



Backup, Restore and Disaster Recovery

NetApp solutions for SAP

NetApp
December 16, 2025

Table of Contents

Backup, Restore and Disaster Recovery	1
SAP HANA backup and recovery with SnapCenter	1
Protect SAP HANA systems with SnapCenter across ONTAP, Azure NetApp Files, and FSx for ONTAP	1
Learn about SAP HANA data protection with NetApp Snapshot technology	1
Learn about the SnapCenter architecture	6
Learn about SnapCenter backup and recovery for SAP HANA	7
Learn about SnapCenter supported configurations for SAP HANA	9
Learn about SnapCenter data protection concepts and best practices	13
Learn about configuring SnapCenter for SAP HANA environments	23
Configure initial SnapCenter settings for SAP HANA	24
Configure SnapCenter resources for individual SAP HANA databases	35
Configure SnapCenter to back up non-data volumes	43
Configure SnapCenter central plug-in host for SAP HANA	46
Learn about backup operations for SAP HANA Snapshot in SnapCenter	50
Execute SAP HANA block consistency checks with SnapCenter	61
Restore and recover SAP HANA databases with SnapCenter	79
Configure advanced SnapCenter options for SAP HANA	100
SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter	103
TR-4926: SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter	103
Backup and recovery using Amazon FSx for ONTAP	103
SnapCenter architecture	108
SnapCenter configuration	113
SnapCenter backup operations	130
Backup of non-data volumes	140
Restore and recover	147
Backup replication with SnapVault	155
Where to find additional information	169
SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster	170
SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster	170
Overview SAP HANA high availability	171
Example configuration overview	174
HANA system provisioning and installation	175
SnapMirror active sync configuration	183
SnapCenter configuration	189
SnapCenter backup operations	194
SnapCenter restore and recovery	197
SAP System refresh operation	199
SnapCenter non-data volumes	199
Failover scenarios	201
Additional information and version history	205

SAP HANA data protection with SnapCenter with VMware VMFS and NetApp ASA systems	205
SAP HANA data protection with SnapCenter with VMware VMFS and NetApp ASA systems	206
Lab setup used for this document	206
HANA system provisioning and installation	207
HANA configuration	214
SnapCenter configuration	215
Backup operations	222
Restore and recovery operations	225
SAP System Refresh	229
Additional information and version history	238
SAP HANA System Replication Backup and Recovery with SnapCenter	239
TR-4719: SAP HANA System Replication - Backup and Recovery with SnapCenter	239
Storage Snapshot backups and SAP System Replication	240
SnapCenter configuration options for SAP System Replication	241
SnapCenter 4.6 configuration using a resource group	243
SnapCenter configuration with a single resource	254
Restore and recovery from a backup created at the other host	267
Where to find additional information	271
Version history	271
SAP HANA Disaster Recovery with Azure NetApp Files	271
TR-4891: SAP HANA disaster recovery with Azure NetApp Files	272
Disaster recovery solution comparison	274
ANF Cross-Region Replication with SAP HANA	278
Disaster recovery testing	290
Disaster recovery failover	303
Update history	314
TR-4646: SAP HANA Disaster Recovery with Storage Replication	314
TR-4711: SAP HANA Backup and Recovery Using NetApp Storage Systems and Commvault Software	315
SnapCenter Integration for SAP ASE Database	315
Introduction	315
Additional information and version history	328
SnapCenter Integration for IBM DB2 Database	329
Introduction	329
Example configuration overview	330
Demo Environment	330
Additional information and version history	337
SnapCenter Integration for SAP MaxDB Database	338
Introduction	338
Example configuration overview	338
Demo Environment	339
Software versions	339
MaxDB Volume Design	339
Steps to Protect Database M02	340
Prerequisites on Host	340
Prerequisites for the Database – Create Backup Templates, Enable Logbackup	340

Deploy SnapCenter Agent to Host sap-lnx25	341
Create SnapCenter Resource Configuration for Database M02	342
Sequence to Recover System M02	349
Recover Instance M02	349
Additional information and version history	355

Backup, Restore and Disaster Recovery

SAP HANA backup and recovery with SnapCenter

Protect SAP HANA systems with SnapCenter across ONTAP, Azure NetApp Files, and FSx for ONTAP

Protect SAP HANA systems with NetApp SnapCenter using Snapshot-based backups and data replication. This solution covers SnapCenter configuration and operational best practices for SAP HANA systems on ONTAP AFF and ASA systems, Azure NetApp Files, and Amazon FSx for ONTAP, including backup strategies, consistency checks, and recovery workflows.

Author: Nils Bauer, NetApp

Additional use case specific details on SAP system refresh operations and SAP HANA system replication can be found at:

- [Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)
- [SAP HANA System Replication - Backup and Recovery with SnapCenter](#)

Best practices for combining SnapCenter data protection and NetApp SnapMirror active sync are described in

- [SAP HANA data protection and high availability with SnapCenter, SnapMirror active sync and VMware Metro Storage Cluster](#)

Additional platform specific best practices documentation is available at

- [SAP HANA data protection with SnapCenter with VMware VMFS and NetApp ASA systems](#)
- [SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter](#)
- [SAP HANA data protection on Azure NetApp Files with SnapCenter \(blog and video\)](#)
- [SAP System Refresh and Cloning operations on Azure NetApp Files with SnapCenter \(blog and video\)](#)

Learn about SAP HANA data protection with NetApp Snapshot technology

Discover how NetApp Snapshot technology protects SAP HANA databases with backups that complete in minutes, regardless of database size. Learn about backup and recovery strategies using Snapshot copies, SnapRestore for fast recovery, and replication with SnapVault or Azure NetApp Files backup for secondary protection.

Companies today require continuous, uninterrupted availability for their SAP applications. They expect consistent performance levels and they need automated daily operations 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 costs to the business.

Backup and recovery using Snapshot backups

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, since all operations are executed at the storage system. Therefore, you can schedule the creation of Snapshot copies without considering peak dialog or batch activity periods. SAP on 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. A restore operation brings back the data in the file system based on the state of a backup. A recovery operation is used to roll forward the database state to a point in time using database log backups.

NetApp SnapRestore technology enables the restoration of an entire database or, alternatively, just a portion of a database, based on the currently available Snapshot backups. The restore process is 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 compared 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. The time needed for restore and recovery is significantly reduced compared to traditional streaming backups.

Because Snapshot backups are stored on the same disk system as the active online data, 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 not available. You can also use the secondary backup if it is necessary to restore a backup that is no longer available on the primary storage.

A backup to a secondary location is based on Snapshot copies created on the primary storage. Hence the data is read directly from the primary storage system without generating load on the SAP database server and its network. The primary storage communicates directly with the secondary storage and replicates the backup data to the destination by using either SnapVault or ANF backup functionality.

SnapVault and ANF backup offer significant advantages compared to traditional backups. After an initial data transfer, where all data is transferred from the source to the destination, all subsequent backups only replicate 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 only the changed blocks are stored at the destination, any additional full database backup consumes 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 database Snapshot operation. The storage Snapshot backup itself is finished in a couple of seconds independent of the HANA database size.

Backup Catalog					
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 8:00:00 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 4:40:03 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 to 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 NetApp Snapshot 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 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 have any impact on the performance of the SAP HANA database. For example, if Snapshot backups are scheduled every six hours, logs would need to be applied in the worst case for the last six hours, if the failure occurs directly before the next Snapshot would have been created. For a daily file-based backup logs for last 24 hours would need to be applied in worst case.

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.

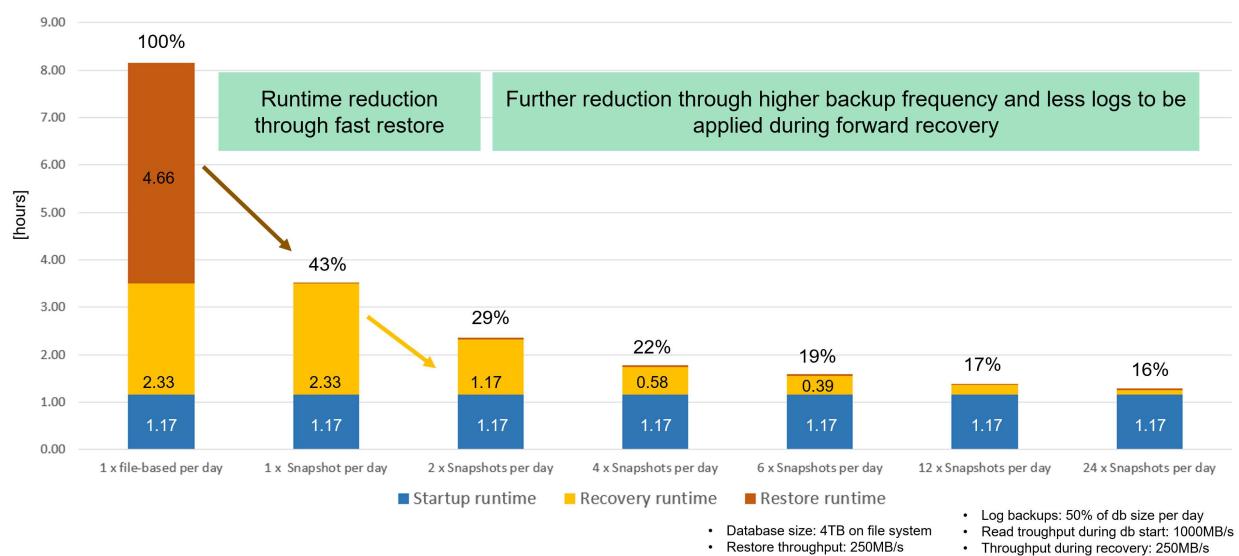
Restore and recovery sample calculation

The following figure shows a comparison between 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

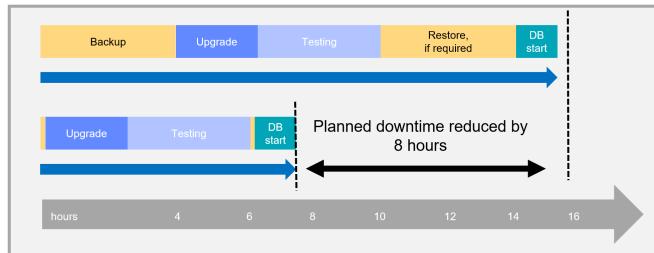
Executing backups is a critical part of any data protection strategy. Backups are scheduled on a regular basis to ensure that you can recover from system failures. This is the most obvious use case, but there are also other SAP lifecycle management tasks, where accelerating backup and recovery operations is crucial.

SAP HANA system upgrade is an example 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, or you have 8 more hours for analyzing and fixing errors by using the Snapshot-based backup and restore operations.

Another use case 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.

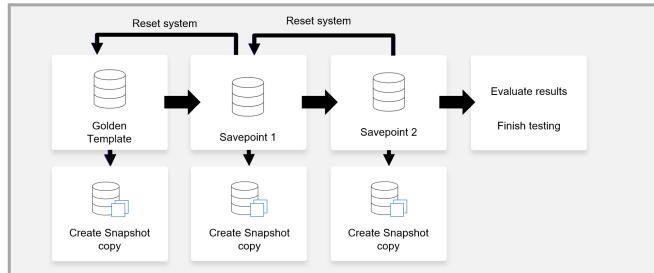
Accelerate HANA system upgrade operations

- 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. 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.

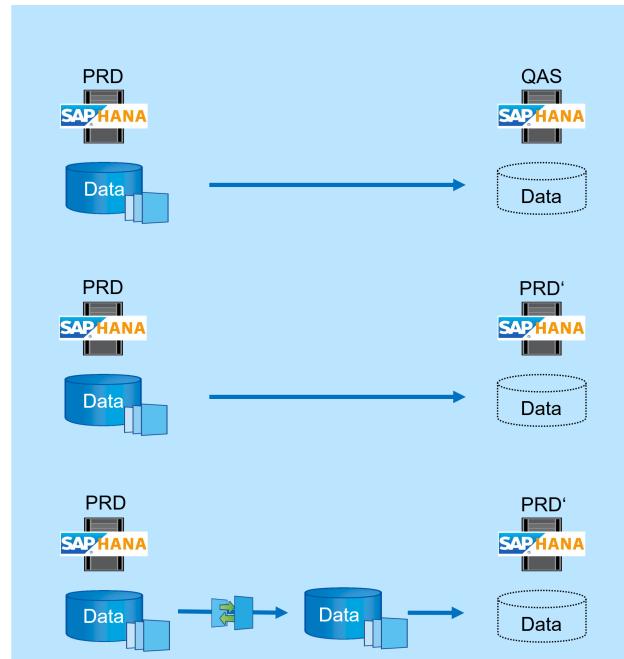
The most popular use case is the SAP System Refresh, where data from the production system needs to be copied to the test or QA system. By leveraging the ONTAP or ANF cloning feature, you can provision the volume for the test system from any Snapshot copy of the production system in a matter of seconds. The new volume then must be attached to the test system and the HANA database must be recovered.

The second use case is the creation of a repair system, which is used to address 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 got 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 ONTAP SnapMirror replication or ANF cross region 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

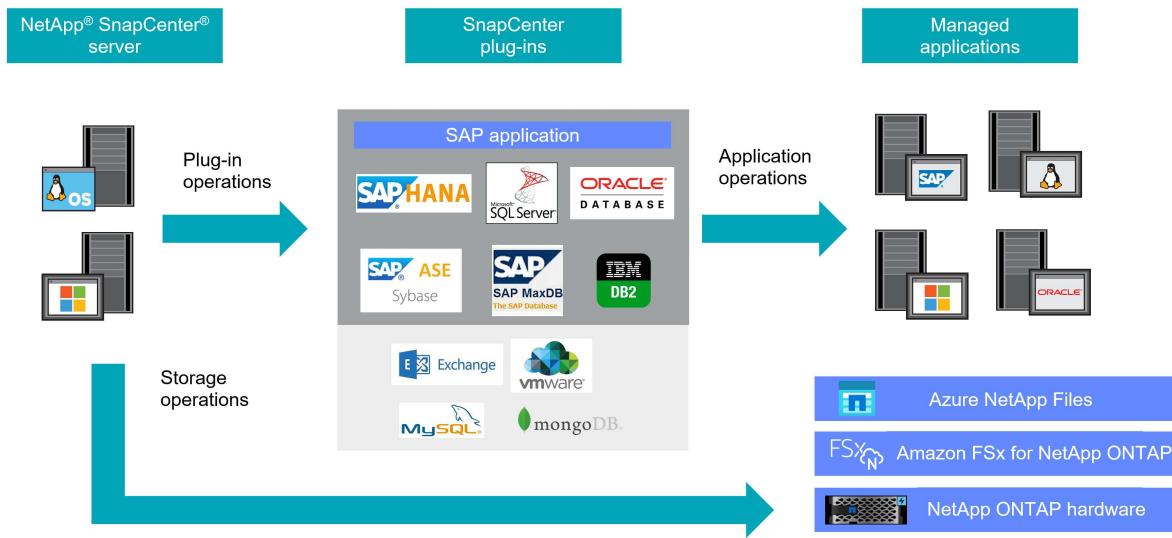


Learn about the SnapCenter architecture

Learn about the SnapCenter architecture for SAP HANA data protection, including the SnapCenter server, plug-in components, and supported storage platforms. SnapCenter provides centralized backup, restore, and clone management for SAP HANA databases on ONTAP systems, Azure NetApp Files, and FSx for ONTAP.

SnapCenter is a unified 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 operations. NetApp SnapCenter is a single tool that can be used by database and storage administrators to manage backup, restore, and cloning operations for a variety of applications and databases. SnapCenter supports NetApp ONTAP storage systems, as well as Azure NetApp Files and FSx for ONTAP. You can also use SnapCenter to replicate data between on-premises environments; between on-premises environments and the cloud; and between private, hybrid, or public clouds.

SnapCenter includes the SnapCenter server and the SnapCenter plug-ins. The plug-ins are available for various applications and infrastructure components. The SnapCenter server can run either on Windows or Linux.



Learn about SnapCenter backup and recovery for SAP HANA

SnapCenter provides comprehensive backup and recovery capabilities for SAP HANA databases using storage-based Snapshot copies, automated retention management, and integration with NetApp ONTAP, Azure NetApp Files, and FSx for NetApp ONTAP. The solution supports application-consistent database backups, non-data volume protection, block integrity checks, and replication to secondary storage using SnapVault or ANF backup.

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 check operations
 - using a file-based backup
 - using the SAP HANA hdbpersdiag tool
- Snapshot backup replication to a secondary backup location
 - using SnapVault/SnapMirror
 - using Azure NetApp Files ANF backup
- 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

Database data backups are executed by SnapCenter in combination with the SnapCenter plug-in for SAP HANA. The plug-in triggers an SAP HANA internal database snapshot so that the snapshots, which are

created on the storage system, are based on an application consistent image of the SAP HANA database.

SnapCenter enables the replication of consistent database images to a secondary backup or disaster recovery location by using SnapVault or the SnapMirror feature. Typically, different retention policies are defined for backups at primary and at secondary storage. SnapCenter handles the retention at primary storage, and ONTAP handles the retention at the secondary 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 consistency check of the persistence layer. You can execute the block consistency check from within SnapCenter either by running a file-based backup or by executing the SAP hdbpersdiag tool.

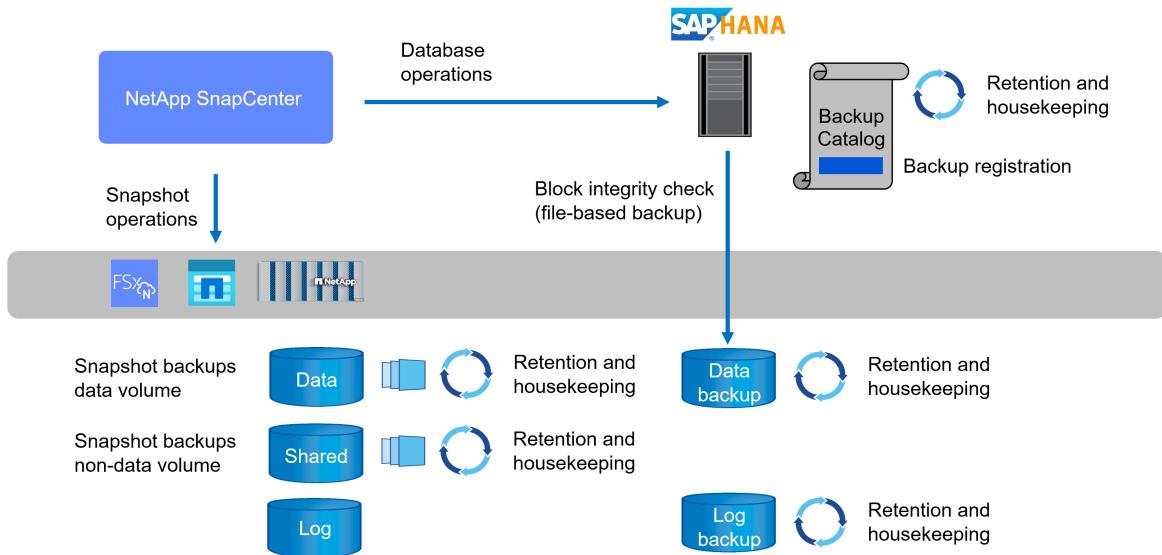
Based on your configured retention policies, SnapCenter manages the housekeeping of data file backups at the primary storage, log file backups, and the SAP HANA backup catalog.

SnapCenter handles the retention at primary storage, while ONTAP manages secondary backup retention.

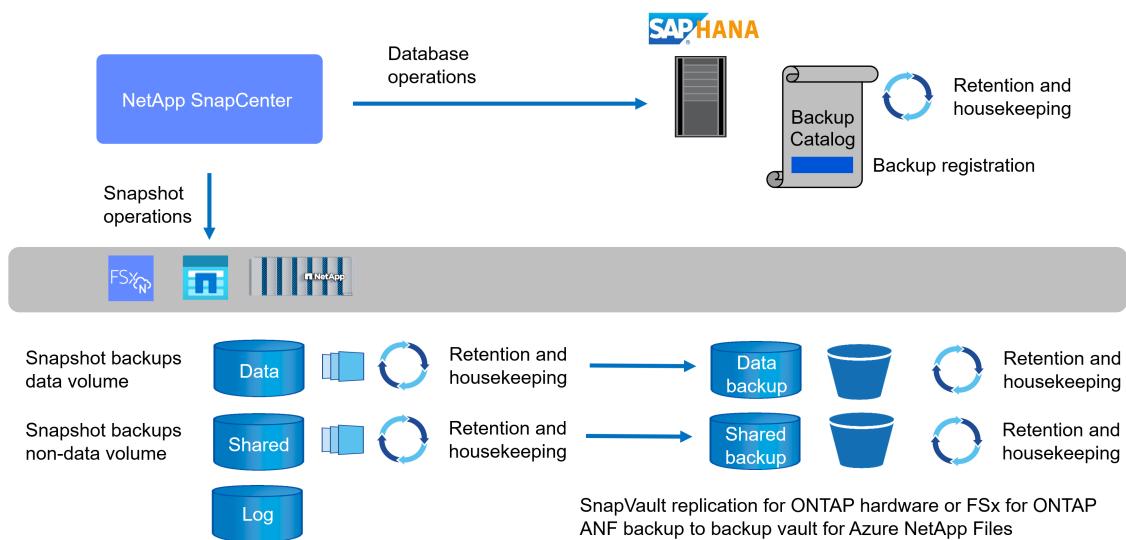
The following figure shows an overview of the SnapCenter backup and retention management operations.

When executing a storage-based Snapshot backup of the SAP HANA database, SnapCenter performs the following tasks:

- Backup operation:
 - Triggers an internal HANA database snapshot to get an application consistent image on the persistence layer
 - Creates a storage-based Snapshot backup of the data volume
 - Closes the internal HANA database snapshot, confirms or abandons the backup operation. This step registers the backup in the HANA backup catalog.
- Retention management:
 - Deletes storage Snapshot backups based on the defined retention
 - Deletes Snapshots on storage layer
 - Deletes SAP HANA backup catalog entries
 - 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



If a secondary backup is configured, either with SnapVault/SnapMirror or with ANF backup, the Snapshot created at the primary volume is replicated to the secondary backup storage. SnapCenter manages HANA backup catalog as well as log backup retention according to availability of secondary backups.



Learn about SnapCenter supported configurations for SAP HANA

SnapCenter supports a wide range of SAP HANA system architectures and deployment scenarios across on-premises and cloud storage platforms. Learn about supported SAP HANA configurations, platform combinations, storage protocols, and available backup and restore operations for each environment.

Supported SAP HANA configurations

SnapCenter supports the following HANA configurations and features:

- SAP HANA single host systems
- SAP HANA multiple host systems

- Requires a central plug-in deployment as described in ["Deployment options for SnapCenter plug-in for SAP HANA"](#).
- SAP HANA MDC systems
 - with a single or with multiple tenants
- SAP HANA systems with multiple partitions
- SAP HANA System Replication
- SAP HANA encryption (data, log, backup)

Supported platform and infrastructure configurations

SnapCenter supports the following combinations of host platforms, file systems and storage platforms.

Host platform	SAP HANA storage connection and file system	Storage platform
VMware	In-guest NFS mounts	ONTAP AFF
VMware	FC datastore with VMFS VM with XFS w/ or w/o Linux LVM	ONTAP AFF or ASA
KVM	In-guest NFS mounts	ONTAP AFF
Bare metal server	NFS mounts	ONTAP AFF
Bare metal server	FC SAN and XFS w/ or w/o Linux LVM	ONTAP AFF or ASA (*)
Azure VM	NFS mounts	Azure NetApp Files
AWS EC2	NFS mounts	FSx for ONTAP

(*): ASA support available starting with SnapCenter 6.2 release



The HANA and Linux plug-ins are only available for the Intel CPU platform. For Linux on IBM Power a central HANA plug-in deployment needs to be setup as described in ["Deployment options for SnapCenter plug-in for SAP HANA"](#).

Supported features and operations

Abbreviation explanation

- VBSR: Volume based SnapRestore
A volume based SnapRestore reverts the volume back to the state of the Snapshot.
- SFSR: Single file SnapRestore
A single file SnapRestore can be used to restore specific file(s) or LUN(s) within a volume.

See also [Types of restore operations for auto discovered SAP HANA databases](#)

ONTAP AFF and FSx for ONTAP



Only column 1 (NFS mounts) of the table below is relevant for FSx for ONTAP.

Operation	NFS mounts Bare metal or in-guest with VMware or KVM	FC SAN Bare metal	FC datastore VMware VMFS
Snapshot backup and restore operations for HANA database			
Snapshot backup	Yes	Yes	Yes
Tamperproof Snapshot	Yes	Yes	Yes
Full restore	VBSR or SFSR (selectable)	SFSR of complete LUN	Clone, mount, copy
Single tenant restore	SFSR	Clone, mount, copy	Clone, mount, copy
SnapVault backup and restore operations for HANA database			
SnapVault replication	Yes	Yes	Yes
Tamperproof Snapshot	Yes	Yes	Yes
Full restore	Yes	Yes	Clone, mount, copy
Single tenant restore	Yes	Clone, mount, copy	Clone, mount, copy
HANA recovery operation from primary Snapshot or SnapVault target			
Automated recovery MDC single tenant	Yes	Yes	Yes
Automated recovery MDC multiple tenants	No	No	No
Backup and restore non-data volumes			
Snapshot backup	Yes	Yes	Yes (*)
Restore from Snapshot	VBSR or SFSR (selectable)	SFSR of complete LUN	VBSR (*)
SnapVault replication	Yes	Yes	Yes (*)
Restore from SnapVault target	Yes	Yes	Yes (*)
SAP System Refresh			
From primary Snapshot	Yes	Yes (**)	Yes (**)
From SnapVault target	Yes	Yes (**)	Yes (**)
HA and DR			
HSR support Snapshots and SnapVault	Yes	Yes	Yes
SnapMirror replication updates with SC	Yes	Yes	Yes
SnapMirror active sync	NA	Yes	Yes

(*): No VMware integration - crash image Snapshot and full volume restore

(**): Workarounds required for SnapCenter releases < 6.2

ONTAP ASA

Operation	FC SAN Bare metal (*)	FC datastore VMware VMFS
Snapshot backup and restore operations for HANA database		
Snapshot backup	Yes	Yes
Tamperproof Snapshot	No	No
Full restore	SFSR of complete LUN	Clone, mount, copy
Single tenant restore	Clone, mount, copy	Clone, mount, copy
SnapVault backup and restore operations for HANA database		
SnapVault replication	Yes	Yes
Tamperproof Snapshot	No	No
Full restore	Yes	Clone, mount, copy
Single tenant restore	Clone, mount, copy	Clone, mount, copy
HANA recovery operation from primary Snapshot or SnapVault target		
Automated recovery MDC single tenant	Yes	Yes
Automated recovery MDC multiple tenants	No	No
Backup and restore non-data volumes		
Snapshot backup	Yes (*)	Yes (*)
Restore from Snapshot	SFSR of complete LUN (*)	SFSR of complete LUN (*)
SnapVault replication	Yes (*)	Yes (*)
Restore from SnapVault target	Yes (*)	Yes (*)
SAP System Refresh		
From primary Snapshot	Yes (**)	Yes (**)
From SnapVault target	Yes (**)	Yes (**)
HA and DR		
HSR support Snapshots and SnapVault	Yes	Yes
SnapMirror replication updates triggered by SnapCenter	Yes	Yes
SnapMirror active sync	Yes	Yes

(*): Support starting with SnapCenter 6.2 release

(**): Workarounds required for SnapCenter releases < 6.2

Azure NetApp Files

Operation	NFS mounts
Snapshot backup and restore operations for HANA database	
Snapshot backup	Yes
Tamperproof Snapshot	No
Full in-place restore	Volume revert or SFSR (selectable)
Single tenant restore	SFSR
ANF backup and restore operations for HANA database	
ANF backup replication	Yes
Tamperproof Snapshot	No
Full in-place restore	Yes
Single tenant restore	Yes
HANA recovery operation from primary Snapshot or ANF backup	
Automated recovery MDC single tenant	Yes
Automated recovery MDC multiple tenants	No
Backup and restore non-data volumes	
Snapshot backup	Yes
Restore from Snapshot	Volume revert
ANF backup replication	Yes
Full in-place restore from ANF backup	No (*)
SAP System Refresh	
From primary Snapshot	Yes
From ANF backup	Yes
HA and DR	
HSR support Snapshots and ANF backup	Yes
Cross region replication update triggered by SnapCenter	No

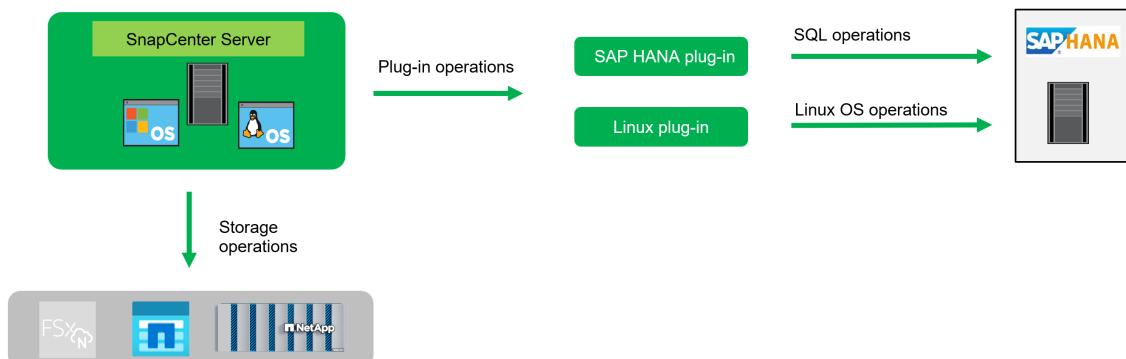
(*): With the current version, a restore operation must be done using Azure portal or CLI

Learn about SnapCenter data protection concepts and best practices

Learn about SnapCenter deployment options, data protection strategies, and backup retention management for SAP HANA environments. SnapCenter supports plug-in deployment on database hosts or central hosts, auto discovery and manual configuration, block consistency checks using file-based backups or hdbpersdiag, and comprehensive retention management across primary and secondary storage.

Deployment options for SnapCenter plug-in for SAP HANA

The following figure shows the logical view of the communication between the SnapCenter server, the SAP HANA database and the storage system. The SnapCenter server leverages the HANA and the Linux plug-ins to communicate with the HANA database and the Linux operating systems.



The recommended and default deployment option for the SnapCenter plug-ins is the installation on the HANA database host. With this deployment option, all configurations and features described in chapter SnapCenter supported configuration are valid. There are a few exceptions where the SnapCenter plug-ins can't be installed on the HANA database host but need to be configured on a central plug-in host, which could be the SnapCenter server itself. A central plug-in host is required for HANA multiple host systems or HANA systems running on the IBM Power platform. Both deployment options can also be mixed, e.g. using the SnapCenter server as a central plug-in host for a multiple host system and deploying the plug-ins on the HANA database hosts for all other single host HANA systems.

In SnapCenter a HANA resource can be either auto discovered or manually configured. A HANA system is auto discovered by default as soon as the HANA and Linux plug-ins are deployed on the database host. SnapCenter auto discovery does not support multiple HANA installations on the same host. HANA systems managed using a central plug-in host must be configured manually in SnapCenter. Also, non-data volumes are by default manual configured resources.

	Plug-in deployed at	SnapCenter resource
HANA database	Database host	Auto discovered
HANA database	Central plug-in host	Manual configured
Non-data volume	N/A	Manual configured

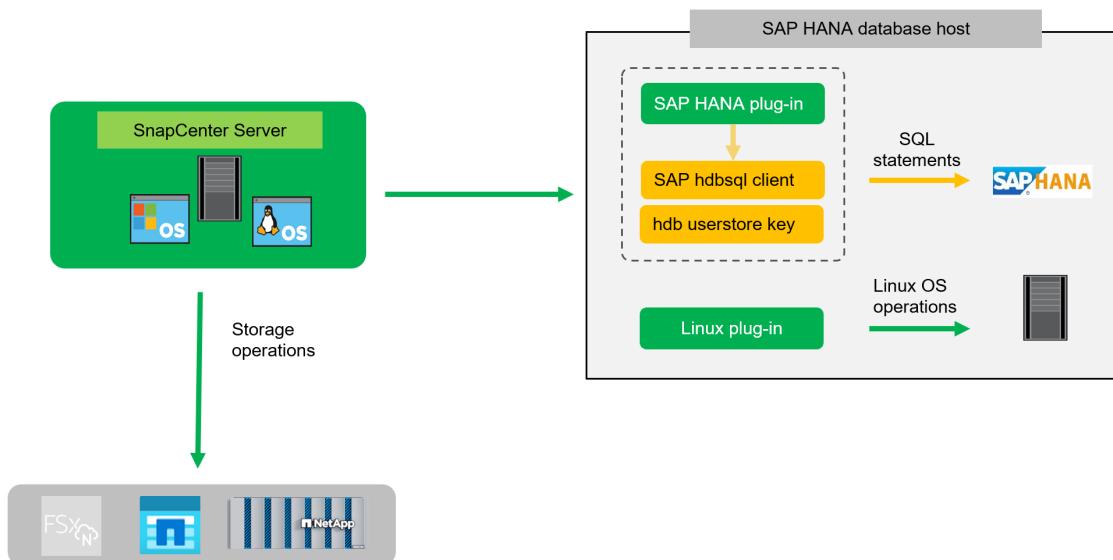
While SnapCenter supports a central plug-in deployment for HANA systems, there are limitations in platform and feature support. The following infrastructure configurations and operations are not supported for HANA systems configured with a central plug-in host:

- VMware with FC datastores
- SnapMirror active sync
- SnapCenter server high availability if used as a central plug-in host
- HANA system auto discovery
- Automated HANA database recovery
- Automated SAP System Refresh
- Single tenant restore

SnapCenter plug-in for HANA deployed on the SAP HANA database host

The SnapCenter server communicates through the HANA plug-in with the HANA databases. The HANA plug-in uses the HANA hdbsql client software to execute SQL commands to the HANA databases. The HANA hdf userstore is used to provide the user credentials, the host name, and the port information to access the HANA databases. The SnapCenter Linux plug-in is used to cover any host file system operations as well as auto discovery of file system and storage resources.

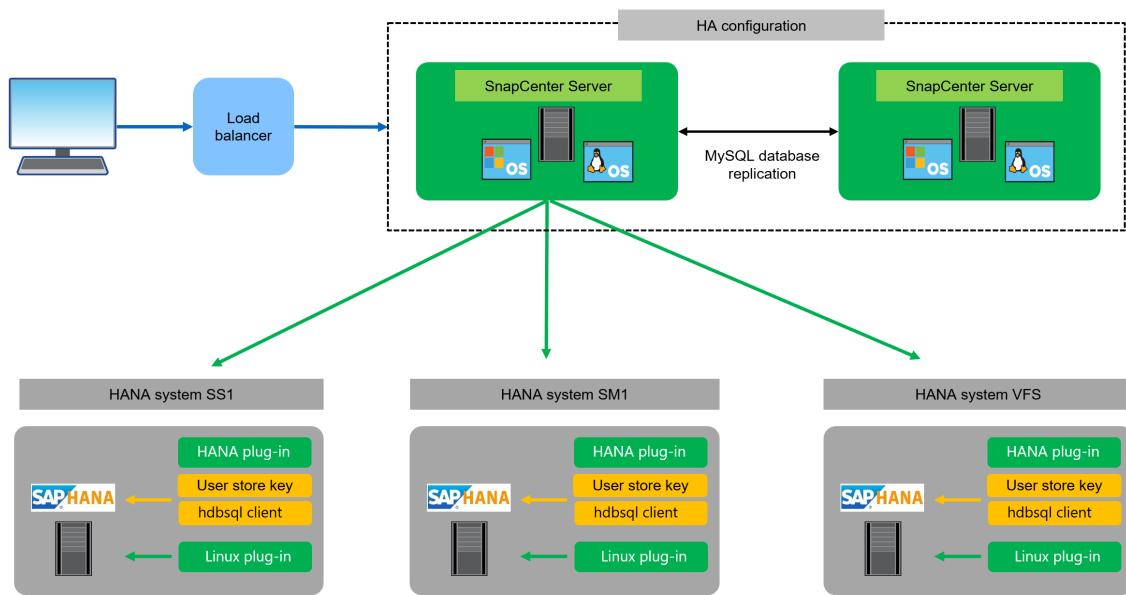
When the HANA plug-in is deployed on the HANA database host, the HANA system is auto discovered by SnapCenter and is flagged as an auto discovered resource in SnapCenter.



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 to access the SnapCenter hosts. 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. More details on SnapCenter HA can be found at [Configure SnapCenter Servers for High Availability](#).



Central plug-in host

As discussed in the chapter before, a central plug-in is required for

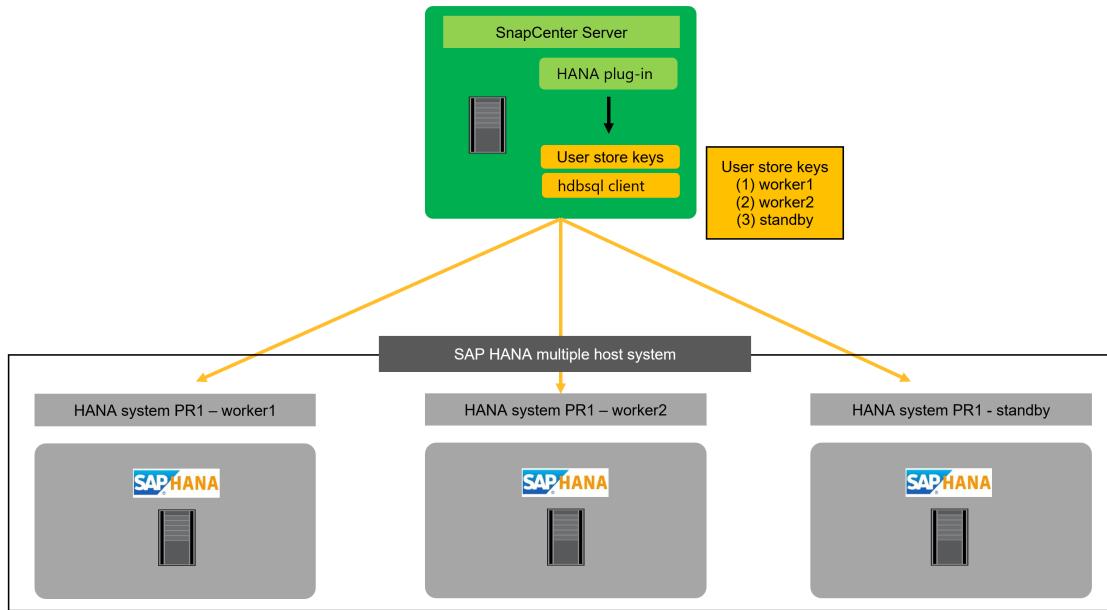
- HANA multiple host systems
- HANA systems running on IBM Power

With a central plug-in host, the HANA plug-in and the SAP HANA hdbsql client must be installed on a host outside of the HANA database hosts. This host can be any Windows or Linux host, for example the SnapCenter server.

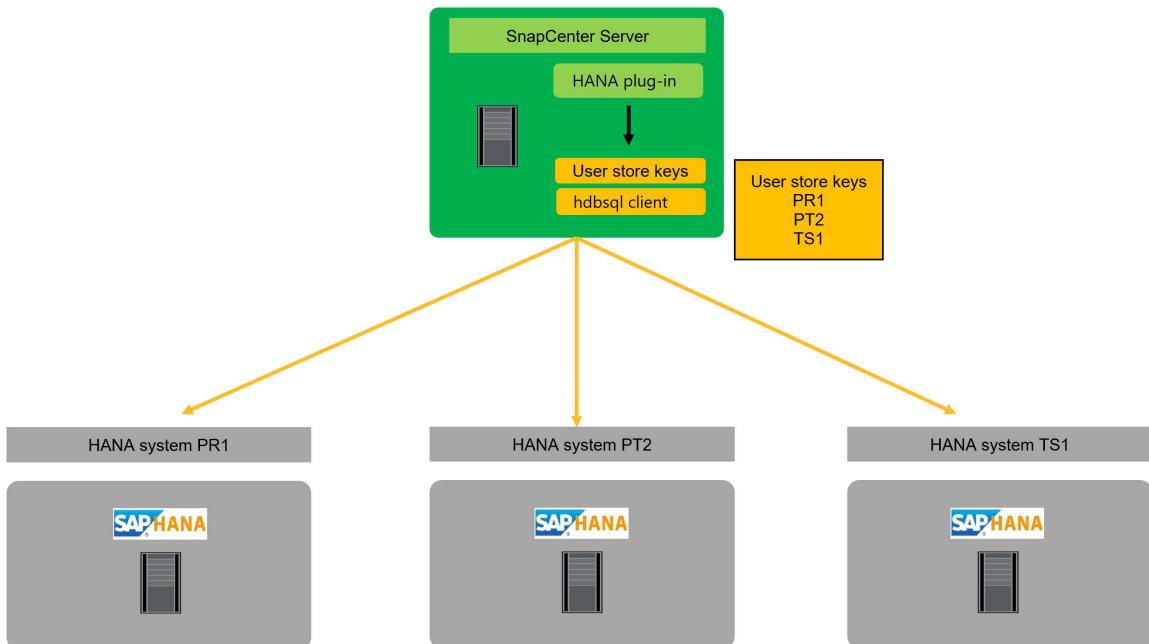


When you run your SnapCenter server on Windows, you can use your Windows system as central plug-in host. When you run your SnapCenter server on Linux, you must use a different host as central plug-in host .

For a HANA multiple host system, SAP HANA user store keys for all worker and standby hosts must be configured at the central plug-in host. SnapCenter tries to connect to the database using each of the provided keys and can therefore operate independently of a failover of the system database (HANA name server) to a different host.



For multiple single host HANA systems managed by a central plug-in host, all individual SAP HANA user store keys of the HANA systems must be configured at the central plug-in host.



SAP HANA block consistency check

SAP recommends including regular HANA block consistency checks into the overall backup strategy. With traditional file-based backups this check is done with each backup operation. With Snapshot backups, the consistency check must be executed in addition to the Snapshot backup operations, for example once per week.

Technically there are two options to execute the block consistency check.

- Executing a standard file-based or backint-based backup
- Executing the HANA tool hdbpersdiag, see also [Persistence Consistency Check | SAP Help Portal](#)

The HANA hdbpersdiag tool is part of the HANA installation and allows to execute block consistency check operations against an offline HANA database. Hence it is a perfect fit to be used in combination with Snapshot backups where existing Snapshot backups can be presented to hdbpersdiag.

When comparing the two approaches, hdbpersdiag has significant advantages compared to the file-based backup for HANA block consistency checks. One dimension is the required storage capacity. With file-based backups at least the size of one backup needs to be available for each HANA system. If you have, for example, 15 HANA systems with a persistence size of 3TB you would need additional 45TB just for the consistency checks. With hdbpersdiag no additional storage capacity is required since the operation is executed against an existing Snapshot backup or a FlexClone of an existing Snapshot backup. The second dimension is the CPU load at the HANA host during the consistency check operation. A file-based backup will require CPU cycles at the HANA database host while the hdbpersdiag processing can be fully offloaded from the HANA host when used in combination with a central verification host. The table below summarizes the key characteristics.

	Required storage capacity	CPU and network load at HANA host
File-based backup	Minimal 1 x data backup size for each HANA system	High
hdbpersdiag using Snapshot directory at HANA host (NFS only)	None	Medium
Central verification host used to run hdbpersdiag with FlexClone volumes	None	None

NetApp recommends using hdbpersdiag to execute HANA block consistency checks. Further details on the implementation are available in chapter "["Block consistency checks with SnapCenter"](#)".

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 a secondary backup site?
- How long should the backups be kept at the secondary 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 a secondary backup site once per day. The test systems have lower requirements and no replication of the backups.

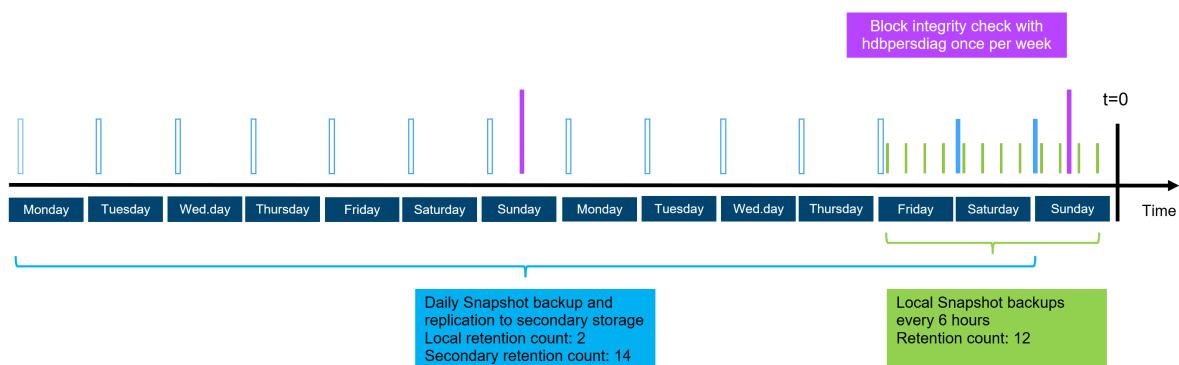
Parameters	Production systems	Development systems	Test systems
Backup frequency	Every 6 hours	Every 6 hours	Every 12 hours
Primary retention	3 days	3 days	6 days
Block integrity check	Once per week	Once per week	No
Replication to secondary backup site	Once per day	Once per day	No
Secondary backup retention	2 weeks	2 weeks	No

The following table shows the policies and the schedules that would need to be configured for the above data protection parameters.

Policy	Backup type	Schedule frequency	Primary retention	SnapVault replication	Secondary retention
LocalSnap	Snapshot based	Every 6 hours	Count=12	No	NA
LocalSnapAndSnapVault	Snapshot based	Once per day	Count=2	Yes	Count=14
SnapAndCallHdbpersdiag	Snapshot based	Once per week	Count=2	No	NA

- ① For ONTAP system or FSx for ONTAP a data protection relationship must be configured in ONTAP for the SnapVault replication, before SnapCenter can execute SnapVault update operations. The secondary retention is defined within the ONTAP protection policy.
- ② For ANF backup, no additional configuration is required outside of SnapCenter. The ANF backup secondary retention is managed by SnapCenter.
- ③ For this example configuration, hdbpersdiag is used for the block integrity check operation. More details can be found in chapter ["Block consistency checks with SnapCenter"](#).

The figure below summarizes the schedules and backup retentions. If SnapCenter is used to manage log backup retention, all log backups which are older than the oldest Snapshot backup will be deleted. In other words, log backups are kept as long as required to enable recovery to current in time for every available backup.



Backup of encryption root keys

When HANA persistence encryption is used it is critical to create backups of the root keys in addition to the standard data backups. Root key backups are required to recover the HANA database in case the data volume and the HANA installation file system are lost. For more information see [SAP HANA Administration Guide](#).



Keep in mind, that if a root key is changed, the new root key can't be used to recover old HANA database backups which have been created before. You always need the root key that has been active at the creation time of the backup.

Backup operations

SnapCenter supports Snapshot backup operations of HANA MDC systems with a single or 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 one 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. 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, SnapCenter identifies 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 to be able to support a single tenant restore operation.

Another step of the application auto discovery is the detection of HANA System Replication (HSR) primary or secondary node. If a HANA system is configured with HSR, SnapCenter must identify the primary node with each backup operation so that the backup SQL commands are executed at the HSR primary node. See also [SAP HANA System Replication - Backup and Recovery with SnapCenter](#).

SnapCenter also detects the HANA data volume configuration and maps it to file system and storage resources. With this approach SnapCenter can handle HANA volume configuration changes, e.g. multiple partitions or storage configuration changes like migrations of volumes.

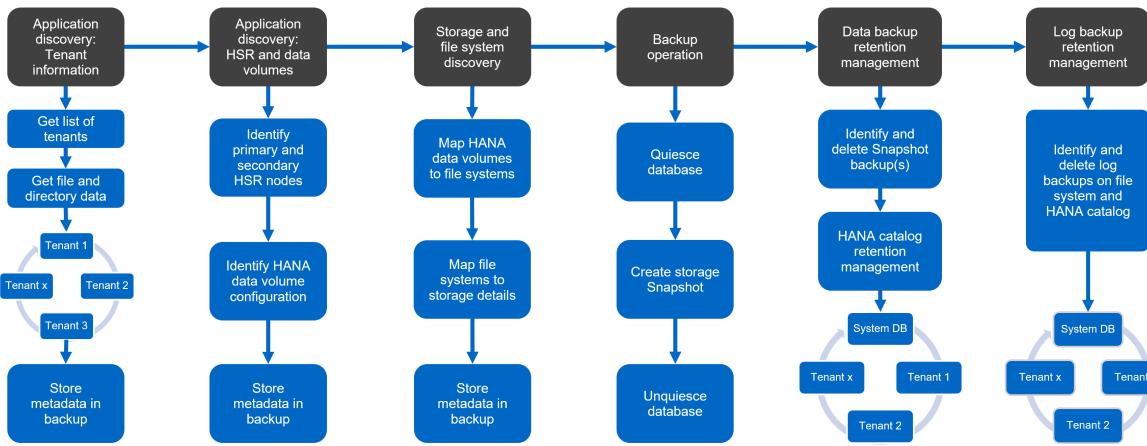
The next step is the Snapshot backup operation itself. This step includes the SQL command to trigger the HANA database snapshot, the storage Snapshot backup, and the SQL command to close the HANA snapshot operation. By using the close command, the HANA database updates the backup catalog of the system 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.

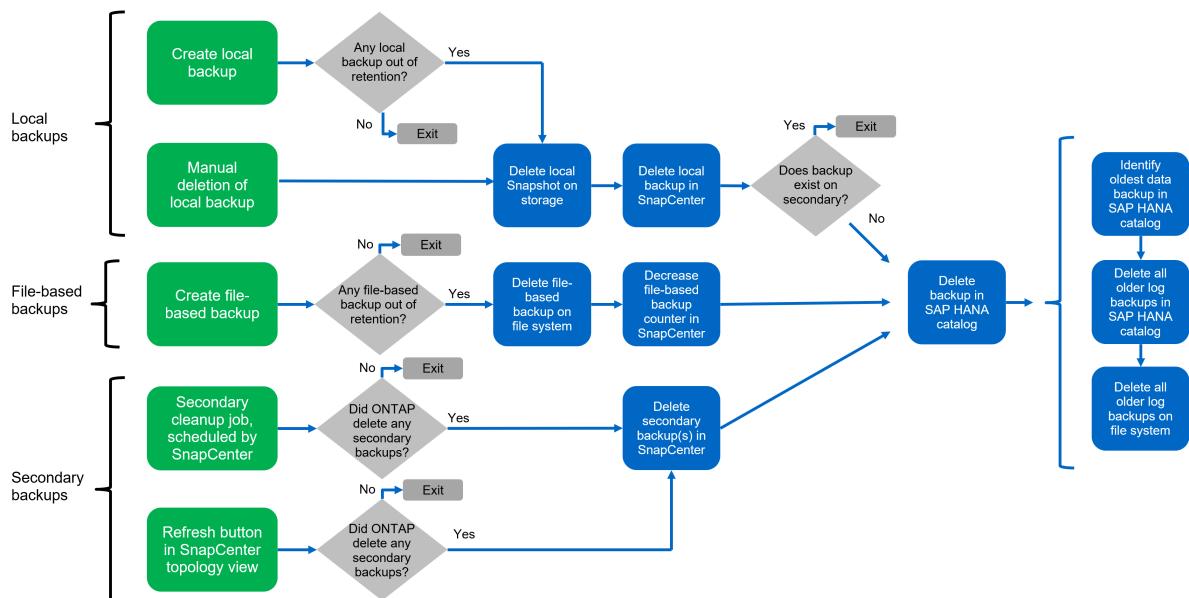


Backup retention management

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 (SnapVault or ANF backup)
- Data backups in the SAP HANA backup catalog
- Log backups in the SAP HANA backup catalog and on 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 is included with each backup workflow in

SnapCenter. 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.

Retention management of backups at the secondary storage (SnapVault)

The retention management of backups at the secondary storage (SnapVault) 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. For details about how to adapt the schedule of the cleanup job or how to trigger a manual refresh, refer to the chapter "["Cleanup of secondary backups"](#)".

Retention management of backups at the secondary storage (ANF backup)

The retention of ANF backups is configured and handled by SnapCenter. SnapCenter handles the housekeeping of ANF backup backups by deleting the backups according to a retention defined in the SnapCenter backup policy. Retention management is included with each backup workflow in SnapCenter.

Retention management of data backups within the SAP HANA backup catalog

When SnapCenter has deleted any backup, local Snapshot or file based or if SnapCenter has identified a 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 operations 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 data 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 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 [Deactivate automated log backup housekeeping](#).

Learn about configuring SnapCenter for SAP HANA environments

Configure SnapCenter for SAP HANA environments using a two-phase approach: initial configuration for shared resources (credentials, storage systems, and policies) and resource-specific configuration for individual HANA systems (host deployment, auto discovery, and protection settings).

The SnapCenter configuration for an SAP HANA environment with multiple HANA systems can be split into two main areas:

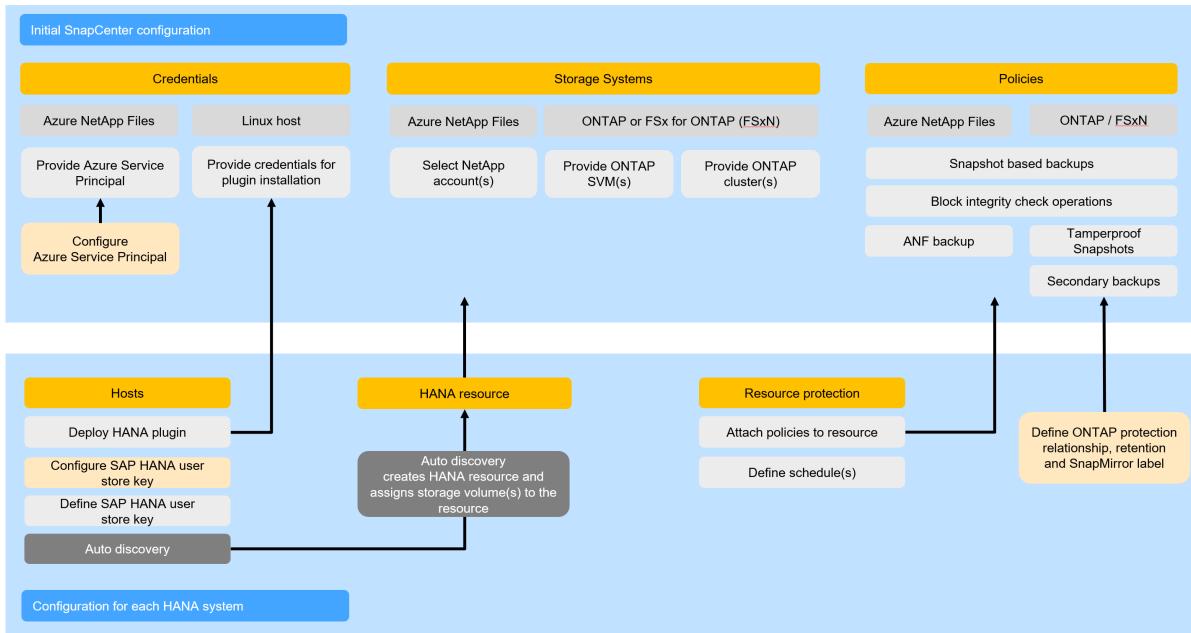
- The initial configuration
 - Credential, storage and policy configurations.
These settings or resources are typically consumed by multiple HANA systems.
- The HANA resource-specific configuration
 - Host, HANA and resource protection configuration must be done for each HANA system individually.

The figure below illustrates the different configuration components and their dependencies.

All configuration steps are described in detail in the following topics.



The descriptions and screenshots in the document are based on SnapCenter auto discovered HANA systems. Additional or different configuration steps for manual configured resources with a central plug-in host are described in "[Central plug-in host configuration](#)".



Configure initial SnapCenter settings for SAP HANA

Configure initial SnapCenter settings for SAP HANA environments by setting up credentials for Azure service principals, adding storage systems, and creating policies for Snapshot backups, block integrity checks, and secondary replication.

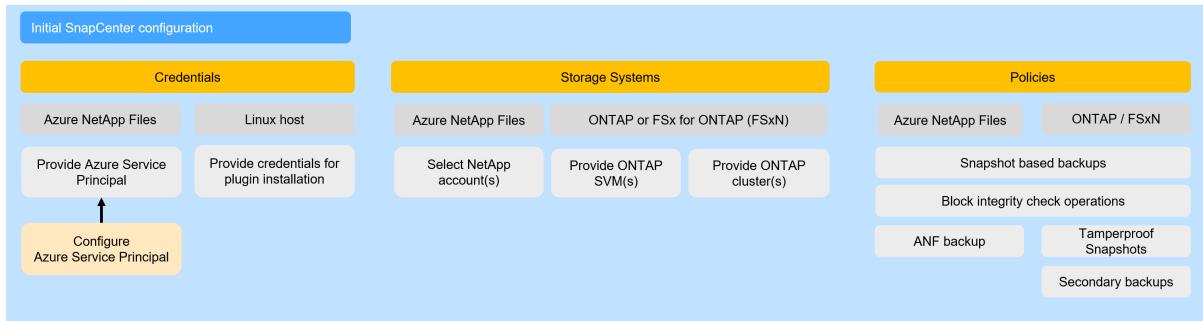
The SnapCenter initial configuration includes the following steps:

1. Credentials configuration
 - a. For HANA systems configured with Azure NetApp Files (ANF), a service principal must be prepared and then configured in SnapCenter.
 - b. Host credentials must be provided to allow the automated installation of the HANA plug-in on the HANA database hosts.
2. Storage system configuration
 - a. For HANA systems configured with ANF, the required NetApp accounts can be selected and added to the SnapCenter configuration.
 - b. For ONTAP or FSx for ONTAP storage systems, either SVMs or the complete storage cluster can be added to SnapCenter.
3. Policies configuration
 - a. Policies for Snapshot based backups as well as block integrity check operations can be configured for ANF as well as for ONTAP and FSx for ONTAP storage systems.
 - b. Policies for tamperproof Snapshots and secondary backups with SnapVault or SnapMirror can only be configured for ONTAP and FSx for ONTAP storage systems.
 - c. For HANA systems configured with ANF, a policy can include [ANF backup](#).



The same Snapshot backup policies can be used for HANA databases as well as for non-data volumes, e.g. the HANA shared volume.

The figure below summarizes the configuration sections.



The following chapters describe the initial configuration steps.

Credentials configuration

Credentials for HANA plug-in deployment

Credentials are configured in the Settings section and by selecting the Credential tab. Credentials can be added by clicking the + icon.

The screenshot shows the 'Credential' tab in the NetApp SnapCenter® interface. The table lists four credentials:

Credential Name	Authentication Mode	Details
installOnBareMetal	Linux	Userid:root
installPluginOnLinux	Linux	Userid:root
installPluginOnWindows	Windows	Userid:sapcc
SCV-sapcc	Linux	Userid:admin

NetApp recommends to configure a user on all HANA database hosts (e.g. scuser) and configure sudo privileges as described in [Prerequisites for adding hosts and installing SnapCenter Plug-in for SAP HANA Database](#).

The 'Credential' dialog box is shown with the following fields:

- Credential Name: `InstallPluginOnLinux`
- Authentication Mode: `Linux`
- Authentication Type: Password Based
- Username: `scuser`
- Password: `.....`
- Use sudo privileges

Buttons at the bottom: `Cancel` and `OK`.

Credentials for Azure NetApp Files

An Azure service principal must be prepared, which enables SnapCenter to execute the required operations for the ANF volumes. The example below shows the minimal required permissions, which must be included.

```

"assignableScopes": [
    "/subscriptions/xxx"
]

```

```
],
  "createdBy": "xxx",
  "createdOn": "2025-05-07T07:12:14.451483+00:00",
  "description": "Restricted Access for SnapCenter",
  "id": "subscriptions/xxx/providers/Microsoft.Authorization/roleDefinitions/xxx"
  ,
  "name": "xxx",
  "permissions": [
    {
      "actions": [
        "Microsoft.NetApp/register/action",
        "Microsoft.NetApp/unregister/action",
        "Microsoft.NetApp/netAppAccounts/read",
        "Microsoft.NetApp/netAppAccounts/getKeyVaultStatus/action",
        "Microsoft.NetApp/netAppAccounts/migrateEncryption/action",
        "Microsoft.NetApp/netAppAccounts/transitionToCmk/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/read",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/read",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/write",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/delete",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/revert/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/poolChange/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/finalizeRelocation/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/revertRelocation/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/breakFileLocks/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/getGroupIdListForLdapUser/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/backups/write",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/backups/restoreFiles/action",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/read",
        "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/write"
      ]
    }
  ]
}
```

```
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/write",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/delete",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/restoreFiles/action",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/subvolumes/read",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/subvolumes/write",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/subvolumes/getMetadata/action",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/volumeQuotaRules/read",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/latestRestoreStatus/current/read",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/mountTargets/read",
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/restoreStatus/read",
",
"Microsoft.NetApp/netAppAccounts/snapshotPolicies/read",
"Microsoft.NetApp/netAppAccounts/snapshotPolicies/write",
"Microsoft.NetApp/netAppAccounts/snapshotPolicies/listVolumes/read",
"Microsoft.NetApp/netAppAccounts/snapshotPolicies/volumes/read",
"Microsoft.NetApp/netAppAccounts/volumeGroups/read",
"Microsoft.NetApp/netAppAccounts/volumeGroups/write",
"Microsoft.NetApp/locations/checknameavailability/action",
"Microsoft.NetApp/locations/checkfilepathavailability/action",
"Microsoft.NetApp/locations/operationresults/read",
"Microsoft.NetApp/Operations/read",
"Microsoft.Resources/resources/read",
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.Resources/subscriptions/resourcegroups/resources/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Network/virtualNetworks/write",
"Microsoft.Network/virtualNetworks/subnets/write",
"Microsoft.NetApp/netAppAccounts/backupVaults/read",
"Microsoft.NetApp/netAppAccounts/backupVaults/write",
```

```

    "Microsoft.NetApp/netAppAccounts/backupVaults/backups/read",
    "Microsoft.NetApp/netAppAccounts/backupVaults/backups/write",
    "Microsoft.NetApp/netAppAccounts/backupVaults/backups/delete",

"Microsoft.NetApp/netAppAccounts/backupVaults/backups/restoreFiles/action"
],
"condition": null,
"conditionVersion": null,
"dataActions": [],
"notActions": [],
"notDataActions": []
}
],
"roleName": "SnapCenter-Restricted-Access",
"roleType": "CustomRole",
"type": "Microsoft.Authorization/roleDefinitions",
"updatedBy": "xxx",
"updatedOn": "2025-05-07T07:12:14.451483+00:00"
}
}

```

Credentials are configured in the Settings section and by selecting the Credential tab. Credentials are configured by clicking the + icon.

Credential Name	Authentication Mode	Details
ANF-EASTUS	AzureCredential	
ANF-EASTUS-Minimal	AzureCredential	
Plugin-Linux	Linux	Userid:root

In the following screen, a credential name must be provided and the Authentication Mode Azure Credentials must be selected. Then tenant ID, client ID and client secret key must be configured.

Credential

Credential Name: ANF-EASTUS-Minimal

Authentication Mode: Azure Credential

Azure Details

Tenant ID: [REDACTED]

Client ID: [REDACTED]

Client Secret Key: Enter client secret key

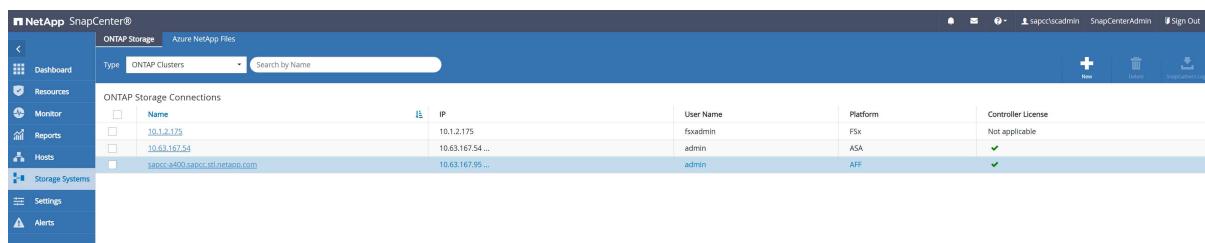
Cancel OK

Storage system configuration

ONTAP systems and FSx for ONTAP

ONTAP system or FSx for ONTAP can be added to SnapCenter either by providing cluster credentials or credentials for each required SVM. When cluster credentials are provided, all SVMs of the cluster are added to SnapCenter.

In our lab setup, we added the storage clusters to SnapCenter. ONTAP clusters are configured in the Storage systems section by selecting the ONTAP storage tab and the ONTAP Cluster type. A new cluster is added by clicking the + icon.

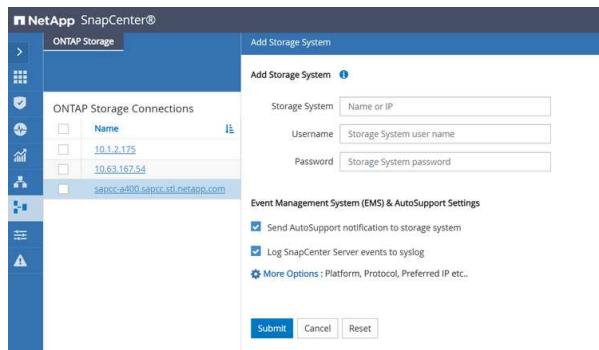


Name	IP	User Name	Platform	Controller License
10.1.2.175	10.1.2.175	fsxadmin	FSx	Not applicable
10.63.167.54	10.63.167.54...	admin	ASA	✓
sapcc-a400.snapc-st1.netapp.com	10.63.167.95...	admin	AFF	✓

In the following screen you need to provide the credentials for a cluster user.



The cluster user admin should not be used. Instead a new user should be created with the required privileges as described in [Create ONTAP cluster roles with minimum privileges](#). Required privileges for ASA system can be found at [Create ONTAP cluster roles for ASA r2 systems](#).



SVMs are configured in the Storage systems section by selecting the ONTAP storage tab and the ONTAP SVMS type. A new SVM is added by clicking the + icon.

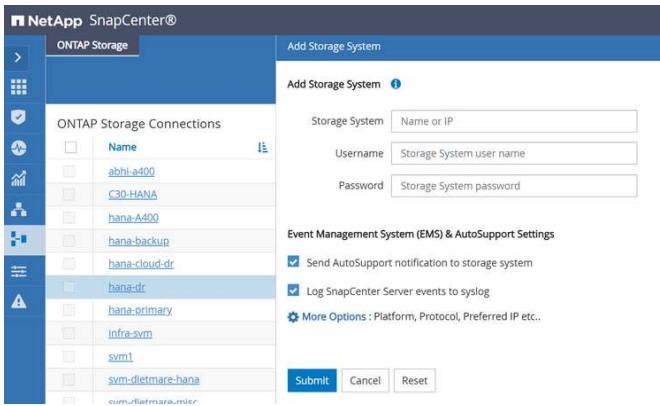
In the following screen you need to provide the credentials for a cluster user.



The SVM user vsadmin should not be used. Instead a new user should be created with the required privileges as described in [Create SVM roles with minimum privileges](#). Required privileges for ASA system can be found at [Create SVM roles for ASA r2 systems](#).

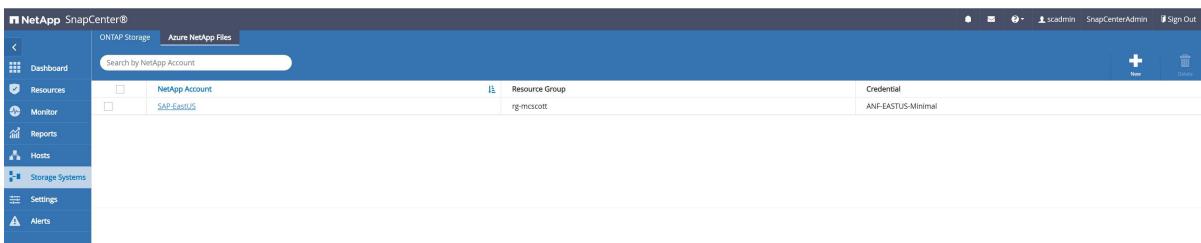


The DNS name for the SVM must match with the SVM name configured at the ONTAP system.

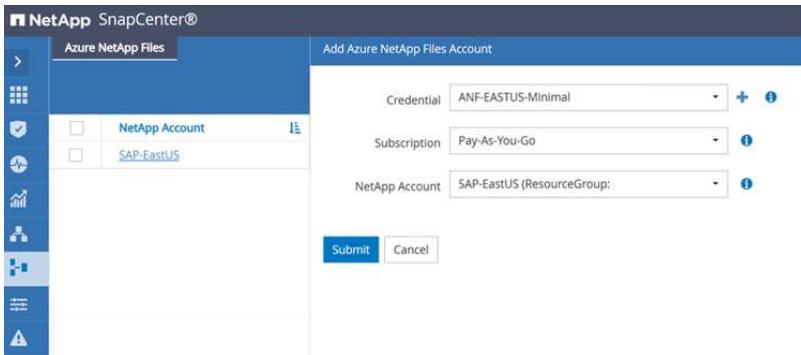


Azure NetApp Files

After the ANF credentials have been configured, ANF NetApp Account(s) can be added to SnapCenter. NetApp Accounts are configured in the Storage systems section and by selecting the Azure NetApp Files tab. A new NetApp Account is added by clicking the + icon.



After selecting the ANF credential and the subscription, a NetApp account can be added to SnapCenter.



Storage configuration when using SnapMirror active sync

Specific storage configuration steps are described at [Storage configuration with SnapMirror active sync](#).

Policies configuration

As discussed in the section Data protection strategy policies are usually configured independently of the resource and can be used for multiple SAP HANA systems.

A typical minimum configuration consists of the following policies:

- Policy for hourly backups without replication
- Policy for daily backups with SnapVault or ANF backup replication
- Policy for weekly block integrity check operation

- using a file-based backup
- using the HANA tool hdbpersdiag

The following sections describe the configuration of these three policies.

Policies are configured in the Settings section and by selecting the Policies tab. A new policy is configured by clicking the + icon. The two screenshots below show the list of policies for HANA systems running with Azure NetApp Files and a second one for HANA systems running with ONTAP storage systems or FSx for ONTAP.

Name	Scope	Schedule Type	Snapshot	Backup	Replication
BlockIntegrityCheck	File Based Backup	Weekly		Copies to keep : 1 copies	
LocalSnap	Data Backup	Hourly	Copies to keep : 12 copies		
LocalSnapAndNFBackup	Data Backup	Daily	Copies to keep : 2 copies	Copies to keep : 7 copies	

Name	Scope	Schedule Type	Snapshot	Backup	Replication
BlockIntegrityCheck	File Based Backup	Weekly		Copies to keep : 1 copies	
LocalSnap	Data Backup	Hourly	Copies to keep : 12 copies		
LocalSnapAndSnapVault	Data Backup	Daily	Copies to keep : 2 copies		
LocalSnapKeep2	Data Backup	Hourly	Copies to keep : 2 copies		SnapVault
LocalSnapOnDemand	Data Backup	On demand	Copies to keep : 2 copies		
LocalSnapTamperProof	Data Backup	Daily	Copies to keep : 4 copies		
PolicyCBA	Data Backup	Daily	Copies to keep : 4 copies		
SnapAndCallHdbpersdiag	Data Backup	Daily	Copies to keep : 2 copies		

Snapshot backups with ONTAP systems and FSx for ONTAP

Snapshot backup policies for ONTAP system or FSx for ONTAP can combine a local Snapshot with replication or Snapshot locking (tamperproof Snapshot) operations. This example shows a policy with replication to a secondary storage using SnapVault.

Provide a policy name and an optional description.

Select ONTAP storage type and Snapshot policy scope.

For this policy a daily schedule type has been configured. A daily Snapshot will be created, and the Snapshot deltas will be replicated to the secondary storage using SnapVault.

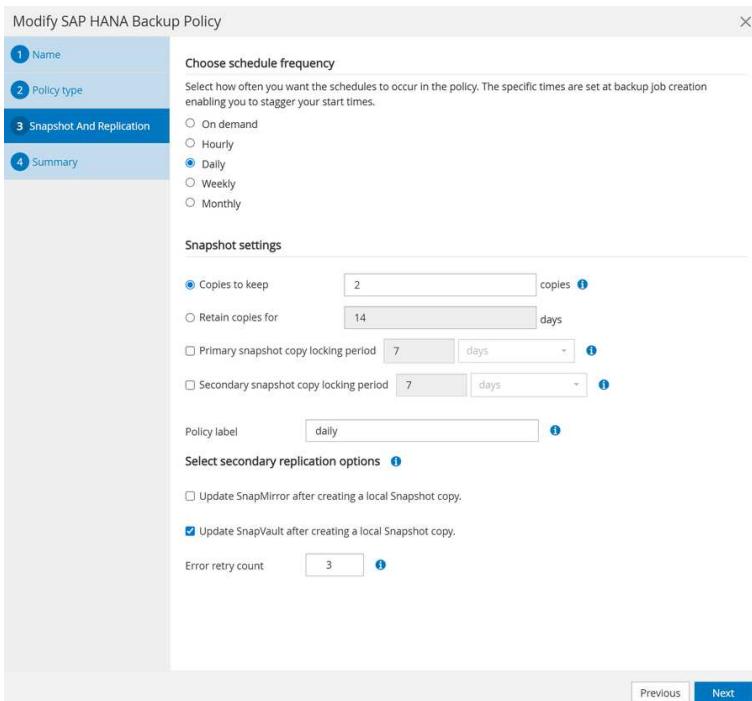


The schedule itself is configured with the individual HANA resource protection configuration.

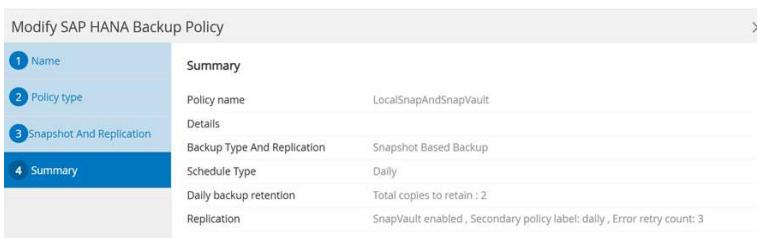
The retention which is configured in the policy is only valid for the primary Snapshots. The retention at the SnapVault target is configured with the ONTAP replication relationship for the individual volume(s) of the HANA database as described in chapter ["SAP HANA Snapshot backup operations"](#). The Snapshot label which is configured in the policy must match with the label configured with the ONTAP replication relationship.

Snapshot locking (tamperproof Snapshots) can be enabled by clicking the check boxes and defining the locking period. This feature requires a SnapLock license at the storage system and the compliance clock being configured.

A policy for local Snapshots only would be configured with an hourly schedule and by disabling the Update SnapVault check box.



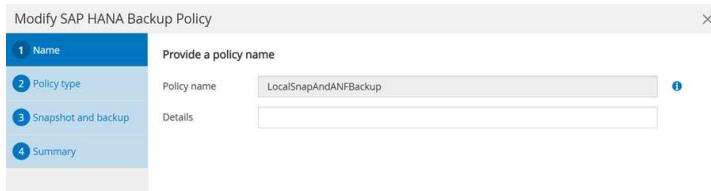
The summary screen shows the configured parameters.



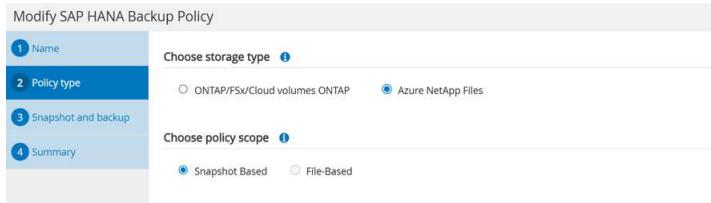
Snapshot backups with Azure NetApp Files

Snapshot backup policies for Azure NetApp Files can combine a local Snapshot with ANF backup, which replicates the Snapshot data to Azure blob. This example shows a policy used for replication with ANF backup.

Provide a policy name and an optional description.



Select Azure NetApp Files storage type and Snapshot policy scope.



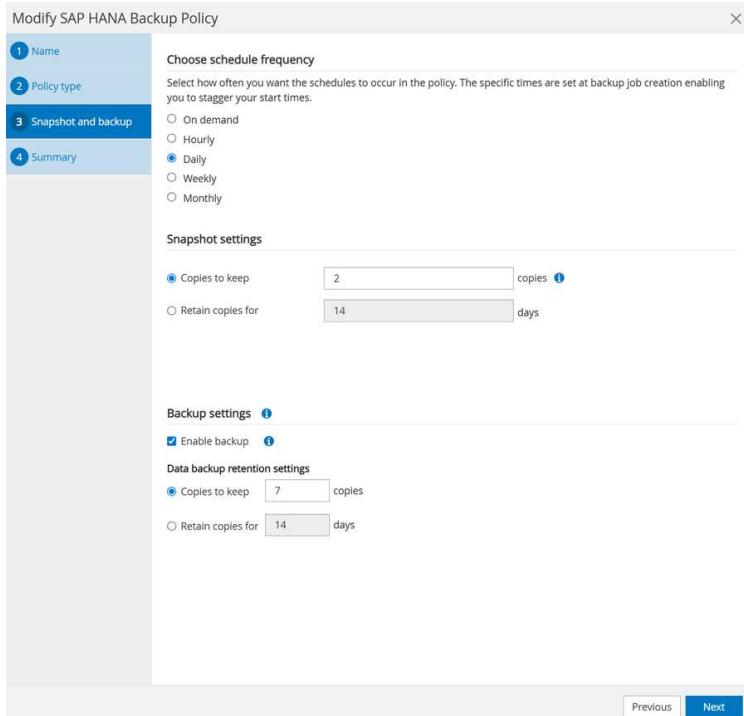
For this policy a daily schedule type has been configured. A daily Snapshot will be created, and the Snapshot deltas will be replicated to the backup vault using ANF backup.



The schedule itself is configured with the individual HANA resource protection configuration.

The Snapshot retention which is configured in the policy is valid for the primary Snapshots at the ANF volume. The retention for the ANF backup is configured with the backup retention settings.

A policy for local Snapshots only would be configured with an hourly schedule and by disabling the Enable backup check box.



The summary screen shows the configured parameters.

Modify SAP HANA Backup Policy

1 Name	Summary
2 Policy type	Policy name LocalSnapAndANFBackup
3 Snapshot and backup	Details
4 Summary	Backup Type And Replication Snapshot Based Backup
	Schedule Type Daily
	Daily backup retention Total copies to retain : 2
	Object store backups Enabled
	Daily object store data backup Total copies to retain : 7
	Retention

Block integrity check operations for all platforms

HANA tool hdbpersdiag

Details are described in chapter ["Block consistency checks with SnapCenter"](#).

File-based backup

Provide a policy name and an optional description.

Modify SAP HANA Backup Policy

1 Name	Provide a policy name
2 Policy type	Policy name BlockIntegrityCheck
3 Snapshot and Replication	Details
4 Summary	

Select ONTAP or Azure NetApp Files storage type, depending on your setup and select File-based policy scope.

Modify SAP HANA Backup Policy

1 Name	Choose storage type
2 Policy type	<input checked="" type="radio"/> ONTAP/FSx/Cloud volumes:ONTAP <input type="radio"/> Azure NetApp Files
3 Snapshot and Replication	Choose policy scope
4 Summary	<input type="radio"/> Snapshot Based <input checked="" type="radio"/> File-Based

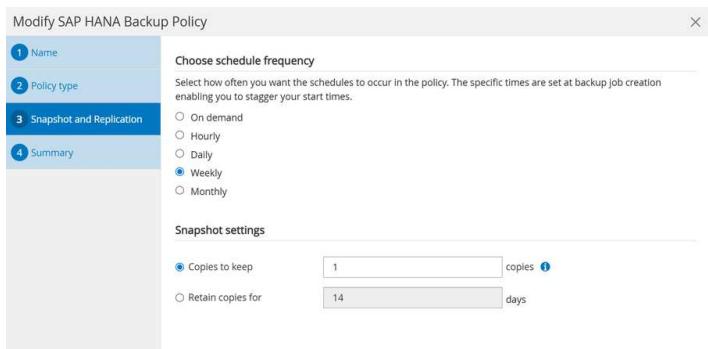
As discussed, it is recommended to execute the block integrity check once per week. Therefore, a weekly schedule is selected.



The schedule itself is configured with the individual HANA resource protection configuration.



The file system where the file-based backup is written to must provide enough capacity for one backup more than defined in the retention settings, because SnapCenter deletes the old backup after the new one has been created. In this example space for two backups is required with a retention of one. The minimal configurable retention is zero.



The summary screen shows the configured parameters.



Policy configuration when using SnapMirror active sync

Specific policy configuration steps are described in the document [Policy configuration SnapMirror active sync](#).

Configure SnapCenter resources for individual SAP HANA databases

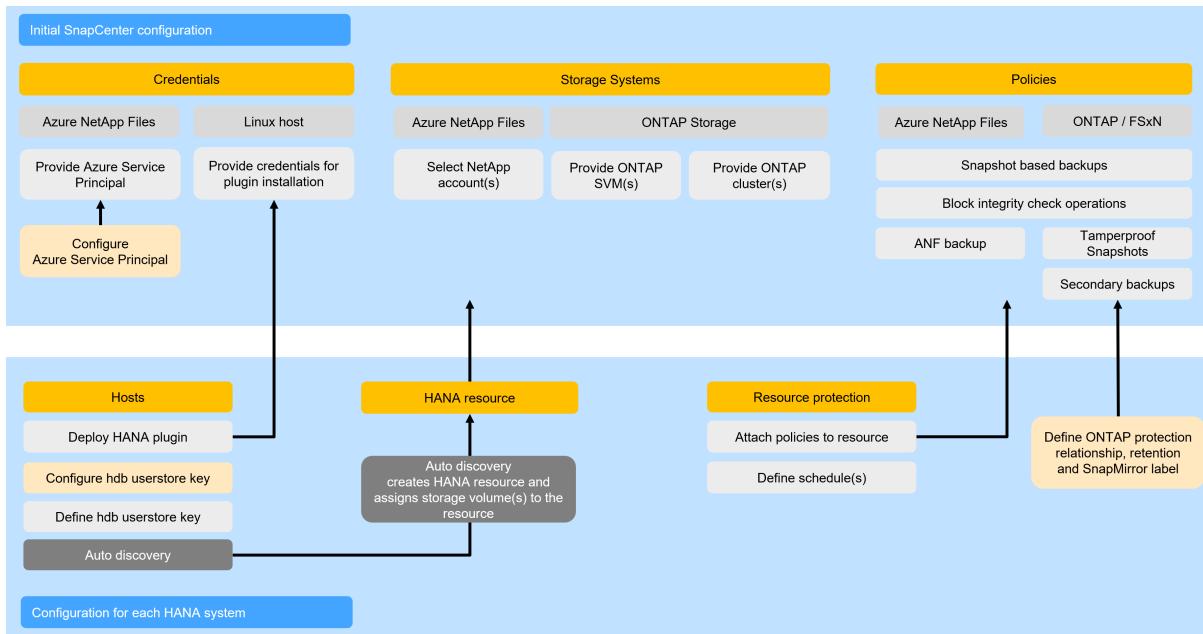
Configure individual SAP HANA databases in SnapCenter by creating backup users and user store keys, setting up storage replication for secondary backups, deploying the HANA plug-in for auto discovery, and configuring resource protection with policies and schedules.

The configuration of a HANA database in SnapCenter is done with the following steps:

1. A SnapCenter backup user must be configured in the HANA system database, and an SAP HANA user store key must be set up at the HANA database host
2. If data replication to a secondary storage is required, the ONTAP storage replication for the HANA data volume must be configured
3. The SnapCenter HANA plug-in must be deployed on the HANA database host
 - a. Auto discovery process gets started
 - b. SAP HANA user store key must be configured in SnapCenter
 - c. Second phase of auto discovery gets started and the HANA resource is added automatically by SnapCenter
4. HANA resource protection must be configured for the new added HANA resource

The initial SnapCenter configuration, as described in the previous topic "[SnapCenter initial configuration](#)" must be done first, since credentials, storage systems and policies are required during the HANA database resource configuration. The figure below summarizes the steps and dependencies.

The figure below visualizes the different configuration components and dependencies.



The following sections provide a detailed description of the required configuration steps.

SAP HANA backup user and SAP HANA user store configuration

NetApp recommends configuring a dedicated user in the HANA database to run the backup operations with SnapCenter. As a second step, an SAP HANA user store key is configured for this backup user, and the SAP HANA user store key is provided in the SnapCenter configuration.

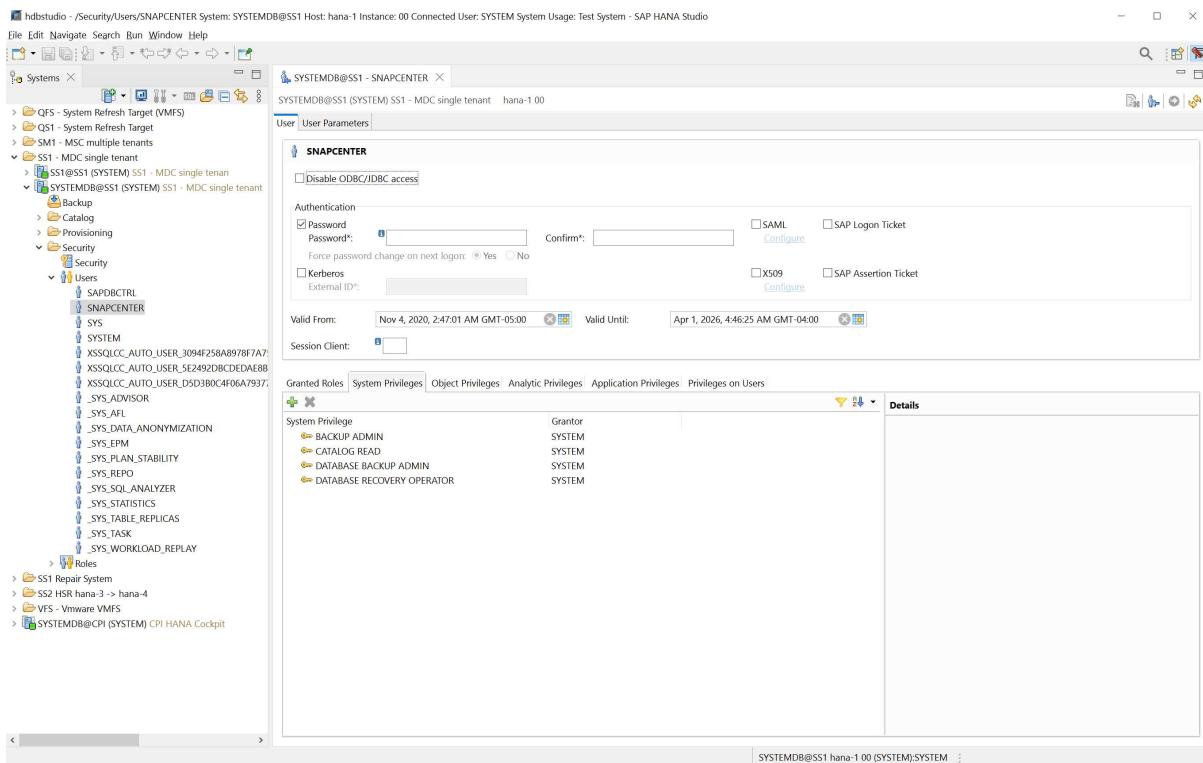
The following figure shows the SAP HANA Studio through which the backup user, in this example SNAPCENTER can be created.



The backup user needs to be configured with the privileges backup admin, catalog read, database backup admin, and database recovery operator.



The backup user must be created in the system database because all backup commands for the system and the tenant databases are executed via the system database.



SAP HANA user store configuration on the HANA database host

SnapCenter uses the <sid>adm user to communicate with the HANA database. Therefore, the SAP HANA user store key must be configured using the <sid>adm user on the database host.

hdbuserstore set <key-name> <host>:<port> <database user> <password>

For an SAP HANA MDC system, the port of the HANA system database is 3<instanceNo>13.

SAP HANA user store configuration examples

The output shows the key SS1KEY which has been configured for the HANA system with instance number = 00.

```

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 SS1SAPDBCTRL
ENV : hana-1:30013
USER: SAPDBCTRL
KEY SS1KEY
ENV : hana-1:30013
USER: SNAPCENTER
KEY SYSTEMKEY
ENV : hana-1:30013
USER: SYSTEM
ACTIVE RECORDS : 10
DELETED RECORDS : 15
NUMBER OF COMPLETE KEY: 3
Operation succeed.
ssladm@hana-1:/usr/sap/SS1/HDB00>

```

The output shows the key SM1KEY which has been configured for the HANA system with instance number = 12.

```

sm1adm@hana-2:/usr/sap/SM1/HDB12> hdbuserstore list
DATA FILE : /usr/sap/SM1/home/.hdb/hana-2/SSFS_HDB.DAT
KEY FILE : /usr/sap/SM1/home/.hdb/hana-2/SSFS_HDB.KEY
KEY SM1SAPDBCTRL
ENV : hana-2:31213
USER: SAPDBCTRL
KEY SM1KEY
ENV : hana-2:31213
USER: SNAPCENTER
ACTIVE RECORDS : 7
DELETED RECORDS : 9
NUMBER OF COMPLETE KEY: 2
Operation succeed.
sm1adm@hana-2:/usr/sap/SM1/HDB12>

```

Storage replication configuration

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 screenshots show a configuration using ONTAP system manager. For FSx for ONTAP systems the replication must be done using the ONTAP CLI as described at [Overview - Backup replication with SnapVault](#).

The following figure shows the configured protection relationship for the data volume of the SAP HANA system

SS1. With this 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_dst.

The screenshot shows the ONTAP System Manager interface under the Replication section. A specific relationship is selected between the source volume 'hana-primary:SS1_data_mnt00001' and the destination volume 'hana-backup:SS1_data_mnt00001_dst'. The policy type is set to 'Asynchronous'. The transfer schedule is listed as 'None'. The protection policy is 'SnapCenterVault'.

The following figure shows the protection policy, which has been created for this lab setup. The protection policy used for the protection relationship defines the SnapMirror label, as well as the retention of backups at the secondary storage. In this example, the used label is Daily, and the retention is set to 5.

- i The SnapMirror label in the replication policy must match the label defined in the SnapCenter policy configuration.
- i The schedule of the relationship must be set to None, because SnapCenter triggers the SnapVault update as part of the backup operation based on the application consistent Snapshot created before.
- i The retention for backups at the secondary backup storage is defined in the policy and controlled by ONTAP.

The screenshot shows the ONTAP System Manager interface under the Policies section. A replication policy named 'SV Daily' is selected. The policy is set to 'Asynchronous' and has a scope of 'hana-primary'. The transfer schedule is 'daily'. The retention count is set to 5, and the retention period is 1 day.

ANF backup configuration

For ANF backup no specific preparation is required. As soon as the first backup with enabled ANF backup is executed an Azure backup vault with the name snapcenter-vault is created by SnapCenter. This backup vault is then used by all following ANF backup operations executed by SnapCenter.

Deployment of SnapCenter plug-in for SAP HANA

The host requirements are listed at [Host requirements for installing the SnapCenter Plug-ins Package for Linux](#).

The HANA plug-in deployment is done by clicking the Add button in the Hosts section of the SnapCenter UI.

In the Add host screen, you need to provide the host type and name and the credentials to be used for the deployment process. In addition, the SAP HANA plug-in must be selected. By clicking submit the deployment process starts.



For this description we didn't add a new host but show the configuration of existing hosts in SnapCenter.

HANA auto discovery

Once the HANA plug-in deployment is finished, the auto discovery process gets started. In the first phase, only basic settings are discovered and SnapCenter creates a new resource which gets listed on the Resources section of the UI marked with a red padlock.

SAP HANA										
System		System ID (SID)	Tenant Databases	Replication	System State	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
QFS	QFS	QFS	QFS	None	Offline	hana-9.sapcc.stl.netapp.com				Not protected
QSI	QSI	QSI	QSI TENANT1 TENANT2	None	Offline	hana-7.sapcc.stl.netapp.com				Not protected
SM1	SM1	SM1	TENANT2 TENANT1	None	Online	hana-2.sapcc.stl.netapp.com	LocalSnap LocalSnapTamperProof	08/18/2025 12:01:04 PM	Backup succeeded	
SS1	SS1	SS1	SS1	None	Online	hana-1.sapcc.stl.netapp.com	BlockIntegrityCheck LocalSnap LocalSnap4EndsVault LocalSnap-OnDemand	08/18/2025 12:01:02 PM	Backup succeeded	
SS2	SS2	SS2	Enabled (Primary)	Online	hana-3.sapcc.stl.netapp.com	SS2 - HANA System Replicat ion	BlockIntegrityCheck LocalSnap4Ends2	08/18/2025 11:57:30 AM	Backup succeeded	
SS2	SS2	SS2	Enabled (Secondary)	Online	hana-4.sapcc.stl.netapp.com	SS2 - HANA System Replicat ion	BlockIntegrityCheck LocalSnap4Ends2	04/11/2022 2:57:21 AM	Backup succeeded	
VFS	VFS	VFS	None	Online	hana-8.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap	08/18/2025 12:16:56 PM	Backup succeeded	

When clicking on the resource, you get asked for the SAP HANA user store key for this HANA database.

Configure Database

Plug-in host	hana-9.sapcc.stl.netapp.com
HDBSQL OS User	qfsadm
HDB Secure User Store Key	<input type="text" value="QFSKEY"/> ?

Cancel OK

After the key has been provided the second phase of the auto discovery process gets started. The auto discovery process detects all tenant databases in the HANA system, log and catalog backup configuration details and HANA system replication roles. In addition, storage footprint details are automatically discovered. These settings can be checked by selecting a resource and clicking on the Details button.



This auto discovery process is executed with each backup operation, so that any changes made to the HANA system, which are relevant for the backup operation will be automatically detected.

SAP HANA		Resource - Details	
Search databases		Details for selected resource	
QFS	System	Type	Multitenant Database Container
QSI	System	HANA System Name	SS1
SM1	System	SID	SS1
SS1	Tenant Databases	SS1	
SS2	Tenant Databases	SS2	
SS2	Tenant Databases	SS2	
VFS	Tenant Databases	VFS	
Plug-in Host		hana-1.sapcc.stl.netapp.com	
HDB Secure User Store Key		SS1KEY	
HDBSQL OS User		ss1adm	
Log backup location		/mnt/log backup	
Backup catalog location		/mnt/catalog	
System Replication		None	
Plug-in name		SAP HANA	
Last backup		08/18/2025 12:01:02 PM (Completed)	
Resource Groups		hana-1.sapcc.stl.netapp.com,hana_MDC_SS1	
Policy		LocalSnap, LocalSnapAndSnapVault, BlockIntegrityCheck, LocalSnap-OnDemand	
Discovery Type		Auto	
Storage Footprint			
SVM	Volume	Junction Path	LUN/Qtree
hana-primary	SS1_data_mnt000001	/SS1_data_mnt000001	

Resource protection configuration

The resource protection configuration screen is opened by clicking on a resource after the auto discovery process has finished. The screenshots in this documentation show the protection configuration of an existing resource.

Configure a custom name format for the Snapshot. NetApp recommends using a custom Snapshot name to easily identify which backups have been created with which policy and schedule type.

In the configuration shown in the following figure, the backup and Snapshot copy names have the following format:

- Scheduled hourly backup:
SnapCenter_<host-name>_LocalSnap_Hourly_<time_stamp>
- Scheduled daily backup:
SnapCenter_<host-name>_LocalSnapAndSnapVault_Daily_<time_stamp>

In the next screen, scripts can be configured, which should be executed at various steps of the backup workflow.

Now policies are attached to the resource and schedules are defined.

In this example we have configured

- A weekly block integrity check, every Sunday
- A local Snapshot backup, every 4 hours

- A daily Snapshot backup with SnapVault replication once per day

Email notification can be configured.

When the resource protection configuration is done, scheduled backups will be executed according to the defined settings.

Configure SnapCenter to back up non-data volumes

Configure SnapCenter to back up non-data volumes such as executables, configuration files, trace files, and application server data.

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, configuration files but also SAP HANA trace files. While executables only change when the SAP HANA database is upgraded, the SAP HANA configuration and trace files might need a higher backup frequency. Also SAP application server volumes can be protected with SnapCenter using non-data volume backups.

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 interaction with SAP HANA database required.

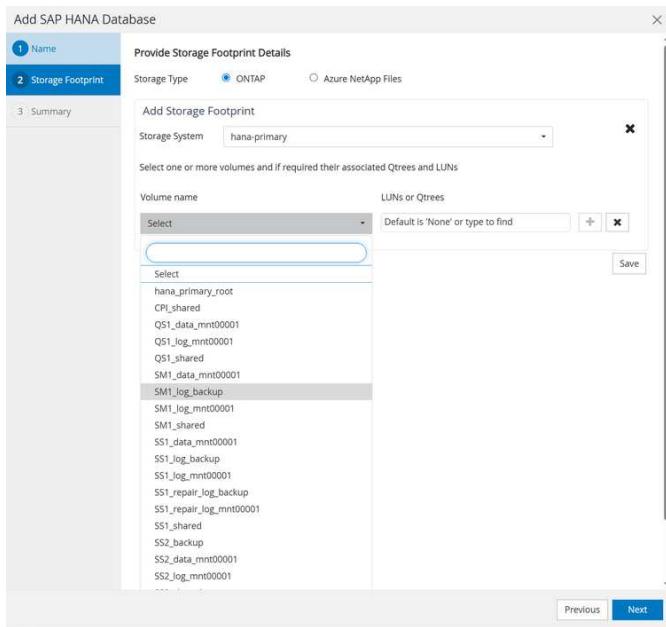
From the Resource tab, select Non-Data-Volume and click Add SAP HANA Database.

System ID (SID)	Tenant Databases	Replication	System State	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
QFS	QFS	None	Offline	hana-9.sapcc.stl.netapp.com				Not protected
QSI	QSI TENANT1 TENANT2	None	Offline	hana-7.sapcc.stl.netapp.com				Not protected
SM1	SM1	TENANT2 TENANT1	None	Online	hana-2.sapcc.stl.netapp.com	LocalSnap LocalSnapTamperProof	08/21/2025 4:01:05 AM	Backup succeeded
SS1	SS1	SS1	None	Online	hana-1.sapcc.stl.netapp.com	BlockIntegrityCheck LocalSnap LocalSnapAndSnapVault LocalSnapOnDemand	08/21/2025 6:01:04 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Primary)	Online	hana-3.sapcc.stl.netapp.com	SS2 - HANA System Replicat ion BlockIntegrityCheck LocalSnapKey2	08/21/2025 6:57:25 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Secondary)	Online	hana-4.sapcc.stl.netapp.com	SS2 - HANA System Replicat ion BlockIntegrityCheck LocalSnapKey2	04/11/2022 2:57:21 AM	Backup succeeded
VFS	VFS	VFS	None	Online	hana-8.sapcc.stl.netapp.com	BlockIntegrityCheck LocalSnap	08/21/2025 6:31:08 AM	Backup succeeded

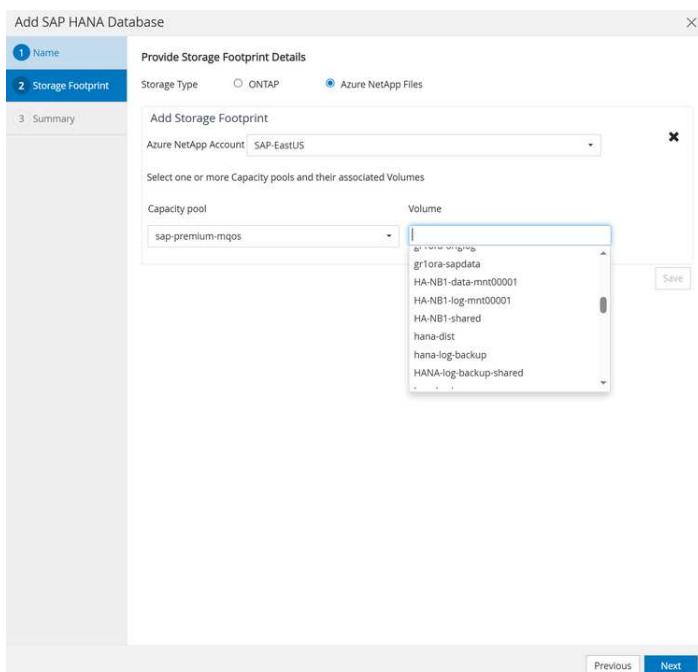
Name	Associated System ID (SID)	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
SM1-Shared-Volume	SM1	hana-2.sapcc.stl.netapp.com		LocalSnap	08/21/2025 6:06:14 AM	Backup succeeded
SS1-Shared-Volume	SS1	hana-1.sapcc.stl.netapp.com		LocalSnap LocalSnapOnDemand	08/21/2025 6:04:13 AM	Backup succeeded

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.

For ONTAP systems and FSx for ONTAP select storage type ONTAP and add the SVM(s) and the storage volume(s) as storage footprint, then click Next.



For ANF select storage type Azure NetApp Files select the NetApp Account and capacity pool and add the ANF volume(s) as storage footprint, then click Next.



In the summary step, click Finish to save the settings.

Repeat these steps for all the required non-data volumes. Continue with the protection configuration of the new resource.



The data protection configuration for non-data volume resources is identical to the workflow for SAP HANA database resources and can be defined on an individual resource level.

Configure SnapCenter central plug-in host for SAP HANA

Deploy the SnapCenter HANA plug-in on a central host to support SAP HANA multiple-host systems or HANA systems on IBM Power. This procedure includes installing the plug-in on a Windows or Linux host, configuring the SAP HANA hdbsql client, and setting up user store keys for each protected HANA system.

As discussed in ["Deployment options for SnapCenter plug-in for SAP HANA"](#), the HANA plug-in can be deployed outside of the HANA database to support a central plug-in configuration which is required SAP HANA multiple host systems or SAP HANA on IBM Power environments.

The central plug-in host can be any Windows or Linux host, but typically the SnapCenter server itself is used as a central plug-in host.

The configuration of a central plug-in host consists of the following steps:

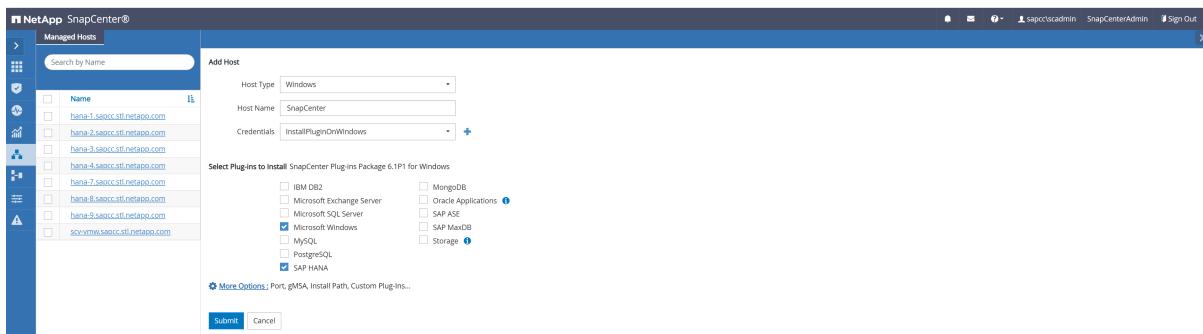
- SnapCenter HANA plug-in deployment
- SAP HANA hdbsql client installation and configuration
- SAP HANA user store configuration for each HANA system which is protected by the central plug-in host

SnapCenter HANA plug-in deployment

The host requirements are listed at [Host requirements for installing the SnapCenter Plug-ins Package for Linux](#).

The central plug-in host on is added as a host, and the SAP HANA plug-in is installed on the host. The screenshot below shows the plug-in deployment on a SnapCenter server running on Windows.

1. Go to Hosts and click Add.
2. Provide the required host information. Click Submit.



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 HANA resource configuration must be able to run the hdbsql executable. The path to the hdbsql executable must be configured in the hana.properties file or in the search path parameters (%PATH%, \$PATH) of the OS user.

Central plug-in host on 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
```

Central plug-in host on Linux:

```
cat /opt/NetApp/snapcenter/scc/etc/hana.properties  
  
HANA_HDBSQL_CMD=/usr/sap/hdbclient/hdbsql
```

SAP HANA user store configuration for a central plug-in host

For each HANA system which is managed by the central plug-in host, a SAP HANA user store key must be configured. Before the key can be configured at the central plug-in host, a database user must be created as described in ["SAP HANA backup user and SAP HANA user store configuration"](#).

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 SAP HANA user store configuration must be done with a different user using the -u <User> option.

```
hdbuserstore.exe -u SYSTEM set <key> <host>:<port> <database user>  
<password>
```

For an SAP HANA multiple-host setup, SAP HANA 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 the system database (HANA name server) to a different host. An SAP HANA user store key is configured for all worker and the standby host. The HANA database user, in this example, SNAPCENTER is the user that has been configured in the system database.

```

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\hdbcclient>hdbuserstore.exe -u SYSTEM list
DATA FILE : C:\ProgramData\.hdb\SNAPCENTER-61\S-1-5-18\SSFS_HDB.DAT
KEY FILE : C:\ProgramData\.hdb\SNAPCENTER-61\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\hdbcclient>

```

HANA manual resource configuration

A manual configured HANA system resource is created in SnapCenter by clicking the Add button in the resource view.

System	System ID (SID)	Tenant Databases	Replication	System State	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
QFS	QFS	QFS	None	Offline	hana-9.sapcc.st.netapp.com				Not protected
QS1	QS1	QS1 TENANT1 TENANT2	None	Offline	hana-7.sapcc.st.netapp.com				Not protected
SM1	SM1	TENANT1 TENANT2	None	Online	hana-2.sapcc.st.netapp.com	LocalSnap		10/10/2025 8:01:01 AM	Backup succeeded
SS1	SS1	SS1	None	Online	hana-1.sapcc.st.netapp.com	BlockIntegrityCheck LocalSnap LocalSnapAndSnapshot LocalSnapshotOnDemand		10/10/2025 8:01:02 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Primary)	Online	hana-3.sapcc.st.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapshot2	10/10/2025 7:57:25 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Secondary)	Online	hana-4.sapcc.st.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapshot2	04/11/2022 2:57:21 AM	Backup succeeded
VFS	VFS	VFS	None	Online	hana-8.sapcc.st.netapp.com		BlockIntegrityCheck LocalSnap	10/10/2025 6:31:29 AM	Backup succeeded

In the next screen you need to provide a couple of system parameters.

- **Plug-in Host:** The central plug-in host must be selected
- **SAP HANA user store key:** For a single host HANA system the key name that has been prepared at the central plug-in host must be provided. For a multiple host HANA system, a comma-separated list of all keys for the system must be provided.

- **HDBSQL OS User:** If the central plug-in host runs on Windows, the user will be pre-select as the SYSTEM user. Otherwise the user which has been used for the SAP HANA user store key must be provided.

Provide Resource Details

Resource Type: Multitenant Database Container

HANA System Name: MCR

SID: MCR

Plug-In Host: Select

HDB Secure User Store Keys: MCRKEY

HDBSQL OS User: [empty]

Previous Next

As a next step the storage footprint needs to be configured. All ONTAP or ANF volumes which belong to the HANA system must be added here.

Provide Storage Footprint Details

Storage Type: ONTAP Azure NetApp Files

Storage System: hana-primary

Add Storage Footprint

Volume name LUNs or Qtrees

Select: [dropdown menu showing volume list]

Default is 'None' or type to find: [input field]

Save

Volume list (partial list):

- hana_primary_root
- CPI_shared
- Q51_data_mnt00001
- Q51_log_mnt00001
- Q51_shared
- SM1_data_mnt00001
- SM1_log_backup
- SM1_log_mnt00001
- SM1_shared
- SS1_data_mnt00001
- SS1_log_backup
- SS1_log_mnt00001
- SS1_repair_log_backup
- SS1_repair_log_mnt00001
- SS1_shared
- SS2_backup
- SS2_data_mnt00001
- SS2_log_mnt00001

Previous Next

Resource protection configuration can now be done in the same way as for auto discovered HANA systems.

Learn about backup operations for SAP HANA Snapshot in SnapCenter

Perform SAP HANA Snapshot backups using SnapCenter. Learn about database Snapshot backups, block integrity checks, non-data volume backups, and backup replication using SnapVault or Azure NetApp Files backup.

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.

SnapCenter supports the following backup operations.

- HANA database Snapshot backup operations
- Block integrity check operations
- Snapshot backups of non-data volumes
- Backup replication using SnapVault or ANF backup for HANA database or non-data volume backups

The following sections describe the different operations for single-host HANA systems which have been auto discovered by SnapCenter (HANA plug-in deployed at the HANA database host)

SAP HANA Snapshot backups in SnapCenter

The SnapCenter resource topology shows the list of backups created by SnapCenter. The following figure shows the backups available on the primary storage and highlights the most recent backup.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-1_LocalSnap_Hourly_08-22-2025_08.00.00.3884		1	08/22/2025 8:01:02 AM
SnapCenter_hana-1_LocalSnapAndSnapshotVault_Daily_08-22-2025_06.00.00.4473		1	08/22/2025 6:01:02 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-22-2025_04.00.00.3795		1	08/22/2025 4:01:05 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-22-2025_06.00.00.3445		1	08/22/2025 12:01:02 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_20.00.00.9333		1	08/21/2025 8:01:01 PM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_16.00.00.4185		1	08/21/2025 4:01:01 PM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_12.00.00.4302		1	08/21/2025 12:01:02 PM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_08.00.00.3794		1	08/21/2025 8:01:02 AM
SnapCenter_hana-1_LocalSnapAndSnapshotVault_Daily_08-21-2025_06.00.00.3982		1	08/21/2025 6:01:04 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_04.00.00.4068		1	08/21/2025 4:01:01 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_00.00.00.3702		1	08/21/2025 12:01:02 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-20-2025_20.00.00.3830		1	08/20/2025 8:01:02 PM
SnapCenter_hana-1_LocalSnap_Hourly_08-20-2025_16.00.00.3981		1	08/20/2025 4:01:01 PM
SnapCenter_hana-1_LocalSnap_Hourly_08-20-2025_12.00.00.4105		1	08/20/2025 12:01:02 PM

Backups at the secondary storage can be listed by clicking on the Vault copies icon.

SS1 Topology

Manage Copies

Secondary Vault Backup(s)

Backup Name

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-22-2025_06.00.00.4473	08/22/2025 6:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-21-2025_06.00.00.3982	08/21/2025 6:01:04 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-20-2025_06.00.00.4188	08/20/2025 6:01:01 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-19-2025_06.00.00.4190	08/19/2025 6:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-18-2025_06.00.00.3538	08/18/2025 6:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-17-2025_06.00.00.4405	08/17/2025 6:00:59 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-01-2025_06.00.00.4123	08/01/2025 6:01:07 AM	1	

Total 7

Activity: The 5 most recent jobs are displayed. 5 Completed, 0 Warnings, 0 Failed, 0 Canceled, 0 Running, 0 Queued.

The following screenshot shows the list of backups for the system SM1, where tamperproof Snapshots have been configured.

SM1 Topology

Manage Copies

Primary Backup(s)

Backup Name

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-2_LocalSnap_Hourly_08-22-2025_08.00.00.4043	08/22/2025 8:01:04 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-22-2025_04.00.00.3939	08/22/2025 4:01:04 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-22-2025_00.00.00.3606	08/22/2025 12:01:04 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-21-2025_20.00.00.4408	08/21/2025 8:01:04 PM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-21-2025_16.00.00.4216	08/21/2025 4:01:04 PM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-21-2025_12.00.00.4306	08/21/2025 12:01:04 PM	1	
SnapCenter_hana-2_LocalSnapTamperProof_Daily_08-21-2025_11.09.00.4474	08/21/2025 11:10:02 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-21-2025_08.00.00.4168	08/21/2025 8:01:04 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-21-2025_04.00.00.4056	08/21/2025 4:01:05 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-21-2025_00.00.00.3702	08/21/2025 12:01:05 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-20-2025_20.00.00.4118	08/20/2025 8:01:04 PM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-20-2025_16.00.00.4319	08/20/2025 4:01:05 PM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-20-2025_12.00.00.4105	08/20/2025 12:01:05 PM	1	
SnapCenter_hana-2_LocalSnapTamperProof_Daily_08-20-2025_11.09.00.4182	08/20/2025 11:10:02 AM	1	
SnapCenter_hana-2_LocalSnap_Hourly_08-19-2025_11.09.00.3746	08/19/2025 11:10:02 AM	1	

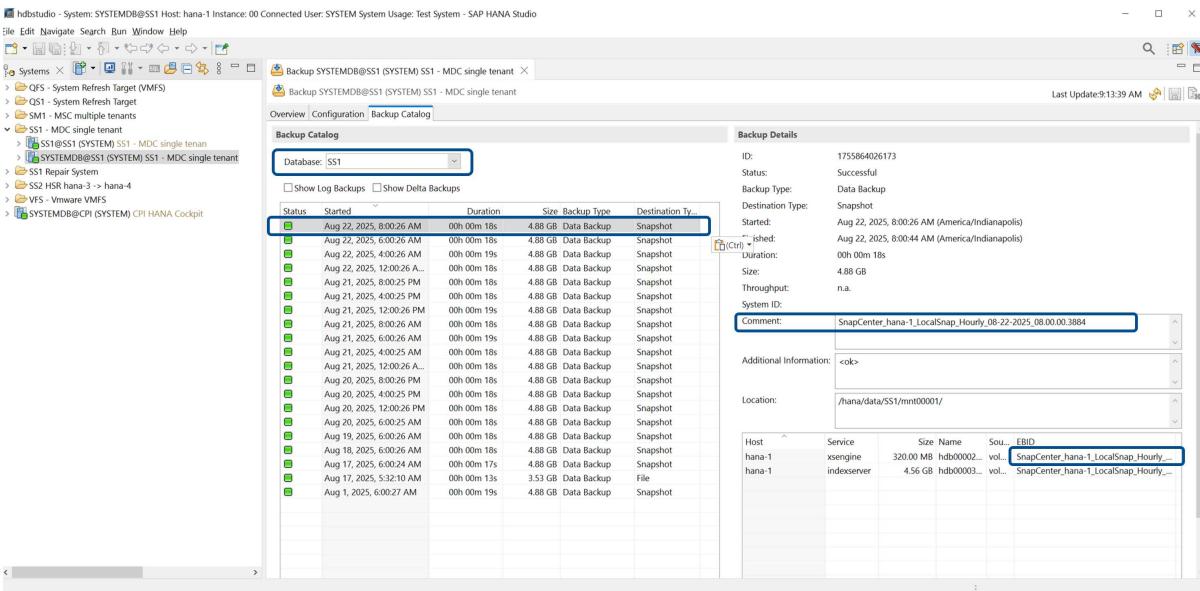
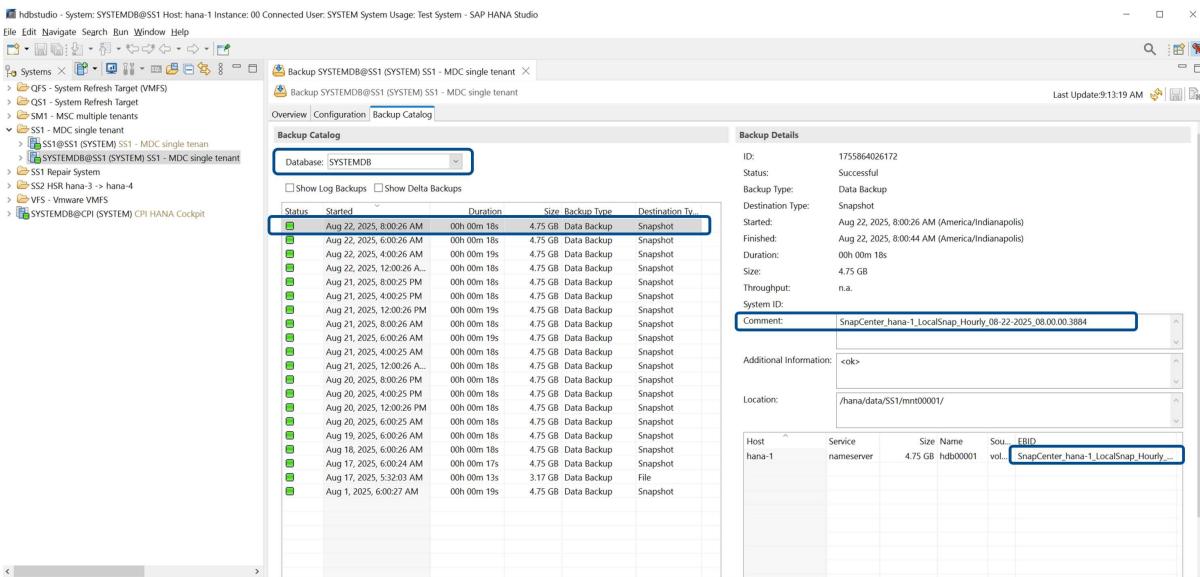
Total 15

Activity: The 5 most recent jobs are displayed. 5 Completed, 0 Warnings, 0 Failed, 0 Canceled, 0 Running, 0 Queued.

SAP HANA Snapshot backups in SAP HANA Studio

When performing a backup using storage Snapshots 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 internal database snapshot 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.

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.



i 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. Also Snapshots created directly on the storage system will not be visible in SnapCenter,

SAP HANA Snapshot backups on storage layer

To view the backups on the storage layer, you can use NetApp System Manager and select the database 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 'Solutions' section of the ONTAP System Manager interface. The left sidebar is collapsed. The main area displays the 'SS1_data_mnt00001' volume details. The 'Solutions' tab is selected, and the 'Snapshots' sub-tab is active. The table lists 17 snapshots, each with a checkbox, name, snapshot creation time, and snapshot restore size. The first snapshot, 'SnapCenter_hana_1_LocalSnap_Hourly_08-22-2025_08.00.00.3884', is highlighted with a blue border. The table has a header row with columns for Name, Snapshot creation time, and Snapshot restore size. The bottom of the table shows a page number '1 - 17 of 17' and navigation icons.

Name	Snapshot creation time	Snapshot restore size
SnapCenter_hana_1_LocalSnap_Hourly_08-22-2025_08.00.00.3884	Aug/22/2025 8:00 AM	8.72 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-22-2025_06.00.00.4473	Aug/22/2025 6:00 AM	8.54 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-22-2025_04.00.00.3795	Aug/22/2025 4:00 AM	8.59 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-22-2025_00.00.00.3445	Aug/22/2025 12:00 AM	8.66 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-21-2025_20.00.00.3933	Aug/21/2025 8:00 PM	8.66 GB
snapshotmirror_1848f339-2e8d-11ec-86b0-d039eaef942_2155360509.2025-08-22_001000	Aug/21/2025 6:10 PM	8.51 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-21-2025_16.00.00.4185	Aug/21/2025 4:00 PM	8.51 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-21-2025_12.00.00.4102	Aug/21/2025 12:00 PM	8.65 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-21-2025_08.00.00.3794	Aug/21/2025 8:00 AM	8.51 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-21-2025_06.00.00.3982	Aug/21/2025 6:00 AM	8.59 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-21-2025_04.00.00.4068	Aug/21/2025 4:00 AM	8.6 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-21-2025_00.00.00.3702	Aug/21/2025 12:00 AM	8.62 GB
SnapCenter_hana_1_LocalSnap_Hourly_08-20-2025_20.00.00.3830	Aug/20/2025 8:00 PM	8.46 GB

The following screenshot shows the available backups for the replication target volume hana_SS1_data_mnt00001_dest at the secondary storage system.

The screenshot shows the 'Solutions' section of the ONTAP System Manager interface. The left sidebar is collapsed. The main area displays the 'SS1_data_mnt00001...' volume details. The 'Solutions' tab is selected, and the 'Snapshots' sub-tab is active. The table lists 7 snapshots, each with a checkbox, name, snapshot creation time, and snapshot restore size. The first snapshot, 'SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-22-2025_06.00.00.4473', is highlighted with a blue border. The table has a header row with columns for Name, Snapshot creation time, and Snapshot restore size. The bottom of the table shows a page number '1 - 7 of 7' and navigation icons.

Name	Snapshot creation time	Snapshot restore size
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-22-2025_06.00.00.4473	Aug/22/2025 6:00 AM	8.52 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-21-2025_06.00.00.3982	Aug/21/2025 6:00 AM	8.61 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-20-2025_06.00.00.4188	Aug/20/2025 6:00 AM	8.67 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-19-2025_06.00.00.4190	Aug/19/2025 6:00 AM	8.64 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-18-2025_06.00.00.3538	Aug/18/2025 6:00 AM	8.64 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-17-2025_06.00.00.4495	Aug/17/2025 6:00 AM	8.48 GB
SnapCenter_hana_1_LocalSnapAndSnapVault_Daily_08-01-2025_06.00.00.4123	Aug/1/2025 6:00 AM	8.65 GB

SAP HANA Snapshot backups with ANF

The following screenshot shows the topology view of a HANA system using Azure NetApp Files. For this HANA system local Snapshot backups as well as backup replication using ANF backup has been configured.

Summary Card

- 22 Backups
- 21 Snapshot based backups
- 1 File-based backup
- 0 Clones

Backup Name	Count	End Date
SnapCenter_vm-pr1_LocalSnap_Hourly_08-22-2025_10_02_13_3629	1	08/22/2025 10:03:47 AM
SnapCenter_vm-pr1_LocalSnapAndANFBackup_Daily_08-22-2025_08_36_12_9809	1	08/22/2025 8:37:41 AM
SnapCenter_vm-pr1_LocalSnapAndANFBackup_Daily_08-22-2025_08_25_09_0607	1	08/22/2025 8:26:48 AM
SnapCenter_vm-pr1_LocalSnap_Hourly_08-21-2025_07_59_43_9453	1	08/21/2025 8:01:19 AM
SnapCenter_vm-pr1_LocalSnap_Hourly_07-07-2025_14_08_02_2735	1	07/07/2025 2:09:24 PM
SnapCenter_vm-pr1_LocalSnap_Hourly_07-07-2025_10_08_02_4762	1	07/07/2025 10:11:00 AM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_06_08_02_2400	1	06/13/2025 6:09:30 AM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_07_08_02_2444	1	06/13/2025 2:09:24 AM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_21_08_02_2583	1	06/12/2025 10:09:23 PM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_18_08_02_2438	1	06/12/2025 6:09:19 PM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_14_08_02_2565	1	06/12/2025 2:09:30 PM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_10_08_02_3048	1	06/13/2025 10:09:19 AM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-12-2025_08_12_18_4446	1	06/12/2025 8:13:45 AM
SnapCenter_vm-pr1_LocalSnap_Hourly_06-11-2025_14_08_02_2870	1	06/11/2025 2:09:16 PM

Total 4

Activity: The 5 most recent jobs are displayed

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

Snapshot backups on the ANF volume can be listed using the Azure portal.

Name	Location	Created	Actions
SnapCenter_vm-pr1_LocalSnapAndANFBackup_Daily_08-22-2025_08_36_12_9809	East US	8/22/2025, 8:26 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-11-2025_14_08_02_3070	East US	6/11/2025, 8:37 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-12-2025_08_12_18_4446	East US	6/12/2025, 8:13 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-12-2025_10_08_02_3048	East US	6/12/2025, 10:08 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-12-2025_14_08_02_2565	East US	6/12/2025, 2:08 PM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-12-2025_18_08_02_4316	East US	6/12/2025, 6:08 PM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-12-2025_22_08_02_2583	East US	6/12/2025, 10:08 PM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_02_08_02_2449	East US	6/13/2025, 2:08 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_06-13-2025_06_08_04_2400	East US	6/13/2025, 6:08 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_07-07-2025_10_08_02_4762	East US	7/7/2025, 10:10 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_07-07-2025_14_08_02_2735	East US	7/7/2025, 2:08 PM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_08-21-2025_07_59_43_9453	East US	8/21/2025, 8:00 AM	...
SnapCenter_vm-pr1_LocalSnap_Hourly_08-22-2025_10_02_13_3629	East US	8/22/2025, 10:03 AM	...

By clicking on the backup icon, you can list the backups which have been replicated with ANF backup.

ANF backups can also be listed in the Azure portal.

Snapshot backups of non-data volumes

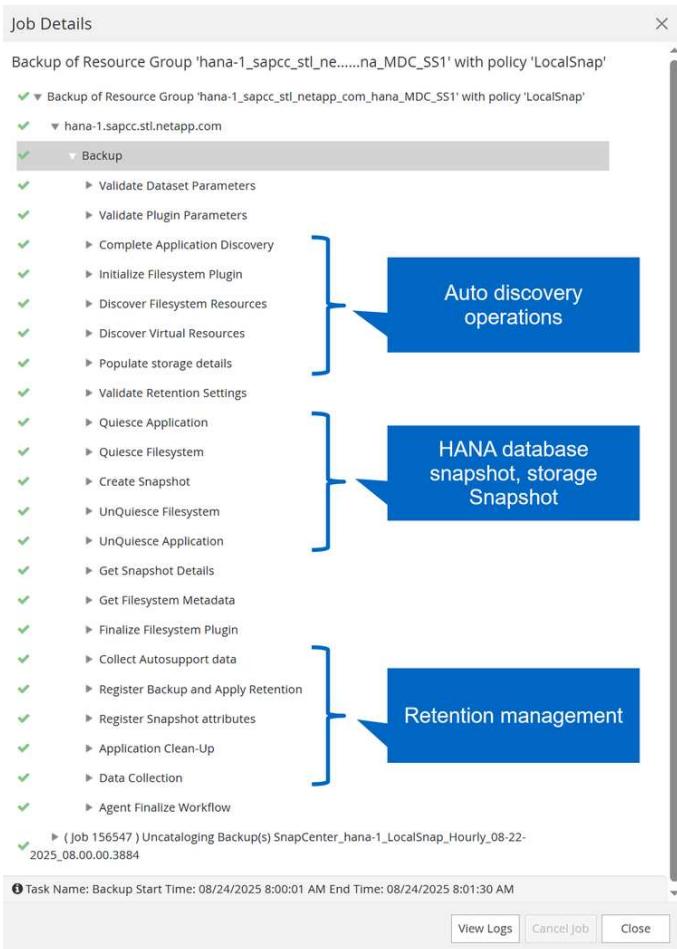
The SnapCenter resource topology shows the list of backups for non-data volumes. In the following figure the backups of the HANA shared volume are listed.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_LocalSnap_Hourly_09-26-2025_06.04.00.3639	09/26/2025 6:04:14 AM	1	
SnapCenter_LocalSnap_Hourly_09-26-2025_06.04.00.3954	09/26/2025 6:04:14 AM	1	
SnapCenter_LocalSnap_Hourly_09-25-2025_22.04.00.3595	09/25/2025 10:04:14 PM	1	
SnapCenter_LocalSnap_Hourly_09-25-2025_18.04.00.4000	09/25/2025 6:04:13 PM	1	
SnapCenter_LocalSnap_Hourly_09-25-2025_14.04.00.3911	09/25/2025 2:04:14 PM	1	
SnapCenter_LocalSnap_Hourly_09-25-2025_10.04.00.3938	09/25/2025 10:04:13 AM	1	
SnapCenter_LocalSnap_Hourly_09-25-2025_06.04.00.4101	09/25/2025 6:04:14 AM	1	
SnapCenter_LocalSnap_Hourly_09-25-2025_02.04.00.3677	09/25/2025 2:04:13 AM	1	
SnapCenter_LocalSnap_Hourly_09-25-2025_22.04.00.4221	09/24/2025 10:04:14 PM	1	
SnapCenter_LocalSnap_Hourly_09-24-2025_18.04.00.4322	09/24/2025 6:04:13 PM	1	
SnapCenter_LocalSnap_Hourly_09-24-2025_14.04.00.4105	09/24/2025 2:04:13 PM	1	
SnapCenter_LocalSnap_Hourly_09-24-2025_10.04.00.4206	09/24/2025 10:04:14 AM	1	

Backup workflow for HANA database backups

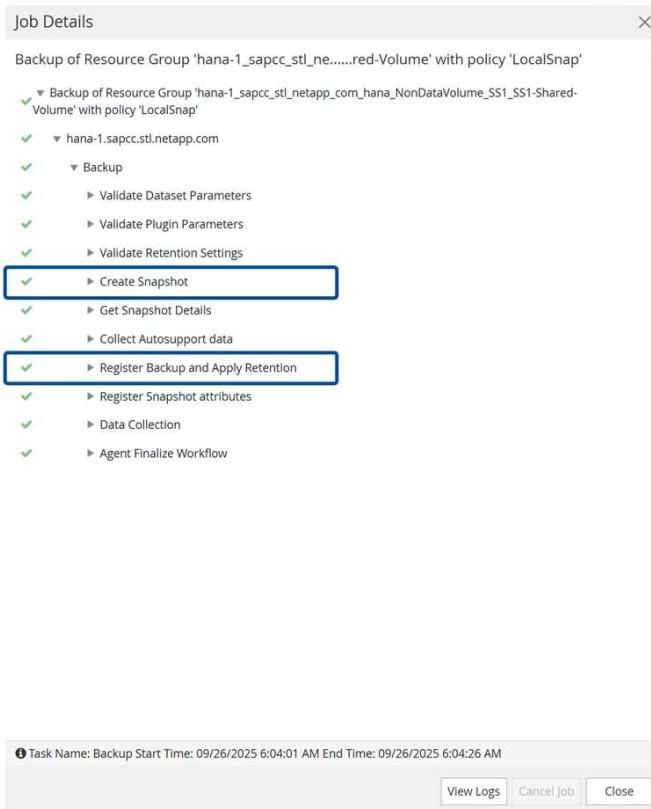
The backup workflow for a HANA database Snapshot backup consists of three main sections.

- Auto discovery
 - Application discovery, e.g.
 - SnapCenter detects any tenant configuration changes
 - SnapCenter detects HANA system replication primary node
 - File system and storage discovery, e.g.
 - SnapCenter detects any changes in volume configuration
 - SnapCenter detects HANA multiple partition configuration
- HANA and Snapshot backup operations
 - Trigger HANA database snapshot
 - Create storage Snapshot
 - Confirm HANA database snapshot and register backup in HANA backup catalog
- Retention management
 - Delete Snapshot backup(s) based on defined retention in
 - SnapCenter repository
 - Storage
 - HANA backup catalog
 - Log backup retention management
 - Delete log backups on file system and HANA backup catalog



Backup workflow for non-data volumes

For a non-data volume, the backup workflow consists of the Snapshot operation and the retention management operation.

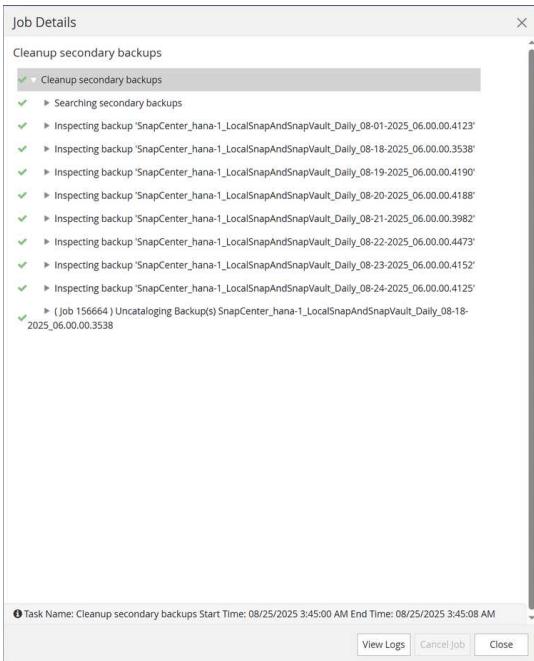


Cleanup of secondary backups

As described in ["Retention management for secondary backups"](#), retention management of data backups to an secondary backup storage is handled by ONTAP. SnapCenter periodically checks if ONTAP has deleted backups at the secondary 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 secondary backup storage have been identified.

ID	Status	Name	Start date	End date	Owner
156670	✓	Backup of Resource Group 'hana-2_sapcc_stl_ne.....ne_MDC_SM1' with policy 'LocalSnap'	08/25/2025 4:00:00 AM	08/25/2025 4:01:35 AM	SAPCCscadmin
156669	✓	Backup of Resource Group 'hana-1_sapcc_stl_ne.....ne_MDC_SS1' with policy 'LocalSnap'	08/25/2025 4:00:00 AM	08/25/2025 4:01:32 AM	SAPCCscadmin
156666	✓	Backup of Resource Group 'SS2 - HANA System R.....ation' with policy 'LocalSnapKeep2'	08/25/2025 3:56:00 AM	08/25/2025 3:57:53 AM	SAPCCscadmin
156663	✓	Cleanup secondary backups	08/25/2025 3:45:00 AM	08/25/2025 3:45:08 AM	SAPCCscadmin
156660	✓	Backup of Resource Group 'SS2 - HANA System R.....ation' with policy 'LocalSnapKeep2'	08/25/2025 2:56:00 AM	08/25/2025 2:57:55 AM	SAPCCscadmin
156657	✓	Backup of Resource Group 'hana-8_sapcc_stl_ne.....ne_MDC_VFS' with policy 'LocalSnap'	08/25/2025 2:29:00 AM	08/25/2025 2:30:43 AM	SAPCCscadmin
156656	✓	Backup of Resource Group 'hana-2_sapcc_stl_ne.....red-Volume' with policy 'LocalSnap'	08/25/2025 2:06:00 AM	08/25/2025 2:06:27 AM	SAPCCscadmin
156655	✓	Backup of Resource Group 'hana-1_sapcc_stl_ne.....red-Volume' with policy 'LocalSnap'	08/25/2025 2:04:00 AM	08/25/2025 2:04:26 AM	SAPCCscadmin



Until this scheduled cleanup has finished, SAP HANA and SnapCenter will still show backups that have already been deleted from the secondary backup storage. This will result in additional log backups that are kept, even if the corresponding storage-based Snapshot backups on the secondary backup storage have already been deleted. NetApp recommends changing the schedule from weekly to daily to avoid keeping log backups, which are not required anymore.

Change the frequency of the SnapCenter cleanup job

SnapCenter executes the cleanup job `SnapCenter_RemoveSecondaryBackup` by default for all resources on a weekly basis. This can be changed using a SnapCenter PowerShell cmdlet.

```

SnapCenterPS C:\> Open-SmConnection

Enter username/password
User: sapcc\scadmin
Password for user sapcc\scadmin: *****

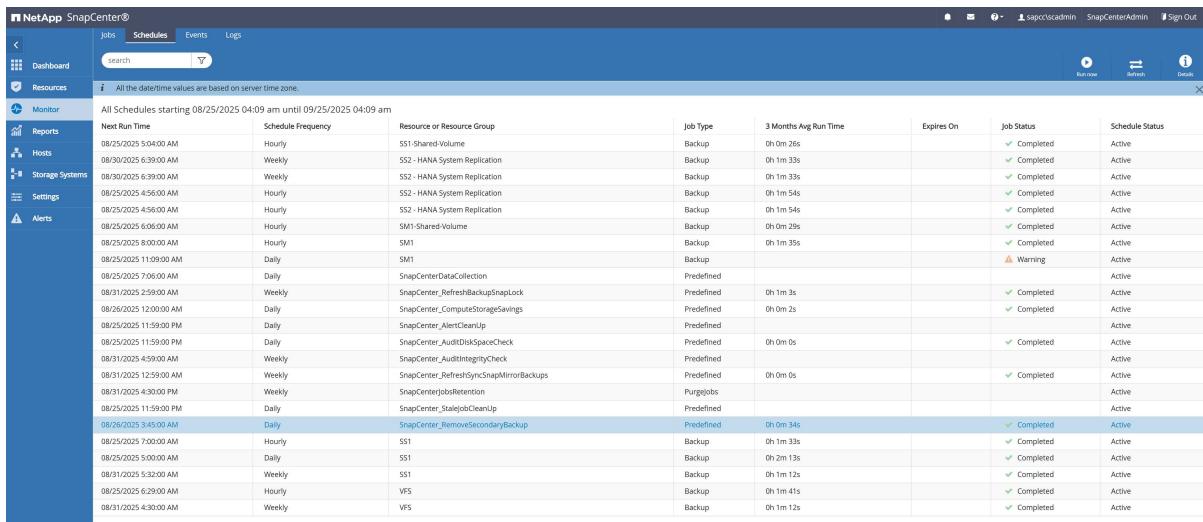
SnapCenterPS C:\> Set-SmSchedule -ScheduleInformation
@{ "ScheduleType"="Daily"; "StartTime"="03:45 AM"; "DaysInterval"="1" }
-TaskName SnapCenter_RemoveSecondaryBackup

TaskName : SnapCenter_RemoveSecondaryBackup
Hosts : {}
StartTime : 8/25/2025 3:45:00 AM
DaysoftheMonth :
MonthsOfTheYear :
DaysInterval : 1
DaysOfTheWeek :
AllowDefaults : False
ReplaceJobIfExist : False

```

```
UserName :  
Password :  
SchedulerType : Daily  
RepeatTask_Every_Hour : 1  
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 :  
ScheduleCategory :  
PolicyId : 0  
PolicyName :  
ProtectionGroupId : 0  
ProtectionGroupName :  
PluginCode : NONE  
PolicyType : None  
ReportTriggerName :  
PolicyScheduleId : 0  
HoursOfTheDay :  
DayStartTime :  
MinuteOffset : ZeroMinutes  
SnapMirrorLabel :  
BackupType :  
SnapCenterPS C:\>
```

The configuration can also be checked in the Monitor - Schedules view in the SnapCenter UI.

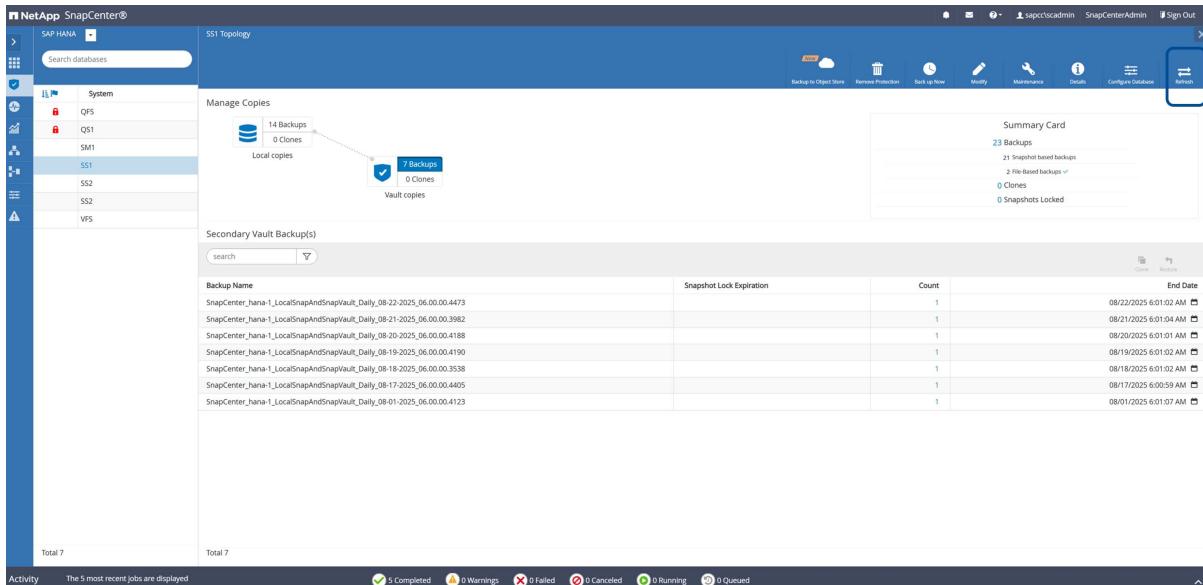


All Schedules starting 08/25/2025 04:09 am until 09/25/2025 04:09 am

Next Run Time	Schedule Frequency	Resource or Resource Group	Job Type	3 Months Avg Run Time	Expires On	Job Status	Schedule Status
08/25/2025 4:04:00 AM	Hourly	S51-Shared-Volume	Backup	0h 0m 26s		✓ Completed	Active
08/25/2025 4:39:00 AM	Weekly	S52 - HANA System Replication	Backup	0h 1m 33s		✓ Completed	Active
08/25/2025 4:39:00 AM	Weekly	S52 - HANA System Replication	Backup	0h 1m 33s		✓ Completed	Active
08/25/2025 4:56:00 AM	Hourly	S52 - HANA System Replication	Backup	0h 1m 54s		✓ Completed	Active
08/25/2025 4:56:00 AM	Hourly	S52 - HANA System Replication	Backup	0h 1m 54s		✓ Completed	Active
08/25/2025 4:56:00 AM	Hourly	S51-Shared-Volume	Backup	0h 0m 29s		✓ Completed	Active
08/25/2025 4:56:00 AM	Hourly	S51	Backup	0h 1m 35s		✓ Completed	Active
08/25/2025 4:59:00 AM	Daily	S51	Backup			⚠ Warning	Active
08/25/2025 5:06:00 AM	Daily	SnapCenter Data Collection	Predefined				Active
08/31/2025 2:59:00 AM	Weekly	SnapCenter_RefreshBackupSnapLock	Predefined	0h 1m 3s		✓ Completed	Active
08/25/2025 12:00:00 AM	Daily	SnapCenter_ComputeStorageSavings	Predefined	0h 0m 2s		✓ Completed	Active
08/25/2025 1:59:00 PM	Daily	SnapCenter_AlertCleanup	Predefined				Active
08/25/2025 1:59:00 PM	Daily	SnapCenter_AuditDiskSpaceCheck	Predefined	0h 0m 0s		✓ Completed	Active
08/31/2025 4:59:00 AM	Weekly	SnapCenter_AuditIntegrityCheck	Predefined				Active
08/31/2025 4:59:00 AM	Weekly	SnapCenter_RefreshSyncSnapMirrorBackups	Predefined	0h 0m 0s		✓ Completed	Active
08/31/2025 4:59:00 AM	Weekly	SnapCenter_JobRetention	PurgeJobs				Active
08/25/2025 1:59:00 PM	Daily	SnapCenter_StageJobCleanup	Predefined				Active
08/25/2025 3:45:00 AM	Daily	SnapCenter_RemoveSecondaryBackup	Predefined	0h 0m 34s		✓ Completed	Active
08/25/2025 7:00:00 AM	Hourly	S51	Backup	0h 1m 33s		✓ Completed	Active
08/25/2025 5:00:00 AM	Daily	S51	Backup	0h 2m 13s		✓ Completed	Active
08/31/2025 3:32:00 AM	Weekly	S51	Backup	0h 1m 12s		✓ Completed	Active
08/25/2025 6:29:00 AM	Hourly	VFS	Backup	0h 1m 41s		✓ Completed	Active
08/31/2025 4:30:00 AM	Weekly	VFS	Backup	0h 1m 12s		✓ Completed	Active

Manual refresh on resource level

If required, a manual cleanup of secondary backups can also be executed in the topology view of a resource. SnapCenter displays the backups on the secondary 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.



Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-22-2025_06.00.00.4473	08/22/2025 6:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-21-2025_06.00.00.3982	08/21/2025 6:01:04 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-20-2025_06.00.00.4188	08/20/2025 6:01:01 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-19-2025_06.00.00.4190	08/19/2025 6:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-18-2025_06.00.00.3538	08/18/2025 6:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-17-2025_06.00.00.4405	08/17/2025 6:00:59 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-01-2025_06.00.00.4123	08/01/2025 6:01:07 AM	1	

Execute SAP HANA block consistency checks with SnapCenter

Execute SAP HANA block consistency checks using the SAP hdbpersdiag tool or by executing file-based backups. Learn about configuration options including local Snapshot directory access, central verification hosts with FlexClone volumes, and SnapCenter integration for scheduling and automation.

The table below summarizes the key parameters helping to decide which method for block consistency checks fits best for your environment.

	HANA hdbpersdiag tool using local Snapshot directory	HANA hdbpersdiag tool with central verification host	File-based backup
Supported configurations	NFS only Bare metal, ANF, FSx ONTAP, VMware or KVM in-guest mounts	All protocols and platforms	All protocols and platforms
CPU load at HANA host	Medium	None	High
Network utilization at HANA host	High	None	High
Runtime	Leverages full read throughput of storage volume	Leverages full read throughput of storage volume	Typically limited by write throughput of target system
Capacity requirements	None	None	At least 1 x backup size per HANA system
SnapCenter integration	Post backup script	Clone create and post cloning script, clone delete	Build-in feature
Scheduling	SnapCenter scheduler	PowerShell script to execute clone create and delete workflow, externally scheduled	SnapCenter scheduler

The following chapters describe the configuration and execution of the different options for block consistency check operations.

Consistency checks with hdbpersdiag using the local snapshot directory

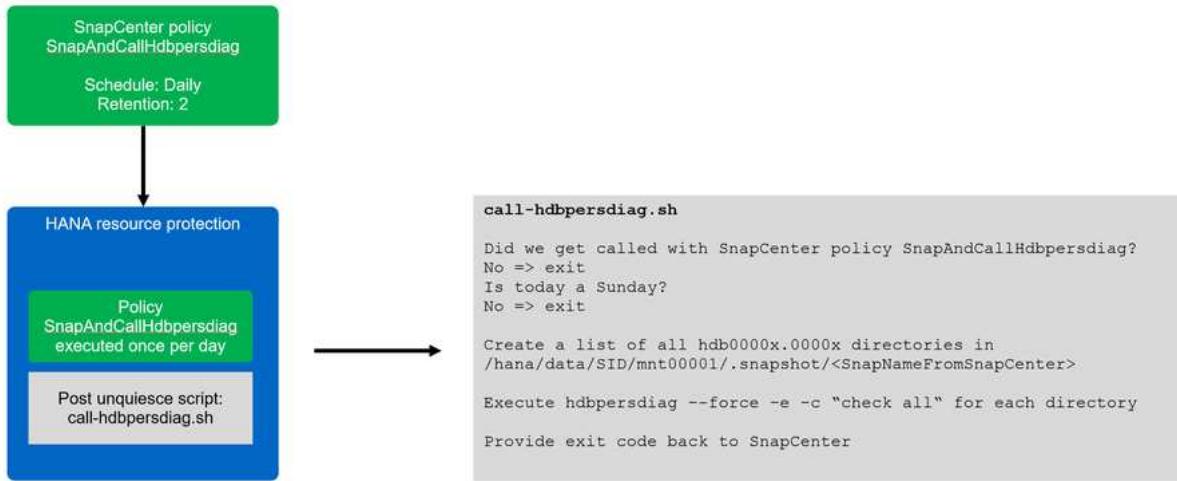
Within SnapCenter a dedicated policy for hdbpersdiag operations is created with a daily schedule and a retention of two. We don't use the weekly schedule, since we would then have at least 2 Snapshot backups (minimum retention=2), where one of them would be up to two weeks old.

Within the SnapCenter resource protection configuration of the HANA system, a post backup script is added, which executes the hdbpersdiag tool. Since the post backup script will be also called with any other policy configured for the resource, we need to check in the script which policy is currently active. Within the script we also check the current day of the week and run the hdbpersdiag operation only once per week on a Sunday. HANA hdbpersdiag is then called for each data volume in the corresponding hdb* directory of the current Snapshot backup directory. If the consistency check with hdbpersdiag reports any error the SnapCenter job will be marked as failed.



The example script call-hdbpersdiag.sh is provided as is and is not covered by NetApp support. You can request the script via email to ng-sapcc@netapp.com.

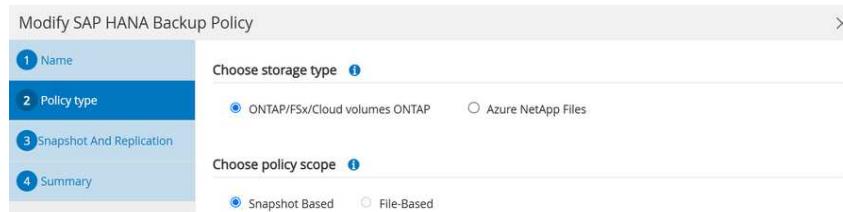
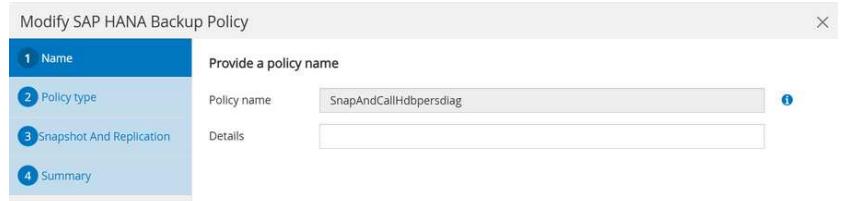
The figure below shows the high-level concept of the consistency check implementation.



As a first step you need to allow access to the snapshot directory, so that the ".snapshot" directory is visible at the HANA database host.

- ONTAP systems and FSX for ONTAP: You need to configure the Snapshot directory access volume parameter
- ANF: You need to configure the Hide Snapshot path volume parameter.

As a next step, you must configure a policy which matches the name that is used in the post backup script. For our script example the name must be SnapAndCallHdbpersdiag. As discussed before a daily schedule is used to avoid keeping old Snapshots with a weekly schedule.



Modify SAP HANA Backup Policy

1 Name

2 Policy type

3 Snapshot And Replication

4 Summary

Choose 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

Snapshot settings

Copies to keep copies 1

Retain copies for days

Primary snapshot copy locking period days 1

Secondary snapshot copy locking period days 1

Policy label 1

Select secondary replication options 1

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Error retry count 1

Within the resource protection configuration, the post backup script is added, and the policy is assigned to the

Post Quiesce

Enter commands to be executed before and after creating Snapshot copies 1

Pre Snapshot Copy

Post Snapshot Copy

Enter commands to be executed before and after returning the application to normal operational state 1

Pre UnQuiesce

Post UnQuiesce

Enter commands to be executed before exiting in the event of a failure 1

Pre Exit

Custom Configurations

Activity: The 5 most recent jobs are displayed

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

Finally, the script must be configured in the `allowed_commands.config` file at the HANA host.

```
hana-1:/ # cat /opt/NetApp/snapcenter/scc/etc/allowed_commands.config
command: mount
command: umount
command: /mnt/sapcc-share/hdbpersdiag/call-hdbpersdiag.sh
```

The Snapshot backup operation will now be executed once per day, and the script handles that the `hdbpersdiag` check is only executed once per week on Sundays.



The script calls `hdbpersdiag` with the “-e” command line option which is required for data volume encryption. If HANA data volume encryption is not used the parameter must be removed.

The output below shows the log file of the script:

```
20251024055824##hana-1##call-hdbpersdiag.sh: Current policy is
SnapAndCallHdbpersdiag
20251024055824##hana-1##call-hdbpersdiag.sh: Executing hdbpersdiag in:
/hana/data/SS1/mnt00001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00001
20251024055827##hana-1##call-hdbpersdiag.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/SS1/HDB00/hana-1/trace
Mounted DataVolume(s)
#0 /hana/data/SS1/mnt00001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00001/ (4.8 GB,
5100273664 bytes)
WARNING: The data volume being accessed is in use by another process, this
is most likely because a running HANA instance is operating on this data
volume
```

Tips:

```
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '||' to redirect the output to a specific command.
INFO: KeyPage loaded and decrypted with success
Default Anchor Page OK
Restart Page OK
Default Converter Pages OK
RowStore Converter Pages OK
Logical Pages (94276 pages) OK
Logical Pages Linkage OK
Checking entries from restart page...
ContainerDirectory OK
ContainerNameDirectory OK
FileIDMappingContainer OK
UndoContainerDirectory OK
LobDirectory OK
MidSizeLobDirectory OK
LobFileIDMap OK
20251024055827##hana-1##call-hdbpersdiag.sh: Consistency check operation
successesful for volume /hana/data/SS1/mnt00001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00001.
20251024055827##hana-1##call-hdbpersdiag.sh: Executing hdbpersdiag in:
/hana/data/SS1/mnt00001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00002.00003
20251024055828##hana-1##call-hdbpersdiag.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/SS1/HDB00/hana-1/trace
Mounted DataVolume(s)
#0 /hana/data/SS1/mnt00001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00002.00003/
(320.0 MB, 335544320 bytes)
WARNING: The data volume being accessed is in use by another process, this
is most likely because a running HANA instance is operating on this data
volume
Tips:
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '||' to redirect the output to a specific command.
INFO: KeyPage loaded and decrypted with success
Default Anchor Page OK
Restart Page OK
Default Converter Pages OK
RowStore Converter Pages OK
Logical Pages (4099 pages) OK
```

```
Logical Pages Linkage OK
Checking entries from restart page...
UndoContainerDirectory OK
DRLoadedTable OK
20251024055828##hana-1##call-hdbpersdiag.sh: Consistency check operation
successesful for volume /hana/data/SS1/mnt0001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00002.00003.
20251024055828##hana-1##call-hdbpersdiag.sh: Executing hdbpersdiag in:
/hana/data/SS1/mnt0001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00003.00003
20251024055833##hana-1##call-hdbpersdiag.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/SS1/HDB00/hana-1/trace
Mounted DataVolume(s)
#0 /hana/data/SS1/mnt0001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00003.00003/
(4.6 GB, 4898947072 bytes)
WARNING: The data volume being accessed is in use by another process, this
is most likely because a running HANA instance is operating on this data
volume
Tips:
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '||' to redirect the output to a specific command.
INFO: KeyPage loaded and decrypted with success
Default Anchor Page OK
Restart Page OK
Default Converter Pages OK
Static Converter Pages OK
RowStore Converter Pages OK
Logical Pages (100817 pages) OK
Logical Pages Linkage OK
Checking entries from restart page...
ContainerDirectory OK
ContainerNameDirectory OK
FileIDMappingContainer OK
UndoContainerDirectory OK
LobDirectory OK
DRLoadedTable OK
MidSizeLobDirectory OK
LobFileIDMap OK
20251024055833##hana-1##call-hdbpersdiag.sh: Consistency check operation
successesful for volume /hana/data/SS1/mnt0001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00003.00003.
20251024060048##hana-1##call-hdbpersdiag.sh: Current policy is
```

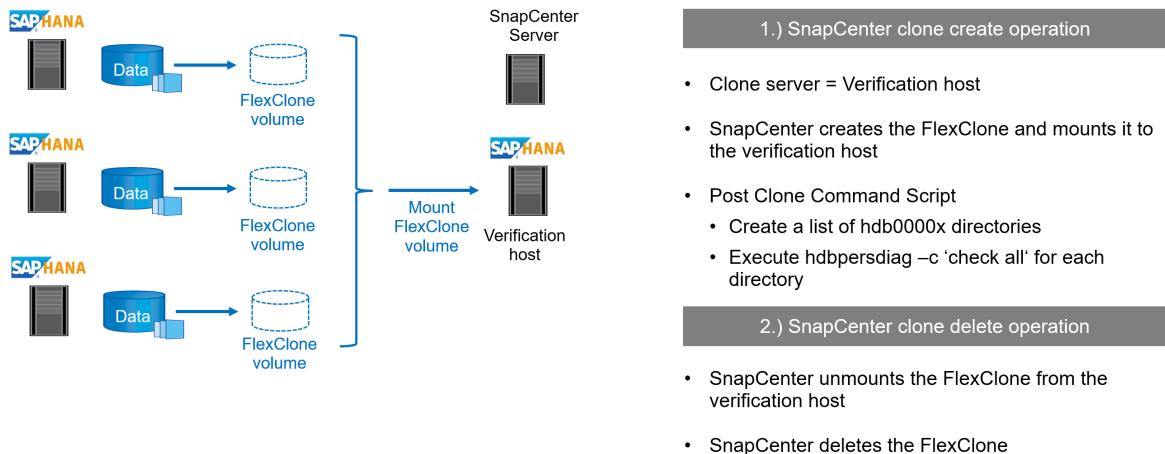
```
LocalSnapAndSnapVault, consistency check is only done with Policy
SnapAndCallHdbpersdiag
20251024080048##hana-1##call-hdbpersdiag.sh: Current policy is
LocalSnap, consistency check is only done with Policy SnapAndHdbpersdiag
```

Consistency checks with hdbpersdiag using a central verification host

The figure below shows a high-level view of the solution architecture and workflow. With a central verification host the verification host can be used to check the consistency of multiple, different HANA systems. The solution leverages the SnapCenter clone create and delete workflows to attach a cloned volume from the HANA system which should be checked to the verification host. A post clone script is used to run the HANA hdbpersdiag tool. As a second step the SnapCenter clone delete workflow is used to unmount and delete the cloned volume.



If the HANA systems are configured with data volume encryption the encryption root keys of the source HANA system must be imported at the verification host before hdbpersdiag is executed. See also [Import Backed-Up Root Keys Before Database Recovery | SAP Help Portal](#)



The HANA tool hdbpersdiag is included in each HANA installation but is not available as a standalone tool. Hence the central verification host must be prepared by installing a normal HANA system.

Initial one-time preparation steps:

- Installation of SAP HANA system to be used as central verification host
- Configuration of SAP HANA system in SnapCenter
 - Deployment of SnapCenter SAP HANA plug-in at verification host. SAP HANA system is auto discovered by SnapCenter.
- The first hdbpersdiag operation after the initial installation is prepared with the following steps:
 - Shutdown target SAP HANA system
 - 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:/mnt/sapcc-share/hdbpersdiag # cat
/opt/NetApp/snapcenter/scc/etc/allowed_commands.config
command: mount
command: umount
command: /mnt/sapcc-share/hdbpersdiag/call-hdbpersdiag-fleclone.sh

```

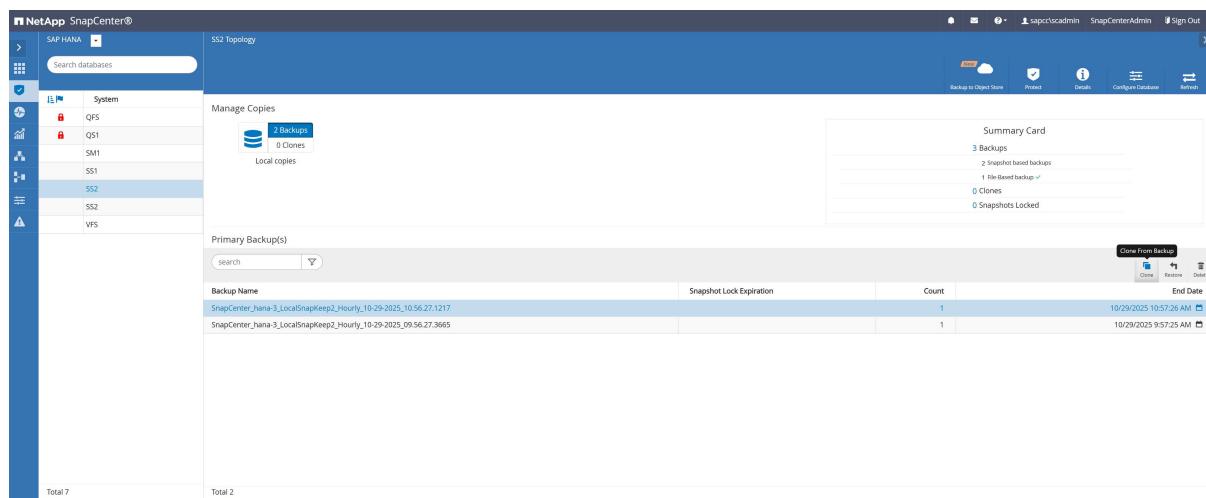


The example script `call-hdbpersdiag-fleclone.sh` is provided as is and is not covered by NetApp support. You can request the script via email to ng-sapcc@netapp.com.

Manual workflow execution

In most cases, the consistency check operation will be run as a scheduled operation as described in the next chapter. However, being aware of the manual workflow is helpful to understand the parameters which are used for the automated process.

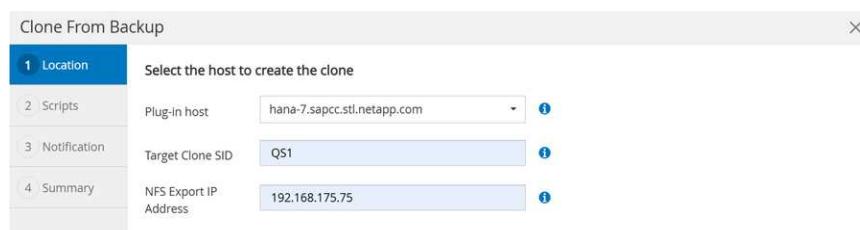
The clone create workflow is started by selecting a backup from the system which should be checked and by clicking on clone from backup.



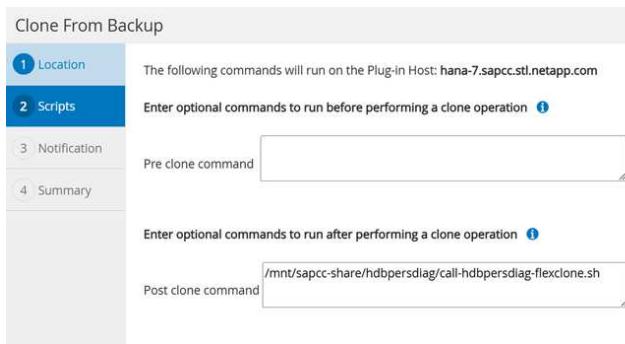
In the next screen the host name, SID and storage network interface of the verification host must be provided.



It is important to always use the SID of the HANA system installed at the verification host, otherwise the workflow will fail.



In the next screen you need to add the `call-hdbpersdiag-fleclone.sh` script as a post clone command.



When the workflow is started, SnapCenter will create a cloned volume based on the selected Snapshot backup and mount it to the verification host.

Note: The example output below is based on HANA systems using NFS as the storage protocol. For HANA system using FC or VMware VMDKs the device will be mounted in the same way to /hana/data/SID/mnt00001.

```

hana-7:/mnt/sapcc-share/hdbpersdiag # df -h
Filesystem Size Used Avail Use% Mounted on
devtmpfs 16G 8.0K 16G 1% /dev
tmpfs 25G 0 25G 0% /dev/shm
tmpfs 16G 474M 16G 3% /run
tmpfs 16G 0 16G 0% /sys/fs/cgroup
/dev/mapper/system-root 60G 9.0G 48G 16% /
/dev/mapper/system-root 60G 9.0G 48G 16% /home
/dev/mapper/system-root 60G 9.0G 48G 16% /.snapshots
/dev/mapper/system-root 60G 9.0G 48G 16% /root
/dev/mapper/system-root 60G 9.0G 48G 16% /opt
/dev/mapper/system-root 60G 9.0G 48G 16% /boot/grub2/i386-pc
/dev/mapper/system-root 60G 9.0G 48G 16% /srv
/dev/mapper/system-root 60G 9.0G 48G 16% /usr/local
/dev/mapper/system-root 60G 9.0G 48G 16% /boot/grub2/x86_64-efi
/dev/mapper/system-root 60G 9.0G 48G 16% /var
/dev/mapper/system-root 60G 9.0G 48G 16% /tmp
/dev/sda1 500M 5.1M 495M 2% /boot/efi
192.168.175.117:/QS1_shared/usr-sap 251G 15G 236G 6% /usr/sap/QS1
192.168.175.86:/sapcc_share 1.4T 858G 568G 61% /mnt/sapcc-share
192.168.175.117:/QS1_log_mnt00001 251G 335M 250G 1% /hana/log/QS1/mnt00001
192.168.175.117:/QS1_shared/shared 251G 15G 236G 6% /hana/shared
tmpfs 3.2G 20K 3.2G 1% /run/user/467
tmpfs 3.2G 0 3.2G 0% /run/user/0
192.168.175.117:/SS2_data_mnt00001_Clone_10292511250337819 250G 6.4G 244G
3% /hana/data/QS1/mnt00001

```

The output below shows the log file of the post clone command call-hdbpersdiag-flexclone.sh.

```
20251029112557##hana-7##call-hdbpersdiag-flexclone.sh: Executing
```

```

hdbpersdiag for source system SS2.
20251029112557##hana-7###call-hdbpersdiag-flexclone.sh: Clone mounted at
/hana/data/QS1/mnt00001.
20251029112557##hana-7###call-hdbpersdiag-flexclone.sh: Executing
hdbpersdiag in: /hana/data/QS1/mnt00001/hdb00001
20251029112600##hana-7###call-hdbpersdiag-flexclone.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/QS1/HDB11/hana-7/trace
Mounted DataVolume(s)
#0 /hana/data/QS1/mnt00001/hdb00001/ (3.1 GB, 3361128448 bytes)

Tips:
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '||' to redirect the output to a specific command.
INFO: KeyPage loaded and decrypted with success
Default Anchor Page OK
Restart Page OK
Default Converter Pages OK
RowStore Converter Pages OK
Logical Pages (65388 pages) OK
Logical Pages Linkage OK
Checking entries from restart page...
ContainerDirectory OK
ContainerNameDirectory OK
FileIDMappingContainer OK
UndoContainerDirectory OK
LobDirectory OK
MidSizeLobDirectory OK
LobFileIDMap OK
20251029112600##hana-7###call-hdbpersdiag-flexclone.sh: Consistency check
operation successful for volume /hana/data/QS1/mnt00001/hdb00001.
20251029112601##hana-7###call-hdbpersdiag-flexclone.sh: Executing
hdbpersdiag in: /hana/data/QS1/mnt00001/hdb00002.00003
20251029112602##hana-7###call-hdbpersdiag-flexclone.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/QS1/HDB11/hana-7/trace
Mounted DataVolume(s)
#0 /hana/data/QS1/mnt00001/hdb00002.00003/ (288.0 MB, 301989888 bytes)

Tips:
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '||' to redirect the output to a specific command.
INFO: KeyPage loaded and decrypted with success
Default Anchor Page OK

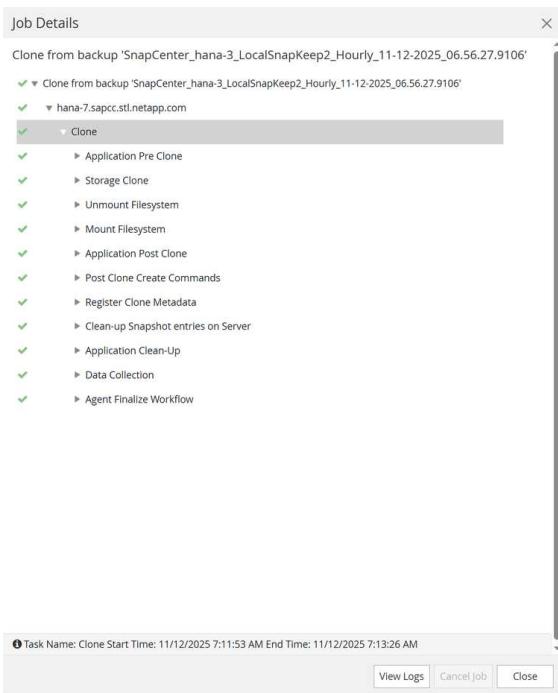
```

```

Restart Page OK
Default Converter Pages OK
RowStore Converter Pages OK
Logical Pages (4099 pages) OK
Logical Pages Linkage OK
Checking entries from restart page...
UndoContainerDirectory OK
DRLoadedTable OK
20251029112602##hana-7##call-hdbpersdiag-flexclone.sh: Consistency check
operation successful for volume /hana/data/QS1/mnt00001/hdb00002.00003.
20251029112602##hana-7##call-hdbpersdiag-flexclone.sh: Executing
hdbpersdiag in: /hana/data/QS1/mnt00001/hdb00003.00003
20251029112606##hana-7##call-hdbpersdiag-flexclone.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/QS1/HDB11/hana-7/trace
Mounted DataVolume(s)
#0 /hana/data/QS1/mnt00001/hdb00003.00003/ (3.7 GB, 3942645760 bytes)
Tips:
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '||' to redirect the output to a specific command.
INFO: KeyPage loaded and decrypted with success
Default Anchor Page OK
Restart Page OK
Default Converter Pages OK
Static Converter Pages OK
RowStore Converter Pages OK
Logical Pages (79333 pages) OK
Logical Pages Linkage OK
Checking entries from restart page...
ContainerDirectory OK
ContainerNameDirectory OK
FileIDMappingContainer OK
UndoContainerDirectory OK
LobDirectory OK
DRLoadedTable OK
MidSizeLobDirectory OK
LobFileIDMap OK
20251029112606##hana-7##call-hdbpersdiag-flexclone.sh: Consistency check
operation successful for volume /hana/data/QS1/mnt00001/hdb00003.00003.

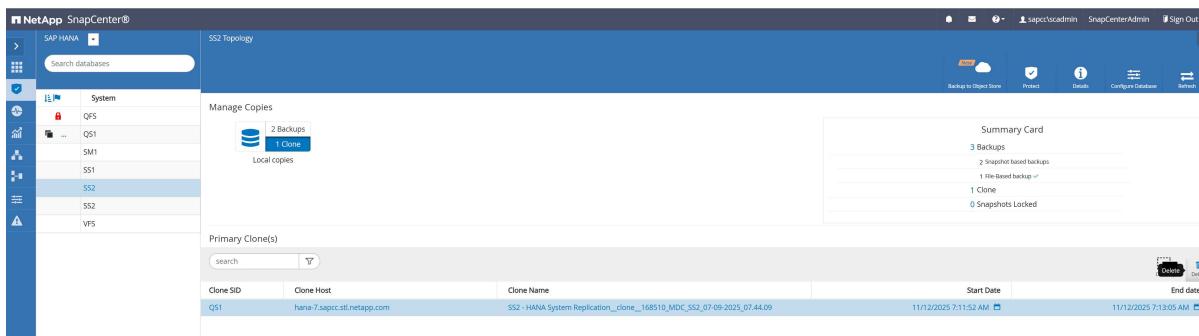
```

 The script calls hdbpersdiag with the “-e” command line option which is required for data volume encryption. If HANA data volume encryption is not used the parameter must be removed. When the post clone script is finished the SnapCenter job is finished as well.



As the next step we will run the SnapCenter clone delete workflow to cleanup the verification host and to delete the FlexClone volume.

In the topology view of the source system, we select the clone and click the delete button.



SnapCenter will now unmount the cloned volume from the verification host and will delete the cloned volume at the storage system.

SnapCenter workflow automation using PowerShell scripts

In the previous section, the clone create and clone delete workflows were executed using SnapCenter UI. All the workflows can also be executed with PowerShell scripts or REST API calls, allowing further automation. The following section describes a basic PowerShell script example to execute the SnapCenter clone create and clone delete workflows.



The example script call-hdbpersdiag-flexclone.sh and clone-hdbpersdiag.ps1 are provided as is and are not covered by NetApp support. You can request the scripts via email to ng-sapcc@netapp.com.

The PowerShell example script executes the following workflow.

- Search for the latest Snapshot backup according to the command line parameter SID and source host

- Executes the SnapCenter clone create workflow using the Snapshot backup defined in the step before. Target host information and hdbpersdiag information is defined in the script. The call-hdbpersdiag-flexclone.sh script is defined as a post clone script and is executed at the target host.
 - \$result = New-SmClone -AppPluginCode hana -BackupName \$backupName -Resources @{\$"Host"="\$sourceHost";"UID"="\$uid"} -CloneToInstance "\$verificationHost" -NFSExportIPs \$exportIpTarget -CloneUid \$targetUid -PostCloneCreateCommands \$postCloneScript

- Executes the SnapCenter clone delete workflow

The text below shows the output of the example script executed at the SnapCenter server.

The text below shows the output of the example script executed at the SnapCenter server.

```
C:\Users\scadmin>pwsh -command "c:\netapp\clone-hdbpersdiag.ps1 -sid SS2
-sourceHost hana-3.sapcc.st1.netapp.com"
Starting verification
Connecting to SnapCenter
Validating clone/verification request - check for already existing clones
Get latest back for [SS2] on host [hana-3.sapcc.st1.netapp.com]
Found backup name [SnapCenter_hana-3_LocalSnapKeep2_Hourly_11-21-
2025_07.56.27.5547]
Creating clone from backup [hana-
3.sapcc.st1.netapp.com/SS2/SnapCenter_hana-3_LocalSnapKeep2_Hourly_11-21-
2025_07.56.27.5547]: [hana-7.sapcc.st1.netapp.com/QS1]
waiting for job [169851] - [Running]
waiting for job [169851] - [Completed]
Removing clone [SS2 - HANA System Replication__clone__169851_MDC_SS2_07-
09-2025_07.44.09]
waiting for job [169854] - [Running]
waiting for job [169854] - [Completed]
Verification completed

C:\Users\scadmin>
```



The script calls hdbpersdiag with the “-e” command line option which is required for data volume encryption. If HANA data volume encryption is not used the parameter must be removed.

The output below shows the log file of the call-hdbpersdiag-flexclone.sh script.

```
20251121085720##hana-7##call-hdbpersdiag-flexclone.sh: Executing
hdbpersdiag for source system SS2.
20251121085720##hana-7##call-hdbpersdiag-flexclone.sh: Clone mounted at
/hana/data/QS1/mnt00001.
20251121085720##hana-7##call-hdbpersdiag-flexclone.sh: Executing
hdbpersdiag in: /hana/data/QS1/mnt00001/hdb00001
20251121085723##hana-7##call-hdbpersdiag-flexclone.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/QS1/HDB11/hana-7/trace
Mounted DataVolume(s)
#0 /hana/data/QS1/mnt00001/hdb00001/ (3.1 GB, 3361128448 bytes)
Tips:
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '|' to redirect the output to a specific command.
INFO: KeyPage loaded and decrypted with success
        Default Anchor Page OK
        Restart Page OK
        Default Converter Pages OK
        RowStore Converter Pages OK
        Logical Pages (65415 pages) OK
        Logical Pages Linkage OK
Checking entries from restart page...
        ContainerDirectory OK
        ContainerNameDirectory OK
        FileIDMappingContainer OK
        UndoContainerDirectory OK
        LobDirectory OK
        MidSizeLobDirectory OK
        LobFileIDMap OK
20251121085723##hana-7##call-hdbpersdiag-flexclone.sh: Consistency check
operation successful for volume /hana/data/QS1/mnt00001/hdb00001.
20251121085723##hana-7##call-hdbpersdiag-flexclone.sh: Executing
hdbpersdiag in: /hana/data/QS1/mnt00001/hdb00002.00003
20251121085724##hana-7##call-hdbpersdiag-flexclone.sh: Loaded library
'libhdbunifiedtable'
Loaded library 'libhdblivecache'
Trace is written to: /usr/sap/QS1/HDB11/hana-7/trace
Mounted DataVolume(s)
#0 /hana/data/QS1/mnt00001/hdb00002.00003/ (288.0 MB, 301989888 bytes)
```

Tips:

```
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '!' to redirect the output to a specific command.
```

INFO: KeyPage loaded and decrypted with success

```
    Default Anchor Page OK
```

```
    Restart Page OK
```

```
    Default Converter Pages OK
```

```
    RowStore Converter Pages OK
```

```
    Logical Pages (4099 pages) OK
```

```
    Logical Pages Linkage OK
```

Checking entries from restart page...

```
    UndoContainerDirectory OK
```

```
    DRLoadedTable OK
```

```
20251121085724##hana-7##call-hdbpersdiag-flexclone.sh: Consistency check
operation successful for volume /hana/data/QS1/mnt00001/hdb00002.00003.
```

```
20251121085724##hana-7##call-hdbpersdiag-flexclone.sh: Executing
hdbpersdiag in: /hana/data/QS1/mnt00001/hdb00003.00003
```

```
20251121085729##hana-7##call-hdbpersdiag-flexclone.sh: Loaded library
'libhdbunifiedtable'
```

Loaded library 'libhdblivecache'

Trace is written to: /usr/sap/QS1/HDB11/hana-7/trace

Mounted DataVolume(s)

```
#0 /hana/data/QS1/mnt00001/hdb00003.00003/ (3.7 GB, 3942645760 bytes)
```

Tips:

```
Type 'help' for help on the available commands
Use 'TAB' for command auto-completion
Use '!' to redirect the output to a specific command.
```

INFO: KeyPage loaded and decrypted with success

```
    Default Anchor Page OK
```

```
    Restart Page OK
```

```
    Default Converter Pages OK
```

```
    Static Converter Pages OK
```

```
    RowStore Converter Pages OK
```

```
    Logical Pages (79243 pages) OK
```

```
    Logical Pages Linkage OK
```

Checking entries from restart page...

```
    ContainerDirectory OK
```

```
    ContainerNameDirectory OK
```

```
    FileIDMappingContainer OK
```

```
    UndoContainerDirectory OK
```

```
    LobDirectory OK
```

```
    DRLoadedTable OK
```

```
    MidSizeLobDirectory OK
```

```
    LobFileIDMap OK
```

```
20251121085729##hana-7##call-hdbpersdiag-flexclone.sh: Consistency check
```

```
operation successful for volume /hana/data/QS1/mnt00001/hdb00003.00003.
hana-7:/mnt/sapcc-share/hdbpersdiag #
```

File-based backup

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 all tenant databases.

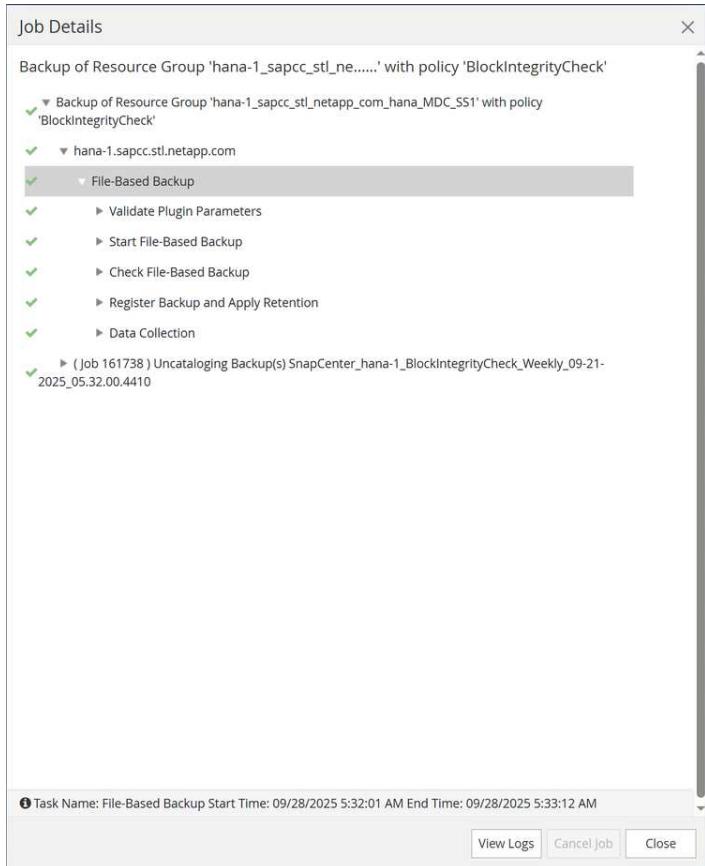
SnapCenter does not display the block integrity check in the same manner as Snapshot copy-based backups. Instead, the summary card shows the number of file-based backups and the status of the previous backup.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-1_Localsnap_Hourly_08-22-2025_08.00.00.3884	08/22/2025 8:01:02 AM	1	08/22/2025 8:01:02 AM
SnapCenter_hana-1_LocalsnapAndSnapshotVault_Daily_08-22-2025_06.00.00.4473	08/22/2025 6:01:02 AM	1	08/22/2025 6:01:02 AM
SnapCenter_hana-1_Localsnap_Hourly_08-22-2025_04.00.00.3795	08/22/2025 4:01:02 AM	1	08/22/2025 4:01:02 AM
SnapCenter_hana-1_Localsnap_Hourly_08-22-2025_00.00.00.3445	08/22/2025 12:01:02 AM	1	08/22/2025 12:01:02 AM
SnapCenter_hana-1_Localsnap_Hourly_08-21-2025_20.00.00.3933	08/21/2025 8:01:01 PM	1	08/21/2025 8:01:01 PM
SnapCenter_hana-1_Localsnap_Hourly_08-21-2025_16.00.00.4185	08/21/2025 4:01:01 PM	1	08/21/2025 4:01:01 PM
SnapCenter_hana-1_Localsnap_Hourly_08-21-2025_12.00.00.4902	08/21/2025 12:01:02 PM	1	08/21/2025 12:01:02 PM
SnapCenter_hana-1_Localsnap_Hourly_08-21-2025_08.00.00.3794	08/21/2025 8:01:02 AM	1	08/21/2025 8:01:02 AM
SnapCenter_hana-1_LocalsnapAndSnapshotVault_Daily_08-21-2025_06.00.00.3982	08/21/2025 6:01:04 AM	1	08/21/2025 6:01:04 AM
SnapCenter_hana-1_Localsnap_Hourly_08-21-2025_04.00.00.4068	08/21/2025 4:01:01 AM	1	08/21/2025 4:01:01 AM
SnapCenter_hana-1_Localsnap_Hourly_08-21-2025_00.00.00.3702	08/21/2025 12:01:02 AM	1	08/21/2025 12:01:02 AM
SnapCenter_hana-1_Localsnap_Hourly_08-20-2025_20.00.00.3830	08/20/2025 8:01:02 PM	1	08/20/2025 8:01:02 PM
SnapCenter_hana-1_Localsnap_Hourly_08-20-2025_16.00.00.3981	08/20/2025 4:01:01 PM	1	08/20/2025 4:01:01 PM
SnapCenter_hana-1_Localsnap_Hourly_08-20-2025_12.00.00.4105	08/20/2025 12:01:02 PM	1	08/20/2025 12:01:02 PM

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.

Status	Started	Duration	Size	Backup Type	Destination Ty...
Success	Aug 22, 2025, 8:00:26 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 22, 2025, 6:00:26 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 22, 2025, 4:00:26 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 22, 2025, 12:00:26 A...	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 8:00:25 PM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 4:00:25 PM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 12:00:25 PM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 8:00:26 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 6:00:26 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 4:00:25 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 12:00:26 A...	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 8:00:26 PM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 4:00:25 PM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 12:00:26 PM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 6:00:25 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 19, 2025, 6:00:26 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 18, 2025, 6:00:26 AM	00h 00m 18s	4.75 GB	Data Backup	Snapshot
Success	Aug 17, 2025, 6:00:24 AM	00h 00m 17s	4.75 GB	Data Backup	Snapshot
Success	Aug 17, 2025, 5:32:03 AM	00h 00m 13s	3.17 GB	Data Backup	File
Success	Aug 1, 2025, 6:00:27 AM	00h 00m 19s	4.75 GB	Data Backup	Snapshot

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:/hana/shared/SS1/HDB00/backup/data # ls -al *
DB_SS1:
total 3717564
drwxr-xr-- 2 ssladm sapsys 4096 Aug 22 11:03 .
drwxr-xr-- 4 ssladm sapsys 4096 Jul 27 2022 ..
-rw-r----- 1 ssladm sapsys 159744 Aug 17 05:32 SnapCenter_SnapCenter_hana-
1_BlockIntegrityCheck_Weekly_08-17-2025_05.32.00.4493_databackup_0_1
-rw-r----- 1 ssladm sapsys 83898368 Aug 17 05:32
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_08-17-
2025_05.32.00.4493_databackup_2_1
-rw-r----- 1 ssladm sapsys 3707777024 Aug 17 05:32
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_08-17-
2025_05.32.00.4493_databackup_3_1
SYSTEMDB:
total 3339236
drwxr-xr-- 2 ssladm sapsys 4096 Aug 22 11:03 .
drwxr-xr-- 4 ssladm sapsys 4096 Jul 27 2022 ..
-rw-r----- 1 ssladm sapsys 163840 Aug 17 05:32 SnapCenter_SnapCenter_hana-
1_BlockIntegrityCheck_Weekly_08-17-2025_05.32.00.4493_databackup_0_1

-rw-r----- 1 ssladm sapsys 3405787136 Aug 17 05:32
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_08-17-
2025_05.32.00.4493_databackup_1_1

```

Restore and recover SAP HANA databases with SnapCenter

Restore and recover SAP HANA systems using SnapCenter with automated or manual recovery options. This includes complete system restores, single tenant restores for HANA databases on ONTAP, Azure NetApp Files, and FSx for ONTAP.

SnapCenter supports the following restore and recovery operations.

- SAP HANA MDC systems with a single tenant
 - End-to-end automated restore and recovery
 - End-to-end automated restore and manual recovery (selectable)
- SAP HANA MDC systems with multiple tenants
 - End-to-end automated restore, recovery needs to be done manually
- Restore of a single tenant
 - End-to-end automated restore, recovery needs to be done manually

 Automated recovery is only supported when the HANA plug-in is deployed on the HANA database host and the HANA system got auto discovered by SnapCenter. With a central plug-in host configuration, recovery needs to be done manually after the restore operation with SnapCenter.



Restore from primary ANF volume is supported. A restore from ANF backup is not yet supported. An in-place restore or a restore to a new volume from an ANF backup must be done manually using the Azure portal or CLI.

Automated restore and recovery for SAP HANA MDC systems with a single tenant

A restore operation is initiated by selecting a Snapshot backup in the resource topology view and by clicking on Restore.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-1_LocalSnap_Hourly_08-25-2025_08.00.00.5626	08/25/2025 8:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-25-2025_06.00.00.4688	08/25/2025 6:01:03 AM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-25-2025_04.00.00.4833	08/25/2025 4:01:02 AM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-25-2025_00.00.3988	08/25/2025 12:01:02 AM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-24-2025_20.00.00.3489	08/24/2025 8:01:02 PM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-24-2025_16.00.00.4463	08/24/2025 4:01:02 PM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-24-2025_12.00.00.4369	08/24/2025 12:01:02 PM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-24-2025_08.00.00.3905	08/24/2025 8:01:02 AM	1	
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_08-24-2025_06.00.00.4125	08/24/2025 6:01:02 AM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-24-2025_04.00.00.4423	08/24/2025 4:01:03 AM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-24-2025_00.00.4290	08/24/2025 12:01:03 AM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-23-2025_20.00.00.2799	08/23/2025 8:01:02 PM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-23-2025_16.00.00.3863	08/23/2025 4:01:02 PM	1	
SnapCenter_hana-1_LocalSnap_Hourly_08-23-2025_12.00.00.5963	08/23/2025 12:01:02 PM	1	

For HANA systems using NFS on ANF, FSx for ONTAP or ONTAP storage systems you can select complete restore with or without a volume revert operation for primary volume Snapshots.

- Complete resource without volume revert uses Single File SnapRestore (SFSR) to restore all files of the database.
- Complete resource with volume revert uses a volume based restore operation (VBSR) to revert the complete volume back to the state of the selected Snapshot.



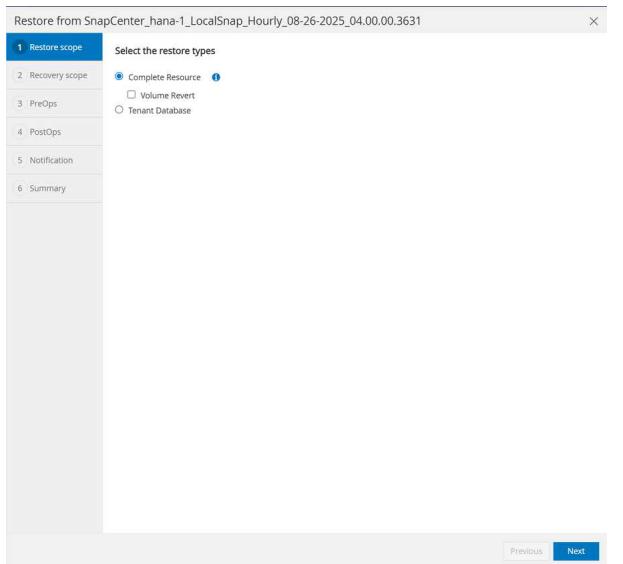
Volume revert can't be used if you need to restore to a Snapshot which is older than the active SnapVault or SnapMirror replication Snapshot.



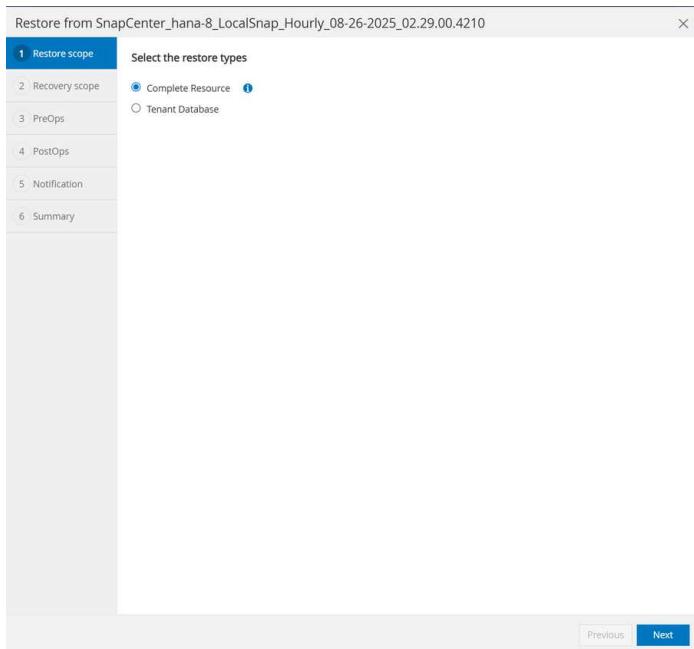
A volume revert operation will delete all Snapshot backups which are newer than the selected Snapshot for the revert operation.



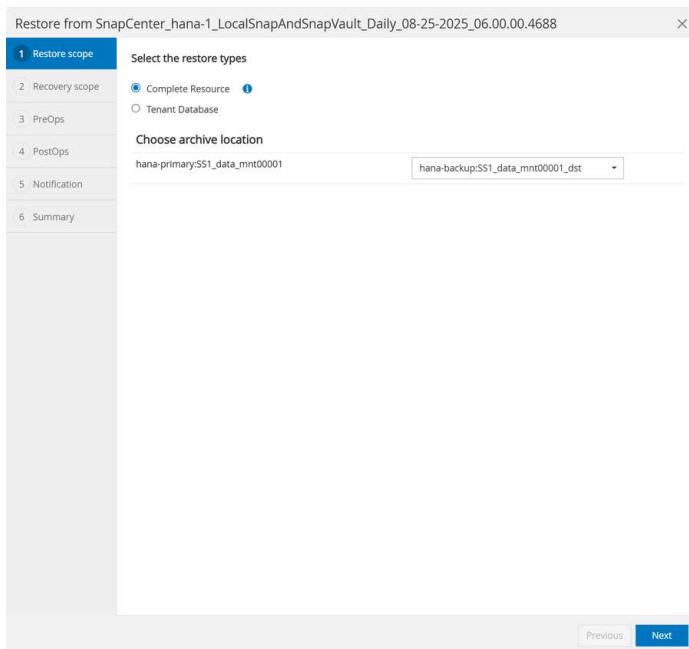
A restore with SFSR is nearly as fast as a volume revert operation but blocks any Snapshot operation until the background process has finished the meta data operations.



For HANA systems on bare metal hosts using FC SAN, a volume revert (VBSR) is not supported, instead SFSR is always used for the restore operation. For HANA systems running on VMware with VMFS a clone, mount, copy operation will be used.



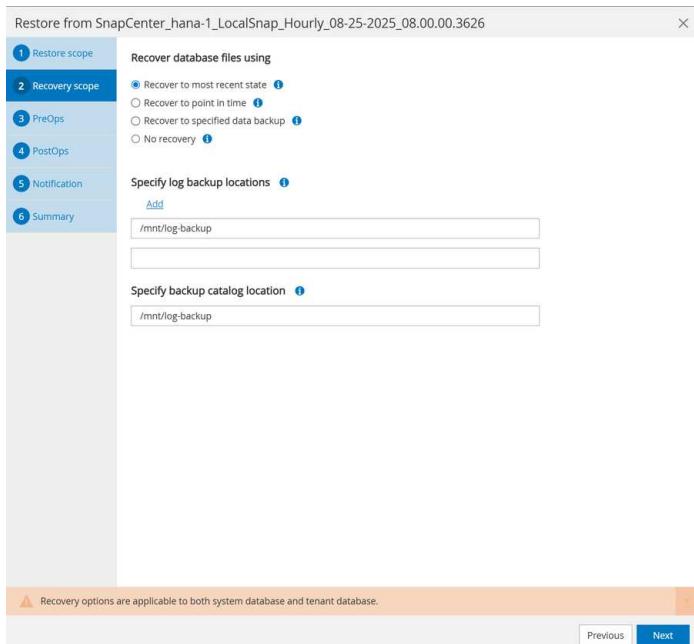
For a restore from a secondary backup you need to select the archive location.



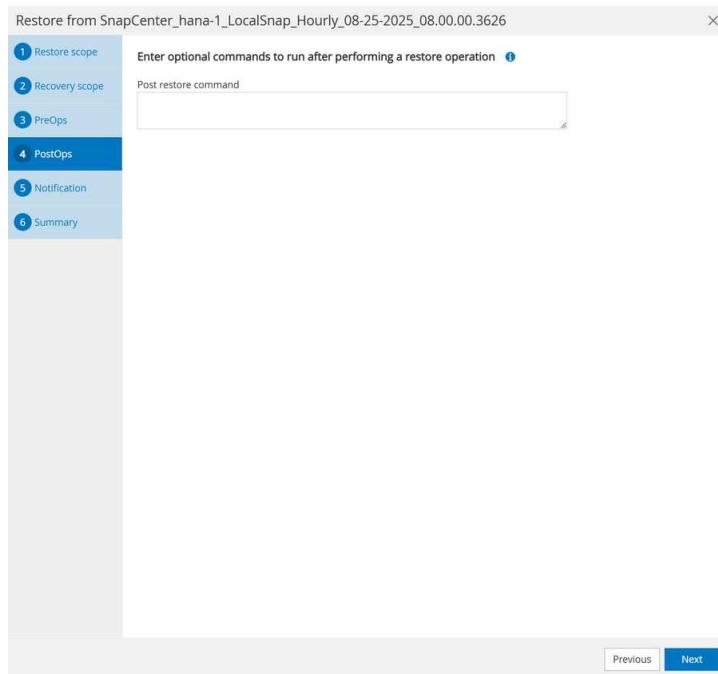
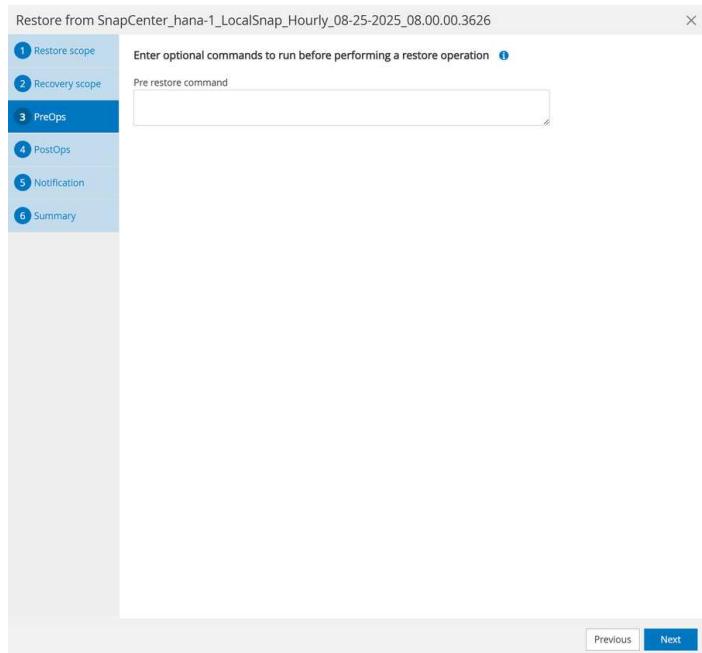
With the recovery scope you can select a 'to most recent state', 'point in time' or a save point recovery without using log backups. If you select no recovery, SnapCenter only executes the restore operation and the recovery needs to be done manually as described ["Manual recovery with HANA Studio"](#).



SnapCenter uses the paths configured in SAP HANA for log backup and catalog backup locations. If you have tiered backups to an additional location, you can add these additional paths.



Optionally you can add pre and post restore scripts.



When clicking on Finish in the summary screen, the restore and recovery operation is started.

Restore from SnapCenter_hana-1_LocalSnap_Hourly_08-25-2025_08.00.00.3626 X

1 Restore scope Summary

2 Recovery scope Backup Name: SnapCenter_hana-1_LocalSnap_Hourly_08-25-2025_08.00.00.3626

3 PreOps Backup date: 08/25/2025 8:01:02 AM

4 PostOps Restore scope: Complete Resource with Volume Revert

5 Notification Recovery scope: Recover to most recent state

6 Summary Log backup locations: /mnt/log-backup

Backup catalog location: /mnt/log-backup

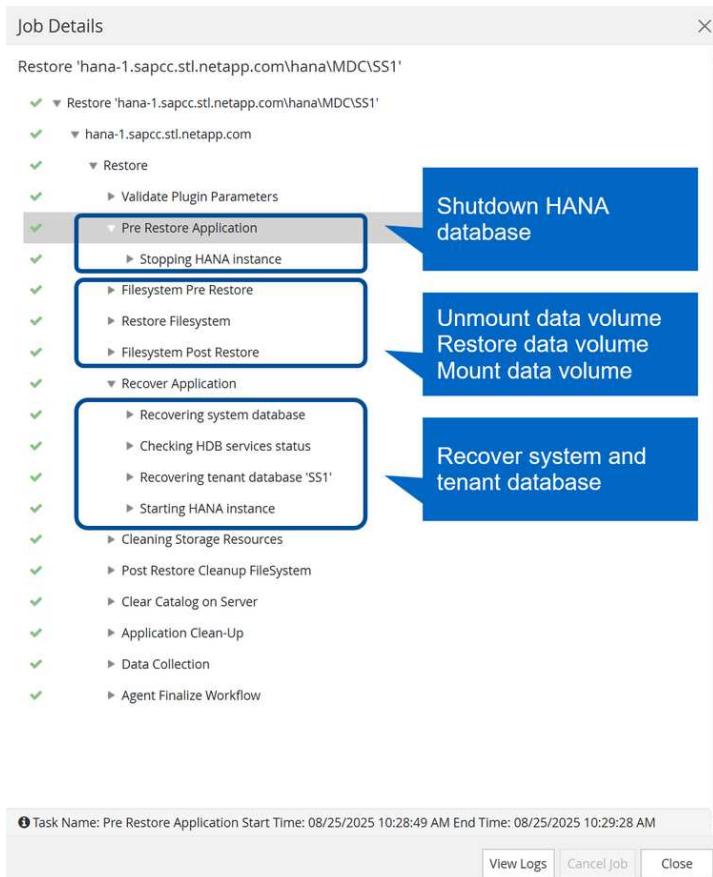
Pre restore command: Send email: No

Post restore command: Send email: No

Previous Finish

The restore and recovery workflow can be divided in three main sections.

- Shutdown of the HANA system
- Restore operation
 - Filesystem specific preparations, e.g. unmount operation
 - Snapshot restore operation
 - Filesystem specific post operations, e.g. mount operation
- HANA recovery
 - System database recovery
 - Tenant database recovery



Manual recovery with HANA Studio

To restore and recover an SAP HANA MDC system with a single or with multiple tenants 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 provide the backup catalog 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 secondary 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.
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.

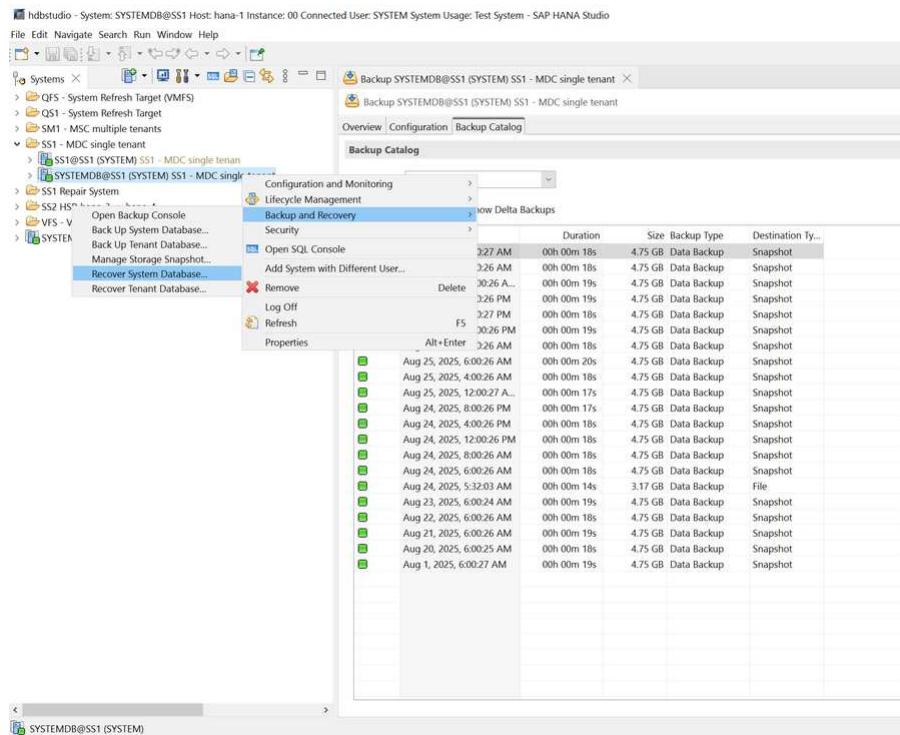
- c. A list of data backups displays. Because the data volume has already been restored, the tenant backup is indicated as available (in green).
 - d. Select this backup and start the recovery process. After the recovery process is finished, the tenant database is started automatically.
5. For a HANA system with multiple tenants repeat step 4 for each tenant.



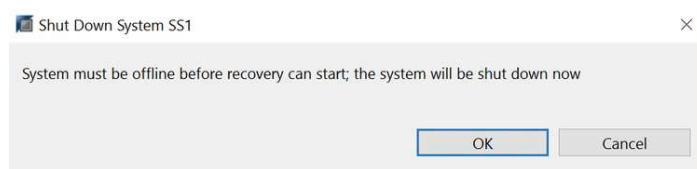
A manual recovery with SAP HANA Cockpit is done with the same steps.

The following section describes the steps of the restore and recovery operations of an SAP HANA MDC system with a single tenant.

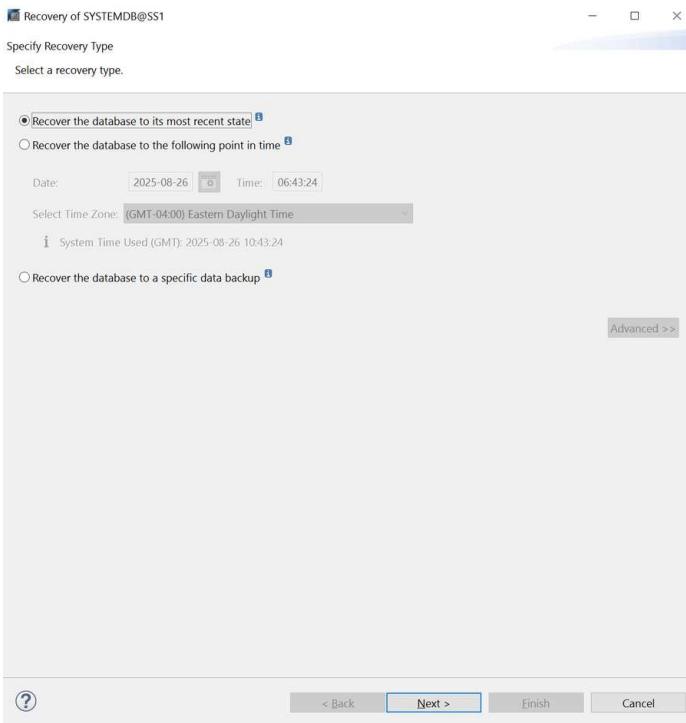
In HANA Studio select Backup and Recovery and Recover System Database.



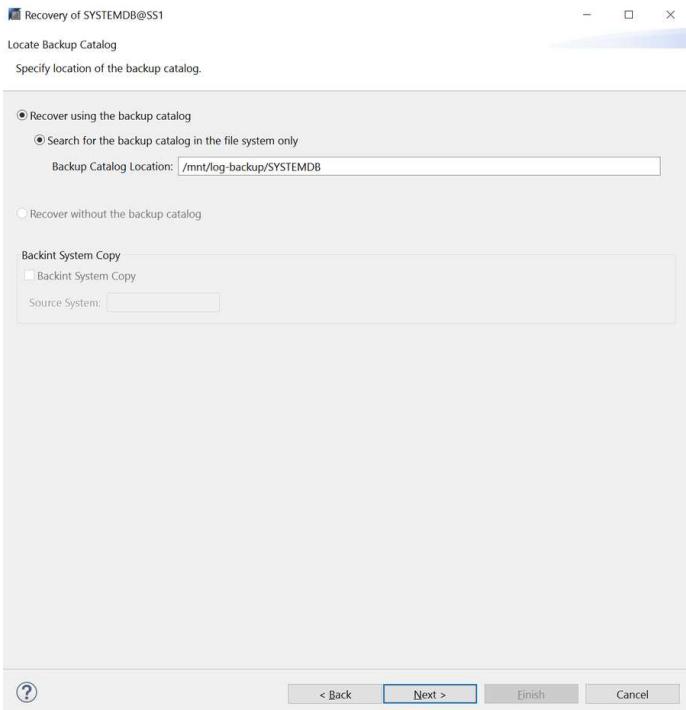
Confirm shutdown operation; only required if the HANA system is still running.



Select recovery operation. In this example we want to recover to the most recent state.

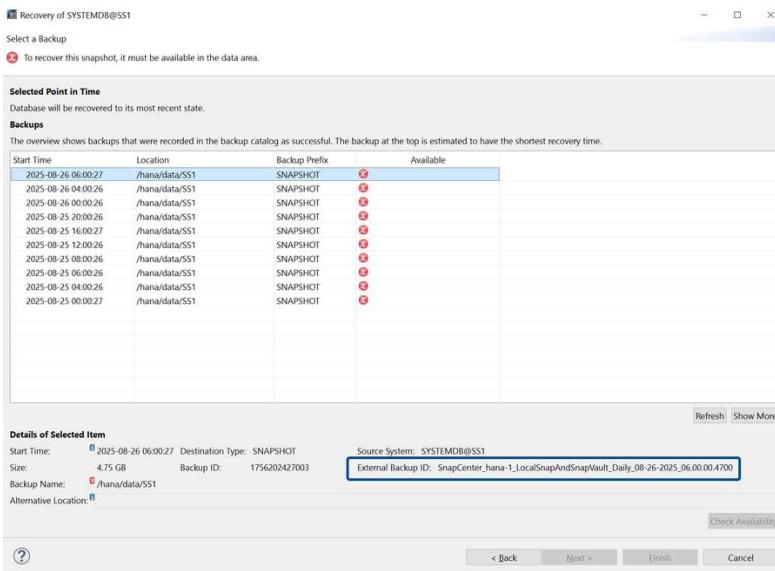


Provide backup catalog location.



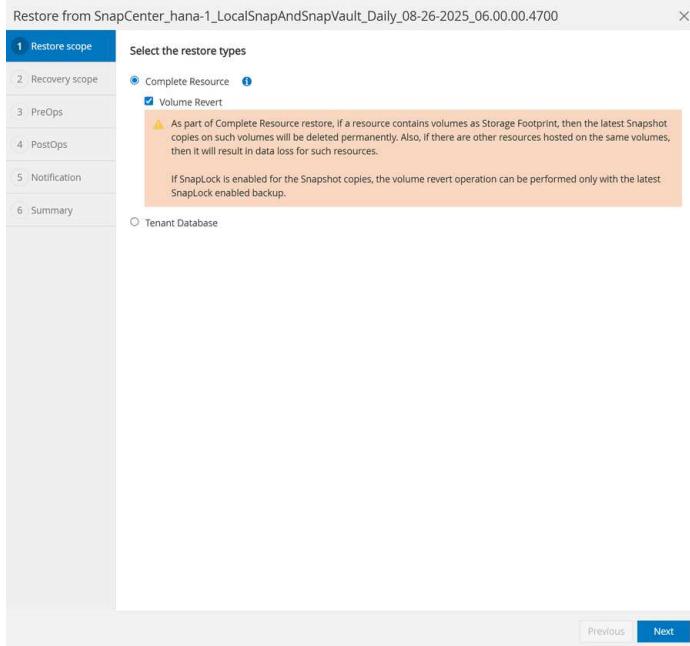
HANA Studio lists the most recent backups stored in the HANA backup catalog.

A list of available backups is shown based on the content of the backup catalog. Choose the required backup and note the external backup ID: in this example, the most recent backup.

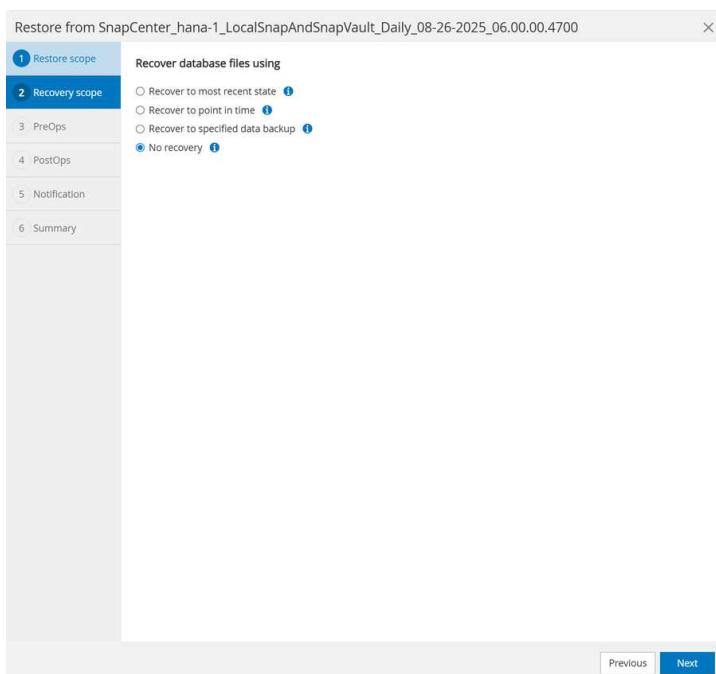


From the SnapCenter GUI, select the resource topology view and select the backup that should be restored, in this example, the most recent primary backup. Click the Restore icon to start the restore.

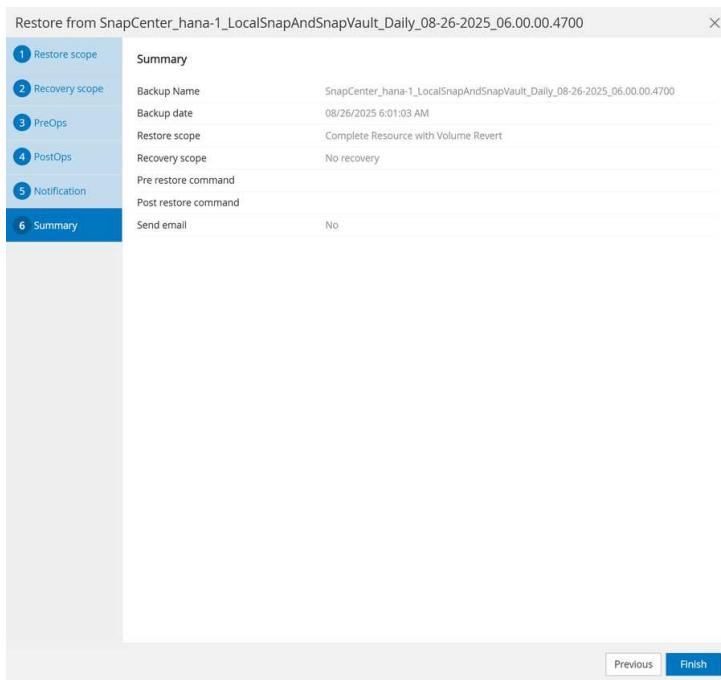
The SnapCenter restore wizard starts. Select the restore type Complete Resource and Volume revert to use a volume-based restore.



Select 'No recovery' to exclude the recovery operations from the SnapCenter workflow.

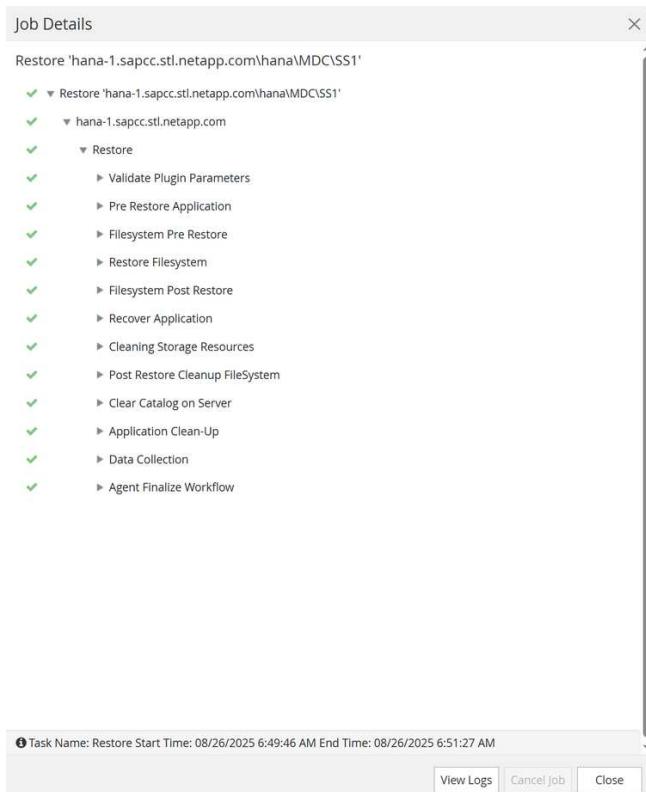


Click on Finish to start the restore operation.



SnapCenter is now executing the restore operation.

- Filesystem specific preparations, e.g. unmount operation
- Snapshot restore operation
- Filesystem specific post operations, e.g. mount operation

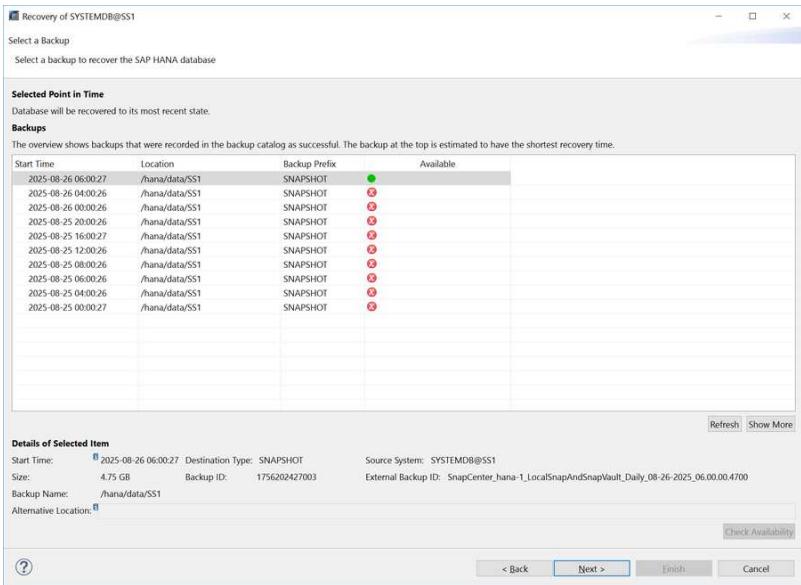


When the Snapshot got restored by SnapCenter a `snapshot_databasebackup_0_1` file is available in the system and tenant database subdirectory of the HANA data volume. This file got created by the HANA database during the HANA database Snapshot creation. HANA deletes the file when the backup operation is finished, so

that the files are only visible within the Snapshot backup. These files are required for any recovery operation. After the recovery the files get deleted by the HANA database.

```
hana-1:~ # cd /hana/data/SS1/mnt00001/
hana-1:/hana/data/SS1/mnt00001 # ls -al *
-rw-r--r-- 1 ssladm sapsys 16 Aug 26 06:00 nameserver.lck
hdb00001:
total 4992236
drwxr-x--- 2 ssladm sapsys 4096 Aug 26 06:00 .
drwxr-x--- 5 ssladm sapsys 4096 Aug 26 06:00 ..
-rw-r----- 1 ssladm sapsys 0 Nov 3 2020
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r----- 1 ssladm sapsys 5100273664 Aug 26 06:00 datavolume_0000.dat
-rw-r----- 1 ssladm sapsys 36 Aug 25 10:30 landscape.id
-rw-r----- 1 ssladm sapsys 163840 Aug 26 06:00 snapshot_databackup_0_1
hdb00002.00003:
total 201420
drwxr-xr-- 2 ssladm sapsys 4096 Nov 3 2020 .
drwxr-x--- 5 ssladm sapsys 4096 Aug 26 06:00 ..
-rw-r--r-- 1 ssladm sapsys 0 Nov 3 2020
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r--r-- 1 ssladm sapsys 335544320 Aug 26 06:00 datavolume_0000.dat
hdb00003.00003:
total 4803140
drwxr-xr-- 2 ssladm sapsys 4096 Aug 26 06:00 .
drwxr-x--- 5 ssladm sapsys 4096 Aug 26 06:00 ..
-rw-r--r-- 1 ssladm sapsys 0 Nov 3 2020
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r--r-- 1 ssladm sapsys 4898947072 Aug 26 06:00 datavolume_0000.dat
-rw-r----- 1 ssladm sapsys 159744 Aug 26 06:00 snapshot_databackup_0_1
hana-1:/hana/data/SS1/mnt00001 #
```

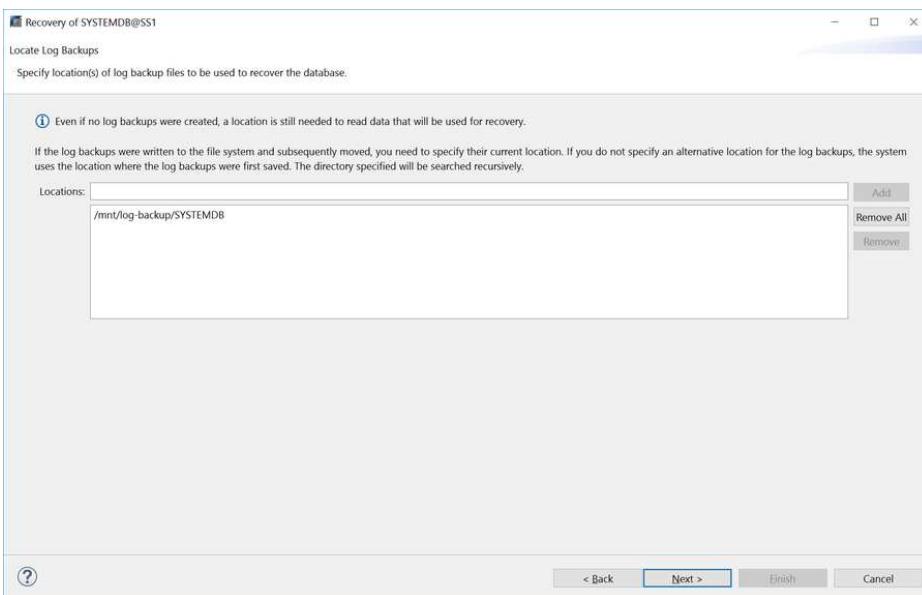
Go to SAP HANA Studio and click Refresh to update the list of available backups. The backup that was restored with SnapCenter is now shown with a green icon in the list of backups. Select the backup and click Next.



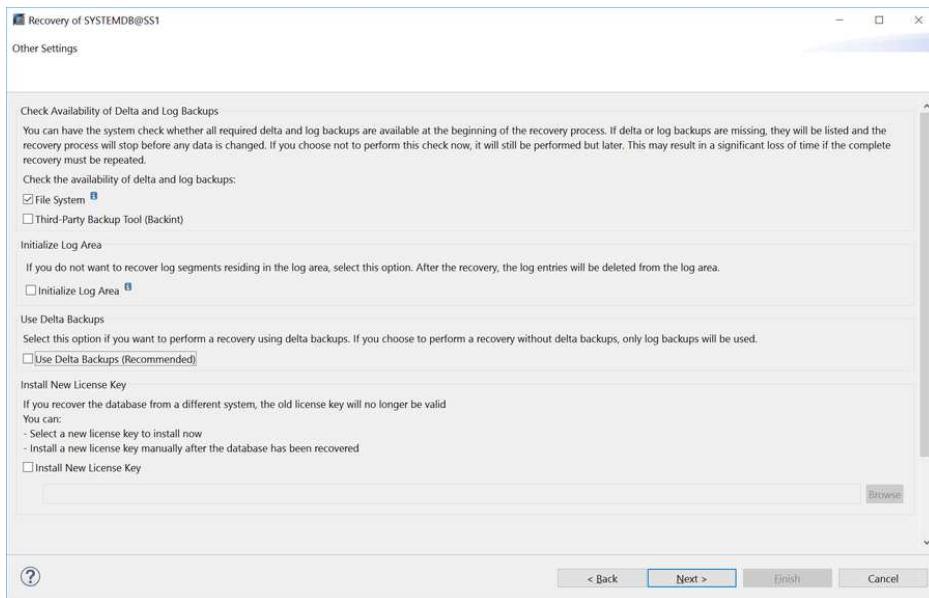
Provide the location of the log backups. Click Next.



SAP HANA Studio uses the paths configured in SAP HANA for log backup and catalog backup locations. If you have tiered backups to an additional location, you can add these additional paths.

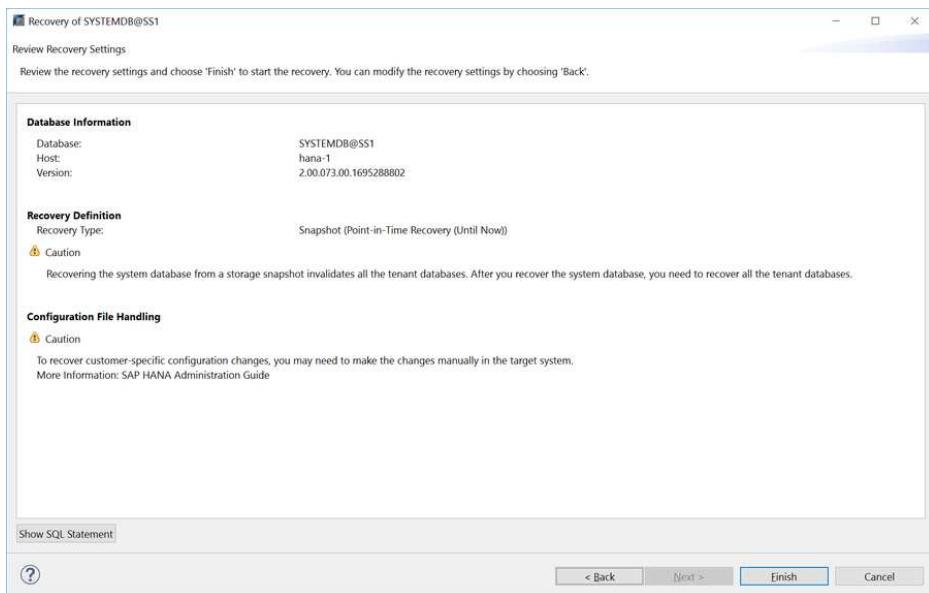


Select other settings as required. Make sure Use Delta Backups is not selected. Click Next.

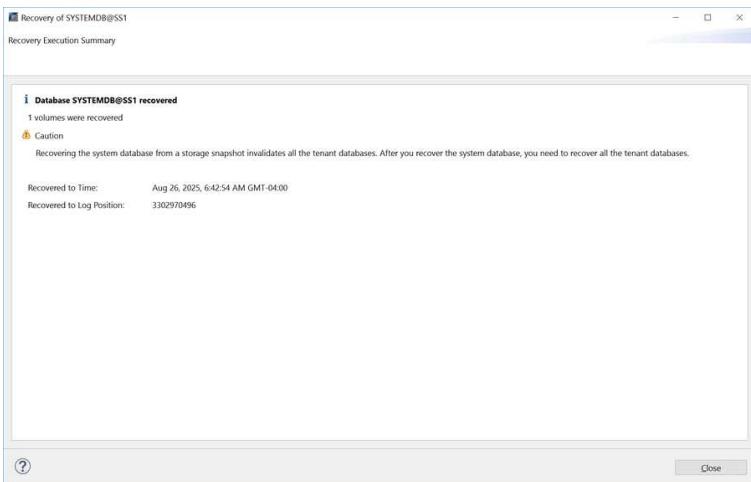


Review the recovery settings and click Finish.

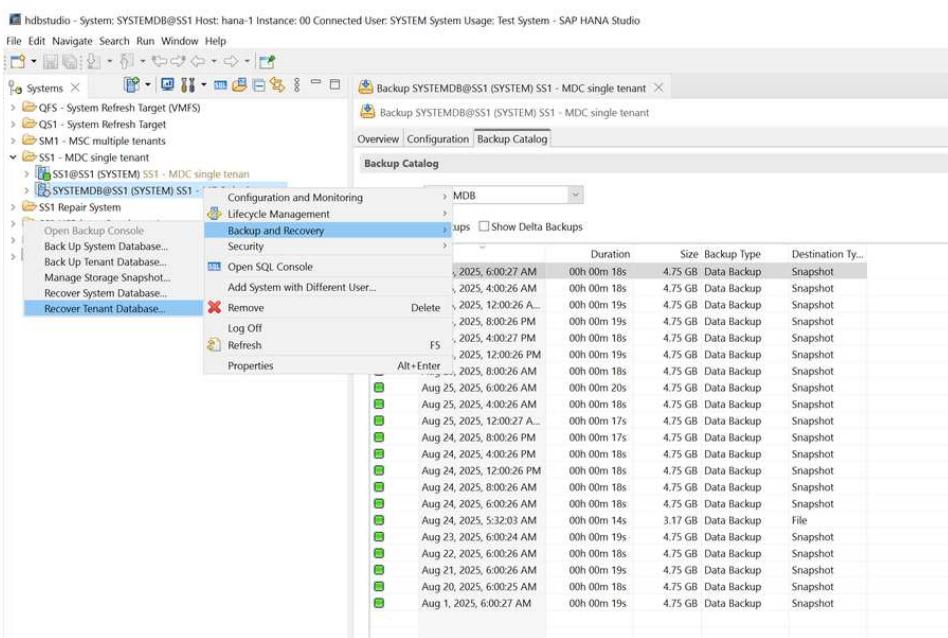
By clicking on show SQL statement, HANA Studio shows the SQL command which is executed for the recovery operation.



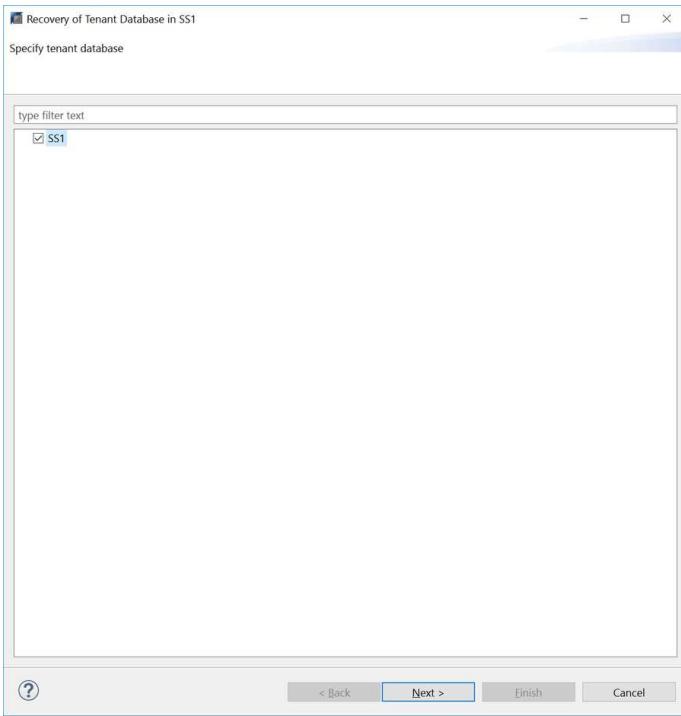
The recovery process starts. Wait until the recovery of the system database is completed.



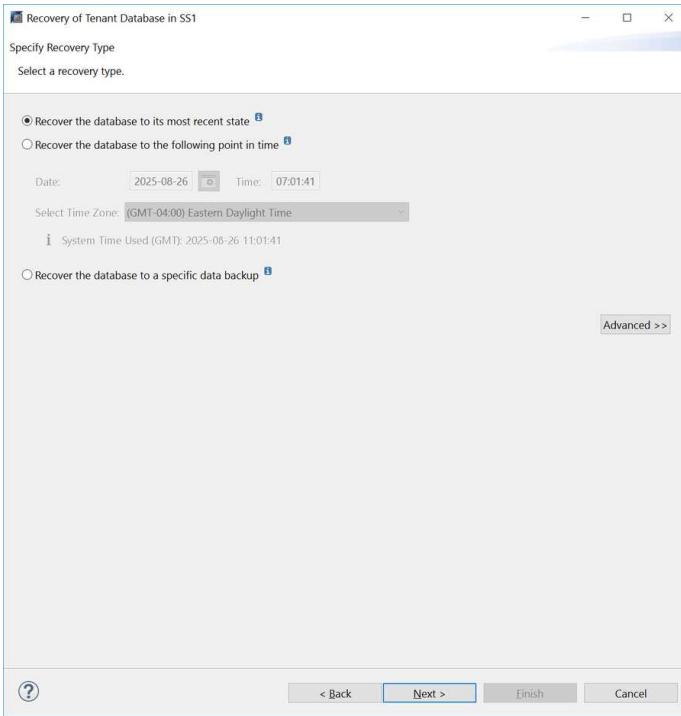
In SAP HANA Studio, select the entry for the system database and start Backup Recovery - Recover Tenant Database.



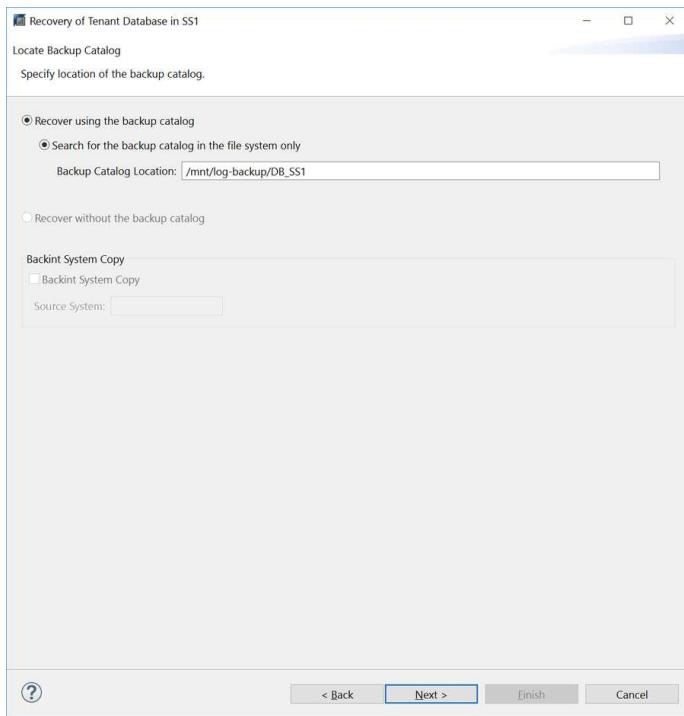
Select the tenant to recover and click Next.



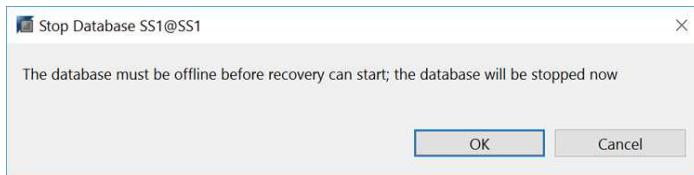
Specify the recovery type and click Next.



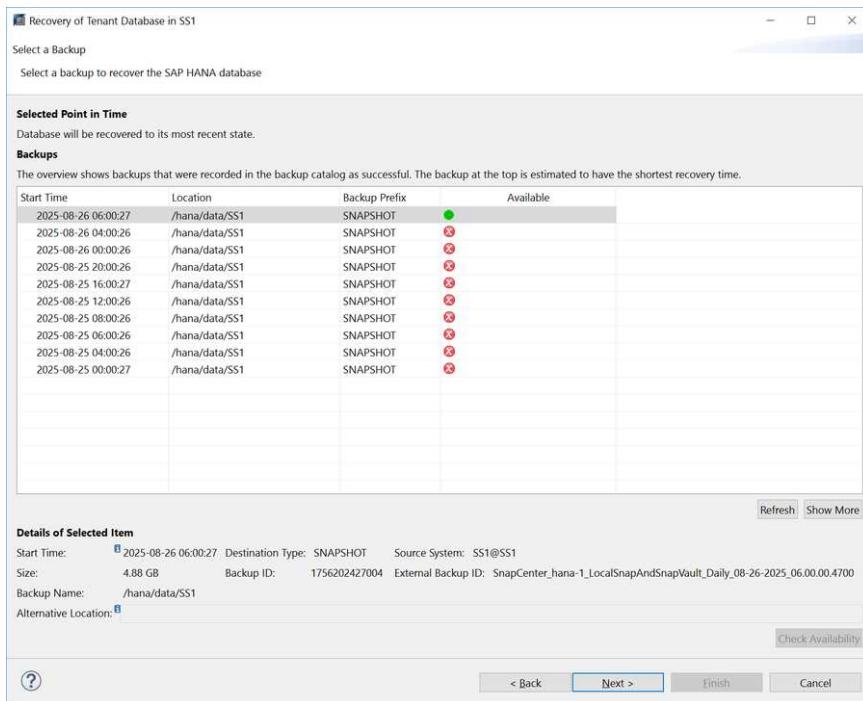
Confirm the backup catalog location and click Next.



Confirm that the shutdown of the tenant database.



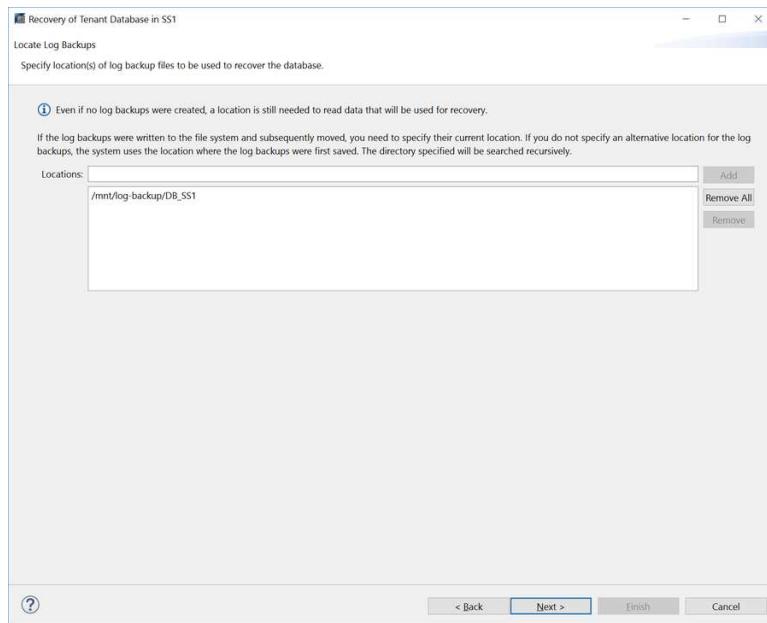
Because the restore of the data volume has been done before the recovery of the system database, the tenant backup is immediately available. Select the backup highlighted in green and click Next.



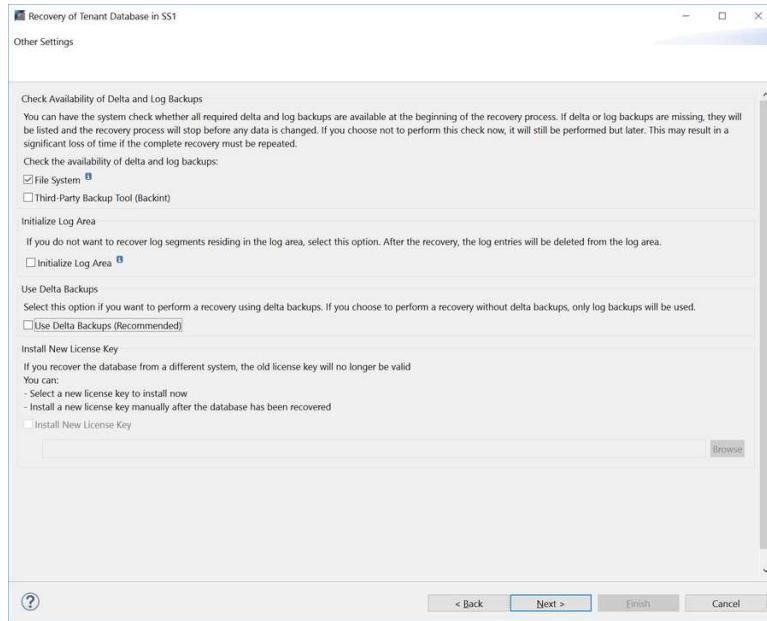
Provide the location of the log backups. Click Next.



SAP HANA Studio uses the paths configured in SAP HANA for log backup and catalog backup locations. If you have tiered backups to an additional location, you can add these additional paths.

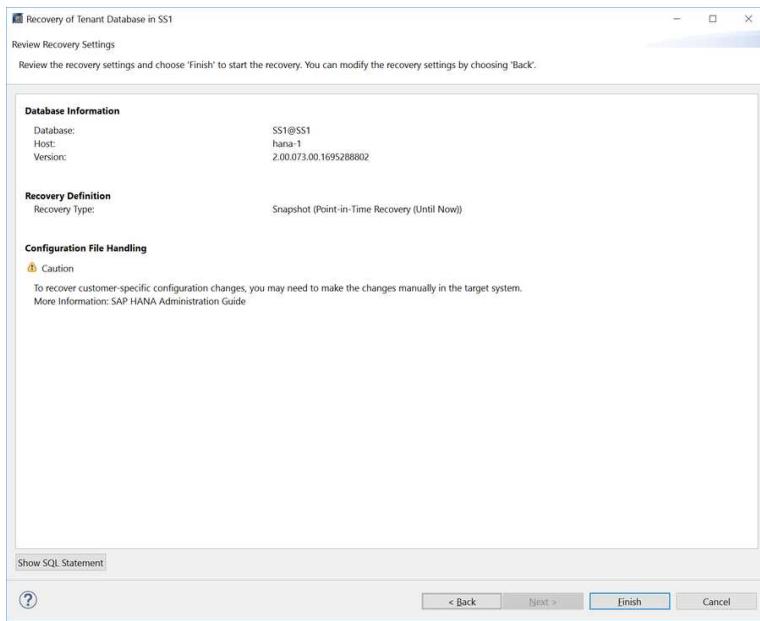


Select other settings as required. Make sure Use Delta Backups is not selected. Click Next.

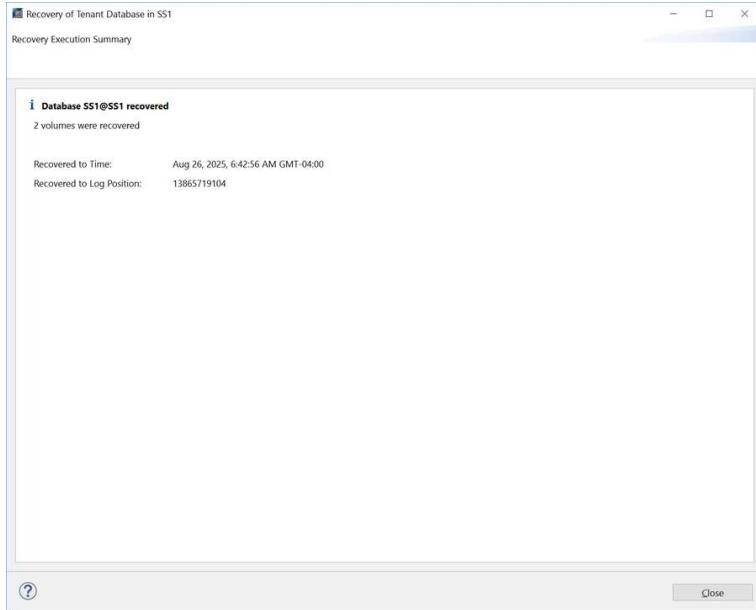


Review the recovery settings and click Finish.

By clicking on show SQL statement, HANA Studio shows the SQL command which is executed for the recovery operation.



Wait until the recovery has finished and the tenant database is started.



When the tenant recovery is finished the SAP HANA system is up and running.



For an SAP HANA MDC system with multiple tenants, you must repeat the tenant recovery for each tenant.

Manual recovery with SQL commands

You can also use SQL statements for the recovery of the HANA system.

First you need to recover the system database.

```
HDBSettings.sh recoverSys.py --command="RECOVER DATABASE UNTIL TIMESTAMP '2026-08-26 10:55:49' USING CATALOG PATH ('mnt/log-backup/SYSTEMDB') USING LOG PATH ('mnt/log-backup/SYSTEMDB') USING SNAPSHOT"
```

As a second step you need to connect to the system database and start the recovery of the tenant database(s). In this example the tenant database is SS1.

```
hdblsql SYSTEMDB=> RECOVER DATABASE FOR SS1 UNTIL TIMESTAMP '2026-08-26 10:55:49' USING CATALOG PATH ('mnt/log-backup/DB_SS1') USING LOG PATH ('mnt/log-backup/DB_SS1') USING SNAPSHOT
```

Single tenant restore and recovery

A single tenant restore and recovery operation with SnapCenter is very similar to the workflow described in the previous topic ["Manual recovery with HANA Studio"](#).

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 Tenant Database and confirm shutdown of the tenant database.
 - b. Select the recovery type and provide the backup catalog 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 secondary 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 of the tenant.
3. Run the recovery process for the tenant 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 tenant database is started.

Restore of non-data volumes

A restore operation for a non-data volume is started by selecting a Snapshot backup in the topology view of the non-data volume resource and by clicking on **Restore**.

For non-data volumes with NFS a complete resource (VBSR) or a file level (SFSR) restore operation can be selected. For the file level restore either all or individual files can be defined for the restore operation.

Configure advanced SnapCenter options for SAP HANA

Configure advanced SnapCenter settings for SAP HANA environments, including suppressing VMware warning messages for in-guest NFS mounts, disabling automated log backup housekeeping, and enabling SSL encryption for HANA database connections.

Warning message with virtualized environments and in-guest mounts

When using for example VMware with NFS in-guest mounts, SnapCenter will issue a warning message, that the SnapCenter VMware plug-in should be used. Since the VMWare plug-in is not required for in-guest mounts the warning message can be ignored and switched off. 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.

Deactivate automated log backup housekeeping

Log backup housekeeping is enabled by default and can be disabled on the HANA plug-in host level. Use the PowerShell command:

The command `Set-SnapCenterConfig -Plugin -HostName <pluginhostname> -PluginCode hana -configSettings @{"LOG_CLEANUP_DISABLE" = "Y"}` disables the log backup housekeeping for this SAP HANA host.

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.

There are various options to configure the SSL communication. By default, SnapCenter uses the `-e ssltrustcert` `hdbsql` command-line option. With this option SSL communication without server certificate validation is done and this option also works for HANA systems where SSL is not enabled.

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.

This can be achieved by using a wrapper script which calls `hdbsql` with the required options. Instead of configuring the `hdbsql` executable in the `hana.properties` files, the wrapper script is added.

```
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 <command line options> $*
```

Disable auto discovery on the HANA plug-in host

To disable auto discovery 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>
```

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 #
```

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.

Backup Catalog				
Stat...	Started	Duration	Size	Backup Ty...
■	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.

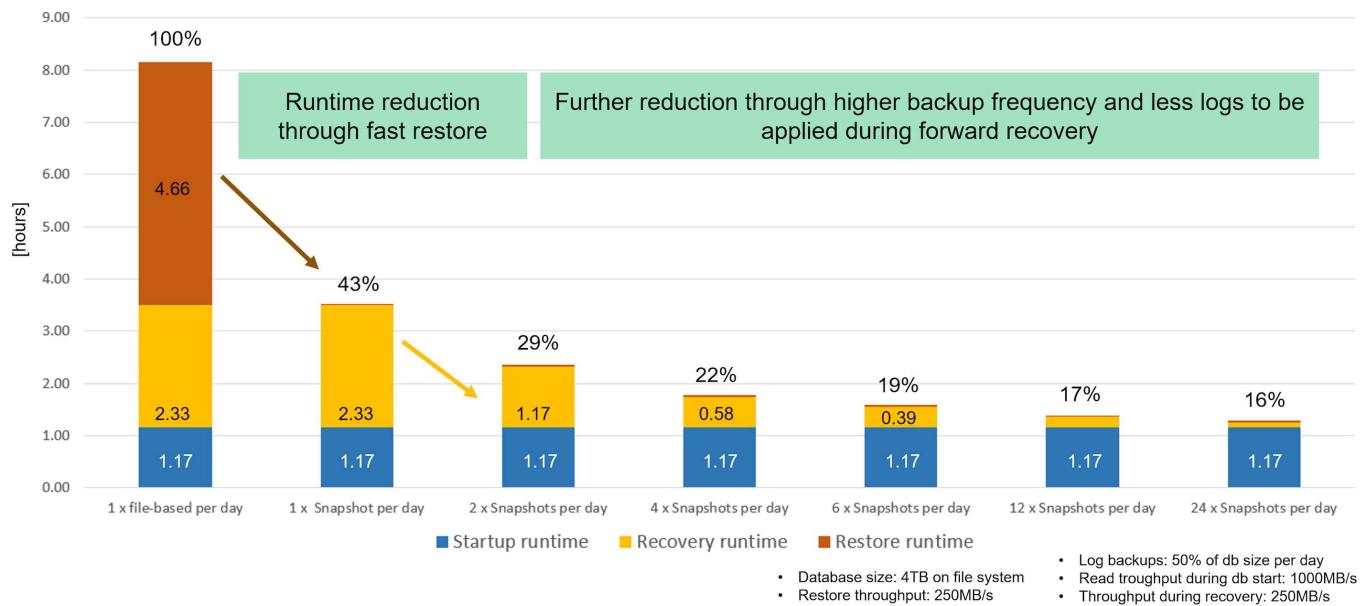
Snapshot backups are typically scheduled with a higher frequency because they do not influence the performance of the SAP HANA database. For example, if Snapshot backups are scheduled every six hours, the recovery time would be, in the worst case, one-fourth of the recovery time for a file-based backup (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

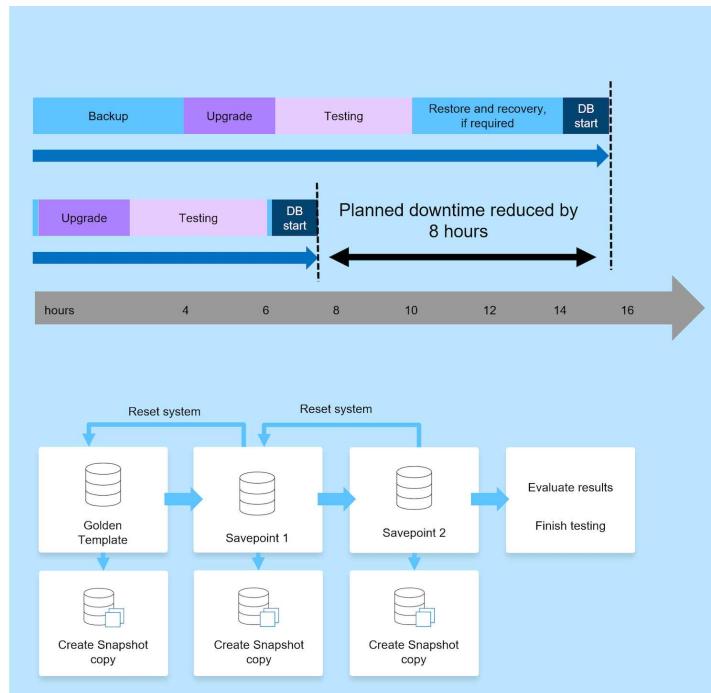
Executing backups is a critical part of any data protection strategy. Backups are scheduled on a regular basis to ensure that you can recover from system failures. This is the most obvious use case, but there are also other SAP lifecycle management tasks, where accelerating backup and recovery operations 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.

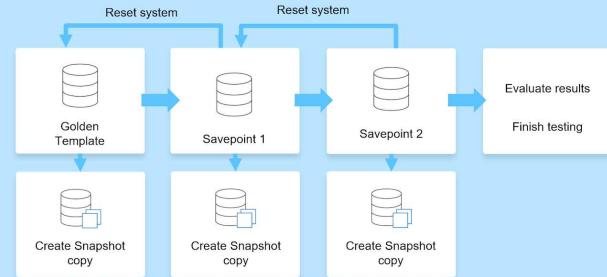
Another use case example would be a typical test cycle, where testing must be done over multiple iterations with different data sets or parameters. When leveraging the fast backup and restore operations, you can easily create save points within your test cycle and reset the system to any of these previous 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.

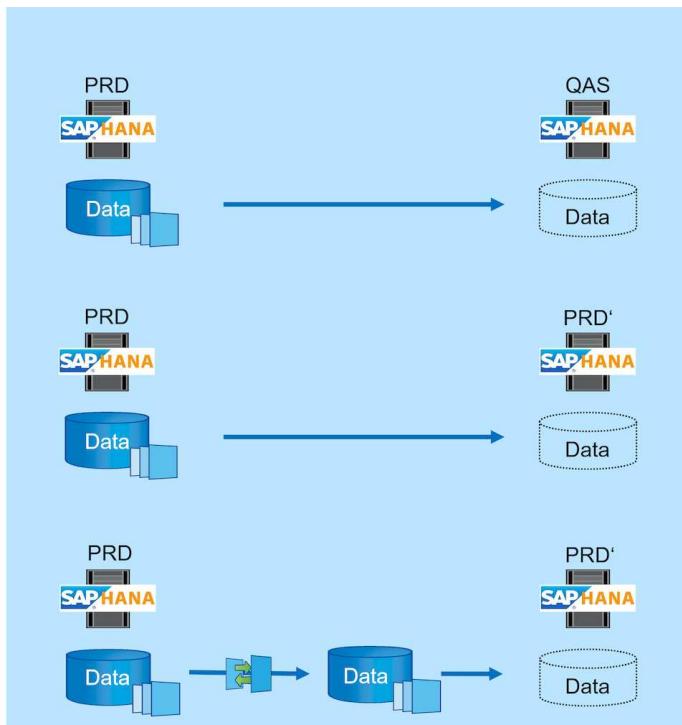
The most popular use case is the SAP System Refresh, where data from the production system needs to be copied to the test or QA system. By leveraging the FSx for ONTAP cloning feature, you can provision the volume for the test system from any Snapshot copy of the production system in a matter of seconds. 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



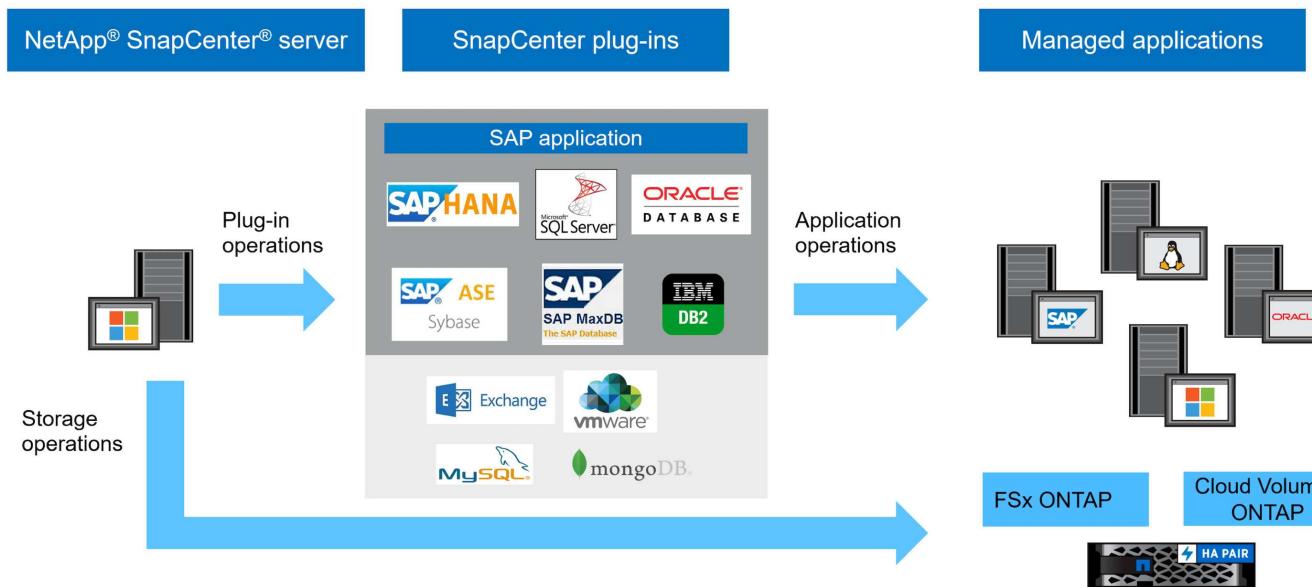
SnapCenter architecture

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.

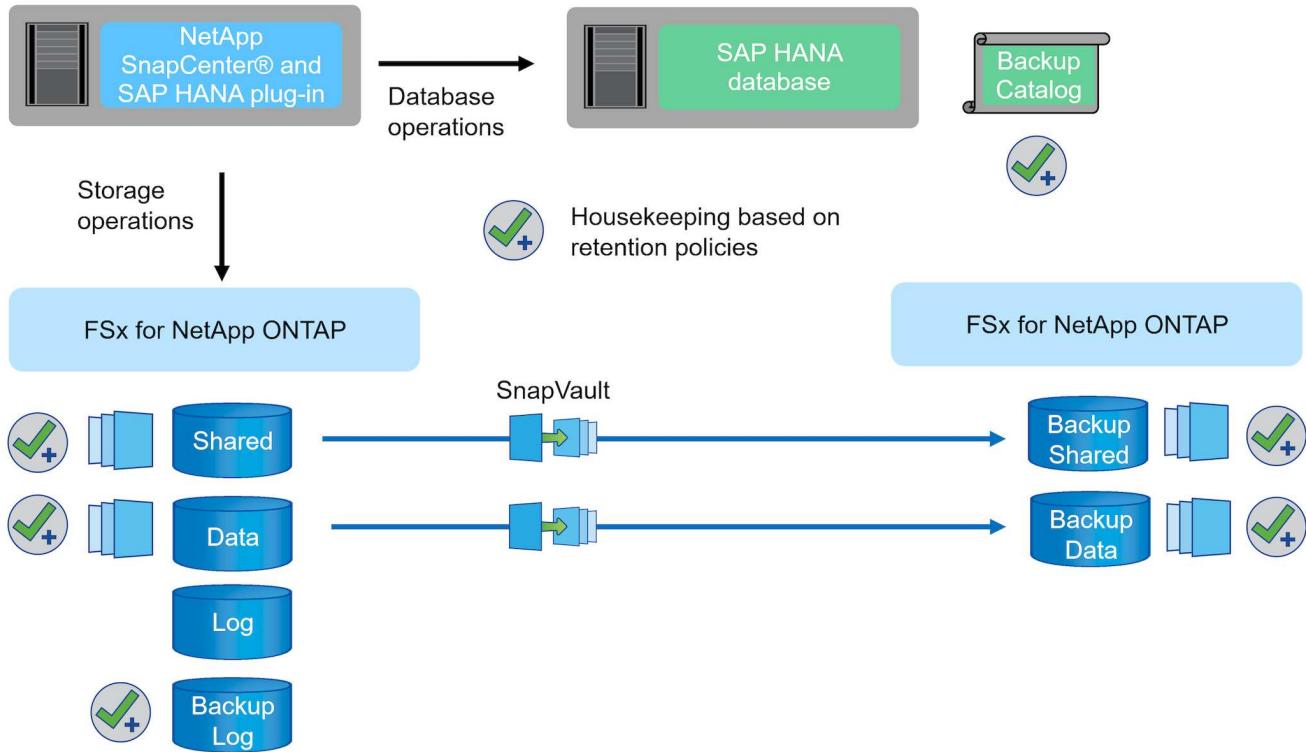
To allow a complete backup of all SAP HANA-related resources, SnapCenter also enables you to back up all non-data volumes by using the SAP HANA plug-in with storage-based Snapshot copies. You can schedule non-data volumes independently from the database data backup to enable individual retention and 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 backup 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

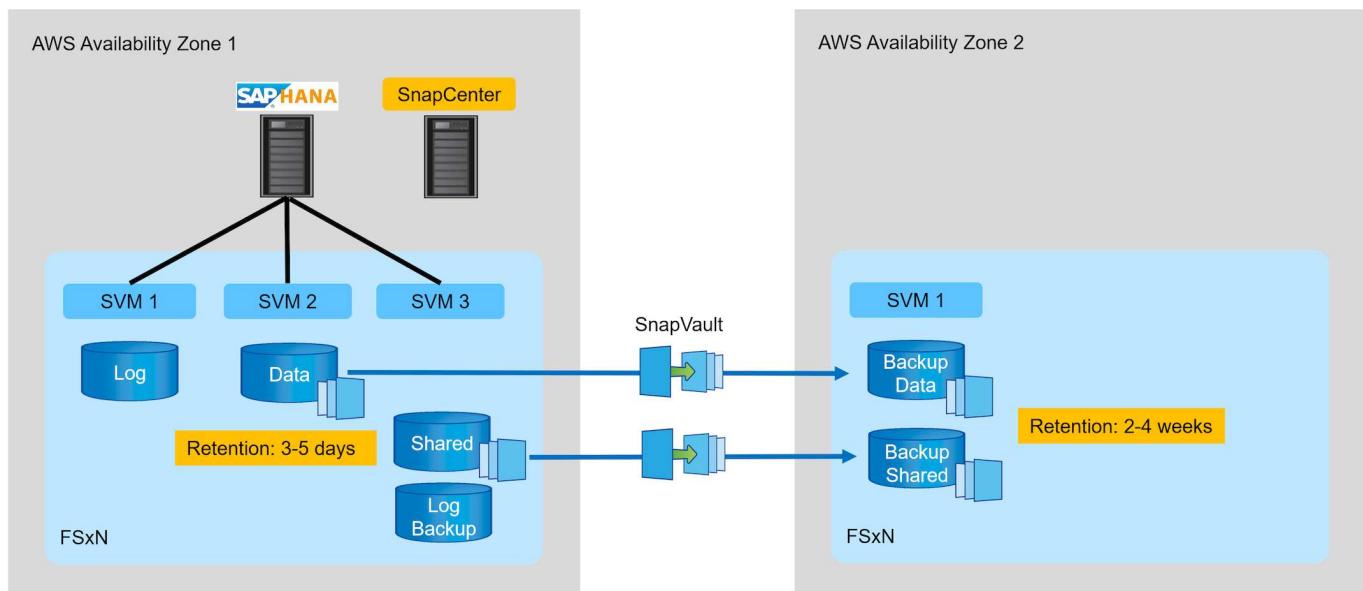
This document describes the most common SnapCenter configuration option for an SAP HANA MDC single host system with a single tenant on FSx for ONTAP. Other configuration options are possible and, in some

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:

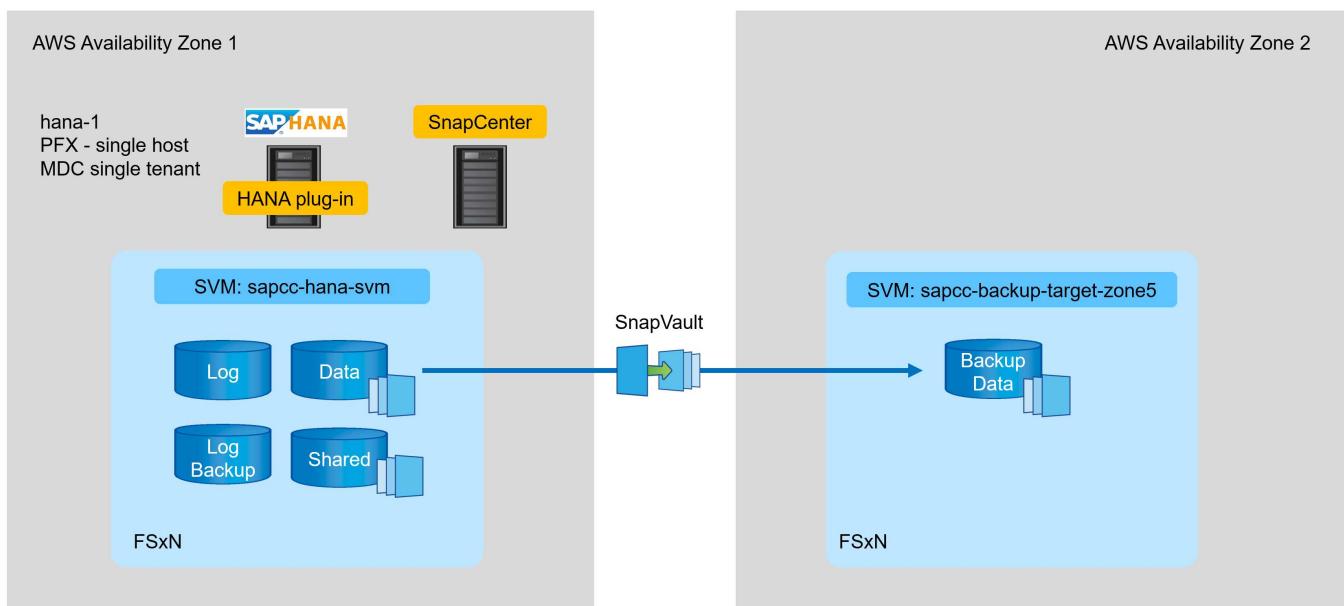
- Single host MDC system with a single tenant
- HANA 2.0 SPS 6 revision 60
- SLES for SAP 15SP3

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.

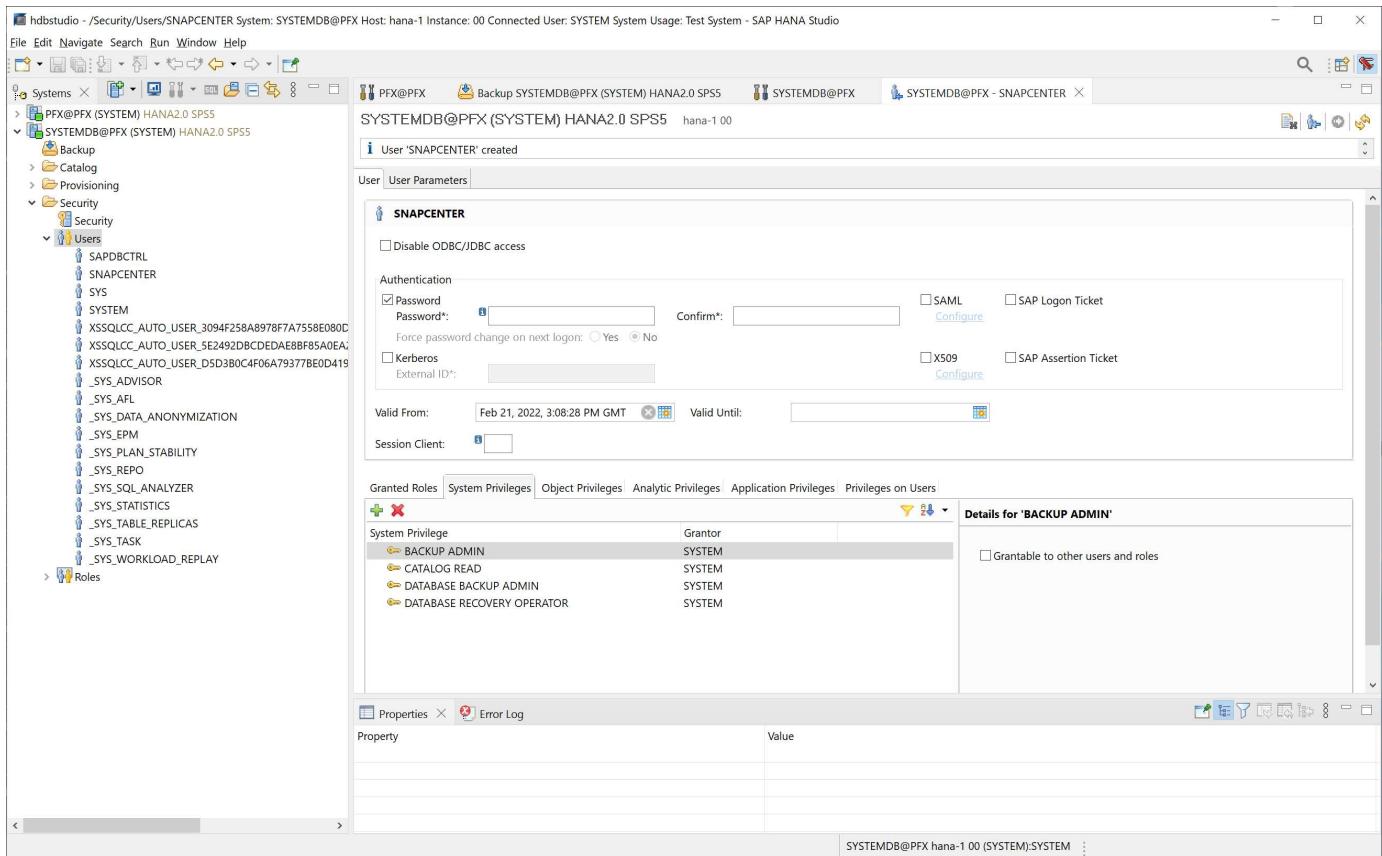
SAP HANA backup user and hdbuserstore configuration

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.

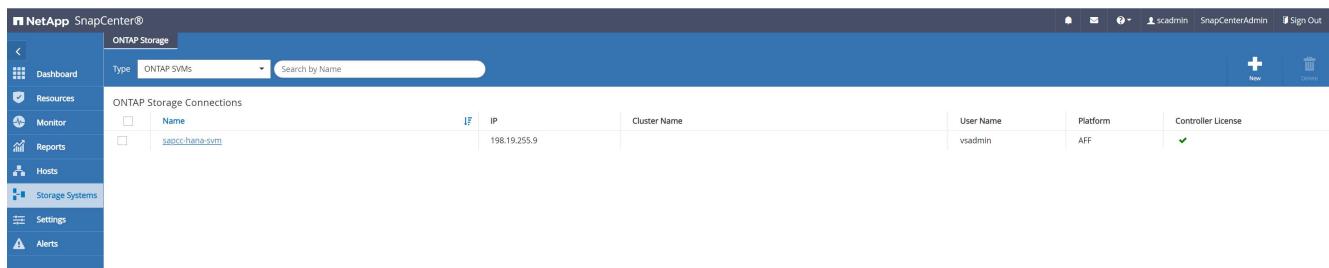
You can select the storage system type, which can be ONTAP SVMs or ONTAP Clusters. In the following example, SVM management is selected.

2. 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”.

3. To configure the storage platform, click **More Options**.
4. Select All Flash FAS as the storage system to ensure that the license, which is part of FSx for ONTAP, is available for SnapCenter.

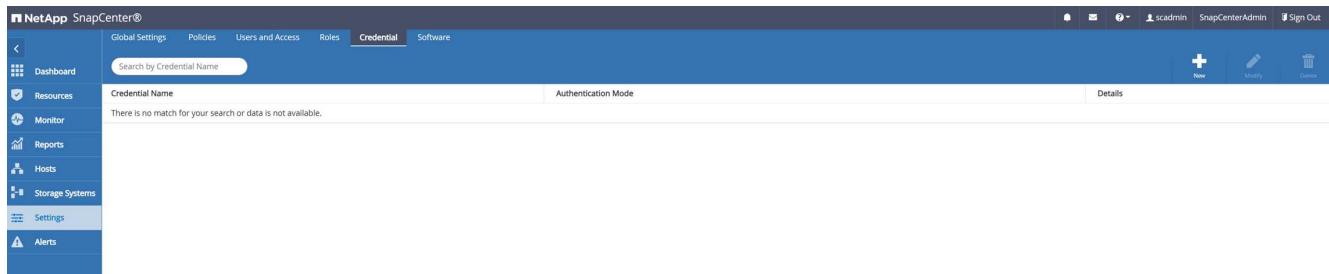
The SVM `sapcc-hana-svm` is now configured in SnapCenter.



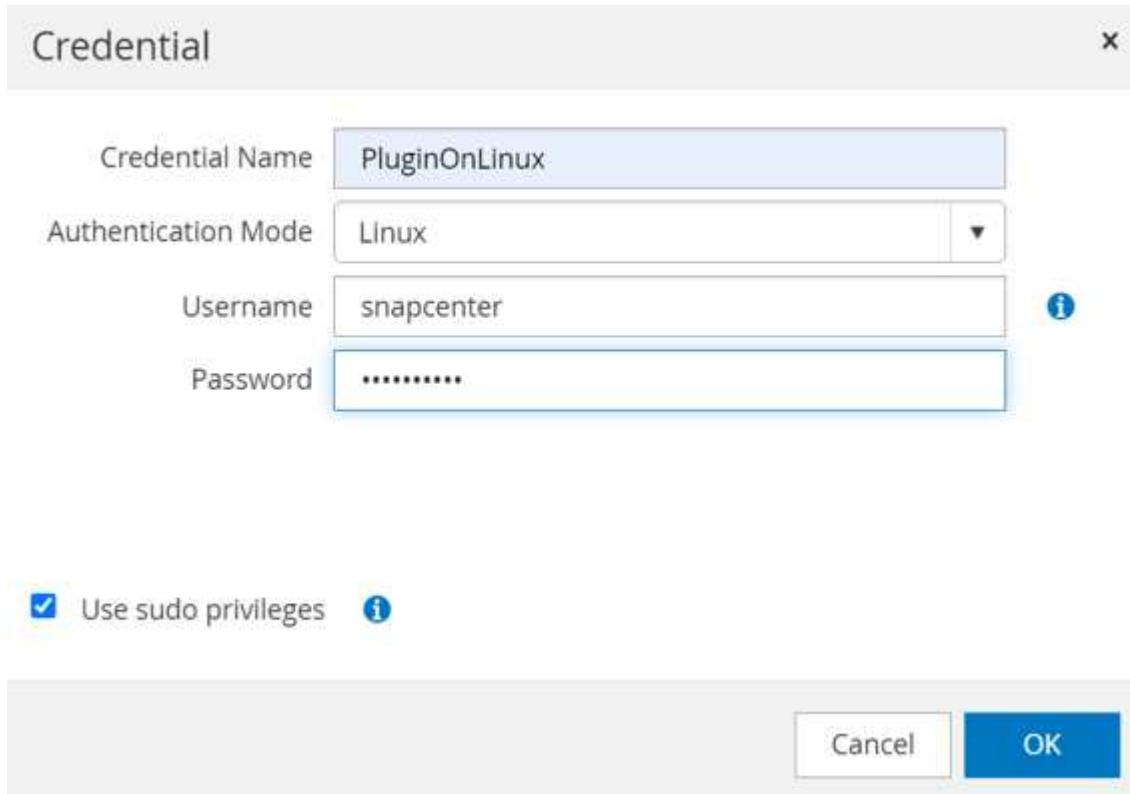
Create credentials for plugin deployment

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.

```

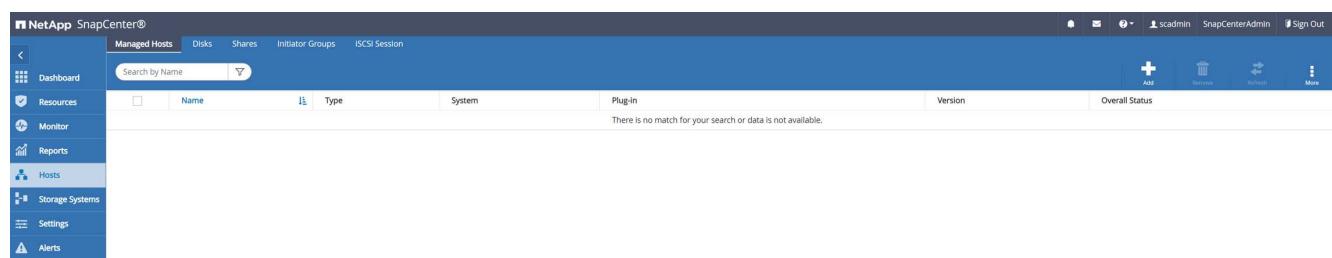
hana-1:/etc/ssh # java -version
openjdk version "1.8.0_312"
OpenJDK Runtime Environment (IcedTea 3.21.0) (build 1.8.0_312-b07 suse-3.61.3-x86_64)
OpenJDK 64-Bit Server VM (build 25.312-b07, mixed mode)
hana-1:/etc/ssh #

```

OpenJDK or Oracle Java is supported with SnapCenter.

To add the SAP HANA host, follow these steps:

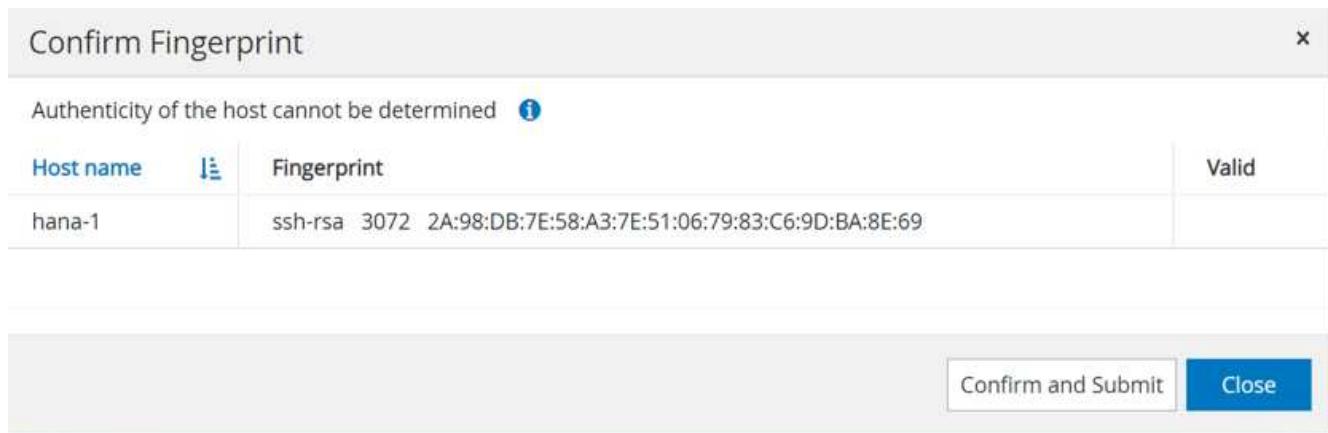
1. From the host tab, click Add.



2. Provide host information and select the SAP HANA plug-in to be installed. Click Submit.



3. Confirm the fingerprint.

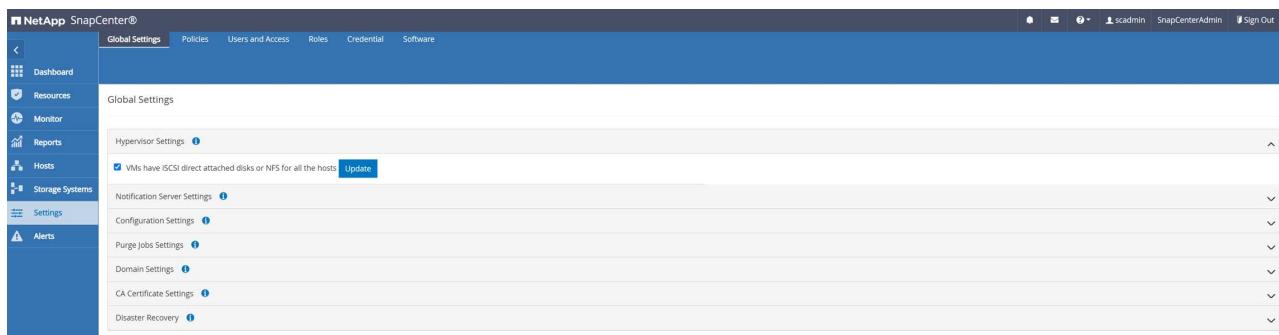


The installation of the HANA and the Linux plug-in starts automatically. When the installation is finished, the status column of the host shows Configure VMware Plug-in. SnapCenter detects if the SAP HANA plug-in is installed on a virtualized environment. This might be a VMware environment or an environment 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.



- a. From the Settings tab, select Global Settings.
- b. 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

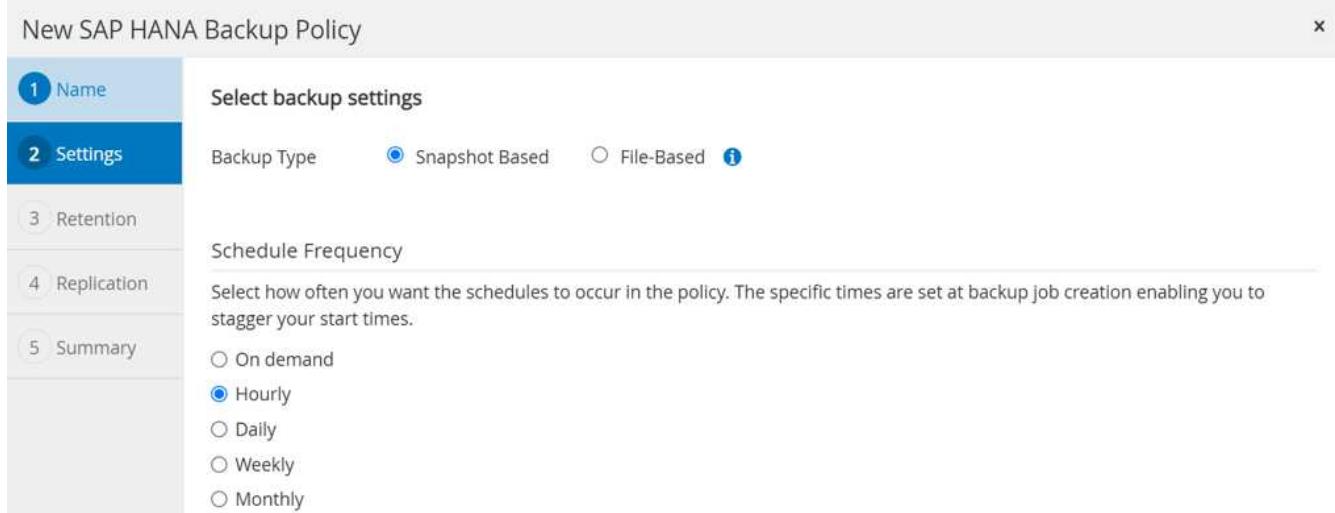
Follow these steps to configure Snapshot backup policies.

1. Go to Settings > Policies and click New.

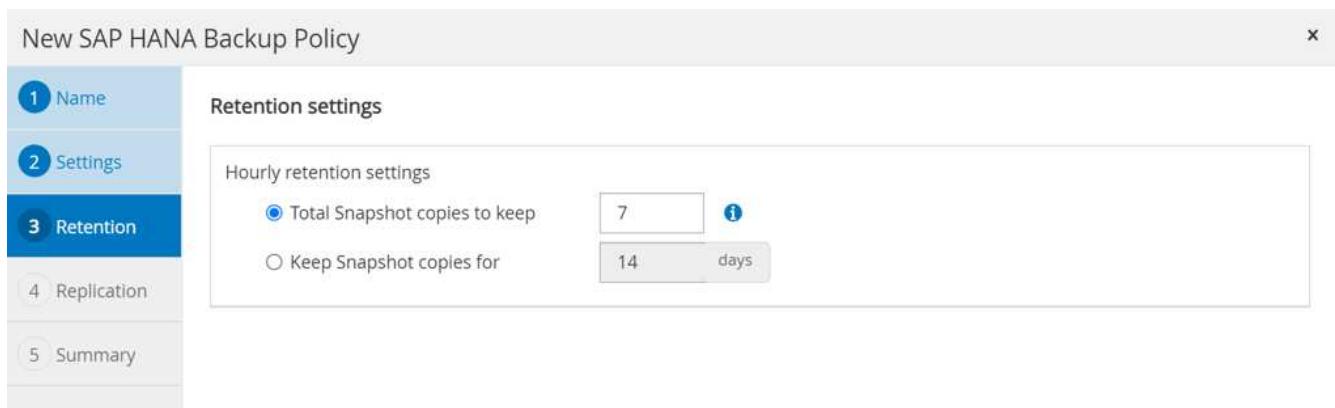
2. Enter the policy name and description. Click Next.

3. Select backup type as Snapshot Based and select Hourly for schedule frequency.

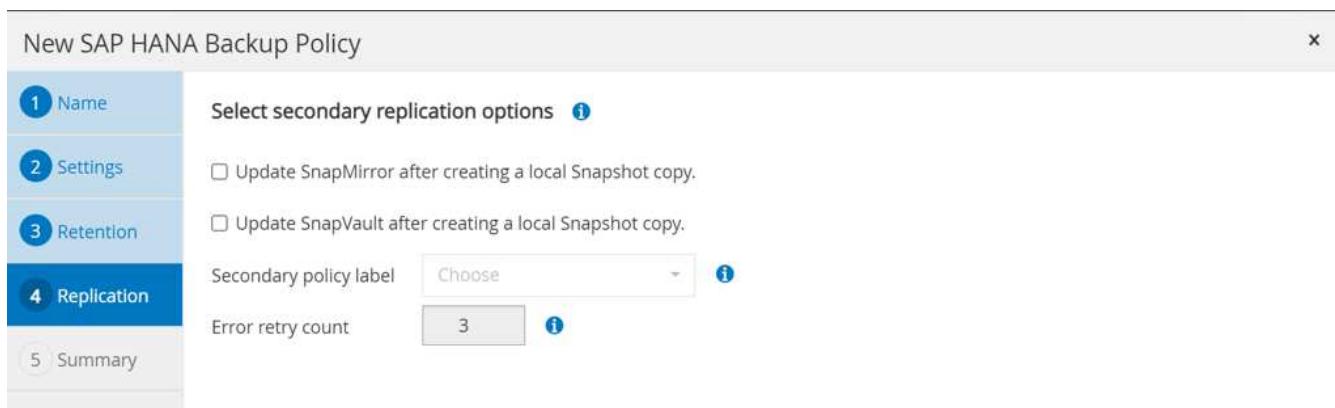
The schedule itself is configured later with the HANA resource protection configuration.



4. Configure the retention settings for on-demand backups.



5. Configure the replication options. In this case, no SnapVault or SnapMirror update is selected.



New SAP HANA Backup Policy

1 Name	Summary	
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.

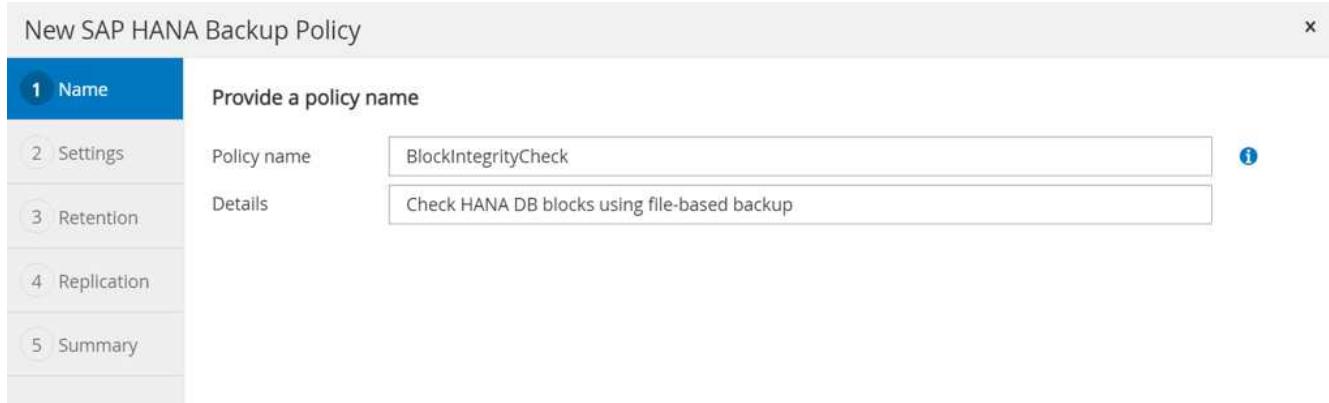


The screenshot shows the NetApp SnapCenter Policies interface. The top navigation bar includes Global Settings, Policies (selected), Users and Access, Roles, Credential, and Software. The left sidebar has links for Dashboard, Resources, Monitor, Reports, Hosts, Storage Systems, Settings (selected), and Alerts. The main content area shows a table with one row for the 'LocalSnap' policy. The columns are Name (LocalSnap), Backup Type (Data Backup), Schedule Type (Hourly), and Replication (none). Action buttons for New, Modify, Copy, Details, and Delete are available for each row.

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.



The screenshot shows the 'New SAP HANA Backup Policy' configuration page. The top navigation bar is identical to the previous screenshot. The left sidebar shows the 'Name' step is selected. The main form has a 'Provide a policy name' section with 'Policy name' set to 'BlockIntegrityCheck' and 'Details' set to 'Check HANA DB blocks using file-based backup'. The right sidebar shows the remaining steps: 2 Settings, 3 Retention, 4 Replication, and 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.

x

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.

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 
 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

NetApp SnapCenter®

Global Settings Policies Users and Access Roles Credential Software

Dashboard Resources Monitor Reports Hosts Storage Systems Settings Alerts

Policies

SAP HANA

Search by Name

Name Backup Type Schedule Type Replication

BlockIntegrityCheck	File Based Backup	Weekly	New
LocalSnap	Data Backup	Hourly	Modify

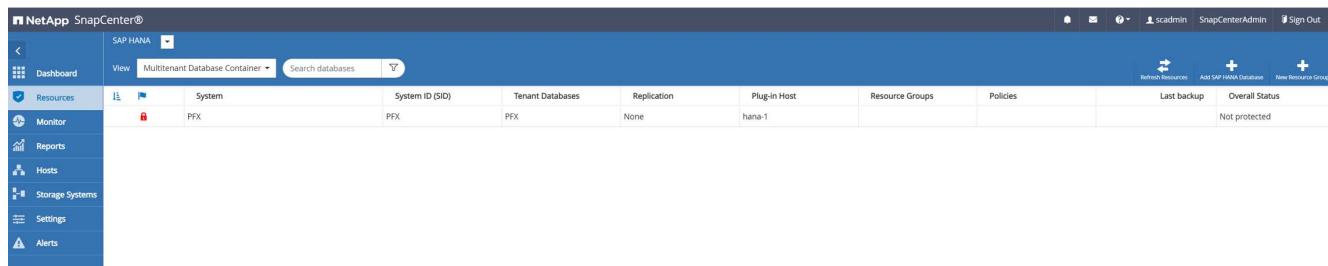
    

Configure and protect a HANA resource

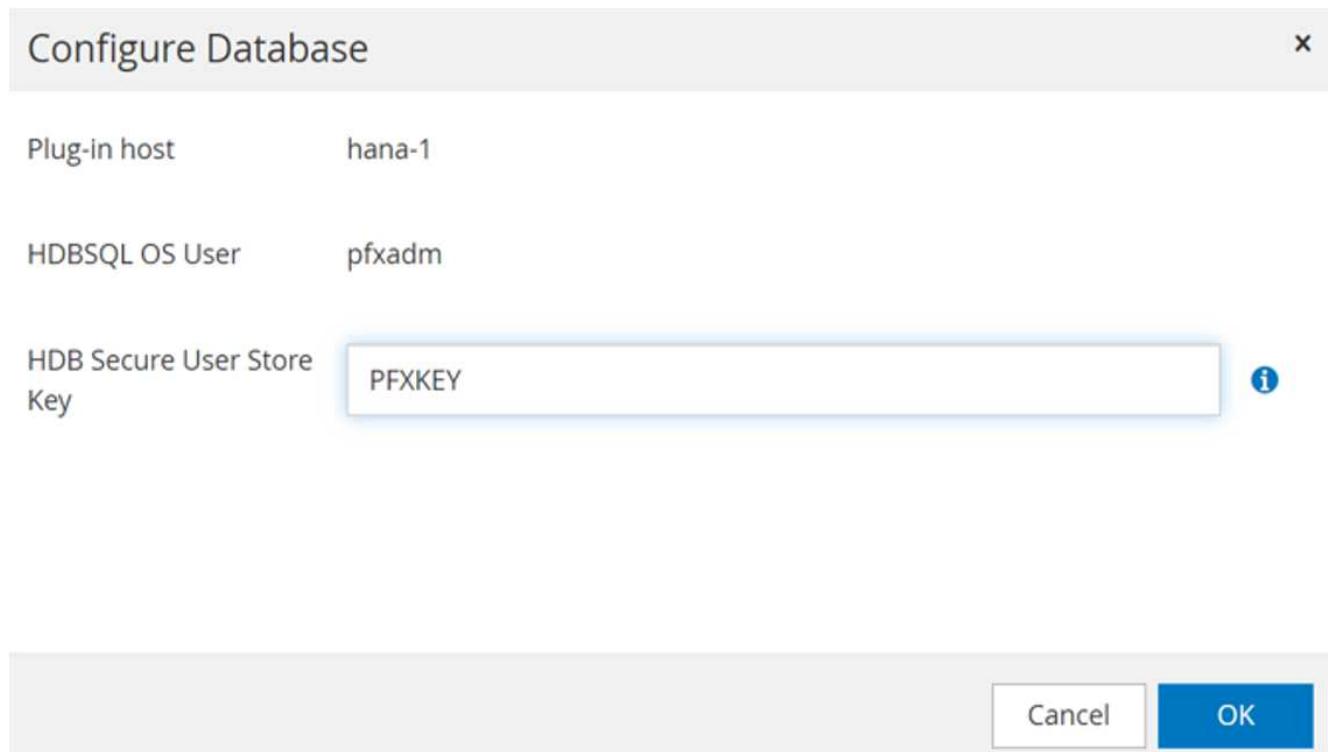
After the plug-in installation, the automatic discovery process of the HANA resource starts automatically. In the Resources screen, a new resource is created, which is marked as locked with the red padlock icon. To configure and protect the new HANA resource, follow these steps:

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.



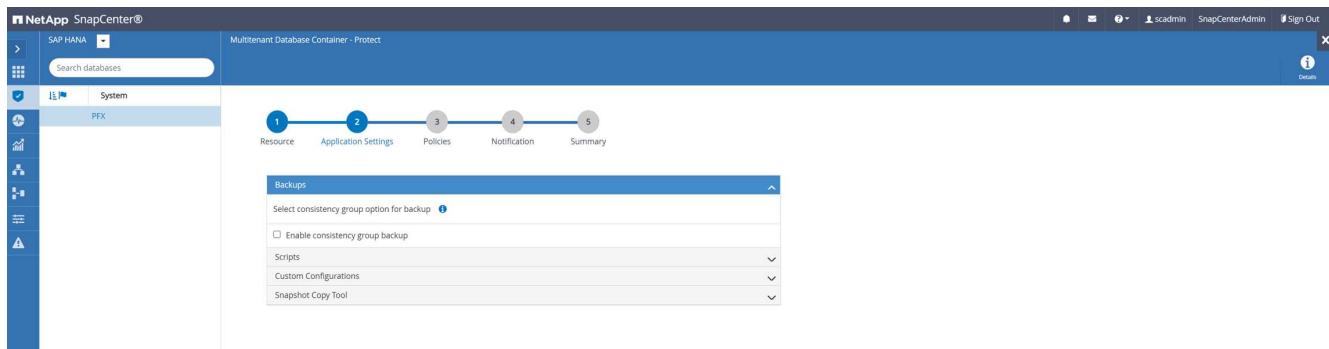
The second level automatic discovery process starts in which tenant data and storage footprint information is discovered.

3. From the Resources tab, double click the resource to configure the resource protection.

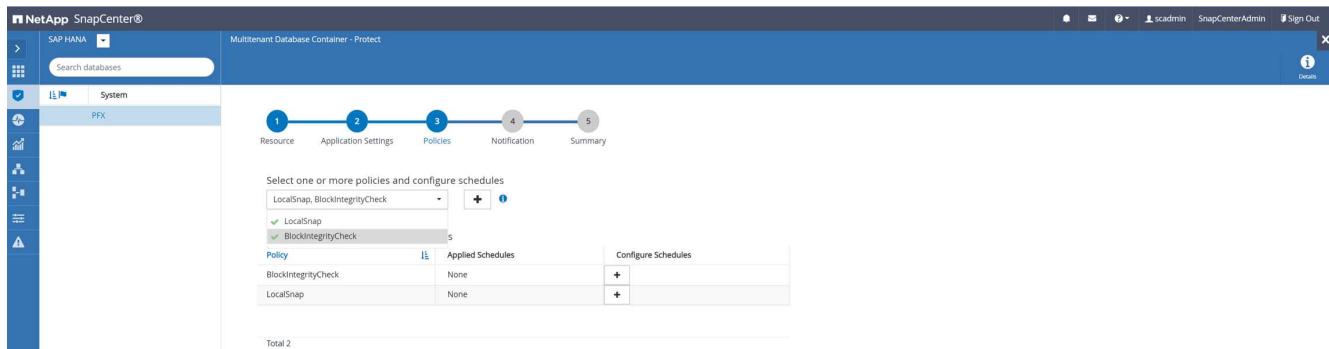
4. Configure a custom name format for the Snapshot copy.

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

X

Weekly

Start date	02/22/2022 12:00 pm	
<input type="checkbox"/> 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.

X

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

x

Hourly

Start date	02/22/2022 02:00 pm			
<input type="checkbox"/> Expires on	04/28/2022 11:57 am			
Repeat every	6	hours	0	mins

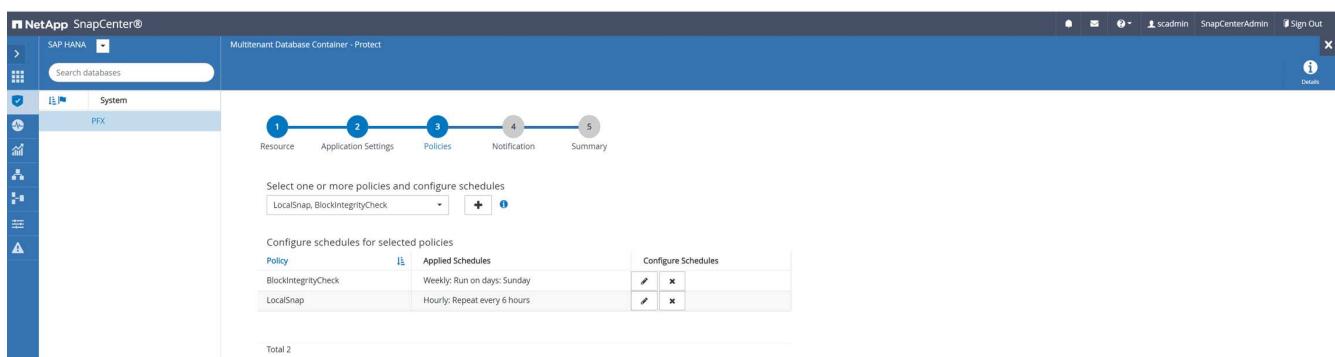


The schedules are triggered in the SnapCenter Server time zone.



Cancel

OK

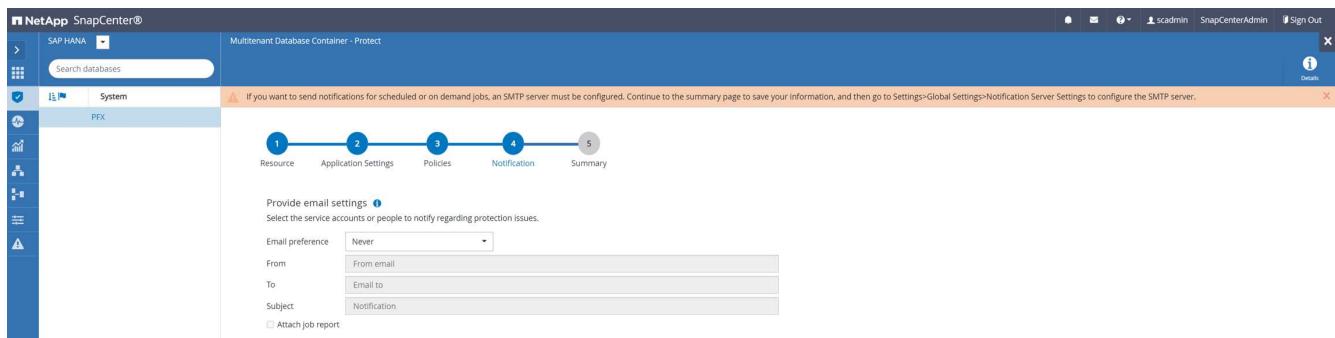


The screenshot shows the 'Multitenant Database Container - Protect' interface in NetApp SnapCenter. The 'Configure Schedules' step is selected. It displays two scheduled tasks for the 'LocalSnap' policy:

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.

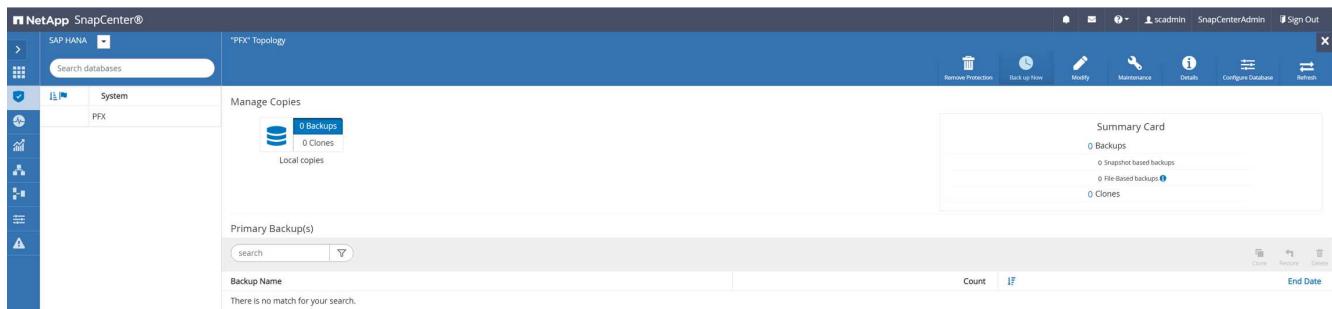
Create an on-demand Snapshot backup

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.



3. From the drop-down list, select the backup policy `LocalSnap`, and then click `Backup` to start the on-demand backup.

Backup

Create a backup for the selected resource

Resource Name: PFX

Policy: LocalSnap ?

Cancel Backup

Confirmation

x



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

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

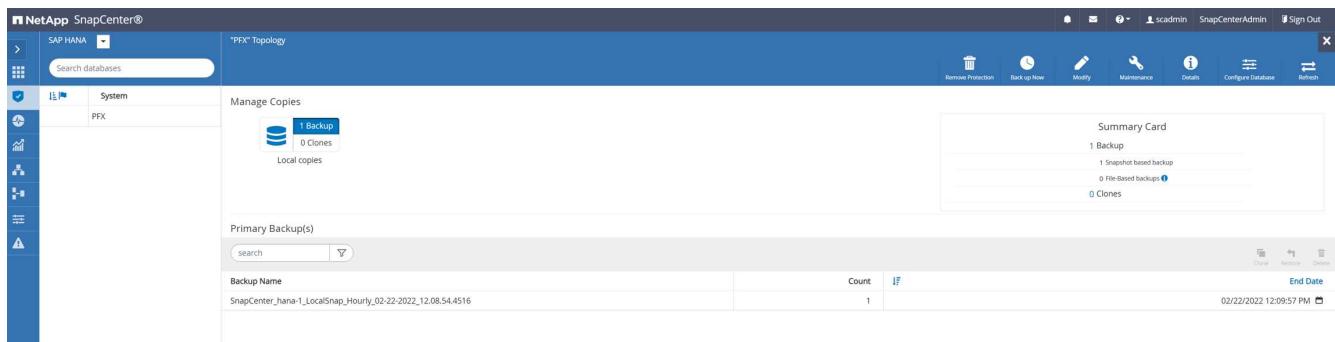
[View Logs](#)

[Cancel Job](#)

[Close](#)

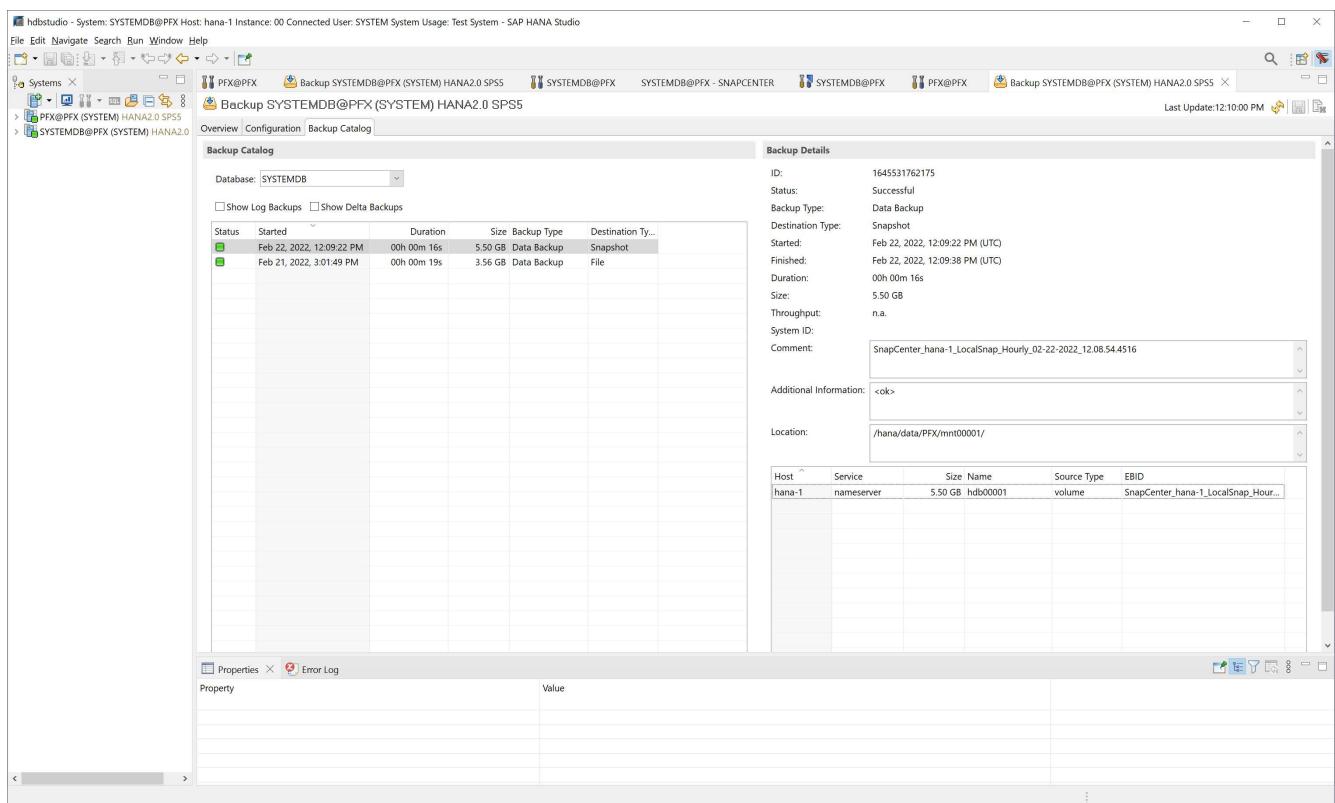
When the backup is finished, a new entry is shown in the topology view. The backup names follow the same naming convention as the Snapshot name defined in the section ["Configure and protect a HANA resource"](#).

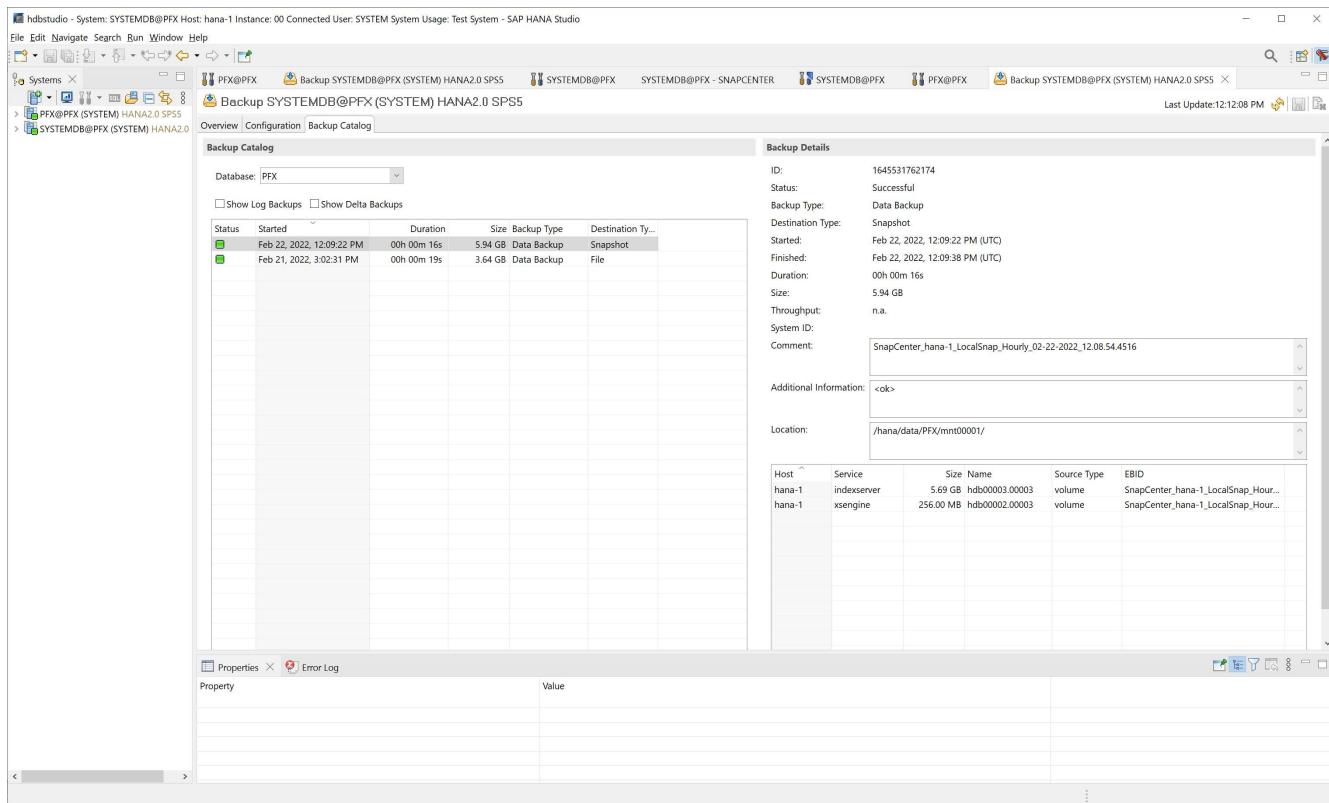
You must close and reopen the topology view to see the updated backup list.



Backup Name	Count	End Date
SnapCenter_hana-1_LocalSnap_Hourly_02-22-2022_12.08.54.4516	1	02/22/2022 12:09:57 PM

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
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

X

Create a backup for the selected resource

Resource Name

PFX

Policy

BlockIntegrityCheck

i

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

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](#)

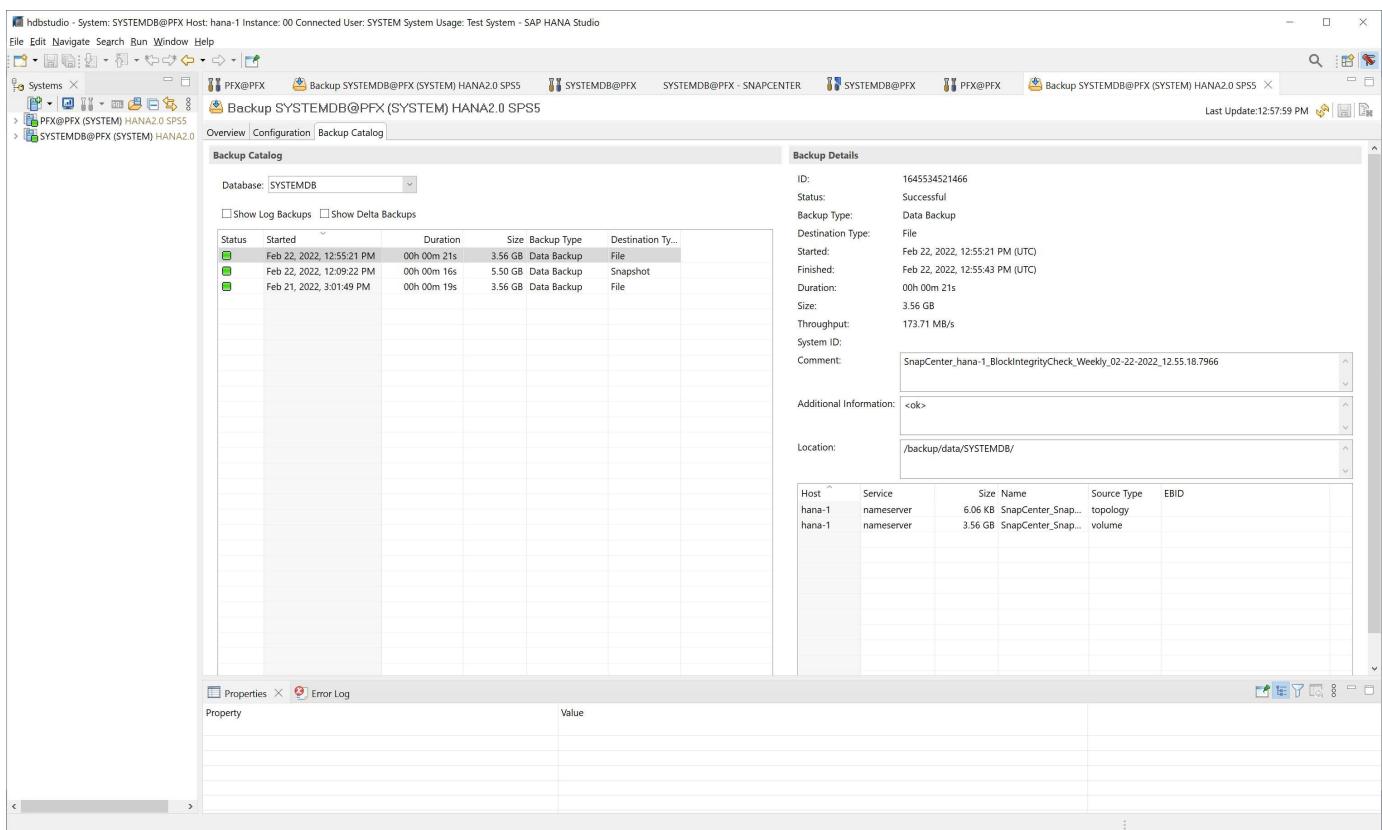
SnapCenter does not display the block integrity check in the same manner as Snapshot copy-based backups.

Instead, the summary card shows the number of file-based backups and the status of the previous backup.



The screenshot shows the NetApp SnapCenter interface. The top navigation bar includes 'SAP HANA', 'PFX Topology', and various administrative buttons. The main area displays a 'Manage Copies' section with a summary card. The summary card indicates '1 Backup' and '0 Clones', with a note that it's a 'Local copies'. The summary card also shows a 'Summary Card' with '2 Backups', '1 Snapshot based backup', and '1 file-based backup ✓'. It also shows the 'Last Backup 2/22/2022 12:56:25 PM' and 'Backup succeeded'. Below the summary card, there are buttons for 'Done', 'Revert', and 'Delete'. A table below the summary card lists 'Backup Name' as 'SnapCenter_hana-1_LocalSnap_Hourly_02-22-2022_12.08.54.4516' with a 'Count' of 1. The bottom right corner shows the date and time as '02/22/2022 12:09:57 PM'.

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.



The screenshot shows the SAP HANA Studio interface with the title 'hdbstudio - System: SYSTEMDB@PFX Host: hana-1 Instance: 00 Connected User: SYSTEM System Usage: Test System - SAP HANA Studio'. The main window is titled 'Backup Catalog' and shows a table of backups for the 'SYSTEMDB' database. The table includes columns for 'Status', 'Started', 'Duration', 'Size', 'Backup Type', and 'Destination Type'. The table lists three backups: one successful data backup and two snapshots. To the right of the table is a 'Backup Details' panel. The 'Backup Details' panel shows the following information for the most recent backup:

ID:	1645534521466
Status:	Successful
Backup Type:	Data Backup
Destination Type:	File
Started:	Feb 22, 2022, 12:55:21 PM (UTC)
Finished:	Feb 22, 2022, 12:55:43 PM (UTC)
Duration:	00h 00m 21s
Size:	3.56 GB
Throughput:	173.71 MB/s
System ID:	
Comment:	SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-2022_12.55.18.7966
Additional Information:	<ok>
Location:	/backup/data/SYSTEMDB/

Below the 'Backup Details' panel is a table showing the destination details for the backup, listing 'Host' (hana-1), 'Service' (nameserver), 'Size' (6.05 KB), 'Name' (SnapCenter_Snap...), 'Source Type' (topology), and 'EBID'.

hdbstudio - System: SYSTEMDB@PFX Host: hana-1 Instance: 00 Connected User: SYSTEM System Usage: Test System - SAP HANA Studio

File Edit Navigate Search Run Window Help

Systems X PFX@PFX Backup SYSTEMDB@PFX (SYSTEM) HANA2.0 SPSS SYSTEMDB@PFX - SNAPCENTER SYSTEMDB@PFX PFX@PFX Backup SYSTEMDB@PFX (SYSTEM) HANA2.0 SPSS

Last Update: 12:58:19 PM

Backup Catalog

Database: PFX

Show Log Backups Show Delta Backups

Status	Started	Duration	Size	Backup Type	Destination Ty...
Success	Feb 22, 2022, 12:55:34 PM	00h 00m 27s	3.64 GB	Data Backup	File
Success	Feb 22, 2022, 12:09:22 PM	00h 00m 16s	5.94 GB	Data Backup	Snapshot
Success	Feb 21, 2022, 3:02:31 PM	00h 00m 19s	3.64 GB	Data Backup	File

Backup Details

ID: 1645534534230
 Status: Successful
 Backup Type: Data Backup
 Destination Type: File
 Started: Feb 22, 2022, 12:55:34 PM (UTC)
 Finished: Feb 22, 2022, 12:56:01 PM (UTC)
 Duration: 00h 00m 27s
 Size: 3.64 GB
 Throughput: 138.07 MB/s
 System ID:
 Comment: SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-2022_12.55.18.7966
 Additional Information: <ok>
 Location: /backup/data/DB_PFX/

Host	Service	Size	Name	Source Type	EBID
hana-1	indexserver	1.58 KB	SnapCenter_Snap...	topology	
hana-1	xengine	80.00 MB	SnapCenter_Snap...	volume	
hana-1	indexserver	3.56 GB	SnapCenter_Snap...	volume	

Properties Error Log

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

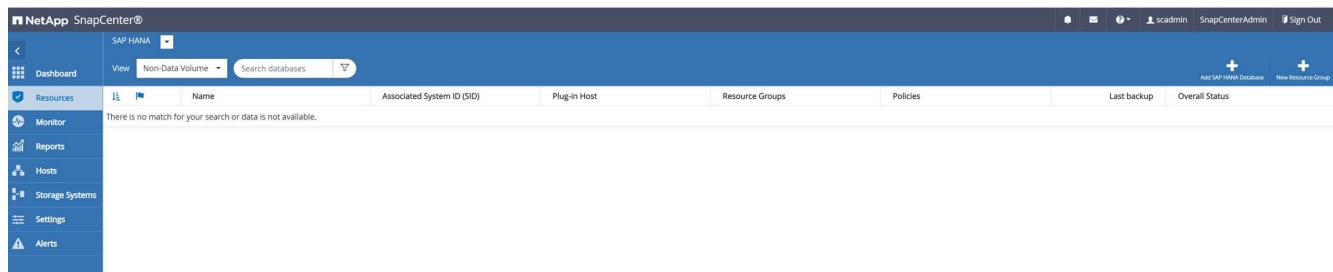
your specific requirements, the backup of non-data volumes might differ in scheduling frequency and retention settings, and you should consider how frequently non-data files are changed. For instance, the HANA volume /hana/shared contains executables but also SAP HANA trace files. While executables only 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

Provide Resource Details

2 Storage Footprint	Resource Type	Non-data Volume
3 Summary	Resource Name	PFX-Shared-Volume
	Associated SID	PFX
	Plug-in Host	hana-1

Previous **Next**

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

Storage Type ONTAP

Add Storage Footprint X

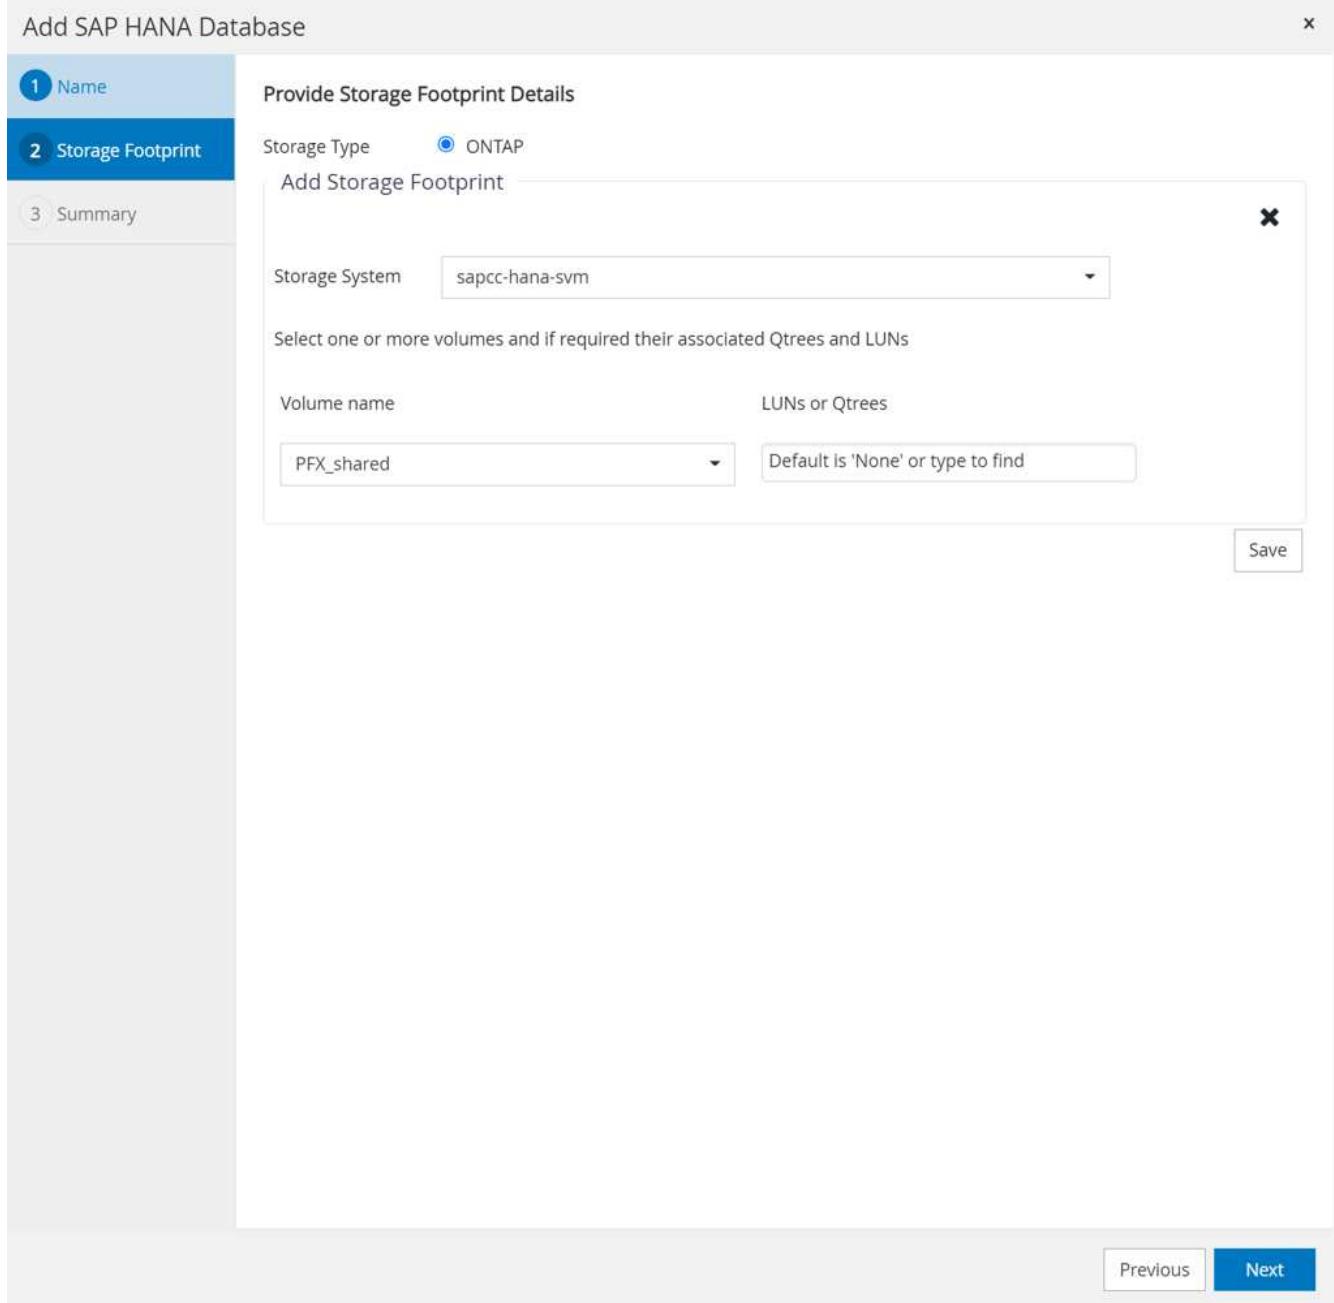
Storage System ▼

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"/>

Save

Previous Next



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.

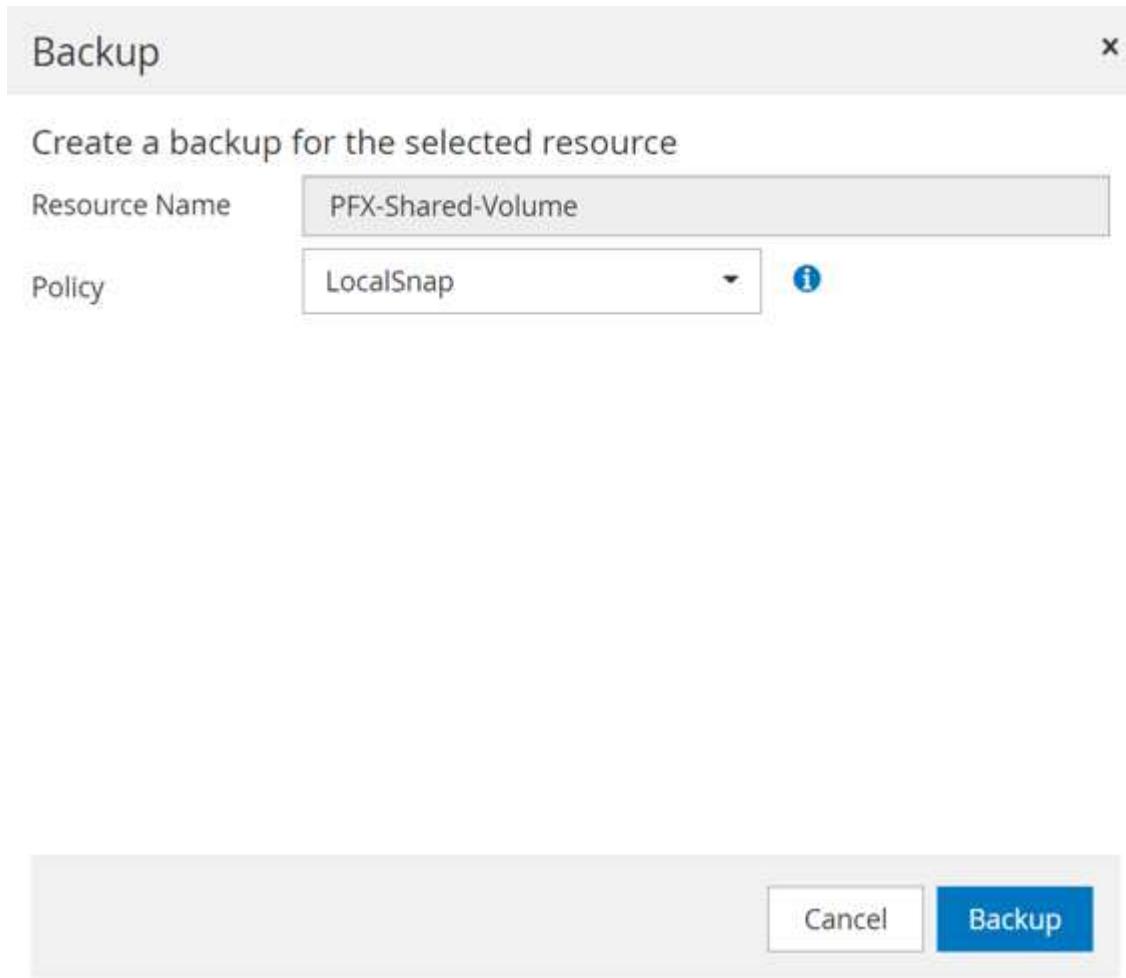
The screenshot shows the NetApp SnapCenter interface. The left sidebar has navigation links: Dashboard, Resources (selected), Monitor, Reports, Hosts, Storage Systems, Settings, and Alerts. The main content area is titled 'SAP HANA' and shows a table of resources. The table has columns: Name, Associated System ID (SID), Plug-in Host, Resource Groups, Policies, Last backup, and Overall Status. One row is visible: 'Name: PFX-Shared-Volume, Associated System ID (SID): PFX, Plug-in Host: hana-1, Overall Status: 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

X

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

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.

The screenshot shows the NetApp SnapCenter interface for a SAP HANA system. The main area displays a summary card with 5 Backups, 4 Snapshot based backups, 1 File-Based backup, and 0 Clones. Below this is a table of Primary Backup(s) with the following data:

Backup Name	Count	End Date
SnapCenter_hana-1_LocalSnap_Hourly_02-23-2022_14.00.05.4361	1	02/23/2022 2:01:11 PM
SnapCenter_hana-1_LocalSnap_Hourly_02-22-2022_20.00.01.4482	1	02/23/2022 8:01:01 PM
SnapCenter_hana-1_LocalSnap_Hourly_02-22-2022_14.00.02.8713	1	02/22/2022 2:01:01 PM
SnapCenter_hana-1_LocalSnap_Hourly_02-22-2022_12.08.54.4516	1	02/22/2022 12:09:57 PM

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://www.netapp.com).

Restore from SnapCenter_hana-1_LocalSnap_Hourly_02-23-2022_14.00.05.4361 x

1 Restore scope Select the restore types

2 Recovery scope Complete Resource i

3 PreOps

4 PostOps

5 Notification

6 Summary

Select the restore types

Complete Resource i

Volume Revert

⚠ 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.

Tenant Database

⚠ 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

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 X

1 Restore scope X

2 Recovery scope X

3 PreOps X

4 PostOps X

5 Notification X

6 Summary X

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

/backup/log

Specify backup catalog location i

/backup/log

⚠ 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 pre-restore commands.

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Enter optional commands to run before performing a restore operation [i](#)

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 x

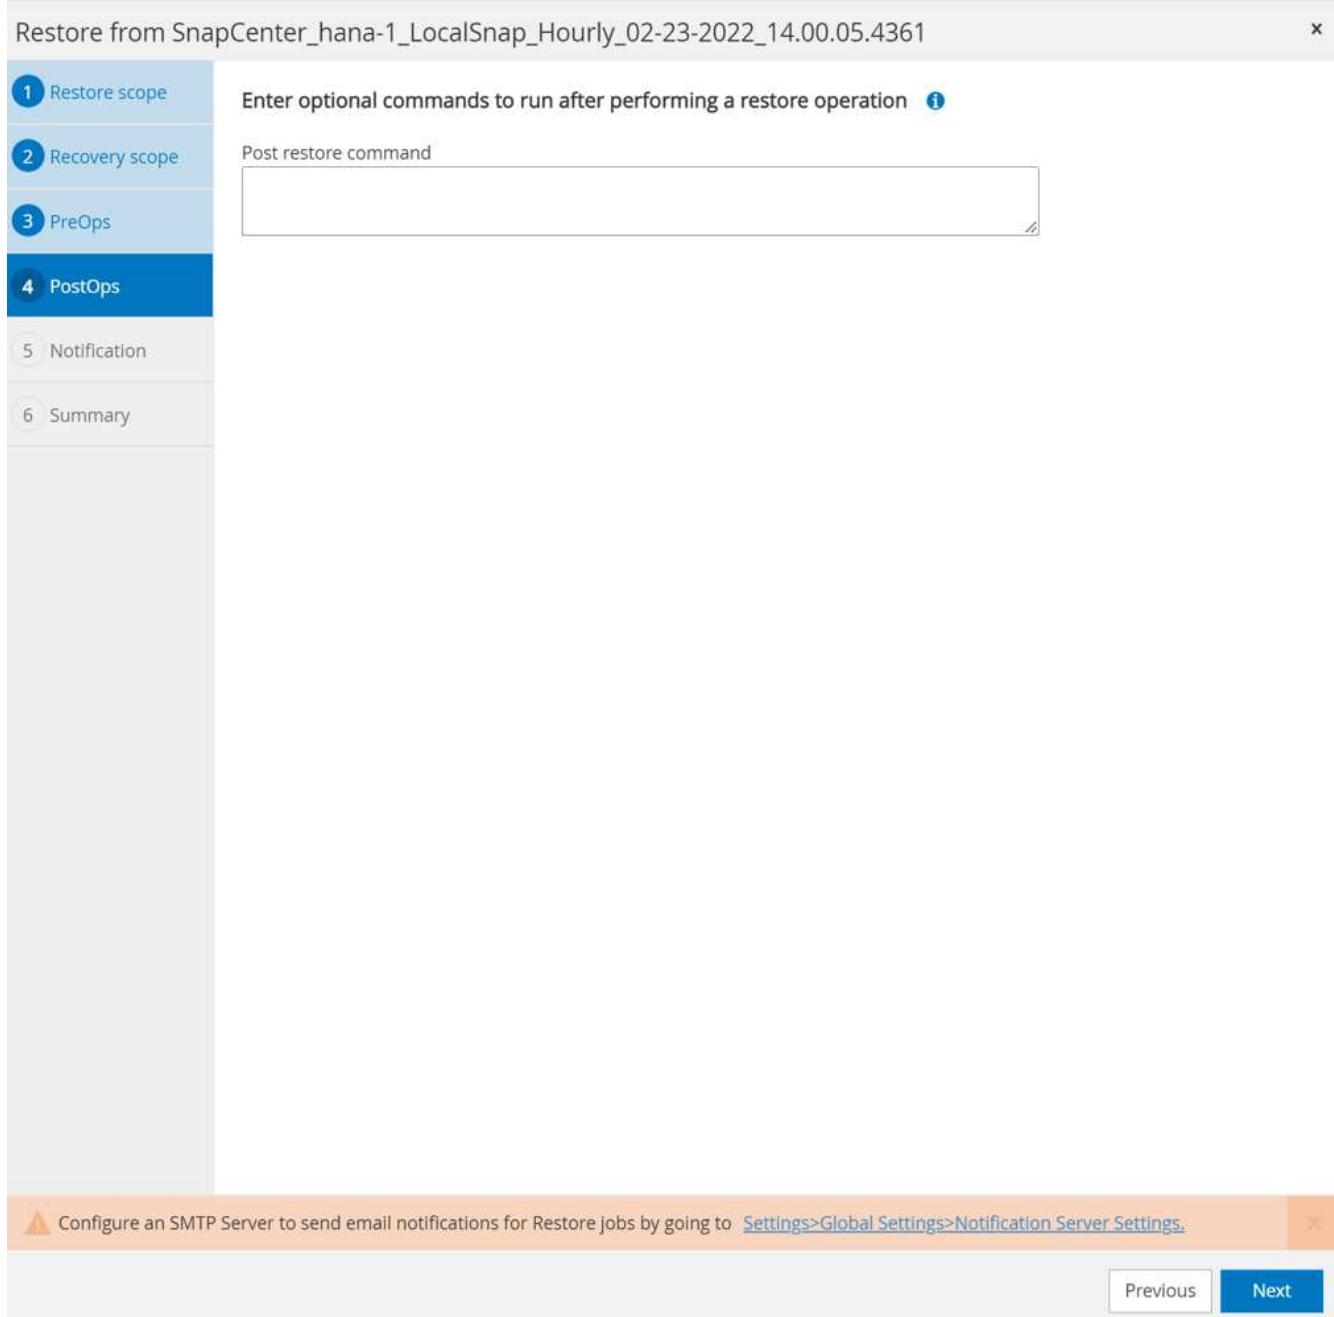
1 Restore scope i
2 Recovery scope
3 PreOps
4 PostOps
5 Notification
6 Summary

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. To start the restore and recovery operation, click Finish.

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. **x**

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

Task Name: Recover Application Start Time: 02/23/2022 2:07:31 PM End Time:

[View Logs](#)

[Cancel Job](#)

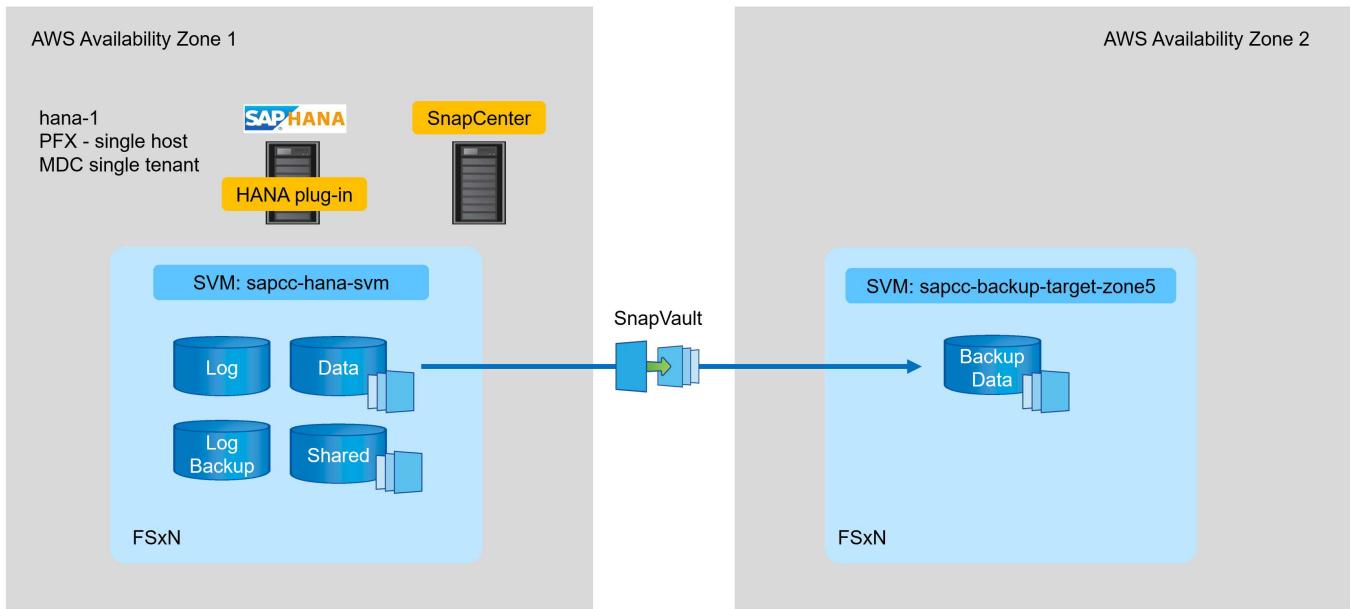
[Close](#)

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://www.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
                           e0e
true
      inter_2      up/up      10.1.2.7/24
FsxId00fa9e3c784b6abbb-02
                           e0e
true
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-addrs are cluster IPs of the destination cluster.

```

FsxId00fa9e3c784b6abbb::> cluster peer create -address-family ipv4 -peer
-addrs 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:PX_data_mnt00001 -source-path sapcc-hana-svm:PX_data_mnt00001
Operation is queued: snapmirror initialize of destination "sapcc-backup-target-zone5:PX_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:PX_data_mnt00001
          XDP   sapcc-backup-target-zone5:PX_data_mnt00001
                           Uninitialized
                           Transferring   1009MB   true
02/24 12:34:28
```

```
FsxId05f7f00af49dc7a3e::> snapmirror show

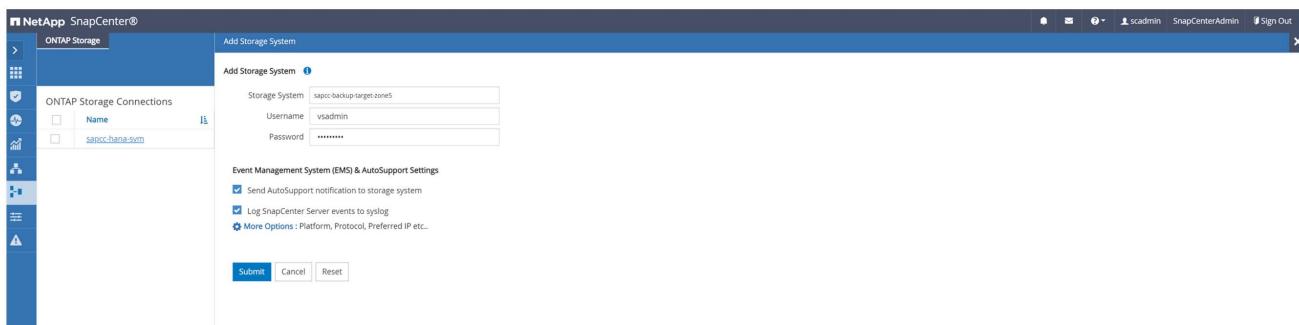
Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status      Progress  Healthy
Updated

sapcc-hana-svm:PX_data_mnt00001
          XDP   sapcc-backup-target-zone5:PX_data_mnt00001
                           Snapmirrored
                           Idle        -        true        -
```

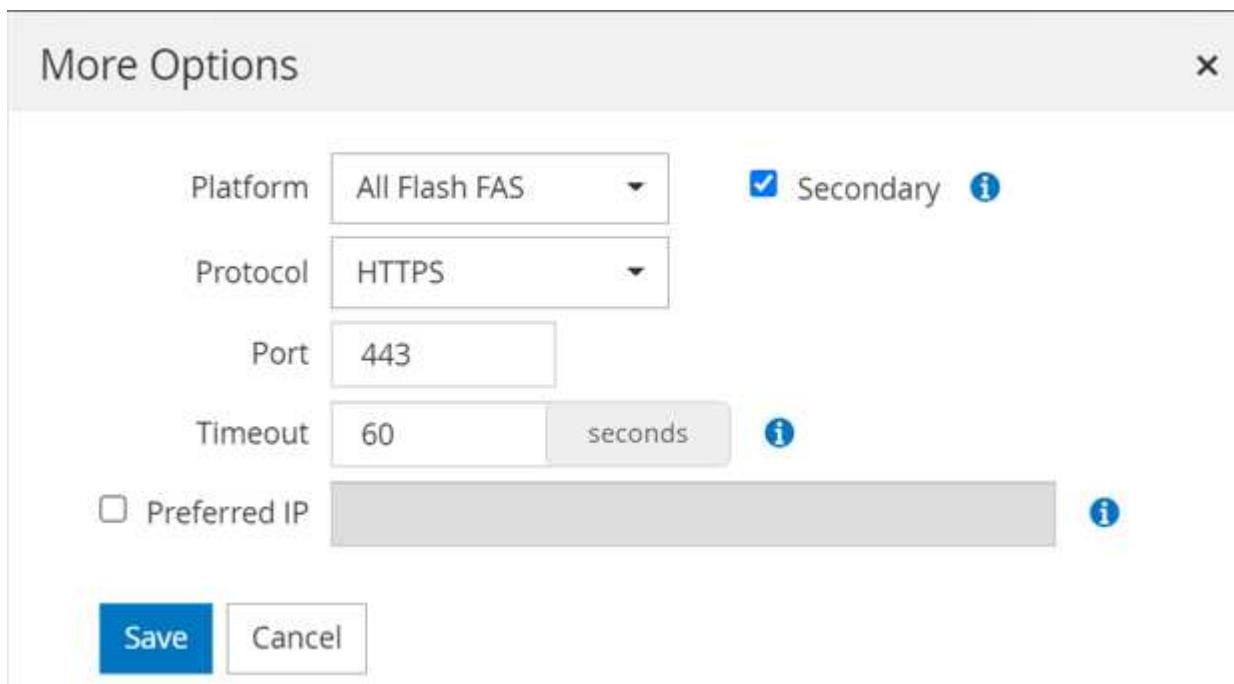
Add a backup SVM to SnapCenter

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.



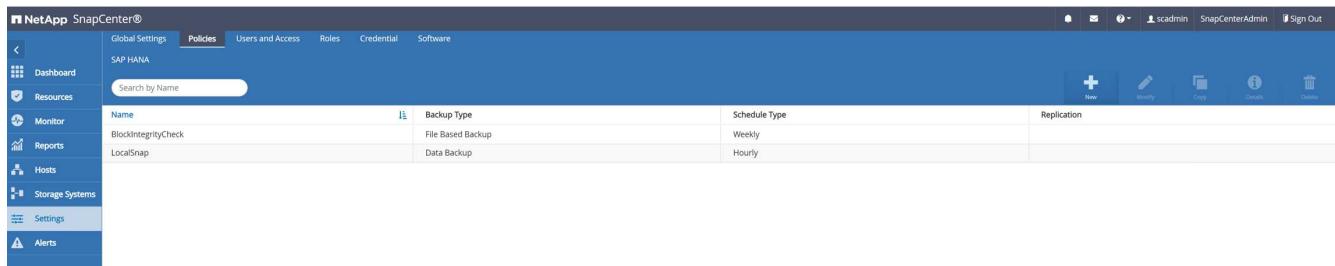
The SVM is now available in SnapCenter.

ONTAP Storage						
ONTAP Storage Connections		Type	Search by Name	Actions		
Name	IP	Cluster Name	User Name	Platform	Controller License	
sapcc-backup-target-zone5	10.1.2.31		vsadmin	AFF	Not applicable	
sapcc-hana-svm	198.19.255.9		vsadmin	AFF	✓	

Create a new SnapCenter policy for backup replication

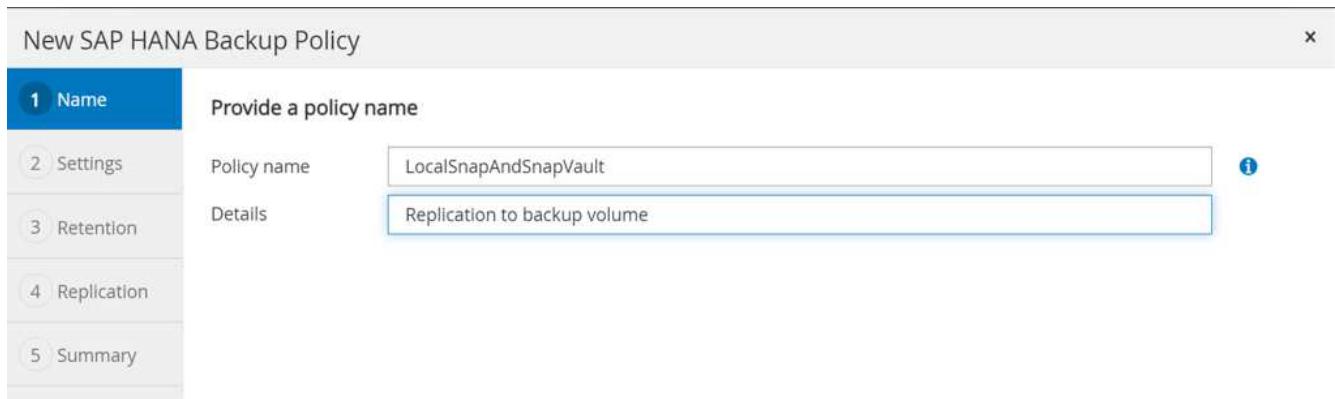
You must configure a policy for the backup replication as follows:

1. Provide a name for the policy.



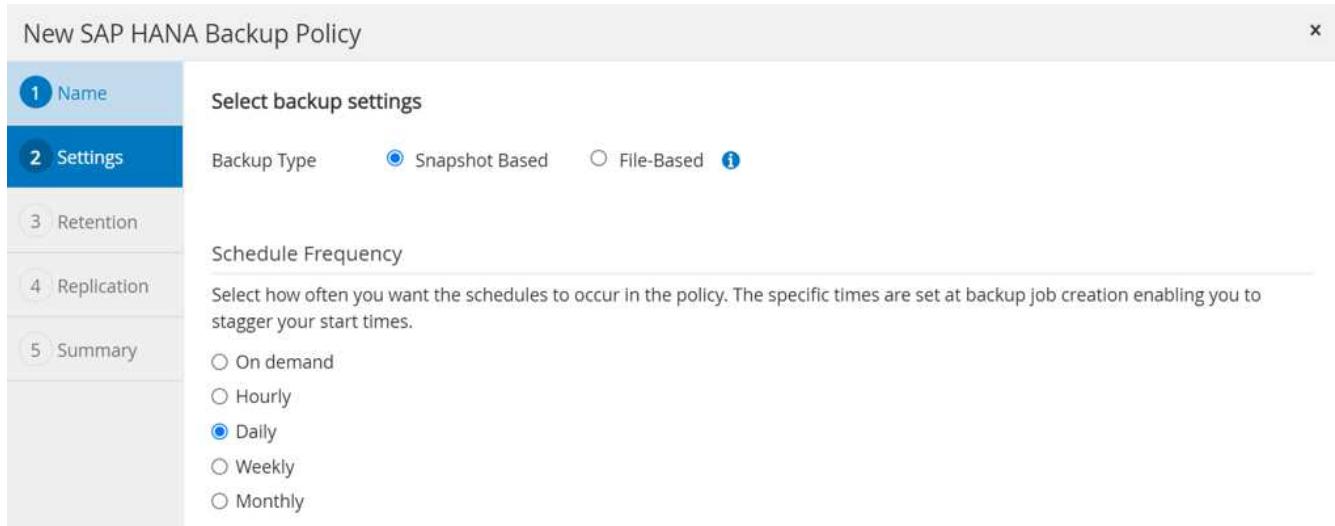
The screenshot shows the NetApp SnapCenter Policies page. The left sidebar includes options like Dashboard, Resources, Monitor, Reports, Hosts, Storage Systems, Settings, and Alerts. The main content area is titled 'SAP HANA' and shows a table of policies. The table columns are 'Name', 'Backup Type', 'Schedule Type', and 'Replication'. There are three entries: 'BlockIntegrityCheck' (File Based Backup, Weekly), 'LocalSnap' (Data Backup, Hourly), and a new row being edited with a placeholder 'Name' and 'Backup Type'.

2. Select Snapshot backup and a schedule frequency. Daily is typically used for backup replication.



The screenshot shows the 'New SAP HANA Backup Policy' dialog. The left sidebar has steps 1-5: 1. Name, 2. Settings, 3. Retention, 4. Replication, 5. Summary. Step 1 is active. The main area is titled 'Provide a policy name' and shows 'Policy name: LocalSnapAndSnapVault' and 'Details: Replication to backup volume'.

3. Select the retention for the Snapshot backups.



The screenshot shows the 'New SAP HANA Backup Policy' dialog. The left sidebar has steps 1-5: 1. Name, 2. Settings, 3. Retention, 4. Replication, 5. Summary. Step 2 is active. The main area is titled 'Select backup settings' and shows 'Backup Type: Snapshot Based' (radio button selected) and 'File-Based' (radio button unselected). Below is a 'Schedule Frequency' section with options: 'On demand', 'Hourly', 'Daily' (radio button selected), 'Weekly', and 'Monthly'.

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

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Retention settings

Daily retention settings

Total Snapshot copies to keep

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

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Select secondary replication options

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label

Error retry count

New SAP HANA Backup Policy

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.

NetApp SnapCenter®

Global Settings Policies Users and Access Roles Credential Software

Dashboard Resources Monitor Reports Hosts Storage Systems Settings Alerts

SAP HANA

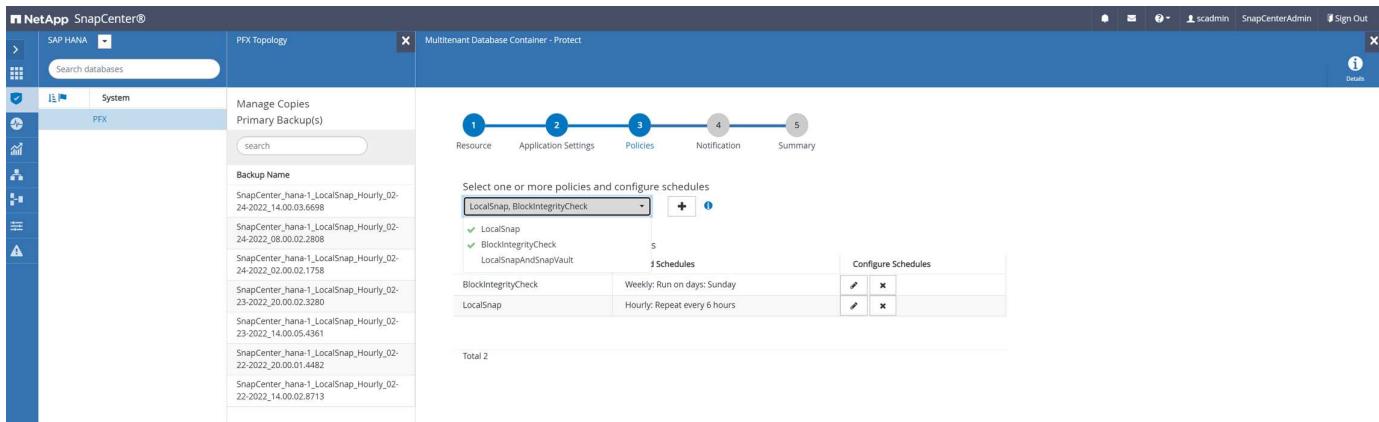
Search by Name

Name Backup Type Schedule Type Replication

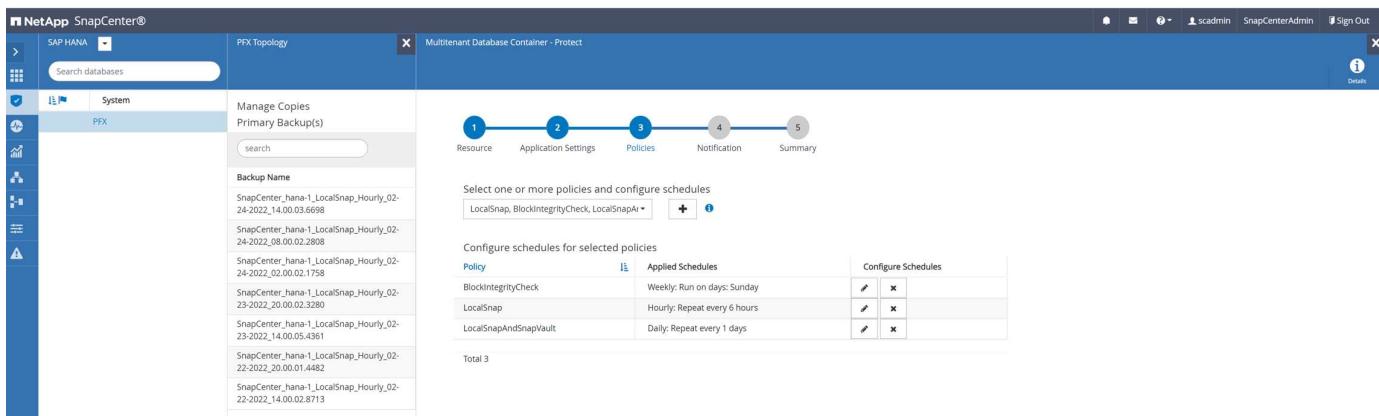
BlockIntegrityCheck	File Based Backup	Weekly	
LocalSnap	Data Backup	Hourly	
LocalSnapAndSnapVault	Data Backup	Daily	SnapVault

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

X

Create a backup for the selected resource

Resource Name

PFX

Policy

LocalSnapAndSnapVault

i

Cancel

Backup

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 X

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,

snapshot, 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...
Green	Apr 28, 2022, 4:22:06 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
Green	Apr 28, 2022, 2:00:26 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
Green	Apr 28, 2022, 8:00:35 AM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
Green	Apr 15, 2022, 5:00:44 PM	00h 06m 59s	5.50 GB	Data Backup	Snapshot
Green	Apr 14, 2022, 8:00:32 PM	00h 00m 16s	5.50 GB	Data Backup	Snapshot
Green	Apr 5, 2022, 2:00:29 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
Green	Apr 5, 2022, 8:00:39 AM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
Green	Mar 31, 2022, 2:00:29 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
Green	Mar 31, 2022, 1:10:57 PM	00h 00m 16s	5.50 GB	Data Backup	Snapshot
Green	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.

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.

The restore dialog shows the secondary locations.

- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Select the restore types

Complete Resource i

Tenant Database

Choose archive location

sapcc-hana-svm:PFX_data_mnt0001

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. x

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

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 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.

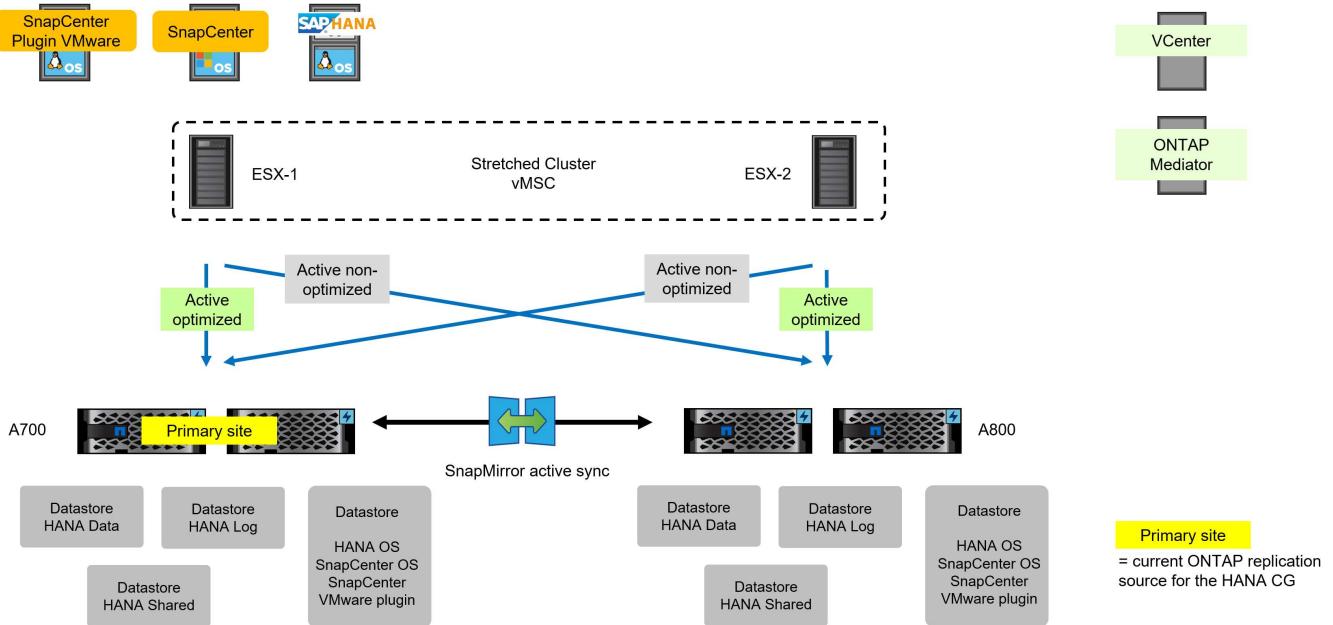
Author: Nils Bauer, NetApp

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

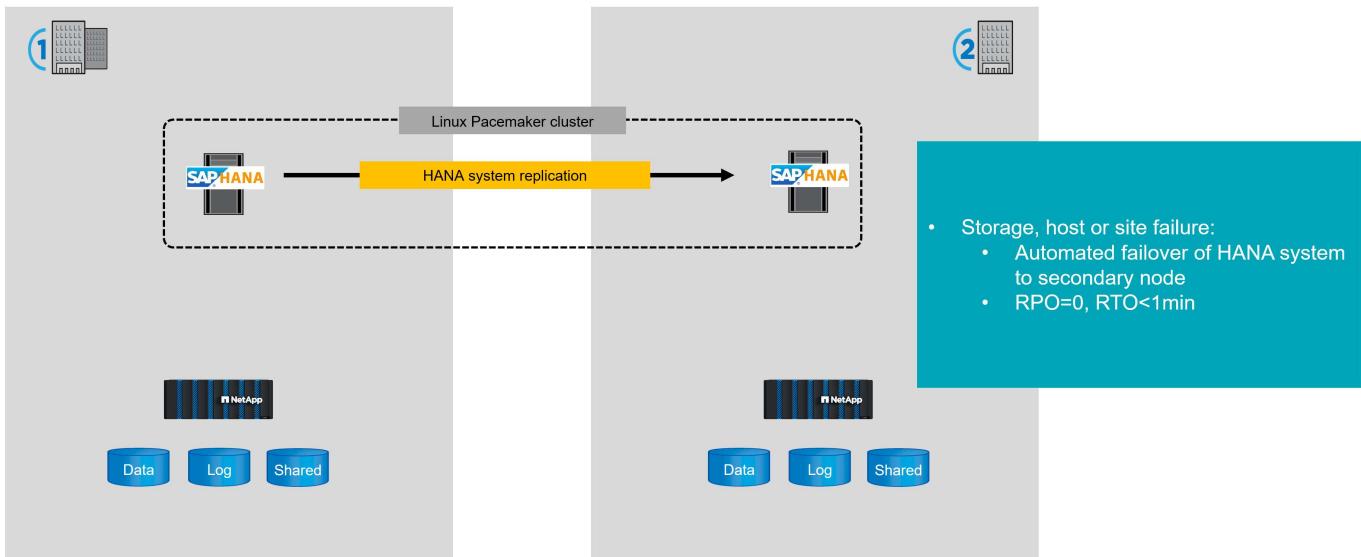
This chapter provides an overview of high availability options for SAP HANA comparing replication on application layer with storage replication.

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.

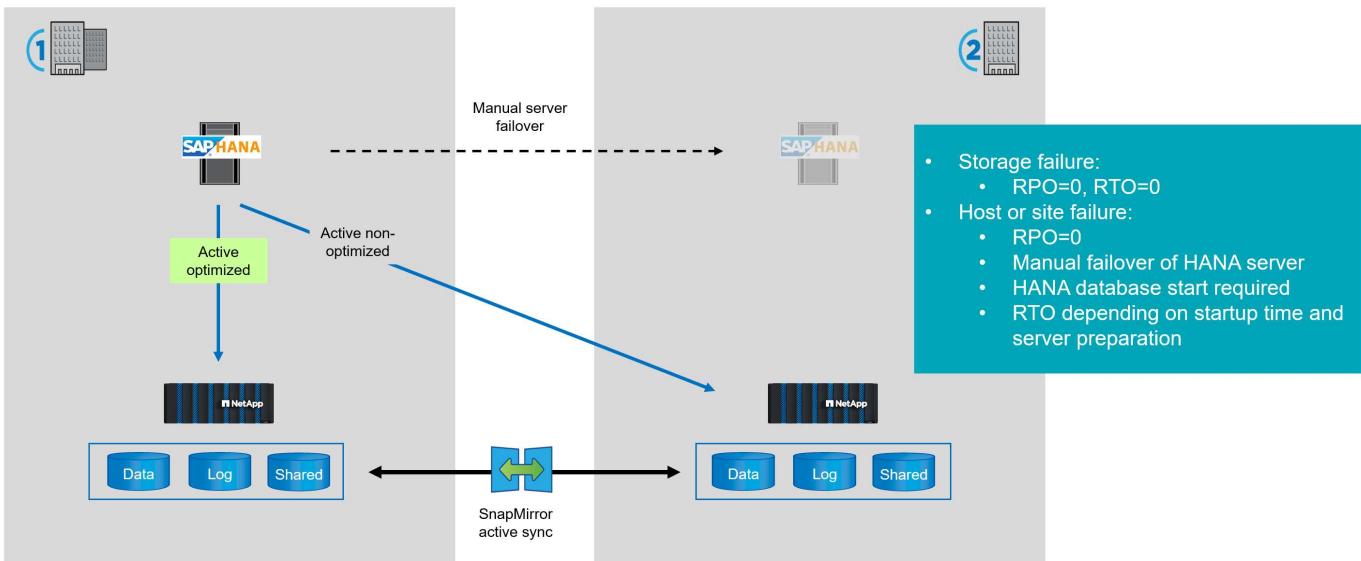
More details can be found at [SnapMirror active sync overview in ONTAP](#).

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.

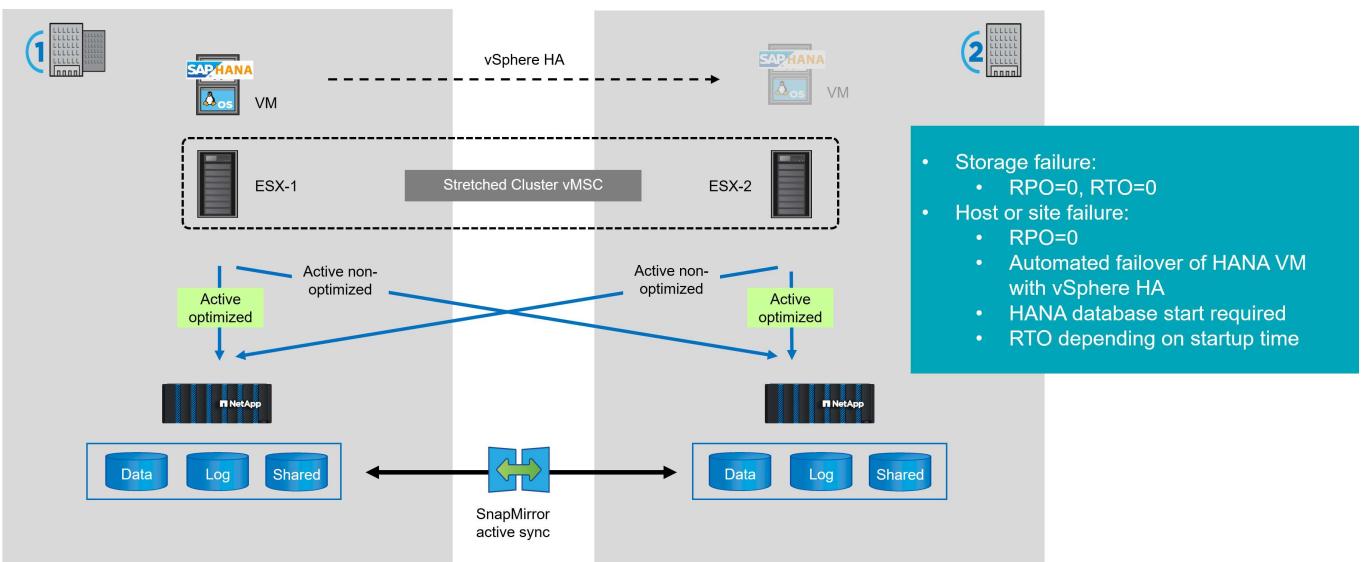


vSphere Metro Storage Cluster (vMSC)

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.



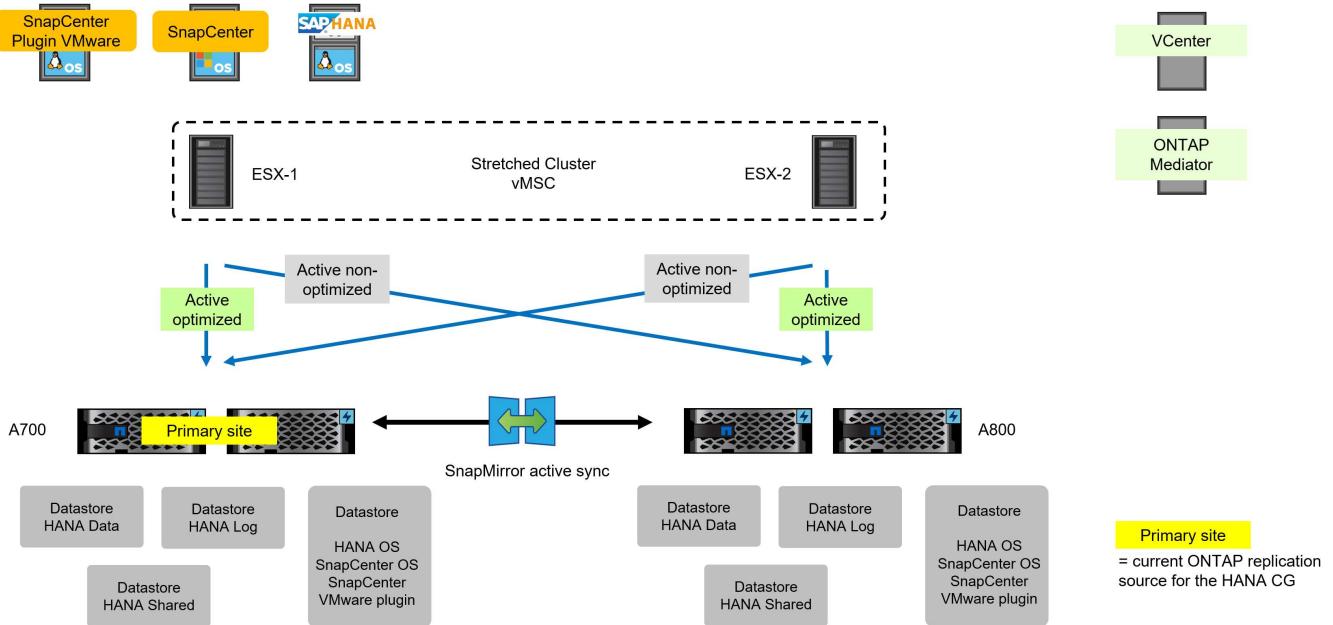
Solution comparison

The following table provides a summary of the key characteristics of the solutions described above.

	HANA System Replication	SnapMirror active sync – bare metal	SnapMirror active sync – VMware vMSC
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

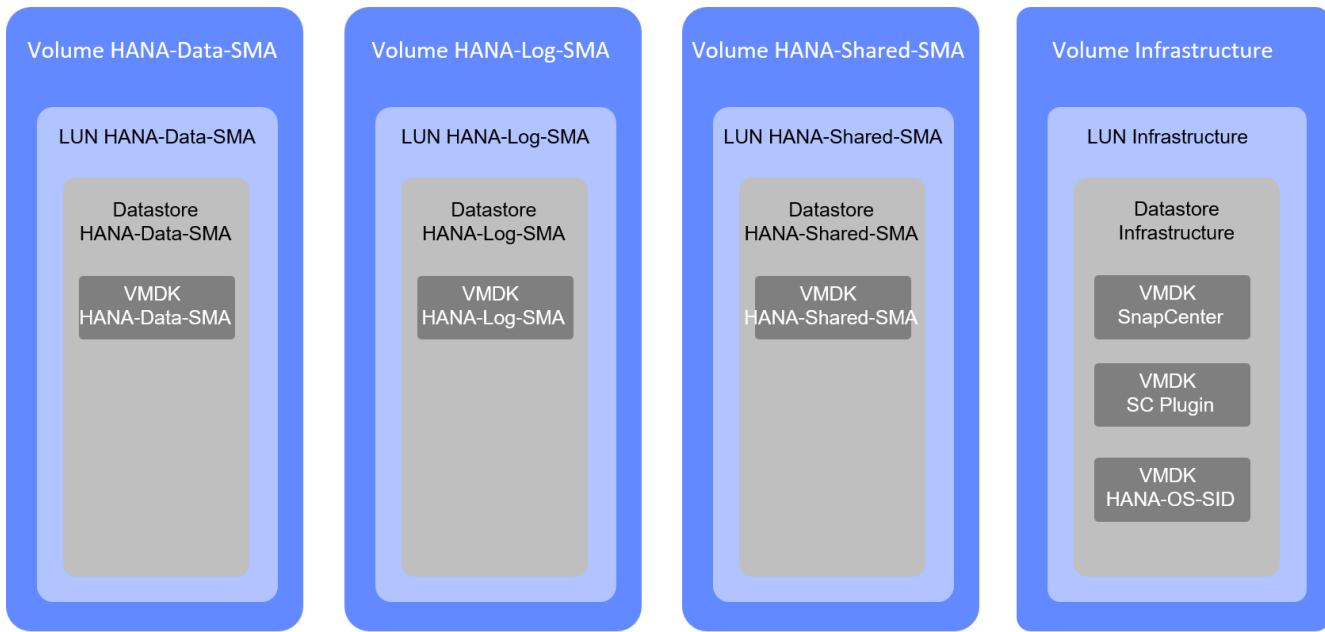
Software	Version
ONTAP	A700: 9.15.1P7, A800: 9.16.1RC1
vSphere client	8.0.3
ESXi	8.0.3
SnapCenter plugin for vSphere	6.0.1
Linux OS	SLES for SAP 15 SP5
SAP HANA	2.0 SPS8
SnapCenter	6.0.1

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_SMA and hana_shared_SMA) as well as the LUN for the OS images and SnapCenter components have been provisioned at the A700 storage cluster.



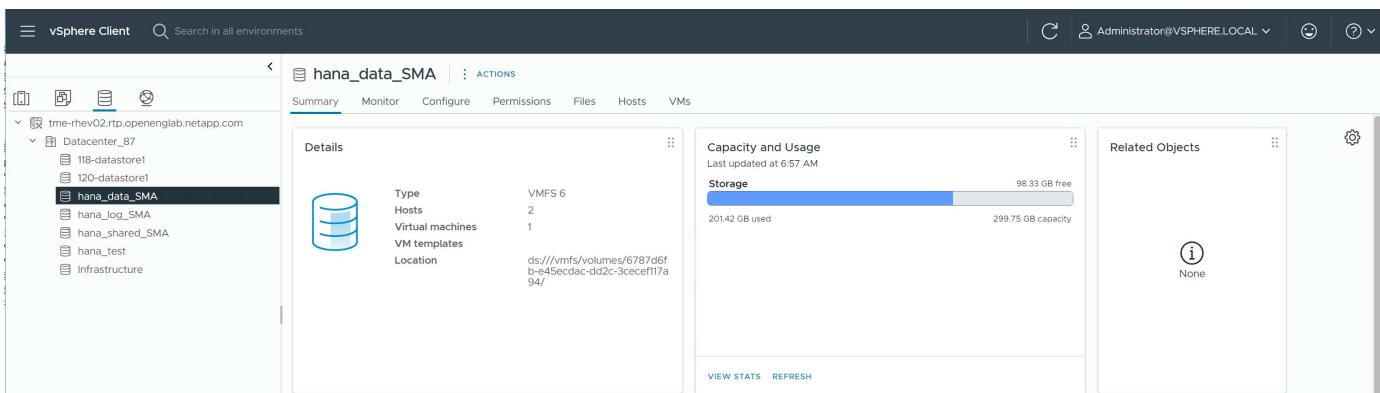
All volumes of the HANA system must be provisioned in the same SVM. In the SnapMirror active sync configuration described later, we will create a consistency group across all three HANA volumes, which requires that the volumes are in the same SVM. The infrastructure volume will be in a different consistency group and could therefore be in a different SVM.

Name	Storage VM	Volume	Size	IOPS	Latency (ms)	Throughput (MB/s)
bluexprdr	svm200_bluexprdr_a700s	vvol_FCoE_2	4 MiB	0	0	0
vvolPE-1724163990635	svm200_bluexprdr_a700s	vvol_FCoE_1	4 MiB	0	0	0
vvolPE-1724163990633	svm200_bluexprdr_a700s	DRAAS_qa_lun1	200 GiB	0	0	0
DRAAS_qa_lun1	svm200_bluexprdr_a700s	DRAAS_qa_lun1	200 GiB	0	0	0
DRAAS_qa_lun2	svm200_bluexprdr_a700s	DRAAS_qa_lun2	100 GiB	0	0	0
Infrastructure	svm200_bluexprdr_a700s	Infrastructure	2 TiB	50	0.31	0.58
hana_data_SMA	svm200_bluexprdr_a700s	hana_data_SMA	300 GiB	0	0.24	0
hana_log_SMA	svm200_bluexprdr_a700s	hana_log_SMA	158 GiB	0	0.24	0
hana_shared_SMA	svm200_bluexprdr_a700s	hana_shared_SMA	210 GiB	1	0.16	0.01
hana_test_lun	svm200_bluexprdr_a700s	hana_test_lun	1 TiB	0	0.39	0

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.

Datastore 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.

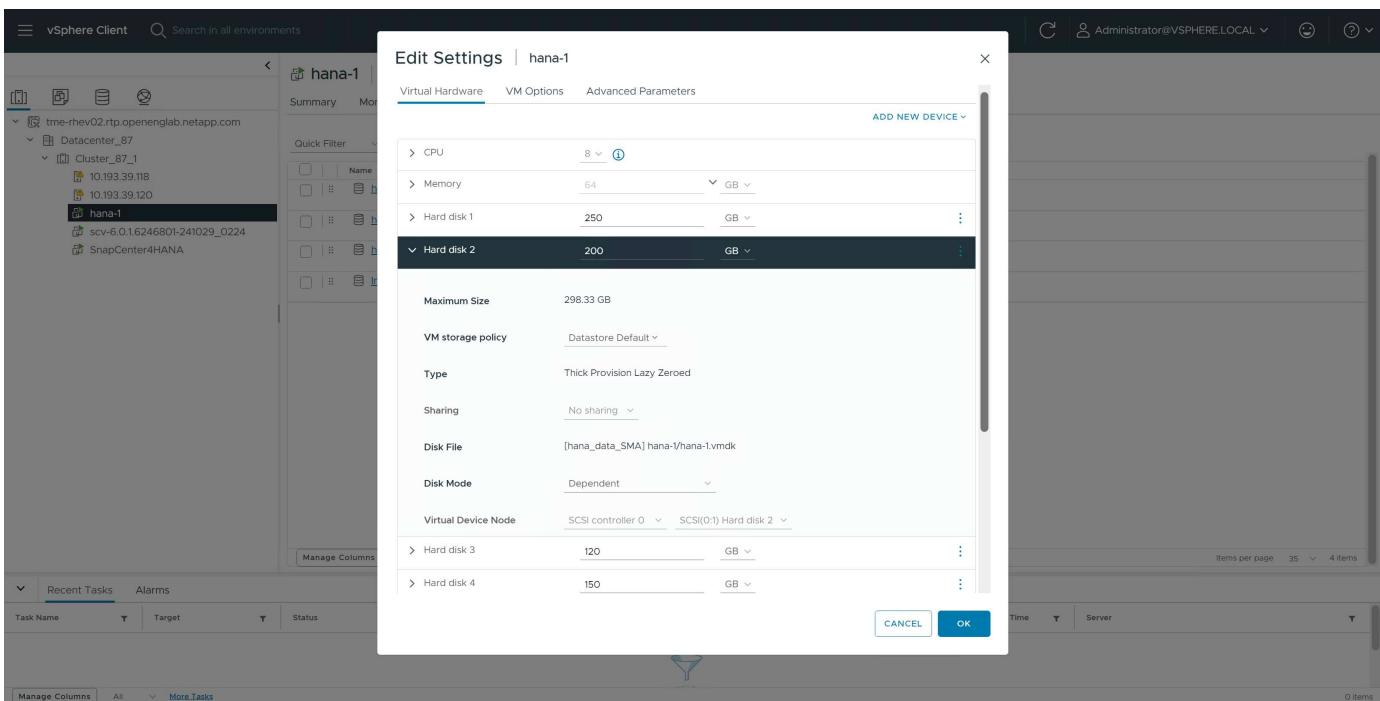


VM provisioning and OS installation

In our lab setup we deployed a new VM and placed the VMDK for the Linux OS in the infrastructure datastore.

VM disk configuration

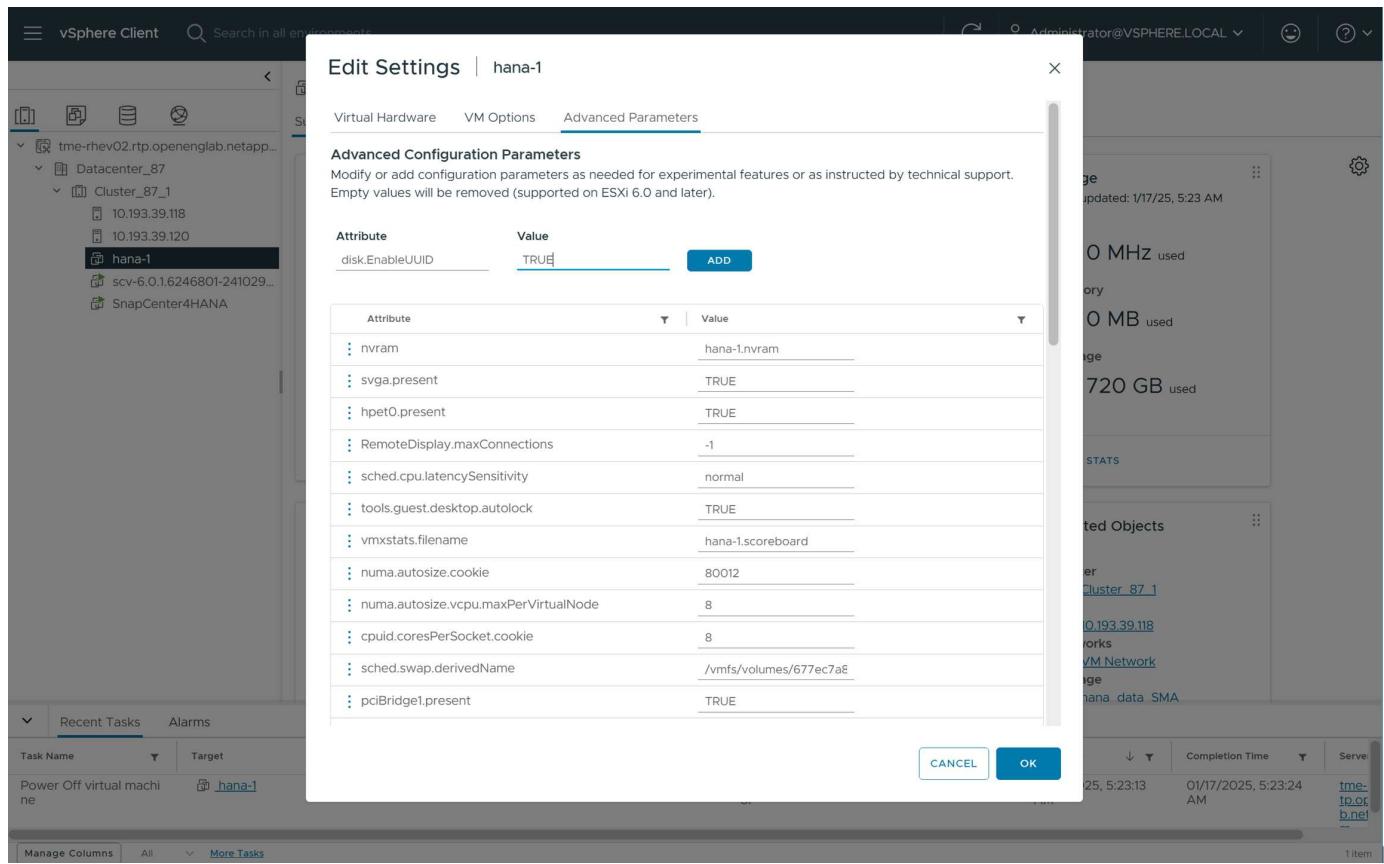
Three new disks have been added to the HANA VM, each disk within one of the datastores which have been created for the HANA system.



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> hdbsql 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.

The screenshot shows the ONTAP System Manager dashboard with the 'PROTECTION' section selected. The 'Protected data' section displays three progress bars: 'Volume protection' (62 volumes, 21 not protected), 'Snapshot copies (local)' (62 volumes, 49 not protected), and 'SnapMirror (local or remote)' (66 volumes, 66 not backed up to cloud). Below these are sections for 'Cluster peers' (Peered cluster name: tme-a800), 'Mediator' (Peered cluster name: 10.61.180.97), and 'Storage VM peers' (Peered storage VMs: 1). On the right, there are three buttons: 'Protect volumes', 'Back up volumes to cloud', and 'Protect for business continuity'. A note for 'Protect volumes' says: 'Lets you select specific volumes for protection if you don't need to protect entire storage VMs.' A note for 'Back up volumes to cloud' says: 'Lets you select which volumes you want to be backed up to a cloud destination.' A note for 'Protect for business continuity' says: 'Lets you protect a consistency group with a zero recovery time objective.' A note at the bottom states: 'NetApp SnapCenter software simplifies backup, restore, and clone management for the applications hosted across ONTAP enabled platforms. [Use NetApp SnapCenter for application-consistent protection](#)'.

The screenshot shows the 'Storage VM peers' section of the ONTAP System Manager. It displays a table with one row, showing a peer storage VM named 'svm200_bluexpdr_a700s' peered with a cluster named 'tme-a800'. The table has columns for 'Storage VM', 'Peered cluster', 'Peered storage VM', 'Status', and 'Applications using this peer' (which is 'SnapMirror').

Storage VM	Peered cluster	Peered storage VM	Status	Applications using this peer
svm200_bluexpdr_a700s	tme-a800	svm200_bluexpdr_a800	Peered	SnapMirror

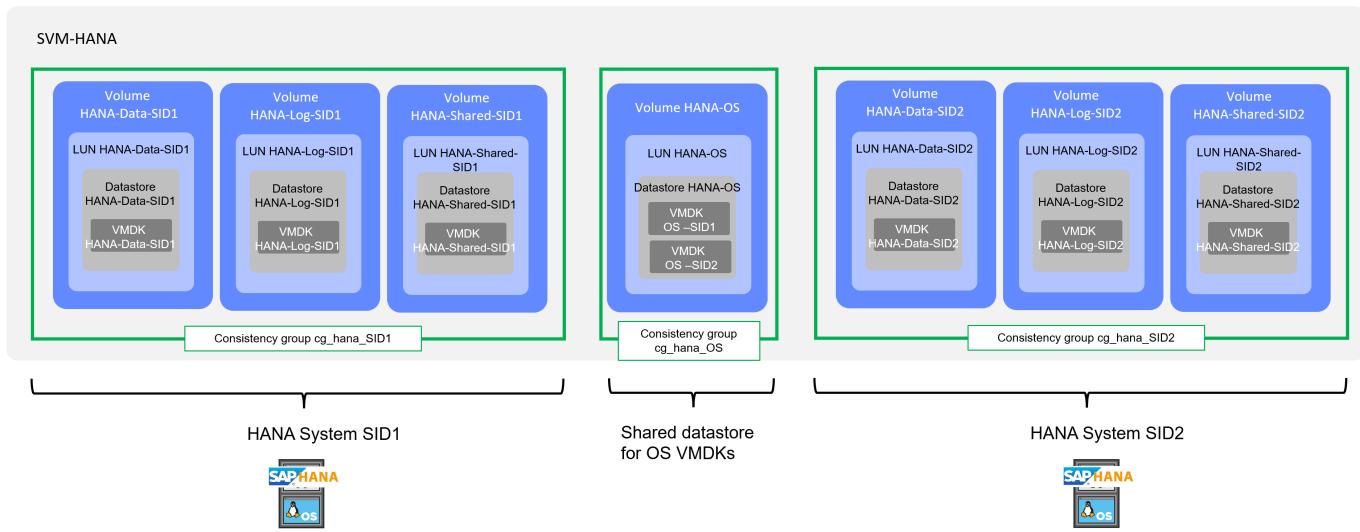
Storage layout and consistency group configuration

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 the cluster_87_1 SAN initiator group. The left sidebar navigation includes DASHBOARD, INSIGHTS, STORAGE (Overview, Volumes, LUNs, NVMe namespaces, Consistency groups, Storage VMs, Tiers), NETWORK (Overview, Ethernet ports, FC ports), EVENTS & JOBS, PROTECTION, and HOSTS (SAN initiator groups, NVMe subsystem, CLUSTER).

The main pane displays the SAN initiator group configuration for cluster_87_1. It shows the following details:

- Storage VM:** svm200_bluepdr_a700s
- Type:** VMware
- Protocol:** Mixed (iSCSI & FC)
- Comment:** -
- Connection Status:** OK
- Replication:** REPLICATED TO SVM: svm200_bluepdr_a800, REPLICATED TO CLUSTER: tme-a800, REPLICATION STATUS: OK
- Initiators:** A table listing initiators with their names, connection status (OK), and proximity to storage SVMs. Two initiators are highlighted with blue boxes and arrows pointing to ESX-1 and ESX-2 respectively:

Name	Description	Connection status	In proximity to
10:00:00:10:9b:17:04:69	-	OK	svm200_bluepdr_a700s
10:00:00:10:9b:17:04:6a	-	OK	svm200_bluepdr_a700s
10:00:00:10:9b:40:b9:7f	-	OK	svm200_bluepdr_a800
10:00:00:10:9b:40:b9:80	-	OK	svm200_bluepdr_a800

cluster_87_1 All SAN initiator groups

cluster_87_1 Nested initiator group

esxi_118

esxi_120

esxi_118

esxi_120

INITIATOR GROUP TYPE: Nested initiator group

REPLICATION: REPLICATED TO SVM: svm200_bluexprdr_a800

REPLICATED TO CLUSTER: tme-a800

REPLICATION STATUS: OK

INITIATORS:

Name	Description	Connection status	In proximity to
10:00:00:10:9b:17...	-	OK	svm200_bluexprdr...
10:00:00:10:9b:17...	-	OK	svm200_bluexprdr...
10:00:00:10:9b:40...	-	OK	svm200_bluexprdr...
10:00:00:10:9b:40...	-	OK	svm200_bluexprdr...

MAPPED LUNS:

Name	ID
hana_data_SMA	1
hana_log_SMA	2
hana_shared_SMA	3
hana_test_lun	4

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

Protected data

Snapshot copies (local)

21 of the 62 volumes aren't protected.

SnapMirror (local or remote)

53 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

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

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 imitator 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.

ONTAP System Manager

Protect Consistency group

PROTECTION POLICY: AutomatedFailOverDuplex

Source: tme-a700s-clus, svm200_bluexdr_z700s

Destination: tme-a800, svm200_bluexdr_a800

CONSISTENCY GROUP: New, NAME: cg_hana_sma

VOLUMES: hana_shared_SMA, hana_data_SMA, hana_log_SMA

Host information: Replicate initiator groups

Save Cancel

By clicking on proximity settings, you can review the configuration done before in the initiator group setup.

ONTAP System Manager

Protect Consistency group

PROTECTION POLICY

Proximity settings

Initiator	Initiator in proximity to
Initiator group: cluster 87.1 Mapped LUNs: 3	
10:00:00:109b:17:04:69	Source
10:00:00:109b:17:04:6a	Source
10:00:00:109b:40:b9:7f	Destination
10:00:00:109b:40:b9:80	Destination

Save Cancel

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.

Source	Destination	Policy type
svm200_bluepdr_a700s:udev1	svm200_bluepdr_a800s:udev1_dest	Asynchronous
svm200_bluepdr_a700s:svvol_nvme_1	svm200_bluepdr_a800:vol_vvol_nvme_1_dest	Asynchronous
svm200_bluepdr_a700s:Draa5_qa_lun1	svm200_bluepdr_a800:Draa5_qa_lun1_dest	Asynchronous
svm200_bluepdr_a700s:svvol_nvme_2	svm200_bluepdr_a800:vol_vvol_nvme_2_dest	Asynchronous
svm200_bluepdr_a700s:NVME_FC_SCV_NEW	svm200_bluepdr_a800:NVME_FC_SCV_NEW_dest	Asynchronous
svm200_bluepdr_a700s:VOLUME_NVME_FC_PLUGIN_PRI	svm200_bluepdr_a800:vol_VOLUME_NVME_FC_PLUGIN_PRI_dest	Asynchronous
svm200_bluepdr_a700s:nvme_vvol_1	svm200_bluepdr_a800:vol_nvme_1_dest	Asynchronous
svm200_bluepdr_a700s/cg/cg_hana_sma	svm200_bluepdr_a800/cg/cg_hana_sma	Synchronous
svm200_bluepdr_a700s/cg/cg_hana_test	svm200_bluepdr_a800/cg/cg_hana_test	Synchronous

At the A800 storage cluster (destination), the new relationship and the status of the replication is listed.

Source	Destination	Protection policy	Relationship health	State	Lag
svm200_bluepdr_a700s:udev1	svm200_bluepdr_a800s:udev1_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_bluepdr_a700s:svvol_nvme_1	svm200_bluepdr_a800:vol_vvol_nvme_1_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_bluepdr_a700s:Draa5_qa_lun1	svm200_bluepdr_a800:Draa5_qa_lun1_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_bluepdr_a700s:svvol_nvme_2	svm200_bluepdr_a800:vol_vvol_nvme_2_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_bluepdr_a700s:NVME_FC_SCV_NEW	svm200_bluepdr_a800:NVME_FC_SCV_NEW_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_bluepdr_a700s:VOLUME_NVME_FC_PLUGIN_PRI	svm200_bluepdr_a800:vol_VOLUME_NVME_FC_PLUGIN_PRI_dest	Asynchronous	Healthy	Mirrored	57 minutes, 7 seconds
svm200_bluepdr_a700s:nvme_vvol_1	svm200_bluepdr_a800:vol_nvme_1_dest	MirrorAndVault	Healthy	Mirrored	141 days, 20 hours, 5 minutes and 4 seconds
svm200_bluepdr_a700s/cg/cg_hana_sma	svm200_bluepdr_a800/cg/cg_hana_sma	AutomatedFailOverDuplex	Healthy	Synchronizing	-
svm200_bluepdr_a700s/cg/cg_hana_test	svm200_bluepdr_a800/cg/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.

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:vvol_VOLUME_NVME_FC_PLUGIN_PRI_dest	Asynchronous
svm200_bluexpdr_a700s:nvme_vvol_1	svm200_bluexpdr_a800:vvol_nvme_1_dest	Asynchronous
svm200_bluexpdr_a700s:nvme_vvol_2	svm200_bluexpdr_a800:vvol_nvme_2_dest	Asynchronous
svm200_bluexpdr_a700s:cg/cg_hana_sma	svm200_bluexpdr_a800:cg/cg_hana_sma	Synchronous

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:

[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">
```

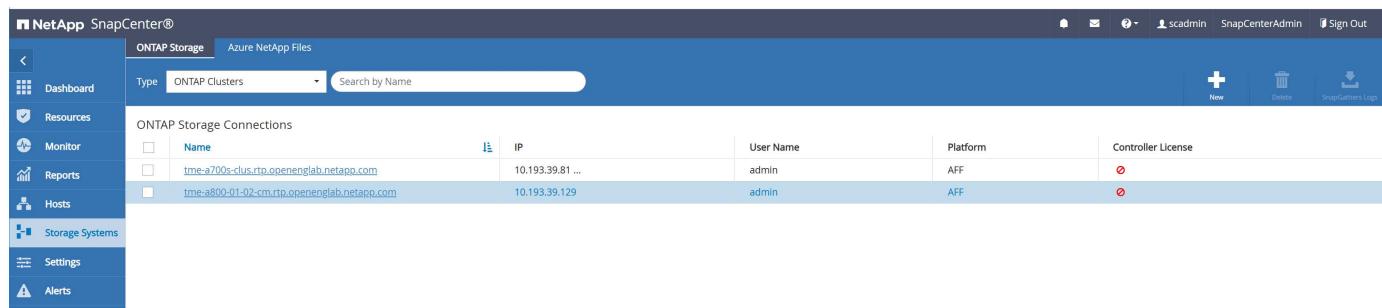
	Name	Date modified	Type	Size
Quick access	App.config	11/8/2024 5:34 AM	CONFIG File	6 KB
Desktop	SMCoreServiceHost.dll.config	1/17/2025 1:59 AM	CONFIG File	5 KB
Downloads	SmJobLauncher.dll.config	11/8/2024 5:35 AM	CONFIG File	1 KB
Documents	SmJobsHelper.dll.config	11/8/2024 5:36 AM	CONFIG File	1 KB
Pictures	SnapDrive.Nsf.Common.Configuration.dll.config	11/8/2024 5:35 AM	CONFIG File	3 KB
This PC	SnapDrive.Nsf.Common.Infrastructure.dll.config	11/8/2024 5:35 AM	CONFIG File	1 KB
3D Objects	SnapDrive.Nsf.Common.Logging.dll.config	11/8/2024 5:35 AM	CONFIG File	3 KB
Desktop	SnapDrive.Nsf.Common.PowerShell.dll.config	11/8/2024 5:35 AM	CONFIG File	3 KB
Documents	SnapDrive.Nsf.Core.dll.config	11/8/2024 5:35 AM	CONFIG File	3 KB
Downloads	SnapDrive.Nsf.Core.Storage.dll.config	11/8/2024 5:35 AM	CONFIG File	3 KB
Music	SnapDrive.Nsf.ServiceProviders.dll.config	11/8/2024 5:35 AM	CONFIG File	3 KB
Pictures	VMWareProxy.dll.config	11/8/2024 5:35 AM	CONFIG File	1 KB
Videos	web.config	11/8/2024 5:44 AM	CONFIG File	1 KB

After the parameter change the SnapCenter SMCore Service must be restarted.

Services (Local)						
	Name	Description	Status	Startup Type	Log On As	
Stop the service	SnapCenter SMCore Service	The startup ...	Running	Automatic	Local System	
Restart the service	Security Accounts Manager	Delivers dat...	Disabled	Local System		
	Sensor Data Service	Monitors va...	Manual (Trigg...	Local Service		
	Sensor Monitoring Service	A service for ...	Manual (Trigg...	Local System		
	Sensor Service	Supports file...	Running	Automatic (Tri...	Local System	
	Server	Manages pr...	Disabled	Local System		
	Shared PC Account Manager	Provides not...	Running	Automatic	Local System	
	Shell Hardware Detection	Creates soft...	Manual (Trigg...	Local Service		
	Smart Card	Allows this s...	Disabled	Local System		
	Smart Card Device Enumerator	Provides ena...	Manual	Local System		
	Smart Card Removal Policy	Manages ac...	Manual	Local System		
	SnapCenter SMCore Service	Receives tra...	Running	Automatic	Local System	
	SNMP Trap	Enables the ...	Manual	Local Service		
	Software Protection	Allows adm...	Automatic (De...	Network Se...		
	Special Administration Cons...	Provides ena...	Manual	Local System		
	Spot Verifier	Provides ena...	Manual (Trigg...	Local System		
	SSDP Discovery	Provides ena...	Disabled	Local Service		
	State Repository Service	Provides req...	Running	Automatic	Local System	
	Still Image Acquisition Events	Launches ap...	Manual	Local System		
	Storage Service	Provides ena...	Running	Automatic (De...	Local System	
	Storage Tiers Management	Optimizes th...	Manual	Local System		
	SysMain	Maintains a...	Running	Automatic	Local System	
	System Event Notification S...	Monitors sy...	Running	Automatic	Local System	
	System Events Broker	Coordinates u...	Running	Automatic (Tri...	Local System	
	System Guard Runtime Mon...	Monitors an...	Manual (Trigg...	Local System		
	Task Scheduler	Enables a us...	Running	Automatic	Local System	
	TCP/IP NetBIOS Helper	Provides sup...	Running	Manual (Trigg...	Local Service	
	Telephony	Provides Tel...	Manual	Network Se...		
	Themes	Provides use...	Running	Automatic	Local System	
	Time Broker	Coordinates ...	Running	Manual (Trigg...	Local Service	
	Touch Keyboard and Handw...	Enables Tou...	Running	Automatic (Tri...	Local System	

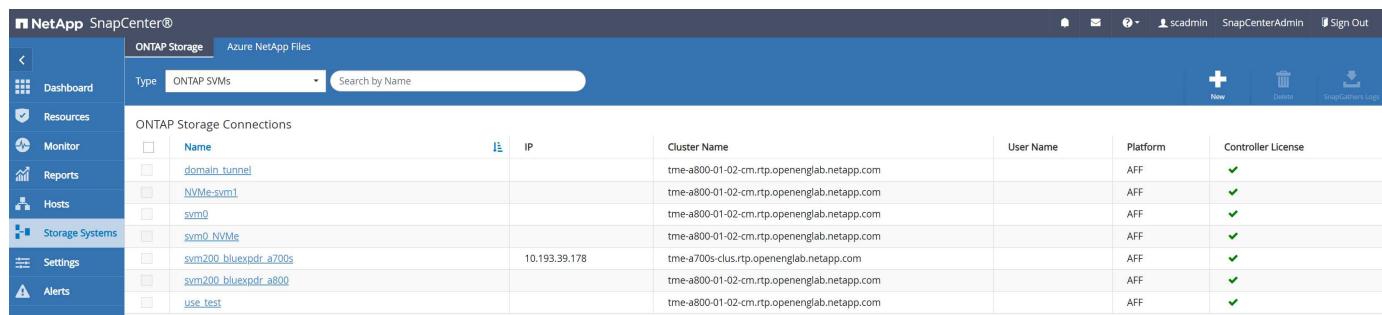
Add storage systems

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 ONTAP Storage section of the NetApp SnapCenter interface. The left sidebar includes options for Dashboard, Resources, Monitor, Reports, Hosts, Storage Systems, Settings, and Alerts. The main content area is titled 'ONTAP Storage' and shows 'ONTAP Storage Connections'. A dropdown menu indicates the type is 'ONTAP Clusters'. A search bar is present. The table lists three entries:

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 ONTAP Storage section of the NetApp SnapCenter interface. The left sidebar includes options for Dashboard, Resources, Monitor, Reports, Hosts, Storage Systems, Settings, and Alerts. The main content area is titled 'ONTAP Storage' and shows 'ONTAP Storage Connections'. A dropdown menu indicates the type is 'ONTAP SVMs'. A search bar is present. The table lists multiple SVMs:

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_bluepdr_a700s	10.193.39.178	tme-a700s-clus.rtp.openenglab.netapp.com		AFF	✓
svm200_bluepdr_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 in the NetApp SnapCenter interface. The left sidebar includes options for Managed Hosts, Resources, Monitor, Reports, Hosts, Storage Systems, Settings, and Alerts. The main content area is titled 'Add Host' and shows fields for 'Host Type' (set to 'vSphere'), 'Host Name' (set to 'Host Name or IP'), and 'Credentials' (set to 'None'). A note 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.' The 'Submit' and 'Cancel' buttons are at the bottom.



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
```

Resource - Details

Details for selected resource

Type	Multitenant Database Container		
HANA System Name	SMA		
SID	SMA		
Tenant Databases	SMA		
Plug-in Host	hana-1		
HDB Secure User Store Key	SMKEY		
HDBSQL OS User	smaadmin		
Log backup location	/usr/sap/SMA/HDB00/backup/log		
Backup catalog location	/usr/sap/SMA/HDB00/backup/log		
System Replication	None		
Plug-in name	SAP HANA		
Last backup	01/29/2025 3:14:18 AM (Completed)		
Resource Groups	hana-1_hana_MDC_SMA		
Policy	SM-AS-Policy		
Discovery Type	Auto		
Storage Footprint			
SVM	Volume	Junction Path	LUN/Qtree
10.193.39.178	hana_data_SMA		hana_data_SMA

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.

Modify SAP HANA Backup Policy

1 Name

2 Policy type

3 Snapshot

4 Replication and backup

5 Summary

Select secondary replication options

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label: Hourly

Error retry count: 3

Global Settings Policies Users and Access Roles Credential Software

SAP HANA

Name Scope Schedule Type Snapshot Backup Replication

SM-AS-Policy Data Backup Hourly Copies to keep : 7 copies SnapMirror

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.

Job Details

Backup of Resource Group 'hana-1_hana_MDC_SMA' with policy 'SM-AS-Policy'

- ✓ Backup of Resource Group 'hana-1_hana_MDC_SMA' with policy 'SM-AS-Policy'
- ✓ hana-1
 - ✓ 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

Task Name: Backup Start Time: 01/29/2025 3:13:07 AM End Time: 01/29/2025 3:14:51 AM

View Logs Cancel Job Close

Job Details

Backup of Resource Group 'hana-1_hana_MDC_SMA' with policy 'SM-AS-Policy'

- ✓ 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
- ✓ Secondary Update
- ✓ Register Backup and Apply Retention
- ✓ Register Snapshot attributes
- ✓ Application Clean-Up
- ✓ Data Collection
- ✓ Agent Finalize Workflow
- ✓ Replicating to Secondary
- ✓ Backup of Resource Group 'hana-1_hana_MDC_SMA' with policy 'SM-AS-Policy'
- ✓ (Job 250) Monitor snapshot in secondary for SCC_BACKUP

Task Name: Backup Start Time: 01/29/2025 3:13:07 AM End Time: 01/29/2025 3:14:51 AM

View Logs Cancel Job Close

Topology view

Within the topology view, SnapCenter shows the backups of both source and destination storage clusters.

NetApp SnapCenter®

SAP HANA

SMA Topology

Manage Copies

Local copies: 3 Backups, 0 Clones

Mirror copies: 3 Backups, 0 Clones

Primary Backup(s)

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

Summary Card

6 Backups

6 Snapshot based backups

0 File-Based backups

0 Clones

0 Snapshots Locked

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.

Details

Source	Relation	Destination	BackupCount
svm200_bluexprdr_a700s:hana_data_SMA	Mirror	svm200_bluexprdr_a800:vol_hana_data_SMA_dest	1

Total 1

Close

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

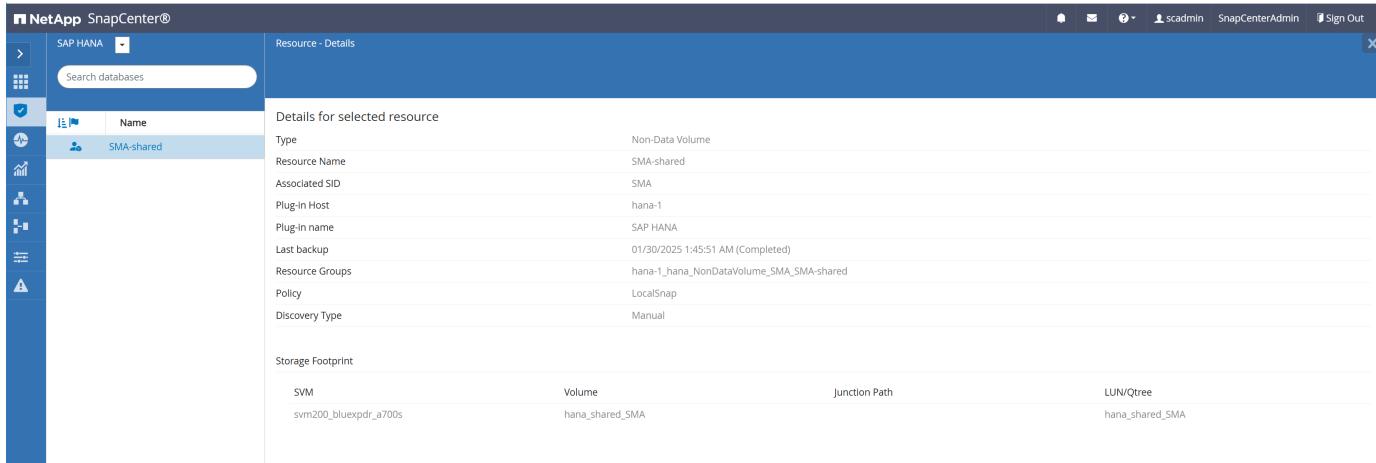
When resources are configured manually in SnapCenter and are not auto discovered, SnapCenter is not aware of VMware and SnapMirror active sync. Therefore, they are not supported natively by SnapCenter.

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.



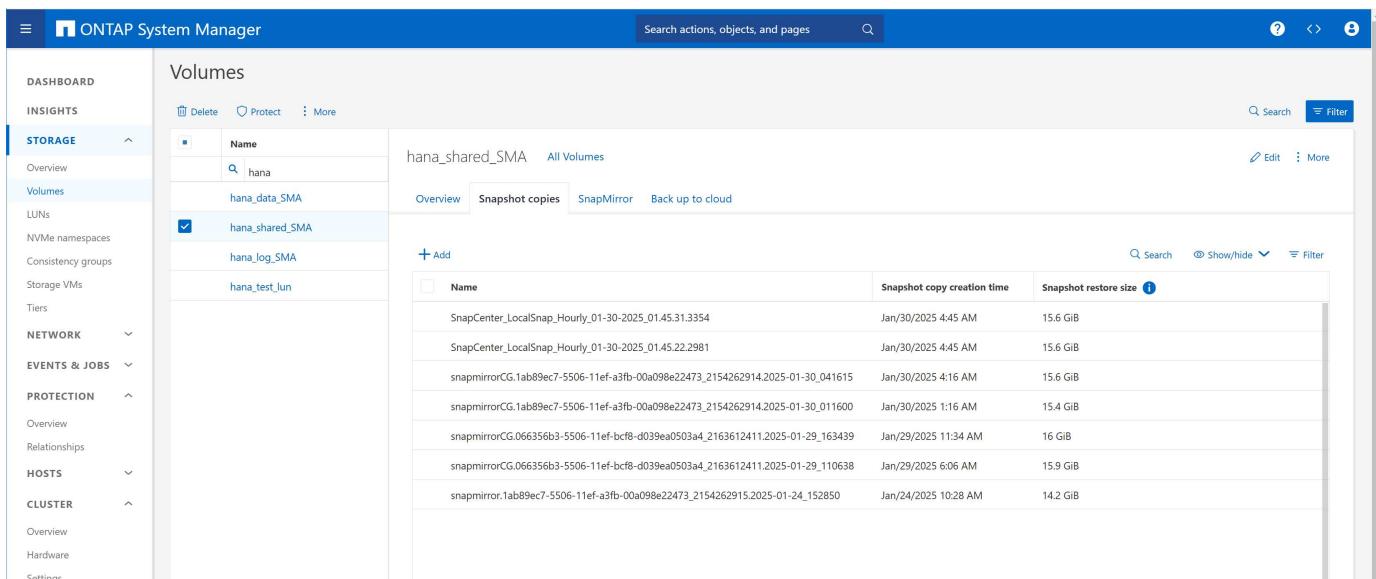
The screenshot shows the NetApp SnapCenter interface. The left sidebar has a 'SAP HANA' dropdown. The main area is titled 'Resource - Details' for a selected resource named 'SMA-shared'. The resource is a 'Non-Data Volume' with the following details:

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

Under 'Storage Footprint', it shows the SVM 'svm200_blueprd_a700s' and the Volume 'hana_shared_SMA' with a Junction Path of 'hana_shared_SMA'.

Backup operations

Even though SnapCenter is not aware of the SnapMirror active sync configuration for the HANA shared volume, Snapshot are replicated to both sites.



The screenshot shows the ONTAP System Manager interface. The left sidebar has a 'STORAGE' section with 'Volumes' selected. The main area shows a list of volumes including 'hana', 'hana_data_SMA', 'hana_shared_SMA' (which is selected), 'hana_log.SMA', and 'hana_test_lun'. The 'hana_shared_SMA' volume is shown in more detail with its 'Snapshot copies' tab selected. It lists several snapshots:

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

The screenshot shows the NetApp ONTAP System Manager interface. The left sidebar is titled 'Storage' and includes sections for Overview, Volumes, LUNs, NVMe namespaces, Consistency groups, Shares, Qtrees, Quotas, Storage VMs, Tiers, Network, Events & jobs, Protection, Hosts, and Cluster. The main content area is for a volume named 'vol_hana_share...'. The 'Schemas' tab is selected, showing a table with columns: Name, Snapshot creation time, and Snapshot restore size. The table lists several snapshots with their respective details.

Name	Snapshot creation time	Snapshot restore size
SnapCenter_LocalSnap_Hourly_01-30-2025_0145313354	Jan/30/2025 4:45 AM	16.2 GiB
SnapCenter_LocalSnap_Hourly_01-30-2025_0145222981	Jan/30/2025 4:45 AM	16.2 GiB
snapshotmirrorCG.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262914.2025-01-30_041615	Jan/30/2025 4:16 AM	16.1 GiB
snapshotmirrorCG.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262914.2025-01-30_011600	Jan/30/2025 1:16 AM	16 GiB
snapshotmirrorCG.066356b3-5506-11ef-bcf8-d039ea0503a4_2163612411.2025-01-29_163439	Jan/29/2025 11:34 AM	15.7 GiB
snapshotmirrorCG.066356b3-5506-11ef-bcf8-d039ea0503a4_2163612411.2025-01-29_110638	Jan/29/2025 6:06 AM	15.8 GiB
snapshotmirror.1ab89ec7-5506-11ef-a3fb-00a098e22473_2154262915.2025-01-24_152850	Jan/24/2025 10:28 AM	14.1 GiB
snapshotmirror.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

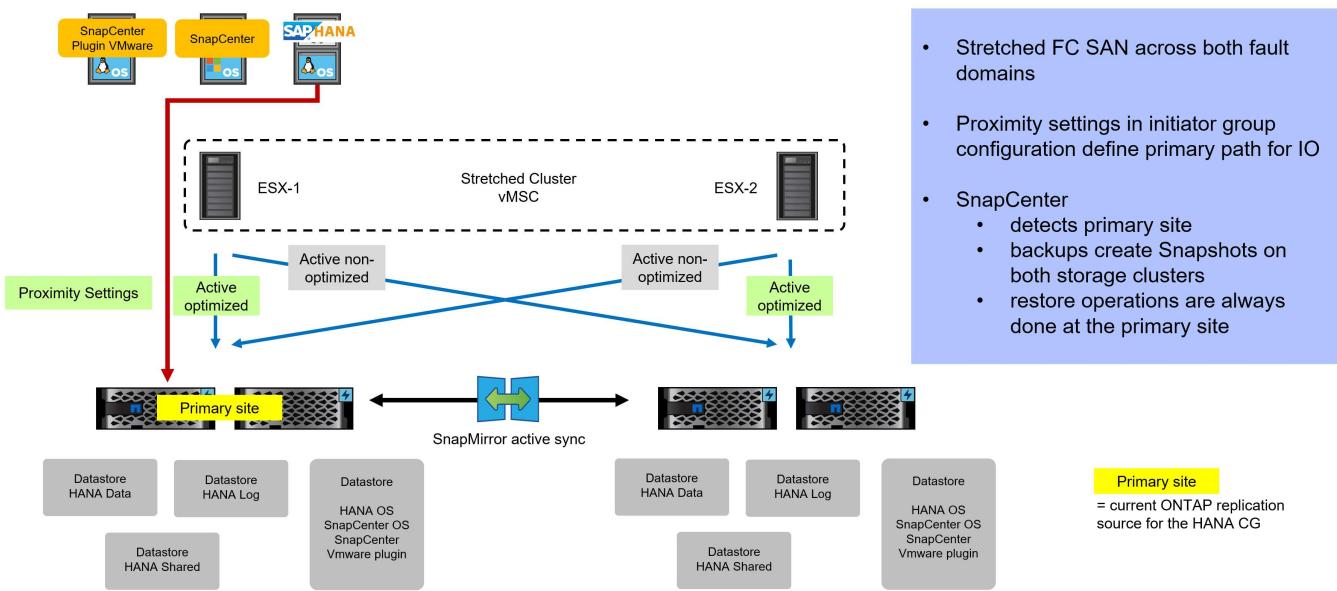
This article will highlight the failover scenarios for this solution.

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.

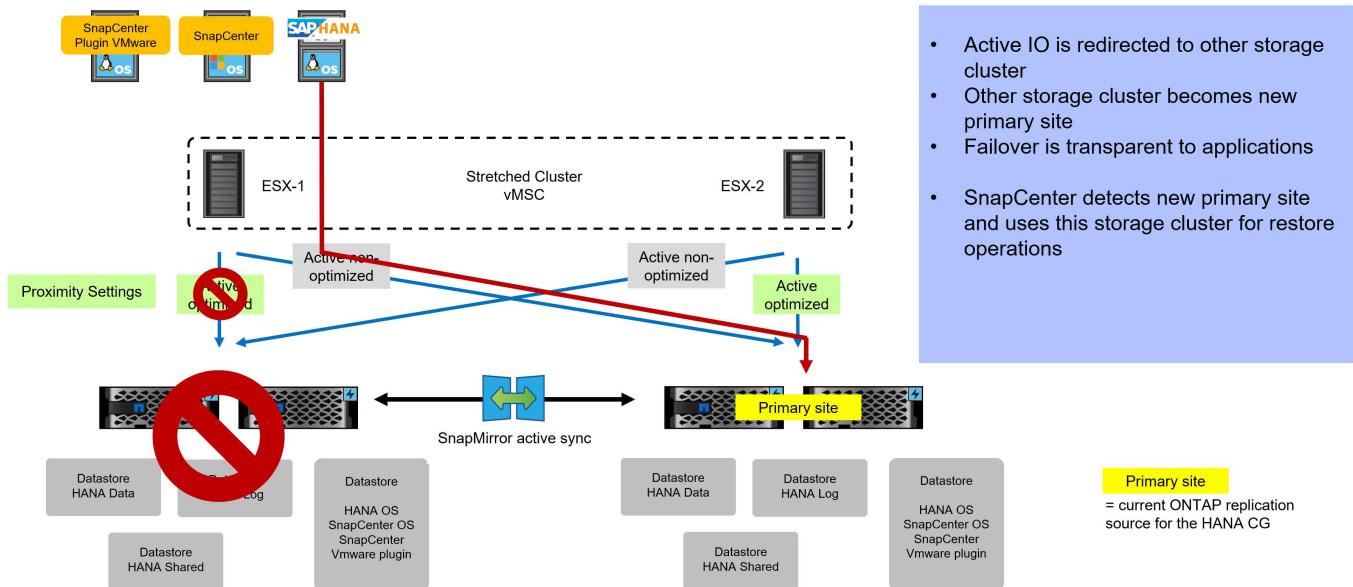
Normal operation

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.



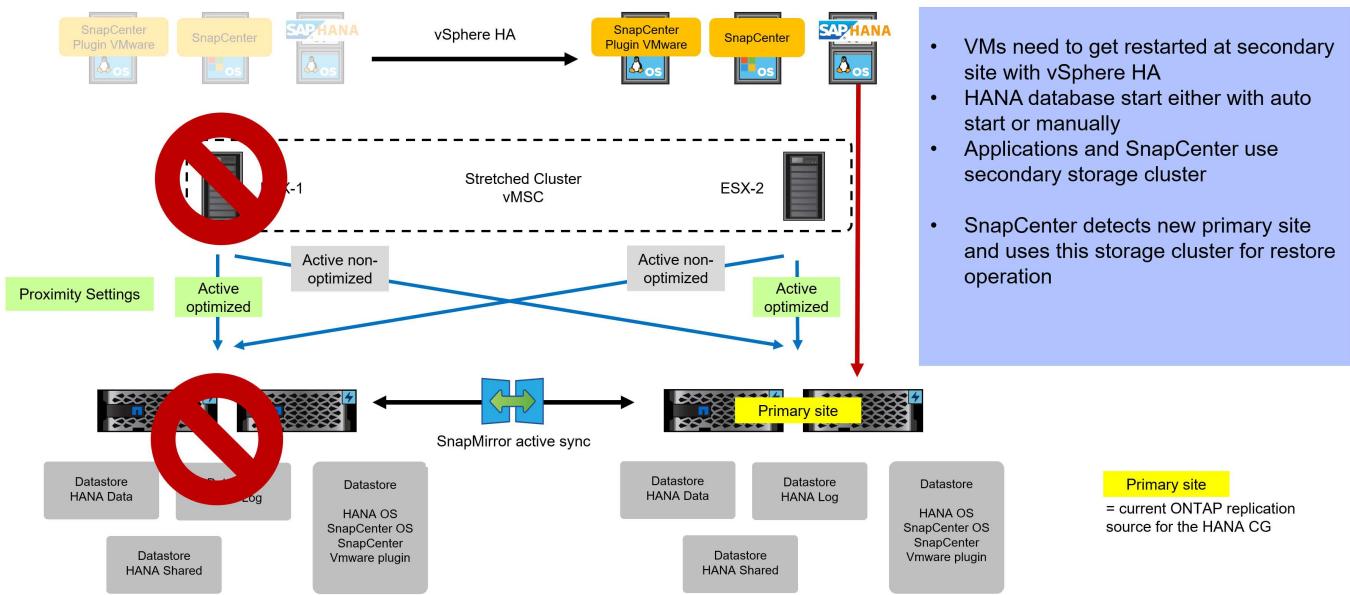
Storage failure

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.



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.

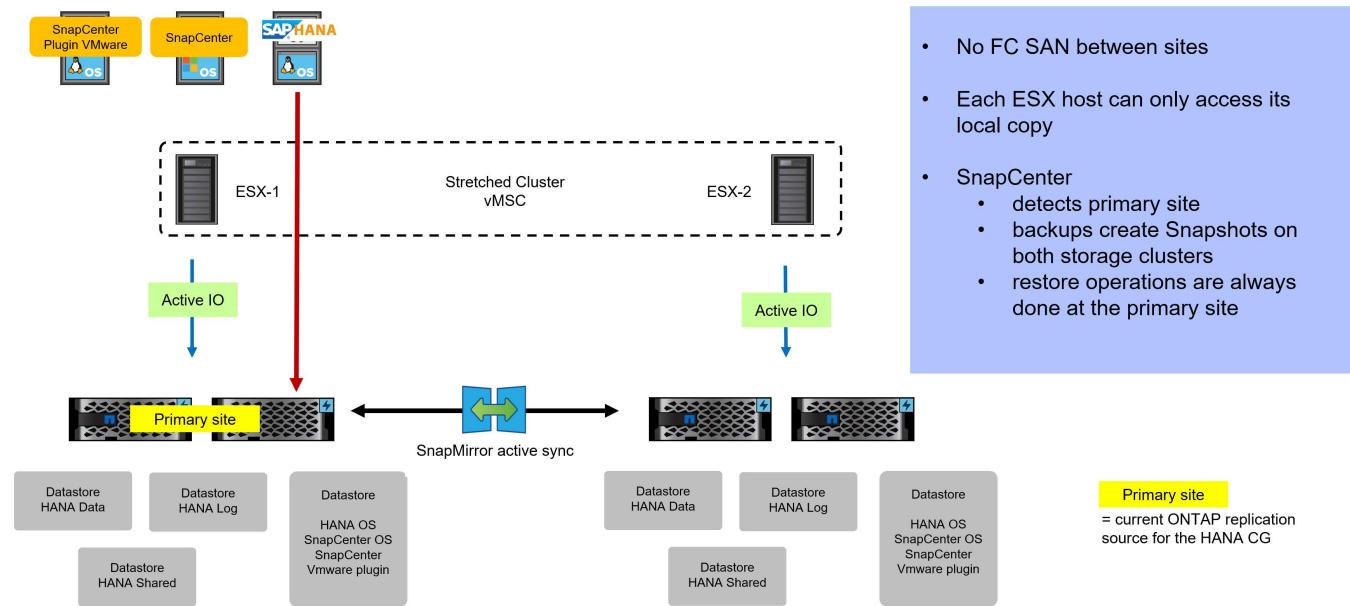


Non-uniform access setup

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

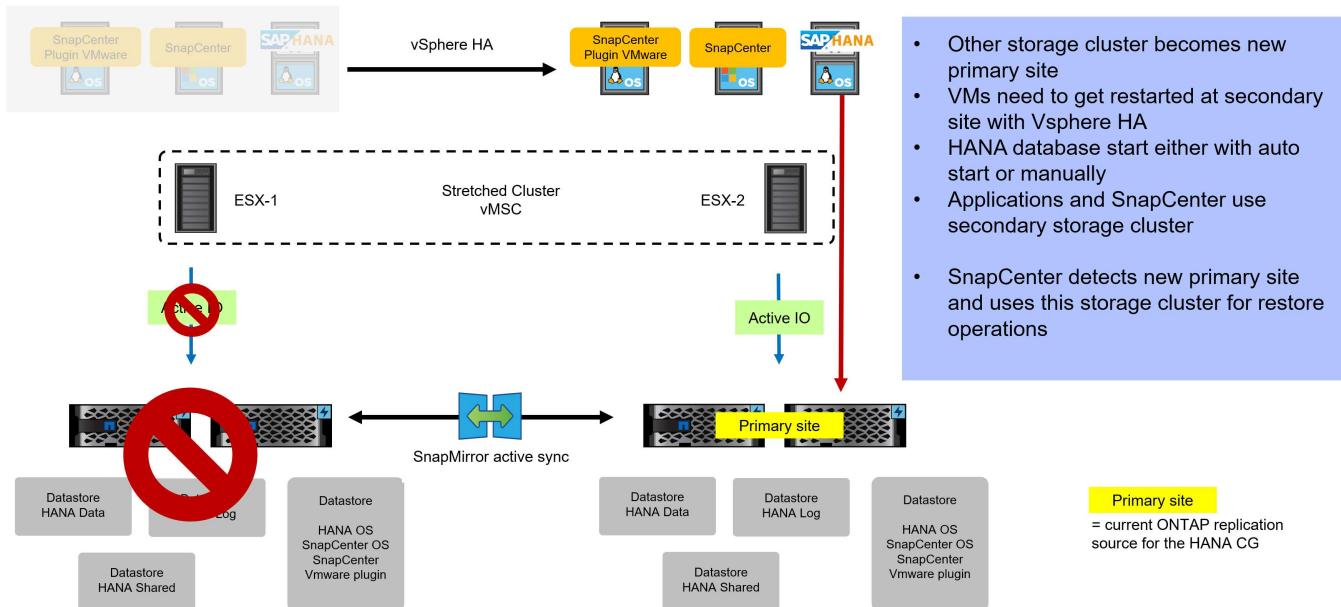
During normal operation the HANA system reads and writes from/to the local copy. With each backup operation, SnapCenter detects the current primary site for the replication relationship and executes the backup operations against the primary site. The Snapshots are replicated to the mirrored copy and are available at both sites. A SnapCenter restore operation would be executed at the primary site.



Storage failure

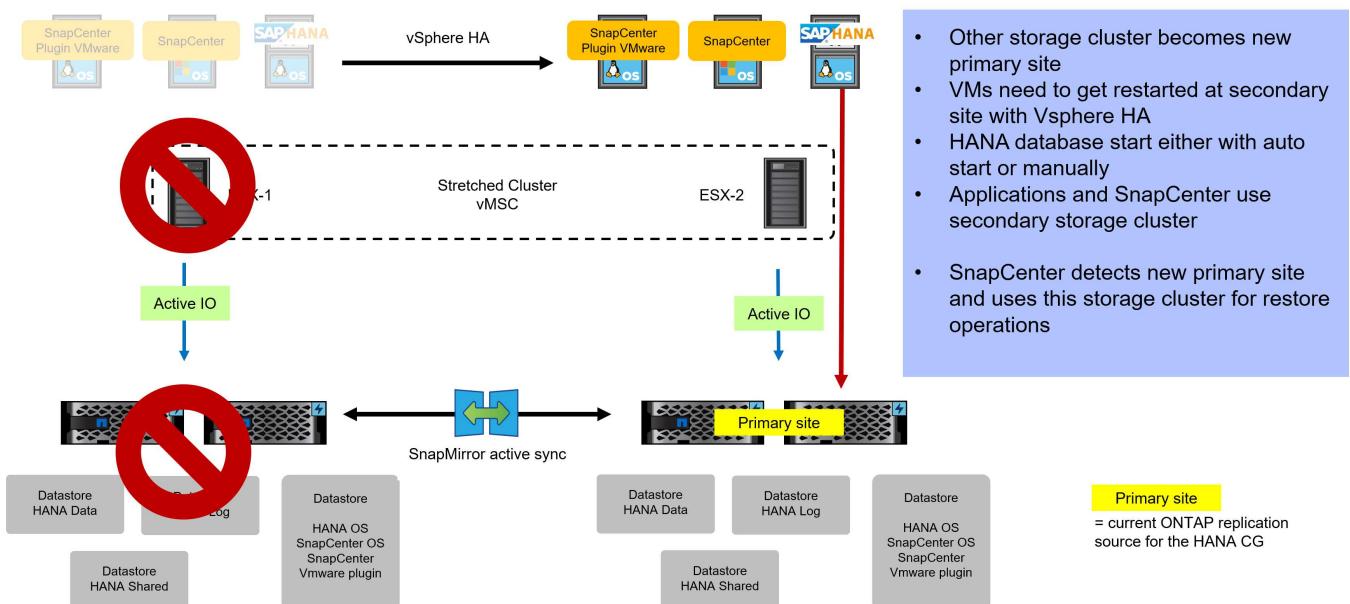
In case of a storage failure, the HANA VM as well as SnapCenter and the SnapCenter for VMware plugin VM will fail over to the ESX host at the secondary site using vSphere HA. The HANA database needs to get started and will then access the mirrored copy at the second site. The primary site switches to the secondary site and

SnapCenter now executes backup and restore operations at the new primary site.



Site failure

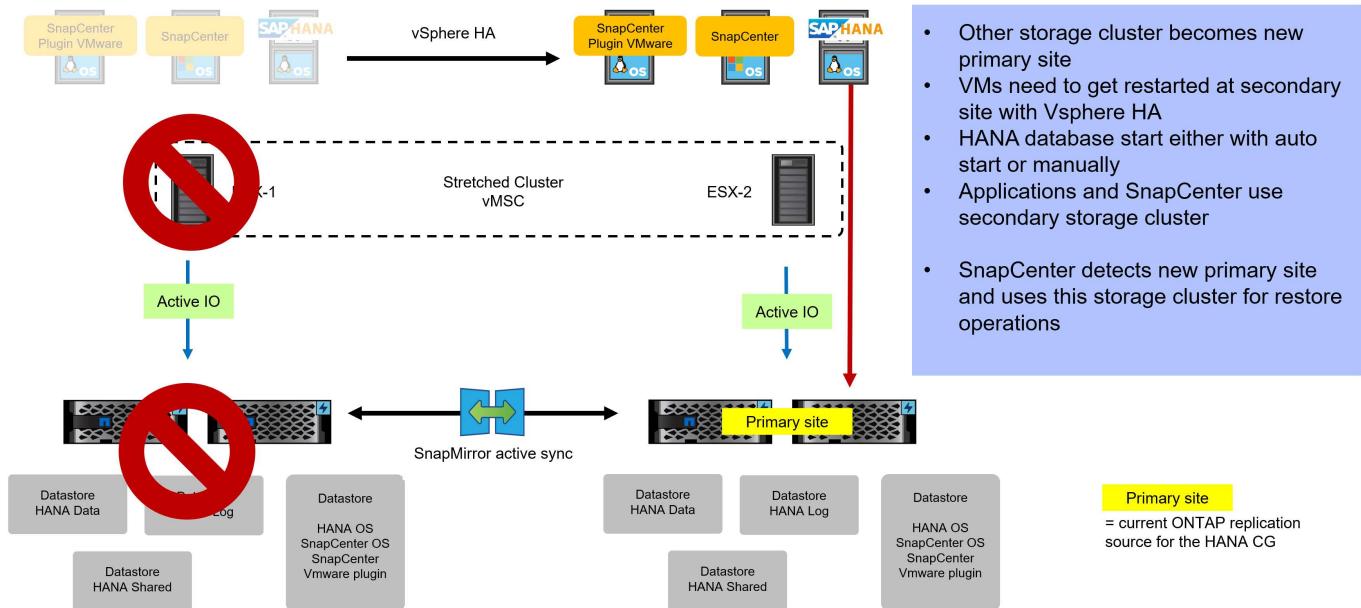
Same as storage failure.



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 able to copy the data.

As a workaround you need to make sure, that the relocation of VM and primary site 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:

[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

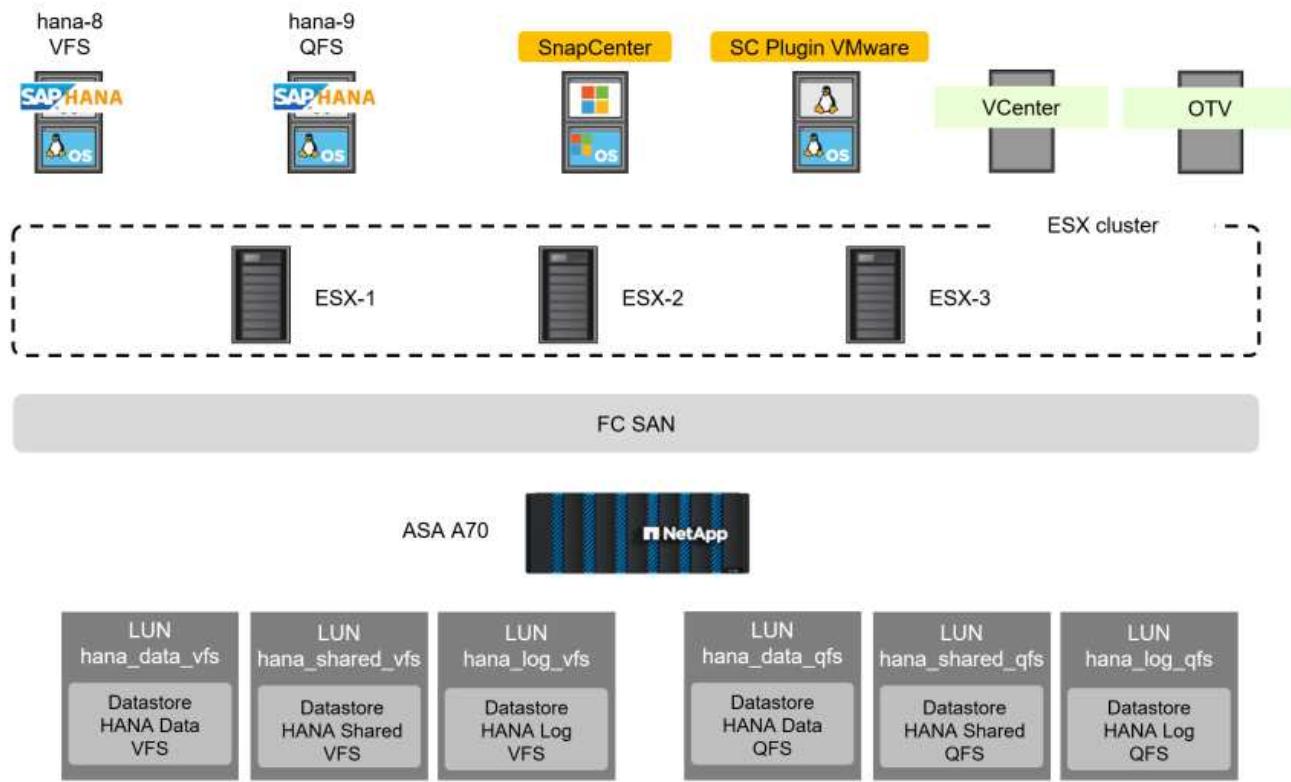
It does not serve as a step-by-step guide for configuring the entire environment but focuses on details specific to SnapCenter and HANA on VMFS, including:

- 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 the SAPCC datacenter. The left sidebar lists datastores: datastore1, datastore1 (1), datastore1 (2), Datastore_C250, Datastore_One, DatastoreA400, hana_data_VFS, hana_log_VFS, hana_shared_VFS, and OS_Image. The main pane displays 'Datacenter Details' with statistics: 3 hosts, 38 virtual machines, 1 cluster, 8 networks, and 10 datastores. It also shows 'Capacity and Usage' for CPU, Memory, and Storage. The 'Custom Attributes' section indicates 'No custom attributes assigned'. Below this is a 'Recent Tasks' table showing three completed tasks: 'Process VMFS datastore e updates' (target 10.63.167.6, initiator System, queued 4 ms, start 05/19/2025 9:20:23, completion 05/19/2025 9:20:23, server vcenter8.sapcc.sti.netapp.com), 'Process VMFS datastore e updates' (target 10.63.167.4, initiator System, queued 5 ms, start 05/19/2025 9:20:23, completion 05/19/2025 9:20:23, server vcenter8.sapcc.sti.netapp.com), and 'Create VMFS datastore' (target 10.63.167.14, initiator SAPCC.VCENTER\Administrator, queued 10 ms, start 05/19/2025 9:20:22, completion 05/19/2025 9:20:23, server vcenter8.sapcc.sti.netapp.com).

Task Name	Target	Status	Initiator	Queued For	Start Time	Completion Time	Server
Process VMFS datastore e updates	10.63.167.6	Completed	System	4 ms	05/19/2025, 9:20:23 AM	05/19/2025, 9:20:23 AM	vcenter8.sapcc.sti.netapp.com
Process VMFS datastore e updates	10.63.167.4	Completed	System	5 ms	05/19/2025, 9:20:23 AM	05/19/2025, 9:20:23 AM	vcenter8.sapcc.sti.netapp.com
Create VMFS datastore	10.63.167.14	Completed	SAPCC.VCENTER\Administrator	10 ms	05/19/2025, 9:20:22	05/19/2025, 9:20:23	vcenter8.sapcc.sti.netapp.com

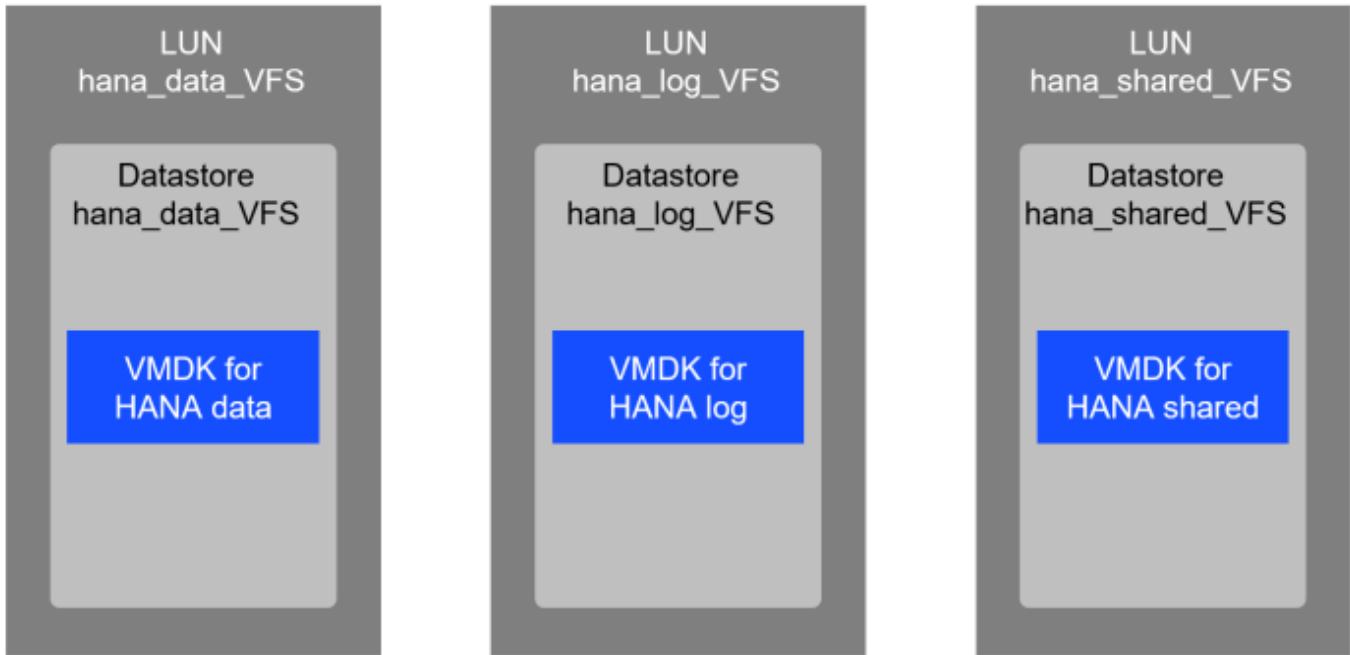
At the storage system three LUNs have been created by OTV.

The screenshot shows the NetApp ONTAP System Manager interface. The left sidebar includes 'Storage' (selected), 'Hosts', 'Network', 'Events & Jobs', 'Protection', and 'Cluster'. The main pane shows 'Storage' details: 19 Storage units, 68.6 TiB Available, and 19 Online units. Below this is a table of LUNs:

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_e3d7e9d4-46f3-4fd1	0	0	0
hana_log_VFS	-	100 GiB	8.69 to 1	otv_host-44_e3d7e9d4-46f3-4fd1	0	0	0
hana_shared_VFS	-	100 GiB	3.13 to 1	otv_host-44_e3d7e9d4-46f3-4fd1	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.



vsphere Client Search in all environments

Administrator@SAPCC.VCENTER

hana-8

Summary Monitor Configuration

Guest OS

Maximum Size: 1.4 TB

VM storage policy: Datastore Default

Type: Thin Provision

Sharing: No sharing

Disk File: [Datastore_One] asa_hana01/asa_hana01.vmdk

Disk Mode: Dependent

Virtual Device Node: SCSI controller 0 SCSI(0:0) Hard disk 1

New Hard disk: 95 GB

Maximum Size: 98.34 GB

VM storage policy: Datastore Default

Location: hana_data_VFS

Disk Provisioning: Thin Provision

Sharing: No sharing

Disk Mode: Dependent

Virtual Device Node: SCSI controller 0 SCSI(0:1) Hard disk 2

PCI Devices: No PCI devices

Related Objects

- Cluster
- Host
- Resource pool
- ASA2
- Networks
- PortGroup-NE5
- ExtAccess
- Storage
- Datastore_One

Storage Policies

VM Storage Policies

VM Storage Policy Compliance

Last Checked Date

VM Replication Groups

Recent Tasks

Task Name	Target	Status	Details
Power On virtual machine	hana-8	Completed	SAPCC.VCENTER/Administrat 7 ms 05/19/2025, 9:41:32 AM
Initialize powering On	SAPCC	Completed	SAPCC.VCENTER/Administrat 7 ms 05/19/2025, 9:41:22
Reconfiguring Virtual Mac	SAPCC.VCENTER/Administrat 7 ms 05/19/2025, 9:41:23	Completed	SAPCC.VCENTER/Administrat 7 ms 05/19/2025, 9:41:23
Rename virtual machine	asa_hana01	Completed	SAPCC.VCENTER/Administrat 7 ms 05/19/2025, 9:41:23

Cancel OK

When the three disk have been added to the VM, they can be listed at the OS level.

```

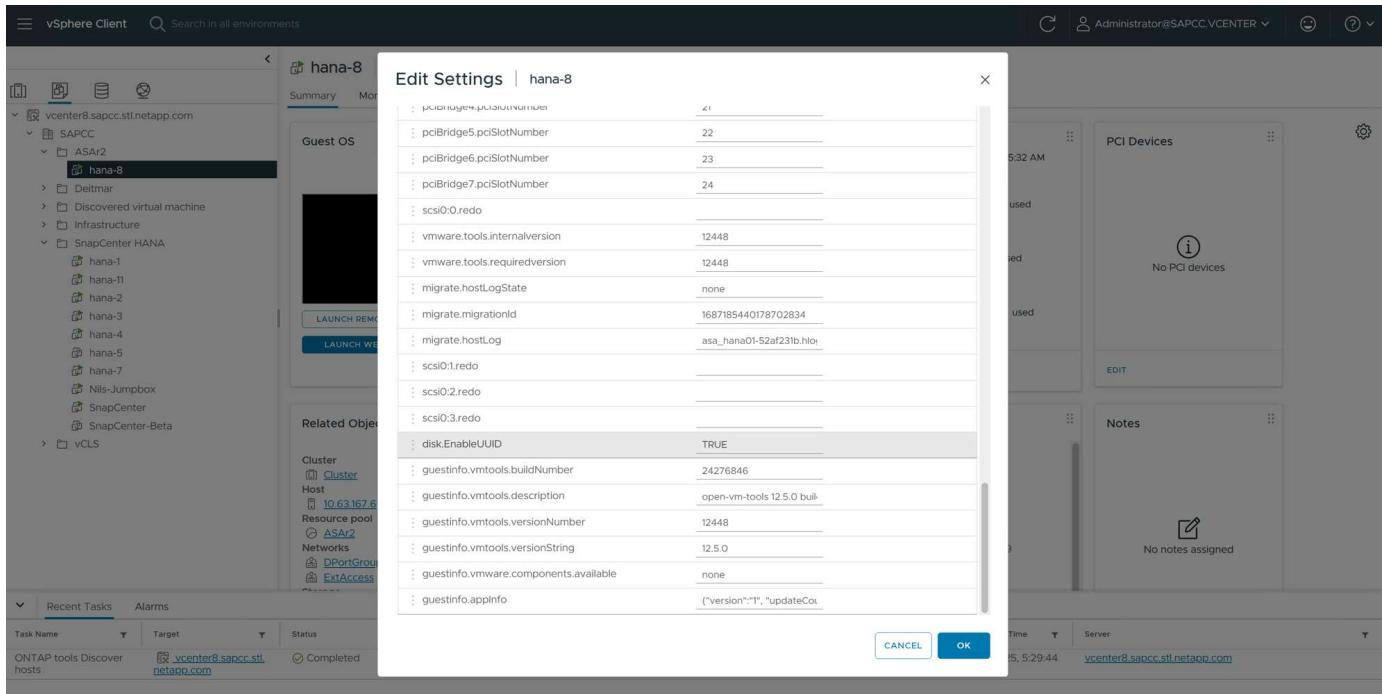
hana-8:~ # lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINTS
sda 8:0 0 100G 0 disk
└─sda1 8:1 0 256M 0 part /boot/efi
└─sda2 8:2 0 82G 0 part
└─system-root 254:0 0 60G 0 lvm /root
  /var
  /usr/local
  /tmp
  /srv
  /opt
  /home
  /boot/grub2/x86++_+64-efi
  /boot/grub2/i386-pc
  /.snapshots
  /
└─system-swap 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 nrext64=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 nrext64=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 nrext64=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,wszie=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.

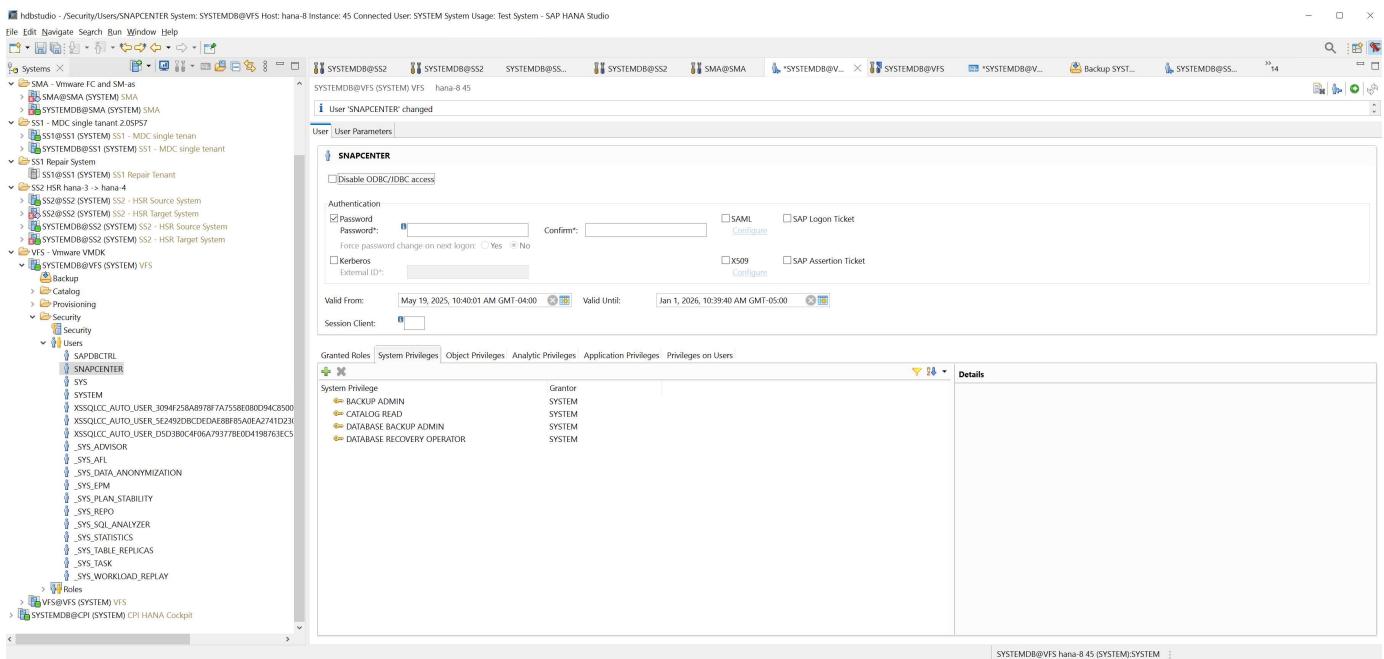


With the described configuration the /usr/sap/VFS directory will be on the OS VMDK. If /usr/sap/VFS should be stored in the shared VMDK, the hana shared disk could be partitioned to provide another file system for /usr/sap/VFS.

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.

```
vfssadm@hana-8:/usr/sap/VFS/HDB45> hdbuserstore SET VFSKEY hana-8:34513
SNAPCENTER <password>
```

Retroactive report: Operation succeed.

Check access with:

```
vfssadm@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

```
vfssadm@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	10.63.150.246	10.63.150.245	AFF	✓	
hana-backup		10.1.2.175	AFF	✓	
hana-cloud-dr		10.63.150.247	FSx	Not applicable	
hana-dr		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.'].

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.]

[View Logs](#)

[Cancel Job](#)

[Close](#)

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" />
```

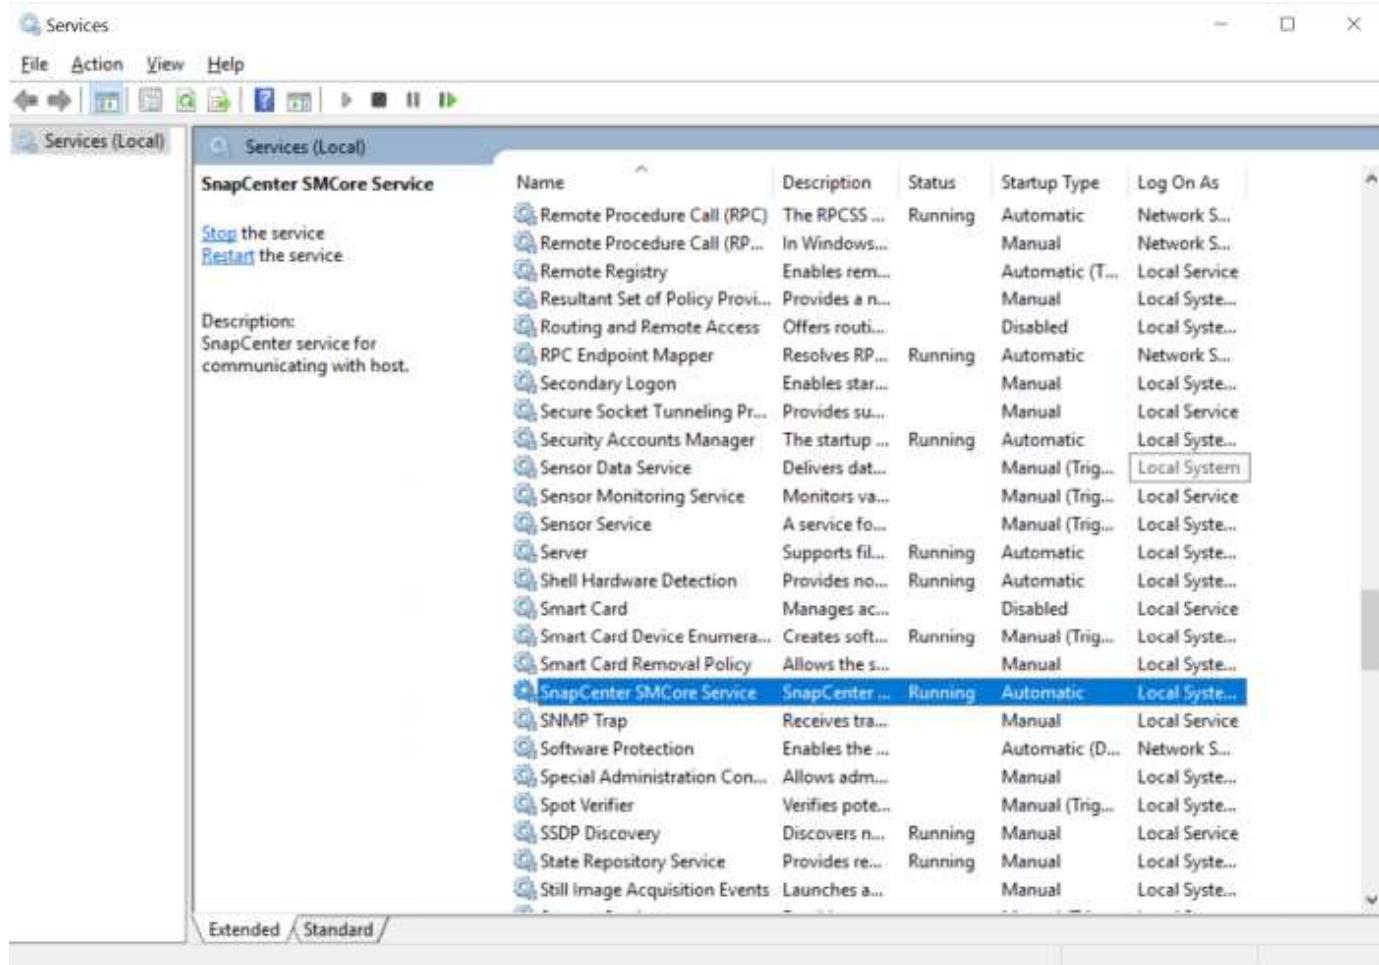
SMCoreServiceHost.dll.config - Notepad

```

<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="OracleAlmActivityParallelUmountInterval" value="20" />
<add key="OracleAlmActivityParallelUmountInterval" 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="IsRestEnabledForStorageConnection" value="true" />
<add key="MinOntapVersionToUseREST" value="9.13.1" />
<add key="IS_COLO_SNAPCENTER_AGENT" value="true" />
<add key="IS_SCM_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.

Host Details

Host Name: scv-vmw.sapcc.stl.netapp.com
Host IP: 10.63.167.24
Overall Status: Running
Host Type: vSphere
System: Stand-alone
Credentials: SCV-sapcc
Push Server: vCenter Host: 10.63.167.20, vCenter Port: 443, vCenter User: administrator@sapcc.vcenter
Plug-ins: SnapCenter Plug-in 6.1.0 for VMware vSphere
VMware vSphere
Submit Cancel Reset

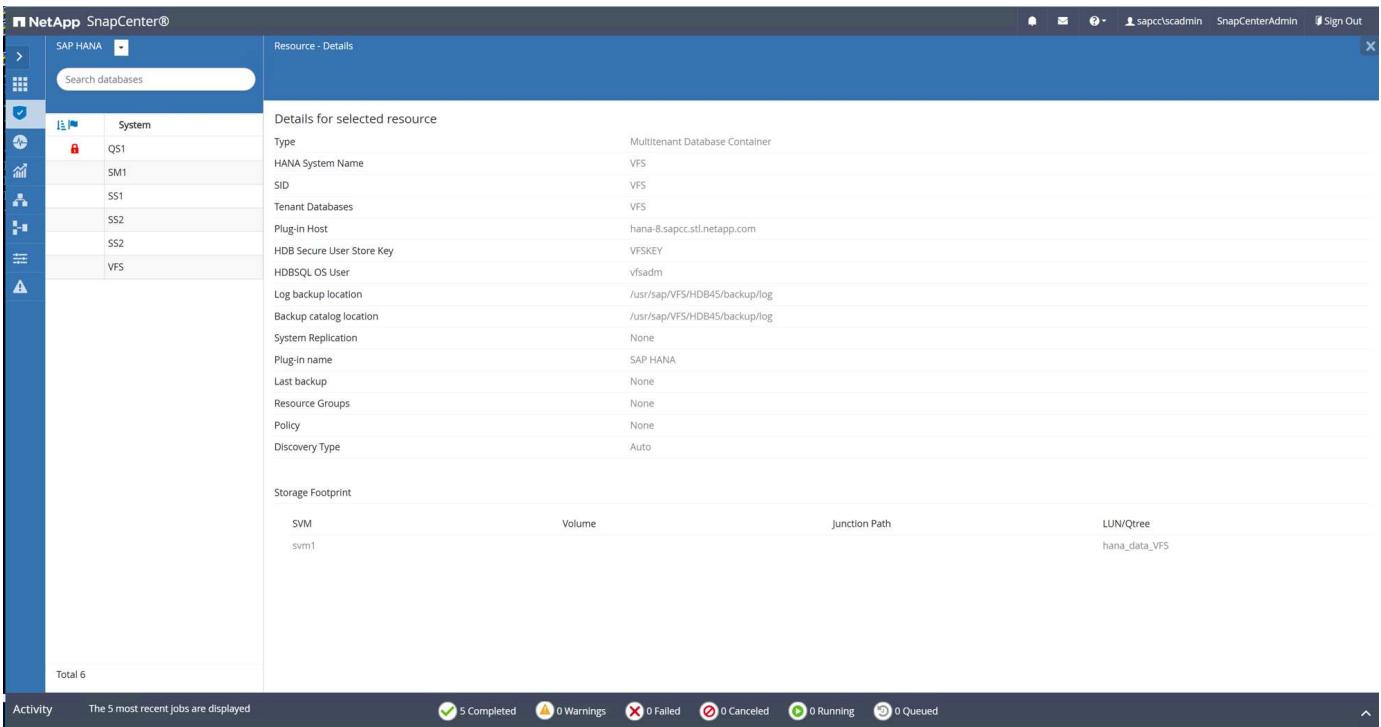
Add HANA host



No specific requirements. Plugin deployment and auto discovery is done as usual.

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	VMware vSphere	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 shows the NetApp SnapCenter interface. The left sidebar has a 'SAP HANA' dropdown and a search bar. The main panel is titled 'Resource - Details' and shows 'Details for selected resource' for a Multitenant Database Container. The resource is named 'QS1' and is of type 'VFS'. It has a 'HANA System Name' of 'hana-8.sapcc.stl.netapp.com' and a 'SID' of 'VFS'. The 'Tenant Databases' are 'SS1', 'SS2', and 'SS2'. The 'Plug-in Host' is 'HDB Secure User Store Key'. The 'HDBSQL OS User' is 'vfsadm'. The 'Log backup location' is '/usr/sap/VFS/HDB45/backup/log' and the 'Backup catalog location' is '/usr/sap/VFS/HDB45/backup/log'. 'System Replication' is set to 'None'. The 'Plug-in name' is 'SAP HANA'. The 'Last backup' and 'Resource Groups' are both 'None'. The 'Policy' is 'None' and the 'Discovery Type' is 'Auto'. Under 'Storage Footprint', there is one SVM named 'svm1' which maps to a Volume and a Junction Path, both labeled 'hana_data_VFS'. A note at the bottom says 'Total 6'. The bottom navigation bar shows 'Activity' with 'The 5 most recent jobs are displayed' and status counts: 5 Completed, 0 Warnings, 0 Failed, 0 Canceled, 0 Running, and 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

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](#)

NetApp SnapCenter®

SAP HANA

VFS Topology

Manage Copies

12 Backups

0 Clones

Local copies

Summary Card

12 Backups

12 Snapshot based backups

0 File-Based backups

0 Clones

0 Snapshots Locked

Primary Backup(s)

Backup Name

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 2:30:12 PM
SnapCenter_hana-8_LocalSnap_Hourly_05-21-2025_10.29.00.3581		1	05/21/2025 10:30:12 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 11:03:03 AM

Total 6

Total 12

Activity The 5 most recent jobs are displayed

5 Completed 0 Warnings 0 Failed 0 Canceled 0 Running 0 Queued

SnapCenter creates a consistency group (CG) and adds the storage unit `hana_data_VFS` to the CG. Snapshots are created at CG level.

NetApp ONTAP System Manager | A70-SAPCC

Storage

The basic unit of storage is a LUN (for SCSI hosts) or NVMe namespace (for NVMe). You can add LUN or NVMe namespace storage units based on your data center configuration. [More](#)

19 Storage units

68.5 TiB Available

19 Online

0 Offline

+ Add

Name	Consistency group	Capacity	Data reduction	Host mapping	IOPS	Latency (ms)	Throughput (MB/s)
<code>hana_data_VFS</code>	<code>sc20250520_110422_689</code>	100 GiB	1 to 1	<code>otv_host-44_e3d7e9d4-46f3-4fd1</code>	1	0.07	0
<code>hana_log_VFS</code>	-	100 GiB	1.19 to 1	<code>otv_host-44_e3d7e9d4-46f3-4fd1</code>	4	0.23	0.41
<code>hana_shared_VFS</code>	-	100 GiB	2.8 to 1	<code>otv_host-44_e3d7e9d4-46f3-4fd1</code>	6	0.23	0.43

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_05-22-2025_02.29.03.2899	May/21/2025 10:36 PM	
SnapCenter_hana-8_LocalSnap_Hourly_05-22-2025_06.29.00.3541	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 X

1 **Restore scope** Select the restore types

2 Recovery scope Complete Resource i

3 PreOps

4 PostOps

5 Notification

6 Summary

Restore from SnapCenter_hana-8_LocalSnap_Hourly_05-22-2025_06.29.00.3706 X

1 **Restore scope** Recover database files using

2 Recovery scope Recover to most recent state i

Recover to point in time i

Recover to specified data backup i

No recovery i

3 PreOps

4 PostOps

5 Notification

6 Summary

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.

Storage

The basic unit of storage is a LUN (for SCSI hosts) or NVMe namespace (for NVMe). You can add LUN or NVMe namespace storage units based on your data center configuration. [More](#)

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_e3d7e9d4-46f3-4fd1-0	0	0	0
hana_data_VFS_Clone_0522250947396031	-	100 GiB	1 to 1	otv_host-57_e3d7e9d4-46f3-4fd1-0	-	-	-
hana_log_VFS	-	100 GiB	1.19 to 1	otv_host-44_e3d7e9d4-46f3-4fd1-0	0	0	0
hana_shared_VFS	-	100 GiB	2.33 to 1	otv_host-44_e3d7e9d4-46f3-4fd1-0	0	0	0

The new LUN (datastore) based on the cloned storage unit gets attached to the ESX cluster.

hana_data_VFS(sc-20250522094807386)

Name	Type	Path
sd.dsf	Folder	[hana_data_VFS(sc-20250522094807386)]/sd.dsf
hana-8	Folder	[hana_data_VFS(sc-20250522094807386)]/hana-8

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Reconfigure virtual machine	hana-8	Completed		SAPCC.VCENTER\Administrator	7 ms	05/22/2025, 9:48:25 AM	05/22/2025, 9:48:26 AM	vcenter8.sapcc.sti.netapp.com
Rename datastore	hana_data_VFS	Completed		SAPCC.VCENTER\Administrator	5 ms	05/22/2025, 9:48:15 AM	05/22/2025, 9:48:21 AM	vcenter8.sapcc.sti.netapp.com
Resizing data store	10.A3.167.6	Completed		SAPCC.VCENTER\Administrator	4 ms	05/22/2025, 9:48:05	05/22/2025, 9:48:05	vcenter8.sapcc.sti.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\MDCVFS'

- ✓ ▾ Restore 'hana-8.sapcc.stl.netapp.com\hana\MDCVFS'
- ✓ ▾ 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

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. The left sidebar lists databases: System, QFS, QSI, SM1, SS1, SS2, SS2, and VFS. The main area displays 'Manage Copies' with 12 Backups and 0 Clones. A 'Primary Backup(s)' table lists 12 backups for 'SnapCenter_hana-8_LocalSnap_Hourly' from 06/16/2025 to 06/14/2025. A 'Summary Card' on the right shows 13 Backups, 12 Snapshot-based backups, 1 File-based backup, 0 Clones, and 0 Snapshots Locked. The bottom status bar shows activity: 3 Completed, 0 Warnings, 0 Failed, 0 Canceled, 2 Running, and 0 Queued.

The target host and SID must be provided.

The 'Clone From Backup' dialog is open. The 'Location' tab is selected, showing 'Select the host to create the clone'. The 'Plug-in host' dropdown is set to 'hana-9.sapcc.stl.netapp.com'. The 'Target Clone SID' dropdown is set to 'QFS'.

The 'Clone From Backup' dialog is open. The 'Location' tab is selected, showing 'LUN Map Settings'. The 'Igroup protocol' dropdown is set to 'FCP', with 'FCP' selected in the list.

In our example we are using a post-script to execute the recovery at the target host.

Clone From Backup

The following commands will run on the Plug-in Host: `hana-9.sapcc.stl.netapp.com`

Enter optional commands to run before performing a clone operation [?](#)

Pre clone command

Enter optional commands to run after performing a clone operation [?](#)

Post clone command

`/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.

NetApp ONTAP System Manager | A70-SAPCC

Storage

The basic unit of storage is a LUN (for SCSI hosts) or NVMe namespace (for NVMe). You can add LUN or NVMe namespace storage units based on your data center configuration. [More](#)

Name	Consistency group	Capacity	Data reduction	Host mapping	iOPS	Latency (ms)	Throughput (MB/s)
<code>hana_data_QFS</code>	-	100 GiB	5.46 to 1	<code>otv_host-44_e3d7e9dd4-46f3-4fd1_4</code>	4	0.11	0.39
<code>hana_data_VFS</code>	<code>sc20250520_110422_689</code>	100 GiB	1 to 1	<code>otv_host-44_e3d7e9dd4-46f3-4fd1_5</code>	5	0.12	0.39
<code>hana_data_VFS_Clone_06172507005037511</code>	-	100 GiB	1 to 1	<code>otv_host-57_e3d7e9dd4-46f3-4fd1_23</code>	23	0.11	1.24
<code>hana_log_QFS</code>	-	100 GiB	4.1 to 1	<code>otv_host-44_e3d7e9dd4-46f3-4fd1_5</code>	5	0.10	0.39
<code>hana_log_VFS</code>	-	100 GiB	1.22 to 1	<code>otv_host-44_e3d7e9dd4-46f3-4fd1_8</code>	8	0.12	0.40
<code>hana_shared_QFS</code>	-	100 GiB	2.81 to 1	<code>otv_host-44_e3d7e9dd4-46f3-4fd1_5</code>	5	0.11	0.39
<code>hana_shared_VFS</code>	-	100 GiB	1.69 to 1	<code>otv_host-44_e3d7e9dd4-46f3-4fd1_5</code>	5	0.13	0.39

SnapCenter then attaches the LUN (datastore) to the ESX host, on which the target HANA VM is running.

vSphere Client | Search in all environments

10.63.167.6 | ACTIONS

Summary Monitor Configure Permissions VMs Datastores Networks Updates

Quick Filter

Name	Status	Type	Datastore Cluster	Capacity	Free
<code>datastore1 (2)</code>	Norm	VMFS 6	al	766 GB	764.58 GB
<code>Datastore_C250</code>	Norm	NFS 3	al	195 TB	195 TB
<code>Datastore_One</code>	Norm	NFS 3	al	2.85 TB	1.22 TB
<code>DatastoreA460</code>	Norm	NFS 3	al	500 GB	271.24 GB
<code>hana_data_QFS</code>	Norm	VMFS 6	al	99.75 GB	87.26 GB
<code>hana_data_VFS</code>	Norm	VMFS 6	al	99.75 GB	90.94 GB
<code>hana_data_VFS/sc2025061_707015334</code>	Norm	VMFS 6	al	99.75 GB	90.94 GB
<code>hana_log_QFS</code>	Norm	VMFS 6	al	99.75 GB	91.31 GB
<code>hana_log_VFS</code>	Norm	VMFS 6	al	99.75 GB	91.3 GB
<code>hana_shared_QFS</code>	Norm	VMFS 6	al	99.75 GB	87 GB
<code>hana_shared_VFS</code>	Norm	VMFS 6	al	99.75 GB	80.55 GB
<code>OS_Image</code>	Norm	NFS 3	al	142.5 GB	55.39 GB

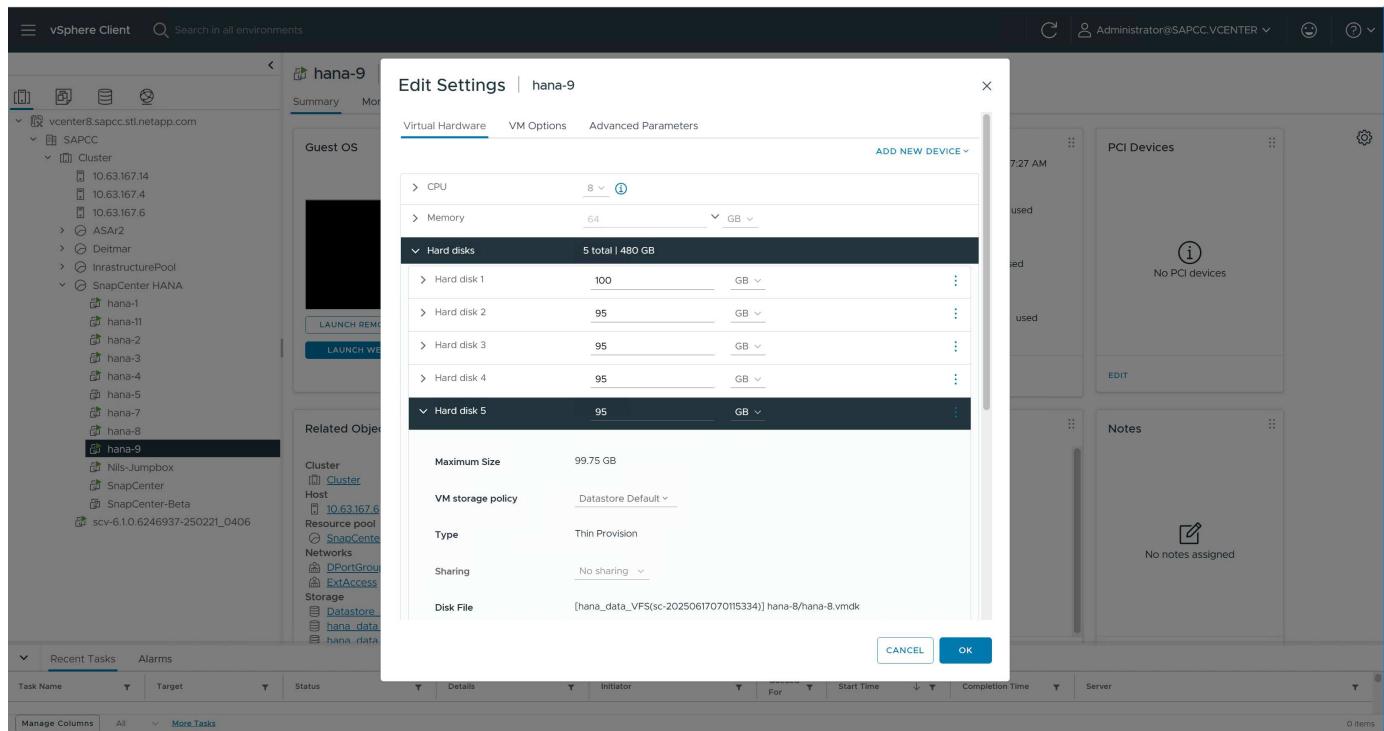
Manage Columns Export Deselect All

Items per page 35 12 Items

Recent Tasks Alarms

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Manage Columns	All	More Tasks	0 Items					

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
efivarsfs 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,wszie=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 X

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.

Resource - Details

Details for selected resource

Type	Multitenant Database Container
HANA System Name	QS1
SID	QS1
Tenant Databases	QFS
Plug-in Host	hana-9.sapcc.stl.netapp.com
HDB Secure User Store Key	QFSKEY
HDBSQL OS User	qfsadm
Log backup location	/usr/sap/QFS/HDB45/backup/log
Backup catalog location	/usr/sap/QFS/HDB45/backup/log
System Replication	None
Plug-in name	SAP HANA
Last backup	None
Resource Groups	None
Policy	None
Discovery Type	Auto
Backup Name	SnapCenter_hana-8_LocalSnap_Hourly_06-17-2025_10.29.00.4260
Backup Name of Clone	SnapCenter_hana-8_LocalSnap_Hourly_06-17-2025_10.29.00.4260

Storage Footprint

SVM	Volume	Junction Path	LUN/Qtree
svml			hana_data_VFS_Clone_06172511013515617

Total 7

Activity The 5 most recent jobs are displayed

3 Completed 1 Warning 1 Failed 0 Canceled 0 Running 0 Queued

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.

VFS Topology

Manage Copies

12 Backups
1 Clone
Local copies

Primary Clone(s)

Clone SID	Clone Host	Clone Name	Start Date	End date
QFS	hana-9.sapcc.stl.netapp.com	hana-8_sapcc_stl_netapp_com_hana_MDC_VFS__clone_146515_MDC_VFS_06-17-2025_10.27.55	06/17/2025 11:01:58 AM	06/17/2025 11:10:22 AM

Total 7 Total 1

Activity The 5 most recent jobs are displayed

3 Completed 1 Warning 1 Failed 0 Canceled 0 Running 0 Queued

In our example we are using a pre-script to shutdown the target HANA database.

Delete Clone

X

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 X

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

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:

- [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

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.

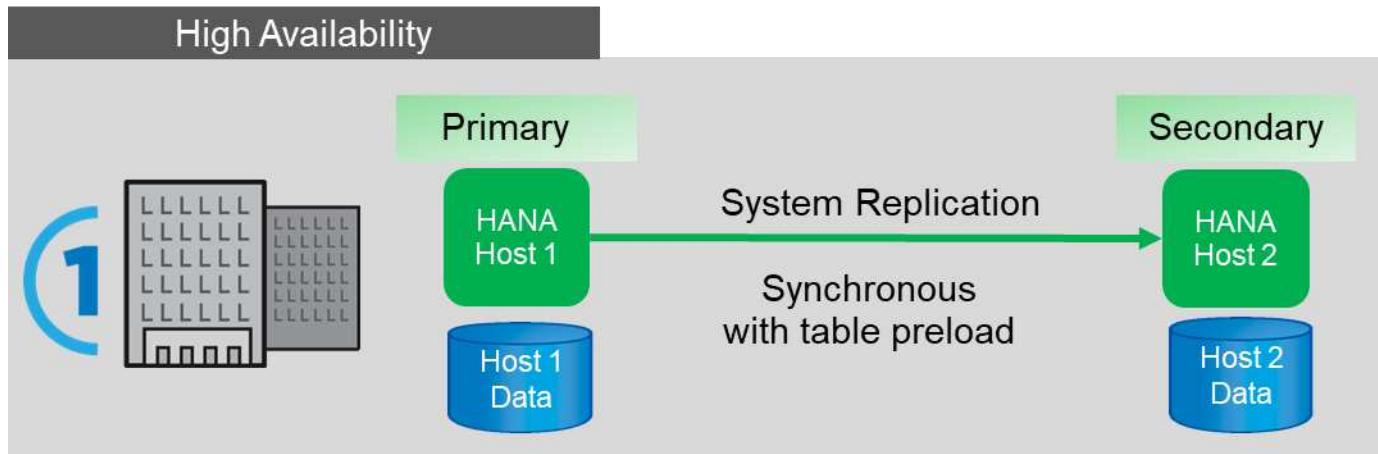
High availability with an RPO of zero and a minimal RTO

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.

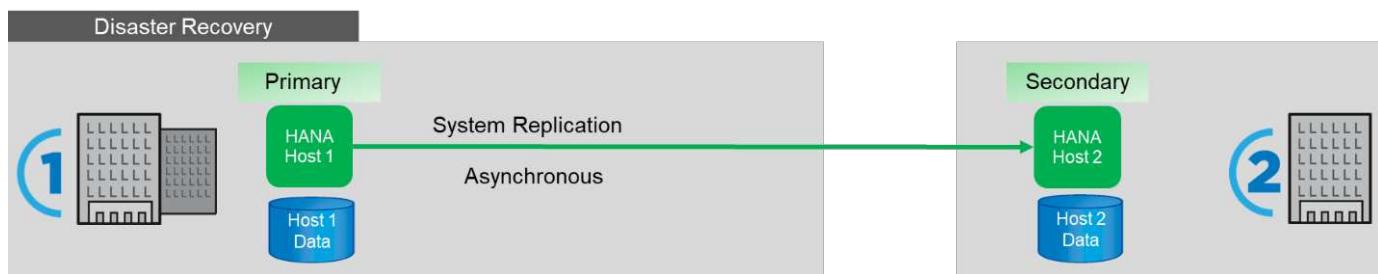
The rest of this document focuses on backup operations with SAP System Replication configured as a high-availability solution.



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.



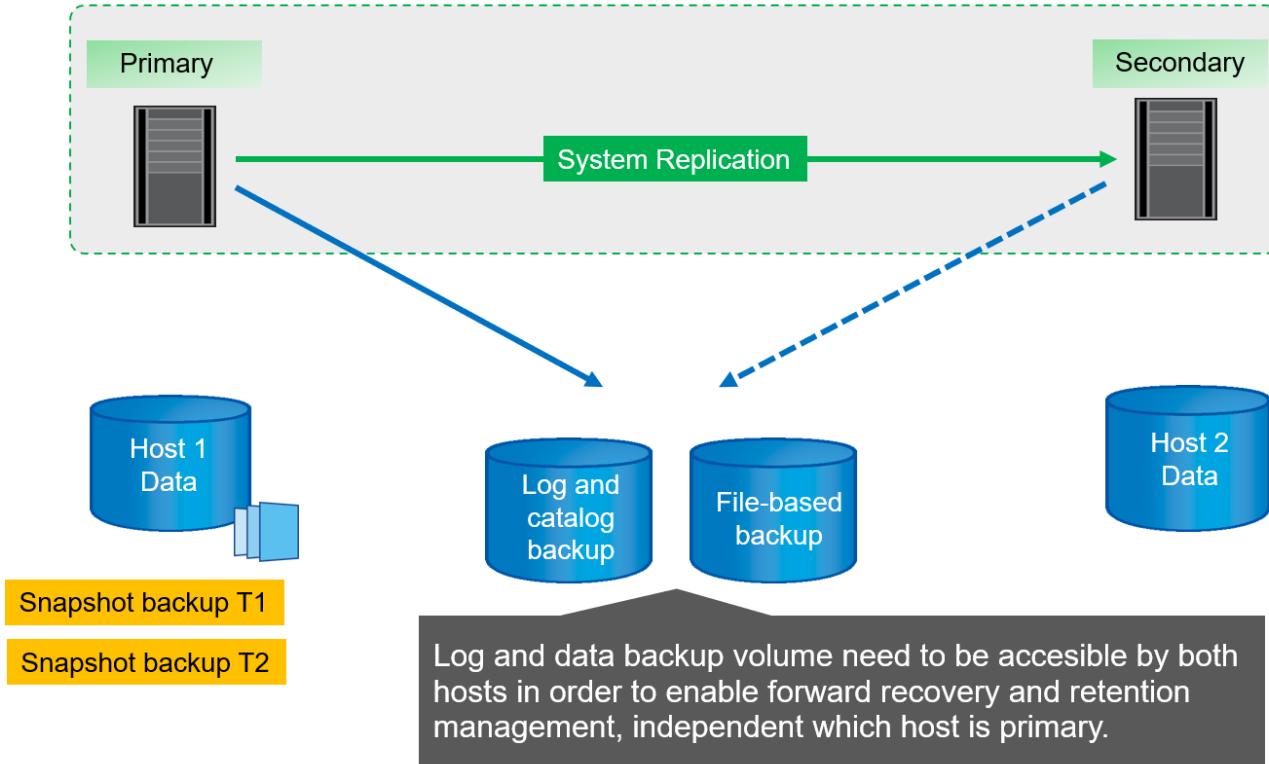
Storage Snapshot backups and SAP System Replication

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.

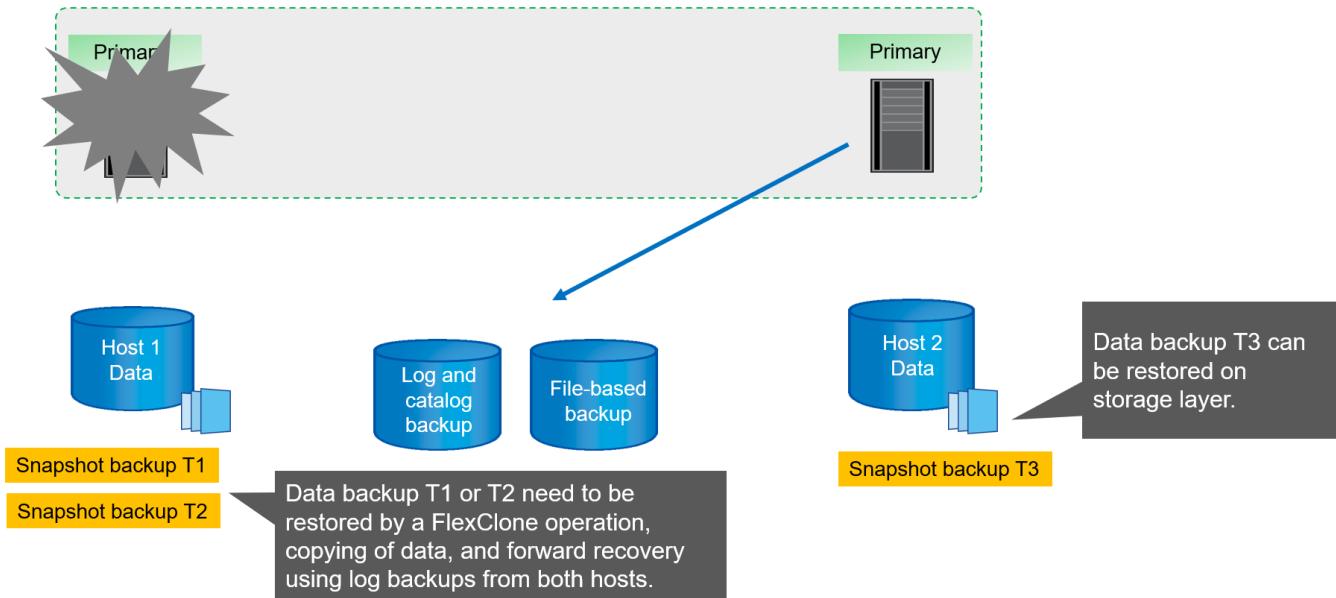
The ability to use any backup for restore and to do forward recovery using log backups from both hosts requires a shared log backup location that is accessible from both hosts. NetApp recommends that you use a shared storage volume. However, you should also separate the log backup destination into subdirectories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

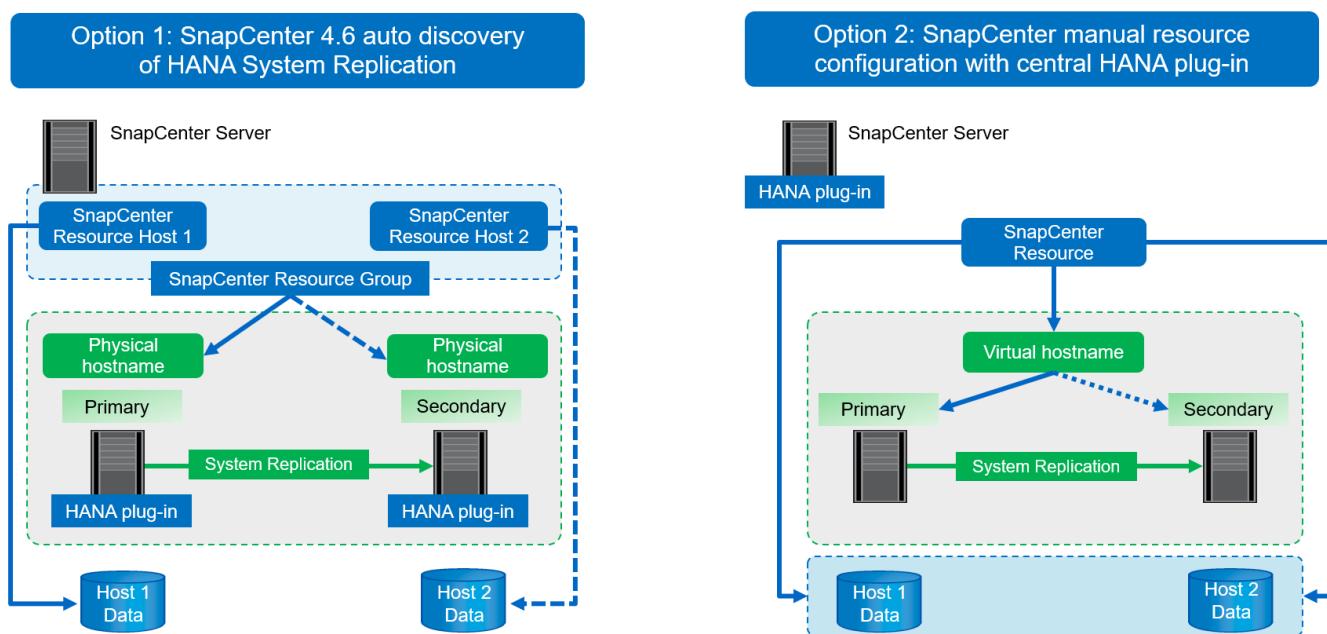
The backup created at host 2 can be restored directly at the storage layer. If you must use a backup created at host 1, then the backup must be copied from the host 1 storage volume to the host 2 storage volume. Forward recovery uses the log backups from both hosts.



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.



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.

	Resource group with SnapCenter 4.6	Single SnapCenter resource and virtual IP address
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.

	Resource group with SnapCenter 4.6	Single SnapCenter resource and virtual IP address
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>
```

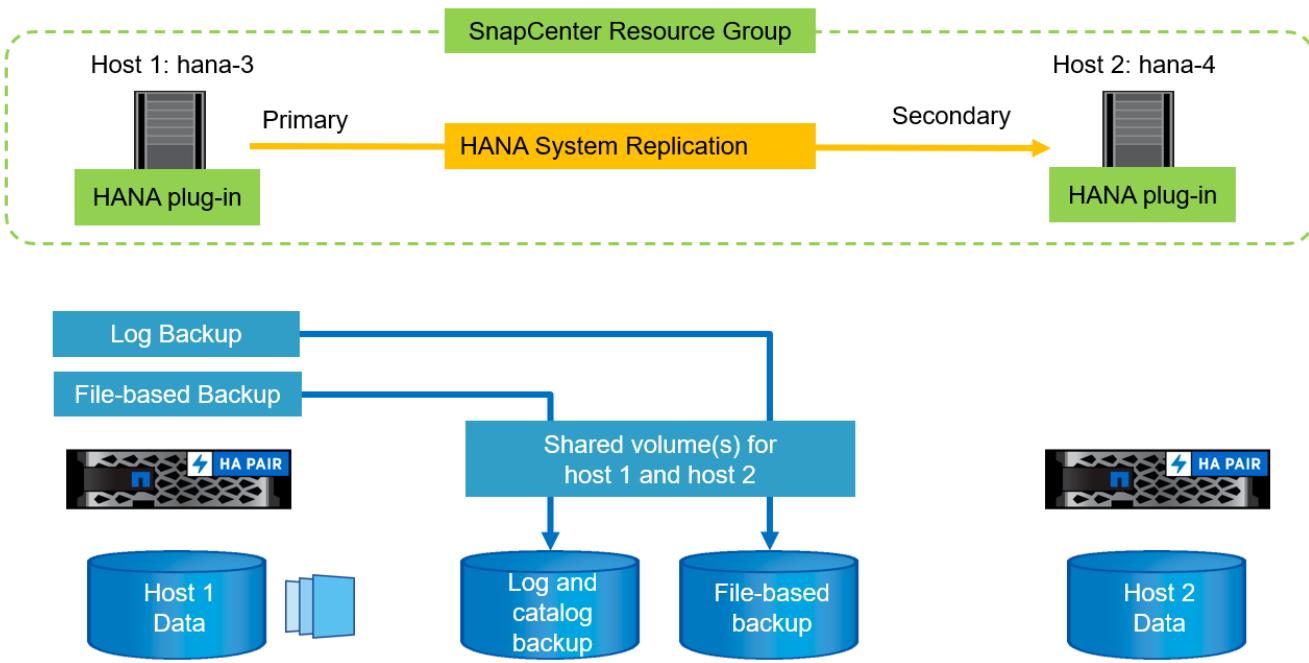
```
ss2adm@hana- 4:/ > hdbuserstore set SS2KEY hana-4:33313 SNAPCENTER
<password>
```

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.

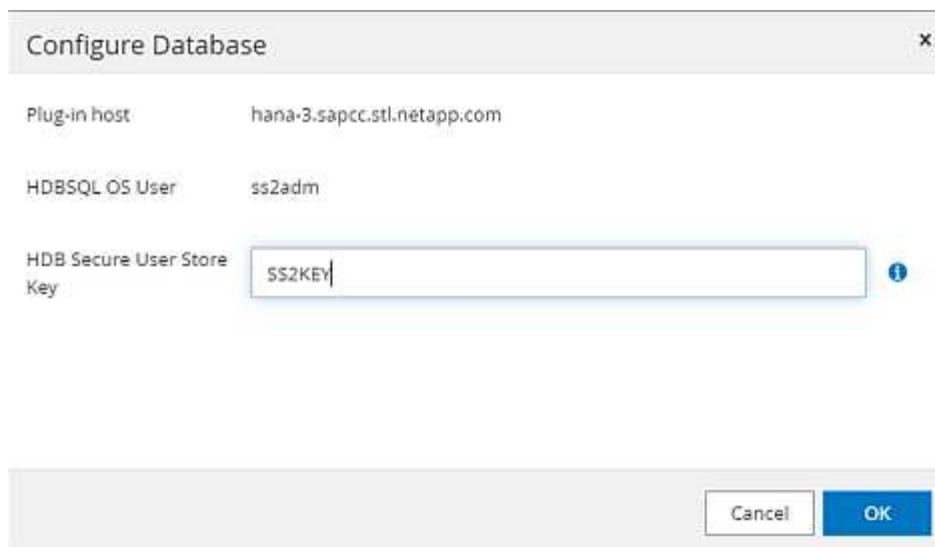


After you have installed the SnapCenter HANA plug-in on both HANA hosts, the HANA systems are shown in the SnapCenter resource view in the same way as other autodiscovered resources. Starting with SnapCenter 4.6, an additional column is displayed that shows the status of HANA system replication (enabled/disabled, primary/secondary).

The screenshot shows the NetApp SnapCenter interface. The left sidebar has 'Resources' selected. The main area shows a table for 'SAP HANA' resources. The table has columns: ID, System, System ID (SID), Tenant Databases, Replication, Plug-in Host, Resource Groups, Policies, Last backup, and Overall Status. Two rows are listed, both with a red error icon. The first row's 'System' column shows 'SS2', 'SS2', 'SS2', 'Enabled (Primary)', 'hana-3.sapcc.stl.netapp.com', and 'Not protected'. The second row's 'System' column shows 'SS2', 'SS2', 'SS2', 'Enabled (Secondary)', 'hana-4.sapcc.stl.netapp.com', and 'Not protected'. The 'Replication' column is highlighted with a blue box.

ID	System	System ID (SID)	Tenant Databases	Replication	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
1	SS2	SS2	SS2	Enabled (Primary)	hana-3.sapcc.stl.netapp.com				Not protected
2	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.



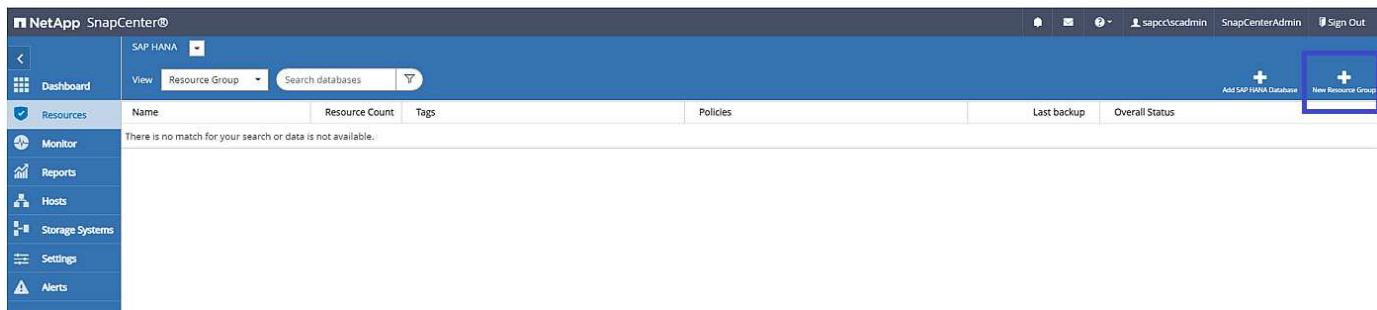
Additional autodiscovery steps are executed, and SnapCenter show the resource details. With SnapCenter 4.6, the system replication status and the secondary server are listed in this view.

System			
System	SS2		
SS2			
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		
plugin 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	

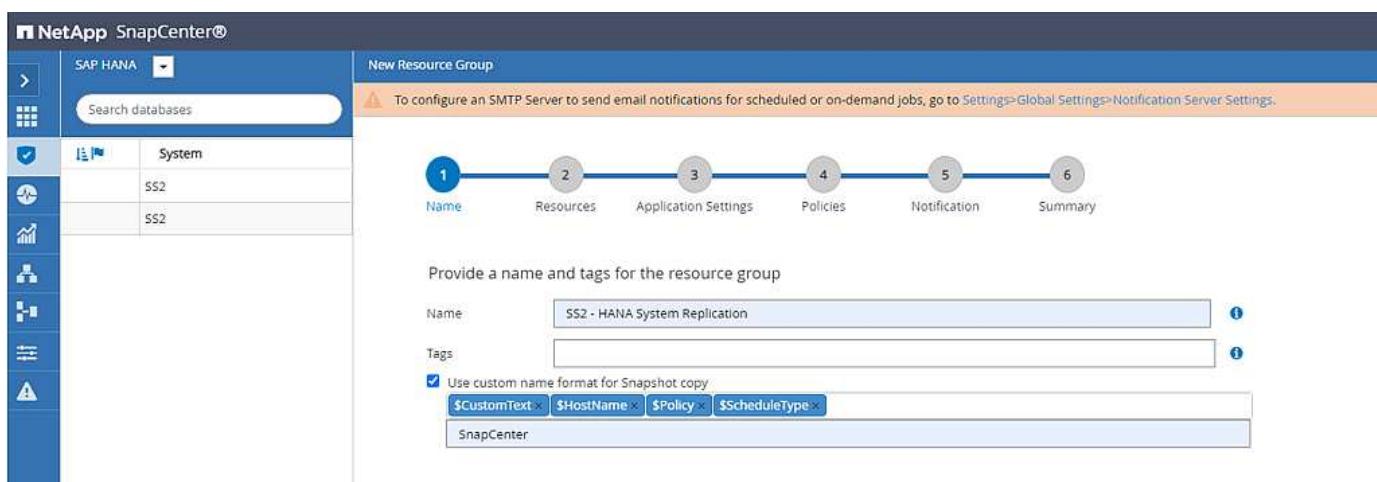
After performing the same steps for the second HANA resource, the autodiscovery process is complete and both HANA resources are configured in SnapCenter.

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.



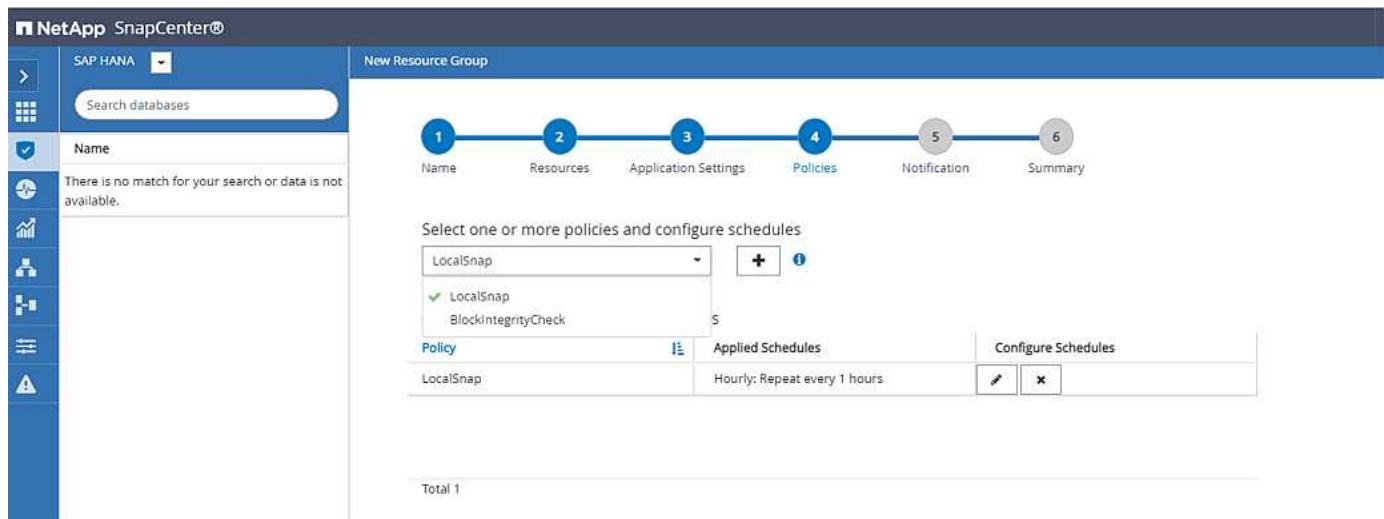
You must add both HANA hosts to the resource group.

>' and '<' are visible between the lists." data-bbox="75 580 919 910"/>

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.

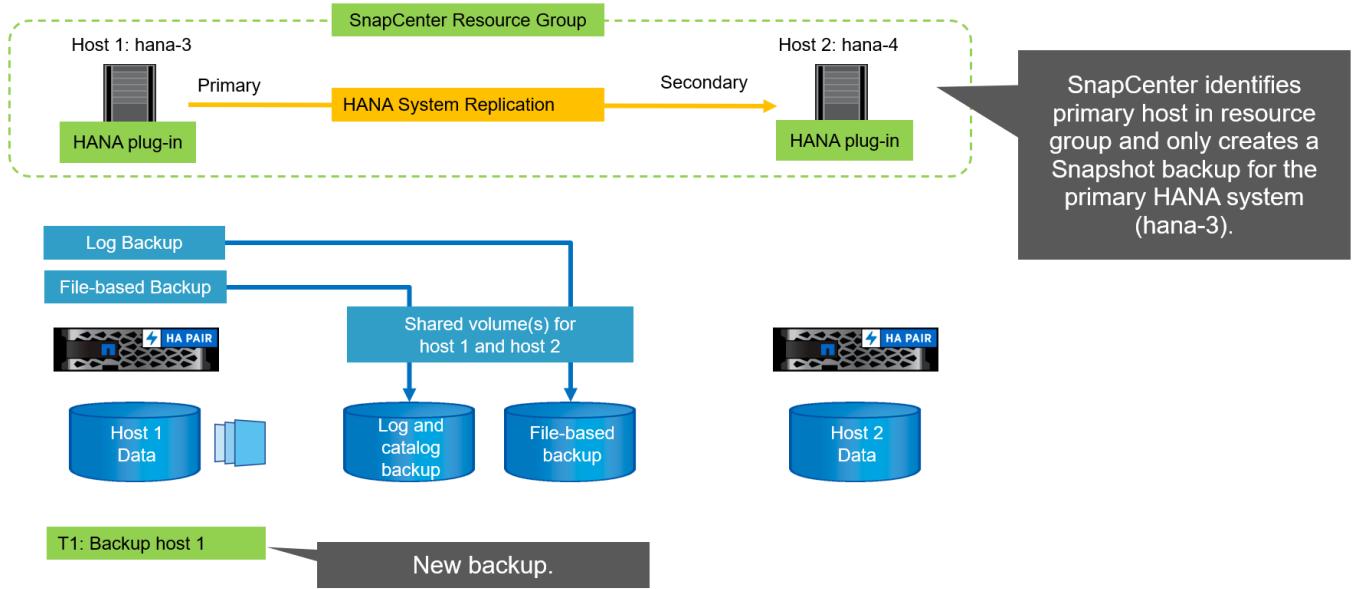


The resource group configuration is now finished and backups can be executed.

The top screenshot shows the 'SS2 - HANA System Replication Details' page with two entries for 'SS2'. The bottom screenshot shows the 'Resources' dashboard with two entries for 'SS2'.

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.

Job Details

Backup of Resource Group 'SS2 - HANA System Replication' with policy 'LocalSnap'

✓ ▾ Backup of Resource Group 'SS2 - HANA System Replication' with policy 'LocalSnap'

 ✓ ▾ Refresh HANA replication resources on host(s): hana-3.sapcc.stl.netapp.com, hana-4.sapcc.stl.netapp.com

 ✓ ▾ hana-3.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

 ✓ ▶ Application Clean-Up

Task Name: Backup Start Time: 12/13/2021 8:35:33 AM End Time:

[View Logs](#) [Cancel Job](#) [Close](#)

A Snapshot backup has now been created at the primary HANA resource. The hostname included in the backup name shows hana-3.

NetApp SnapCenter®

SAP HANA

SS2 - HANA System Replication ...

SS2 Topology

Name: SS2 - HANA System Replication

Manage Copies

1 Backup 0 Clones

Local copies

Primary Backup(s)

Backup Name: SnapCenter_hana-3_LocalSnap_Hourly_12-13-2021_08:35:30.7075

Count: 1 End Date: 12/13/2021 8:36:32 AM

Summary Card

1 Backup

1 Snapshot based backup

0 File-Based backups

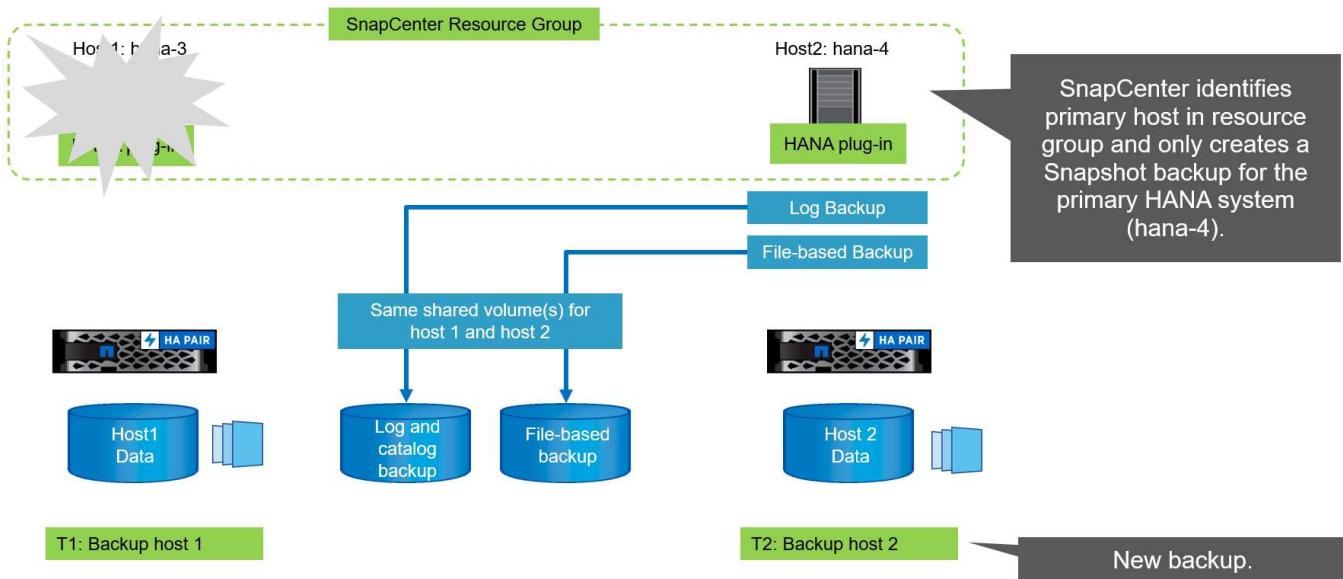
0 Clones

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.

Job Details

Backup of Resource Group 'SS2 - HANA System Replication' with policy 'LocalSnap'

- ✓ ▾ Backup of Resource Group 'SS2 - HANA System Replication' with policy 'LocalSnap'
 - ✓ ▾ Refresh HANA replication resources on host(s): hana-3.sapcc.stl.netapp.com, hana-4.sapcc.stl.netapp.com
 - ✓ ▾ hana-4.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
 - ✓ ▶ Application Clean-Up

Task Name: Backup Start Time: 12/13/2021 8:56:44 AM End Time:

[View Logs](#) [Cancel Job](#) [Close](#)

A Snapshot backup has now been created at the primary HANA resource. The hostname included in the backup name shows hana-4.

NetApp SnapCenter®

SAP HANA SS2 - HANA System Replication ... SS2 Topology

System Resource Name SS2

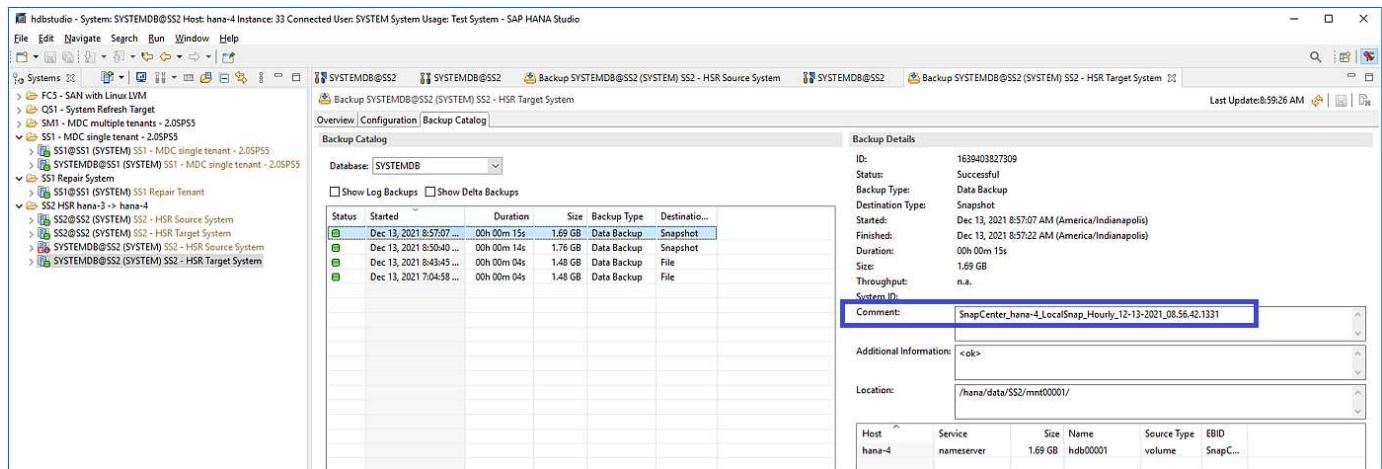
Manage Copies 1 Backup 0 Clones Local copies

Summary Card 1 Backup 1 Snapshot-based backup 0 File-based backups 0 Clones

Primary Backup(s)

Backup Name	Count	End Date
SnapCenter_hana-4_LocalSnap_Hourly_12-13-2021_08:56:42.1331	1	12/13/2021 8:57:41 AM

The same Snapshot backup is also visible in the HANA backup catalog.



Status	Started	Durations	Size	Backup Type	Destinatio...
Dec 13, 2021 8:57:07 ...	00h 00m 15s	1.69 GB	Data Backup	Snapshot	
Dec 13, 2021 8:56:40 ...	00h 00m 14s	1.76 GB	Data Backup	Snapshot	
Dec 13, 2021 8:43:45 ...	00h 00m 04s	1.48 GB	Data Backup	File	
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.

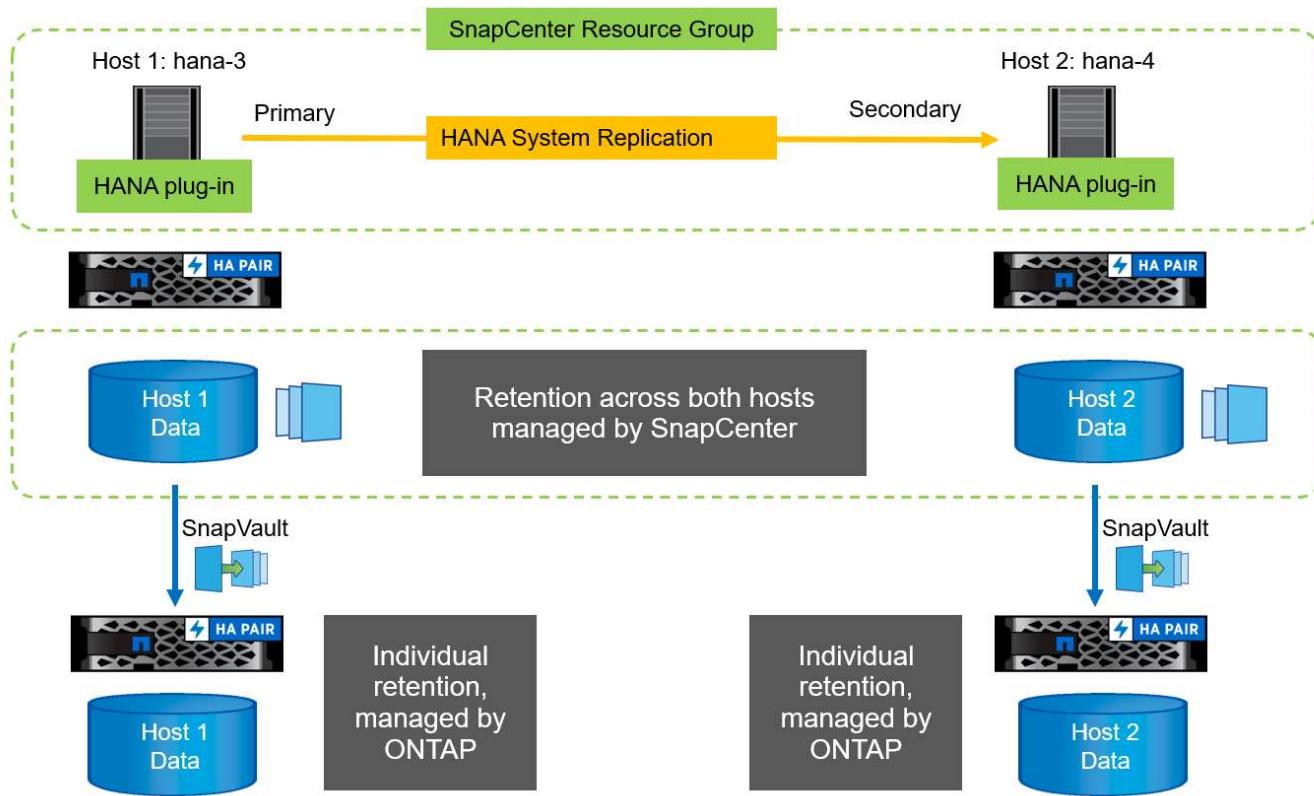


If a takeover to the secondary host is not performed for a long time, the number of changed blocks for the first SnapVault update at the secondary host will be high.

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.



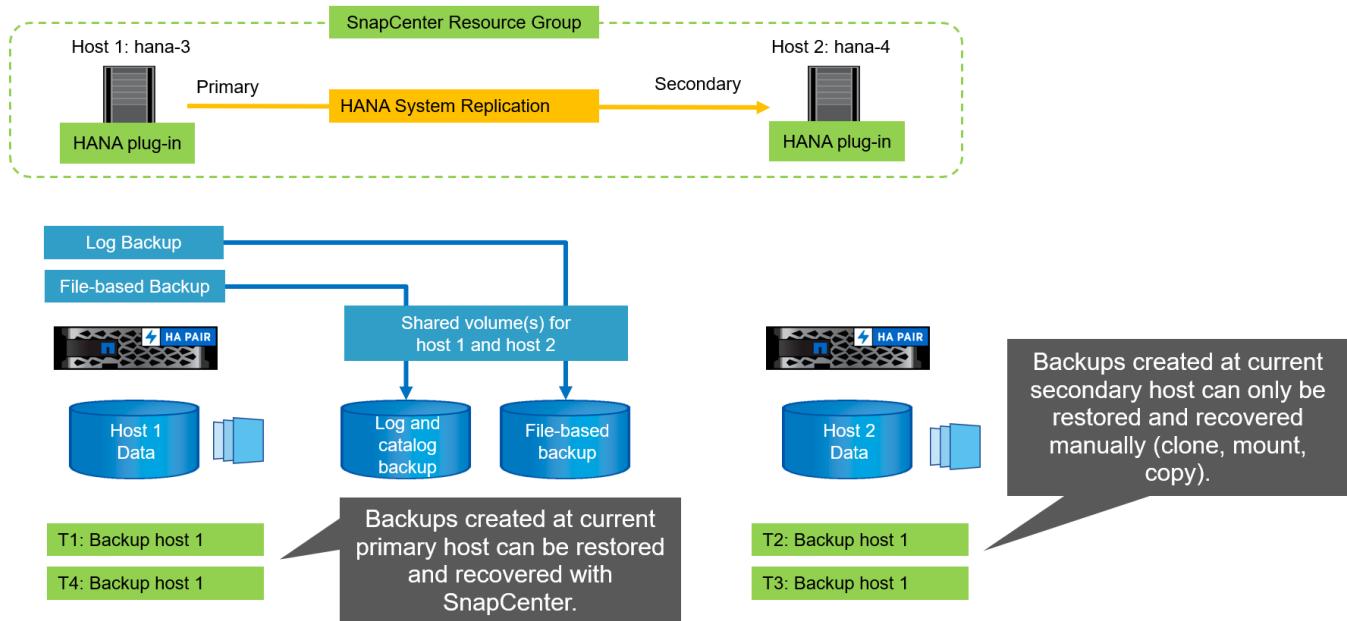
Retention management

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 [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

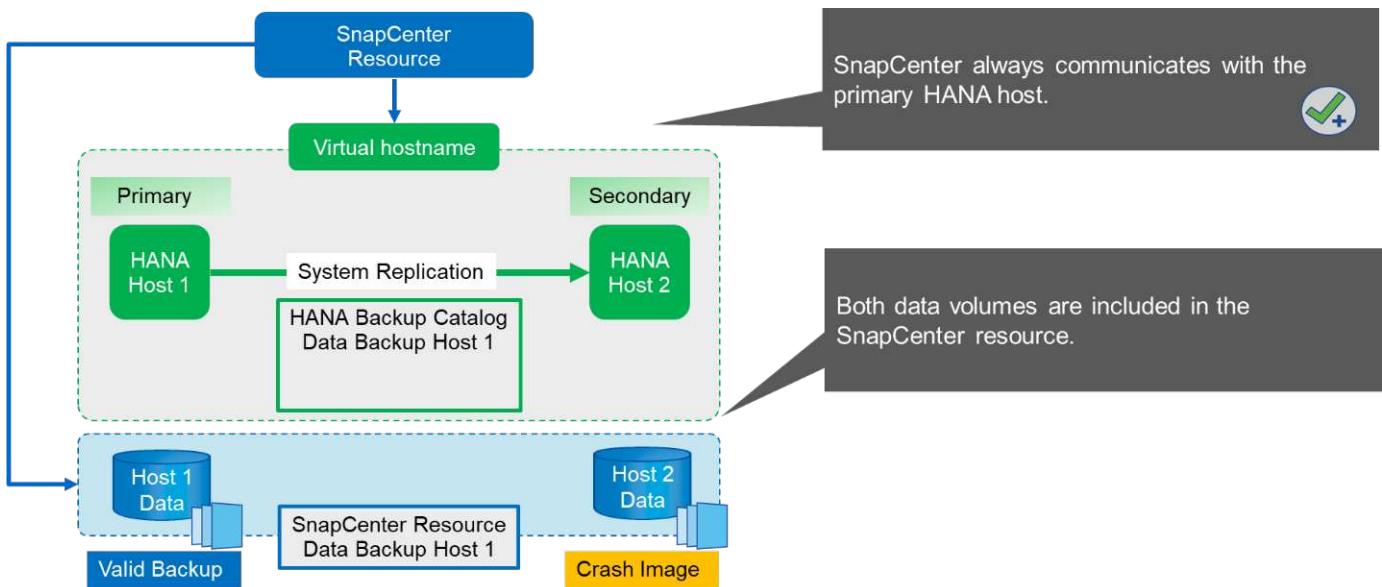
A SnapCenter resource is configured with the virtual IP address (host name) of the HANA System Replication environment. With this approach, SnapCenter always communicates with the primary host, regardless of whether host 1 or host 2 is primary. The data volumes of both SAP HANA hosts are included in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

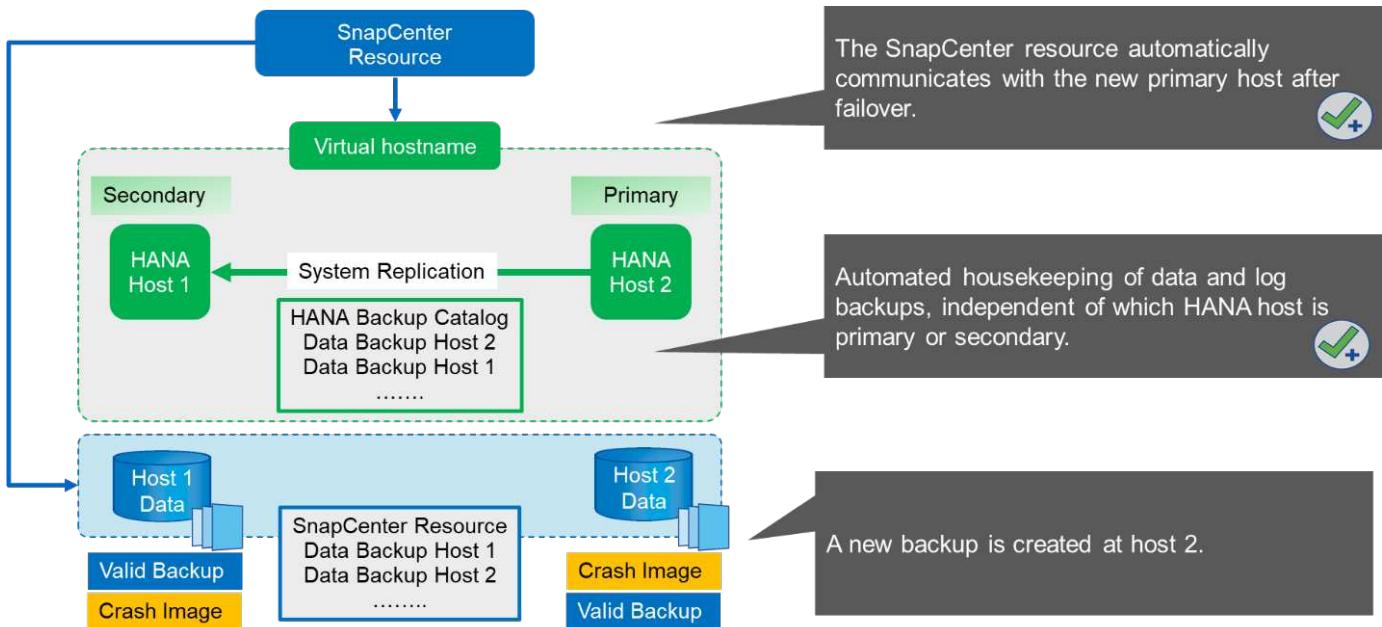
When a backup is executed with host 1 as the primary host, a database-consistent Snapshot backup is created at the data volume of host 1. Because the data volume of host 2 is part of the SnapCenter resource, another Snapshot copy is created for this volume. This Snapshot copy is not database consistent;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

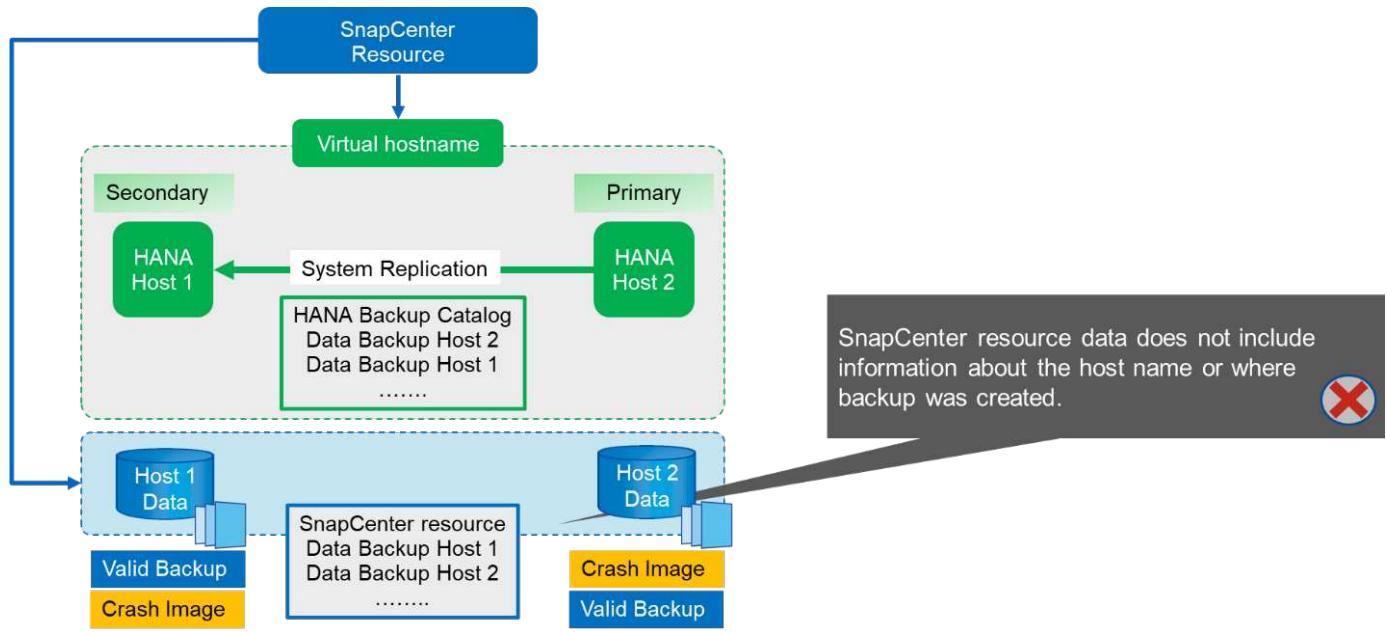


As discussed in the section [Storage Snapshot Backups and SAP System Replication](#), a restore operation with storage-based Snapshot backups is different, depending on which backup must be restored. It is important to identify which host the backup was created at to determine if the restore can be performed 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

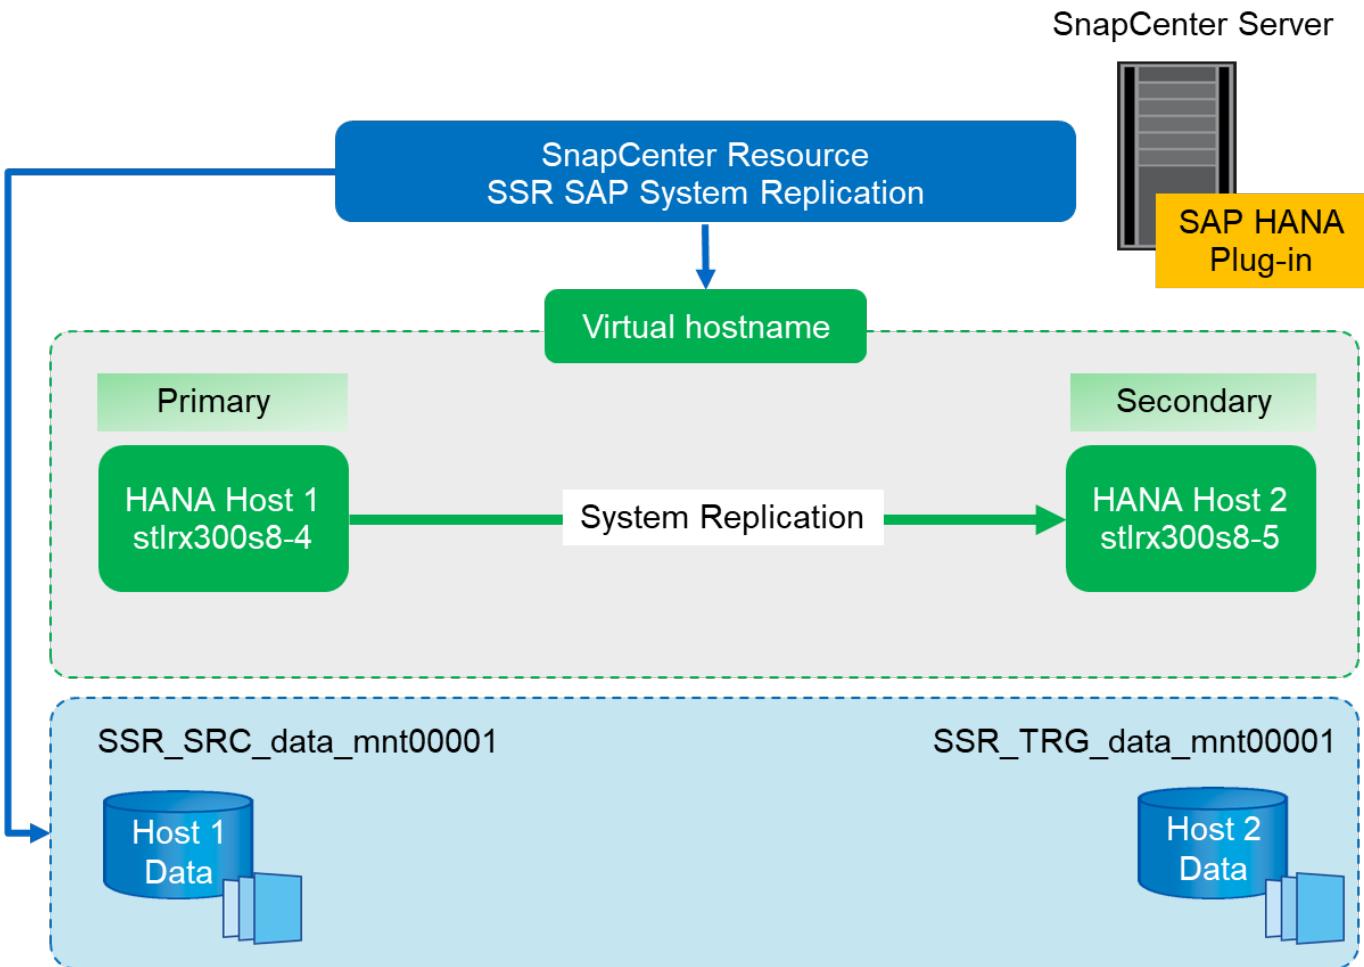
identify which host is currently the primary SAP HANA host.

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

Provide Resource Details

Resource Type Multitenant Database Container (MDC) - Single Tenant

HANA System Name

SID i

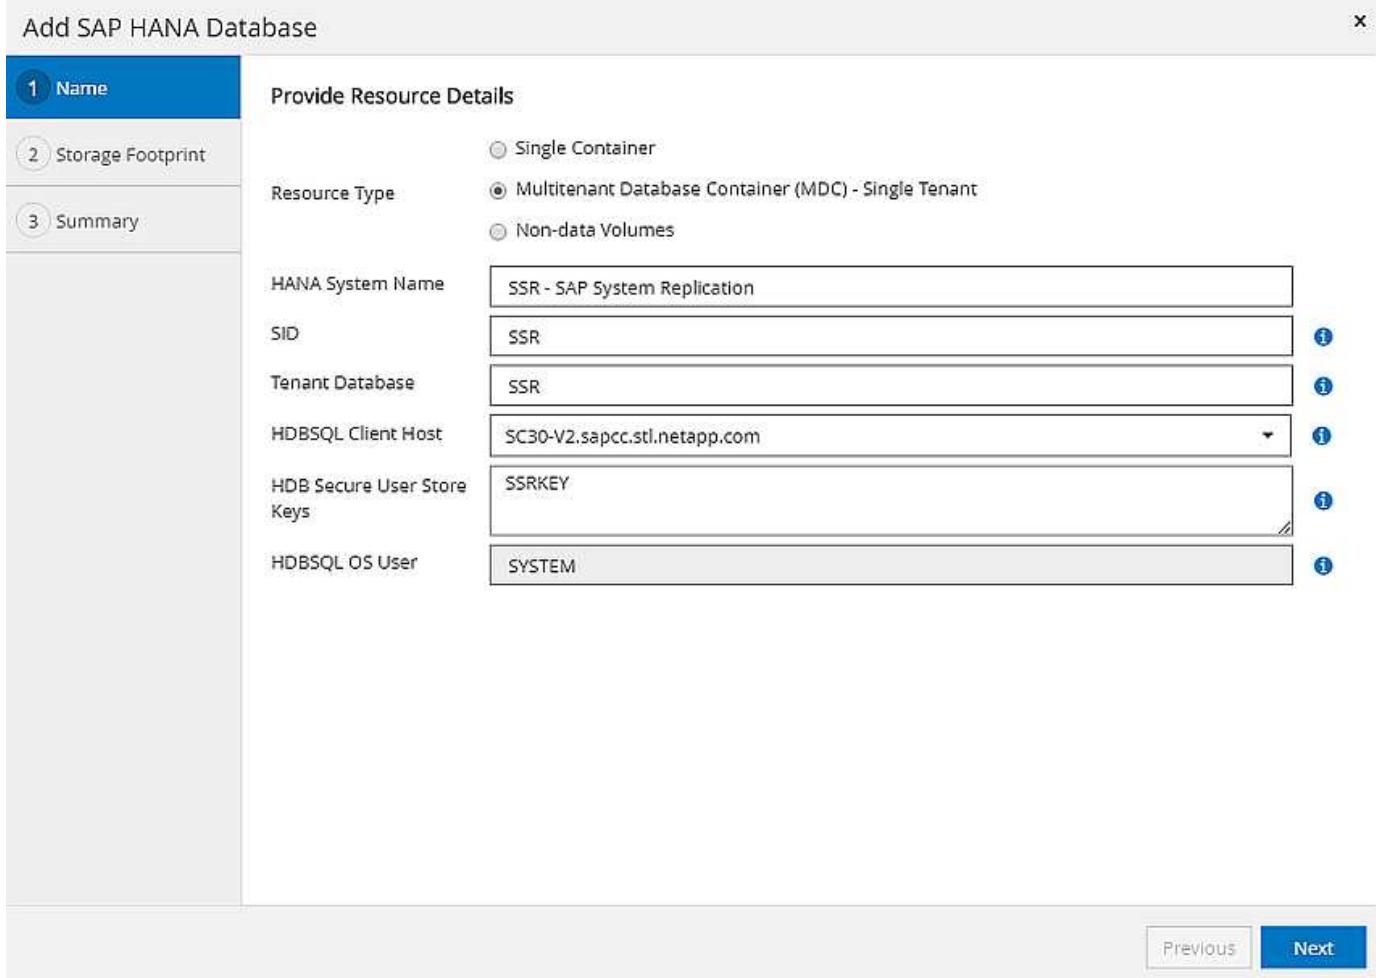
Tenant Database i

HDBSQL Client Host i

HDB Secure User Store Keys i

HDBSQL OS User i

Previous Next



The data volumes of both SAP HANA hosts are included in the storage footprint configuration, as the following figure shows.

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

Select one or more volumes and if required their associated Qtrees and LUNs

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

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.

i 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.

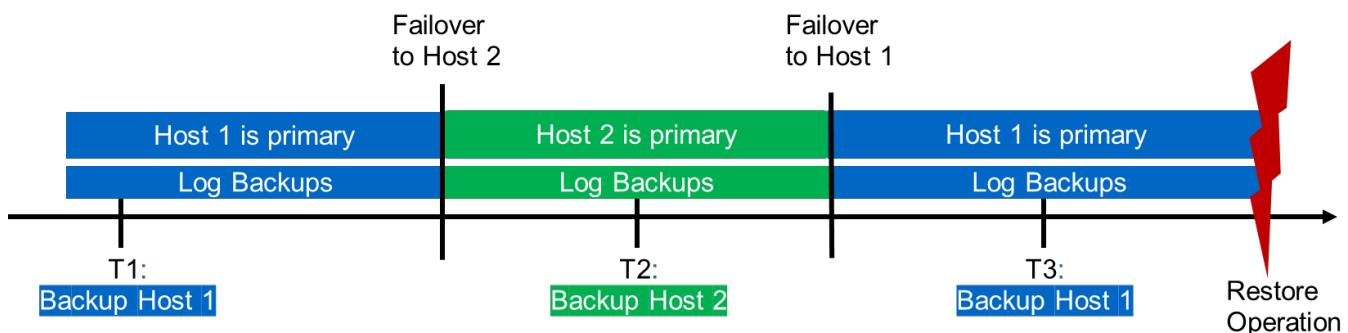
Restore and recovery

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

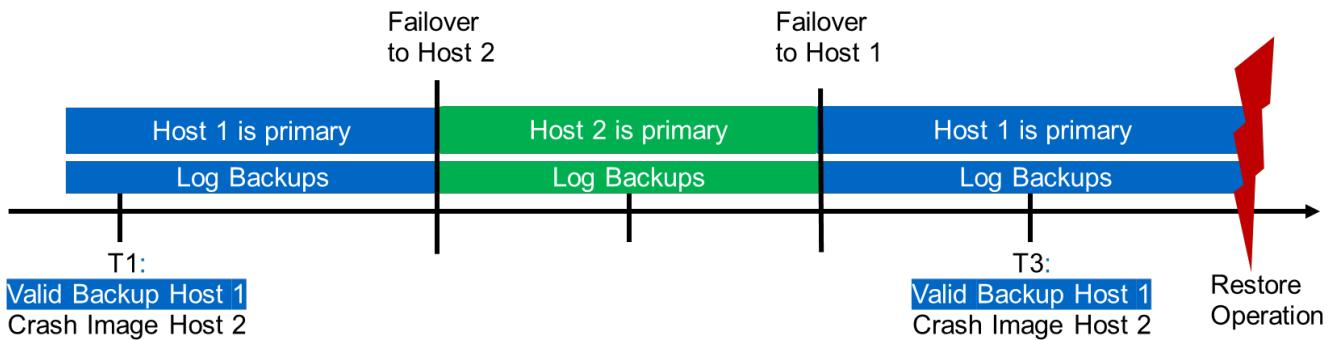
With a single SnapCenter resource configuration, Snapshot copies are created at both storage volumes of both SAP HANA System Replication hosts. Only the Snapshot backup that is created at the storage volume of the primary SAP HANA host is valid to use for forward recovery. The Snapshot copy created at 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.

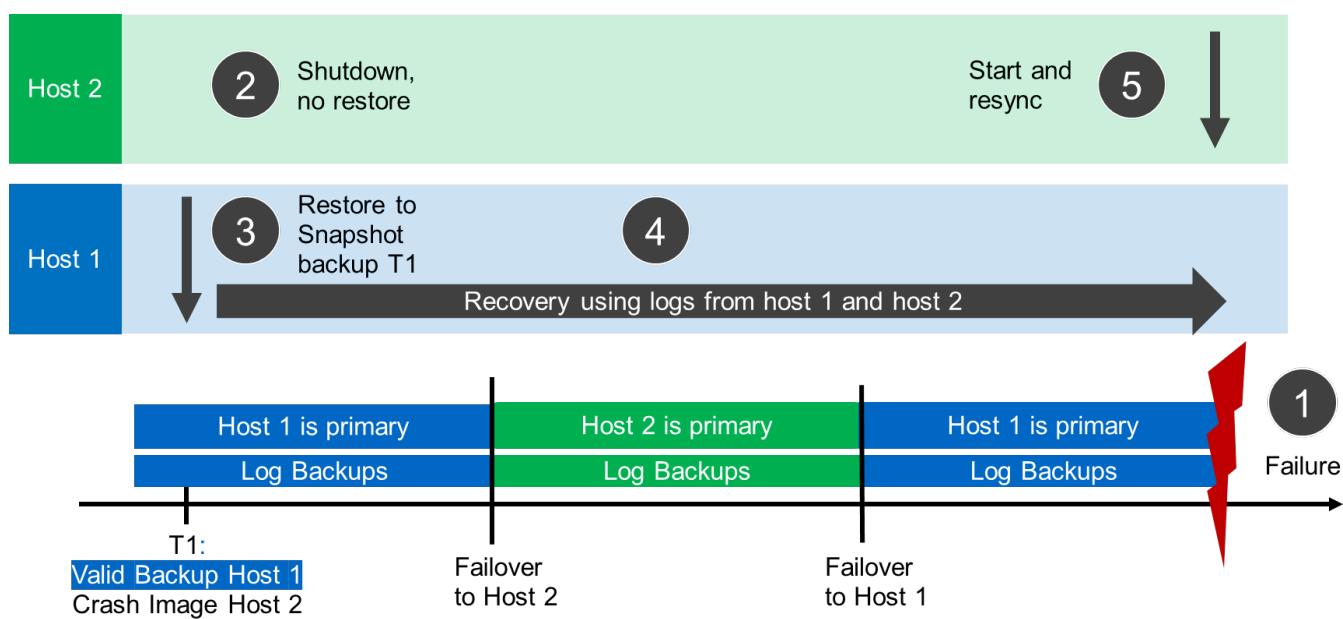


SnapCenter restore of the valid backup only

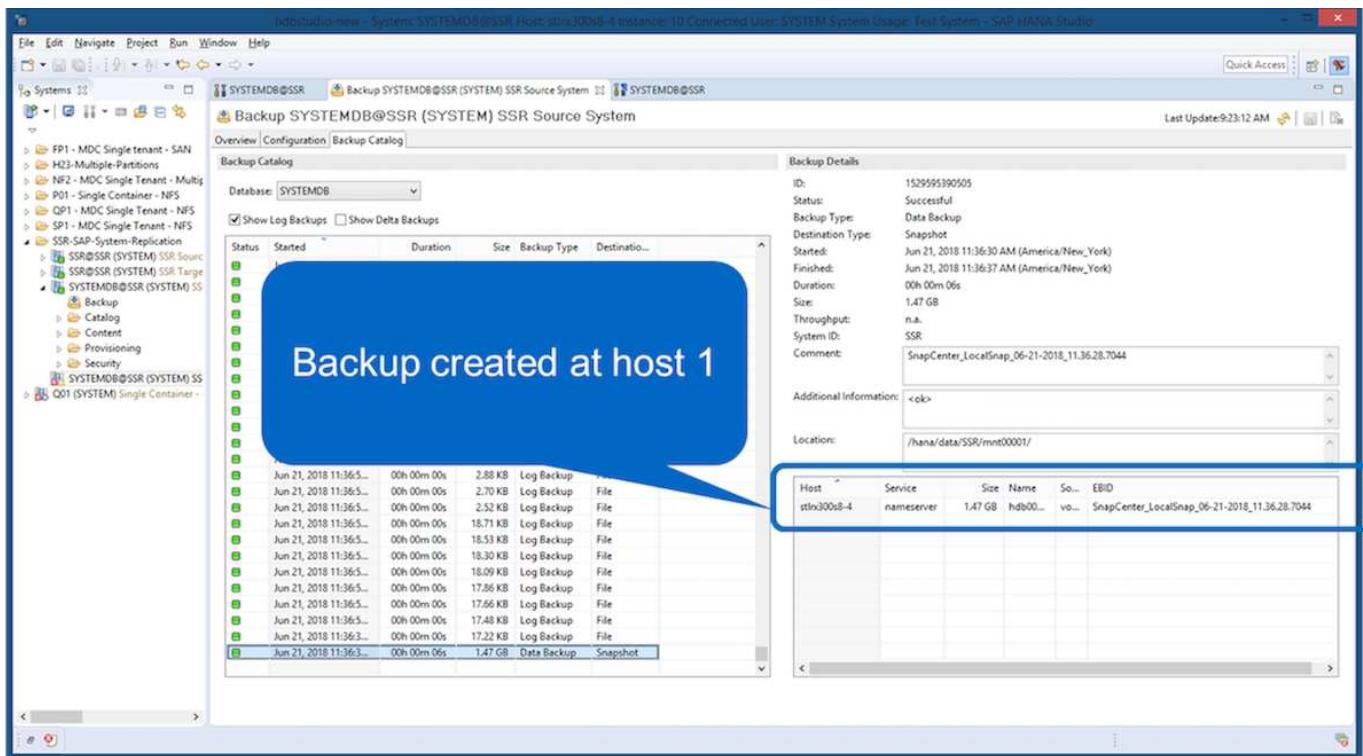
The following figure shows an overview of the restore and recovery scenario described in this section.

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.

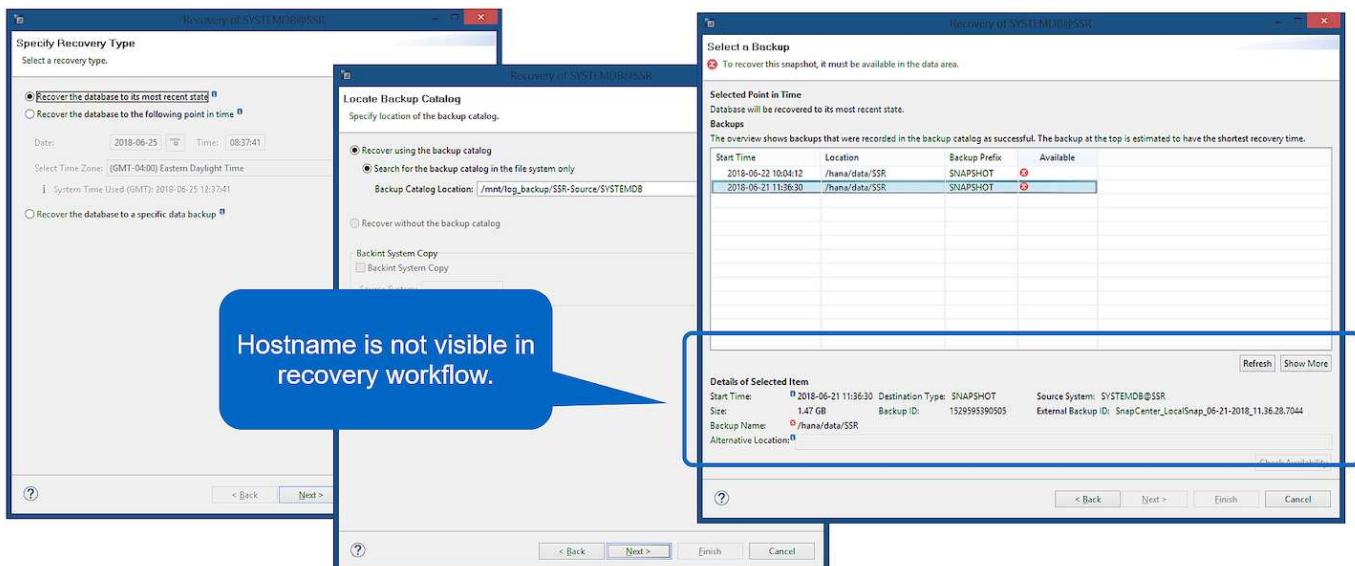


25

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.

The screenshot shows the SAP HANA Studio interface with the 'SSR - SAP System Replication Topology' tab selected. On the left, there's a sidebar with 'System' and 'SSR - SAP System Replication' sections. The main area shows 'Manage Copies' with 2 Backups and 0 Clones. Below that is a 'Primary Backup(s)' table with two entries: 'Backup Name' (SnapCenter_LocalSnap_06-22-2018_10.04.03.2739) and 'SnapCenter_LocalSnap_06-21-2018_11.36.28.7044'. A blue callout box points to the 'Restore from SnapCenter...' button in a modal dialog. The dialog has tabs for 'Restore Scope' (selected), 'PreOps', 'PostOps', 'Notification', and 'Summary'. Under 'Select the restore types', 'File Level' is selected. Under 'Select files to restore', 'File Path' is chosen, and a checkbox for 'hana/vol/SSR_SRC_data_mnt0001' is checked. A note says 'Provide one or more file paths separated by comma'.

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.

The screenshot shows the 'Recovery of SYSTEMDB@SSR' wizard. The first step, 'Select a Backup', shows a table of backups with one entry highlighted in green. A blue callout box points to this entry with the text 'Backup available after SnapCenter restore operation.' The second step, 'Locate Log Backups', shows a 'Locations' field containing '/mnt/log_backup/SSR-Source/SYSTEMDB'. A blue callout box points to this field with the text 'Log backup location is included in backup catalog. No changes are required here.'

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



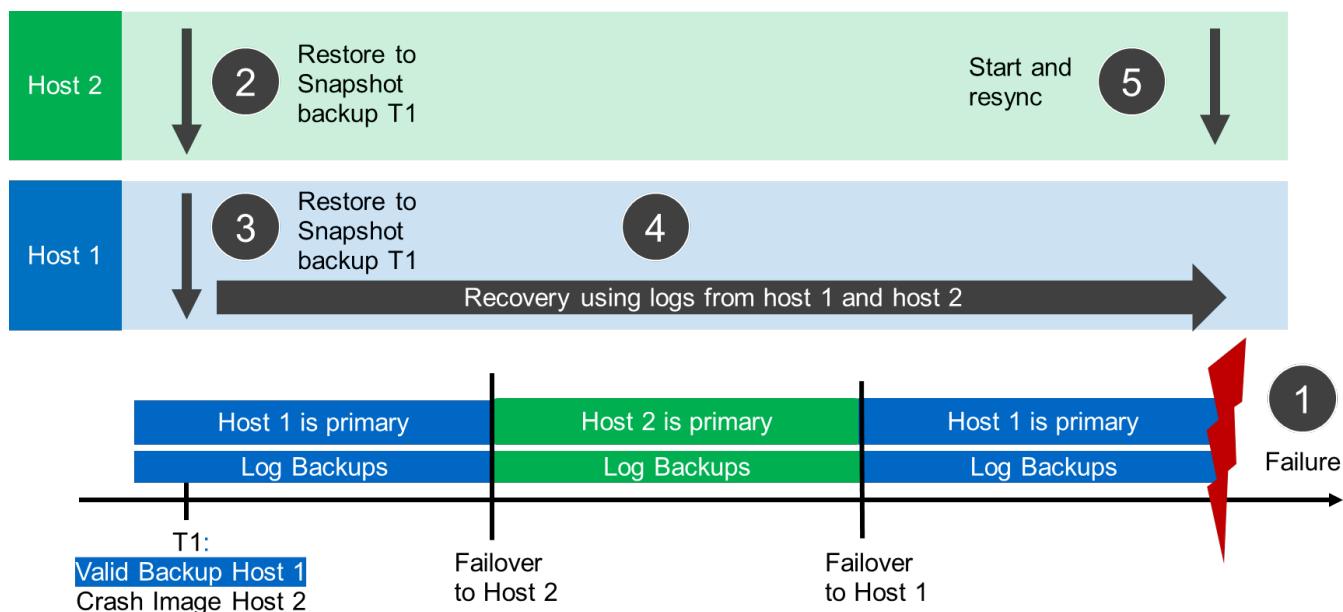
Even though the secondary host is up-to-date (no restore operation was performed for host 2), SAP HANA executes a full replication of all data. This behavior is standard after a restore and recovery operation with SAP HANA System Replication.

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.

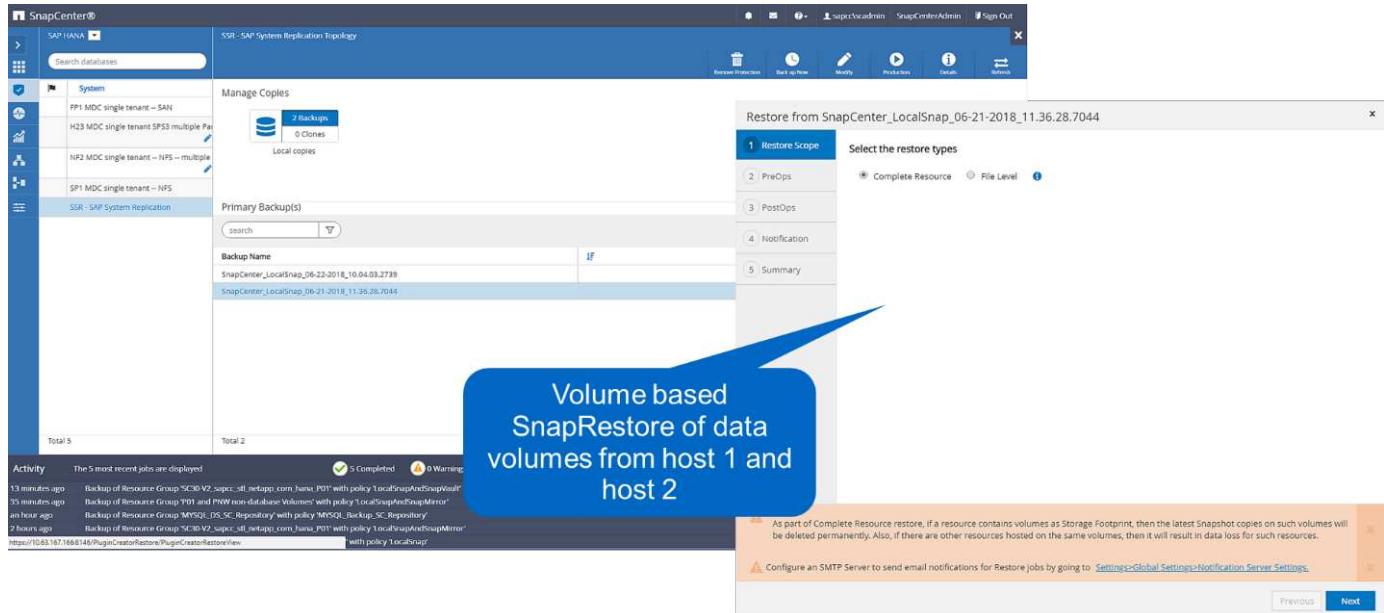
1. A failure occurred and you must restore to the backup created at T1 at host 1.
2. The secondary host (host 2) is shut down and the T1 crash image is restored.
3. The storage volume of host 1 is restored to the backup created at T1.
4. A forward recovery is performed with logs from host 1 and host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



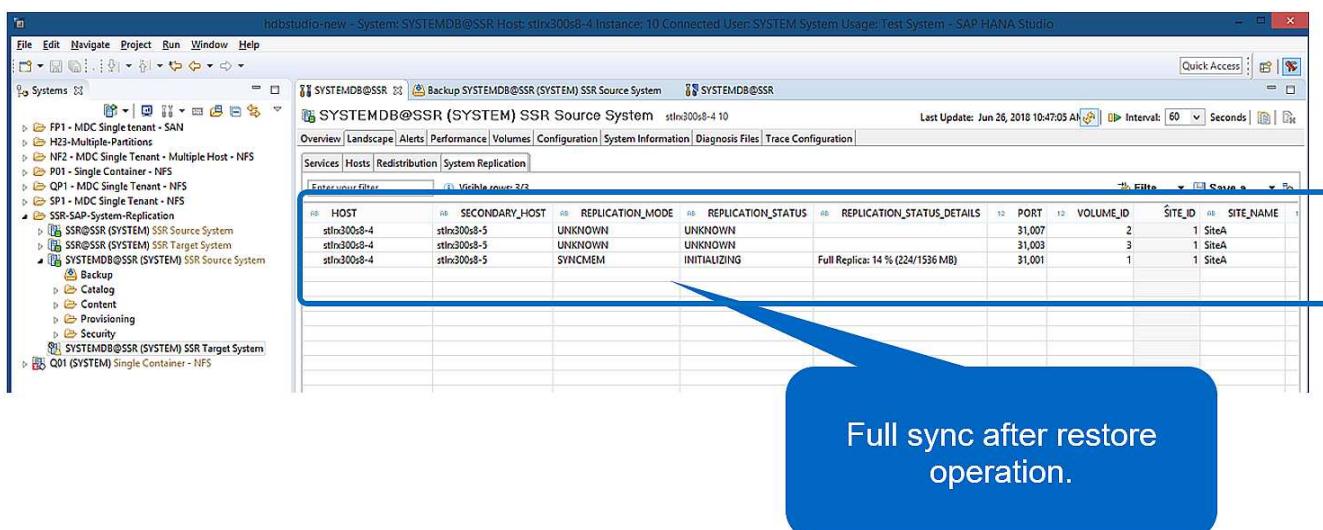
The restore and recovery operation with SAP HANA Studio is identical to the steps described in the section

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.



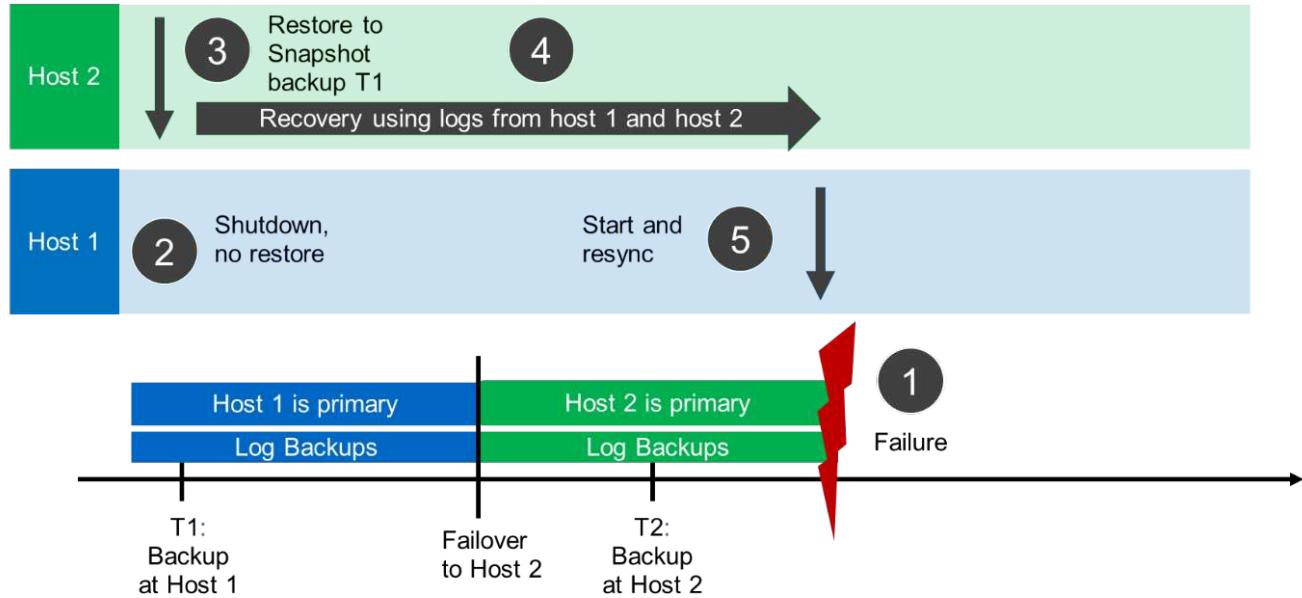
Restore and recovery from a backup created at the other host

A restore 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 screenshot shows the SAP HANA Studio interface with the following details:

Backup Catalog

Status	Started	Duration	Size	Backup Type	Destination
Success	Jun 28, 2018 9:23:46 ...	00h 00m 07s	1.53 GB	Data Backup	File
Success	Jun 27, 2018 7:45:56 ...	00h 00m 03s	1.52 GB	Data Backup	Snapshot
Success	Jun 27, 2018 7:12:37 ...	00h 00m 06s	1.55 GB	Data Backup	Snapshot

Backup Details

ID:	153097957115
Status:	Successful
Backup Type:	Data Backup
Destination Type:	Snapshot
Started:	Jun 27, 2018 7:12:37 AM (America/New_York)
Finished:	Jun 27, 2018 7:12:43 AM (America/New_York)
Duration:	00h 00m 06s
Size:	1.55 GB
Throughput:	n.a.
System ID:	SSR
Comment:	Snapshot LocalSnap_06-27-2018_07.12.29.1232
Additional Information:	<ok>
Location:	/hana/data/SSR/mnt00001/

Restored Volumes

Host	Service	Size	Name	Source Type	EBID
stln3008-4	nameserver	1.55 GB	hdb00001	volume	Snapshot

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.

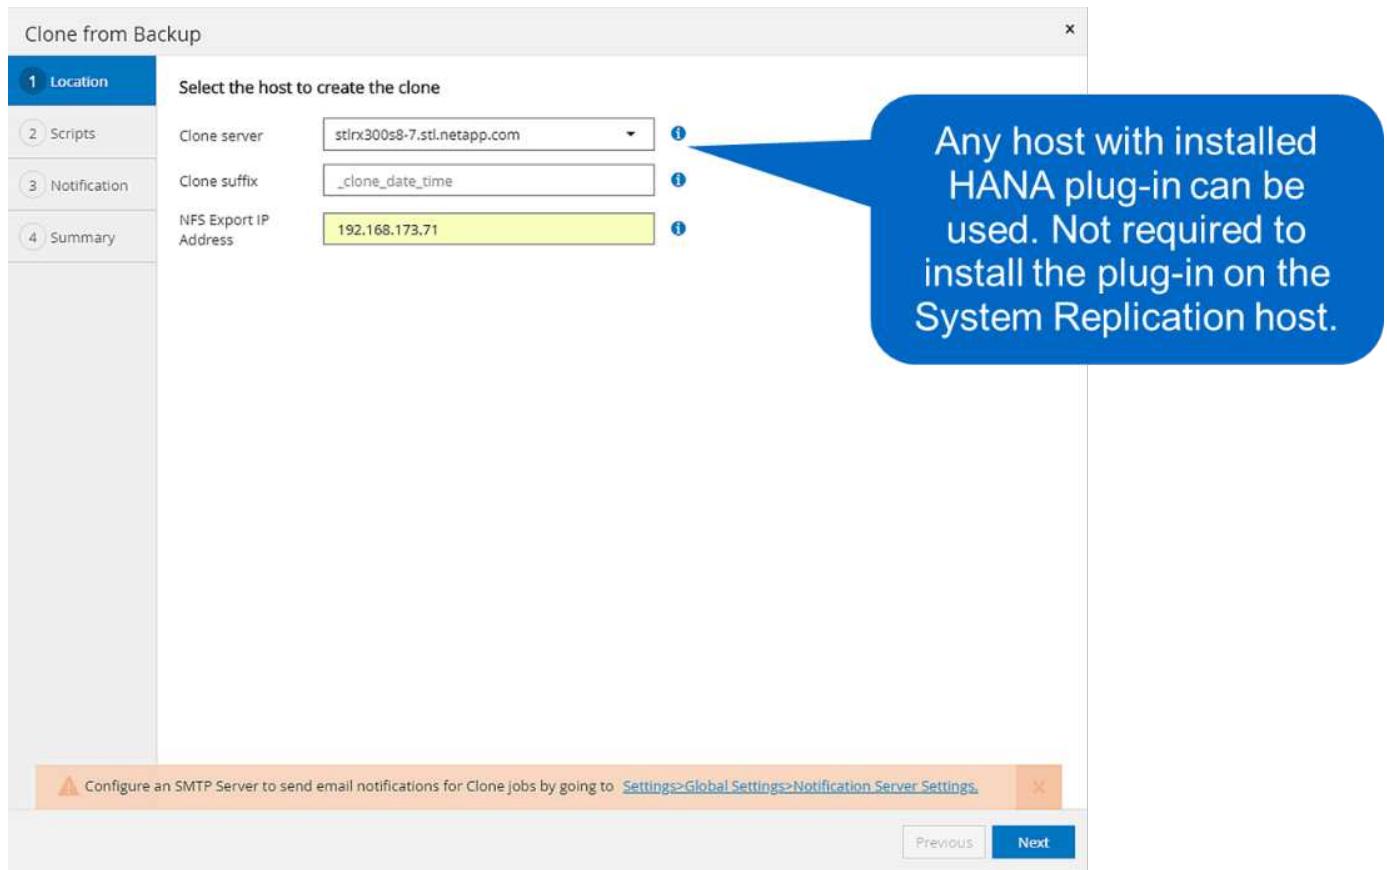
In SnapCenter, the backup is selected and the clone operation is started.

Backup Name	End Date
SnapCenter_Local5Snap_06-27-2018_07.12.29.123	6/27/2018 7:12:49 AM

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

i 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:

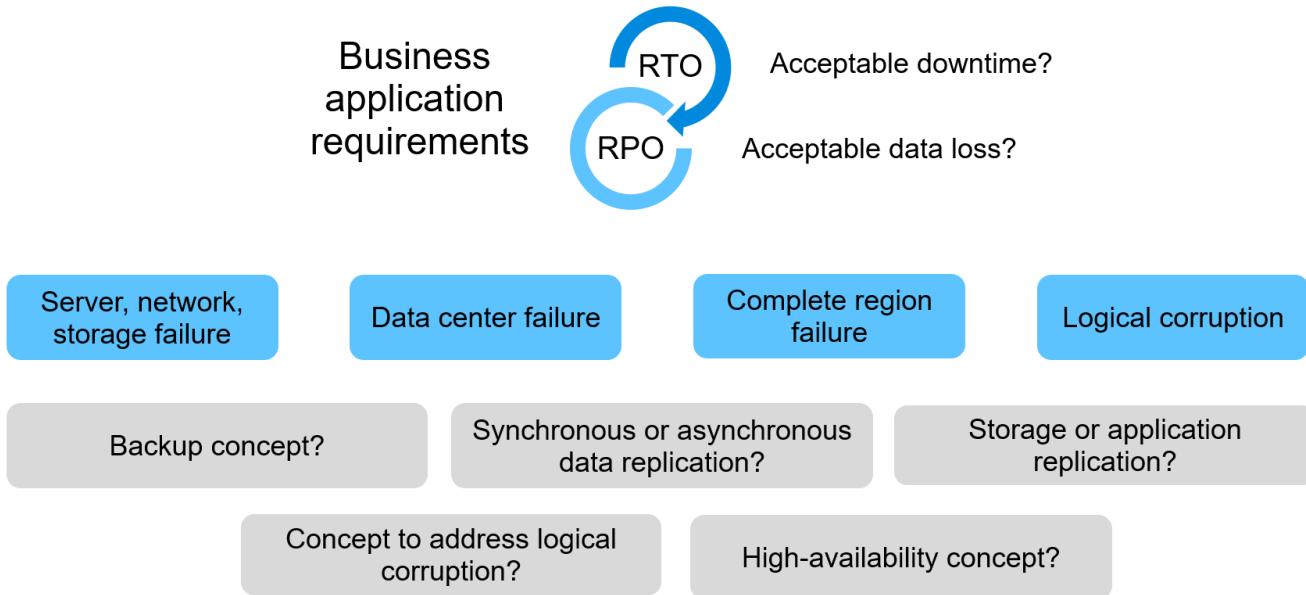
Nils Bauer, NetApp

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.

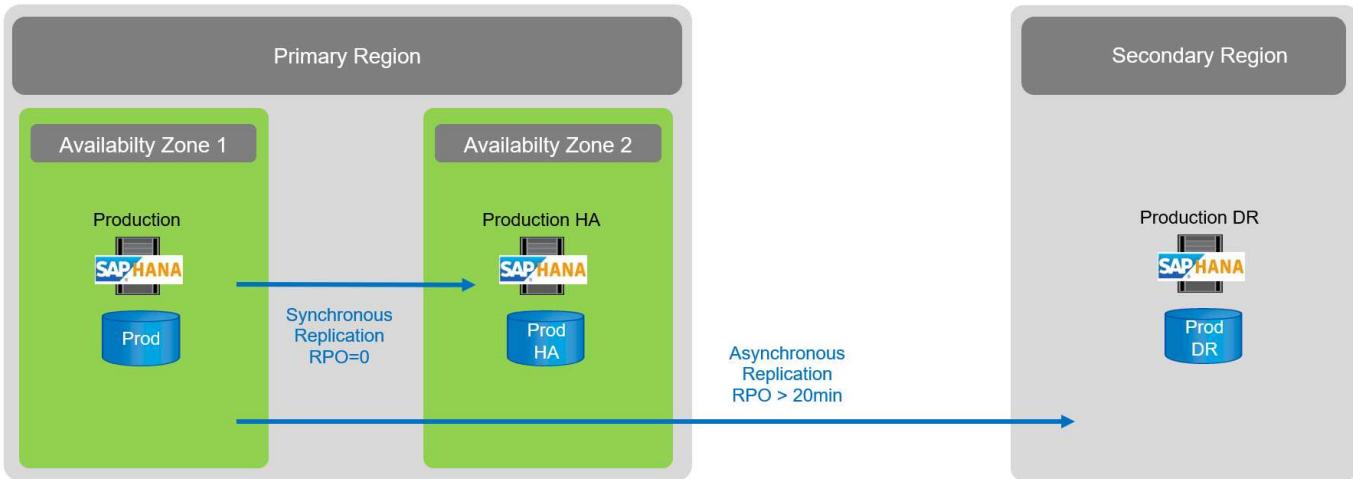
Disaster recovery solution comparison

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.

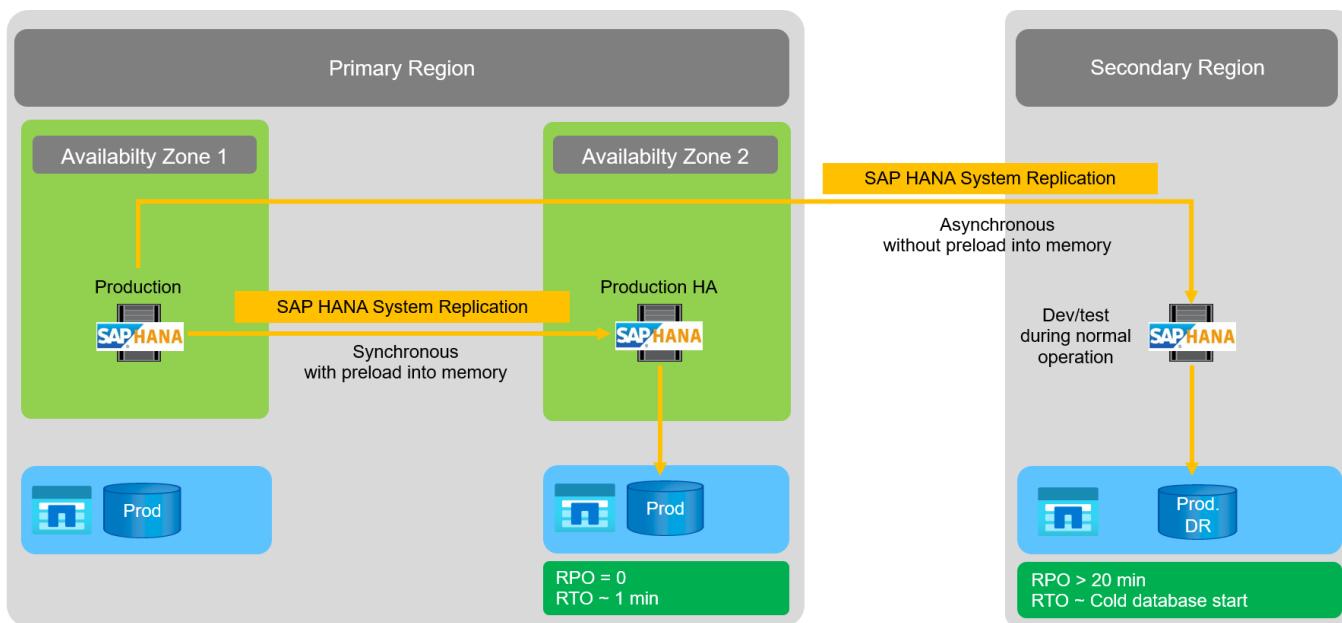
SAP HANA System Replication can be operated in one of two modes:

- 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

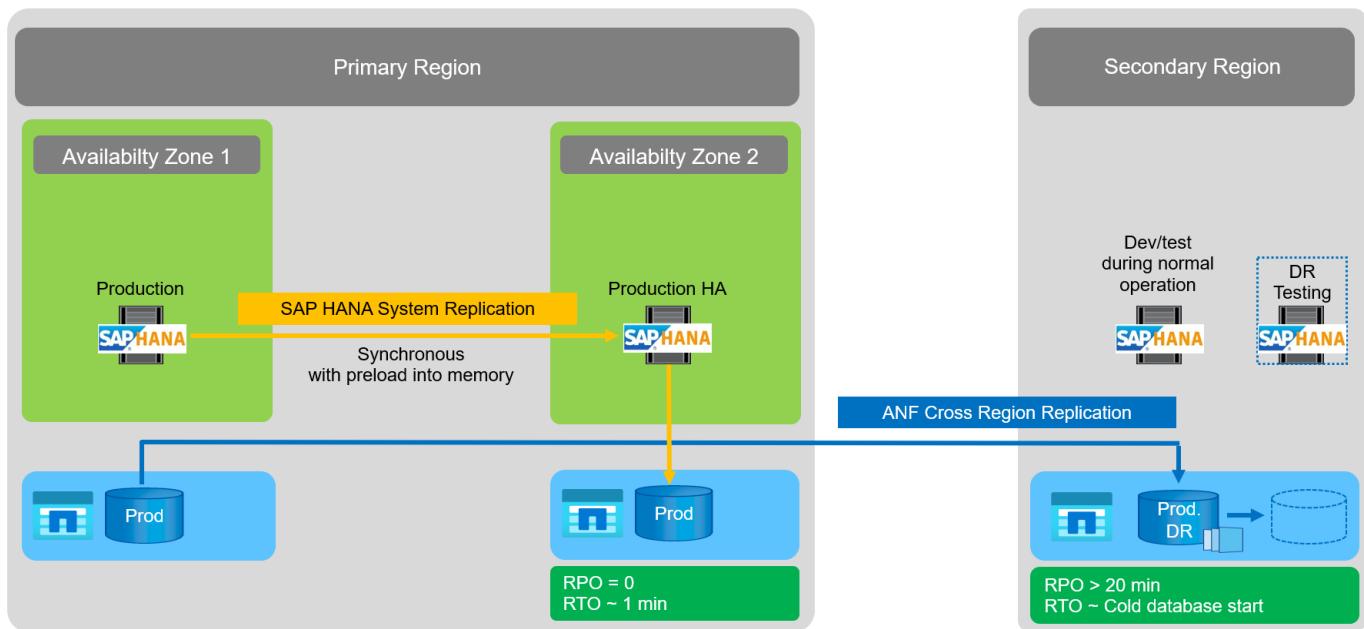
cross region replication. The RPO is defined by the replication schedule and replication options.

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.

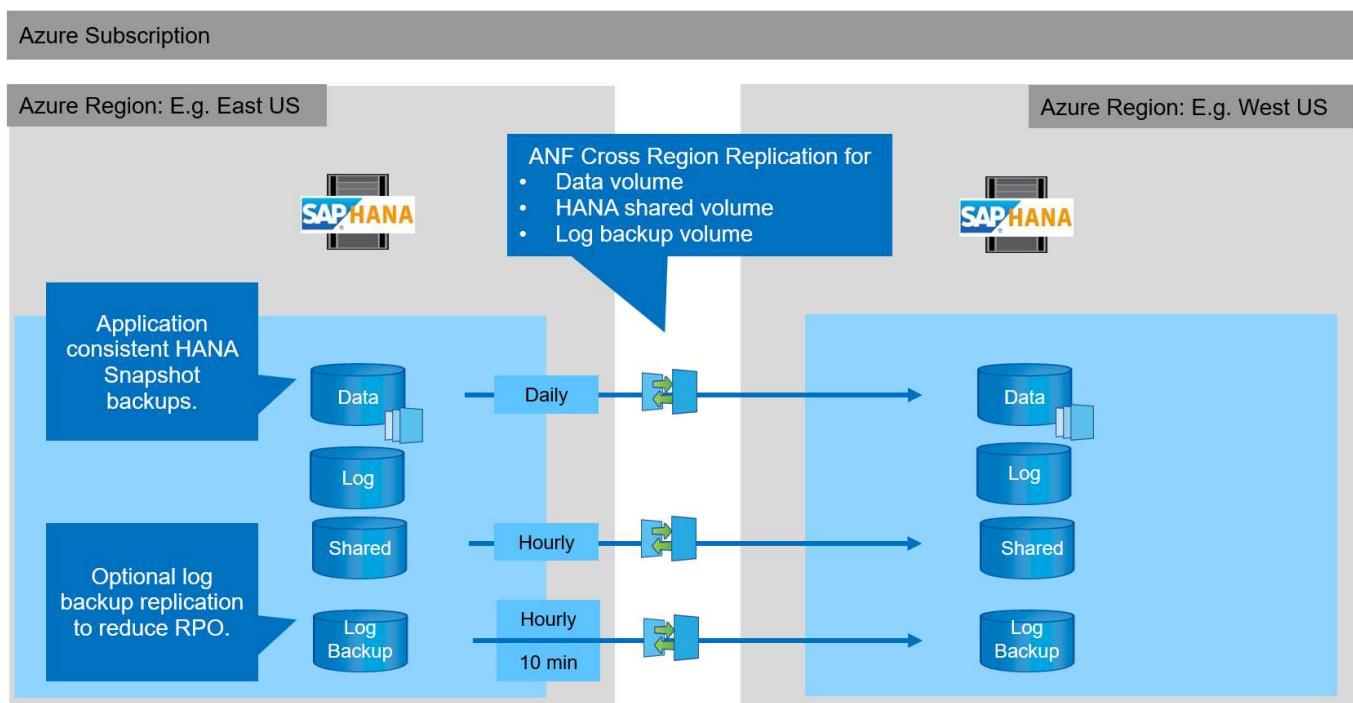


The following description and the lab setup focus on the HANA database. Other shared files, for example the SAP transport directory would be protected and replicated in the same way as the HANA shared volume.

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

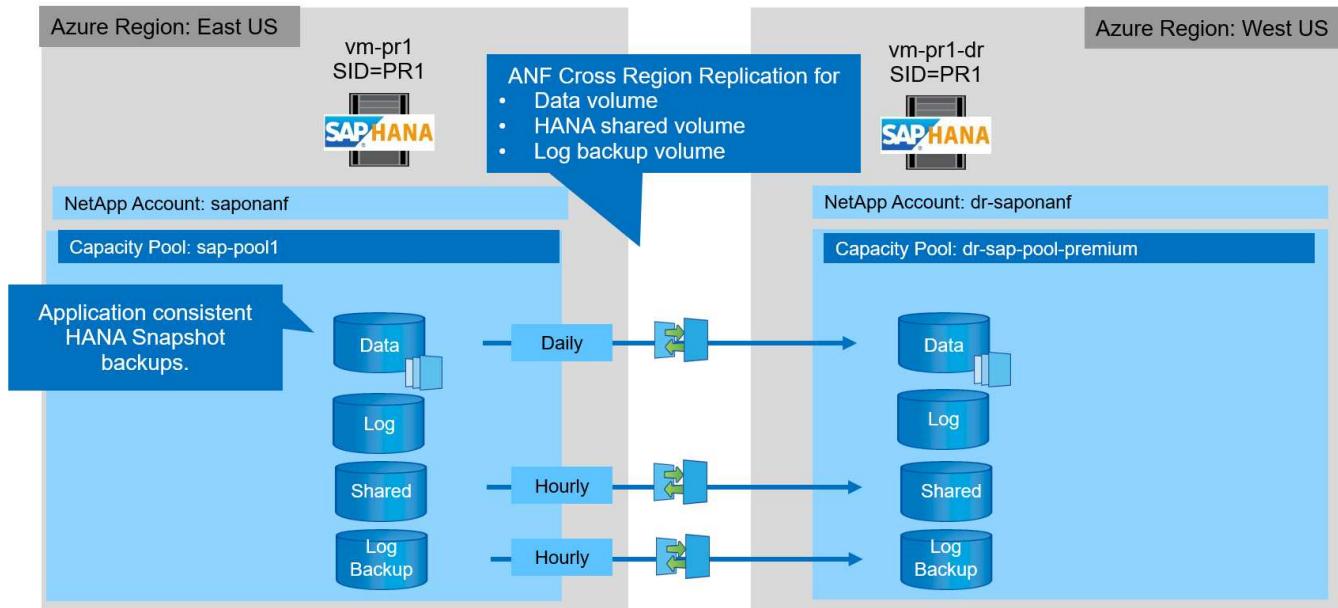
Solution validation has been performed with an SAP HANA single-host system. The Microsoft AzAcSnap Snapshot backup tool for ANF has been used to configure HANA application-consistent Snapshot backups. A daily data volume, hourly log backup, and shared volume replication were all configured. Disaster 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

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

Search snapshots

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.

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.

i 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 *	PR1-data-mnt00001-sm-dest	
Capacity pool *	dr-sap-pool1	
Available quota (GiB) 	4096	4 TiB
Quota (GiB) * 	500	 500 GiB
Virtual network *	dr-vnet (10.2.0.0/16,10.0.2.0/24)	
	Create new	
Delegated subnet *	default (10.0.2.0/28)	
	Create new	
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](#)

<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 & Write"/> <input type="button" value="▼"/>	<input type="text" value="On"/> <input type="button" value="▼"/>	<input type="button" value="..."/>

[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 ⓘ

/subscriptions/28cf403-f3f6-4b07-9847-4eb16109e870/resourceGroups/rg...✓

Replication schedule ⓘ

Daily

Every 10 minutes

Hourly

Daily

Review + create

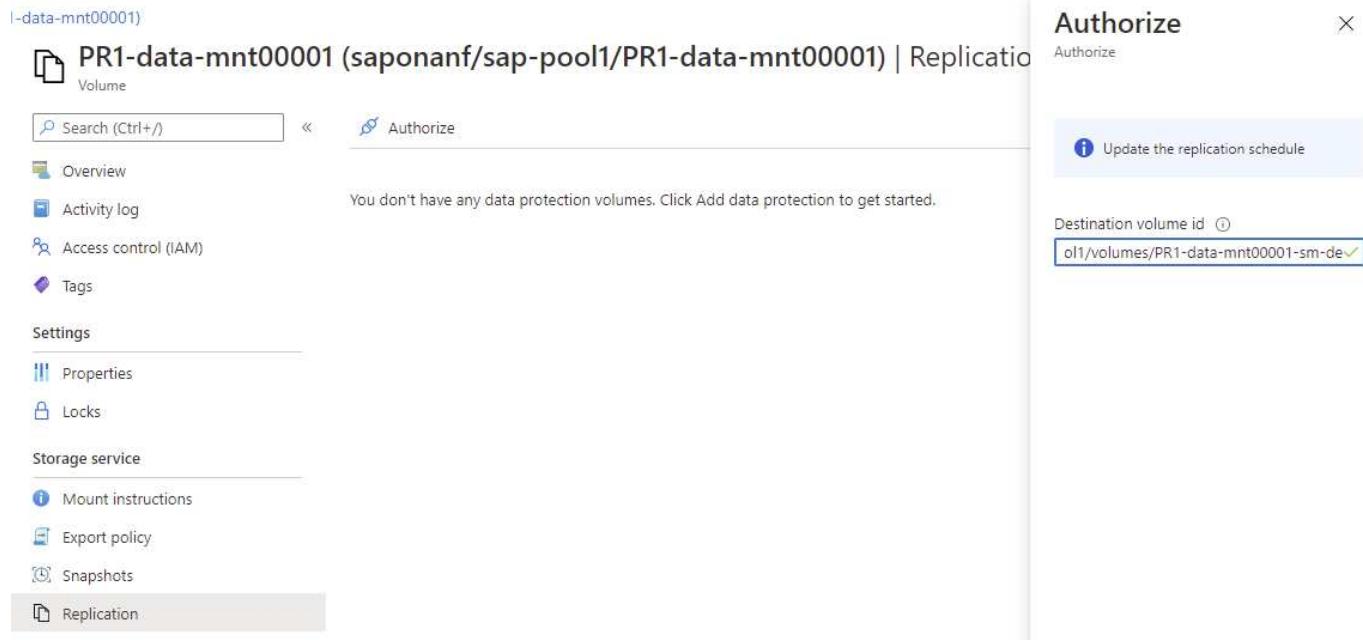
< Previous

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 screenshot shows the SAP HANA Cloud Platform Volume Management interface. The left sidebar shows navigation options: Overview, Activity log, Access control (IAM), Tags, Properties, Locks, Mount instructions, Export policy, Snapshots, and Replication. The Replication option is selected and highlighted in grey. The main content area shows a message: 'You don't have any data protection volumes. Click Add data protection to get started.' On the right, an 'Authorize' dialog is open with the title 'Authorize'. It contains a button 'Authorize' and a link 'Update the replication schedule'. Below the dialog, a text input field shows the destination volume ID: 'ol1/volumes/PR1-data-mnt0001-sm-de' with a checkmark.

The same steps must be performed for the HANA shared and the log backup volume.

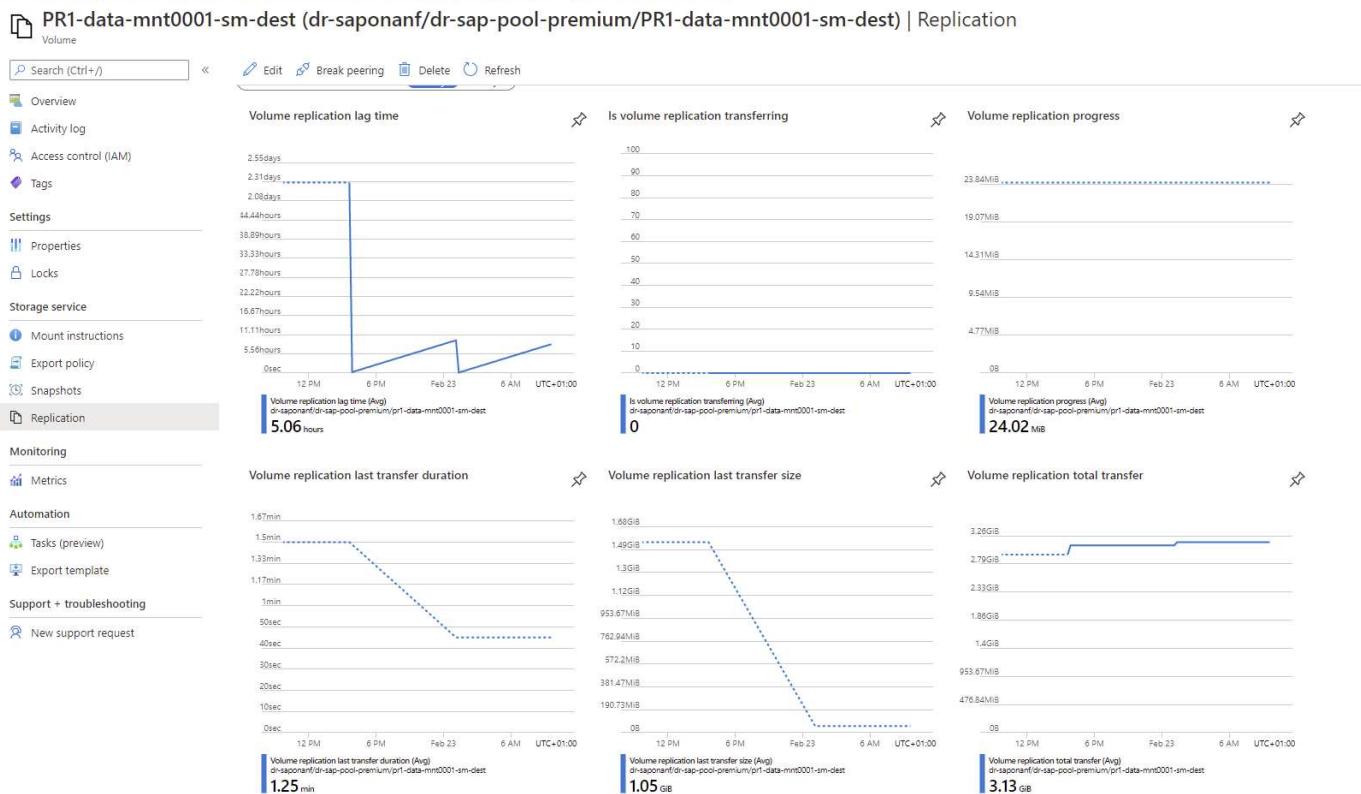
Monitoring ANF Cross-Region Replication

The following three screenshots show the replication status for the data, log backup, and shared volumes.

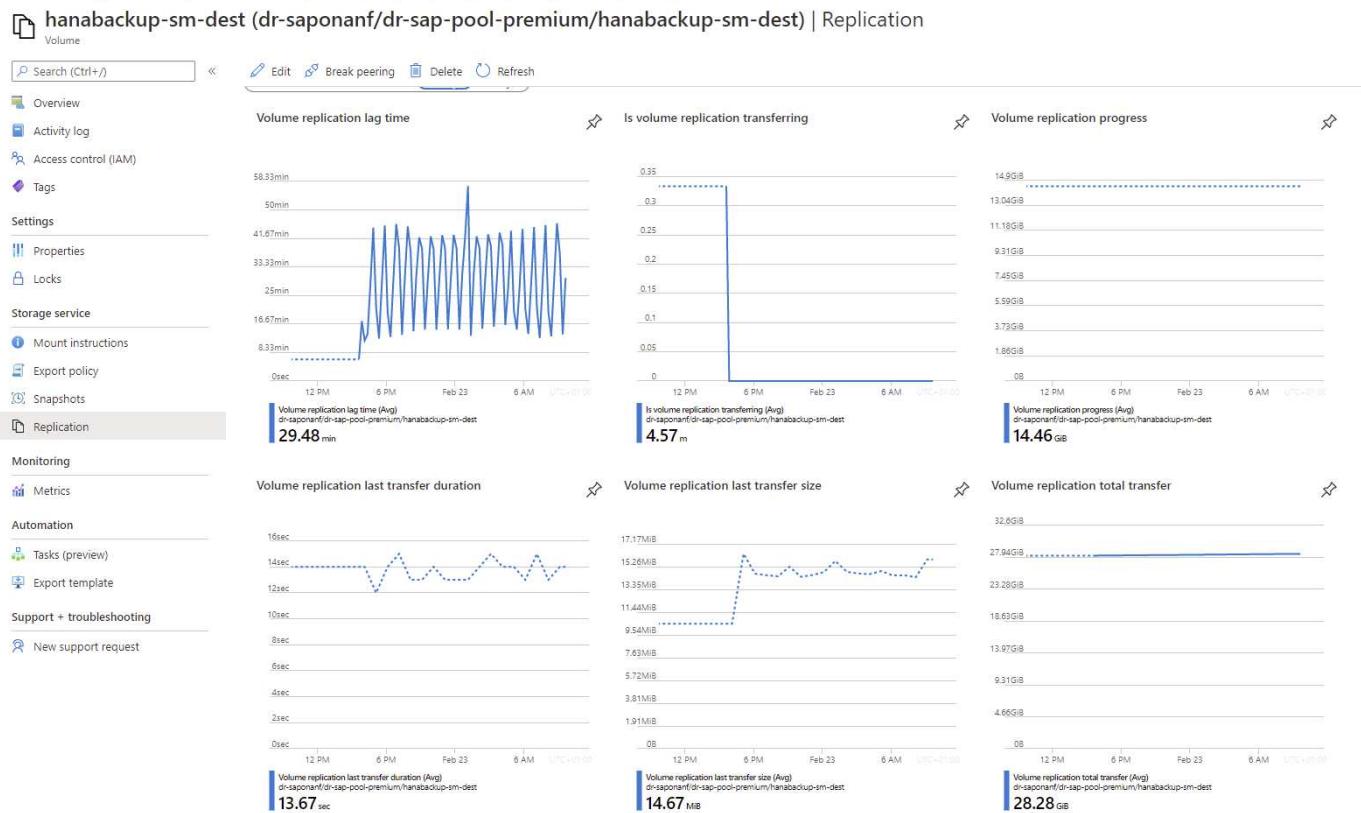
The volume replication lag time is a useful value to understand RPO expectations. For example, the log backup volume replication shows a maximum lag time of 58 minutes, which means that the maximum RPO has the same value.

The transfer duration and transfer size provide valuable information on bandwidth requirements and change the rate of the replicated volume.

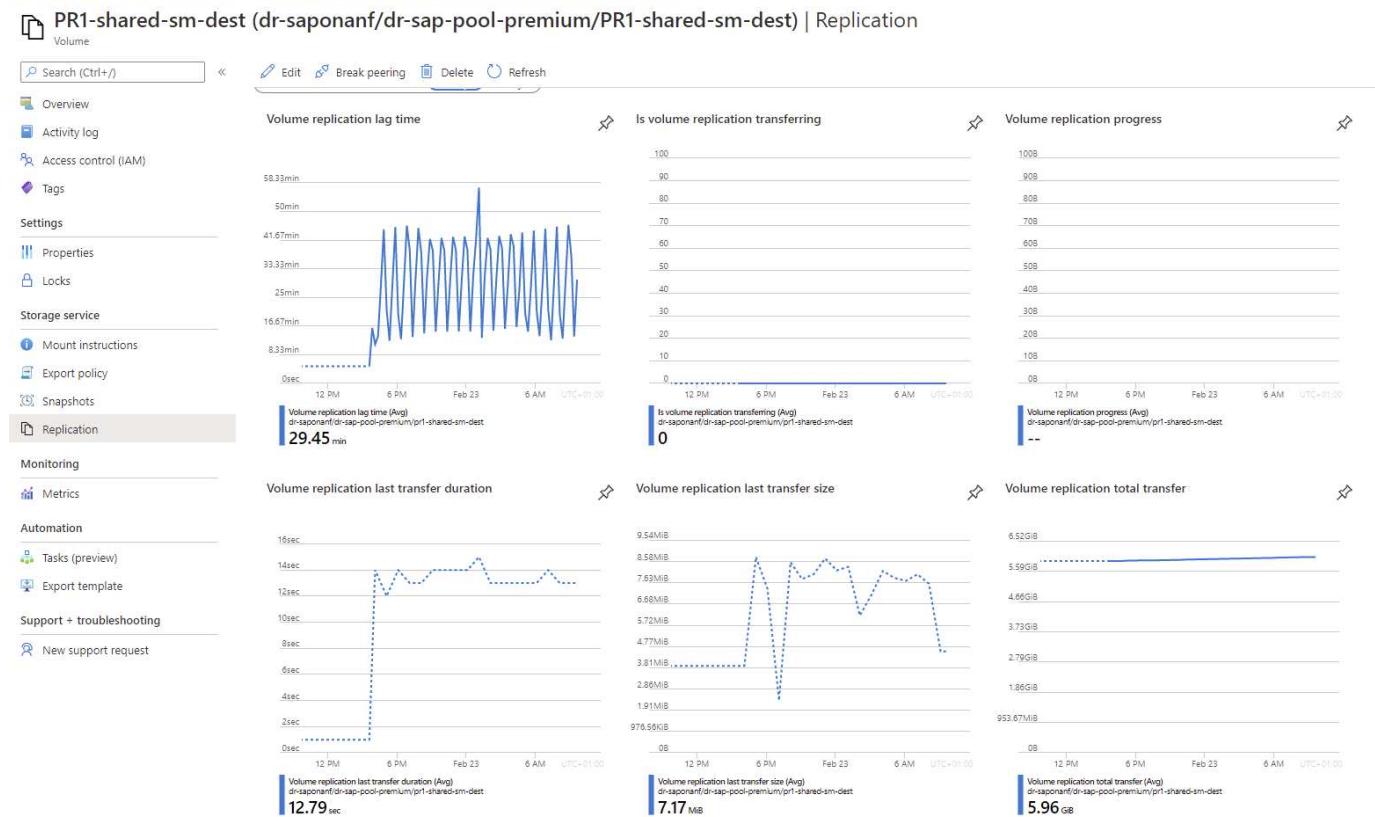
The following screenshot shows the replication status of HANA data volume.



The following screenshot shows the replication status of HANA log backup volume.



The following screenshot shows the replication status of HANA shared volume.



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+F) < + 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.b1e048d-7114-11eb-b147-d039ea1e211e_2155791247.2021-02-22_143159	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_2021-02-23T000002-0313792Z	West US	02/23/2021, 01:00:05 AM	...
snapmirror.b1e048d-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](#).

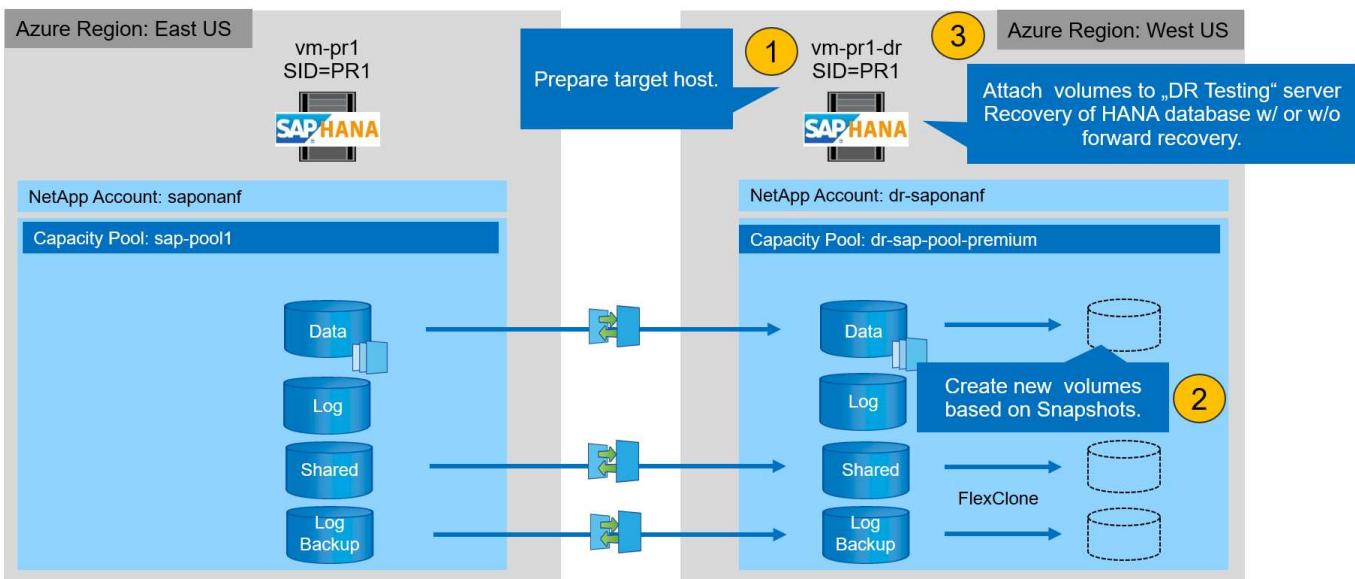
Depending on whether log backup replication is part of the disaster recovery setup or not, the steps for disaster recovery are slightly different. This section describes the disaster recovery testing for data-backup-only replication as well as for data volume replication combined with log backup volume 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.

Azure Subscription



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-mnt0001-sm-dest	/hana/data/PR1/mnt0001
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-mnt0001-sm-dest /hana/data/PR1/mnt0001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noatime,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-log-mnt0001-dr /hana/log/PR1/mnt0001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noatime,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,wsize=262144,intr,noatime,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,wsize=262144,intr,noatime,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,wsize=262144,nconnect=8,bg,noatime,nolock 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)

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 PM	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 01:10:00 PM	...

Restore to new volume

Revert volume

Delete

2. The new volume name and quota must be provided in the user interface.

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

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 *

PR1-data-mnt00001-sm-dest-clone



Restoring from snapshot ⓘ

azacsnap__2021-02-18T000001-7955243Z

Available quota (GiB) ⓘ

2096

2.05 TiB

Quota (GiB) * ⓘ

500



500 GiB

Virtual network ⓘ

dr-vnet (10.2.0.0/16,10.0.2.0/24)



Delegated subnet ⓘ

default (10.0.2.0/28)



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 * [\(i\)](#)

PR1-data-mnt00001-sm-dest-clone

Versions

NFSv4.1

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	0.0.0.0/0	Read & Write	On

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

NetApp account

Search (Ctrl+ /) < Add volume Add data replication Refresh

Overview Activity log Access control (IAM) Tags Properties Locks

Azure NetApp Files Active Directory connections

Storage service Capacity pools **Volumes**

Data protection Snapshot policies

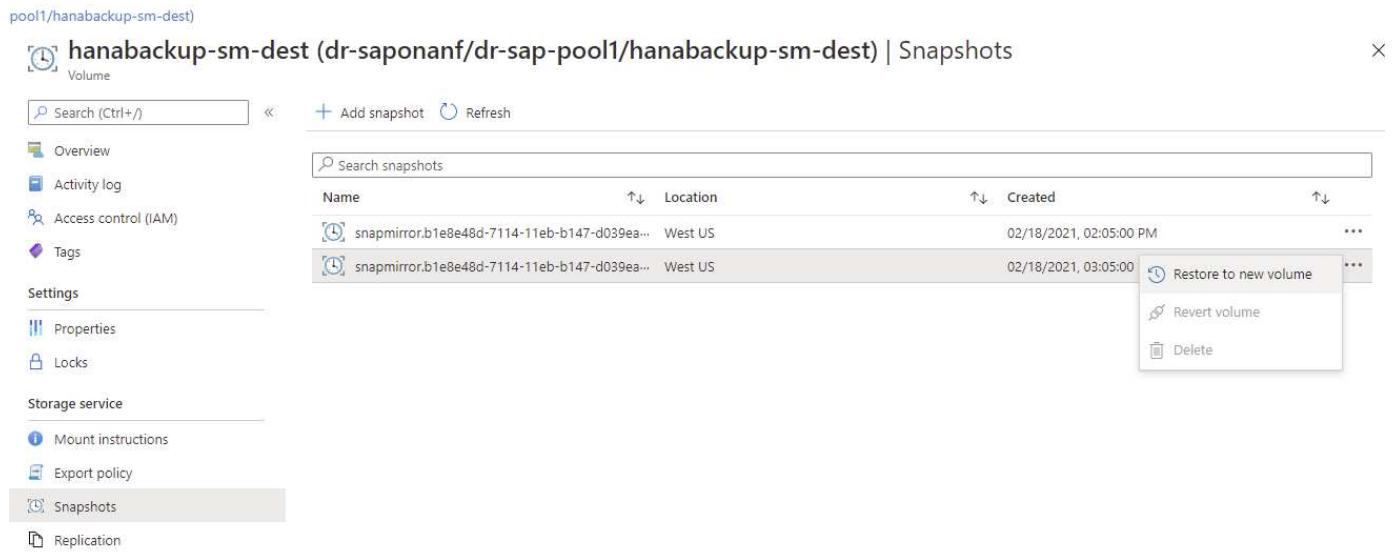
Automation Tasks (preview) Export template

Support + troubleshooting New support request

Search volumes

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-sm-dest	Standard	dr-sap-pool1
PR1-data-mnt00001-sm-dest-clone	500 GiB	NFSv4.1	10.0.2.4:/PR1-data-mnt00001-sm-dest-clone	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.



pool1/hanabackup-sm-dest

hanabackup-sm-dest (dr-saponanf/dr-sap-pool1/hanabackup-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

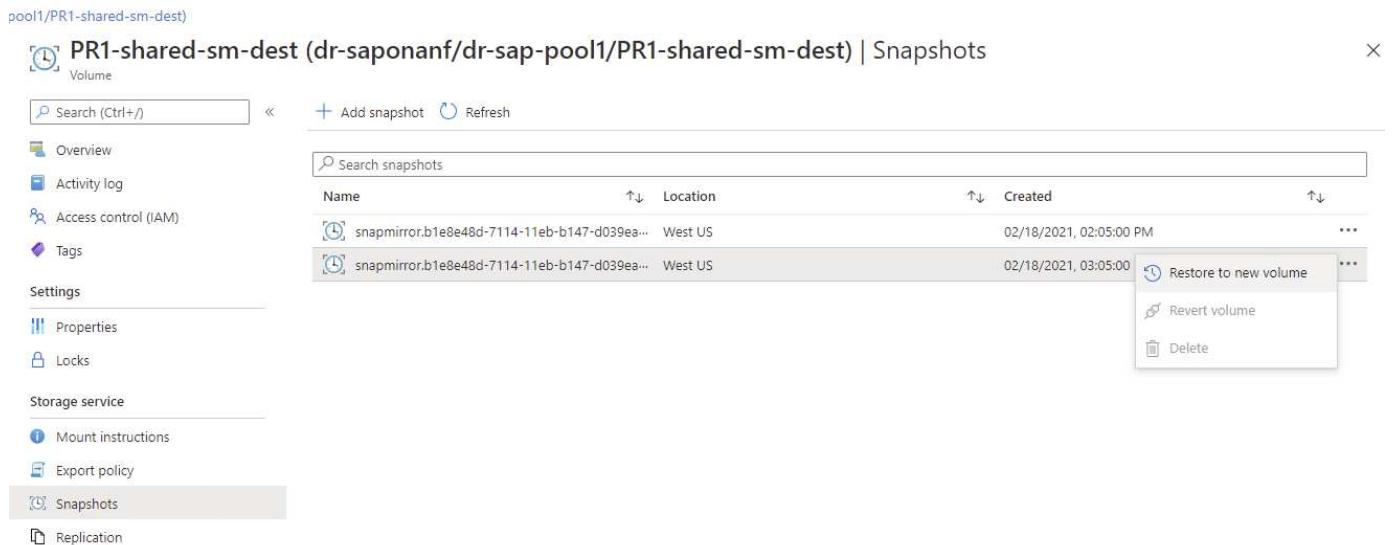
Name	Location	Created	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 02:05:00	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 03:05:00	...

Restore to new volume

Revert volume

Delete

The following screenshot shows the HANA shared volume restored to new volume.



pool1/PR1-shared-sm-dest

PR1-shared-sm-dest (dr-saponanf/dr-sap-pool1/PR1-shared-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

Name	Location	Created	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 02:05:00	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 03:05:00	...

Restore to new volume

Revert volume

Delete



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

```

pr1adm@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

```

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 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 'libhdbcaccessor'
Loaded library 'libhdblivecache'
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

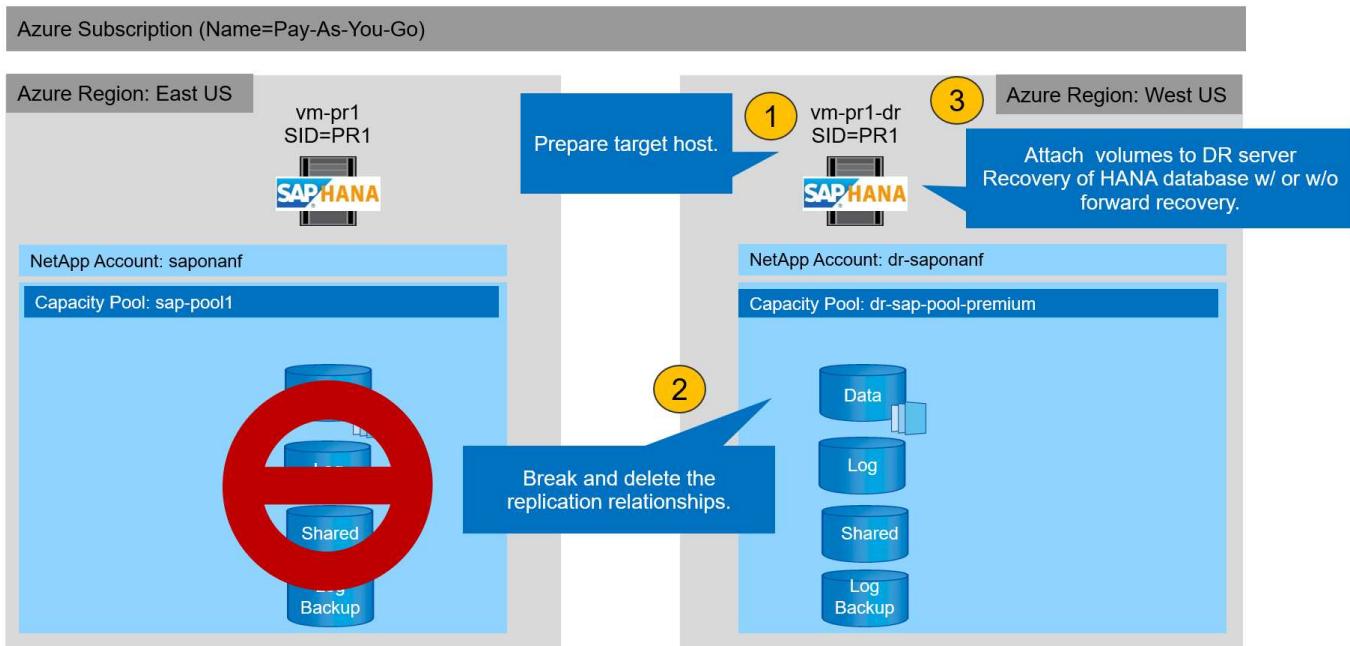
Disaster recovery failover

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-mnt0001-sm-dest	/hana/data/PR1/mnt0001
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-mnt0001-sm-dest /hana/data/PR1/mnt0001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noatime,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-log-mnt0001-dr /hana/log/PR1/mnt0001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noatime,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,wsize=262144,intr,noatime,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,wsize=262144,intr,noatime,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,wsize=262144,nconnect=8,bg,noatime,nolock 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.

PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest)

Volume

Search (Ctrl+ /) Edit Break peering Delete Refresh

Essentials

End point type : Destination Source

Healthy : Healthy Relationship state

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

9.72 hours 8.33 hours 6.94 hours 5.56 hours

Is volume replication transfer

100 90 80 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

X

PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest)

Volume

Search (Ctrl+ /) Resync Delete Refresh

Essentials

End point type : Destination Source

Healthy : Healthy Relationship state

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

1.67min 1.5min 1.33min 1.17min 1min 50sec

Is volume replication transfer

100 90 80 70 60 50

Delete replication

Delete replication object

⚠ Warning this operation will delete the connection between PR1-data-mnt0001 and PR1-data-mnt0001-sm-dest

This will delete the replication object of PR1-data-mnt0001, type 'yes' to proceed

yes

X

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 X

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_2021-02-18T120002-2150721Z	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	...
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	...
azacsnap_2021-02-23T040001-7114-11eb-b147-d039ea...	West US	02/23/2021, 01:10:00 AM	...

Restore to new volume Revert volume Delete

PR1-data-mnt0001-sm-dest (dr-saponanf/dr-sap-pool-premium/PR1-data-mnt0001-sm-dest) X

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
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
azacsnap_2021-02-22T160002-0144647Z	West US
azacsnap_2021-02-22T200002-0649581Z	West US
azacsnap_2021-02-23T000002-0311379Z	West US
azacsnap_2021-02-23T040001-7114-11eb-b147-d039ea...	West US

Revert volume to snapshot X

Revert volume PR1-data-mnt0001-sm-dest to snapshot azacsnap_2021-02-23T000002-0311379Z

⚠ 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

```

pr1adm@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

```

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 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 'libhdbcaccessor'
Loaded library 'libhdblivecache'
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

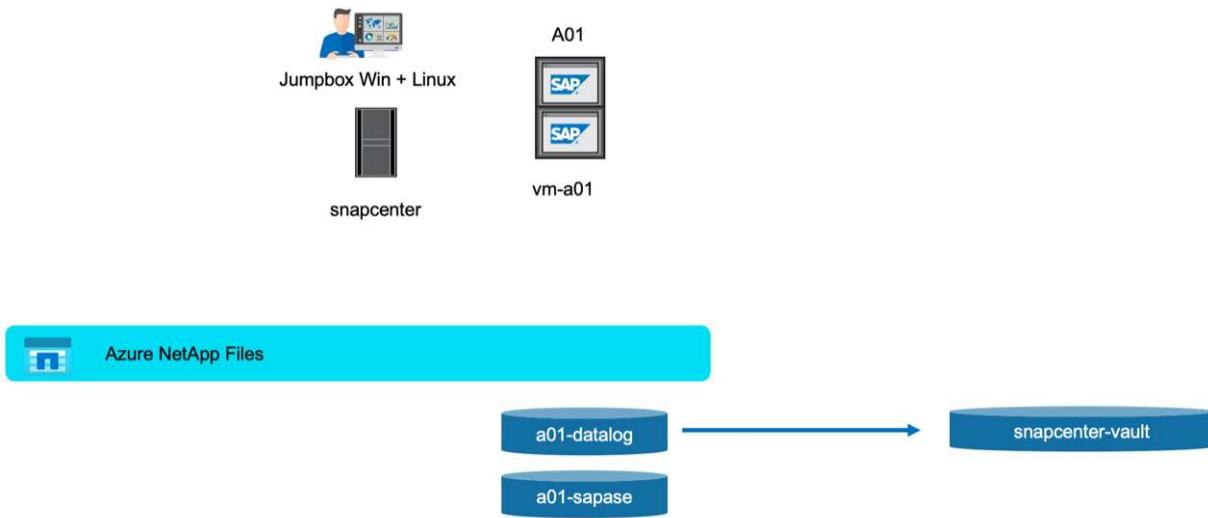
Example Implementation for SnapCenter ASE Plugin for an SAP System running on the Azure Platform.



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
	sybsystem	/sybase/<SID>/sybsystem	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

+ **i**

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 i |
| <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 i |
| <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 X

1 Name

Provide Resource Details

Name	A01	i
Host Name	vm-a01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net	i
Type	Instance	i
Credential Name	None	+ i

i Add information for the credential

Credential Name	sapsa-A01
Username	sapsap
Password

Add

Previous Next



If Password contains Special Characters, they must be masked with a backslash.
E.g. Test!123! → Test!\123\!

Add SAP ASE Resource X

1 Name

Provide Resource Details

Name	A01	i
Host Name	vm-a01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net	i
Type	Instance	i
Credential Name	sapsa-A01	+ i

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		
------------	-------------	---	---

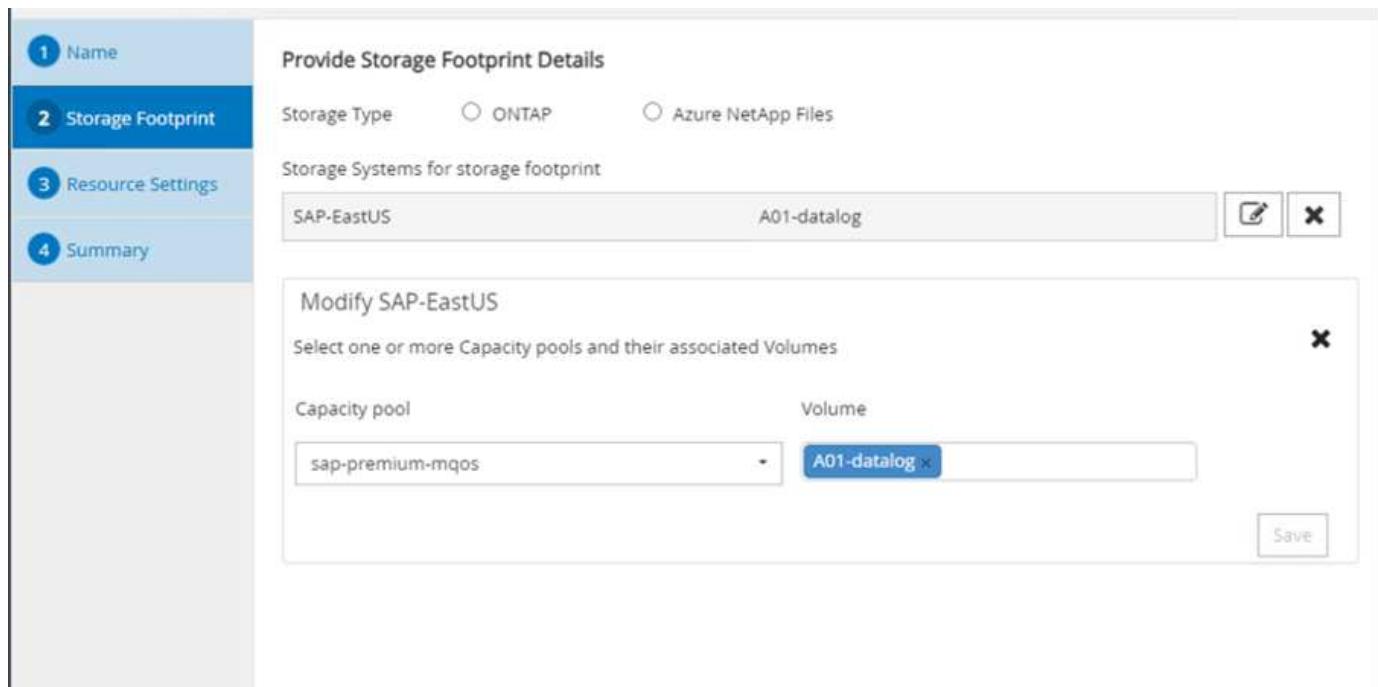
Modify SAP-EastUS

Select one or more Capacity pools and their associated Volumes

Capacity pool Volume

sap-premium-mqos	A01-datalog
------------------	-------------

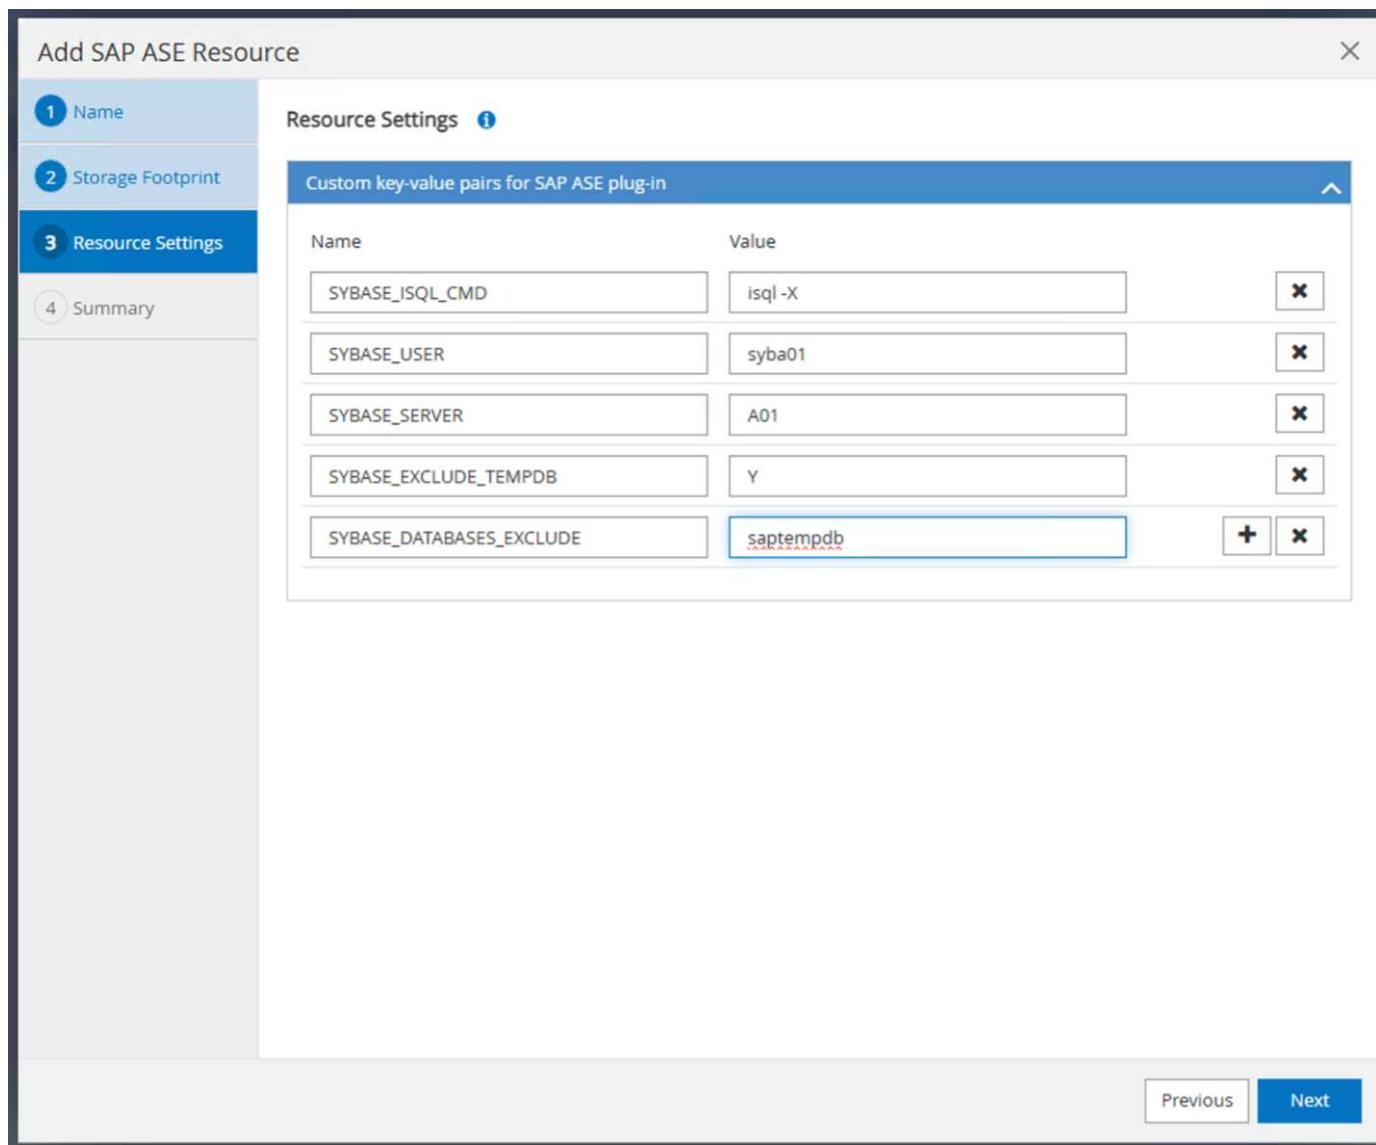




If you are using the volume design out of the [MS Technical Community Blog](#).

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).



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: https://infocenter.sybase.com/help/index.jsp?topic=/com.sybase.infocenter.dc34237.1500/html/mvsiinst/CIHHFDGC.htm
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 Agent start 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, sybsystemdb, sybsystemprocs, 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_E_XCLUDE	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_TEMPORARY_DATABASES	(Y / N)	Enables automatic exclusion of user-created temporary databases.

Sequence to Recover System A01

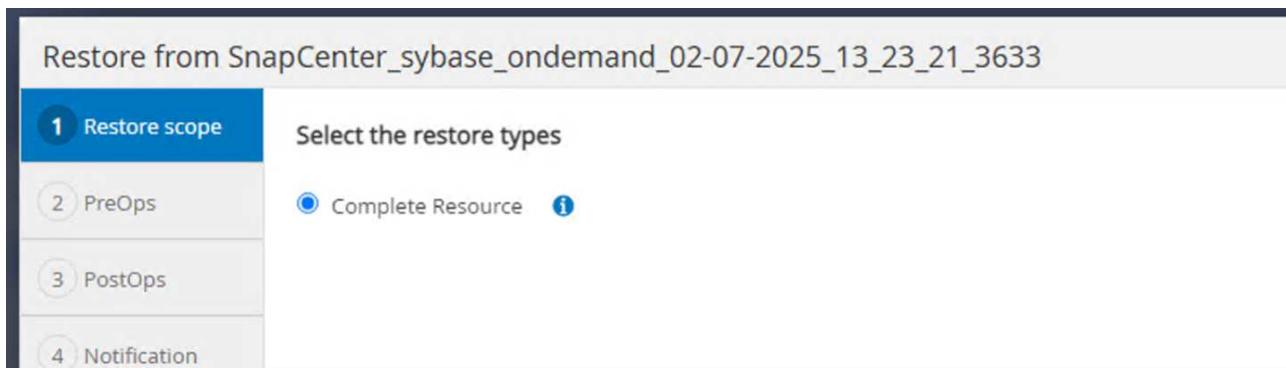
1. stop SAP System A01 (including database), stop sapinit
2. umount Filesystems
3. restore Volumes A01-datalog (using SnapCenter)
4. mount Filesystems
5. start Database A01 (with option –q, to avoid automatic online and keep database forward recoverable – according to SAP Note 1887068)
6. start BackupServer 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

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_sybase_ondemand_02-07-2025_13_21_3633	02/07/2025 1:23:58 PM	1	
SnapCenter_sybase_daily_02-07-2025_11_08_28_9176	02/07/2025 11:09:07 AM	1	
SnapCenter_sybase_ondemand_02-07-2025_09_31_42_2639	02/07/2025 9:32:23 AM	1	
SnapCenter_sybase_daily_02-06-2025_16_35_19_5734	02/06/2025 4:36:32 PM	1	
SnapCenter_sybase_ondemand_02-06-2025_16_34_01_6115	02/06/2025 4:34:36 PM	1	
SnapCenter_sybase_ondemand_02-06-2025_15_41_33_6630	02/06/2025 3:42:21 PM	1	

- For ANF Deployment – only Complete Resource is available



i Selecting Complete Resource will trigger a Volume Based Snap Restore (VBSR). Within Azure it is called [volume revert](#).

① **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 X

1 Restore scope 2 PreOps 3 PostOps 4 Notification 5 Summary

Select the restore types

Complete Resource i

File Level i

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 X

1	Restore scope	Summary
2	PreOps	Backup Name: SnapCenter_sybase_onDemand_02-07-2025_13_23_21_3633
3	PostOps	Backup date: 02/07/2025 1:23:58 PM
4	Notification	Restore scope: Complete Resource
5	Summary	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

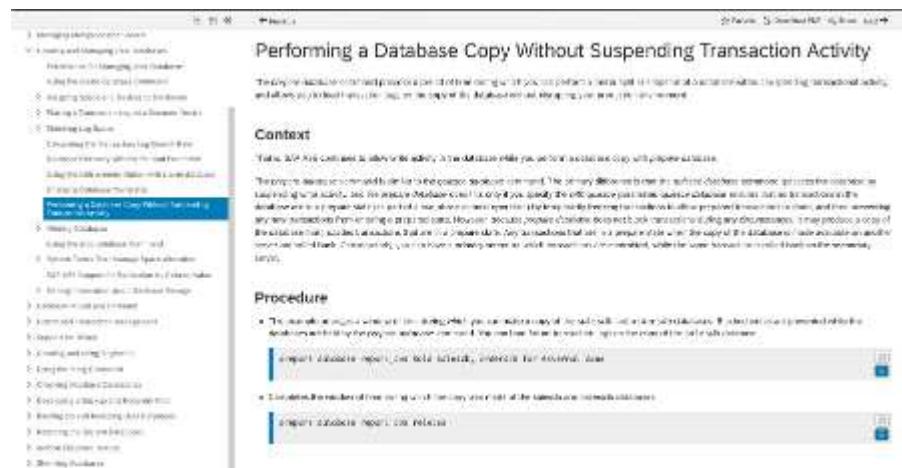
- Mount Filesystems (vm-a01)
 - User root: mount -a -t nfs
- Start Database A01 + BackupServer
 - Modify RUN_A01 and add -q \ (according to SAP Note 1887068)
 - User syba01: RUN_A01 &
 - User syba01: RUN_A01_BS&
- Online databases saptools, sybsecurity , sybmgmtdb
 - User syba01: isql -S A01 -U sapsa -X -w 1024
 - online database saptools
 - go
 - online database sybsecurity
 - go
 - online database sybmgmtdb
 - go

- recover Database A01
 - sp_dump_history (to show the transaction log dumps)
 - go
 - Load transaction log dumps according your needs – for more information see documentation:
<https://infocenter.sybase.com/help/index.jsp?topic=/com.sybase.infocenter.dc36272.1572/html/commands/X75212.htm>
 - Example:
LOAD TRAN A01 FROM '/sybase/A01/log_archives/A01.TRAN.20250207.140248.6.000'
 - go
 - online database A01
 - go
 - remove -q from RUN_A01
 - start SAP System
 - User root: /etc/init.d/sapinit start
 - User a01adm: startsap

Additional information and version history

Quiesce vs. prepare

See the documentation on [xref:./backup/ SAP help page](#).



SnapCenter SAP ASE Plugin uses the quiesce database command, however it could be replaced by the prepare command. If required, it must be changed in the SYBASE.pm in line 473, 475, 479, 481, 673, 675 e.g.

Recorded Demos

Following recoded Demos are available to support the documentation.

Installation and Configuration ASE Plugin, Backup of ASE 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](#)
- [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

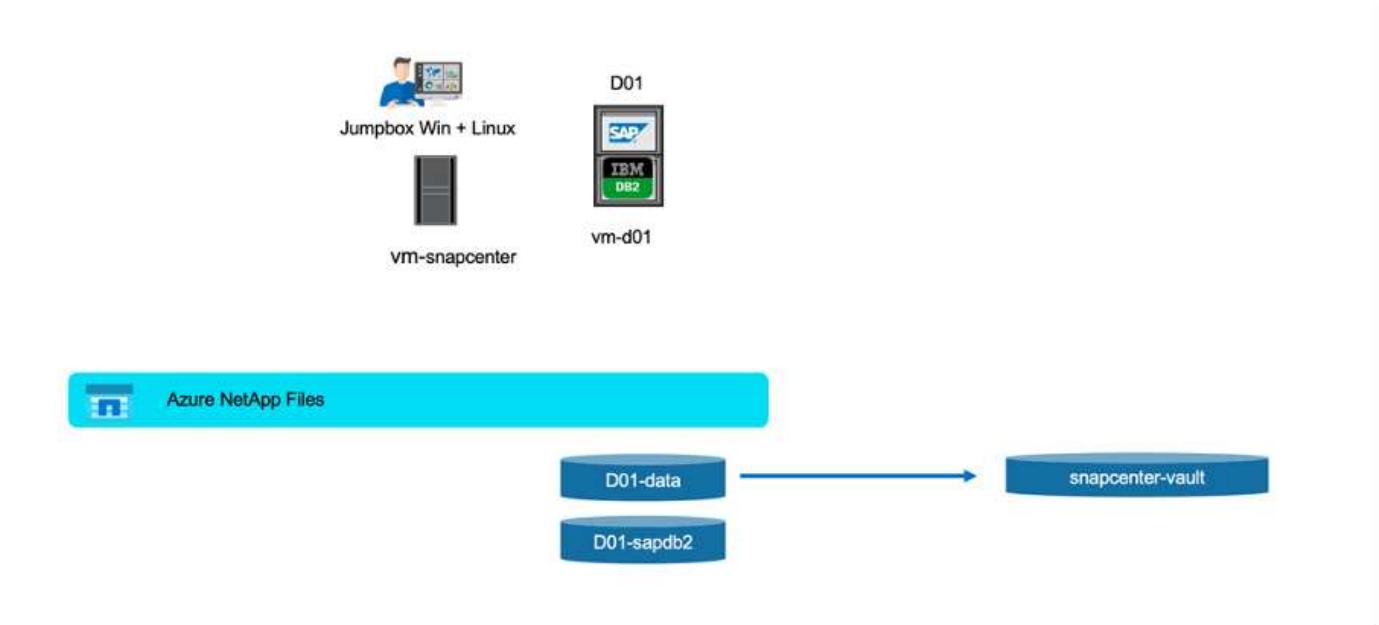
Example configuration overview

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-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: Linux

Host Name: vm-d01

Credentials: linux-snapcenter + i

Select Plug-ins to Install SnapCenter Plug-ins Package 6.1 for Linux

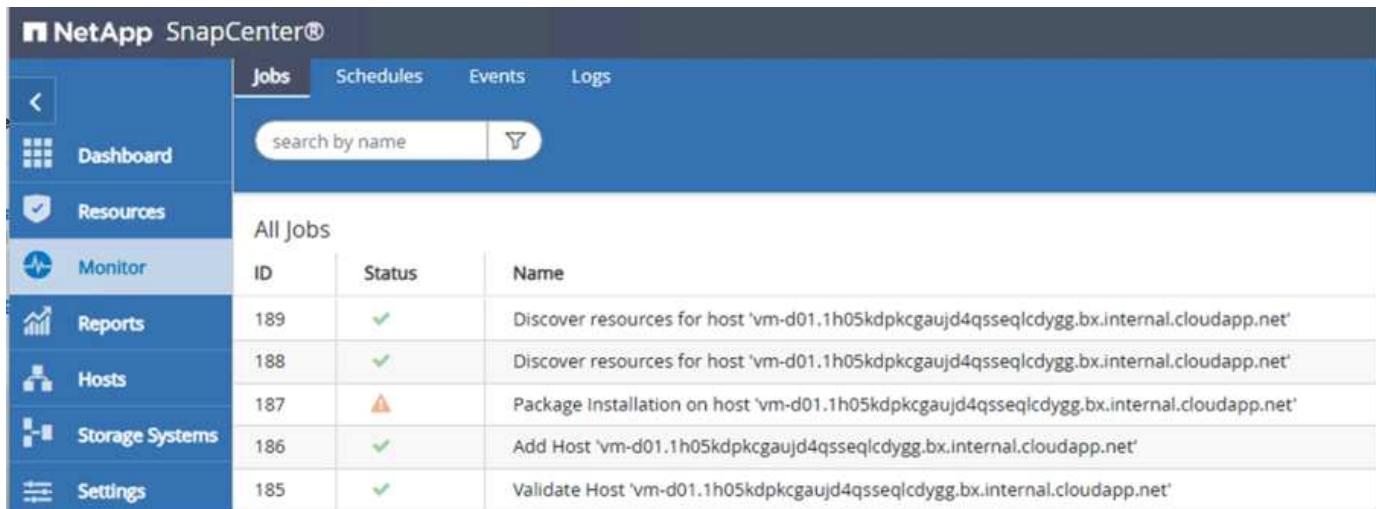
<input checked="" type="checkbox"/> IBM DB2	<input type="checkbox"/> MongoDB
<input type="checkbox"/> MySQL	<input type="checkbox"/> Oracle Applications i
<input type="checkbox"/> Oracle Database	<input type="checkbox"/> SAP ASE
<input type="checkbox"/> PostgreSQL	<input type="checkbox"/> SAP MaxDB
<input type="checkbox"/> SAP HANA	<input type="checkbox"/> Storage i
<input checked="" type="checkbox"/> Unix File Systems	

i [More Options](#) : Port, Install Path, Custom Plug-Ins...

Submit Cancel



After the installation a discovery of the Databases on the host is triggered.



The screenshot shows the NetApp SnapCenter interface. The left sidebar has icons for Dashboard, Resources, Monitor, Reports, Hosts, Storage Systems, and Settings. The 'Jobs' tab is selected. A search bar at the top says 'search by name'. Below it is a table titled 'All Jobs' with columns 'ID', 'Status', and 'Name'. The table lists five entries:

ID	Status	Name
189	✓	Discover resources for host 'vm-d01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net'
188	✓	Discover resources for host 'vm-d01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net'
187	⚠	Package Installation on host 'vm-d01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net'
186	✓	Add Host 'vm-d01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net'
185	✓	Validate Host 'vm-d01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net'

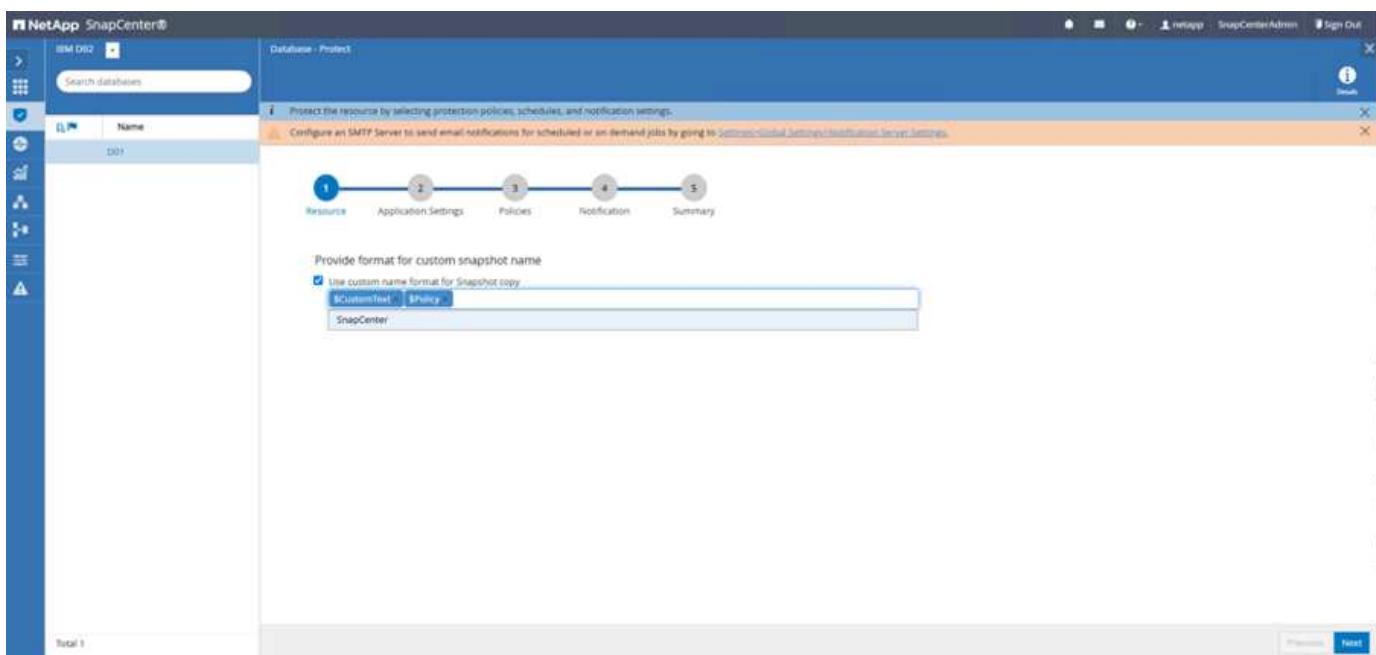
Create Resource Configuration for Database D01

Select discovered Resource D01



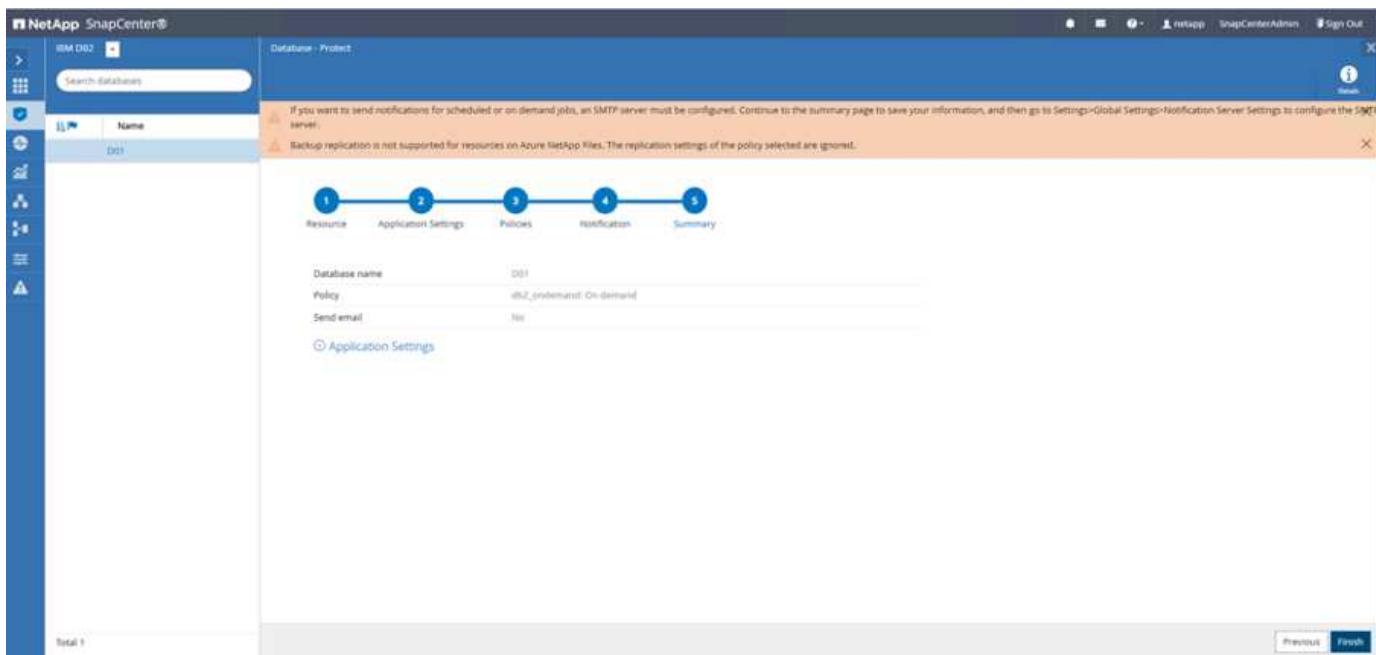
The screenshot shows the 'Resources' section of the NetApp SnapCenter interface. The left sidebar shows 'IBM DB2' selected. A table lists resources with columns: Name, Type, Instance, Host, Resource Groups, Policies, Last backup, and Overall Status. One row is selected: 'D01' (Database, db2001, vm-d01.1h05kdpkcgaujd4qsseqldygg.bx.internal.cloudapp.net, empty, empty, Not protected).

Configure Snapshot Name



The screenshot shows the 'Database - Protect' page. The left sidebar shows 'IBM DB2' selected. The main area has a progress bar with steps 1-5: Resource, Application Settings, Policies, Notification, Summary. Step 1 is active. A note says 'Protect the resource by selecting protection policies, schedules, and notification settings.' A warning says 'Configure an SMTP Server to send email notifications for scheduled or on-demand jobs by going to [System Control > Notifications > SMTP Server Settings](#).' A section for 'Provide format for custom snapshot name' has a checked checkbox 'Use custom name format for Snapshot copy' with radio buttons 'Customized' (selected) and 'Policy'. The 'Customized' field contains 'SnapCenter'.

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 X

1	Restore scope	Summary
2	PreOps	Backup Name: SnapCenter_db2_onDemand_03-26-2025_14_42_07_4144
3	PostOps	Backup date: 03/26/2025 2:43:50 PM
4	Notification	Restore scope: Complete Resource without Volume Revert
5	Summary	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

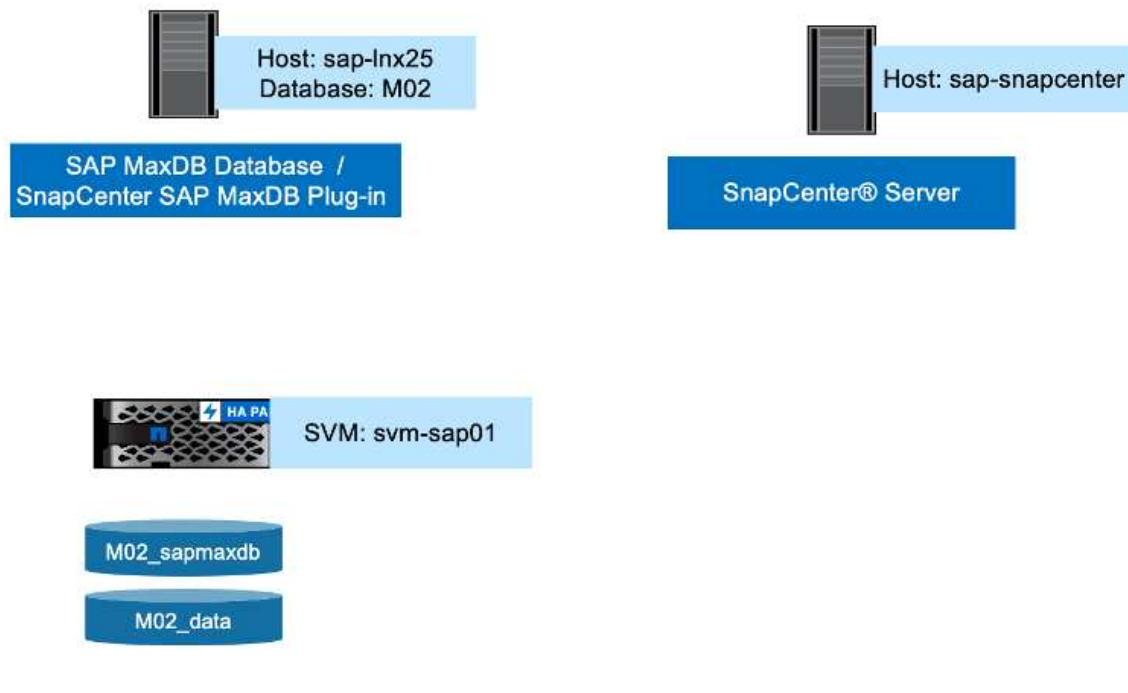
Example Implementation for SnapCenter MaxDB Plugin for an SAP System running in our Demo Center.



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>/backu p	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	<input type="text" value="Linux"/>
Host Name	<input type="text" value="sap-lnx25"/>
Credentials	<input type="text" value="linux-snapcenter"/> + i

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 i |
| <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 i |
| <input checked="" type="checkbox"/> Unix File Systems | |

⚙ [More Options](#) : Port, Install Path, Custom Plug-Ins...

Submit

Cancel

Create SnapCenter Resource Configuration for Database M02

Resources → SAP MaxDB → Add Resources

Add SAP MaxDB Resource X

1 Name **Provide Resource Details**

2 Storage Footprint **3 Resource Settings** **4 Summary**

Name: M02 i

Host Name: sap-lnx25.muccbc.hq.netapp.com ▼

Type: Database ▼

Credential Name: None + i

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 X

1 Name **Provide Resource Details**

2 Storage Footprint **3 Resource Settings** **4 Summary**

Name: M02 i

Host Name: sap-lnx25.muccbc.hq.netapp.com ▼

Type: Database ▼

Credential Name: control-M02 + i

Add SAP MaxDB Resource

1 Name

2 Storage Footprint

3 Resource Settings

4 Summary

Provide Storage Footprint Details

Storage Type ONTAP Azure NetApp Files

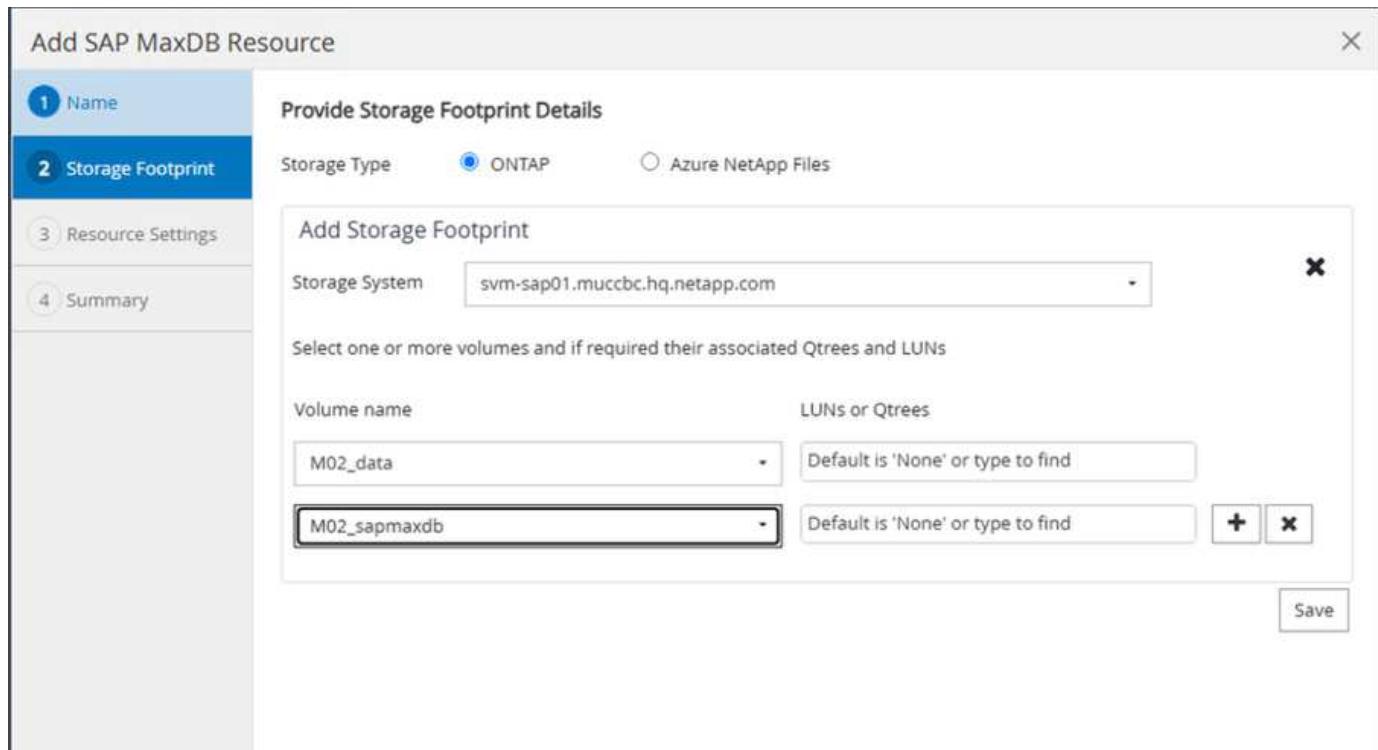
Add Storage Footprint

Storage System

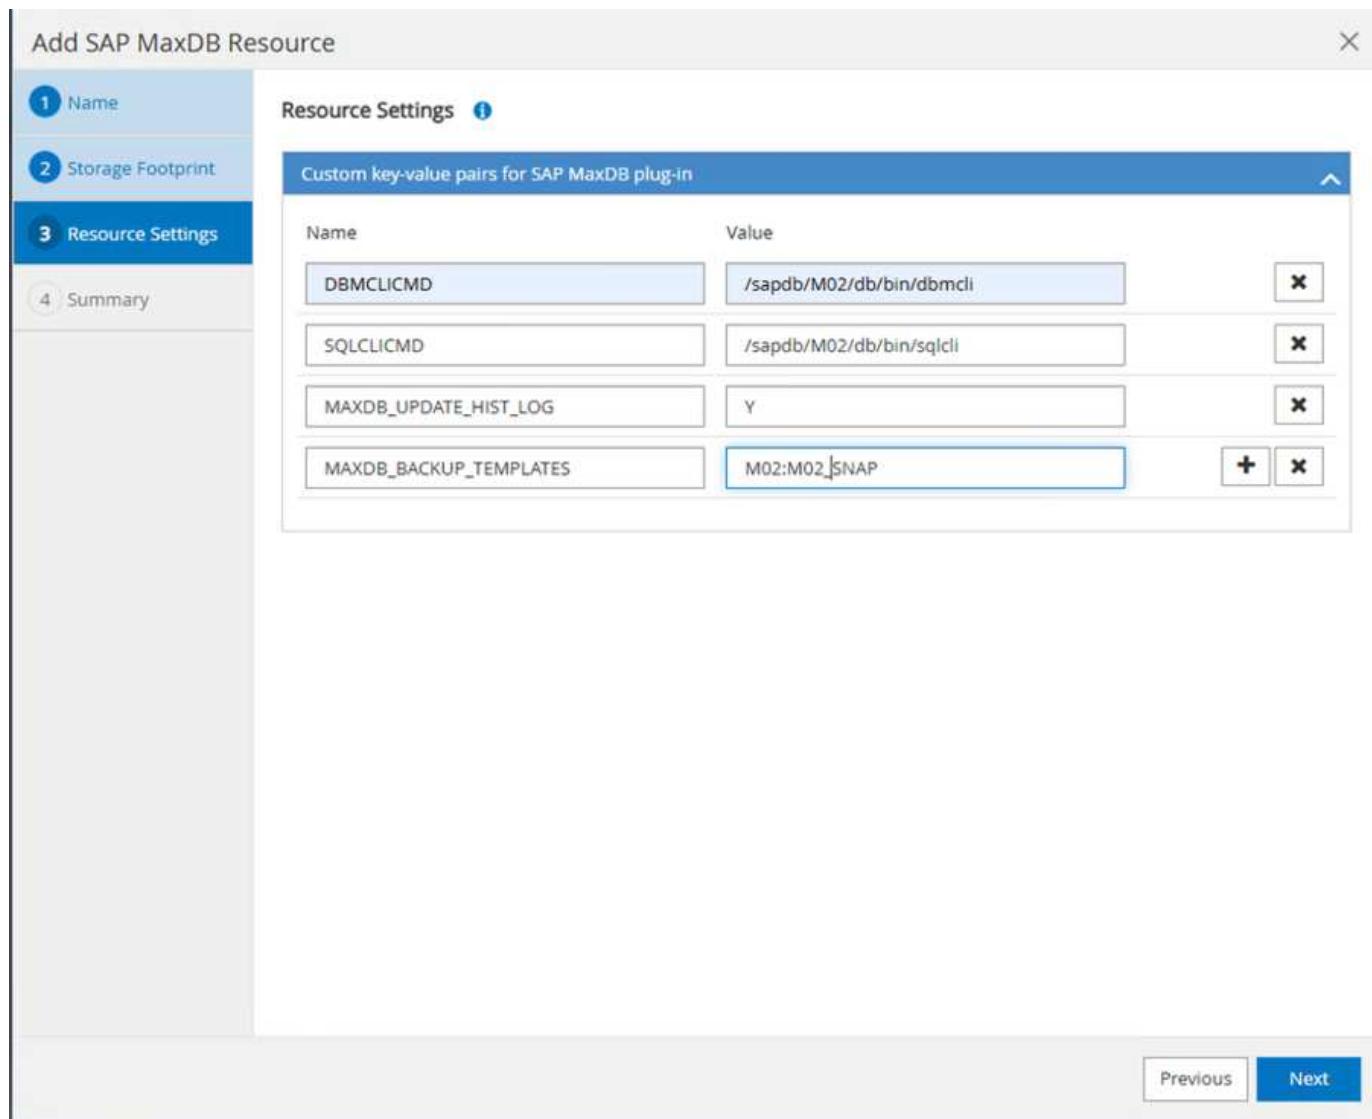
Select one or more volumes and if required their associated Qtrees and LUNs

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"/>

Save



Following Resource Settings Custom key-value pairs must be made (at least).



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.
MAXDB_BG_SERVER_PREFIX	bg_server_prefix (e.g. na_bg)	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®

SAP MaxDB

View Database Search databases

Dashboard Resources Monitor Reports Hosts Storage Systems Settings Alerts

Name	Type	Host	Resource Groups	Policies	Last backup	Overall Status
MO2	Database	sap-in026.muc02sapdemo.netapp.com				Not protected

NetApp SnapCenter®

SAP MaxDB

Database - Project

Configure an SMTP Server to send email notifications for scheduled or on demand jobs by going to [Settings > Global Settings > Notification](#) (see [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:
\$CustomText \$Policy
SnapCenter

Total 1

NetApp SnapCenter®

SAP MaxDB

Database - Protect

1 Resource 2 Application Settings 3 Policies 4 Notification 5 Summary

Backups

Select consistency group option for backup [?](#)

Enable consistency group backup

Scripts

Custom Configurations

Snapshot Copy Tool

Total 1

NetApp SnapCenter®

SAP MaxDB

Database - Protect

Resource Application Settings Policies Notifications Summary

Select one or more policies and configure schedules

maxdb_on-demand +

Configure schedules for selected policies

Policy	Applied Schedules	Configure Schedules	Secondary Protection
maxdb_on-demand	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

Previous Next

NetApp SnapCenter®

SAP MaxDB

Database - Protect

Resource Application Settings Policies Notifications 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.

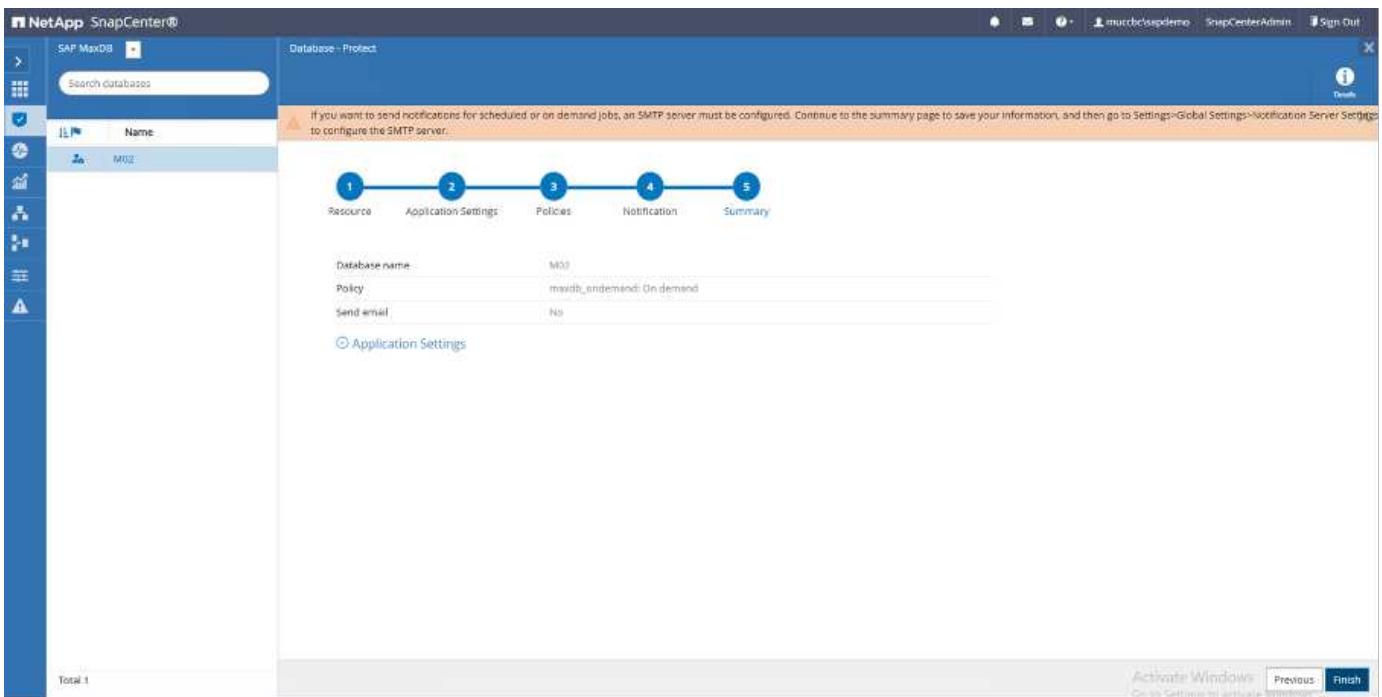
Email preference	Never
From	From: email
To	Email to
Subject	Notification

Attach job report

Total 1

Activate Windows

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

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_maxdb_on-demand_05-20-2025_13:29:34.4378	05/20/2025 1:30:03 PM	1	05/20/2025 1:28:21 PM
SnapCenter_maxdb_on-demand_05-20-2025_13:27:47.2824		1	



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**.

① 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 X

1 Restore scope

Select the restore types

Complete Resource i

File Level i

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...		

⚠ 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 X

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-lnx25)
 - User root: mount /sapdb/M02/sapdata
- Start Database M02 in admin mode and 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_00000008|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 backp
 - 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 recorded 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

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.