAVOLUMES; do
datavolumename="$I"
snapshotpostfix="$PROP_SnapPostFix"
clonepostfix="$PROP_ClonePostFix"
create_yml_file
run_ansible_playbook
done
else
echo "[DEBUG]: Doing nothing .... Volume deleted in different Task"
fi
#Cleanup Temporary Files
rm $xmlFile
fi
#HookOperationName - ClearMountConfig
#####
if [ $HookOperationName = ClearMountConfig ] ;then
    #Assure that Properties ClonePostFix and SnapPostfix has been
configured through the provisioning process
    if [ -z $PROP_ClonePostFix ]; then echo "[ERROR]: Propertiy
ClonePostFix is not handed over - please investigate";exit 5;fi
    if [ -z $PROP_SnapPostFix ]; then echo "[ERROR]: Propertiy
SnapPostFix is not handed over - please investigate";exit 5;fi
    #Instance 00 + 01 share the same volumes - clone delete needs to
be done once
    if [ $SAPSYSTEM != 01 ]; then
        #generating Volume List - assuming usage of qtrees - "IP-
Adress:/VolumeName/qtree"
        xmlFile=$MOUNT_XML_PATH
        if [ -e $TMPFILE ];then rm $TMPFILE;fi
        numMounts=`xml_grep --count "/mountconfig/mount" $xmlFile
| grep "total: " | awk '{ print $2 }'`
        i=1
        while [ $i -le $numMounts ]; do
            xmllint --xpath
"/mountconfig/mount[$i]/exportpath/text()" $xmlFile |awk -F"/" '{print

```

```

$2}' >>$TMPFILE
        i=$((i + 1))
    done
    DATAVOLUMES=`cat $TMPFILE |sort -u| awk -F
$PROP_ClonePostFix '{ print $1 }'`
    #Create yml file and rund playbook for each volume
    for I in $DATAVOLUMES; do
        datavolumename="$I"
        snapshotpostfix="$PROP_SnapPostFix"
        clonepostfix="$PROP_ClonePostFix"
        create_yml_file
        run_ansible_playbook
    done
else
    echo "[DEBUG]: Doing nothing .... Volume deleted in
different Task"
    fi
    #Cleanup Temporary Files
    rm $xmlFile
fi
#Cleanup
#####
#Cleanup Temporary Files
if [ -e $TMPFILE ];then rm $TMPFILE;fi
if [ -e $YAML_TMP ];then rm $YAML_TMP;fi
exit 0

```

### Ansible Playbook netapp\_lama\_CloneVolumes.yml

The playbook that is executed during the CloneVolumes step of the LaMa system clone workflow is a combination of `create_snapshot.yml` and `create_clone.yml` (see [NetApp Ansible modules - YAML files](#)). This playbook can be easily extended to cover additional use cases like cloning from secondary and clone split operations.

```

root@sap-jump:~# cat /usr/sap/scripts/ansible/netapp_lama_CloneVolumes.yml
---
- hosts: ontapserver
  connection: local
  collections:
    - netapp.ontap
  gather_facts: false
  name: netapp_lama_CloneVolumes
  tasks:
  - name: Create SnapShot
    na_ontap_snapshot:
      state: present
      snapshot: "{{ datavolumename }}{{ snapshotpostfix }}"
      use_rest: always
      volume: "{{ datavolumename }}"
      vsserver: "{{ svmname }}"
      hostname: "{{ inventory_hostname }}"
      cert_filepath: "{{ certfile }}"
      key_filepath: "{{ keyfile }}"
      https: true
      validate_certs: false
  - name: Clone Volume
    na_ontap_volume_clone:
      state: present
      name: "{{ datavolumename }}{{ clonepostfix }}"
      use_rest: always
      vsserver: "{{ svmname }}"
      junction_path: '/{{ datavolumename }}{{ clonepostfix }}'
      parent_volume: "{{ datavolumename }}"
      parent_snapshot: "{{ datavolumename }}{{ snapshotpostfix }}"
      hostname: "{{ inventory_hostname }}"
      cert_filepath: "{{ certfile }}"
      key_filepath: "{{ keyfile }}"
      https: true
      validate_certs: false

```

### Ansible Playbook netapp\_lama\_ServiceConfigRemoval.yml

The playbook that is executed during the ServiceConfigRemoval phase of the LaMa system destroy workflow is combination of delete\_clone.yml and delete\_snapshot.yml (see [NetApp Ansible modules - YAML files](#)). It must be aligned to the execution steps of the netapp\_lama\_CloneVolumes playbook.

```

root@sap-jump:~# cat
/usr/sap/scripts/ansible/netapp_lama_ServiceConfigRemoval.yml
---
- hosts: ontapservers
  connection: local
  collections:
    - netapp.ontap
  gather_facts: false
  name: netapp_lama_ServiceConfigRemoval
  tasks:
    - name: Delete Clone
      na_ontap_volume:
        state: absent
        name: "{{ datavolumename }}{{ clonepostfix }}"
        use_rest: always
        vsserver: "{{ svmname }}"
        wait_for_completion: True
        hostname: "{{ inventory_hostname }}"
        cert_filepath: "{{ certfile }}"
        key_filepath: "{{ keyfile }}"
        https: true
        validate_certs: false
    - name: Delete SnapShot
      na_ontap_snapshot:
        state: absent
        snapshot: "{{ datavolumename }}{{ snapshotpostfix }}"
        use_rest: always
        volume: "{{ datavolumename }}"
        vsserver: "{{ svmname }}"
        hostname: "{{ inventory_hostname }}"
        cert_filepath: "{{ certfile }}"
        key_filepath: "{{ keyfile }}"
        https: true
        validate_certs: false
root@sap-jump:~#

```

### Ansible Playbook netapp\_lama\_ClearMountConfig.yml

The playbook, which is executed during the netapp\_lama\_ClearMountConfig phase of the LaMa system refresh workflow is combination of delete\_clone.yml and delete\_snapshot.yml (see [NetApp Ansible modules - YAML files](#)). It must be aligned to the execution steps of the netapp\_lama\_CloneVolumes playbook.

```

root@sap-jump:~# cat
/usr/sap/scripts/ansible/netapp_lama_ServiceConfigRemoval.yml
---
- hosts: ontapservers
  connection: local
  collections:
    - netapp.ontap
  gather_facts: false
  name: netapp_lama_ServiceConfigRemoval
  tasks:
    - name: Delete Clone
      na_ontap_volume:
        state: absent
        name: "{{ datavolumename }}{{ clonepostfix }}"
        use_rest: always
        vsserver: "{{ svmname }}"
        wait_for_completion: True
        hostname: "{{ inventory_hostname }}"
        cert_filepath: "{{ certfile }}"
        key_filepath: "{{ keyfile }}"
        https: true
        validate_certs: false
    - name: Delete SnapShot
      na_ontap_snapshot:
        state: absent
        snapshot: "{{ datavolumename }}{{ snapshotpostfix }}"
        use_rest: always
        volume: "{{ datavolumename }}"
        vsserver: "{{ svmname }}"
        hostname: "{{ inventory_hostname }}"
        cert_filepath: "{{ certfile }}"
        key_filepath: "{{ keyfile }}"
        https: true
        validate_certs: false
root@sap-jump:~#

```

### Sample Ansible inventory.yml

This inventory file is dynamically built during workflow execution, and it is only shown here for illustration.

```
ontapservers:
  hosts:
    grenada:
      ansible_host: "grenada"
      keyfile: "/usr/sap/scripts/ansible/certs/ontap.key"
      certfile: "/usr/sap/scripts/ansible/certs/ontap.pem"
      svmname: "svm-sap01"
      datavolumename: "HN9_sap"
      snapshotpostfix: " _snap_20221115"
      clonepostfix: " _clone_20221115"
```

## Conclusion

The integration of a modern automation framework like Ansible into SAP LaMa provisioning workflows gives customers a flexible solution to address standard or more complex infrastructure requirements.

### Where to find additional information

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

- Collections in the NetApp Namespace

<https://docs.ansible.com/ansible/latest/collections/netapp/index.html>

- Documentation about Ansible Integration and Sample Ansible Playbooks

[https://github.com/sap-linuxlab/demo.netapp\\_ontap](https://github.com/sap-linuxlab/demo.netapp_ontap)

- General Ansible and NetApp Integration

<https://www.ansible.com/integrations/infrastructure/netapp>

- Blog on integrating SAP LaMa with Ansible

<https://blogs.sap.com/2020/06/08/outgoing-api-calls-from-sap-landscape-management-lama-with-automation-studio/>

- SAP Landscape Management 3.0, Enterprise Edition Documentation

<https://help.sap.com/doc/700f9a7e52c7497cad37f7c46023b7ff/3.0.11.0/en-US/4df88a8f418c5059e1000000a42189c.html#loio4df88a8f418c5059e1000000a42189c>

- SAP LaMa Documentation – Provider Definitions

<https://help.sap.com/doc/700f9a7e52c7497cad37f7c46023b7ff/3.0.11.0/en-US/bf6b3e43340a4cbcb0c0f3089715c068.html>

- SAP LaMa Documentation - Custom Hooks

<https://help.sap.com/doc/700f9a7e52c7497cad37f7c46023b7ff/3.0.11.0/en-US/139eca2f925e48738a20dbf0b56674c5.html>

- SAP LaMa Documentation - Configuring SAP Host Agent Registered Scripts

<https://help.sap.com/doc/700f9a7e52c7497cad37f7c46023b7ff/3.0.11.0/en-US/250dfc5eef4047a38bab466c295d3a49.html>

- SAP LaMa Documentation - Parameters for Custom Operations and Custom Hooks

<https://help.sap.com/doc/700f9a7e52c7497cad37f7c46023b7ff/3.0.11.0/en-US/0148e495174943de8c1c3ee1b7c9cc65.html>

- SAP LaMa Documentation - Adaptive Design

<https://help.sap.com/doc/700f9a7e52c7497cad37f7c46023b7ff/3.0.11.0/en-US/737a99e86f8743bdb8d1f6cf4b862c79.html>

- NetApp Product Documentation

<https://www.netapp.com/support-and-training/documentation/>

## Version history

Version	Date	Document version history
Version 1.0	January 2023	Initial release

# 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

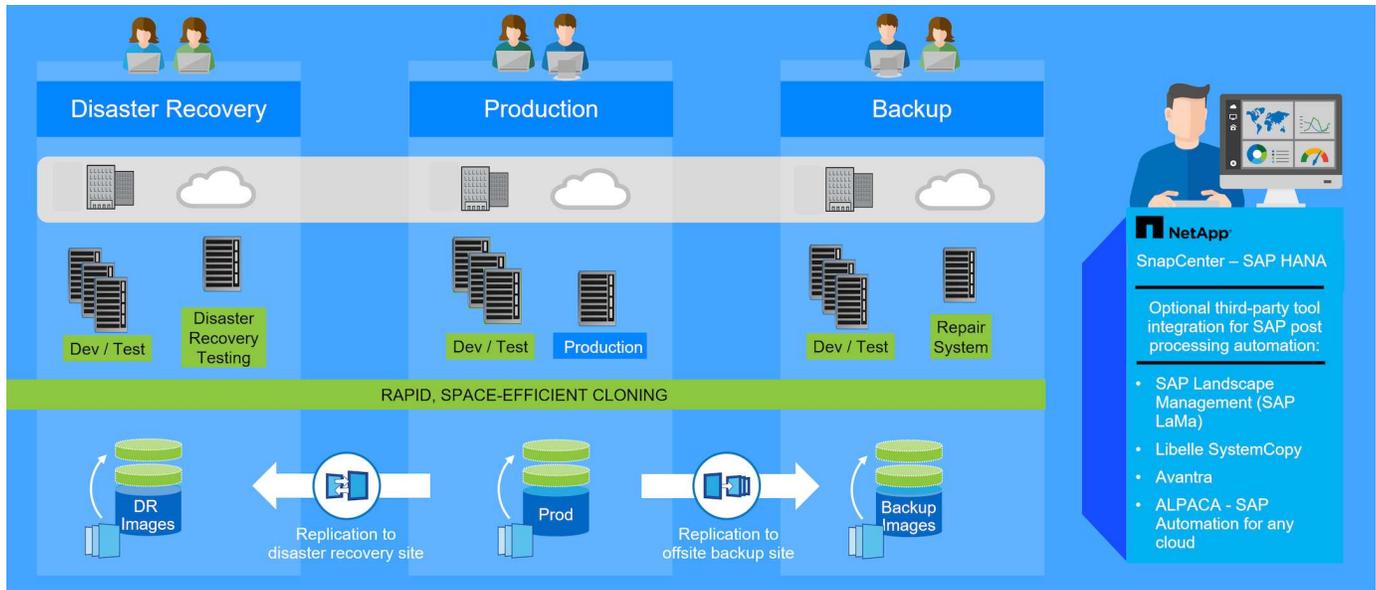
### Introduction

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.

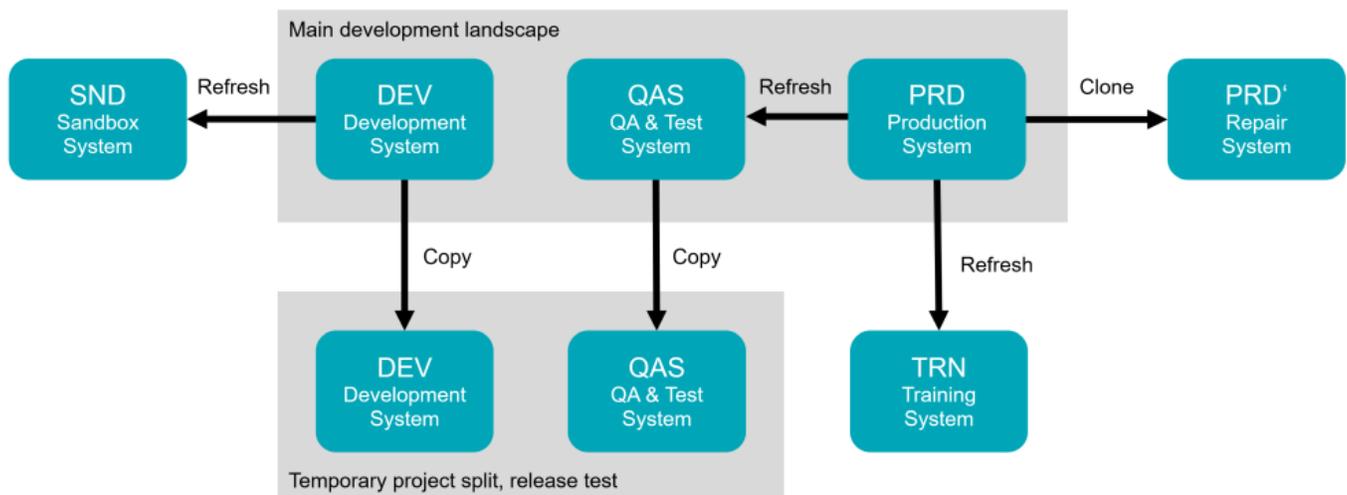


## Use cases for system refresh and cloning

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.

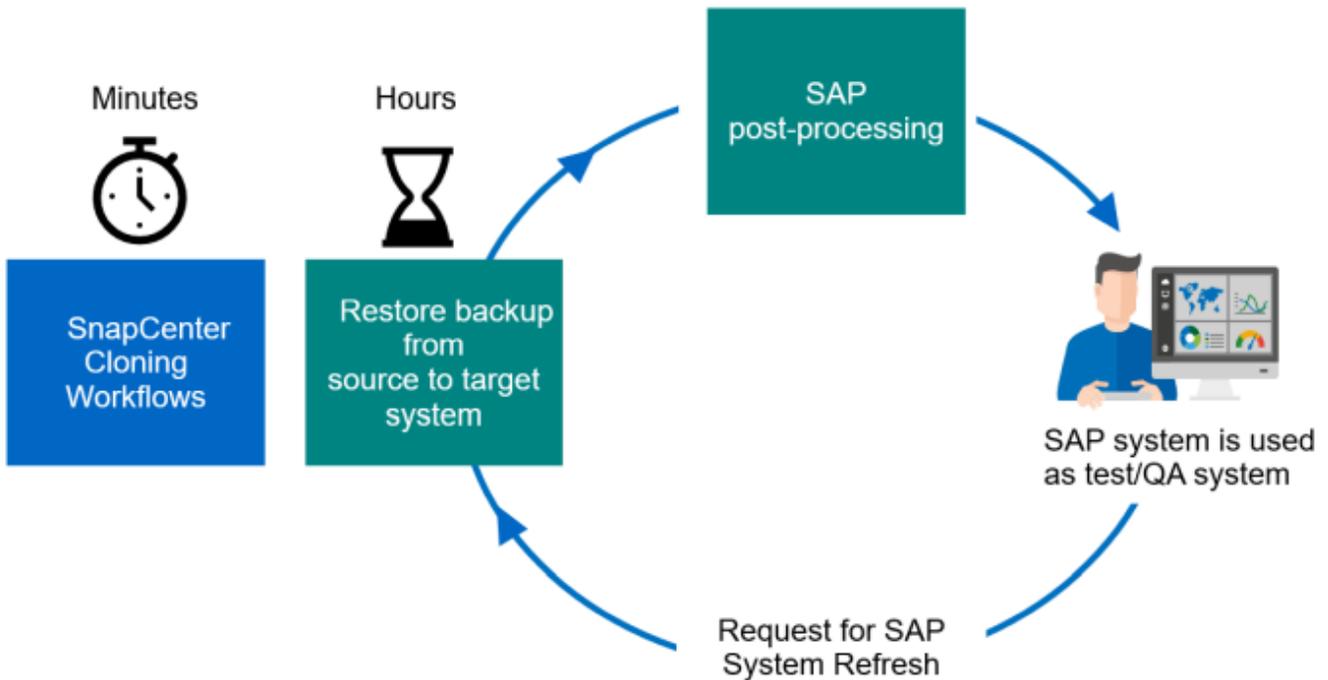
### Data refresh of QA, test, sandbox, or training systems

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

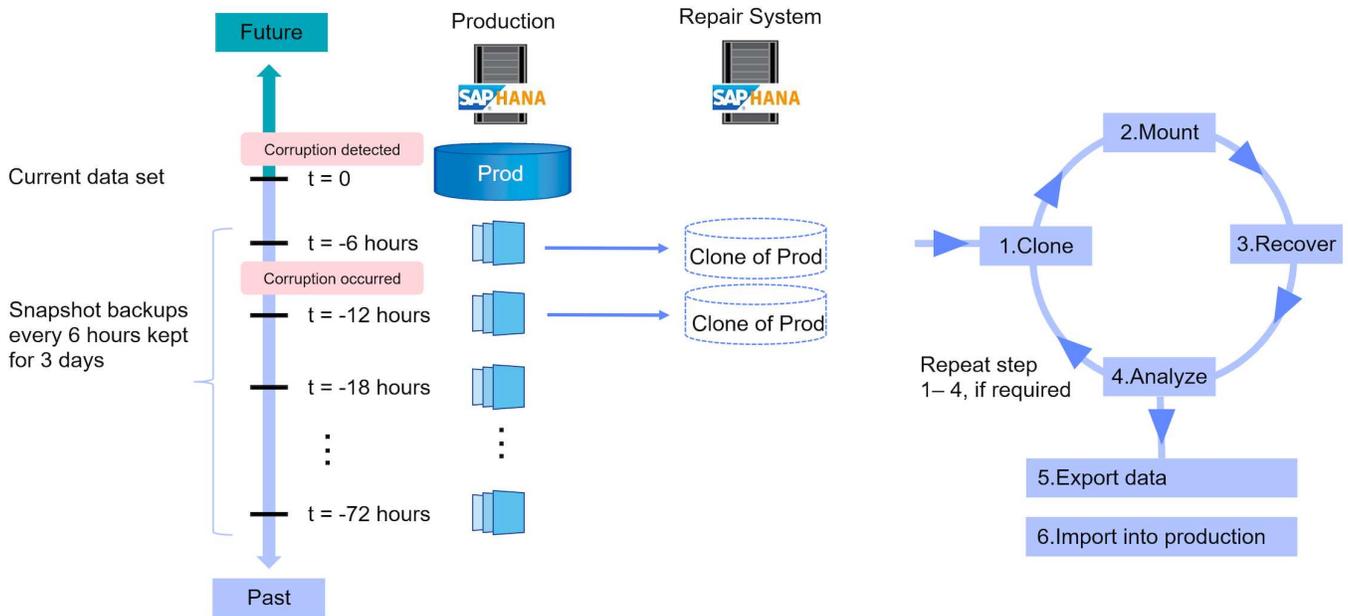
Logical corruption can be caused by software errors, human errors, or sabotage. Unfortunately, logical corruption often cannot be addressed with standard high-availability and disaster recovery solutions. As a result, depending on the layer, application, file system, or storage where the logical corruption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



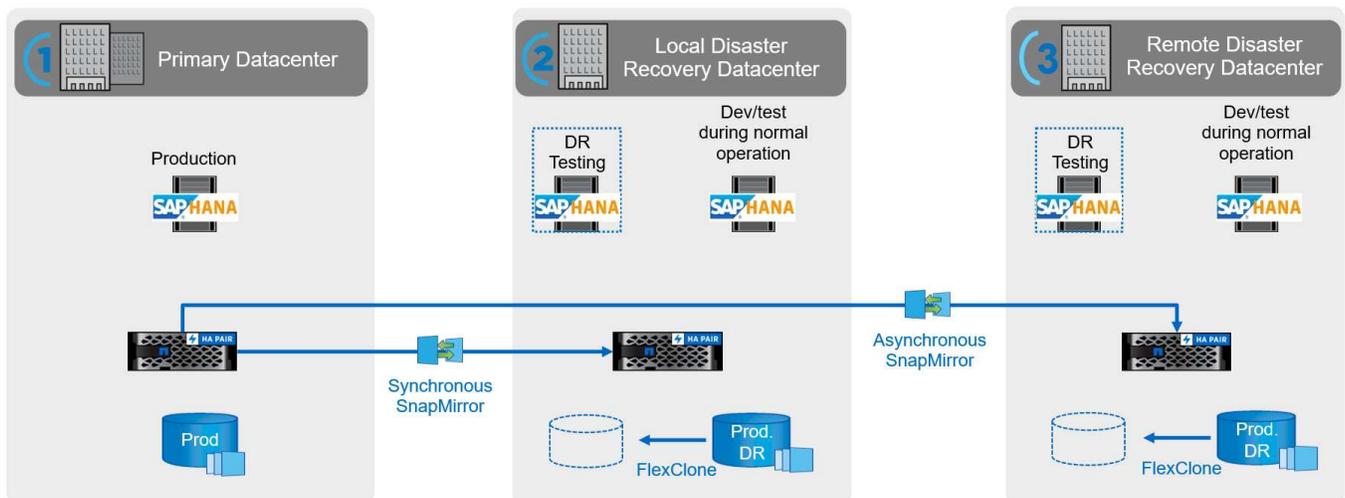
The workflow of the repair system creation is described in the section “SAP system clone with SnapCenter.”

### 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 \(netapp.com\)](https://netapp.com/tr-4646)

## 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

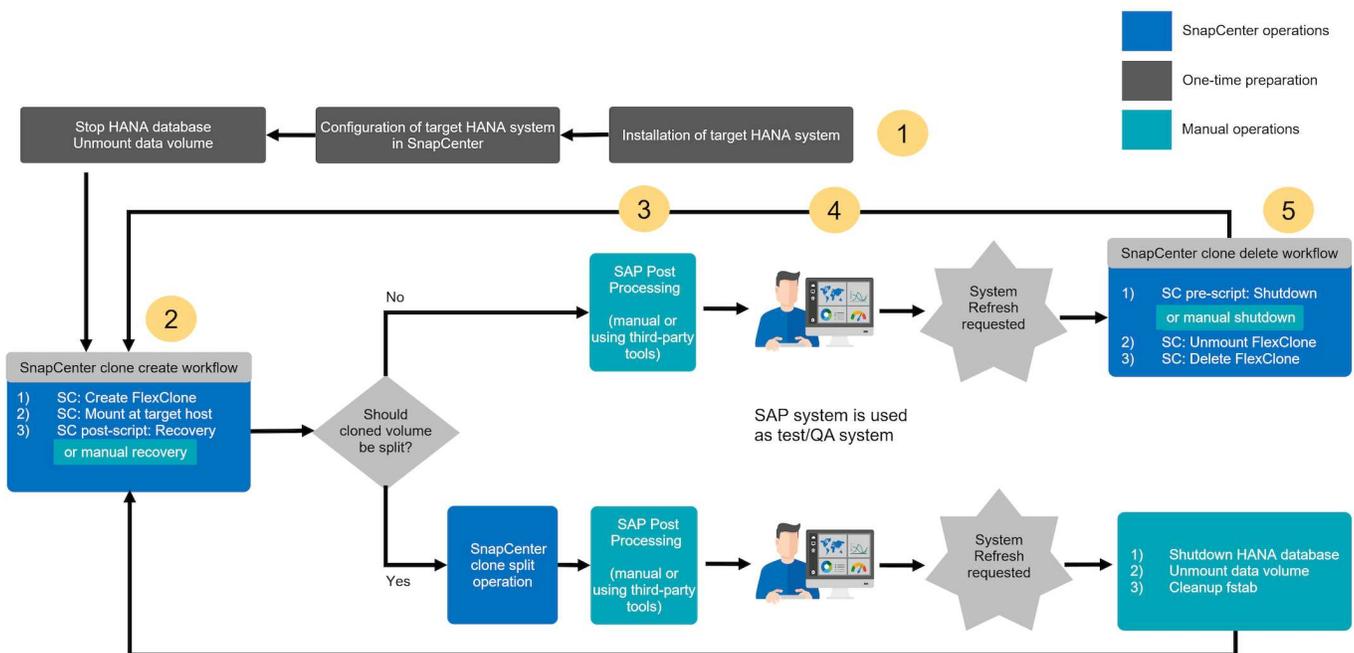
SnapCenter provides workflows that allow you to manage clones of data sets from any existing Snapshot backup. This cloned data set, a FlexClone volume, can be used to rapidly provision a HANA data volume from a source system and attach it to a target system. It is therefore a perfect fit for executing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



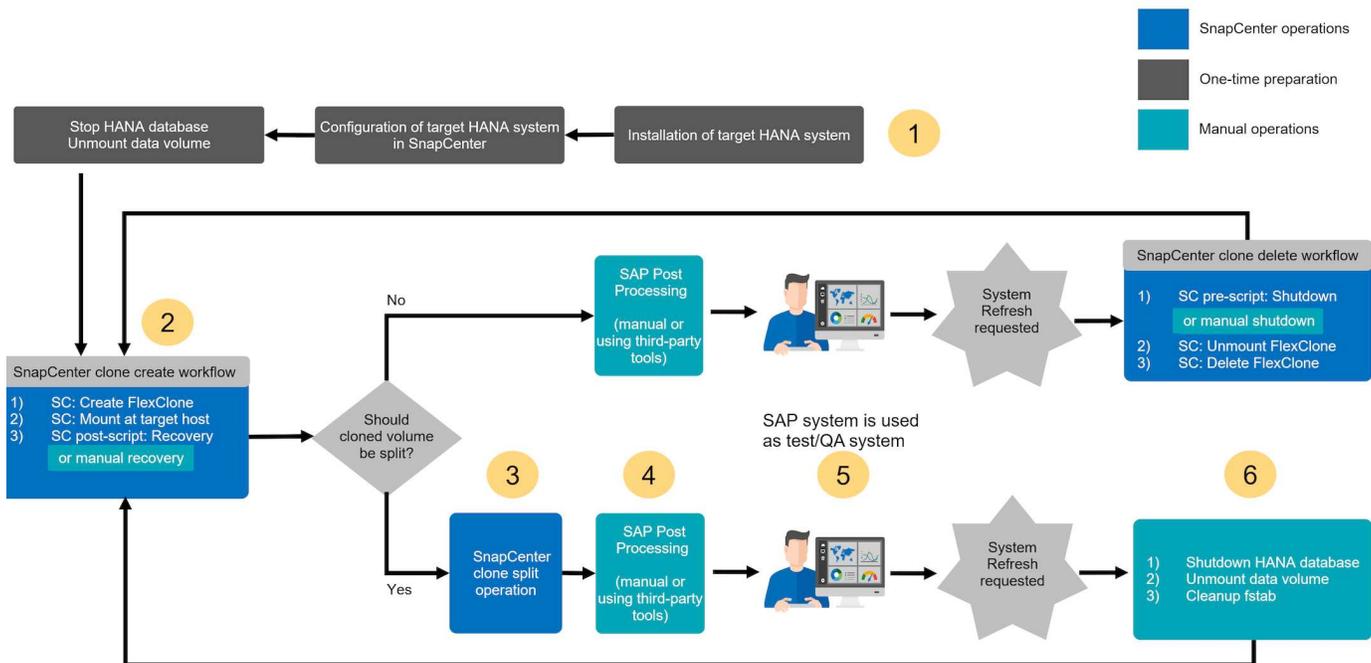
With SnapCenter and Azure NetApp Files, the clone split operation is not available, since Azure NetApp Files always splits the clone after creation.

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

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. The FlexClone volume is split using the SnapCenter clone split workflow.
4. SAP post processing (manual or with a third-party tool)
5. 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.

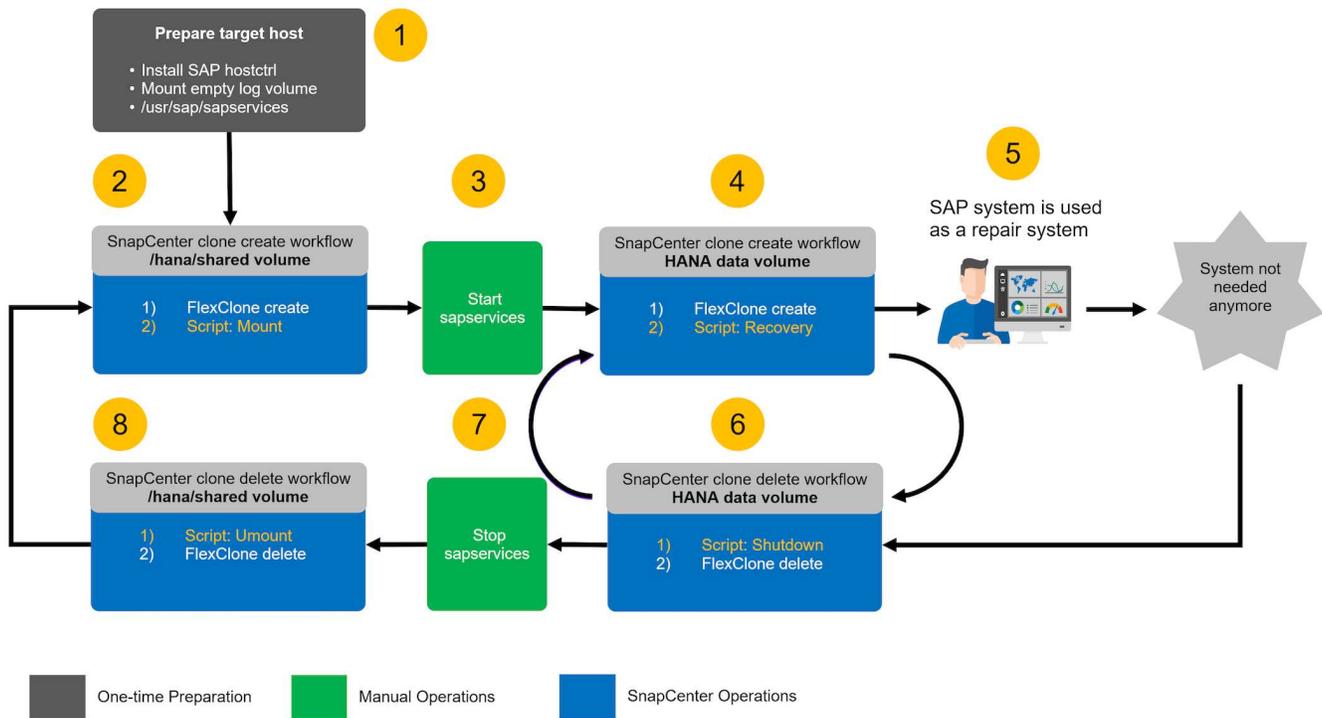
1. Prepare the target host.
2. SnapCenter clone create workflow for the SAP HANA shared volume.
3. Start SAP HANA services.
4. SnapCenter clone create workflow for the SAP HANA data volume including database recovery.
5. The SAP HANA system can now be used as a repair system.

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

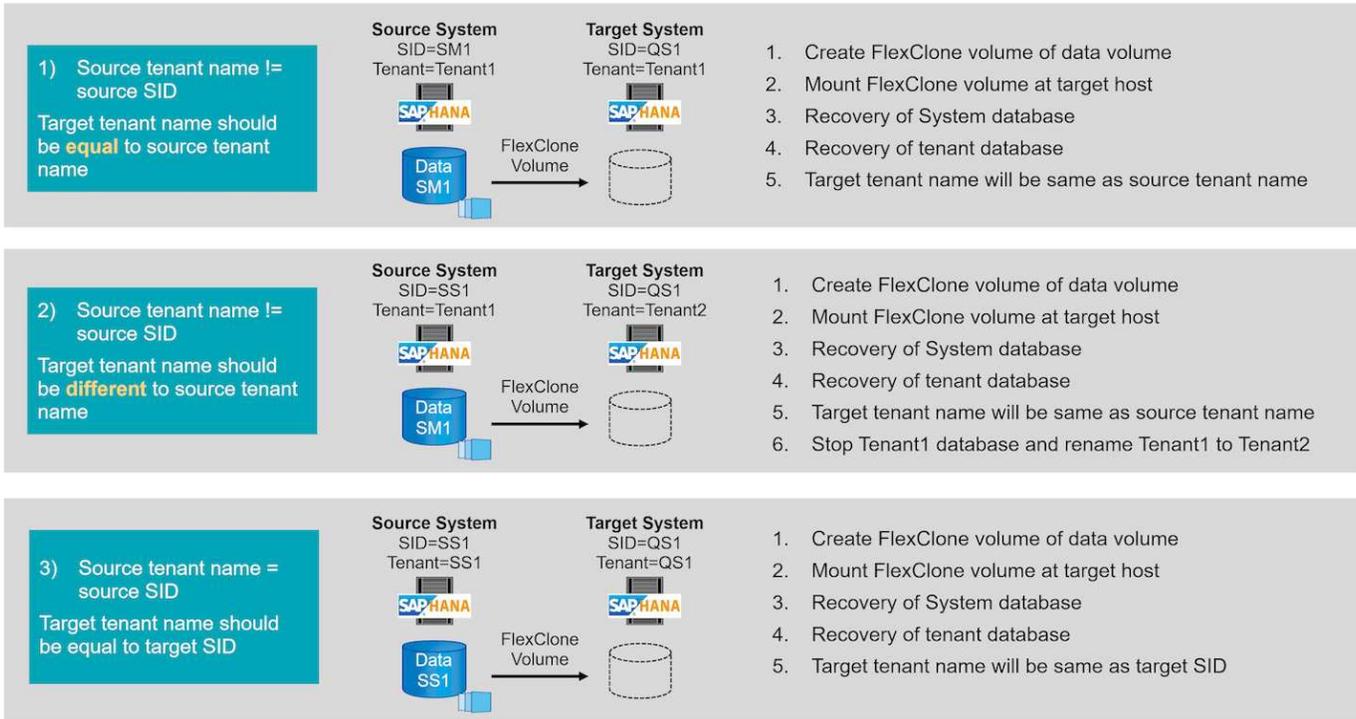
The steps required to perform an SAP HANA system refresh depend on the source system tenant configuration and the required tenant name at the target system, as shown in the figure below.

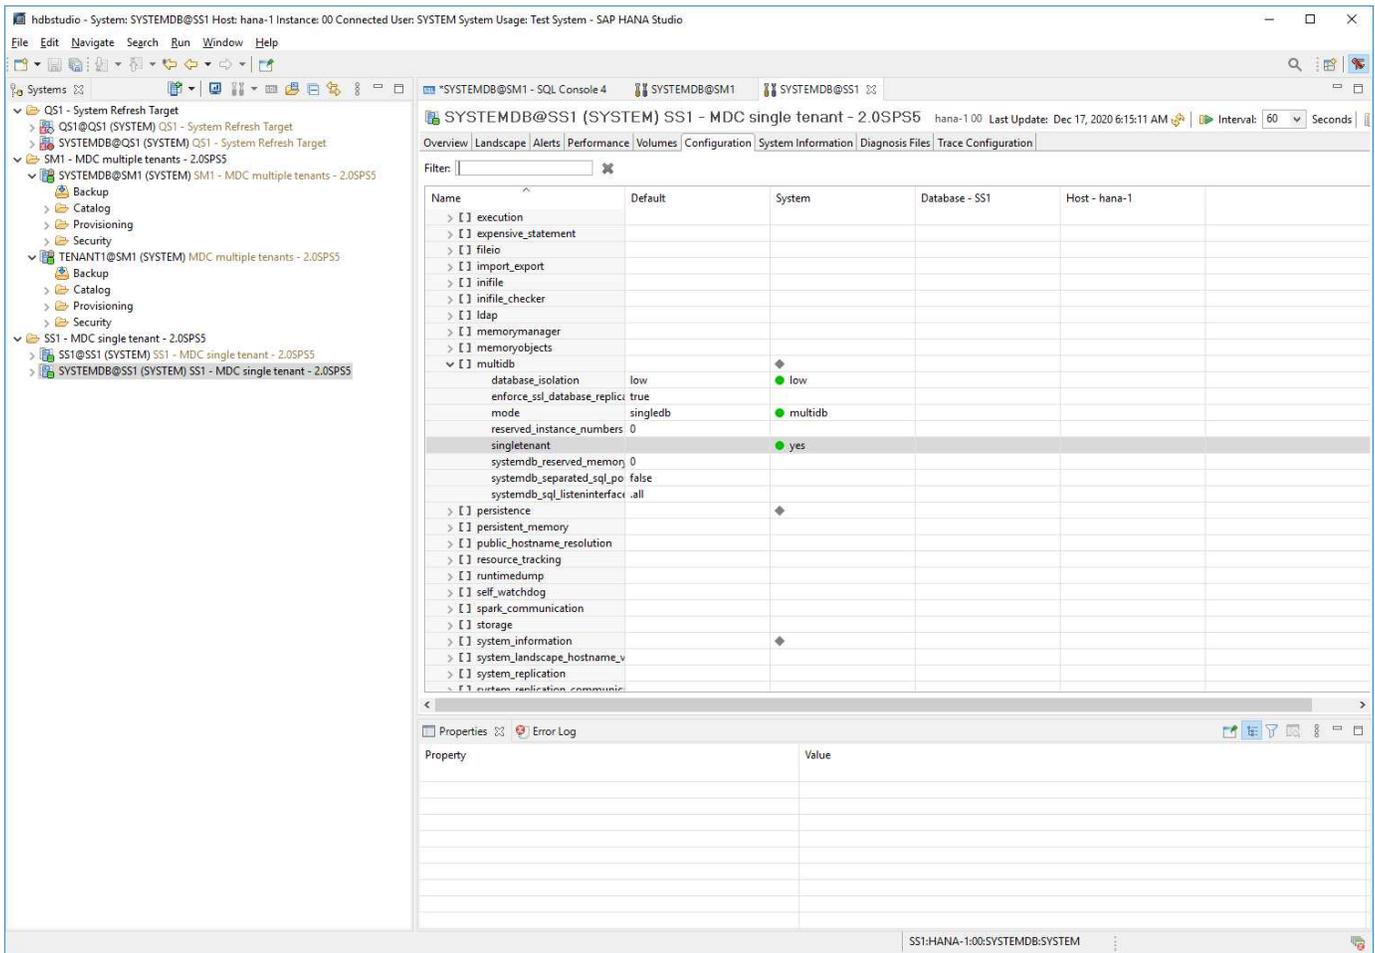
Since the tenant name is configured in the system database, the tenant name of the source system is also available at the target system after the recovery of the system database. Therefore, the tenant at the target system can only be recovered with the same name as the source tenant as shown in option 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

```

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.

The screenshot shows a SQL query in a client window. The query is a SELECT statement that uses a CASE WHEN clause to map database names to their respective DBID values. The result set below the query shows three rows: TENANT1 with DBID 4, SYSTEMDB with DBID 1, and TENANT2 with DBID 3.

```

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	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`

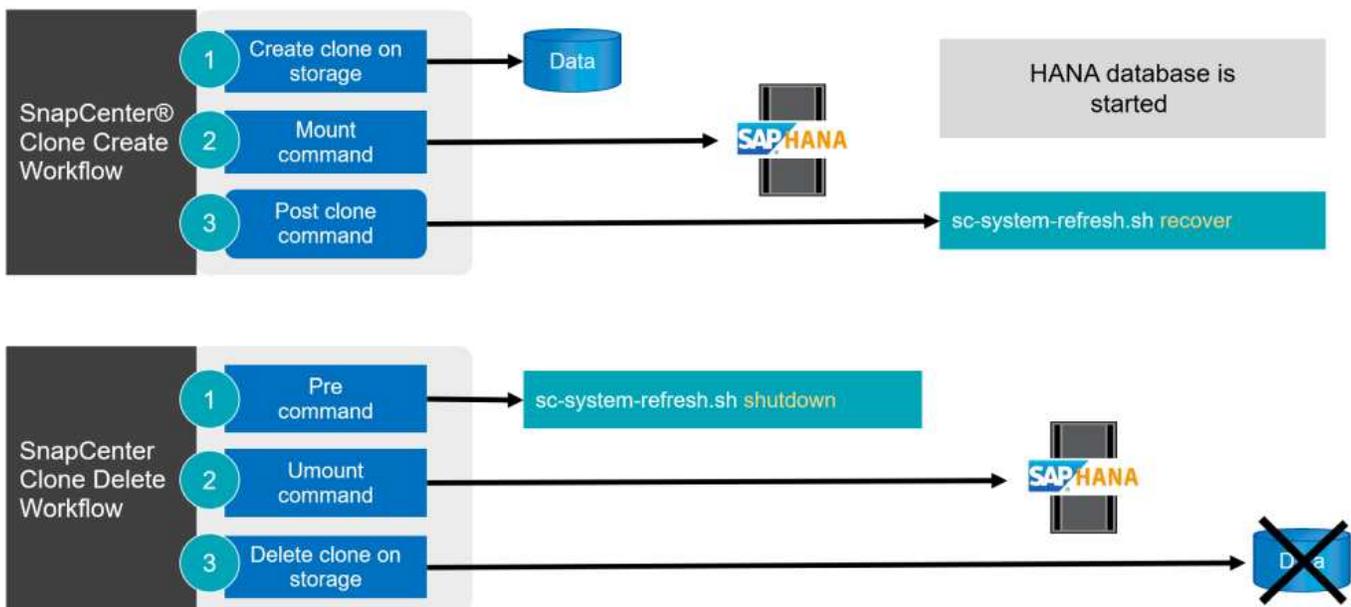
The example script `sc-system-refresh.sh` is used to execute recovery and shutdown operations. The script is called with specific command-line options within the SnapCenter workflows clone create and clone delete, as shown in the figure below.

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.

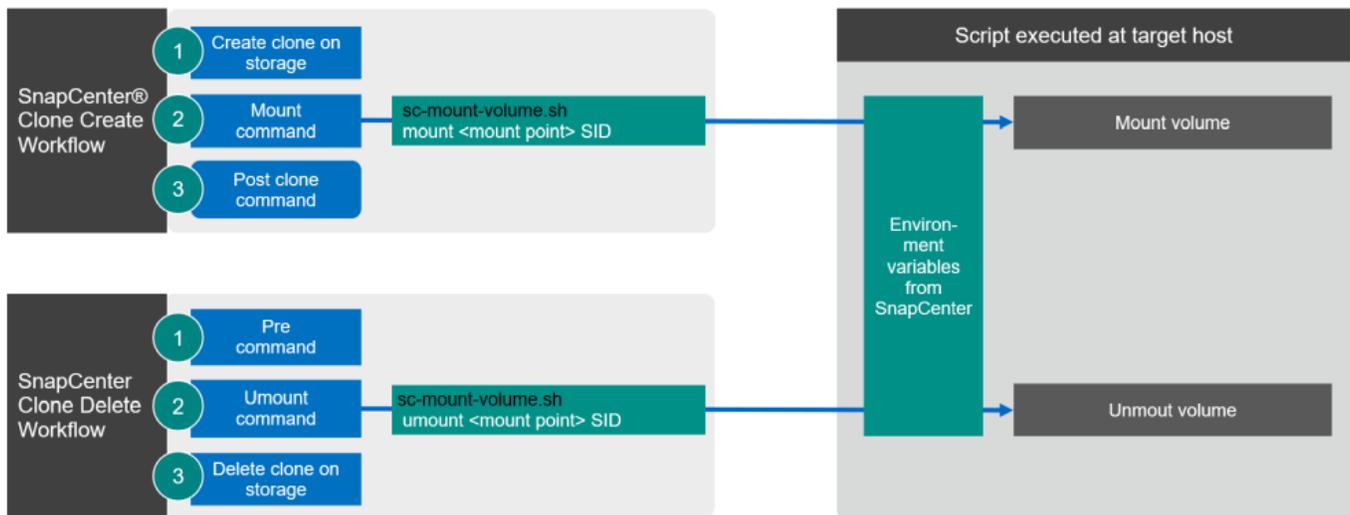
SAP HANA system	Tenant configuration at source system	Resulting tenant configuration at target system
MDC single tenant	Source tenant name equal to source SID	Target tenant name is equal to target SID
MDC single tenant	Source tenant name not equal to source SID	Target tenant name is equal to source tenant 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:

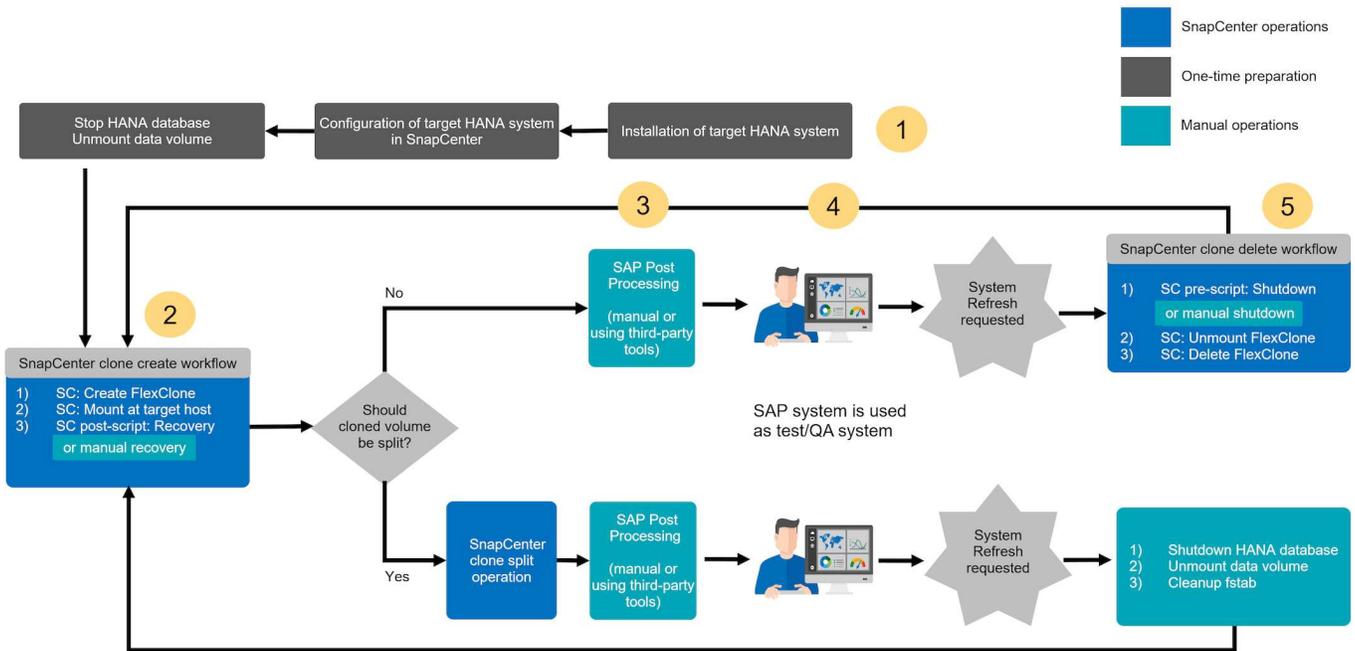
```

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.

Source system	Source system configuration	SnapCenter and SAP HANA operations
	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.

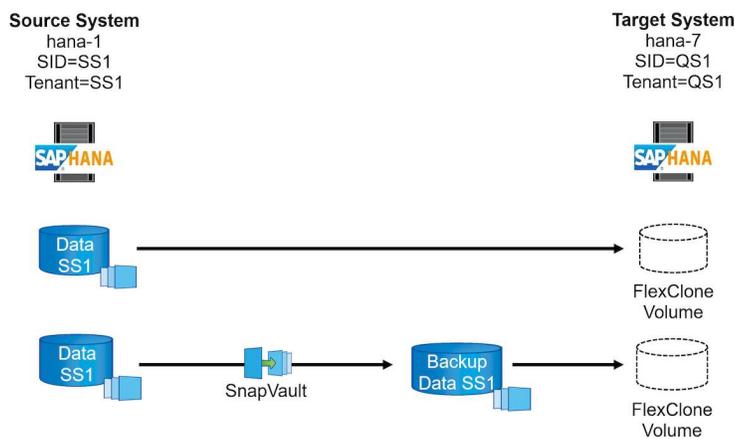
- Cloning from primary storage or off-site backup storage; tenant name is equal to the SID.
  - Source SAP HANA system: SS1 with Tenant SS1
  - Target SAP HANA system: QS1 with Tenant QS1
- Cloning from primary storage; multiple tenants.
  - Source SAP HANA system: SM1 with Tenant1 and Tenant2
  - Target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

The following software versions were used:

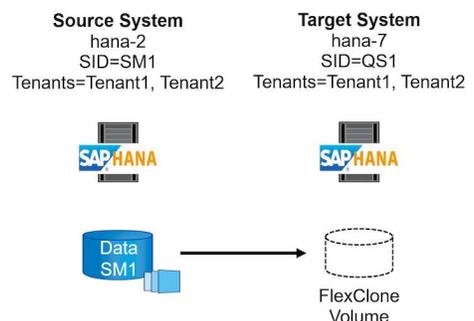
- SnapCenter 5.0
- SAP HANA systems: HANA 2.0 SPS7 rev.73
- SLES 15
- ONTAP 9.14P1

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`.

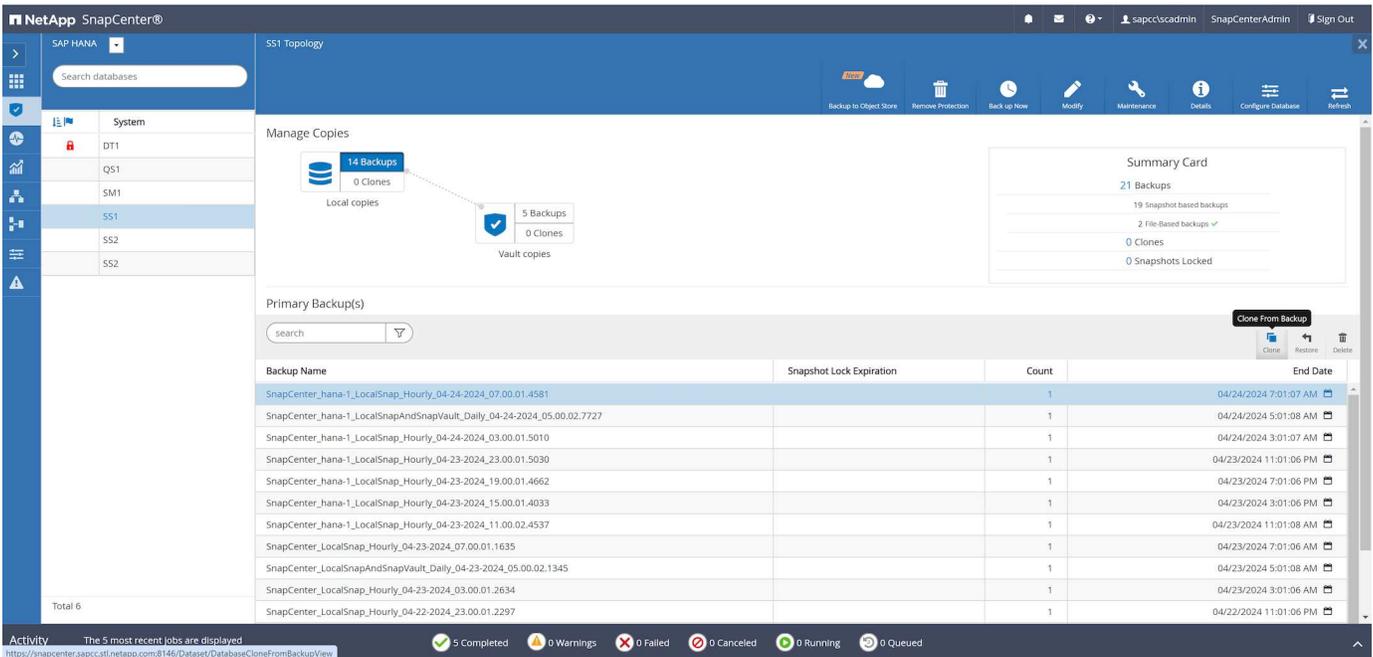


The workflow consists of the following steps:

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.



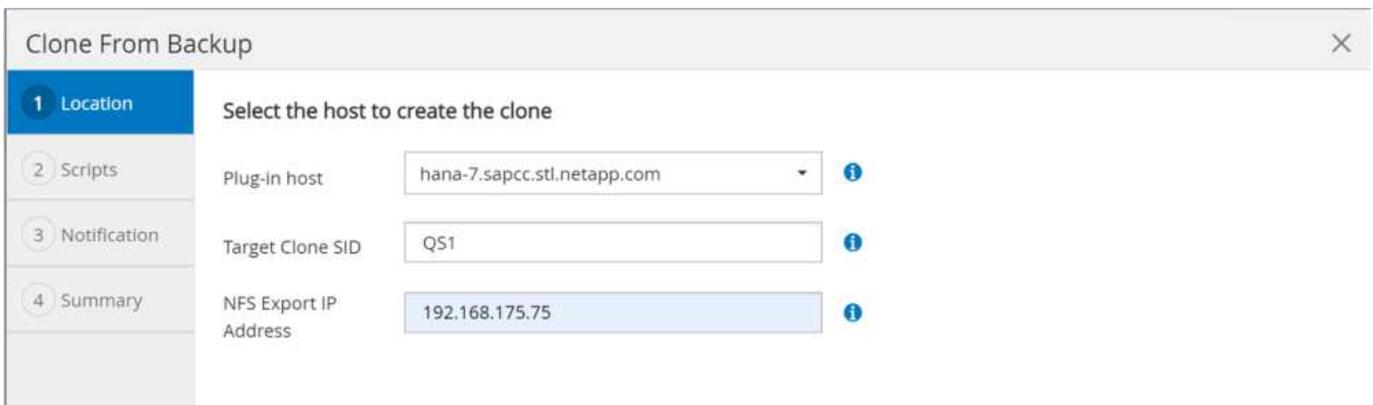
2. Select the host where the target system QS1 is installed. Enter QS1 as the target SID. The NFS export IP address must be the storage network interface of the target host.



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.



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** **Select the host to create the clone**

2 Settings Plug-in host  ⓘ

3 Scripts Target Clone SID  ⓘ

4 Notification NFS Export IP Address  ⓘ

5 Summary

Clone From Backup ✕

**1 Location** **LUN Map Settings**

**2 Settings** Igroup protocol

3 Scripts

4 Notification

5 Summary

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** **Select the host to create the clone**

2 Scripts Plug-in host  ⓘ

3 Notification Target Clone SID  ⓘ

4 Summary NFS Export IP Address  ⓘ

Capacity Pool Max. Throughput (MIB/s)  ⓘ

3. 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
 The following commands will run on the Plug-in Host: `hana-7.sapcc.stl.netapp.com`

2 Scripts
 Enter optional commands to run before performing a clone operation i

3 Notification
 Pre clone command

4 Summary
 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
```



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.

4. 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.

Job Details
✕

Clone from backup 'SnapCenter\_hana-1\_LocalSnap\_Hourly\_04-25-2024\_11.00.01.5630'

- ✓ ▾ Clone from backup 'SnapCenter\_hana-1\_LocalSnap\_Hourly\_04-25-2024\_11.00.01.5630'
- ✓ ▾ hana-7.sapcc.stl.netapp.com
- ✓ ▾ Clone
- ✓ ▶ Application Pre Clone
- ✓ ▶ Storage Clone
- ✓ ▶ Unmount Filesystem
- ✓ ▶ Mount Filesystem
- ✓ ▶ Application Post Clone
- ✓ ▶ Post Clone Create Commands
- ✓ ▶ Register Clone Metadata
- ✓ ▶ Clean-up Snapshot entries on Server
- ✓ ▶ Application Clean-Up
- ✓ ▶ Data Collection
- ✓ ▶ Agent Finalize Workflow

**i** Task Name: Clone Start Time: 04/25/2024 11:22:40 AM End Time: 04/25/2024 11:25:29 AM

View Logs
Cancel Job
Close

5. 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 displays the NetApp SnapCenter interface for managing SAP HANA databases. The top navigation bar shows the user is logged in as 'sapcc/scadmin'. The main content area is titled 'SS1 Topology' and features a 'Manage Copies' section. This section shows a hierarchy of copies: 'Local copies' with 14 Backups and 1 Clone, and 'Vault copies' with 5 Backups and 0 Clones. A 'Summary Card' on the right provides a quick overview: 21 Backups, 19 Snapshot based backups, 2 File Based backups, 1 Clone, and 0 Snapshots Locked. Below this, the 'Primary Clone(s)' section contains a table with one entry:

Clone SID	Clone Host	Clone Name	Start Date	End date
QS1	hana-7.sapcc.stl.netapp.com	hana-1_sapcc_stl_netapp_com_hana_MDC_SS1_clone_102162_MDC_SS1_04-22-2024_09.54.34	04/24/2024 9:47:10 AM	04/24/2024 9:48:00 AM

The bottom status bar shows activity: 'The 5 most recent jobs are displayed' with 1 Completed, 2 Warnings, 0 Failed, 0 Canceled, 2 Running, and 0 Queued jobs.

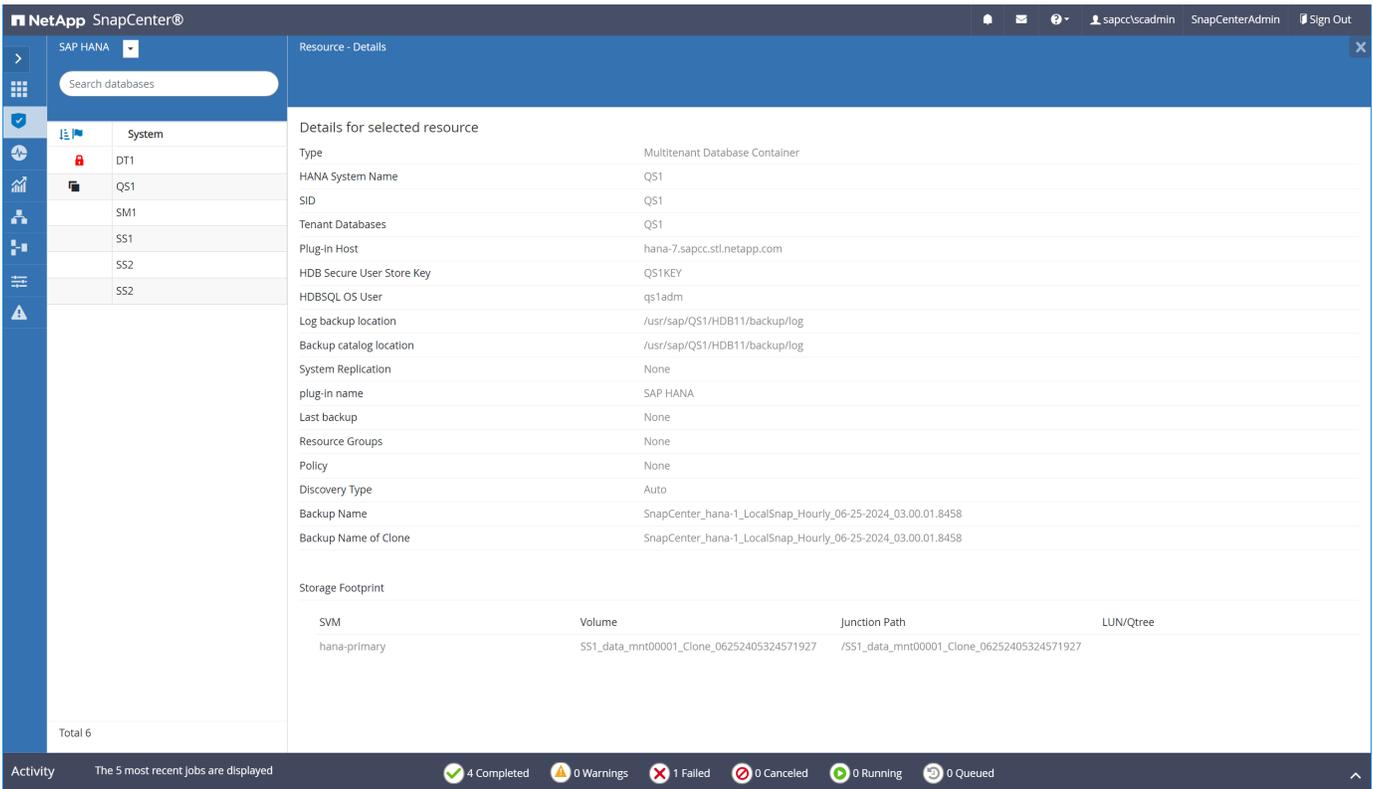
7. The SAP HANA database is now running.
8. If you want to protect the target SAP HANA system, you need to run the auto discovery by clicking on the target system resource.

The 'Configure Database' dialog box is shown with the following configuration:

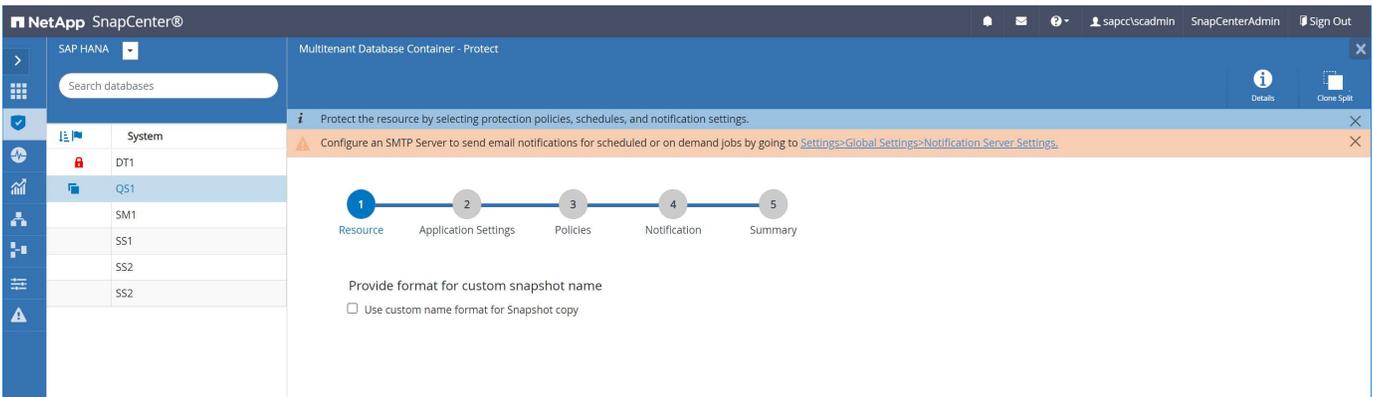
- Plug-in host: hana-7.sapcc.stl.netapp.com
- HDBSQL OS User: qs1adm
- HDB Secure User Store Key: QS1KEY

Buttons for 'Cancel' and 'OK' are visible at the bottom right of the dialog.

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



By clicking on the resource again, data protection can be configured for the refreshed QS1 system.

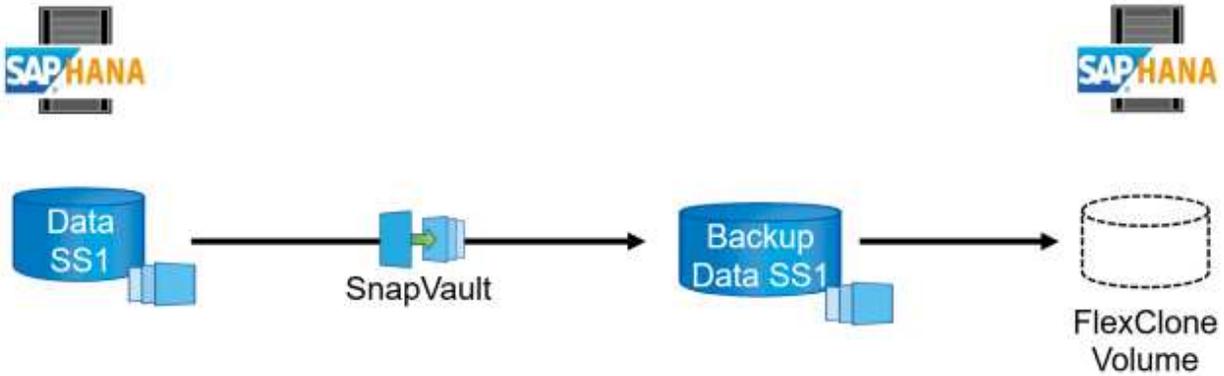


## Cloning from off-site backup storage

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

**Target System**  
hana-7  
SID=QS1  
Tenant=QS1



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. On the left, a sidebar lists systems: QS1, SM1, SS1, SS2, and SS3. The main area is titled 'SS1 Topology' and 'Manage Copies'. It shows 'Local copies' with 14 Backups and 0 Clones, and 'Vault copies' with 9 Backups and 0 Clones. A 'Summary Card' on the right indicates 25 Backups, including 23 Snapshot based backups and 2 File Based backups. Below this is a table of 'Secondary Vault Backup(s)'. The table has columns for 'Backup Name', 'Count', and 'End Date'. A 'Clone From Backup' button is located at the top right of the table area.

Backup Name	Count	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.

x
Clone From Backup

1 Location

2 Scripts

3 Notification

4 Summary

Select the host to create the clone

Plug-in host  i

Target Clone SID  i

NFS Export IP Address  i

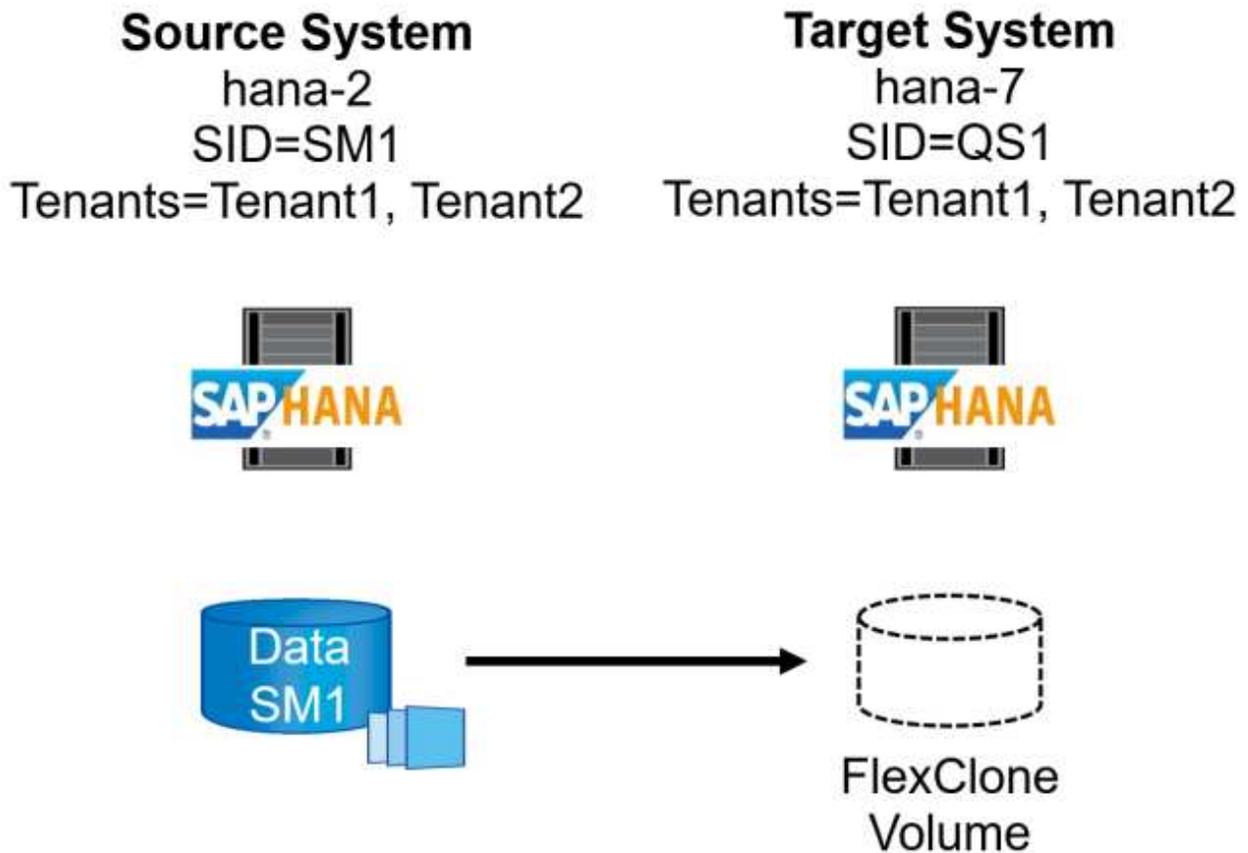
Secondary storage location : Snap Vault / Snap Mirror

Source Volume	Destination Volume
hana-primary.sapcc.stl.netapp.com:SS1_data_mnt0000 1	<input style="width: 150px;" type="text" value="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: Successfully 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 successfully 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 successfully finished.
20240430070350###hana-7###sc-system-refresh.sh: Status: GREEN
20240430070350###hana-7###sc-system-refresh.sh: *****
Finished script: recovery operation *****

```

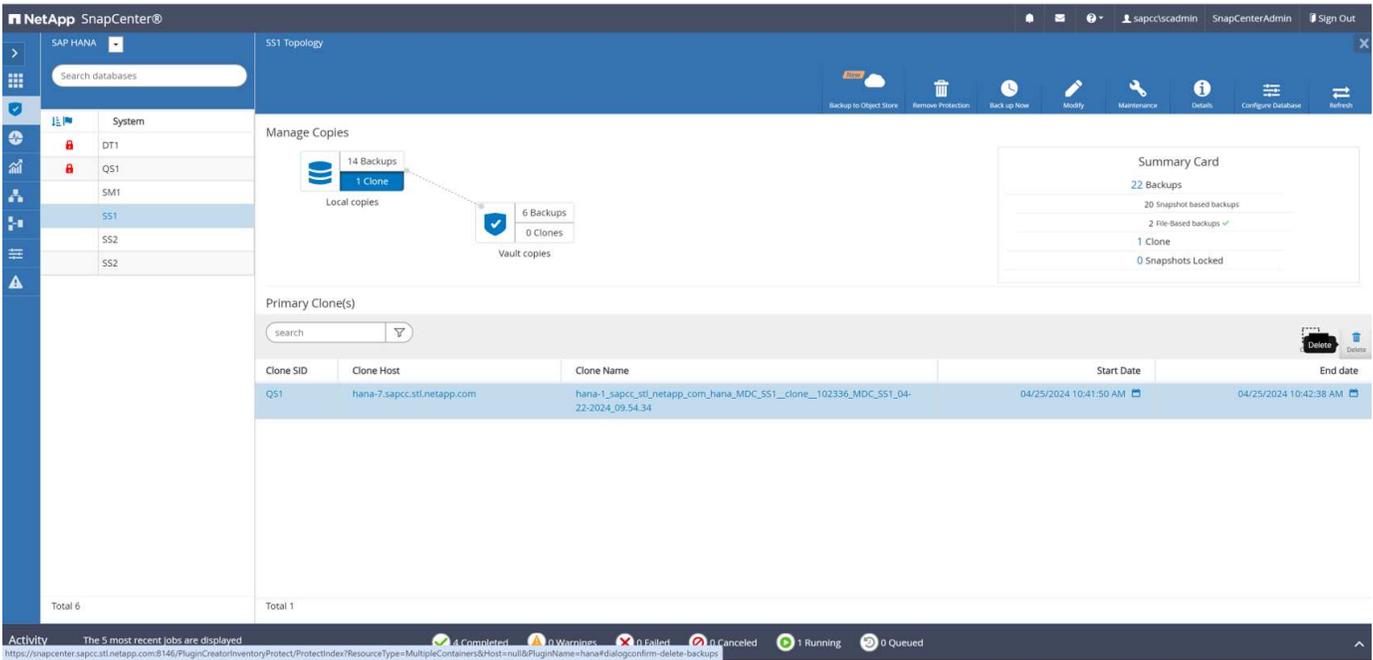
## Clone delete operation

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

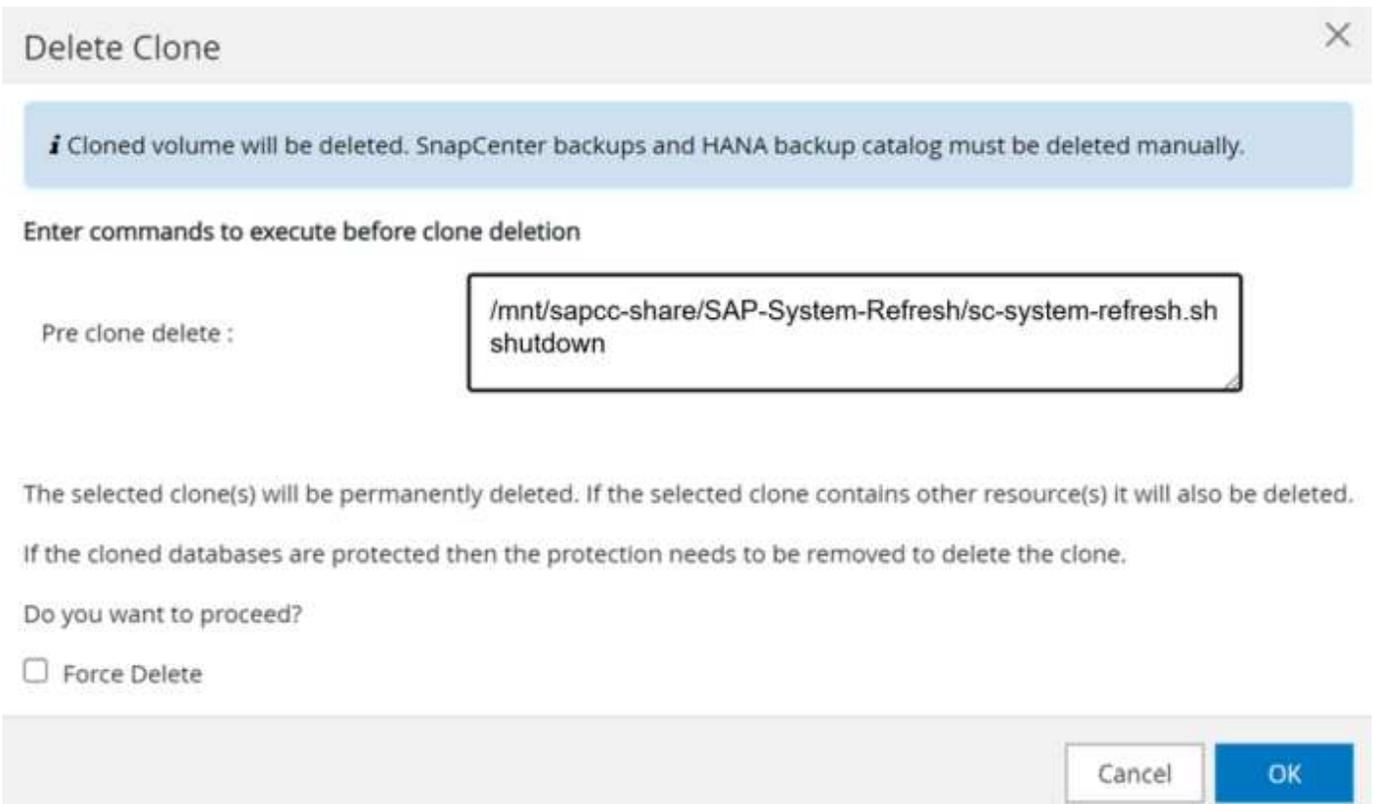
If the target SAP HANA system has been protected in SnapCenter, the protection must be removed first. Within the topology view of the target system, click Remove Protection.

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

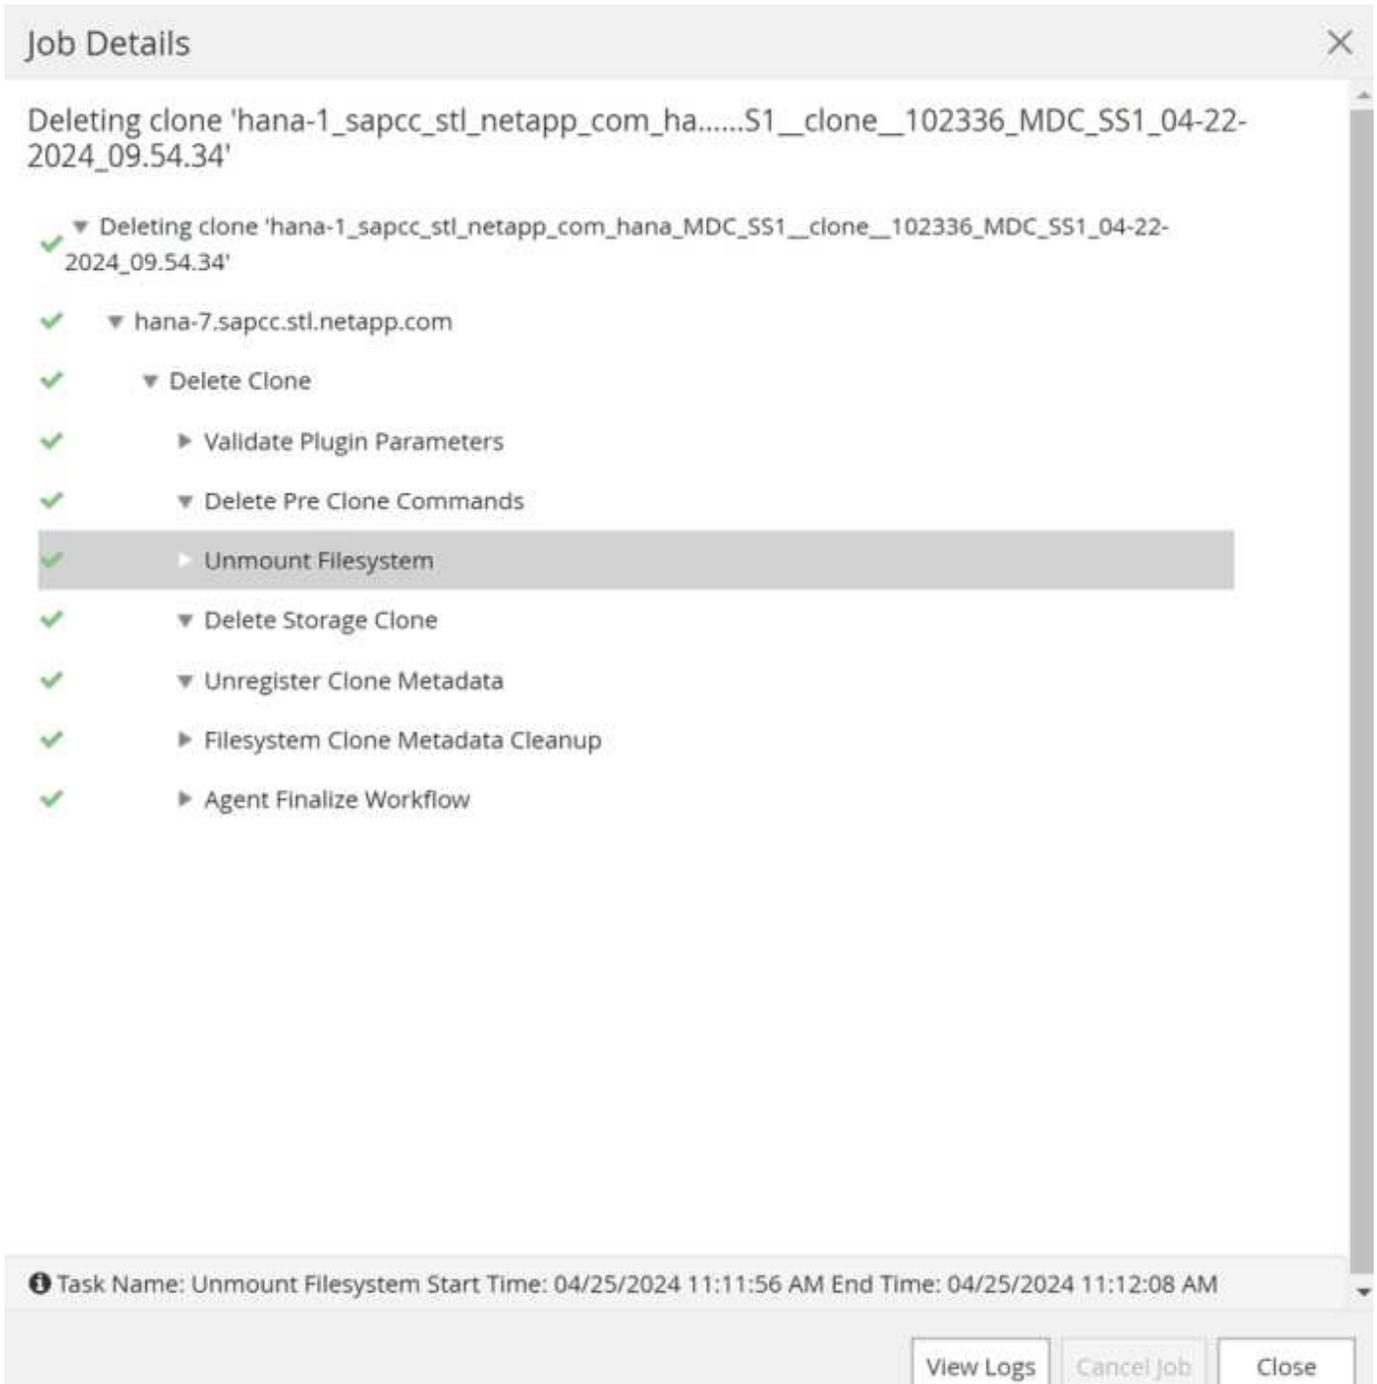
1. Select the clone within the topology view of the source system and click Delete.



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



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



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.

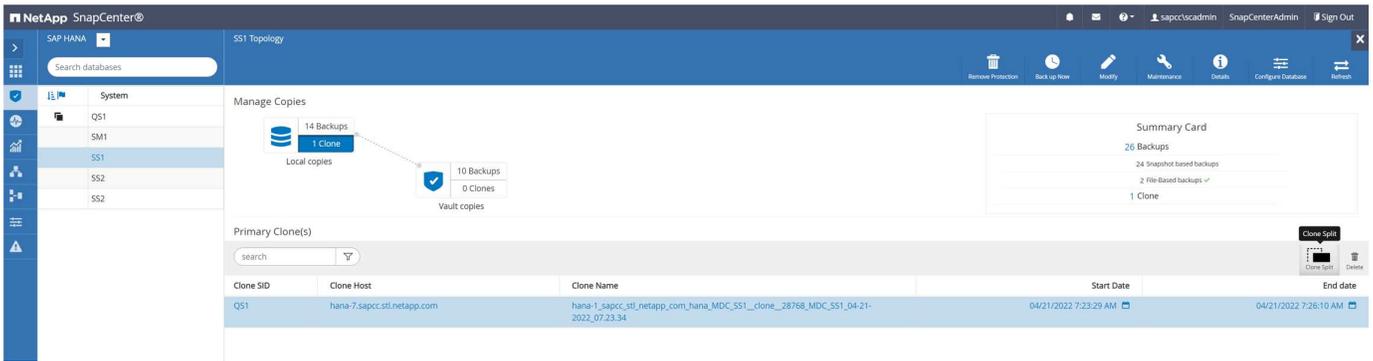


The clone split operation does not block the use of the cloned volume and can therefore be executed at any time while the SAP HANA database is in use.

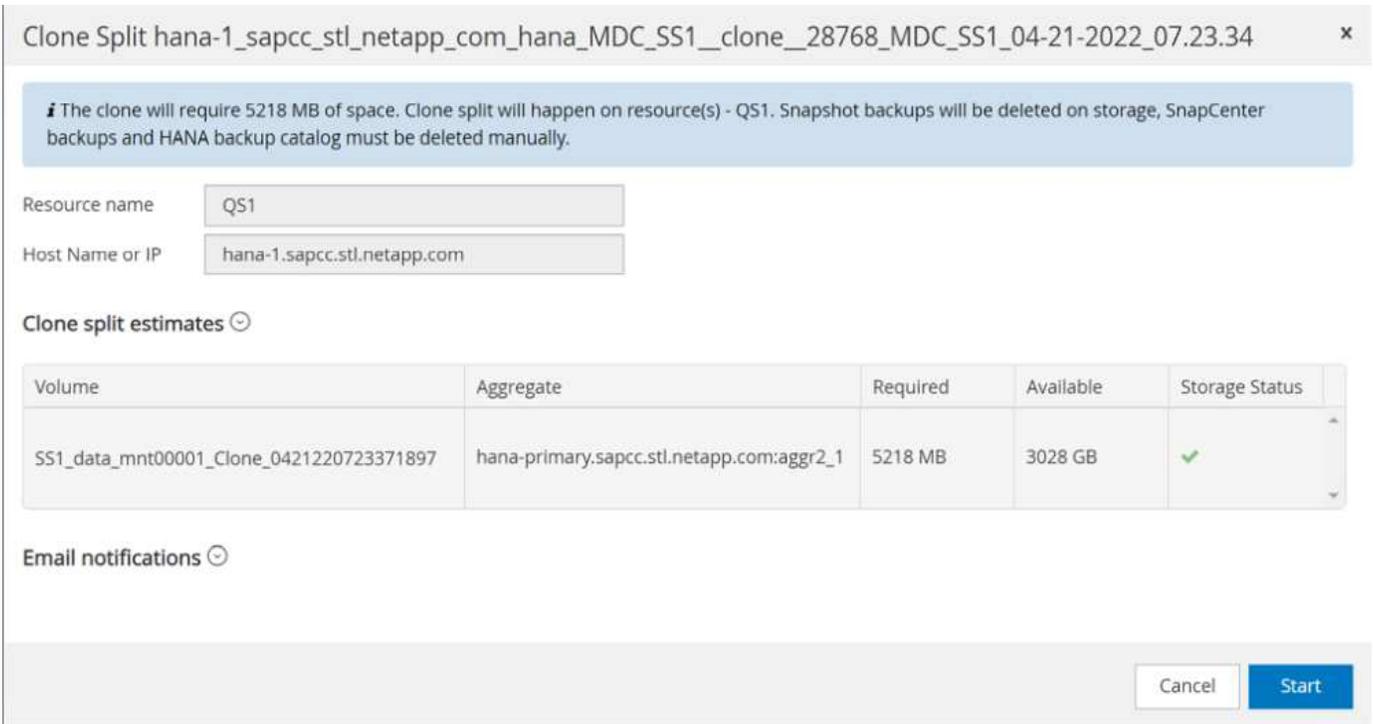


With Azure NetApp Files, the clone split operation is not available, since Azure NetApp Files always splits the clone after creation.

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

**i** 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.

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

Manage Copies

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

Primary Backup(s)

Backup Name	Count	IF	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 finished."
```

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
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 finished."  
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.

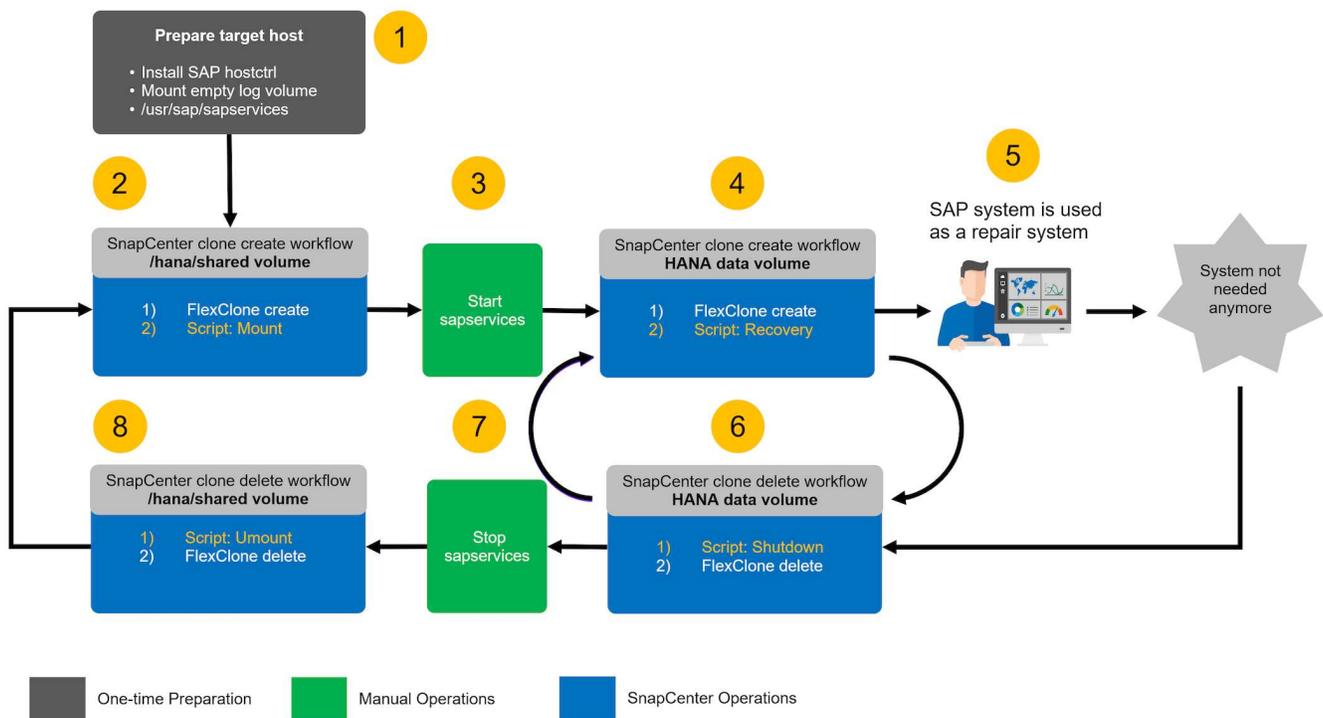
1. Prepare the target host.
2. SnapCenter clone create workflow for the SAP HANA shared volume.
3. Start SAP HANA services.
4. SnapCenter clone create workflow for the SAP HANA data volume including database recovery.
5. The SAP HANA system can now be used as a repair system.



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.

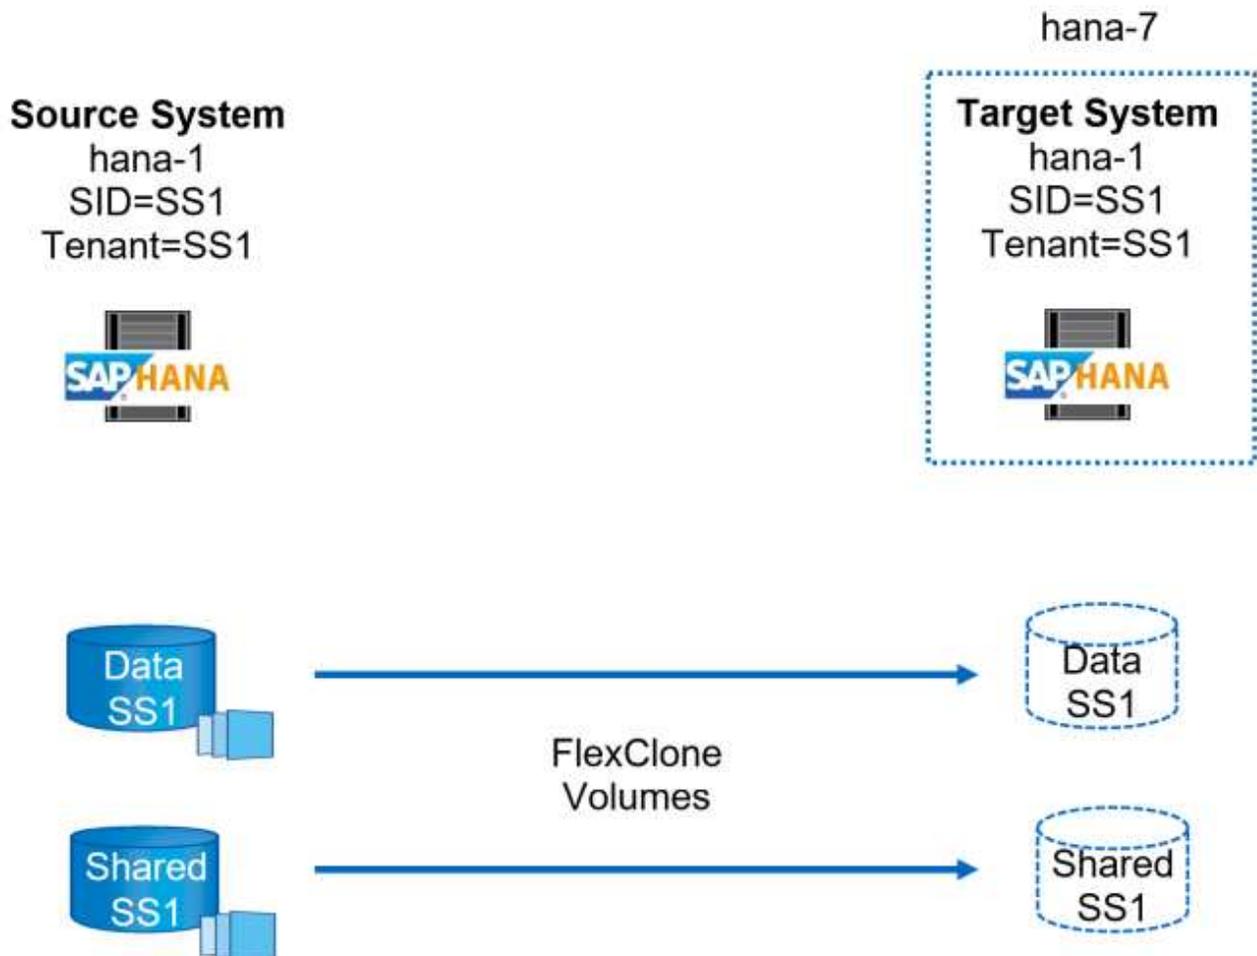
## Lab setup

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,rsize=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,rsize=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	IF	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.](#) ✕

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

Job Details x

Clone from backup 'SnapCenter\_LocalSnap\_Hourly\_05-13-2022\_05.04.01.8012'

- ✓ ▾ Clone from backup 'SnapCenter\_LocalSnap\_Hourly\_05-13-2022\_05.04.01.8012'
- ✓ ▾ hana-7.sapcc.stl.netapp.com
  - ✓ ▾ Clone
    - ▶ Storage Clone
    - ▶ Register Clone Metadata
    - ▶ Data Collection
    - ▶ Agent Finalize Workflow

**i** Task Name: Clone Start Time: 05/13/2022 5:14:02 AM End Time: 05/13/2022 5:14:16 AM

[View Logs](#) [Cancel Job](#) [Close](#)

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,rsiz=1048576,wsiz=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,rsiz=1048576,wsiz=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.

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 interface for managing SAP HANA backups. The left sidebar lists systems: QS1, SM1, SS1 (selected), SS2, and SS2. The main area displays the 'SS1 Topology' with a 'Manage Copies' section showing 15 Backups (Local copies) and 11 Backups (Vault copies). A 'Summary Card' on the right indicates 28 Backups, including 26 Snapshot based backups and 2 File Based backups. Below this is a table of 'Primary Backup(s)' with columns for Backup Name, Count, IF, and End Date.

Backup Name	Count	IF	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

Select the host to create the clone

2

Scripts

Plug-in host

3

Notification

Target Clone SID

4

Summary

NFS Export IP Address

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

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

2

Scripts

Enter optional commands to run before performing a clone operation i

Pre clone command

3

Notification

Enter optional commands to run after performing a clone operation i

Post clone command

4

Summary

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

**i** 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.

```

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

## Automating SAP system copy operations with Libelle SystemCopy

### TR-4929: Automating SAP system copy operations with Libelle SystemCopy

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

Authors:

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.

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 AnyDBs and SAP HANA databases. In addition, NetApp integrates into SAP lifecycle management tools, combining efficient application-integrated data protection with the flexible provisioning of SAP test systems.

While these NetApp solutions solve the issue of efficiently managing enormous amounts of data even for the largest databases, full end-to-end SAP system- copy and refresh operations have to include pre- and post-copy activities to completely change the identity of the source SAP system to the target system. SAP describes the required activities in their [SAP homogenous system copy guide](#). To further reduce the number of manual processes and to improve the quality and stability of a SAP system copy process, our partner [Libelle](#) has developed the [Libelle SystemCopy \(LSC\)](#) tool. We have jointly worked with Libelle to integrate the NetApp solutions for SAP system copies into LSC to provide [full end-to-end automated system copies in record time](#).

### **Application-integrated Snapshot copy operation**

The ability to create application-consistent NetApp Snapshot copies on the storage layer is the foundation for the system copy and system clone operations described in this document. Storage-based Snapshot copies are created with the NetApp SnapCenter Plug-In for SAP HANA or SAP Any DBs on native NetApp ONTAP systems or by using the [Microsoft Azure Application Consistent Snapshot tool \(AzAcSnap\)](#) and interfaces provided by the SAP HANA and Oracle database running in Microsoft Azure. When using SAP HANA, SnapCenter and AzAcSnap register Snapshot copies 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 copies can be replicated on the storage layer to an off-site backup site or a disaster recovery site controlled by SnapCenter on-premises. Replication is based on block changes and is therefore space and bandwidth efficient. The same technology is available for SAP HANA and Oracle systems running in Azure with Azure NetApp Files by using the Cross Region Replication (CRR) feature to efficiently replicate Azure NetApp Files volumes between Azure regions.

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

NetApp technology and software integration allows you to use any Snapshot copy of a source system for an SAP system copy or clone operation. This Snapshot copy 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 (such as Azure NetApp Files backup in Azure), or the storage at the disaster recovery site (Azure NetApp Files CRR target volumes). 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.

## 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 and other third-party vendors (for example, Libelle) to support different scenarios and levels of automation.

NetApp SnapCenter with the plug-in for SAP HANA and SAP AnyDBs or AzAcSnap on Azure is used to provision the required storage- volume clones based on an application-consistent Snapshot copy and to execute all required host and database operations up to a started SAP 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 provisioning or refresh of SAP test systems can be performed by using Libelle SystemCopy (LSC) automation. The integration of SnapCenter or AzAcSnap into LSC is described in more detail in this document.

## Libelle SystemCopy

Libelle SystemCopy is a framework-based software solution to create fully automated system and landscape copies. With the proverbial touch of a button, QA and test systems can be updated with fresh production data. Libelle SystemCopy supports all conventional databases and operating systems, provides its own copy mechanisms for all platforms but, at the same time, integrates backup/restore procedures or storage tools such as NetApp Snapshot copies and NetApp FlexClone volumes. The activities that are necessary during a system copy are controlled from outside the SAP ABAP stack. In this way, no transports or other changes are required in the SAP applications. Generally, all steps necessary to successfully complete a system copy procedure can be categorized into four steps:

- **Check phase.** Check the involved system environments.
- **Pre phase.** Prepare the target system for a system copy.
- **Copy phase.** Provide a copy of the actual production database to the target system from the source.
- **Post phase.** All tasks after the copy to complete the homogeneous system copy procedure and to provide an updated target system.

During the copy phase, NetApp Snapshot and FlexClone functionality is used to minimize the time needed to a couple of minutes even for the largest databases.

For the Check, Pre, and Post phases, LSC comes with 450+ preconfigured tasks covering 95% of typical refresh operations. As a result, LSC embraces automation following SAP standards. Due to the software-defined nature of LSC, system refresh processes can be easily adjusted and enhanced to meet the specific needs of customer SAP environments.

## Use cases for SAP system refresh and cloning

There are multiple scenarios in which data from a source system must be made available to a target system:

- Regular refresh of quality assurance and test and training systems
- Creating break fix or repair system environments to address logical corruption
- Disaster recovery test scenarios

Although repair systems and disaster recovery test systems are typically provided using SAP system clones (which don't require extensive post-processing operations) for refreshed test and training systems, these post-processing steps must be applied to enable coexistence with the source system. Therefore, this document

focuses on SAP system refresh scenarios. More details about the different use cases can be found in the technical report [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#).

The remainder of this document is separated into two parts. The first part describes the integration of NetApp SnapCenter with Libelle SystemCopy for SAP HANA and SAP AnyDBs systems running on NetApp ONTAP systems on-premises. The second part describes the integration of AzAcSnap with LSC for SAP HANA systems running in Microsoft Azure with Azure NetApp Files provided. Although the underlying ONTAP technology is identical, Azure NetApp Files provides different interfaces and tool integration (for example, AzAcSnap) compared to native ONTAP installation.

## SAP HANA system refresh with LSC and SnapCenter

This section describes how to integrate LSC with NetApp SnapCenter. The integration between LSC and SnapCenter supports all SAP- supported databases. Nevertheless, we must differentiate between SAP AnyDBs and SAP HANA because SAP HANA provides a central communication host that is not available for SAP AnyDBs.

The default SnapCenter agent and database plug-in installation for SAP AnyDBs is a local installation from the SnapCenter agent in addition to the corresponding database plug-in for the database server.

In this section, the integration between LSC and SnapCenter is described using an SAP HANA database as an example. As previously stated for SAP HANA, there are two different options for the installation of the SnapCenter agent and SAP HANA database plug-in:

- **A standard SnapCenter agent and SAP HANA Plug-in installation.** In a standard installation, the SnapCenter agent and the SAP HANA Plug-in are locally installed on the SAP HANA database server.
- **A SnapCenter installation with a central communication host.** A central communication host is installed with the SnapCenter agent, the SAP HANA Plug-in, and the HANA database client that handles all database-related operations needed to back up and restore an SAP HANA database for several SAP HANA systems in the landscape. Therefore, a central communication host does not need to have a complete SAP HANA database system installed.

For more details regarding these different SnapCenter agents and SAP HANA database plug-in installation options, see the technical report [TR-4614: SAP HANA backup and recovery with SnapCenter](#).

The following sections highlight the differences between integrating LSC with SnapCenter using either the standard installation or the central communication host. Notably, all configuration steps that are not highlighted are the same regardless of the installation option and the database used.

To perform an automated Snapshot copy-based backup from the source database and create a clone for the new target database, the described integration between LSC and SnapCenter uses the configuration options and scripts described in [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#).

### Overview

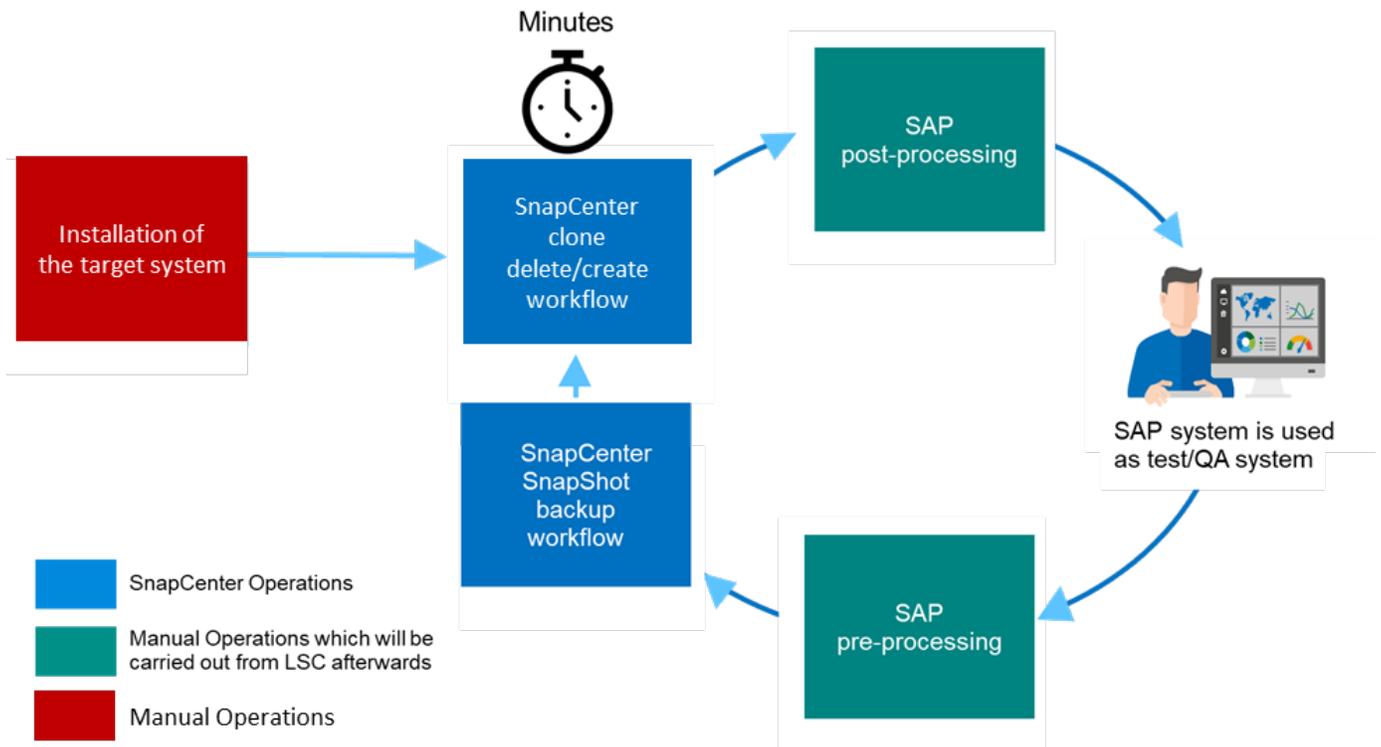
The following figure shows a typical high-level workflow for an SAP system refresh lifecycle with SnapCenter without LSC:

1. A one-time, initial installation and preparation of the target system.
2. Manual preprocessing (exporting licenses, users, printers, and so on).
3. If necessary, the deletion of an already existing clone on the target system.

4. The cloning of an existing Snapshot copy of the source system to the target system performed by SnapCenter.
5. Manual SAP post-processing operations (importing licenses, users, printers, disabling batch jobs, and so on).
6. The system can then be used as test or QA system.
7. When a new system refresh is requested, the workflow restarts at step 2.

SAP customers know that the manual steps colored in green in the figure below are time consuming and error prone. When using LSC and SnapCenter integration, these manual steps are carried out with LSC in a reliable and repeatable manner with all necessary logs needed for internal and external audits.

The following figure provides an overview of the general SnapCenter-based SAP system refresh procedure.



### Prerequisites and limitations

The following prerequisites must be fulfilled:

- SnapCenter must be installed. The source and target system must be configured in SnapCenter, either in a standard installation or by using a central communication host. Snapshot copies can be created on the source system.
- The storage backend must be configured properly in SnapCenter, as shown in the image below.

Storage Connections

<input type="checkbox"/>	Name	IP	Cluster Name	User Name	Controller License
<input type="checkbox"/>	<a href="#">svm-trident</a>		grenada.muccbc.hq.netapp.com		✓
<input type="checkbox"/>	<a href="#">svm-sap02</a>	10.65.58.253	grenada.muccbc.hq.netapp.com		✓
<input type="checkbox"/>	<a href="#">svm-sap01</a>	10.65.58.252	grenada.muccbc.hq.netapp.com		✓

The next two images cover the standard installation in which the SnapCenter agent and the SAP HANA Plug-in

are installed locally on each database server.

The SnapCenter agent and the appropriate database plug-in must be installed on the source database.

<input type="checkbox"/>	Name	Type	System	Plug-in	Version	Overall Status
<input type="checkbox"/>	<a href="#">sap-lnx35.muccbc.hq.netapp.com</a>	Linux	Stand-alone	UNIX, SAP HANA	4.3.1	<span style="color: green;">●</span> Running

The SnapCenter agent and the appropriate database plug-in must be installed on the target database.

<input type="checkbox"/>	<a href="#">sap-lnx36.muccbc.hq.netapp.com</a>	Linux	Stand-alone	UNIX, SAP HANA	4.3.1	<span style="color: green;">●</span> Running
--------------------------	--	-------	-------------	----------------	-------	--

The following image portrays central communication-host deployment in which the SnapCenter agent, the SAP HANA Plug-in, and the SAP HANA database client are installed on a centralized server (such as the SnapCenter Server) to manage several SAP HANA systems in the landscape.

The SnapCenter agent, the SAP HANA database plug-in, and the HANA database client must be installed on the central communication host.

<input type="checkbox"/>	Name	Type	System	Plug-in	Version	Overall Status
<input type="checkbox"/>	<a href="#">dlbh03.muccbc.hq.netapp.com</a>	Linux	Stand-alone	UNIX, SAP HANA	4.4	<span style="color: orange;">●</span> Upgrade available (optional)
<input type="checkbox"/>	<a href="#">sap-sc-demo-dev.muccbc.hq.netapp.com</a>	Windows	Stand-alone	Microsoft Windows Server, SAP HANA	4.5	<span style="color: green;">●</span> Running
<input type="checkbox"/>	<a href="#">sap-win02.muccbc.hq.netapp.com</a>	Windows	Stand-alone	Microsoft Windows Server	4.5	<span style="color: green;">●</span> Running

The backup for the source database must be configured properly in SnapCenter so that the Snapshot copy can be successfully created.

Backup Name	Count	End Date
SnapCenter__sap-lnx35_SAPhana_hourly_07-09-2020_13.00.02.4519	1	07/09/2020 1:01:42 PM
SnapCenter__sap-lnx35_SAPhana_hourly_07-09-2020_11.20.15.2146	1	07/09/2020 11:22:01 AM
Total 3		
Total 27		

The LSC master and the LSC worker must be installed in the SAP environment. In this deployment, we also installed the LSC master on the SnapCenter Server and the LSC worker on the target SAP database server, which should be refreshed. More details are described in the following section “[Lab setup.](#)”

Documentation resources:

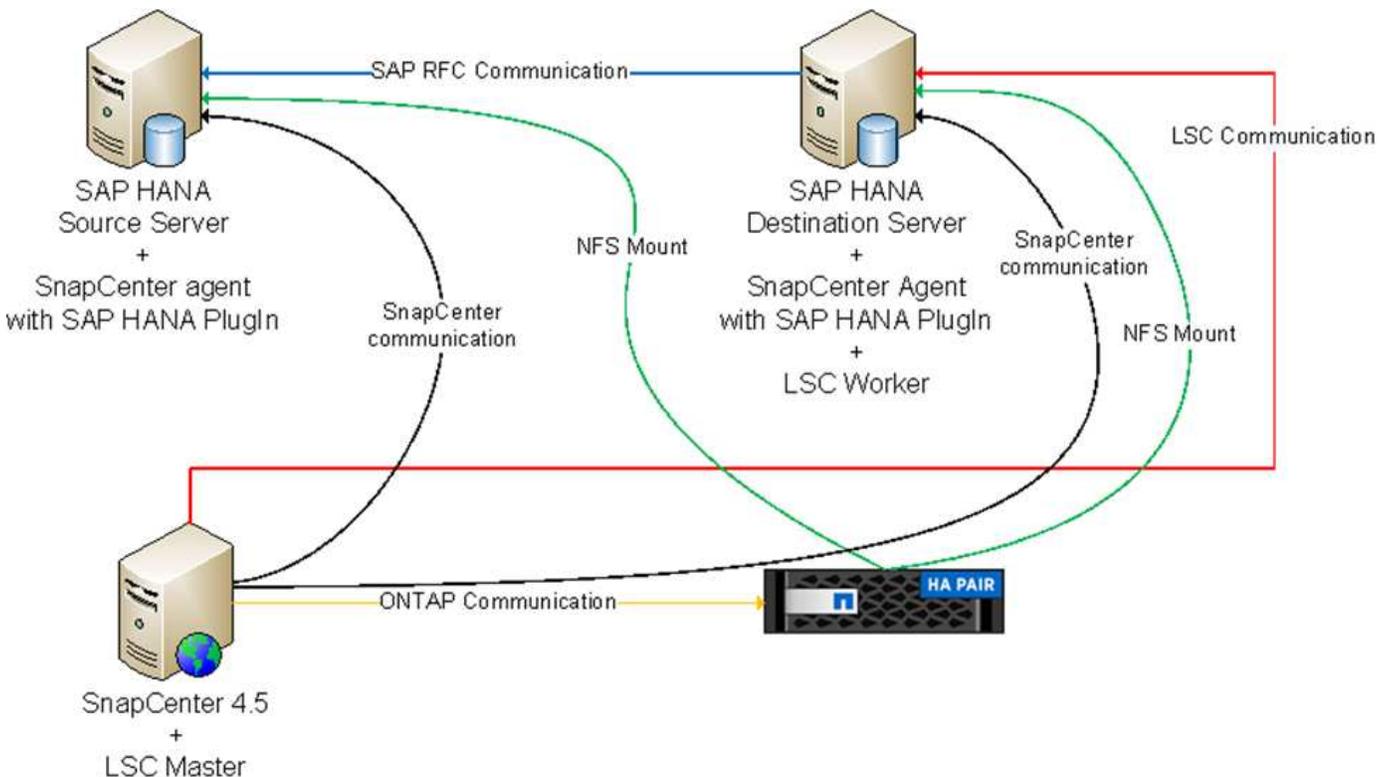
- [SnapCenter Documentation Center](#)
- [TR-4700: SnapCenter Plug-In for Oracle Database](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- [TR-4769 -SnapCenter Best Practices and Sizing Guidelines](#)
- [SnapCenter 4.6 Cmdlet Reference Guide](#)

## Lab setup

This section describes an example architecture that was set up in a demo data center. The setup was divided into a standard installation and an installation using a central communication host.

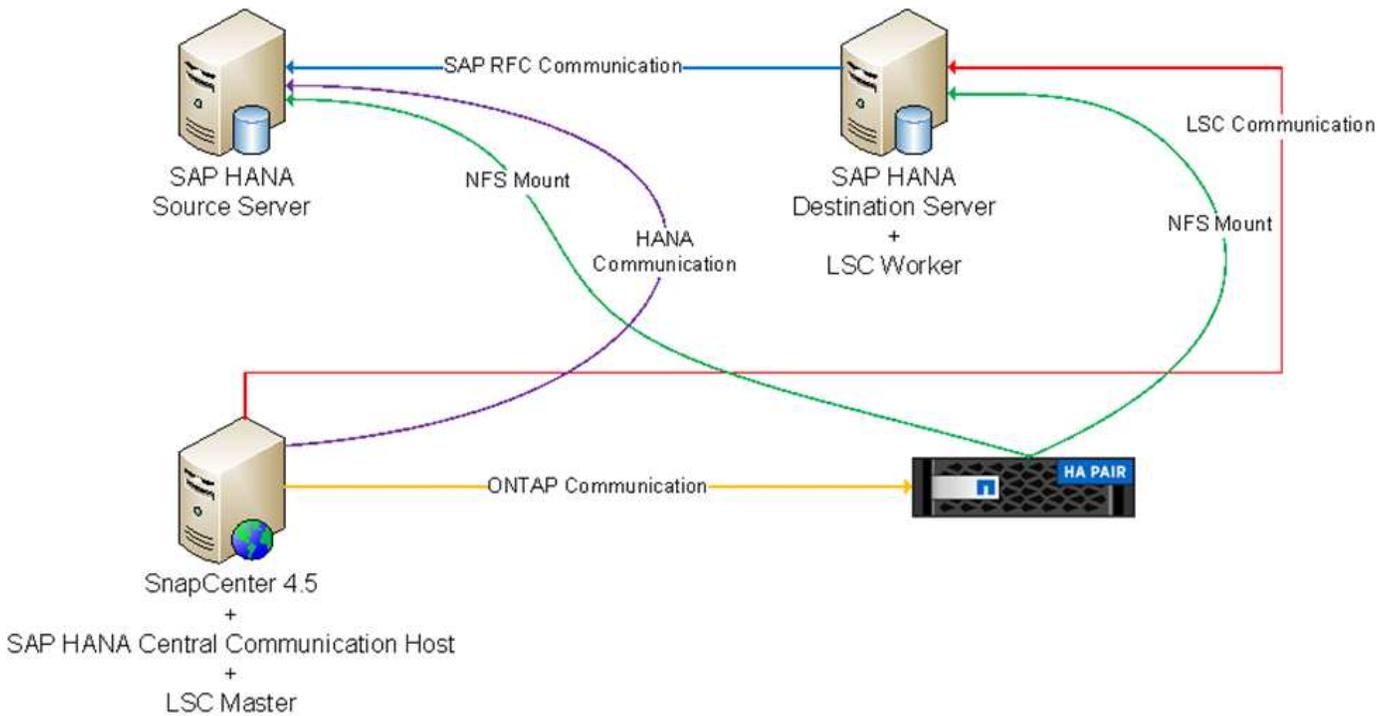
### Standard installation

The following figure shows a standard installation in which the SnapCenter agent together with the database plug-in was installed locally on the source and the target database server. In the lab setup, we installed the SAP HANA Plug-in. In addition, the LSC worker was also installed on the target server. For simplification and to reduce the number of virtual servers, we installed the LSC master on the SnapCenter Server. The communication between the different components is illustrated in the following figure.



### Central communication host

The following figure shows the setup using a central communication host. In this configuration, the SnapCenter agent together with the SAP HANA Plug-in and the HANA database client was installed on a dedicated server. In this setup, we used the SnapCenter Server to install the central communication host. In addition, the LSC worker was again installed on the target server. For simplification and to reduce the number of virtual servers, we decided to also install the LSC master on the SnapCenter Server. The communication between the different components is illustrated in the figure below.



### Initial one-time preparation steps for Libelle SystemCopy

There are three main components of an LSC installation:

- **LSC master.** As the name suggests, this is the master component that controls the automatic workflow of a Libelle-based system copy. In the demo environment, the LSC master was installed on the SnapCenter Server.
- **LSC worker.** An LSC worker is the part of the Libelle software that typically runs on the target SAP system and executes the scripts required for the automated system copy. In the demo environment, the LSC worker was installed on the target SAP HANA application server.
- **LSC satellite.** An LSC satellite is a part of the Libelle software that runs on a third-party system on which further scripts must be executed. The LSC master can also fulfill the role of an LSC satellite system at the same time.

We first defined all the involved systems inside LSC, as shown in the following image:

- **172.30.15.35.** The IP address of the SAP source system and the SAP HANA source system.
- **172.30.15.3.** The IP address of the LSC master and the LSC satellite system for this configuration. Because we installed the LSC master on the SnapCenter Server, the SnapCenter 4.x PowerShell Cmdlets are already available on this Windows host because they were installed during the SnapCenter Server installation. So, we decided to enable the LSC satellite role for this system and execute all SnapCenter PowerShell Cmdlets on this host. If you use a different system, make sure you install the SnapCenter PowerShell Cmdlets on this host according to the SnapCenter documentation.
- **172.30.15.36.** The IP address of the SAP destination system, the SAP HANA destination system, and the LSC worker.

Instead of IP addresses, host names, or fully qualified domain names can also be used.

The following image shows the LSC configuration of the master, worker, satellite, SAP source, SAP target, source database, and target database.

System Identifier	Worker	Source SAP	Source Database	Target SAP	Target Database	Satellite System
172.30.15.35		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
172.30.15.3	172.30.15.3:9000	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.30.15.36	172.30.15.36:9000	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

For the main integration, we must again separate the configuration steps into the standard installation and the installation using a central communication host.

### Standard installation

This section describes the configuration steps needed when using a standard installation where the SnapCenter agent and the necessary database plug-in are installed on the source and target systems. When using a standard installation, all tasks needed to mount the clone volume and to restore and recover the target system are carried out from the SnapCenter agent that is running on the target database system on the server itself. This allows access to all the clone-related details that are available through environmental variables from the SnapCenter agent. Therefore, you only need to create one additional task in the LSC copy phase. This task carries out the Snapshot copy process on the source database system and the clone and restore and recovery process on the target database system. All SnapCenter related tasks are triggered by using a PowerShell script that is entered in the LSC task `NTAP_SYSTEM_CLONE`.

The following image shows LSC task configuration in the copy phase.

copy	Copy Phase		phase
copy 1	NTAP_SYSTEM_CLONE	NetApp SnapShot and Clone	psh
copy 2	NTAP_SYSTEM_CLONE_CP	NetApp SnapShot and Clone	psh
copy 3	NTAP_MNT_RECOVER_CP	Mount Volume and Recover HANA Database	cmd
copy 4	LPDBBCKP	Backup Source DB in Filesystem	ish
copy 5	LPDBCOPYFLS	Copy DB Backup Files From Source to Target System	ish
copy 6	LTDBRESTORE	Restore DB Files	ish
copy 7	LTDBRESTORE_TENANT	Restore DB Files for Tenant Database	ish
post	Post Phase		phase

The following image highlights the configuration of the `NTAP_SYSTEM_CLONE` process. Because you are executing a PowerShell script, this Windows PowerShell script is executed on the satellite system. In this instance, this is the SnapCenter Server with the installed LSC master that also acts as a satellite system.

Task: NTAP\_SYSTEM\_CLONE Version: 0

**Configuration Data**

Main Attributes

Comment

Category

**Execution Attributes**

Parameters

Return Codes

Code

Activated:  Wait after execution:

Type: Windows PowerShell Script

**Systems**

- Execute task for all systems with any of the roles:
  - Source SAP
  - Source Database
  - Target SAP
  - Target Database
  - Satellite System
- Execute task for the following systems (selected by their IDs):

**Clients**

- Execute task with the system's default client.
- Execute task with every client having the copy flag set.
- Execute task with each client defined in the system.
- Execute task with the following clients:

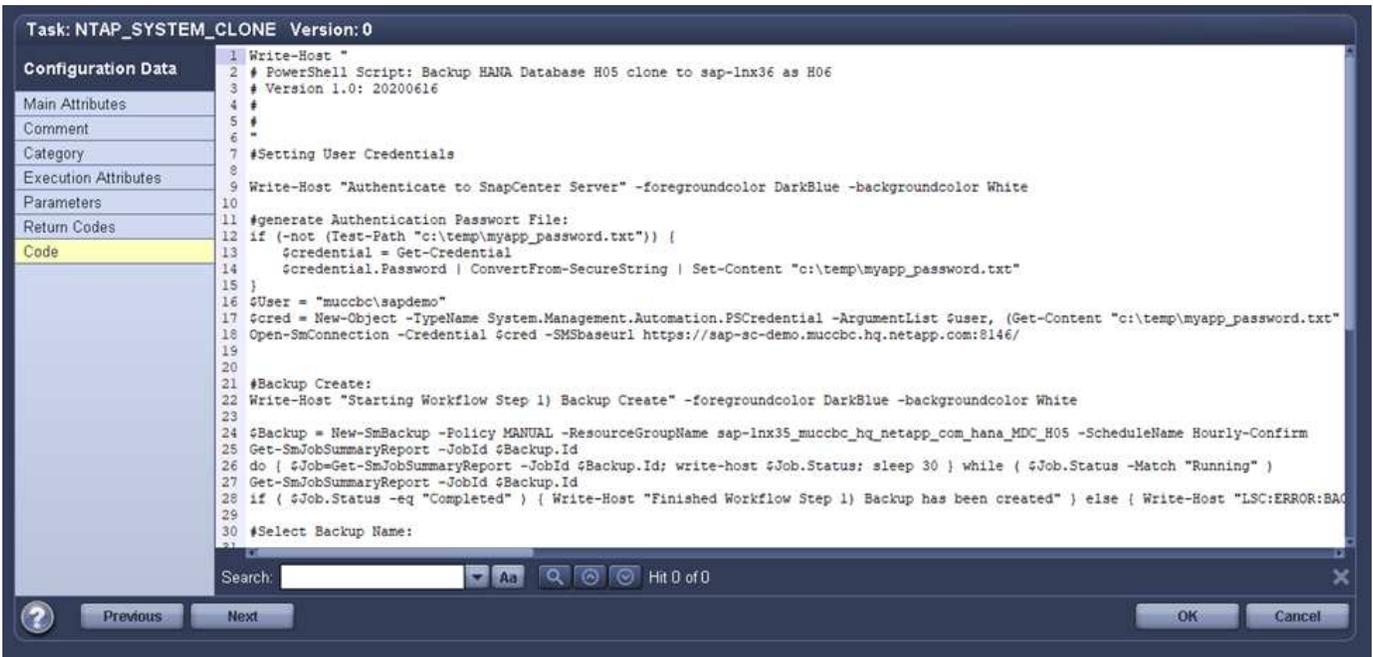
Buttons: Previous, Next, OK, Cancel

Because LSC must be made aware of whether the Snapshot copy, cloning, and recovery operation has been successful, you must define at least two return code types. One code is for a successful execution of the script, and the other code is for a failed execution of the script, as shown in the following image.

- LSC:OK must be written from the script to standard out if the execution was successful.
- LSC:ERROR must be written from the script to standard out if the execution has failed.



The following image shows part of the PowerShell script that must run to execute a Snapshot-based backup on the source database system and a clone on the target database system. The script is not intended to be complete. Rather, the script shows how integration between LSC and SnapCenter can look and how easy it is to set it up.



Because the script is executed on the LSC master (which is also a satellite system), the LSC master on the SnapCenter Server must be run as a Windows user that has appropriate permissions to execute backup and

cloning operations in SnapCenter. To verify whether the user has appropriate permission, the user should be able to execute a Snapshot copy and a clone in the SnapCenter UI.

There is no need to run the LSC master and the LSC satellite on the SnapCenter Server itself. The LSC master and the LSC satellite can run on any Windows machine. The prerequisite for running the PowerShell script on the LSC satellite is that the SnapCenter PowerShell cmdlets have been installed on the Windows Server.

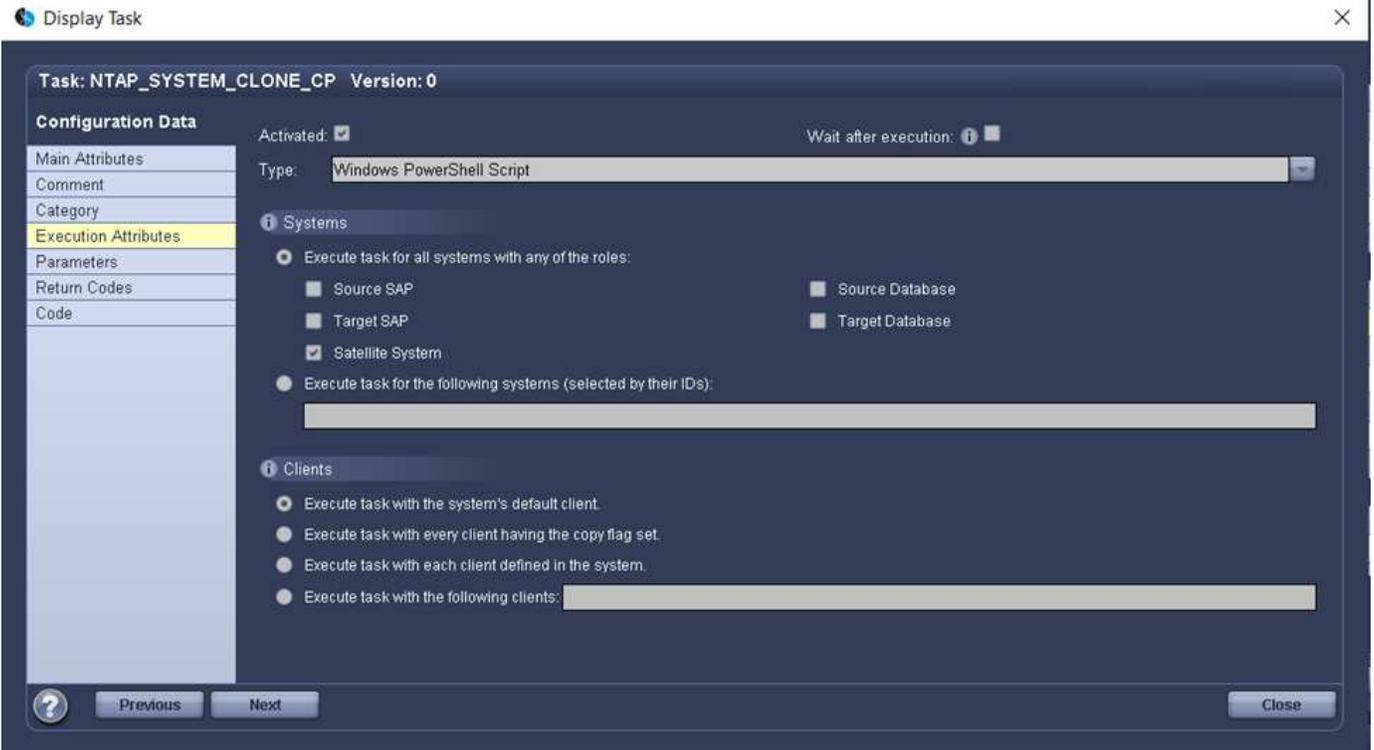
### Central communication host

For integration between LSC and SnapCenter using a central communication host, the only adjustments that must be made are performed in the copy phase. The Snapshot copy and the clone are created using the SnapCenter agent on the central communication host. Therefore, all details about the newly created volumes are only available on the central communication host and not on the target database server. However, these details are needed on the target database server to mount the clone volume and to carry out the recovery. This is the reason why two additional tasks are needed in the copy phase. One task is executed on the central communication host and one task is executed on the target database server. These two tasks are shown in the image below.

- **NTAP\_SYSTEM\_CLONE\_CP.** This task creates the Snapshot copy and the clone using a PowerShell script that executes the necessary SnapCenter functions on the central communication host. This task therefore runs on the LSC satellite, which in our instance is the LSC master that runs on Windows. This script collects all details about the clone and the newly created volumes and hands it over to the second task **NTAP\_MNT\_RECOVER\_CP**, which runs on the LSC worker that runs on the target database server.
- **NTAP\_MNT\_RECOVER\_CP.** This task stops the target SAP system and the SAP HANA database, unmounts the old volumes, and then mounts the newly created storage clone volumes based on the parameters that were passed through from the previous task **NTAP\_SYSTEM\_CLONE\_CP**. The target SAP HANA database is then restored and recovered.

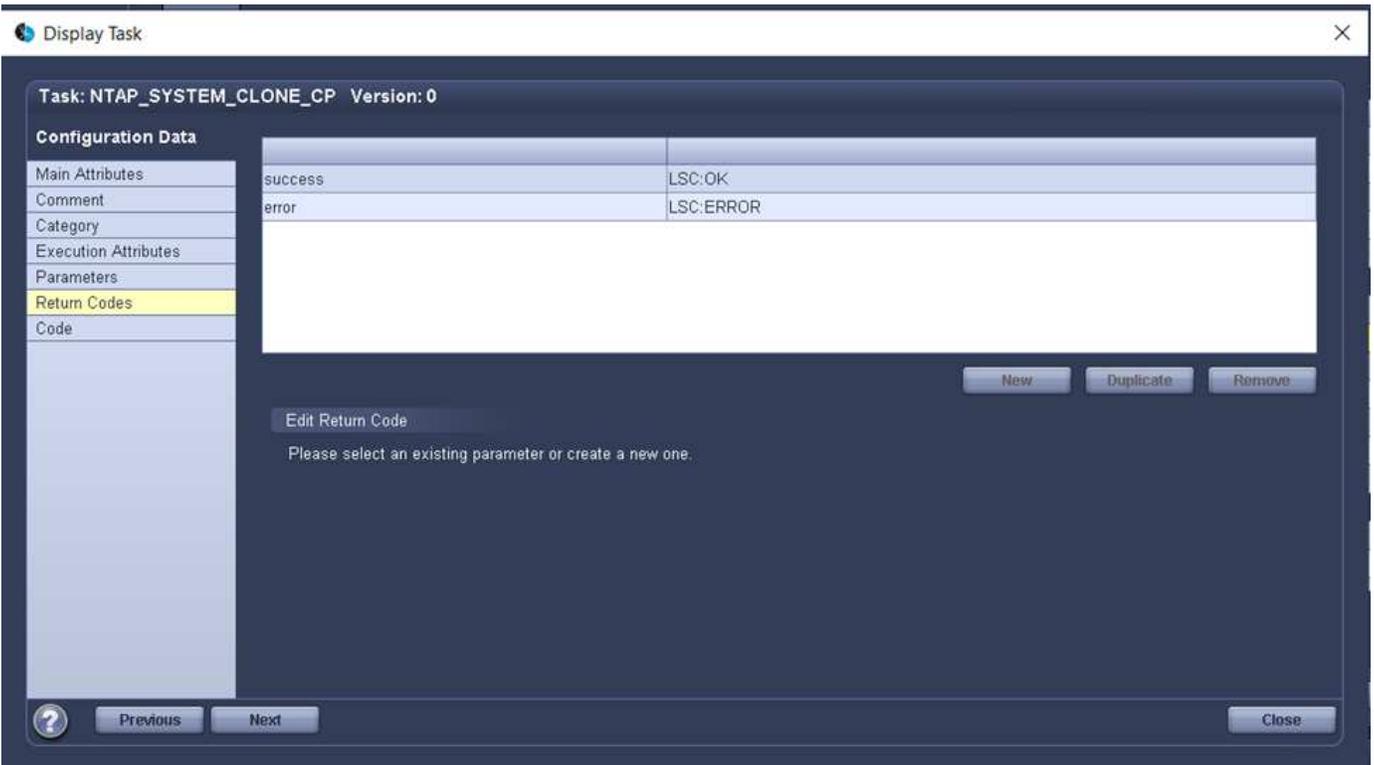
copy	Copy Phase		phase
copy 1	NTAP_SYSTEM_CLONE	NetApp Snapshot and Clone	psh
copy 2	NTAP_SYSTEM_CLONE_CP	NetApp Snapshot and Clone	psh
copy 3	NTAP_MNT_RECOVER_CP	Mount Volume and Recover HANA Database	cmd
copy 4	LPDBBCKP	Backup Source DB in Filesystem	lsh
copy 5	LPDBCOPYFLS	Copy DB Backup Files From Source to Target System	lsh
copy 6	LTDBRESTORE	Restore DB Files	lsh
copy 7	LTDBRESTORE_TENANT	Restore DB Files for Tenant Database	lsh
post	Post Phase		phase

The following image highlights the configuration of the task **NTAP\_SYSTEM\_CLONE\_CP**. This is the Windows PowerShell script that is executed on the satellite system. In this instance, the satellite system is the SnapCenter Server with the installed LSC master.

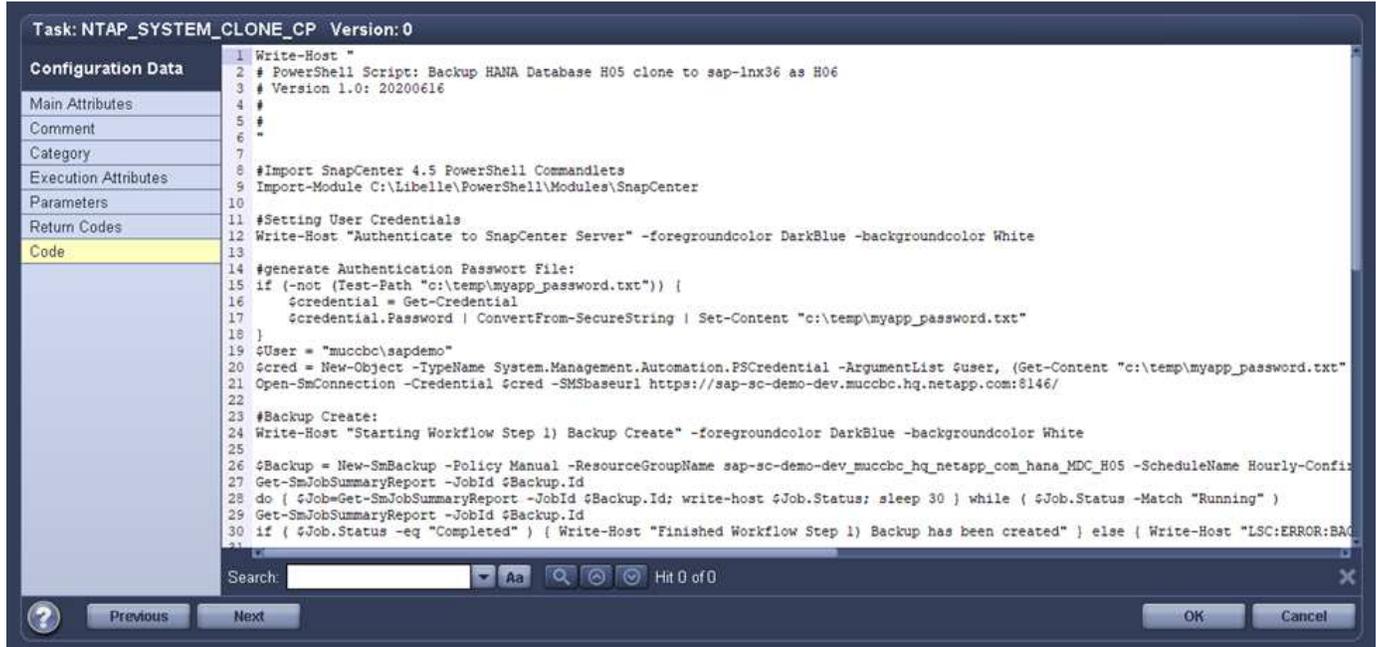


Because LSC must be aware of whether the Snapshot copy and cloning operation was successful, you must define at least two return code types: one return code for a successful execution of the script and the other for a failed execution of the script, as shown in the image below.

- LSC :OK must be written from the script to standard out if the execution was successful.
- LSC :ERROR must be written from the script to standard out if the execution failed.



The following image shows part of the PowerShell script that must run to execute a Snapshot copy and a clone using the SnapCenter agent on the central communication host. The script is not meant to be complete. Rather, the script is used to show how integration between LSC and SnapCenter can look and how easy it is to set it up.

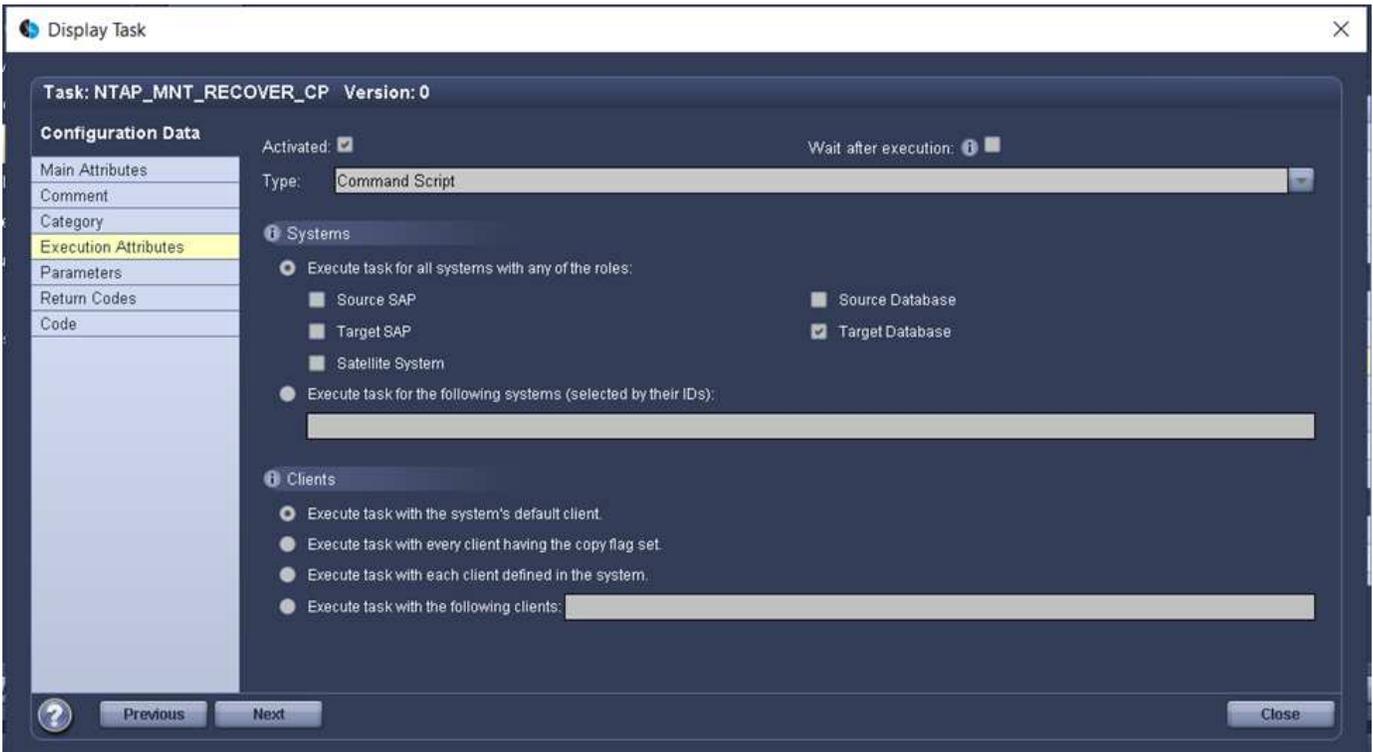


As previously mentioned, you must hand over the name of the clone volume to the next task NTAP\_MNT\_RECOVER\_CP to mount the clone volume on the target server. The name of the clone volume, also known as the junction path, is stored in the variable \$JunctionPath. The handover to a subsequent LSC task is achieved through a custom LSC variable.

```
echo $JunctionPath > $_task(current, custompath1)_$_
```

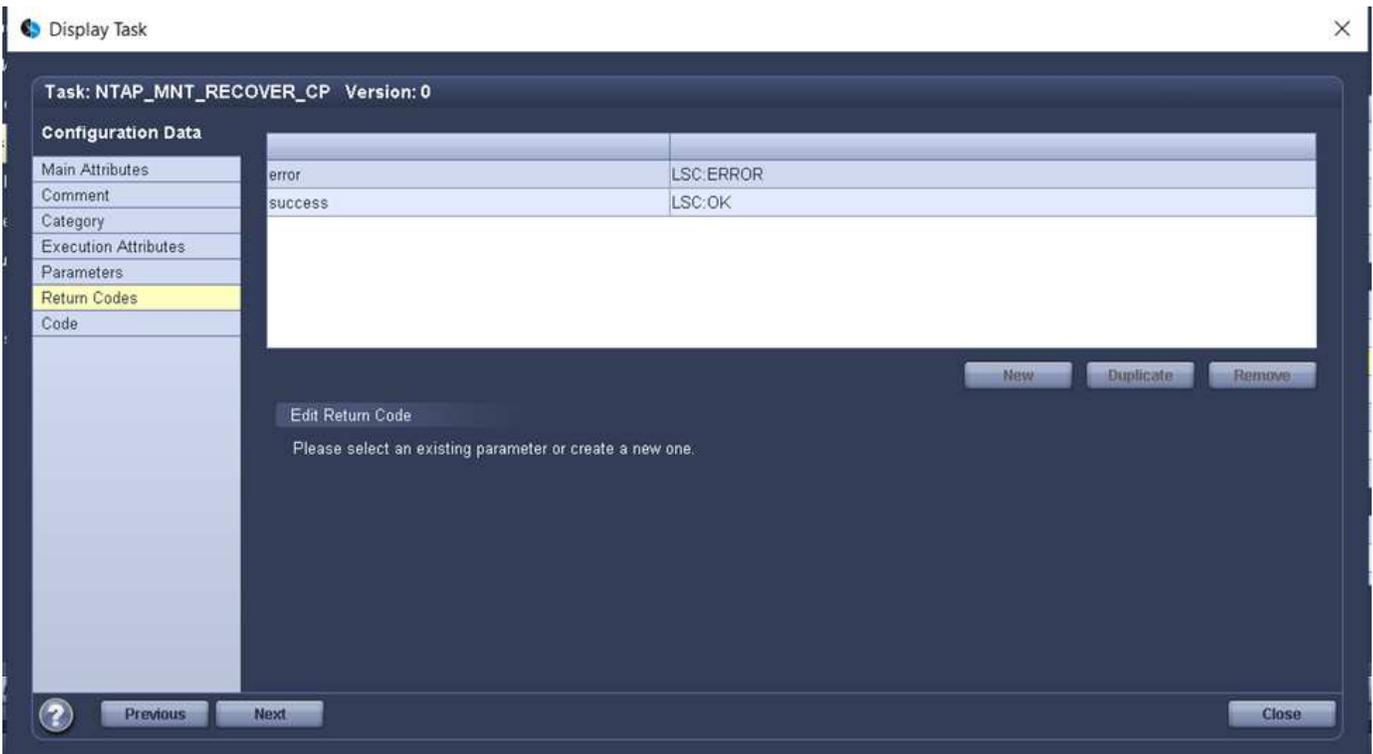
Because the script is executed on the LSC master (which is also a satellite system), the LSC master on the SnapCenter Server must run as a Windows user that has appropriate permissions to execute the backup and cloning operations in SnapCenter. To verify whether it has the appropriate permissions, the user should be able to execute a Snapshot copy and clone in the SnapCenter GUI.

The following figure highlights the configuration of the task NTAP\_MNT\_RECOVER\_CP. Because we want to execute a Linux Shell script, this is a command script executed on the target database system.



Because LSC must be made aware of mounting the clone volumes and whether restoring and recovering the target database was successful, we must define at least two return code types. One code is for a successful execution of the script, and one is for a failed execution of the script, as is shown in the following figure.

- LSC:OK must be written from the script to standard out if the execution was successful.
- LSC:ERROR must be written from the script to standard out if the execution failed.



The following figure shows part of the Linux Shell script used to stop the target database, unmount the old

volume, mount the clone volume, and restore and recover the target database. In the previous task, the junction path was written into an LSC variable. The following command reads this LSC variable and stores the value in the `$JunctionPath` variable of the Linux Shell script.

```
JunctionPath=$_include($_task(NTAP_SYSTEM_CLONE_CP, custompath1)_$, 1, 1)_$_$
```

The LSC worker on the target system runs as `<sideadm>`, but mount commands must be run as the root user. This is why you must create the `central_plugin_host_wrapper_script.sh`. The script `central_plugin_host_wrapper_script.sh` is called from the task `NTAP_MNT_RECOVERY_CP` using the `sudo` command. Using the `sudo` command, the script runs with UID 0 and we are able to carry out all subsequent steps, such as unmounting the old volumes, mounting the clone volumes, and restoring and recovering the target database. To enable script execution using `sudo`, the following line must be added in `/etc/sudoers`:

```
hn6adm ALL=(root)
NOPASSWD:/usr/local/bin/H06/central_plugin_host_wrapper_script.sh
```



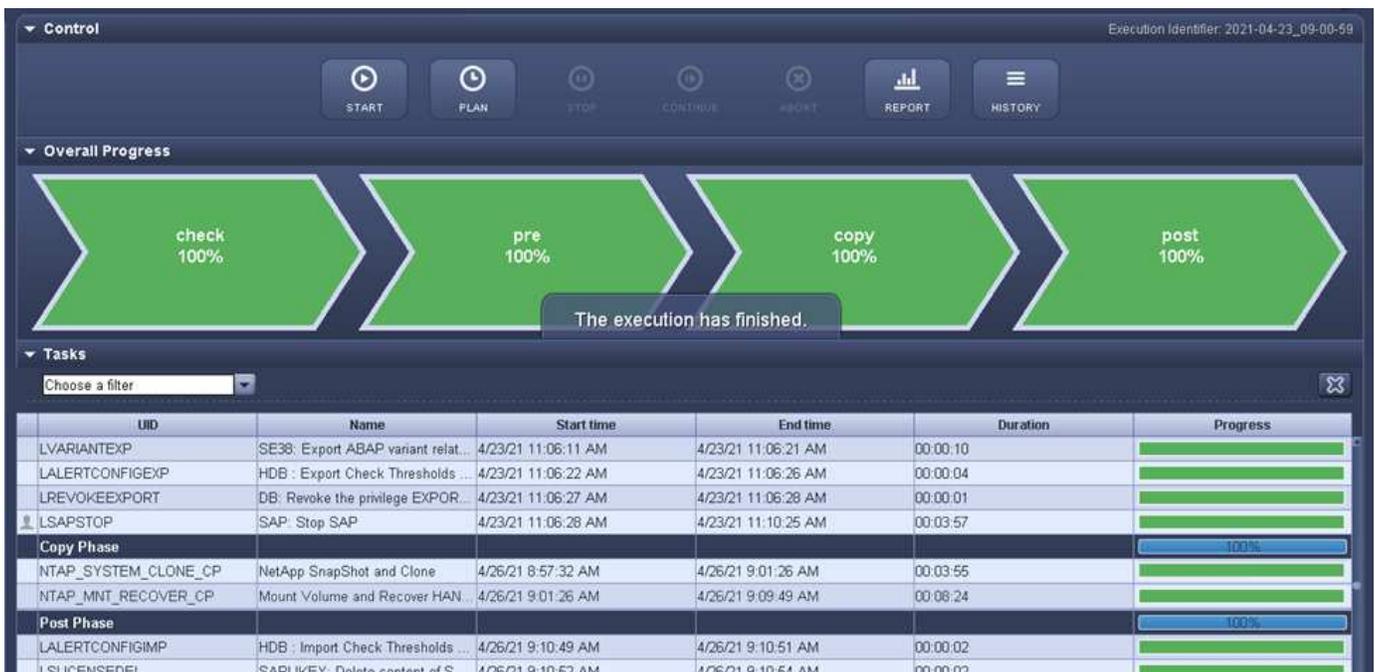
## SAP HANA system refresh operation

Now that all necessary integration tasks between LSC and NetApp SnapCenter have been carried out, starting a fully automated SAP system refresh is a one-click task.

The following figure shows the task `NTAP` ``SYSTEM` `CLONE` in a standard installation. As you can see, creating a Snapshot copy and a clone, mounting the clone volume on the target database server, and restoring and recovering the target database took approximately 14 minutes. Remarkably, with Snapshot and NetApp FlexClone technology, the duration of this task remains nearly the same, independent of the size of the source database.



The following figure shows the two tasks NTAP\_SYSTEM\_CLONE\_CP and NTAP\_MNT\_RECOVERY\_CP when using a central communication host. As you can see, creating a Snapshot copy, a clone, mounting the clone volume on the target database server, and restoring and recovering the target database took approximately 12 minutes. This is more or less the same time needed to carry out these steps when using a standard installation. Again, Snapshot and NetApp FlexClone technology enables the consistent, rapid completion of these tasks, independent of the size of the source database.



## SAP HANA system refresh with LSC, AzAcSnap, and Azure NetApp Files

Using [Azure NetApp Files for SAP HANA](#), Oracle, and DB2 on Azure provides customers with the advanced data management and data protection features of NetApp ONTAP with the native Microsoft Azure NetApp Files service. [AzAcSnap](#) is the foundation for very fast SAP system refresh operations to create application-consistent NetApp Snapshot copies

of SAP HANA and Oracle systems (DB2 is not currently supported by AzAcSnap).

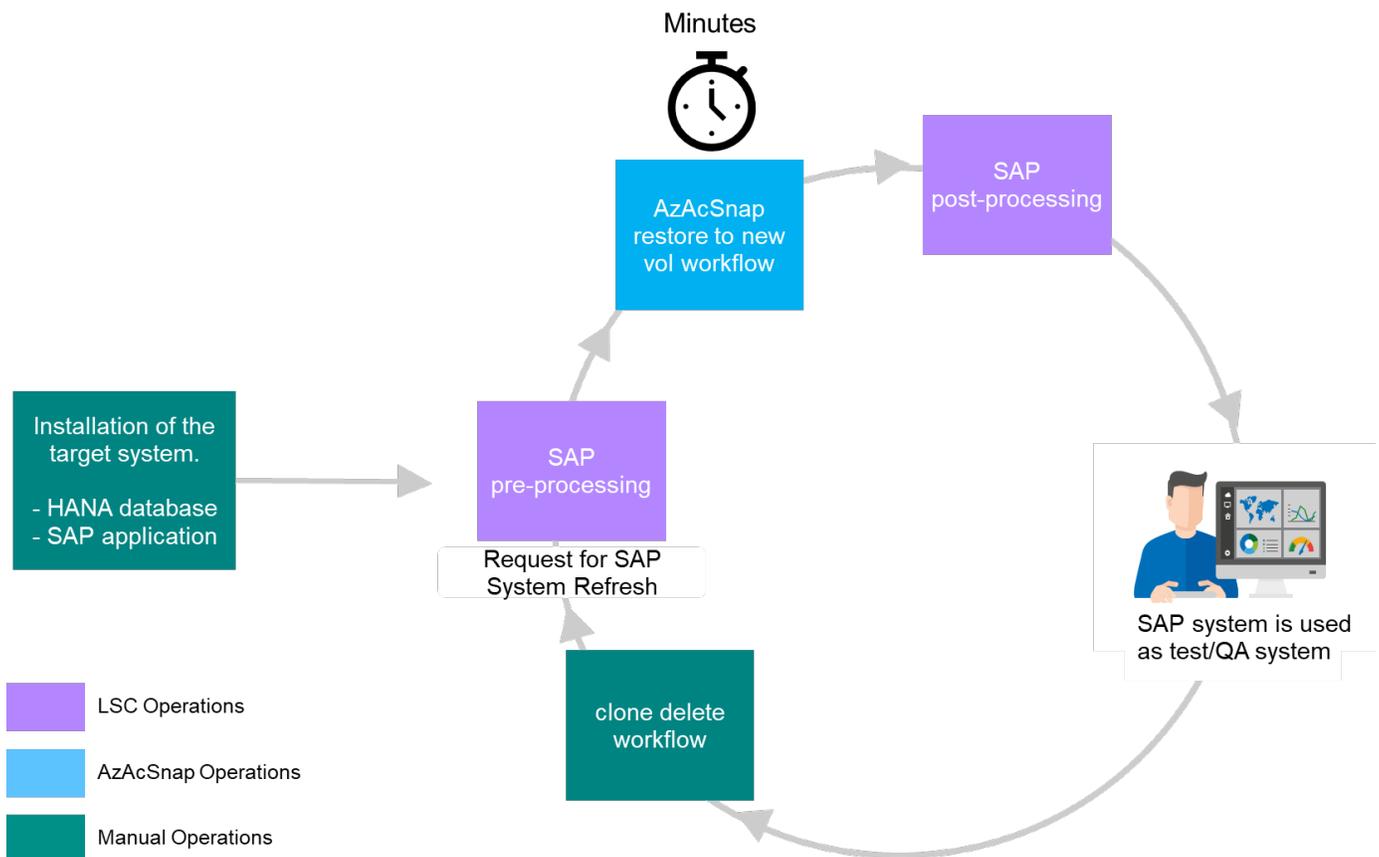
Snapshot copy backups, which are created either on-demand or on a regular basis as part of the backup strategy, can then be efficiently cloned to new volumes and used to quickly refresh target systems. AzAcSnap provides the workflows necessary to create backups and clone them to new volumes, while Libelle SystemCopy performs the pre- and post-processing steps necessary for a full end-to-end system refresh.

In this chapter, we describe an automated SAP system refresh using AzAcSnap and Libelle SystemCopy using SAP HANA as the underlying database. Because AzAcSnap is also available for Oracle, the same procedure can also be implemented using AzAcSnap for Oracle. Other databases might be supported by AzAcSnap in the future, which would then enable system copy operations for those databases with LSC and AzAcSnap.

The following figure shows a typical high-level workflow of an SAP system refresh lifecycle with AzAcSnap and LSC:

- A one-time, initial installation and preparation of the target system.
- SAP preprocessing operations performed by LSC.
- Restoring (or cloning) an existing Snapshot copy of the source system to the target system performed by AzAcSnap.
- SAP post-processing operations performed by LSC.

The system can then be used as a test or QA system. When a new system refresh is requested, the workflow restarts with step 2. Any remaining cloned volumes must be deleted manually.

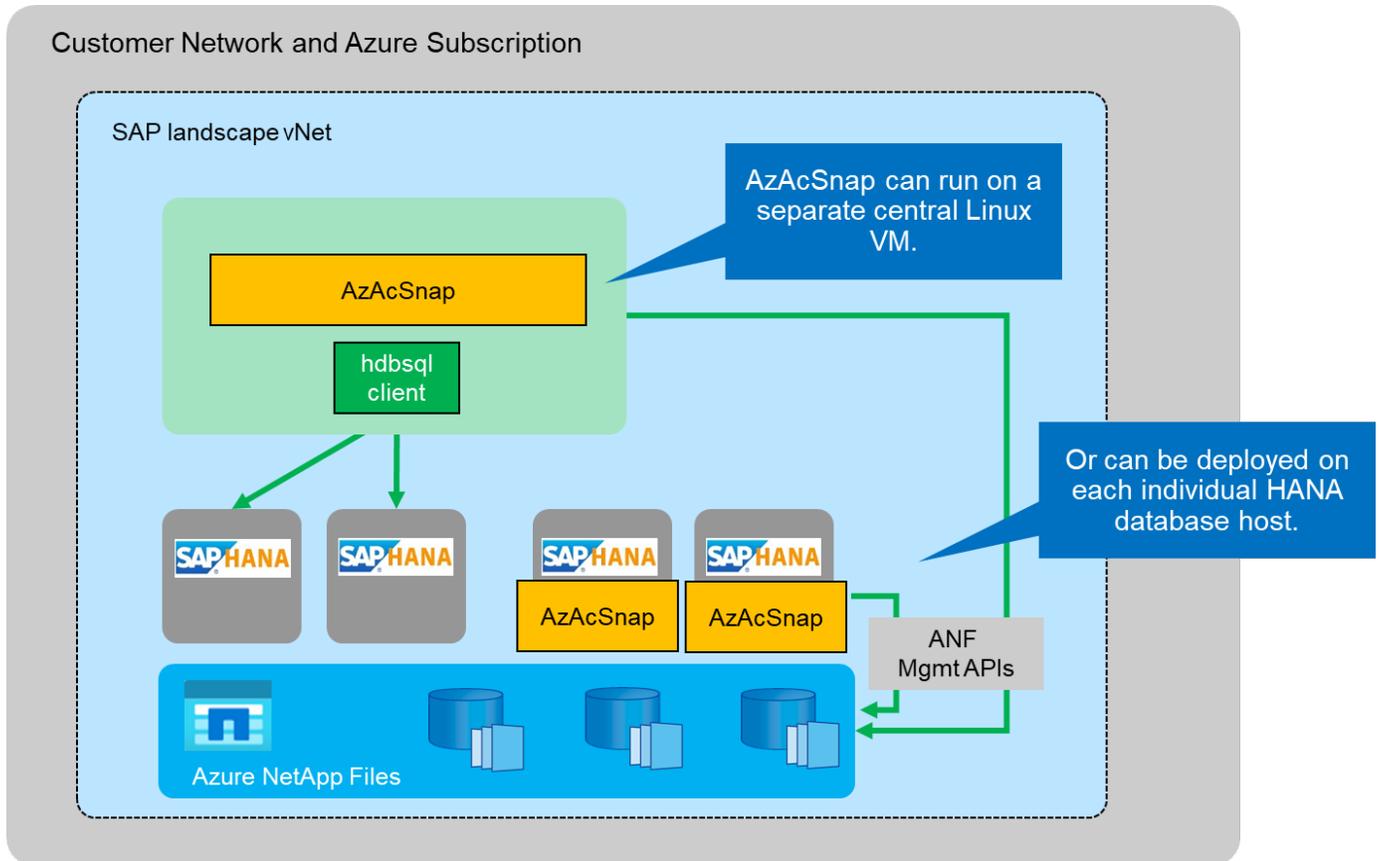


## Prerequisites and limitations

The following prerequisites must be fulfilled.

### AzAcSnap installed and configured for the source database

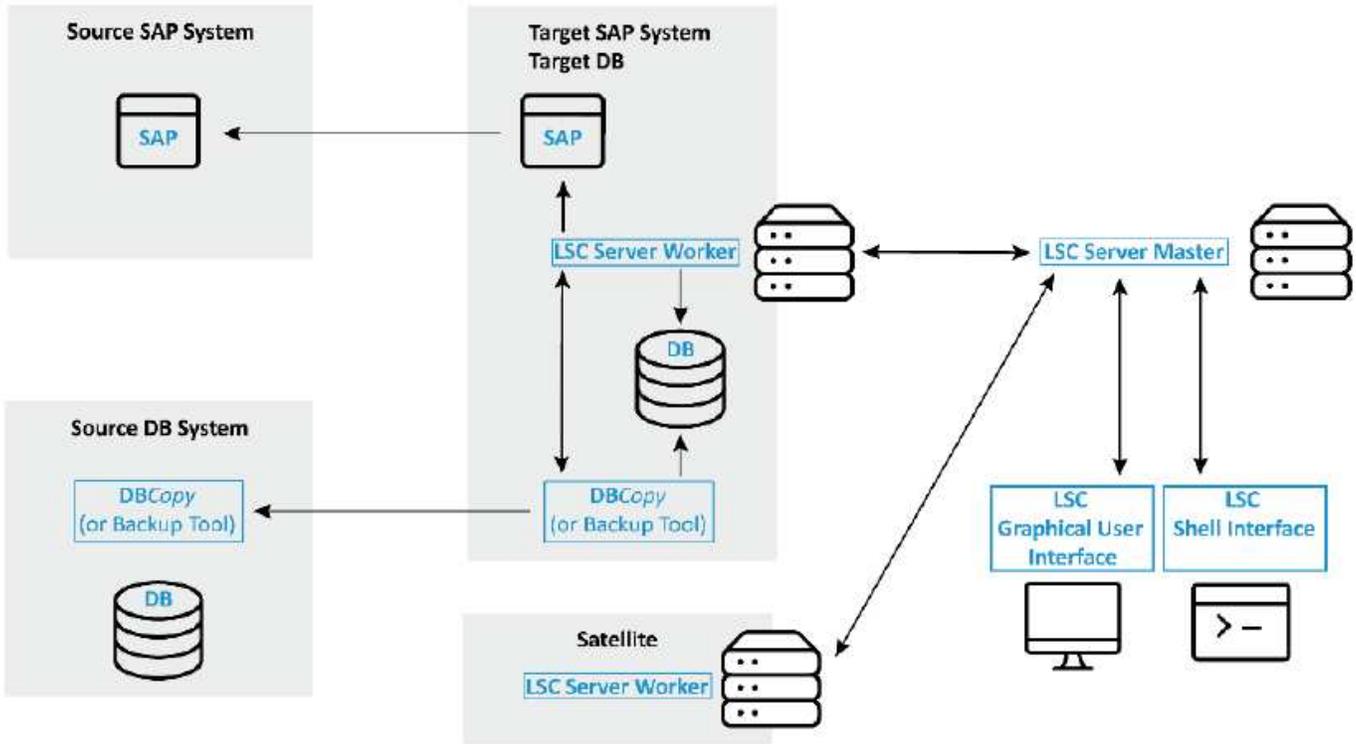
In general, there are two deployments options for AzAcSnap, as is shown in the following picture.



AzAcSnap can be installed and run on a central Linux VM for which all DB configuration files are stored centrally and AzAcSnap has access to all databases (through the hdbsql client) and the configured HANA userstore keys for all these databases. With a decentralized deployment, AzAcSnap is installed individually on each database host where typically only the DB configuration for the local database is stored. Both deployment options are supported for LSC integration. However, we followed a hybrid approach in the lab setup for this document. AzAcSnap was installed on a central NFS share along with all DB configuration files. This central installation share was mounted on all VMs under `/mnt/software/AZACSNAP/snapshot-tool`. The execution of the tool was then performed locally on the DB VMs.

### Libelle SystemCopy installed and configured for source and target SAP system

Libelle SystemCopy deployments consist of the following components:



- **LSC Master.** As the name suggests, this is the master component that controls the automatic workflow of a Libelle-based system copy.
- **LSC Worker.** An LSC worker usually runs on the target SAP system and executes the scripts required for the automated system copy.
- **LSC Satellite.** An LSC satellite runs on a third-party system on which further scripts must be executed. The LSC master can also fulfill the role of an LSC satellite system.

The Libelle SystemCopy (LSC) GUI must be installed on a suitable VM. In this lab setup, the LSC GUI was installed on a separate Windows VM, but it can also run on the DB host together with the LSC worker. The LSC worker must be installed at least on the VM of the target DB. Depending on your chosen AzAcSnap deployment option, additional LSC worker installations might be required. You must have an LSC worker installation on the VM where AzAcSnap is executed.

After LSC is installed, the basic configuration for the source and the target database must be performed according to the LSC guidelines. The following images shows the configuration of the lab environment for this document. See the next section for details about the source and the target SAP systems and databases.



You should also configure a suitable standard task list for the SAP systems. For more details about the installation and configuration of LSC, consult the LSC user manual that is part of the LSC installation package.

### Known limitations

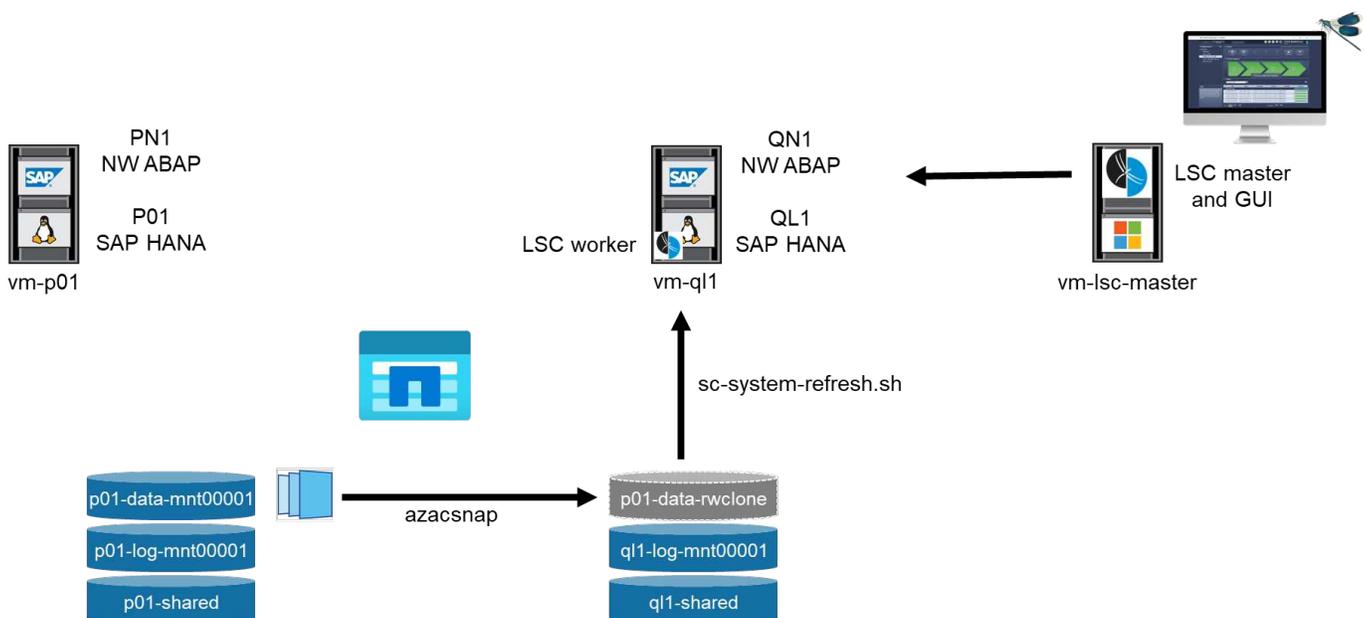
The AzAcSnap and LSC integration described here only works for SAP HANA single-host databases. SAP HANA multiple-host (or scale-out) deployments can also be supported, but such deployments require a few adjustments or enhancements to the LSC custom tasks for the copy phase and the underlying scripts. Such enhancements are not covered in this document.

SAP system refresh integration always uses the latest successful Snapshot copy of the source system to perform the refresh of the target system. If you would like to use other older Snapshot copies, the corresponding logic in the [ZAZACSNAPRESTORE](#) custom task must be adjusted. This process is out of scope for this document.

### Lab setup

The lab setup consists of a source SAP system and a target SAP system, both running on SAP HANA single-host databases.

The following picture shows the lab setup.



It contains the following systems, software versions, and Azure NetApp Files volumes:

- **P01.** SAP HANA 2.0 SP5 database. Source database, single host, single user tenant.
- **PN1.** SAP NetWeaver ABAP 7.51. Source SAP system.
- **vm-p01.** SLES 15 SP2 with AzAcSnap installed. Source VM hosting P01 and PN1.
- **QL1.** SAP HANA 2.0 SP5 database. System refresh target database, single host, single-user tenant.
- **QN1.** SAP NetWeaver ABAP 7.51. System refresh target SAP system.
- **vm-ql1.** SLES 15 SP2 with LSC worker installed. Target VM hosting QL1 and QN1.
- LSC master version 9.0.0.0.052.
- **vm-lsc-master.** Windows Server 2016. Hosts LSC master and LSC GUI.

- Azure NetApp Files volumes for data, log, and shared for P01 and QL1 mounted on the dedicated DB hosts.
- Central Azure NetApp Files volume for scripts, AzAcSnap installation, and configuration files mounted on all VMs.

### Initial one-time preparation steps

Before the first SAP system refresh can be executed, you must integrate Azure NetApp Files Snapshot copy-and-cloning-based storage operations executed by AzAcSnap. You must also execute an auxiliary script for starting and stopping the database and mounting or unmounting the Azure NetApp Files volumes. All required tasks are performed as custom tasks in LSC as part of the copy phase. The following picture shows the custom tasks in the LSC task list.

	Phase	UID	Name	Type
pre 76	LALERTCONFIGEXP		HDB : Export Check Threshold...	ish
pre 77	LREVOKEEXPORT		DB: Revoke the privilege EXPO...	cmd
pre 78	LJAVACONFEXP		JAVA: Backup java config files...	cmd
pre 79	LSTOPSLTJOBS		LTRC: Stop all replication jobs ...	ish
pre 80	LSAPSTOP		SAP: Stop SAP	intv
pre 81	LSTOPSAPSYSTEM		Stops all SAP instances (appli...	ish
<b>copy</b>	<b>Copy Phase</b>			<b>phase</b>
copy 1	ZSCCOPYSHUTDOWN		Shutdown HANA DB	cmd
copy 2	ZSCCOPYUMOUNT		Unmount data volumes	cmd
copy 3	ZAZACSNAPRESTORE		Restore snapshot backup of so...	cmd
copy 4	ZSCCOPYMOUNT		Mount data volumes	cmd
copy 5	ZSCCOPYRECOVER		Recover target DB based on sn...	cmd
<b>post</b>	<b>Post Phase</b>			<b>phase</b>
post 1	LCHNGHDBPWD		HDB : Restore the password fo...	cmd
post 2	LHDBLICIMP		HANA DB License Import	ish
post 3	LALERTCONFIGIMP		HDB : Import Check Threshold...	ish

All five copy tasks are described here in more detail. In some of these tasks, a sample script `sc-system-refresh.sh` is used to further automate the required SAP HANA database recovery operation and the mount and unmount of the data volumes. The script uses an `LSC: success` message in the system output to indicate a successful execution to LSC. Details about custom tasks and available parameters can be found in the LSC user manual and the LSC developer guide. All tasks in this lab environment are executed on the target DB VM.



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

### Sc-system-refresh.sh configuration file

As mentioned before, an auxiliary script is used to start and stop the database, to mount and unmount the Azure NetApp Files volumes, and to recover the SAP HANA database from a Snapshot copy. The script `sc-system-refresh.sh` is stored on the central NFS share. The script requires a configuration file for each target database that must be stored in the same folder as the script itself. The configuration file must have the following name: `sc-system-refresh-<target DB SID>.cfg` (for example `sc-system-refresh-QL1.cfg` in this lab environment). The configuration file used here uses a fixed/hard-coded source DB SID. With a few changes, the script and the config file can be enhanced to take the source DB SID as an input parameter.

The following parameters must be adjusted according to the specific environment:

```
# hdbuserstore key, which should be used to connect to the target database
KEY="QL1SYSTEM"
# single container or MDC
export P01_HANA_DATABASE_TYPE=MULTIPLE_CONTAINERS
# source tenant names { TENANT_SID [, TENANT_SID]* }
export P01_TENANT_DATABASE_NAMES=P01
# cloned vol mount path
export CLONED_VOLUMES_MOUNT_PATH=`tail -2
/mnt/software/AZACSNAP/snapshot_tool/logs/azacsnap-restore-azacsnap-
P01.log | grep -oe "[0-9]*\.[0-9]*\.[0-9]*\.[0-9]*:/*.* "`
```

## ZSCCOPYSHUTDOWN

This task stops the target SAP HANA database. The Code section of this task contains the following text:

```
$_include_tool(unix_header.sh)_$
sudo /mnt/software/scripts/sc-system-refresh/sc-system-refresh.sh shutdown
$_system(target_db, id)_$ > $_logfile_
```

The script `sc-system-refresh.sh` takes two parameters, the `shutdown` command and the DB SID, to stop the SAP HANA database using `sapcontrol`. The system output is redirected to the standard LSC logfile. As mentioned before, an LSC: success message is used to indicate successful execution.

Task: ZSCCOPYSHUTDOWN Version: 0	
<b>Configuration Data</b>	
Main Attributes	success LSC:success
Comment	
Category	
Execution Attributes	
Parameters	
Return Codes	
Code	

## ZSCCOPYUMOUNT

This task unmounts the old Azure NetApp Files data volume from the target DB operating system (OS). The code section of this task contains the following text:

```
$_include_tool(unix_header.sh)_$
sudo /mnt/software/scripts/sc-system-refresh/sc-system-refresh.sh umount
$_system(target_db, id)_$ > $_logfile_
```

The same scripts as in the previous task is used. The two parameters passed are the `umount` command and the DB SID.

## ZAZACSNAPRESTORE

This task runs AzAcSnap to clone the latest successful Snapshot copy of the source database to a new

volume for the target database. This operation is equivalent to a redirected restore of backup in traditional backup environments. However, the Snapshot copy and cloning functionality enables you to perform this task within seconds even for the largest databases, whereas, with traditional backups, this task could easily take several hours. The code section of this task contains the following text:

```

$_include_tool(unix_header.sh)_$
sudo /mnt/software/AZACSNAP/snapshot_tool/azacsnap -c restore --restore
snaptovol --hanasid $_system(source_db, id)_$
--configfile=/mnt/software/AZACSNAP/snapshot_tool/azacsnap
-$_system(source_db, id)_$.json > $_logfile_$

```

Full documentation for the AzAcSnap command line options for the `restore` command can be found in the Azure documentation here: [Restore using Azure Application Consistent Snapshot tool](#). The call assumes that the json DB configuration file for the source DB can be found on the central NFS share with the following naming convention: `azacsnap-<source DB SID>.json`, (for example, `azacsnap-P01.json` in this lab environment).



Because the output of the AzAcSnap command cannot be changed, the default `LSC: success` message cannot be used for this task. Therefore, the string `Example mount instructions` from the AzAcSnap output is used as a successful return code. In the 5.0 GA version of AzAcSnap, this output is only generated if the cloning process was successful.

The following figure shows the AzAcSnap restore to new volume success message.

Task: ZAZACSNAPRESTORE Version: 0	
<b>Configuration Data</b>	
Main Attributes	success
Comment	Example mount instructions
Category	
Execution Attributes	
Parameters	
Return Codes	
Code	

### ZSCCOPYMOUNT

This task mounts the new Azure NetApp Files data volume on the OS of the target DB. The code section of this task contains the following text:

```

$_include_tool(unix_header.sh)_$
sudo /mnt/software/scripts/sc-system-refresh/sc-system-refresh.sh mount
$_system(target_db, id)_$ > $_logfile_$

```

The `sc-system-refresh.sh` script is used again, passing the `mount` command and the target DB SID.

### ZSCCOPYRECOVER

This task performs an SAP HANA database recovery of the system database and the tenant database based on the restored (cloned) Snapshot copy. The recovery option used here is to specific database backup, such as no additional logs, are applied for forward recovery. Therefore, the recovery time is very short (a few minutes at most). The runtime of this operation is determined by the startup of the SAP HANA database that

happens automatically after the recovery process. To speed up the startup time, the throughput of the Azure NetApp Files data volume can be increased temporarily if needed as described in this Azure documentation: [Dynamically increasing or decreasing volume quota](#). The code section of this task contains the following text:

```

$_include_tool(unix_header.sh)_$
sudo /mnt/software/scripts/sc-system-refresh/sc-system-refresh.sh recover
$_system(target_db, id)_$ > $_logfile_$

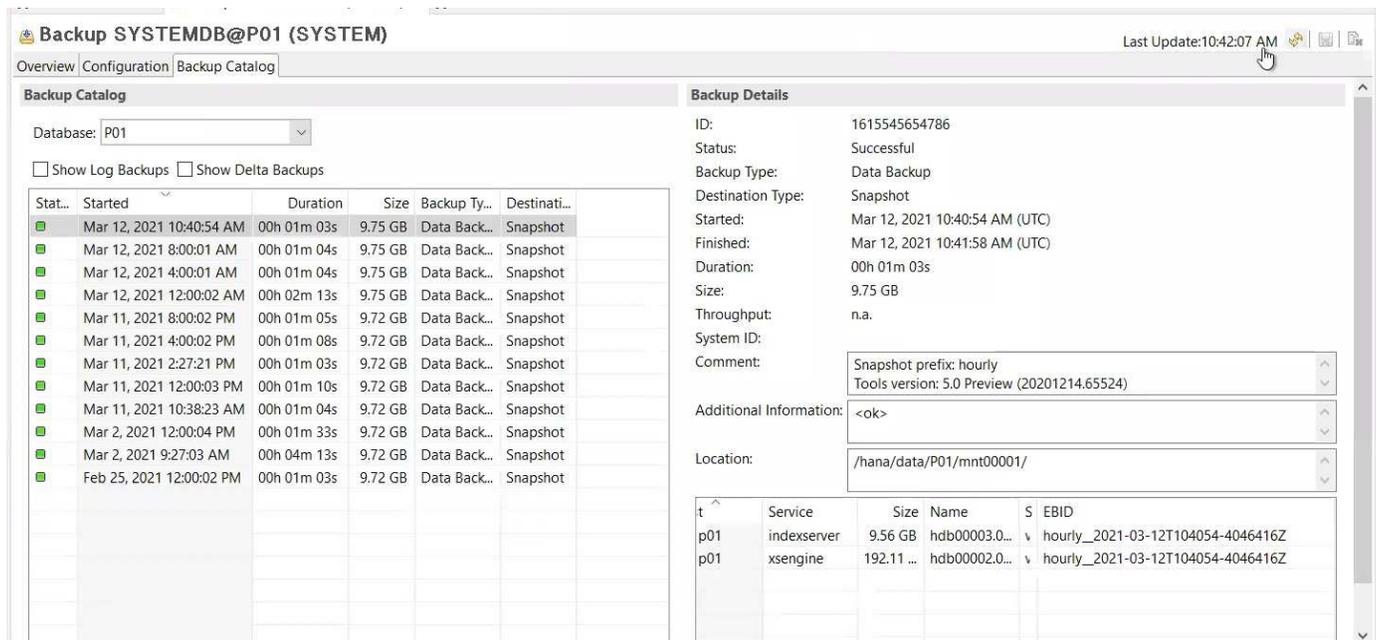
```

This script is used again with the `recover` command and the target DB SID.

### SAP HANA system refresh operation

In this section a sample refresh operation of lab systems shows the main steps of this workflow.

Regular and on-demand Snapshot copies have been created for the P01 source database as listed in the backup catalog.



For the refresh operation, the latest backup from March 12th was used. In the backup details section, the external backup ID (EBID) for this backup is listed. This is the Snapshot copy name of the corresponding Snapshot copy backup on the Azure NetApp Files data volume as shown in the following picture.

(mcScott-EastUS/mcScott-Premium/p01-data-mnt00001) | ...

+ Add snapshot Refresh

Search snapshots

Name	Location	Created
hourly_2021-02-25T120001-8350005Z	East US	02/25/2021, 11:59:37 AM
offline-20210226	East US	02/26/2021, 01:09:40 PM
hourly_2021-03-02T092702-8909509Z	East US	03/02/2021, 09:27:20 AM
hourly_2021-03-02T120003-4067821Z	East US	03/02/2021, 11:59:38 AM
hourly_2021-03-11T103823-2185089Z	East US	03/11/2021, 10:37:55 AM
hourly_2021-03-11T120003-0695010Z	East US	03/11/2021, 11:59:23 AM
hourly_2021-03-11T142720-7544262Z	East US	03/11/2021, 02:26:35 PM
hourly_2021-03-11T160002-4458098Z	East US	03/11/2021, 03:59:17 PM
hourly_2021-03-11T200001-9577603Z	East US	03/11/2021, 07:59:17 PM
hourly_2021-03-12T000001-7550954Z	East US	03/11/2021, 11:59:51 PM
hourly_2021-03-12T040001-5101399Z	East US	03/12/2021, 03:59:16 AM
hourly_2021-03-12T080001-5742724Z	East US	03/12/2021, 07:59:34 AM
hourly_2021-03-12T104054-4046416Z	East US	03/12/2021, 10:40:26 AM

1615545654786  
 Successful  
 Data Backup  
 Snapshot  
 Mar 12, 2021 10:40:54 AM (UTC)  
 Mar 12, 2021 10:41:58 AM (UTC)  
 00h 01m 03s  
 9.75 GB  
 n.a.

Snapshot prefix: hourly  
 Tools version: 5.0 Preview (20201214.65524)

Location: /hana/data/P01/mnt00001/

Size	Name	EBID
9.56 GB	hdb00003.0...	hourly_2021-03-12T104054-4046416Z
192.11 ...	hdb00002.0...	hourly_2021-03-12T104054-4046416Z

To start the refresh operation, select the correct configuration in the LSC GUI, and then click Start Execution.

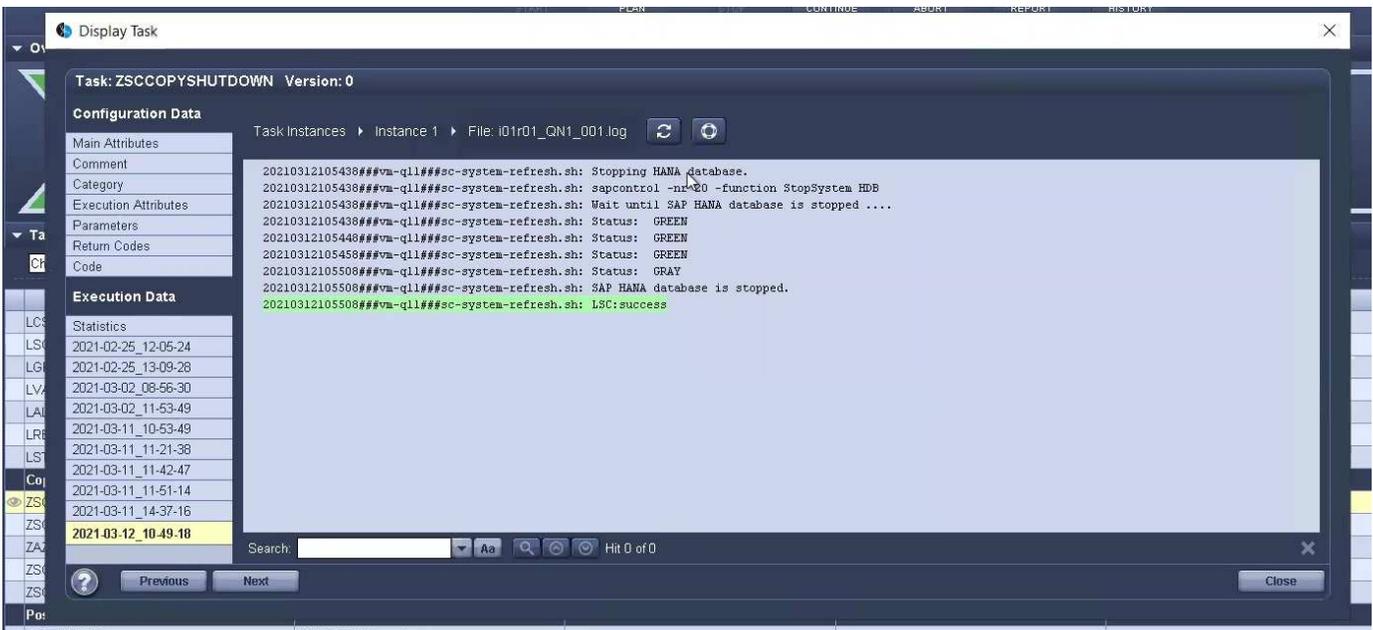
The screenshot shows the Libelle SystemCopy GUI with a task execution progress bar at 100% for 'check', 'pre', 'copy', and 'post' phases. A 'Start Execution' dialog box is open, showing 'Execute Start Checks' and 'Perform only the start checks' options. The 'Start Execution' button is highlighted, with a tooltip that says 'The execution will be started immediately.' Below the dialog, a table shows the execution progress for various tasks.

Task	End time	Duration	Progress
Execute Start Checks	PM	00:00:04	100%
Perform only the start checks	PM	00:00:03	100%
Start Execution	PM	00:00:03	100%
Task 1	PM	00:00:04	100%
Task 2	PM	00:00:03	100%
Task 3	PM	00:00:02	100%
Task 4	PM	00:00:02	100%
Task 5	PM	00:00:01	100%
Task 6	PM	00:00:02	100%
Task 7	PM	00:00:05	100%
Task 8	PM	00:00:01	100%
Task 9	PM	00:00:03	100%
Task 10	PM	00:00:01	100%

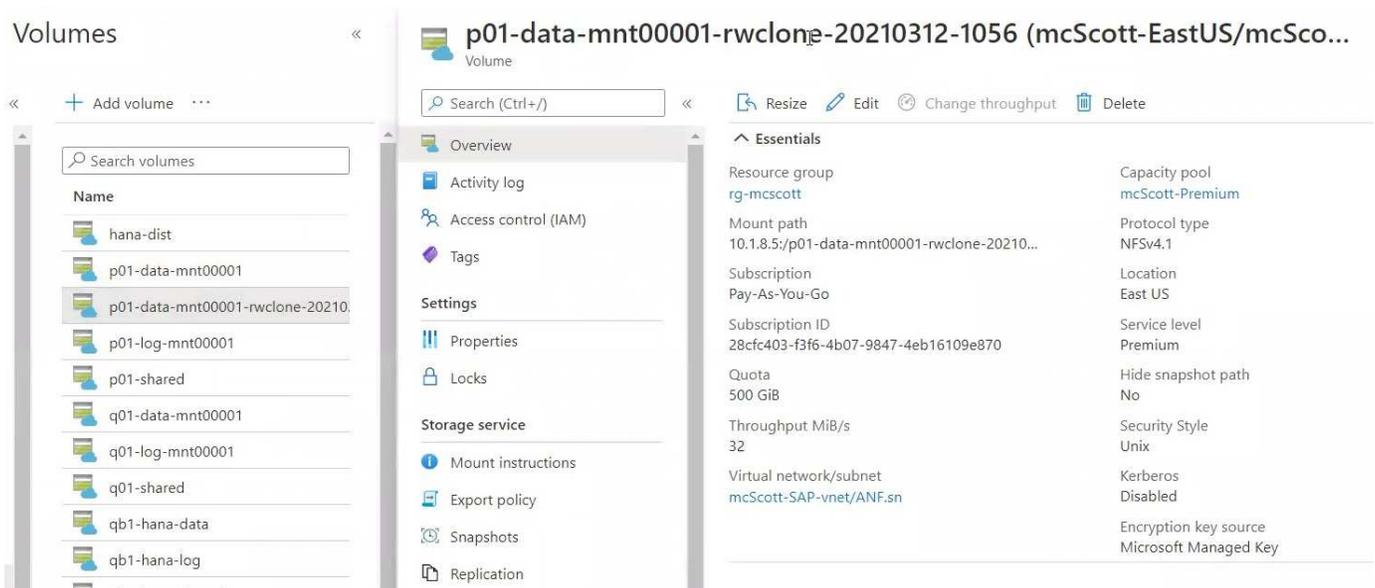
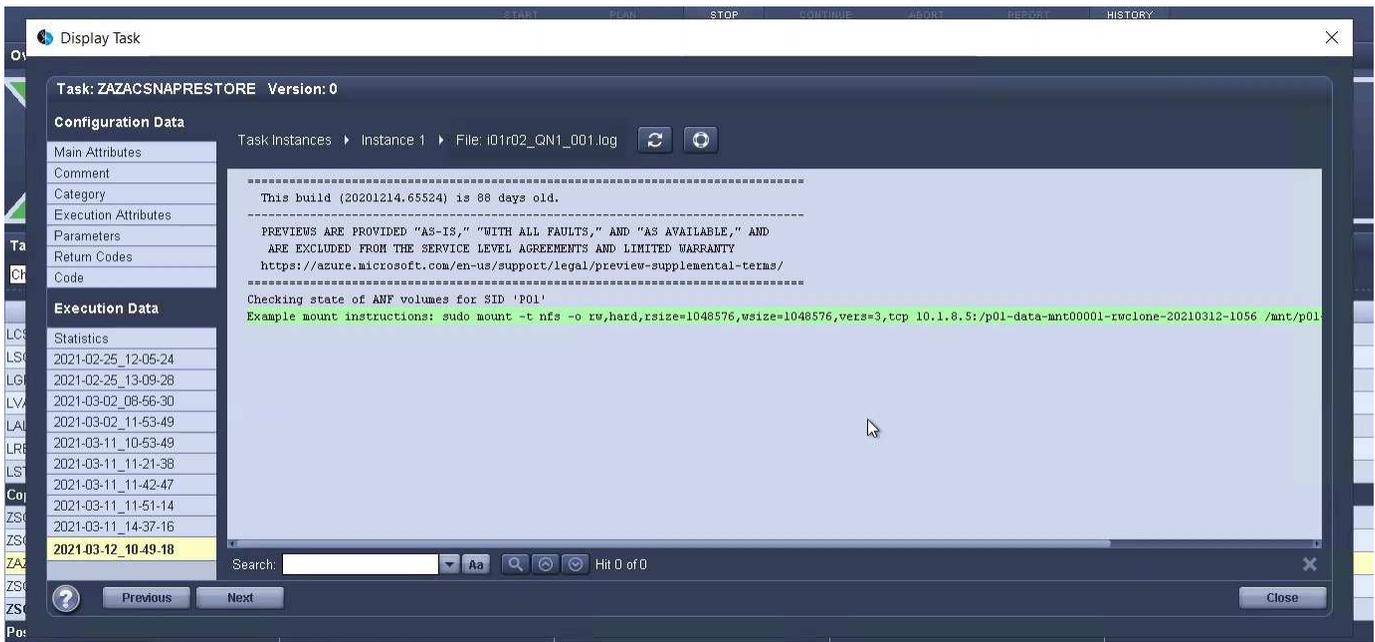
LSC starts to execute the tasks of the Check phase followed by the configured tasks of the Pre phase.



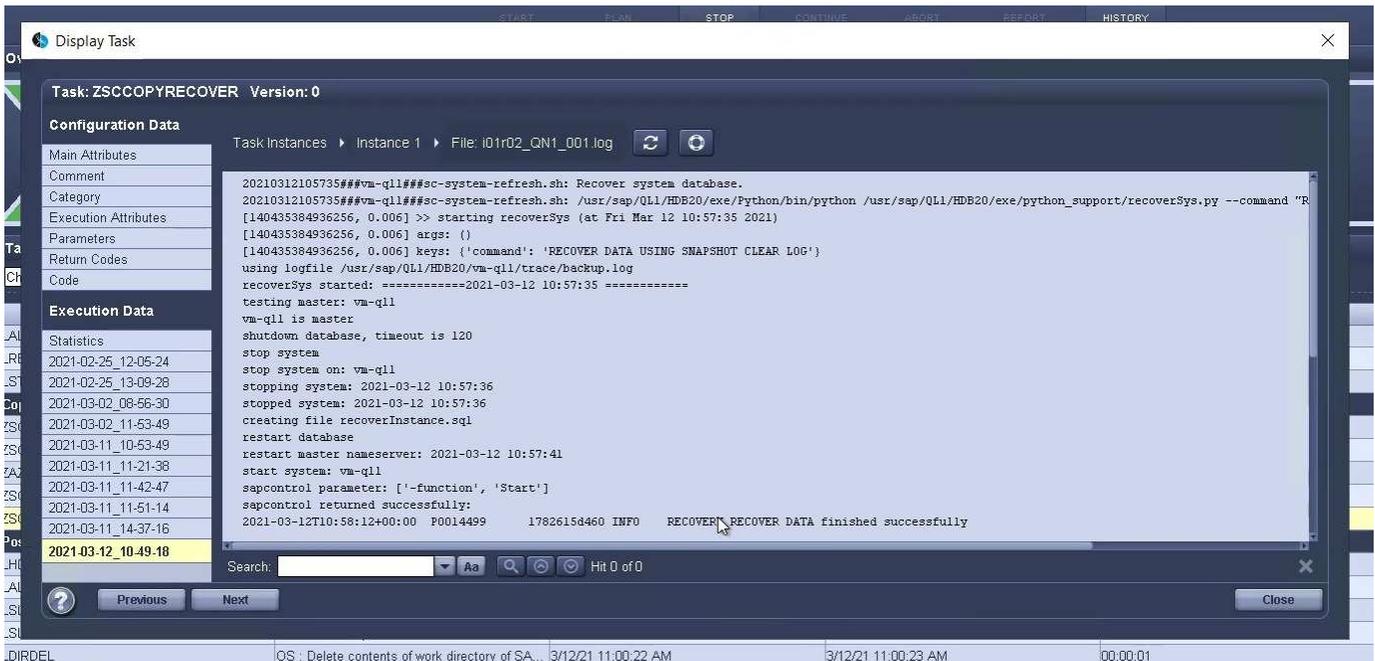
As the last step of the Pre phase, the target SAP system is stopped. In the following Copy phase, the steps described in the previous section are executed. First, the target SAP HANA database is stopped, and the old Azure NetApp Files volume is unmounted from the OS.



The ZAZACSNAPRESTORE task then creates a new volume as a clone from the existing Snapshot copy of the P01 system. The following two pictures show the logs of the task in the LSC GUI and the cloned Azure NetApp Files volume in the Azure portal.



This new volume is then mounted on the target DB host and the system database and the tenant database are recovered using the containing Snapshot copy. After successful recovery, the SAP HANA database is started automatically. This startup of the SAP HANA database occupies most of the time of the Copy phase. The remaining steps typically finish in a few seconds to a few minutes, regardless of the size of the database. The following picture shows how the system database is recovered using SAP- provided python recovery scripts.



After the Copy phase, LSC continues with all the defined steps of the Post phase. When the System Refresh process finishes completely, the target system is up and running again and fully usable. With this lab system, the total runtime for the SAP system refresh was roughly 25 minutes, of which the Copy phase consumed just under 5 minutes.



### Where to find additional information and version history

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

- NetApp Product Documentation

<https://docs.netapp.com>

**Version history**

<b>Version</b>	<b>Date</b>	<b>Document Version History</b>
Version 1.0	April 2022	Initial release.

# Solution Briefs

## SB-3978: Lifecycle Management for SAP HANA

NetApp is addressing the challenges of the slow implementation of features, a lack of automation, and a loss of productivity by providing a lifecycle management solution that is fully integrated into the tools that SAP administrators use for day-to-day operations, such as SAP Landscape Management (SAP LaMa). The goal is to simplify the provisioning workflow from preprocessing to postprocessing, including all the software and storage layer tasks needed to create a copy of the production system. With this solution, administrators can create a development and test environment in a couple of mouse clicks, which results in improved lifecycle management.

<https://www.netapp.com/pdf.html?item=/media/6996-sb-3978pdf.pdf>

## SB-3965: Backup and Recovery for SAP HANA

### The challenge

With SAP HANA backup and restore operations, your organization faces the following challenges:

- Long backup operations with performance degradation on production SAP systems
- Unacceptable system downtime due to long restore and recovery operations
- Shrinking backup windows because of the criticality of the applications
- The need for a flexible solution to mitigate logical corruption

### The Solution

With NetApp® storage solutions that run NetApp ONTAP® data management software, in combination with NetApp SnapCenter® data protection software, you can meet all those challenges. And with the NetApp Snapshot™ technology that is included in ONTAP software, you can create backups or execute restore operations of any size dataset in a matter of seconds. SAP HANA supports the use of storage-based Snapshot copies as a valid backup operation.

<https://www.netapp.com/pdf.html?item=/media/6997-sb-3965pdf.pdf>

## SB-3968: Disaster Recovery for SAP HANA

### The challenge

Business continuity is essential in IT organizations. They must be able to provide high availability services for the mission-critical applications that their customers require to run their businesses. Otherwise, their customers will face productivity decrease, and eCommerce organizations could face a direct impact on their revenue.

## The Solution

NetApp has developed a full portfolio of technologies and tools to help IT organizations build or adapt their disaster recovery plans to respond to all business demands:

<https://www.netapp.com/pdf.html?item=/media/6998-sb-3968pdf.pdf>

## SB-4292: SAP automation with Ansible

This document focuses on integrating NetApp® storage systems—whether they are operated on premises, in a public cloud infrastructure-as-a-service (IaaS) environment, or in hybrid cloud—into SAP Landscape Management (LaMa) by using Ansible Playbooks and custom scripts.

### Solution overview

SAP systems are very complex. But for the companies that use SAP, these systems are central to their business processes. By automating recurring daily operational tasks, SAP system administrators can manage more systems with less effort, produce repeatable results, and reducing human error.

This document focuses on integrating NetApp® storage systems—whether they are operated on premises, in a public cloud infrastructure-as-a-service (IaaS) environment, or in hybrid cloud—into SAP Landscape Management (LaMa) by using Ansible Playbooks and custom scripts. This integration enables SAP administrators to speed up SAP system refresh tasks by using NetApp Snapshot™ and NetApp FlexClone® technology.

### Target audience

This document is targeted to SAP system administrators who haven't had much (or any) experience with Ansible automation. It should help you get started with Ansible, run your first playbooks, and configure and run your first SAP LaMa-based system refresh operation.

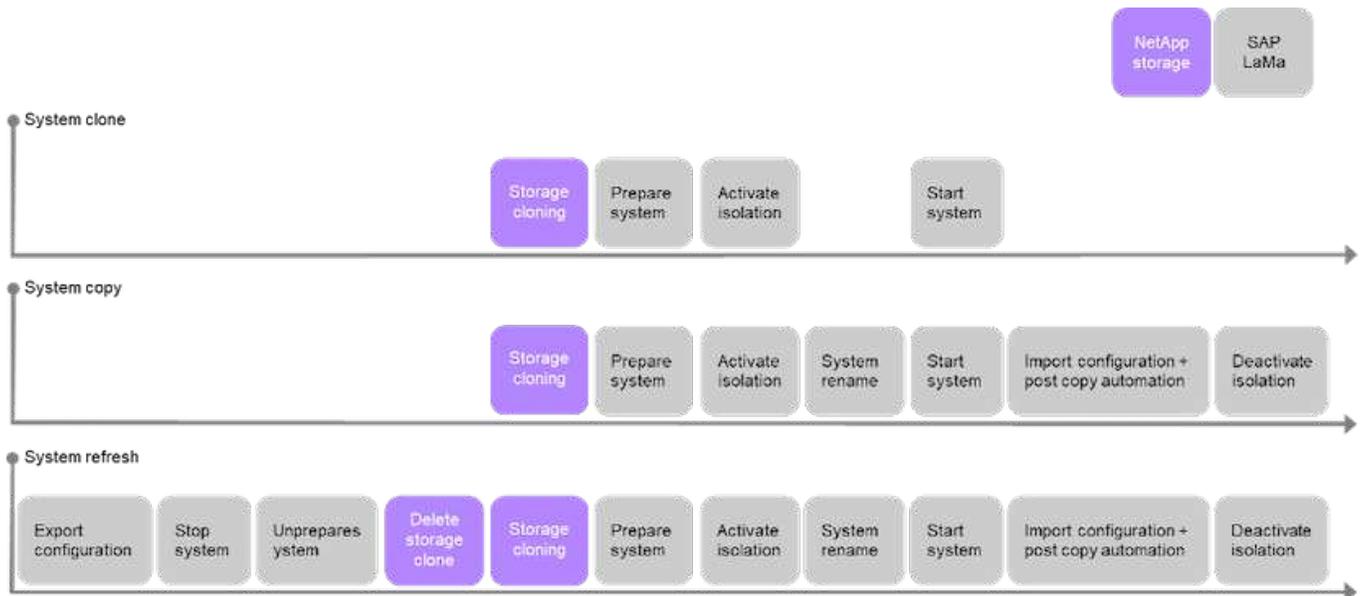
### SAP system clone, copy, and refresh scenarios

The term SAP system copy is often used as a synonym for three different processes: SAP system clone, SAP system copy, and SAP system refresh. It is important to distinguish between the different operations, because the workflows and use cases differ.

- **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 secure identifier (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.
- **SAP system copy.** An SAP system copy is a setup of a new target SAP system with data from a source SAP system. For example, the new target system could be 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 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.

The following image shows SAP system clone, copy, and refresh LaMa workflow steps that are related to NetApp storage.

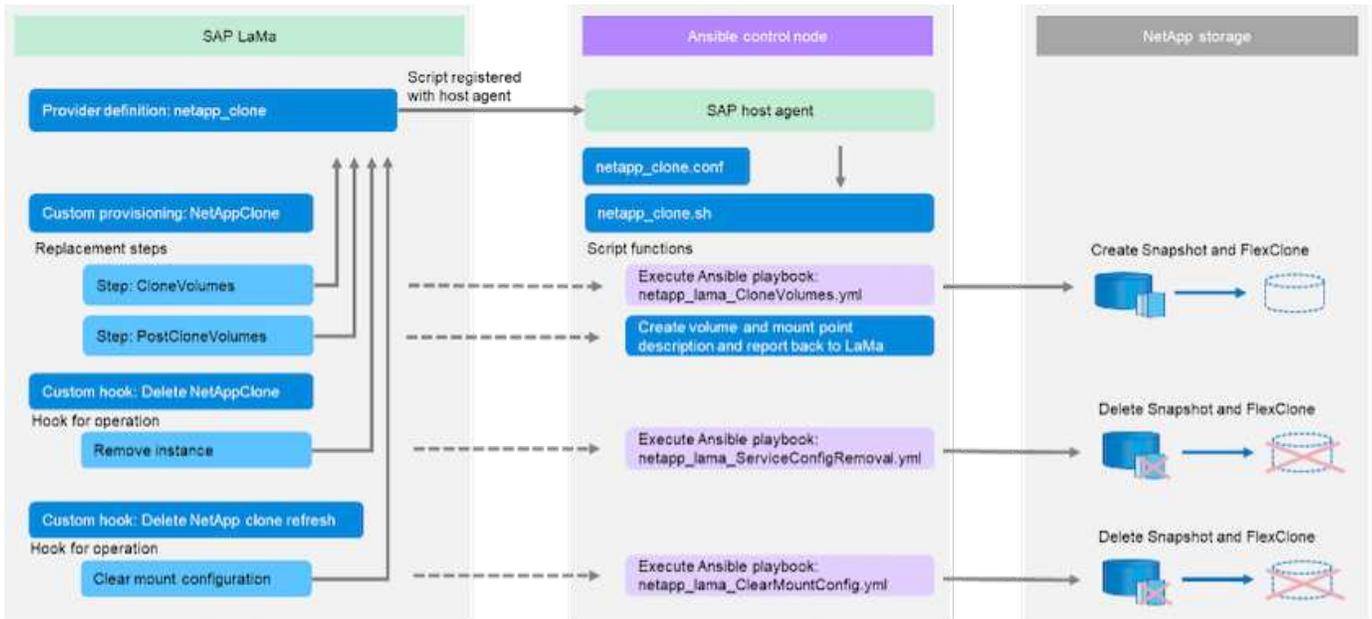


## Solution technology

The overall solution consists of these main components:

- SAP LaMa system
- NetApp storage system
- Ansible control node with installed SAP Host Agent. We recommend using Red Hat Ansible Automation Platform, because it provides additional benefits such as:
  - Using AI to generate code recommendations for automation tasks
  - Reducing manual tasks with event-driven automation
  - Being defined, consistent, and portable
  - Scaling automation across environments
  - Accelerating automation with prepackaged content
  - Tracking and managing automation with rich reporting and observability metrics
  - Creating tasks, modules, and playbooks

The following image shows how SAP LaMa and NetApp storage systems integrate through Ansible Playbooks on a dedicated Ansible host, triggered by shell scripts executed from SAP Host Agent.

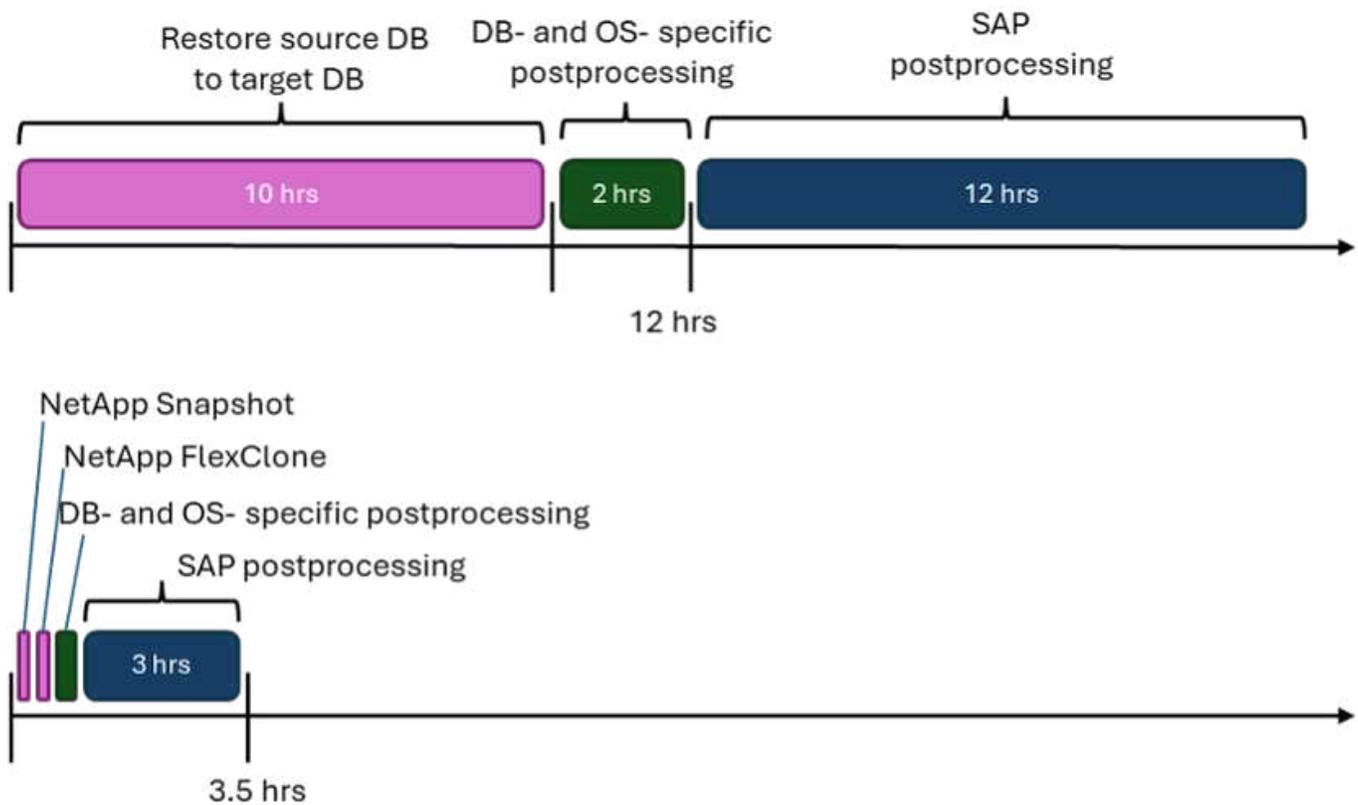


## Use case summary

There are several 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 regularly with data from the source system to make sure that testing and training are performed with the current dataset. These system refresh operations consist of multiple tasks on the infrastructure, database, and application layers, and they can take several days, depending on the level of automation.

To accelerate and automate the required tasks at the infrastructure and database layers, you can use SAP LaMa and NetApp cloning workflows. Instead of restoring a backup from the source system to the target system, SAP LaMa uses NetApp Snapshot and FlexClone technology so that required tasks to get a database started can be performed in minutes instead of hours, as shown in the following figure. The time needed for the cloning process does not depend on the size of the database; therefore, even very large systems can be created in a couple of minutes. You can further reduce the run time by automating tasks on the operating system and database layer as well as on the SAP postprocessing side.

The following image shows possible operational efficiency improvements when you use automation.



## Integrating the different technology components

To integrate SAP LaMa with NetApp storage systems by using Ansible, you need a node on which you can run Ansible Playbooks. We recommend using Ansible Automation Platform. To run shell scripts and Ansible Playbooks on this host, started from SAP LaMa, you need a running SAP Host Agent on this server. SAP Host Agent takes over the bidirectional communication with SAP LaMa and executes shell scripts that will trigger the actual playbooks.

This loosely coupled architecture gives you the freedom to start workflows from SAP LaMa and also outside SAP LaMa. Playbooks and corresponding logic needs to be configured only once and can be used for different scenarios and use cases.

## Conclusion

The combination of NetApp, SAP LaMa, and Ansible Automation Platform provides a powerful solution that can dramatically reduce the time and effort needed for the most complex and time-consuming tasks related to SAP system administration. This combination can also help avoid the configuration drift that human error can cause between the systems.

Because system refreshes, copies, clones, and disaster recovery testing are very sensitive procedures implementing such a solution can free up precious administration time. It can also reinforce the trust that the rest of the organization will have in the SAP system administrators: They will see how much easier it is to copy systems for testing or other purposes, and how much troubleshoot time can be saved.

## Where to find additional information

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

- [Automating ongoing day 1 and day 2 operations by using Ansible Playbooks for NetApp ONTAP®](#)
- [NetApp specific Ansible documentation](#)
- [NetApp ONTAP Ansible modules and full documentation](#)
- [Red Hat Ansible Automation Platform](#)

## Version history

Version	Date	Update summary
Version 0.1	03.2023	1st draft.
Version 0.2	01.2024	Review and some minor corrections
Version 0.3	06.2024	Converted to html format

# SB-4293: Automating SAP system copy, refresh, and clone workflows with ALPACA and NetApp SnapCenter

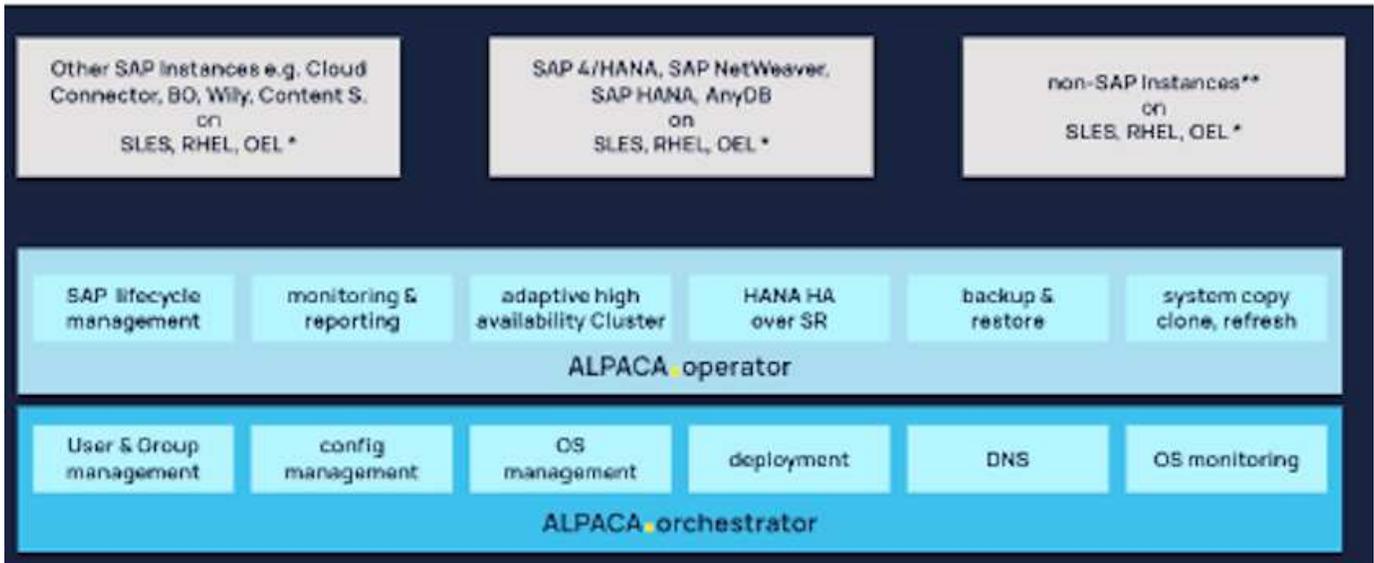
This document focuses on integrating NetApp® Snapshot™ and FlexClone® technologies into ALPACA automation workflows.

## Solution overview

SAP systems and solutions operations are very complex. However, for companies that use SAP, the systems and services are central to their business processes. By automating recurring daily operational tasks, like system copy and refresh operations, SAP system administrators can manage more systems with less effort, produce repeatable results, and reduce human error.

This document focuses on integrating NetApp® Snapshot™ and FlexClone® technologies into ALPACA automation workflows.

The Automating Landscapes Proactively—Cloud and Anywhere (ALPACA) suite is a comprehensive management interface that enables detailed oversight and monitoring across your SAP landscapes. ALPACA streamlines and expedites SAP infrastructure operations, ensuring optimal availability and transparency. It provides a comprehensive array of tools for managing the entire landscape, including infrastructure, and proactively notifies about anomalies such as service disruptions, job halts, and congestion. The suite is designed to operate seamlessly in on-premises, hybrid, and all-cloud environments, including multicloud scenarios, ensuring adaptability to any infrastructure. This module-based framework automates standard and regular SAP admin tasks as well as complex scenarios like failover during an outage. administrators/experts, operators, and managers, ALPACA gives these professionals a high degree of control and automation.



This document describes how ALPACA integrates with NetApp SnapCenter®, the tool to orchestrate Snapshot based backups, perform restores, and create FlexClone volumes. This integration allows SAP administrators to significantly accelerate SAP system daily operational tasks. NetApp Snapshot, FlexClone, and SnapRestore® technologies accelerate backup, restore, and clone operations because NetApp’s storage technology is pointer-based. This approach is fast, and it also reduces the storage overhead during clone operations, because only new and changed data (not existing data) must be written to the storage medium. This is true regardless of whether it is an on-premises NetApp storage system or a NetApp storage solution at one of the three major cloud providers.

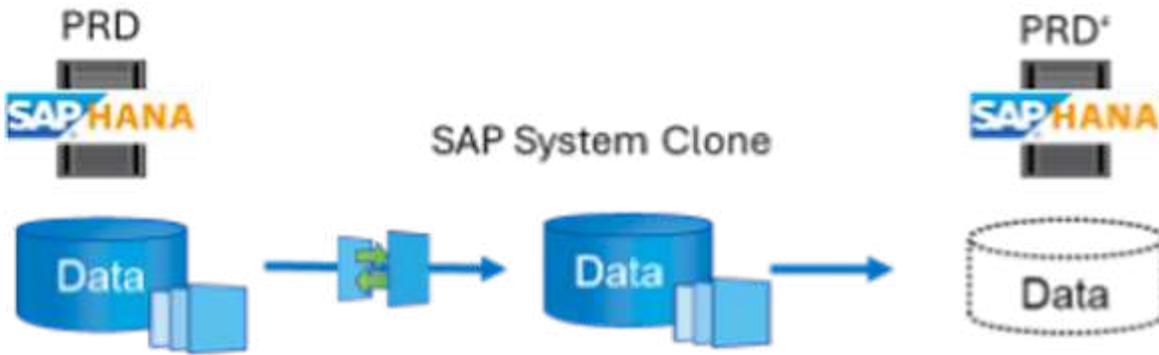
**Target audience**

This document is aimed at SAP system administrators who have carried out SAP system copies manually and would like to automate this activity with ALPACA. The intended goal of combining NetApp Snapshot and FlexClone technologies, orchestrated by NetApp SnapCenter, with ALPACA workflows is to reduce the duration of fully automated SAP system copies.

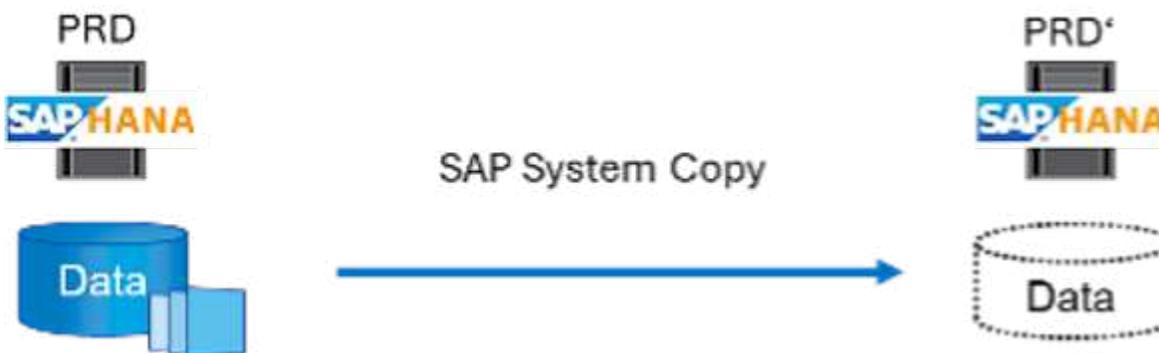
**SAP system clone, copy, and refresh scenarios**

The term SAP system copy is often used as a synonym for three different processes: SAP system clone, SAP system copy, and SAP system refresh. It is important to distinguish between these operations, because the workflows and use cases differ.

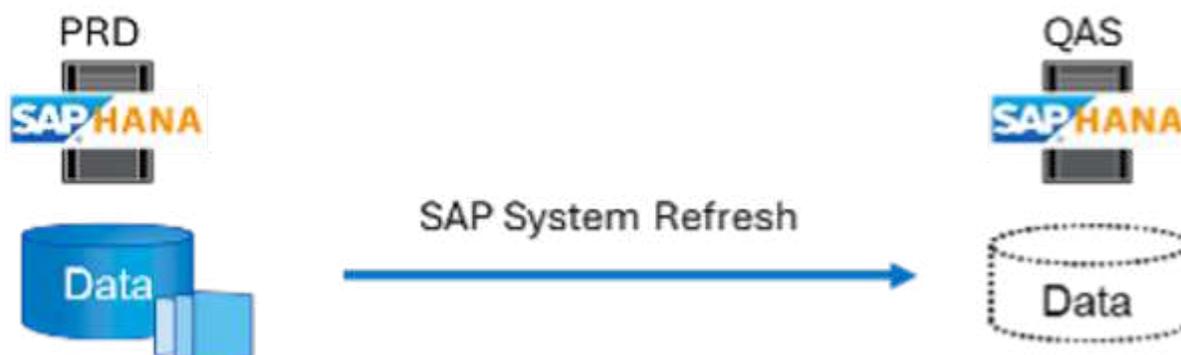
- **SAP system clone.** An SAP system clone is an identical clone of a source SAP system. SAP system clones are typically used to address logical corruption or to test disaster recovery scenarios. With a system clone operation, the host name, instance number, and secure identifier (SID) remain the same. 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.



- **SAP system copy.** An SAP system copy is a setup of a new target SAP system with data from a source SAP system. For example, the new target system could be an additional test system with data from the production system. The host name, instance number, and SID are different for the source and target systems. The new system is not isolated from the source system.



- **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 sandbox system—that is refreshed with data from the production system. The hostname, instance number, and SID are different for the source and target systems.



Even though these are three different use cases, the data management process stays the same. All three uses cases use the same underlying data management technology—NetApp Snapshot and FlexClone.

### Solution technology

The overall solution consists of these main components:

- SAP source system with installed SnapCenter agent and SnapCenter database Plug-In

- SAP target system with installed SnapCenter agent and SnapCenter database Plug-In
- ALPACA system with configured SAP source and SAP target system
- NetApp SnapCenter Server
- NetApp storage system:
  - Physical on-premises hardware: AFF-A, AFF-C, ASA-A, ASA-C, or FAS series
  - Software-defined storage on premises: ONTAP® Select
  - NetApp cloud storage:
    - Cloud Volumes ONTAP for AWS, Google Cloud, or Azure
    - Azure NetApp Files
    - Amazon FSx for NetApp ONTAP

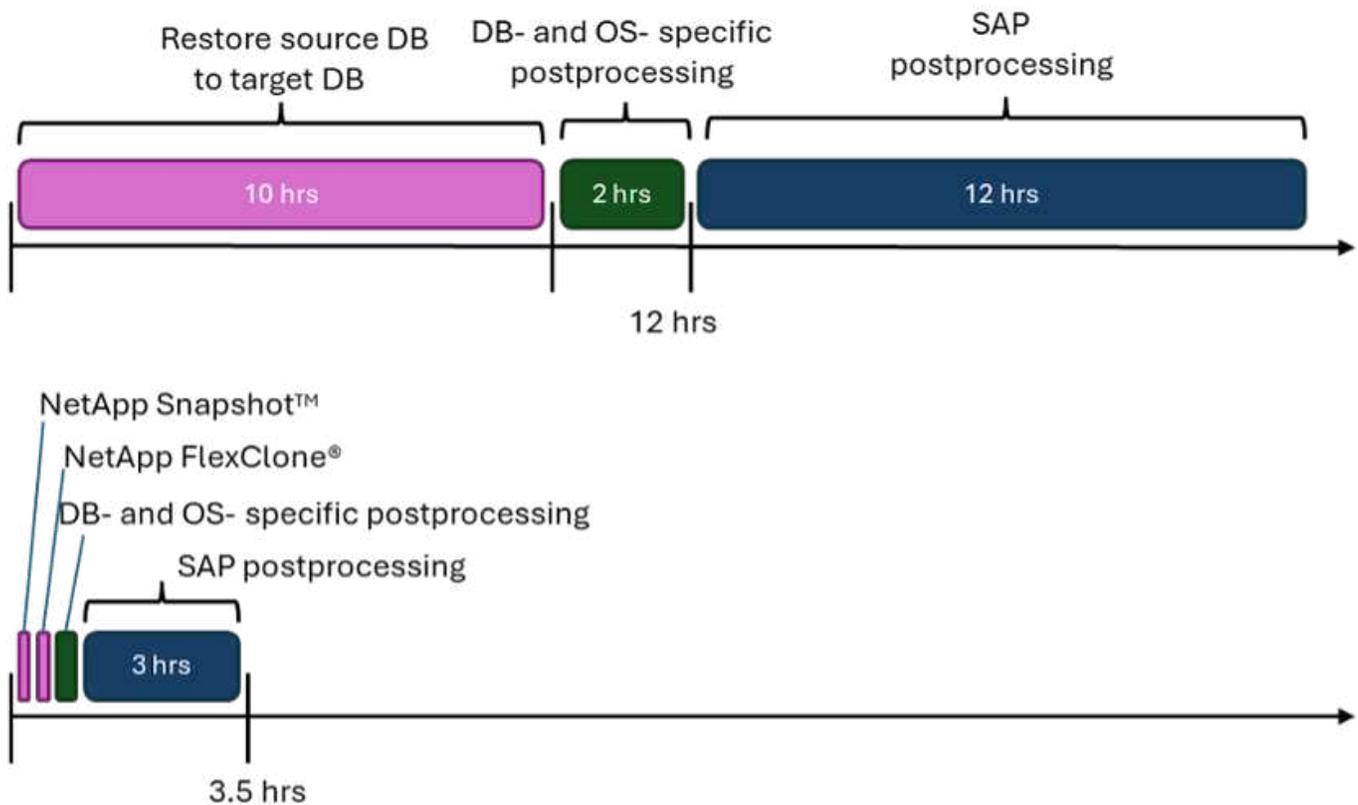
The following image shows the ALPACA server, the NetApp SnapCenter server, the NetApp storage system, the SAP source and SAP target systems, and how everything is integrated. The goal is to make the integration as flexible as possible by using the SnapCenter REST API to ensure maximum reuse of configuration work that has already been done inside existing components.

### **Use case summary**

There are several 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 regularly with data from the source system to make sure that testing and training are performed with the current dataset. These system refresh operations consist of multiple tasks on the infrastructure, database, and application layers, and they can take several days, depending on the level of automation.

To speed up operations, automate tasks, and eliminate human error at the infrastructure, database, and application levels, you can use ALPACA Workflows. Instead of restoring a backup from the source system to the target system, which is time consuming and involves high resource consumption, this integration uses NetApp Snapshot and FlexClone technologies. All the tasks required to spin up a database are completed in minutes rather than hours. The time required for the cloning process does not depend on the size of the database; therefore, even very large systems can be created in just a few minutes. ALPACA further reduces run time by automating tasks at the operating system and database levels as well as on the SAP postprocessing side.

The following image shows possible operational efficiency improvements when you use automation.



### Integrating the technology components

The actual integration of SnapCenter in an ALPACA workflow consists of using shell scripts to access the NetApp SnapCenter REST API. This REST API-based integration creates a Snapshot copy of the SAP source system, creates a FlexClone volume, and mounts it onto the SAP target system. Storage and SAP administrators know how to develop scripts that are triggered by SnapCenter and executed by the SnapCenter agent to automate recurring daily operation tasks. This loosely coupled architecture, which triggers SnapCenter tasks via shell scripts, enables them to reuse their existing automation procedures to achieve the desired results faster using ALPACA as a workflow engine for end-to-end automation.

### Conclusion

The combination of ALPACA and NetApp data management technology provides a powerful solution that can dramatically reduce the time and effort needed for the most complex and time-consuming tasks related to SAP system administration. This combination can also help avoid the configuration drift that human error can cause between the systems.

Because system refreshes, copies, clones, and disaster recovery testing are very sensitive procedures, implementing such a solution can free up precious administration time. It can also reinforce the trust that the line-of-business staff members have in the SAP system administrators. They will see how much troubleshooting time can be saved and how much easier it is to copy systems for testing or other purposes. This is true regardless of where the source and target systems are operated—on premises, in a public cloud, hybrid cloud, or hybrid multicloud.

### Where to find additional information

To learn more about the information contained in this document, review the following documents and websites:

- [ALPACA](#)
- [Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- [REST APIs supported for SnapCenter Server and plug-Ins](#)

## Version history

Version	Date	Update summary
Version 0.1	04.2024	1st draft.
Version 0.2	06.2024	Converted to html format

# SB-4294: Automating SAP system copy, refresh, and clone workflows with Avanza and NetApp SnapCenter

This document describes how Avanza integrates with the NetApp SnapCenter® platform.

## Solution overview

The operations of SAP systems and solutions are very complex. However, for companies that use SAP, these systems and services are central to their business processes. By automating recurring daily operational tasks—like system copy and refresh operations—SAP system administrators can manage more systems with less effort, produce repeatable results, and reduce human error.

This document focuses on integrating NetApp® Snapshot™ and FlexClone® technologies into Avanza automation workflows. Avanza is an IT management platform that focuses on automated management of IT operations and services. It provides solutions for monitoring, automating, and managing IT infrastructures to improve the efficiency and reliability of IT systems. Avanza allows businesses to proactively monitor their IT environments, detect issues early, and perform automated actions for troubleshooting or optimizing system performance. The platform typically integrates with other IT management tools and can be deployed in various environments such as cloud, on-premises, and hybrid infrastructures.

This document describes how Avanza integrates with the NetApp SnapCenter® platform. NetApp SnapCenter is the tool for orchestrating Snapshot based backups, performing restores, and creating FlexClone volumes. This integration allows SAP administrators to significantly speed up daily operational tasks for SAP systems by using NetApp techniques. Snapshot, FlexClone, and NetApp SnapRestore® software accelerate backup, restore, and clone operations because NetApp storage technology is pointer based. This approach is fast. It also reduces the storage overhead during clone operations, because only new and changed data is written to the storage medium, regardless of whether it is an on-premises NetApp storage system or a NetApp storage solution at one of the three major cloud providers.

## Target audience

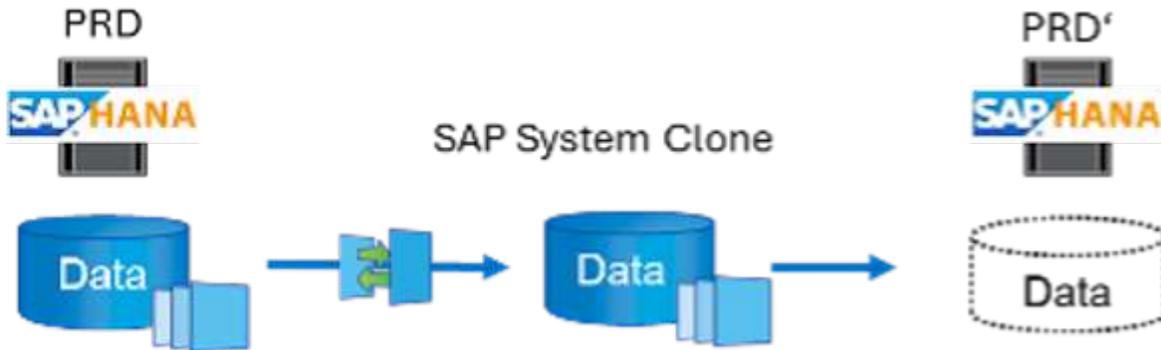
This document is aimed at SAP system administrators who have previously carried out SAP system copies manually and would like to automate this activity with Avanza. The intended goal of combining NetApp Snapshot and FlexClone technology—orchestrated by NetApp SnapCenter—with Avanza workflows is to speed up SAP system copies by fully automating them.

## SAP system clone, copy, and refresh scenarios

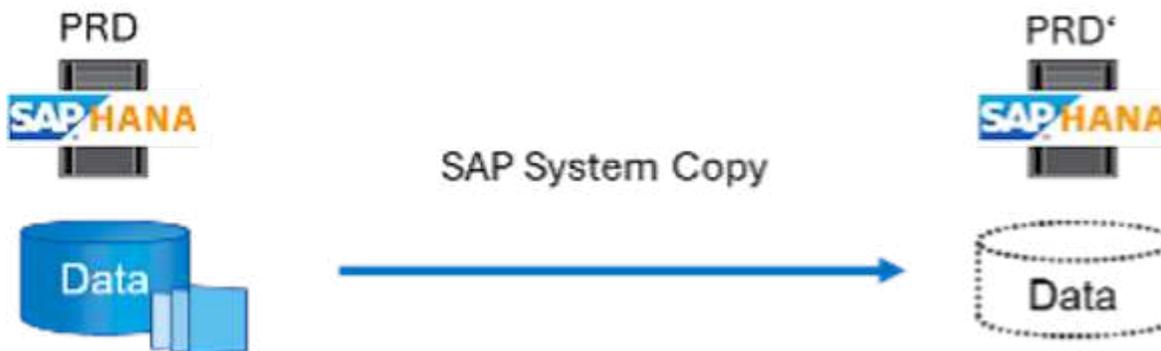
The term SAP system copy is often used as an umbrella term for three different processes: SAP system clone, SAP system copy, and SAP system refresh. It is important to distinguish between the different operations,

because the workflows and use cases differ.

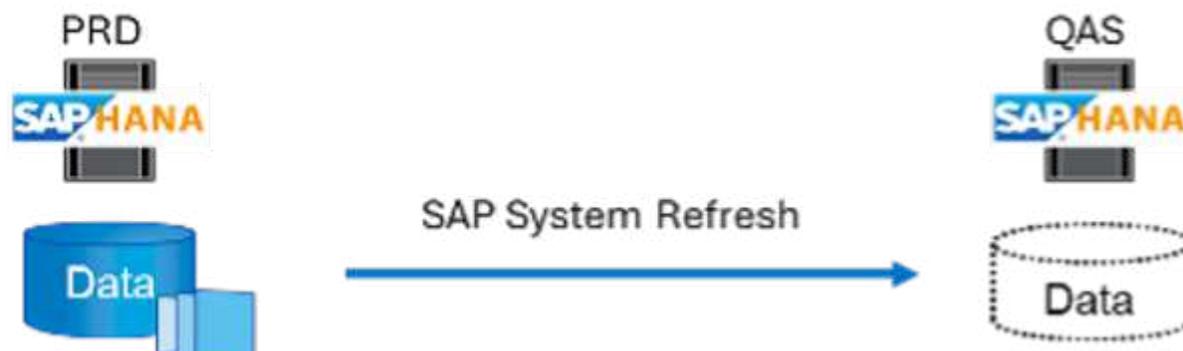
- **SAP system clone.** An SAP system clone is an identical clone of a source SAP system. SAP system clones are typically used to address logical corruption or to test disaster recovery scenarios. With a system clone operation, the host name, instance number, and secure identifier (SID) remain the same. 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.



- **SAP system copy.** An SAP system copy is a setup of a new target SAP system with data from a source SAP system. For example, the target system could be an additional test system with data from the production system. The host name, instance number, and SID are different for the source and target systems. The new system is not isolated from the source system.



- **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 sandbox system—that is refreshed with data from the production system. The host name, instance number, and SID are different for the source and target system.



Even though we have three different use cases, the data management process stays the same. All three use cases are leveraging the same underlying data management technology: NetApp Snapshot and FlexClone.

## **Solution technology**

The overall solution consists of these main components:

- SAP source system with installed SnapCenter agent and SnapCenter database plug-in
- SAP target system with installed SnapCenter agent and SnapCenter database plug-in
- Avanza system with configured SAP source and SAP target system
- NetApp SnapCenter Server
- NetApp storage system:
  - Physical on-premises hardware: NetApp AFF A-Series, AFF C-Series, ASA A-Series, ASA C-Series, or FAS series
  - Software-defined storage on premises: NetApp ONTAP® Select
  - NetApp cloud storage:
    - NetApp Cloud Volumes ONTAP® in AWS, Google Cloud, or Azure
    - Azure NetApp Files
    - Amazon FSx for NetApp ONTAP (AWS)

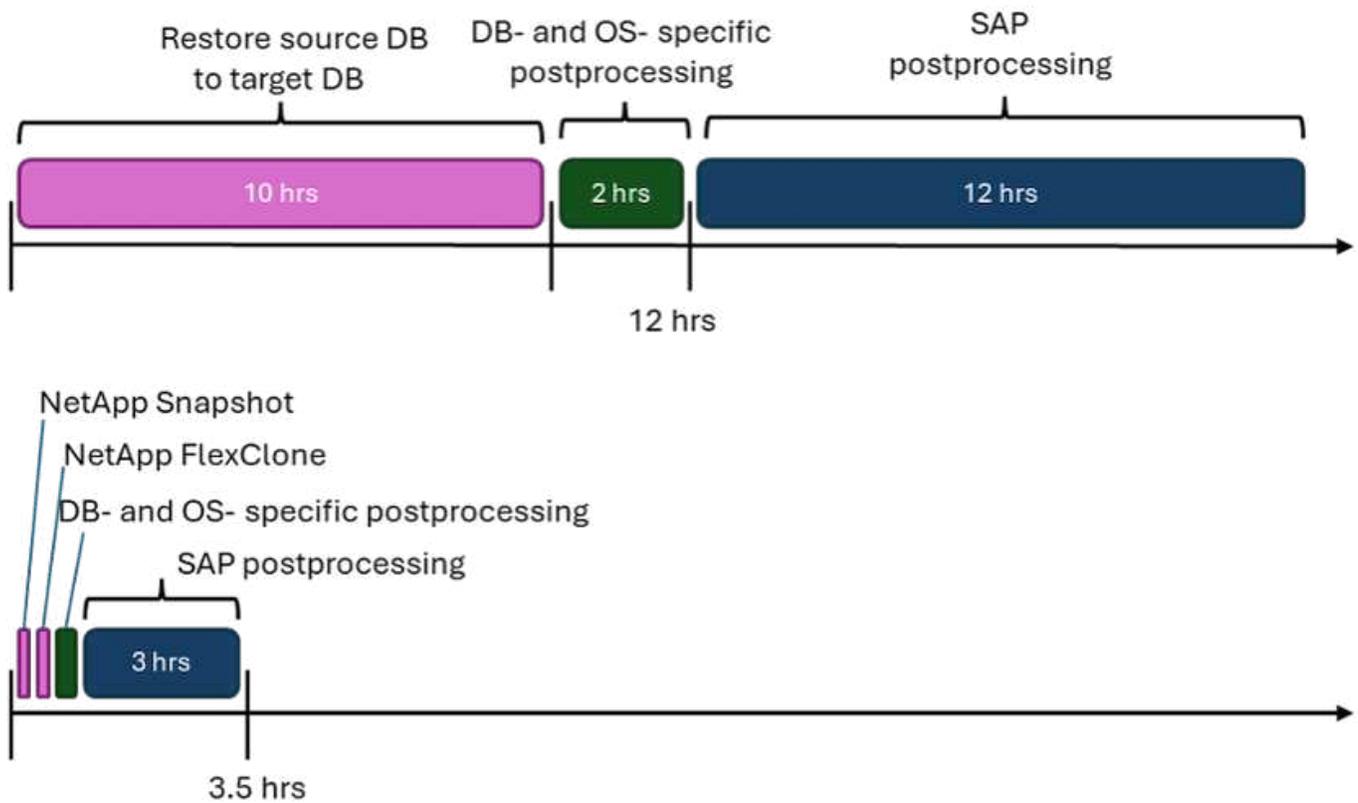
The following image shows the Avanza server, the NetApp SnapCenter Server, the NetApp storage system, the SAP source and SAP target systems, and how everything integrates. The goal was to make the integration as flexible as possible by using the SnapCenter REST API for maximum reuse of configuration work that has already been done inside existing components.

## **Use case summary**

There are several 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 regularly with data from the source system to make sure that testing and training are performed with the current dataset. These system refresh operations consist of multiple tasks on the infrastructure, database, and application layers, and they can take several days, depending on the level of automation.

To shorten time, automate operational tasks, and eliminate human error at the infrastructure, database, and application level, you can use Avanza workflows. Instead of restoring a backup from the source system to the target system—which is time consuming and involves high resource consumption—this integration uses NetApp Snapshot and FlexClone technology. All the tasks required to spin up a database are completed in minutes instead of hours. The time required for the cloning process does not depend on the size of the database; therefore, even very large systems can be created in just a few minutes. Avanza further reduces run time by automating tasks at the operating system and database level as well as on the SAP postprocessing side.

The following image shows possible operational efficiency improvements when you use automation.



### Integrating the different technology components

The actual integration of SnapCenter in an Avantra workflow consists of using JavaScript to access the NetApp SnapCenter REST API. This REST API–based integration creates a Snapshot copy of the SAP source system, creates a FlexClone volume, and mounts it onto the SAP target system.

Storage and SAP administrators have invested time and know-how to develop scripts that are triggered by SnapCenter and executed by the SnapCenter agent to automate reoccurring daily operation tasks. This loosely coupled architecture—which uses JavaScript to trigger SnapCenter tasks—enables them to reuse their existing automation procedures to achieve the desired results faster using Avantra as a workflow engine for end-to-end automation.

### Conclusion

The combination of Avantra and NetApp data management technology provides a powerful solution that can dramatically reduce the time and effort needed for the most complex and time-consuming tasks related to SAP system administration. This combination can also help avoid the configuration drift that human error can cause between the systems.

Because system refreshes, copies, clones, and disaster recovery testing are very sensitive procedures, implementing such a solution can free up precious administration time. It can also reinforce the trust that line-of-business staff members have in SAP system administrators: They will see how much troubleshooting time can be saved and how much easier it is to copy systems for testing or other purposes. The solution offers these advantages regardless of where the source and target systems are operated—on premises, in a public cloud, or in a hybrid or hybrid multicloud environment.

## Where to find additional information

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

- [Avantra](#)
- [Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- [REST APIs supported for SnapCenter Server and plug-ins](#)

## Version history

Version	Date	Update summary
Version 0.1	03.2024	1st draft.
Version 0.2	03.2024	Integration of feedback from NetApp colleagues.
Version 0.3	04.2024	Integrated requested changes to be NetApp branding compliant
Version 0.4	06.2024	Converted to html format

# Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

## Copyright

<https://www.netapp.com/company/legal/copyright/>

## Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

## Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

## Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

## 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.