



SAP HANA backup and recovery with SnapCenter

NetApp solutions for SAP

NetApp
February 25, 2026

Table of Contents

- 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
 - Backup and recovery using Snapshot backups 2
 - Runtime of Snapshot backup and restore operations 2
 - Recovery time objective comparison 3
 - Use cases and values of accelerated backup and cloning operations 4
 - 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
 - Supported SAP HANA configurations 9
 - Supported platform and infrastructure configurations 10
 - Supported features and operations 10
 - Learn about SnapCenter data protection concepts and best practices 13
 - Deployment options for SnapCenter plug-in for SAP HANA 14
 - SAP HANA block consistency check 17
 - Data protection strategy 18
 - Backup of encryption root keys 20
 - Backup operations 20
 - Backup retention management 21
 - Learn about configuring SnapCenter for SAP HANA environments 23
 - Configure initial SnapCenter settings for SAP HANA 24
 - Credentials configuration 25
 - Storage system configuration 29
 - Policies configuration 30
 - Configure SnapCenter resources for individual SAP HANA databases 35
 - SAP HANA backup user and SAP HANA user store configuration 36
 - Storage replication configuration 38
 - ANF backup configuration 39
 - Deployment of SnapCenter plug-in for SAP HANA 40
 - HANA auto discovery 40
 - Resource protection configuration 41
 - Configure SnapCenter to back up non-data volumes 43
 - Configure SnapCenter central plug-in host for SAP HANA 46
 - SnapCenter HANA plug-in deployment 46
 - SAP HANA hdbsql client software installation and configuration 46
 - SAP HANA user store configuration for a central plug-in host 47
 - HANA manual resource configuration 48
 - Learn about backup operations for SAP HANA Snapshot in SnapCenter 50
 - SAP HANA Snapshot backups in SnapCenter 50
 - SAP HANA Snapshot backups in SAP HANA Studio 51
 - SAP HANA Snapshot backups on storage layer 52
 - SAP HANA Snapshot backups with ANF 53

Snapshot backups of non-data volumes	55
Backup workflow for HANA database backups	56
Backup workflow for non-data volumes	57
Cleanup of secondary backups	58
Execute SAP HANA block consistency checks with SnapCenter	61
Consistency checks with hdbpersdiag using the local snapshot directory	62
Consistency checks with hdbpersdiag using a central verification host	68
File-based backup	77
Restore and recover SAP HANA databases with SnapCenter	79
Automated restore and recovery for SAP HANA MDC systems with a single tenant	80
Manual recovery with HANA Studio	85
Manual recovery with SQL commands	98
Single tenant restore and recovery	99
Restore of non-data volumes	99
Configure advanced SnapCenter options for SAP HANA	100
Warning message with virtualized environments and in-guest mounts	100
Deactivate automated log backup housekeeping	101
Enable secure communication to HANA database	101
Disable auto discovery on the HANA plug-in host	102

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.

Stat...	Started	Duration	Size	Backup Ty...	Destinati...
Jan 11, 2022 10:26:59 AM	00h 01m 17s	4.51 TB	Data Back...	Snapshot	
Jan 11, 2022 8:40:02 AM	00h 27m 11s	4.51 TB	Data Back...	Snapshot	
Jan 11, 2022 1:00:58 AM	04h 05m 39s	3.82 TB	Data Back...	File	
Jan 9, 2022 4:40:03 PM	00h 01m 23s	4.51 TB	Data Back...	Snapshot	
Jan 9, 2022 8:00:02 AM	02h 39m 04s	3.82 TB	Data Back...	File	
Jan 9, 2022 12:40:03 AM	00h 01m 18s	4.51 TB	Data Back...	Snapshot	
Jan 8, 2022 4:40:03 PM	00h 01m 18s	4.51 TB	Data Back...	Snapshot	
Jan 8, 2022 8:40:03 AM	00h 01m 22s	4.51 TB	Data Back...	Snapshot	
Jan 8, 2022 12:40:03 AM	00h 01m 19s	4.51 TB	Data Back...	Snapshot	
Jan 7, 2022 4:40:03 PM	00h 01m 19s	4.51 TB	Data Back...	Snapshot	
Jan 7, 2022 8:40:02 AM	00h 01m 19s	4.51 TB	Data Back...	Snapshot	
Jan 7, 2022 12:40:02 AM	00h 01m 20s	4.51 TB	Data Back...	Snapshot	
Jan 6, 2022 4:40:02 PM	00h 01m 18s	4.51 TB	Data Back...	Snapshot	
Jan 6, 2022 8:40:03 AM	00h 01m 17s	4.51 TB	Data Back...	Snapshot	
Jan 6, 2022 12:40:03 AM	00h 01m 19s	4.51 TB	Data Back...	Snapshot	
Jan 5, 2022 4:40:03 PM	00h 01m 19s	4.51 TB	Data Back...	Snapshot	

File-based backup: **4 hours 05 min**
 (~270 MB/s throughput)

04h 05m 39s	3.82 TB	Data Back...	File
-------------	---------	--------------	------

Snapshot backup: **1 min 20 sec**

00h 01m 18s	4.51 TB	Data Back...	Snapshot
00h 01m 22s	4.51 TB	Data Back...	Snapshot
00h 01m 19s	4.51 TB	Data Back...	Snapshot

Backup runtime reduced by 99%

Recovery time objective comparison

This section provides a recovery time objective (RTO) comparison of file-based and storage-based Snapshot backups. The RTO is defined by the sum of the time needed to restore, recover, and then 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.

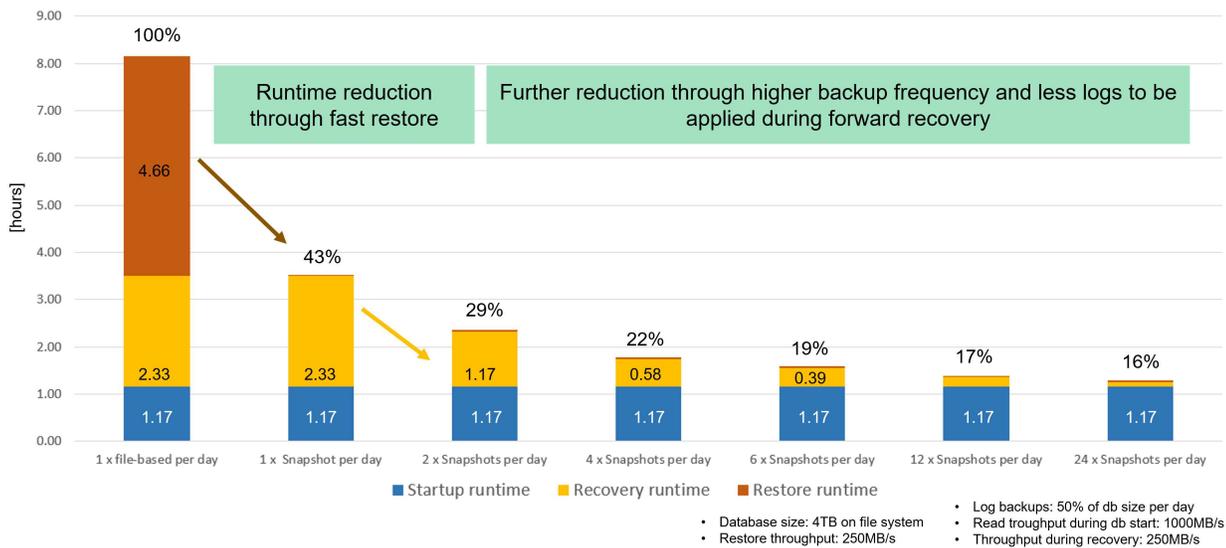
Restore and recovery sample calculation

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

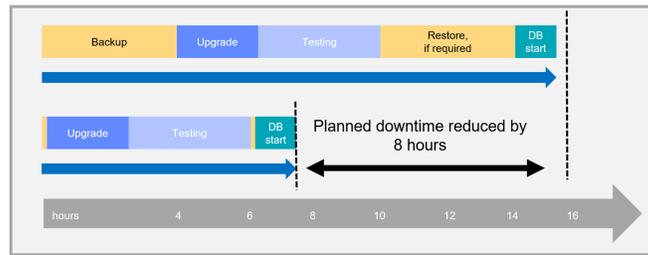
Executing backups is a critical part of any data protection strategy. Backups are scheduled on a regular basis to ensure that you can recover from system failures. This is the most obvious use case, but there are also other SAP lifecycle management tasks, where accelerating backup and recovery operations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

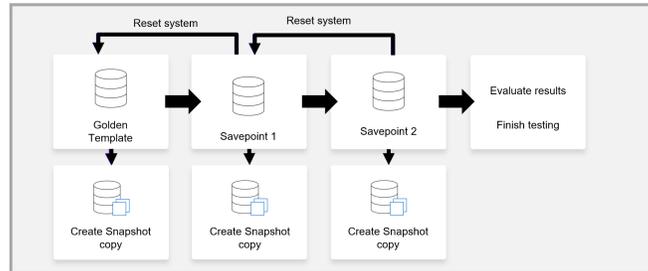
Accelerate HANA system upgrade operations

- 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.

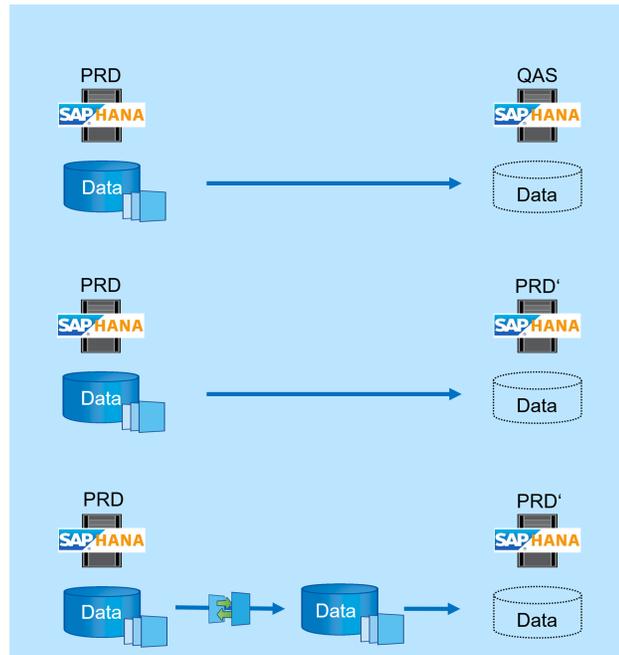
The most popular use case is the SAP System Refresh, where data from the production system needs to be copied to the test or QA system. By leveraging the ONTAP or ANF cloning feature, you can provision the volume for the test system from any Snapshot copy of the production system in a matter of seconds.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

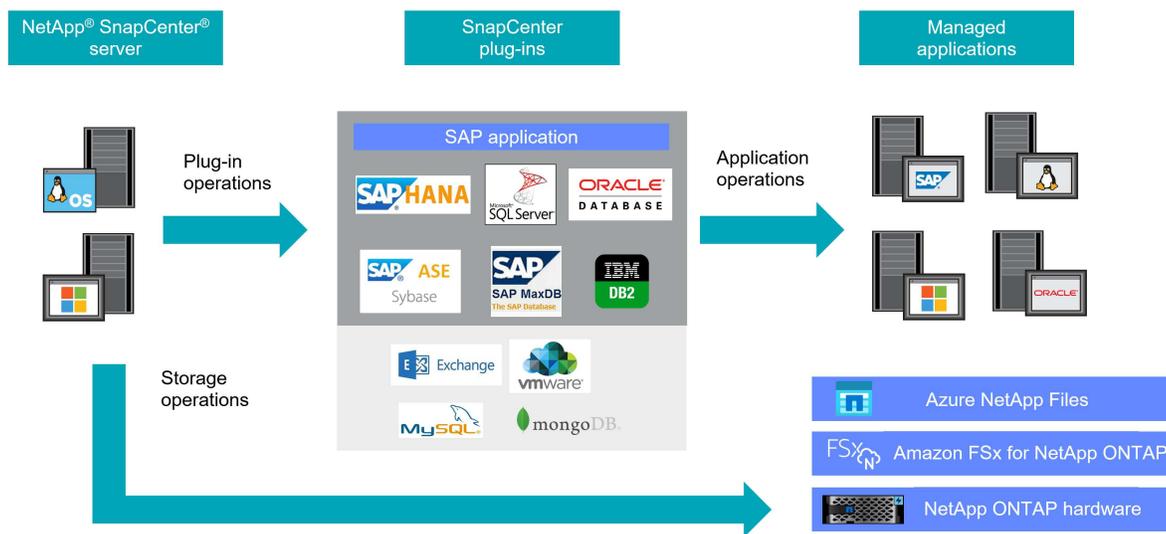


Learn about the SnapCenter architecture

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.

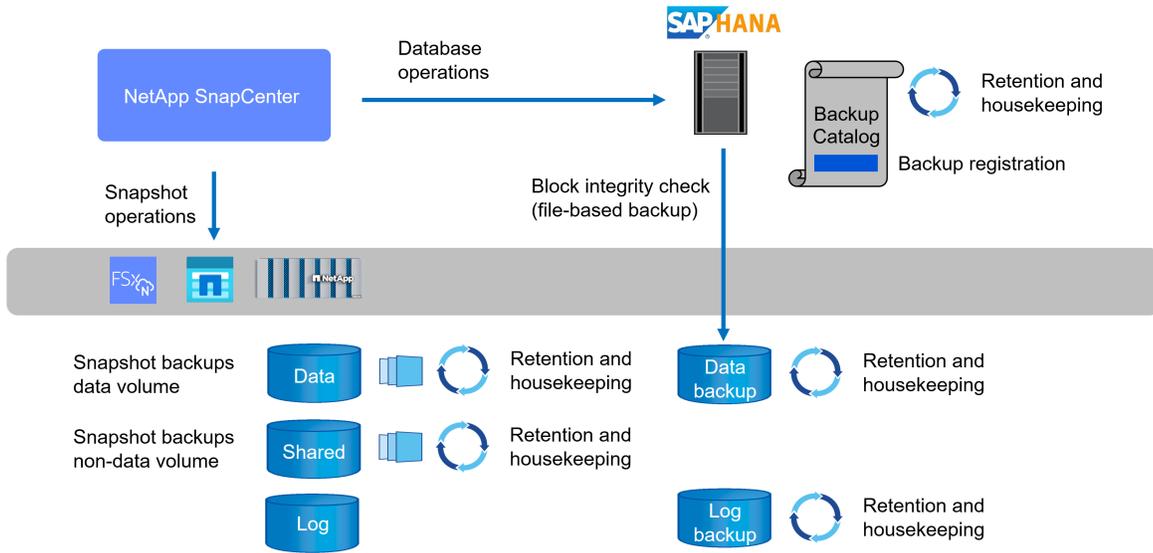
Based on your configured retention policies, SnapCenter manages the housekeeping of data file backups at the primary storage, log file backups, and the SAP HANA backup catalog.

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

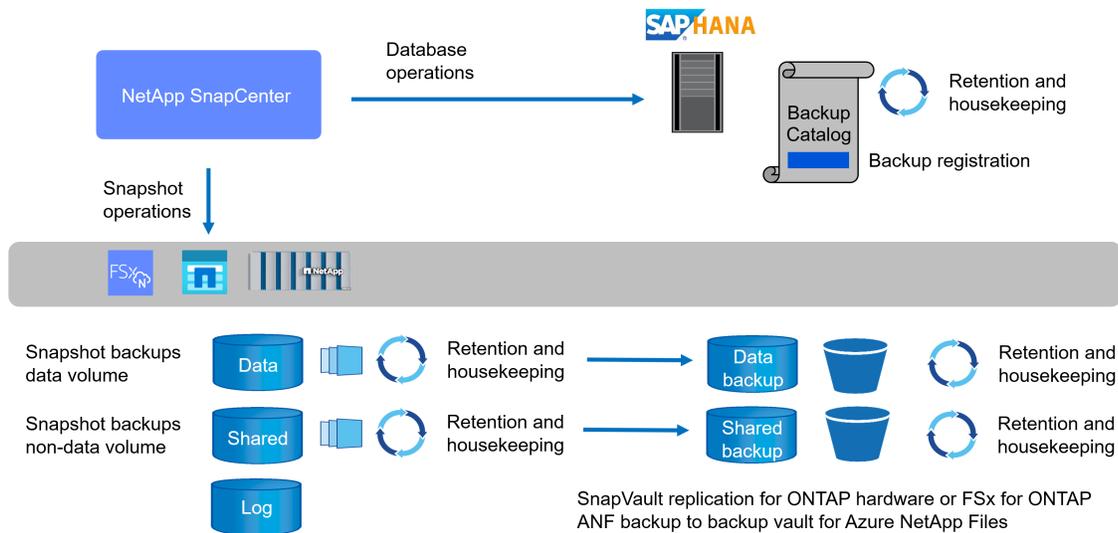
The following figure shows an overview of the SnapCenter backup and retention management operations.

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 SF SR (selectable)	SF SR of complete LUN	Clone, mount, copy
Single tenant restore	SF SR	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 SF SR (selectable)	SF SR 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 SF SR (selectable)
Single tenant restore	SF SR
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

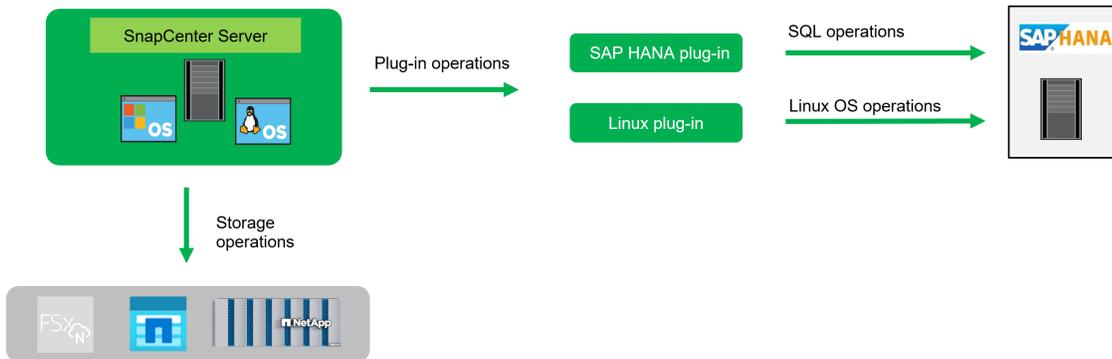
Learn about SnapCenter data protection concepts and best practices

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:

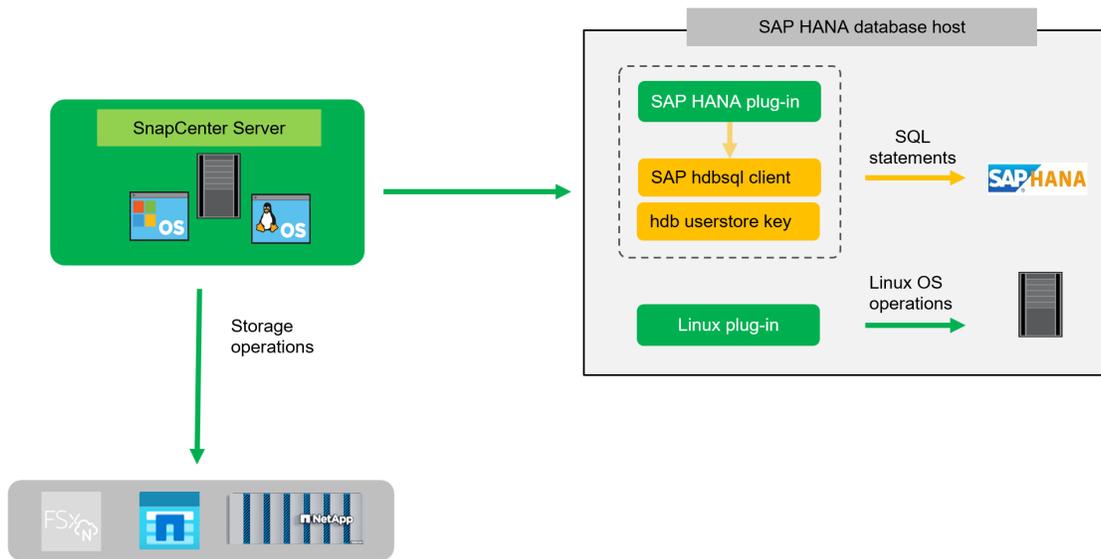
- VMware with FC datastores
- SnapMirror active sync
- SnapCenter server high availability if used as a central plug-in host
- HANA system auto discovery
- Automated HANA database recovery

- 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 hdb 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.

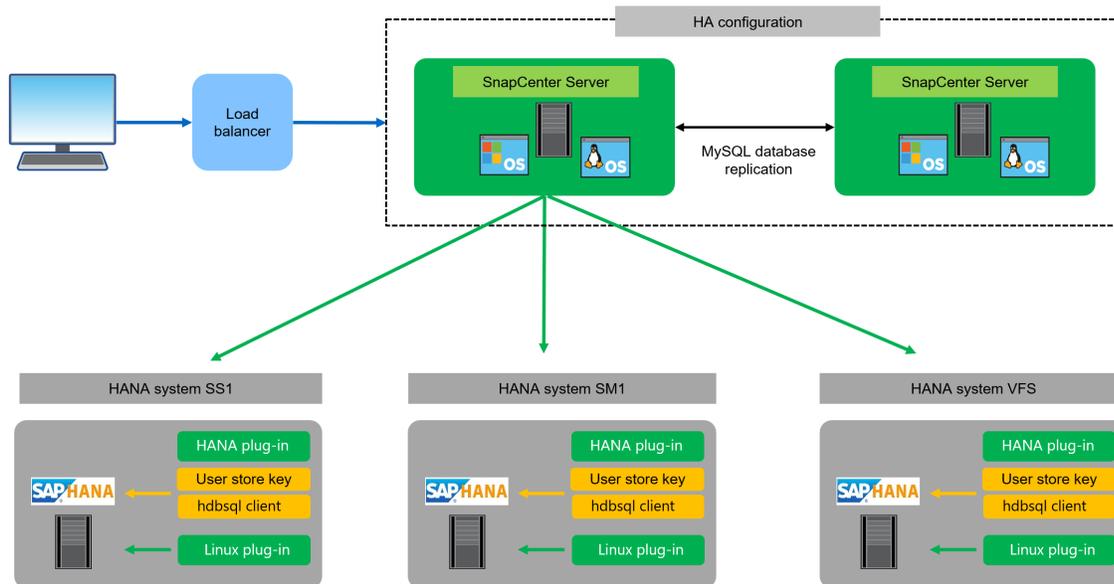
When the HANA plug-in is deployed on the HANA database host, the HANA system is auto discovered by SnapCenter and is flagged as an auto discovered resource in SnapCenter.



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

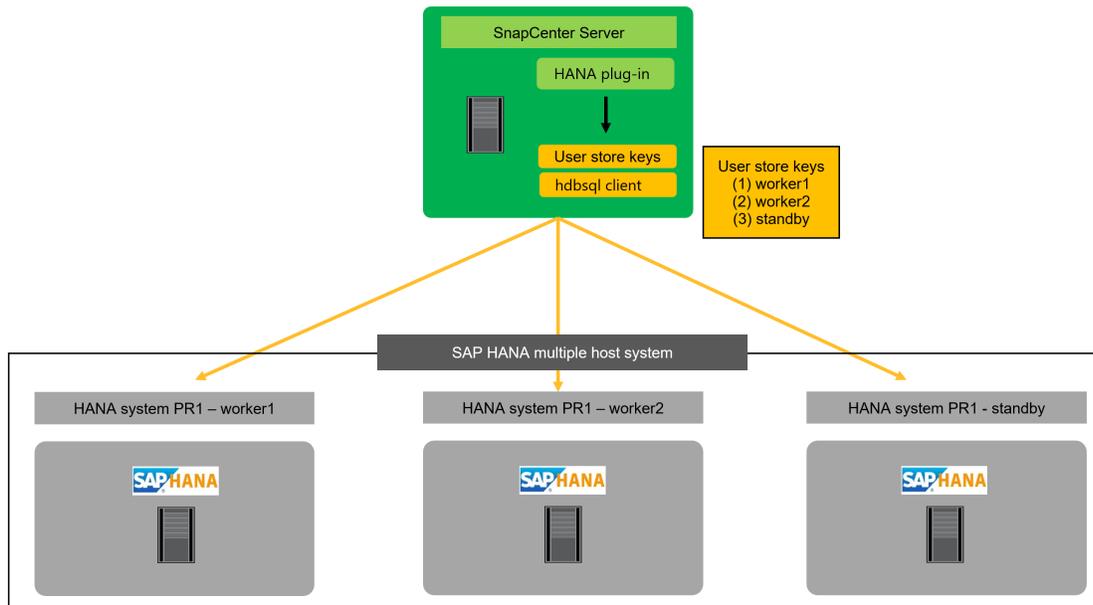
- HANA multiple host systems
- HANA systems running on IBM Power

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.

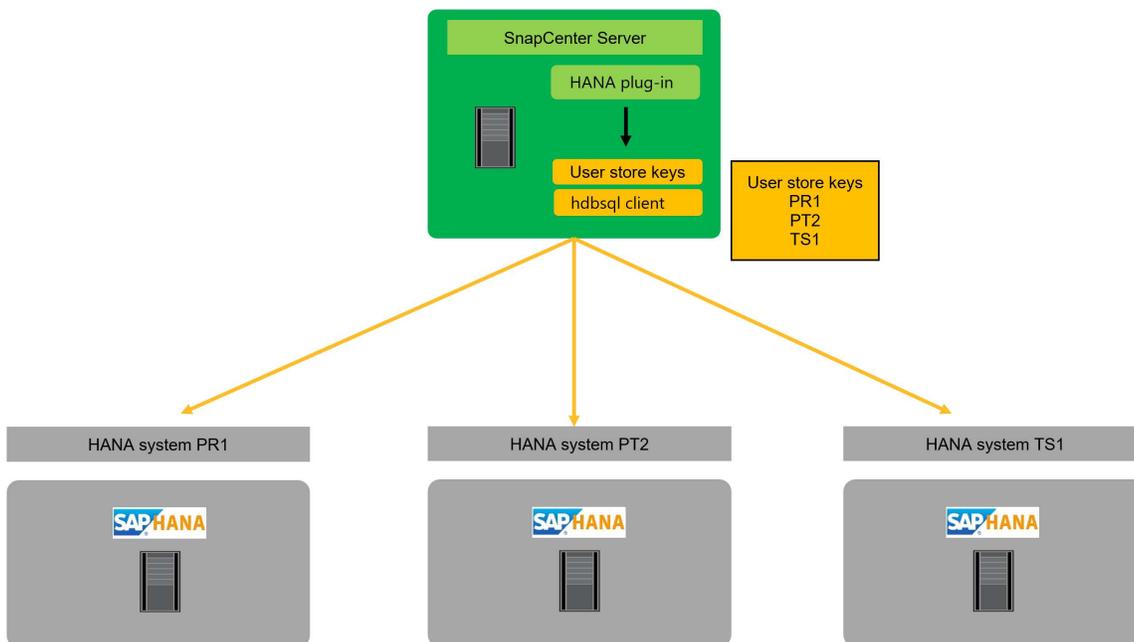


When you run your SnapCenter server on Windows, you can use your Windows system as central plug-in host. When you run your SnapCenter server on Linux, you must use a different host as central plug-in host .

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



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.

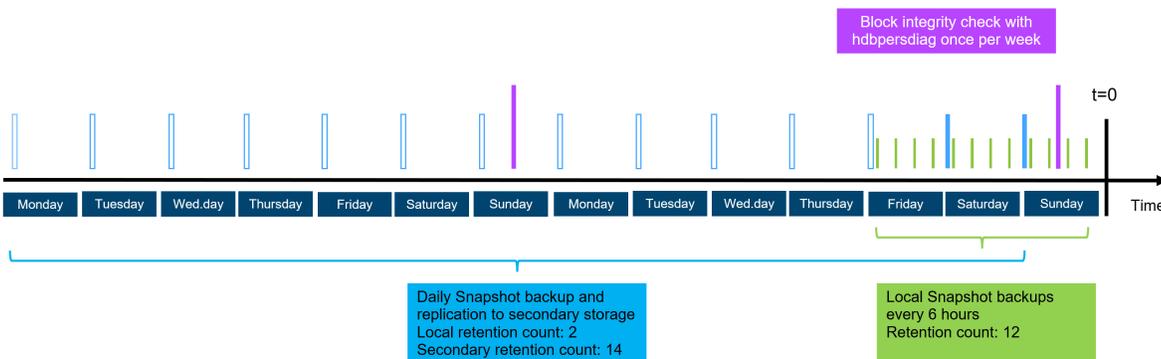


For ANF backup, no additional configuration is required outside of SnapCenter. The ANF backup secondary retention is managed by SnapCenter.



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.

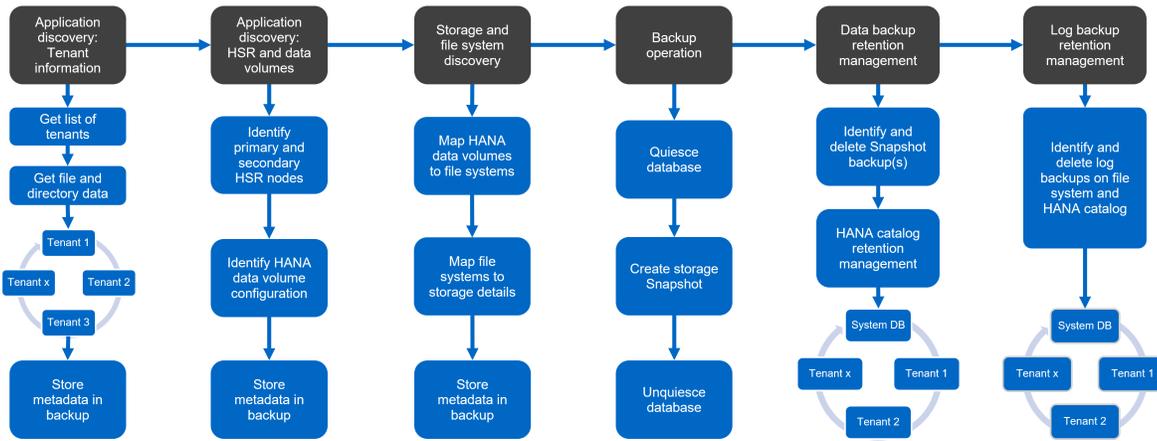
The next step is the Snapshot backup operation itself. This step includes the SQL command to trigger the HANA database snapshot, the storage Snapshot backup, and the SQL command to close the HANA snapshot operation. By using the close command, the HANA database updates the backup catalog of the system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

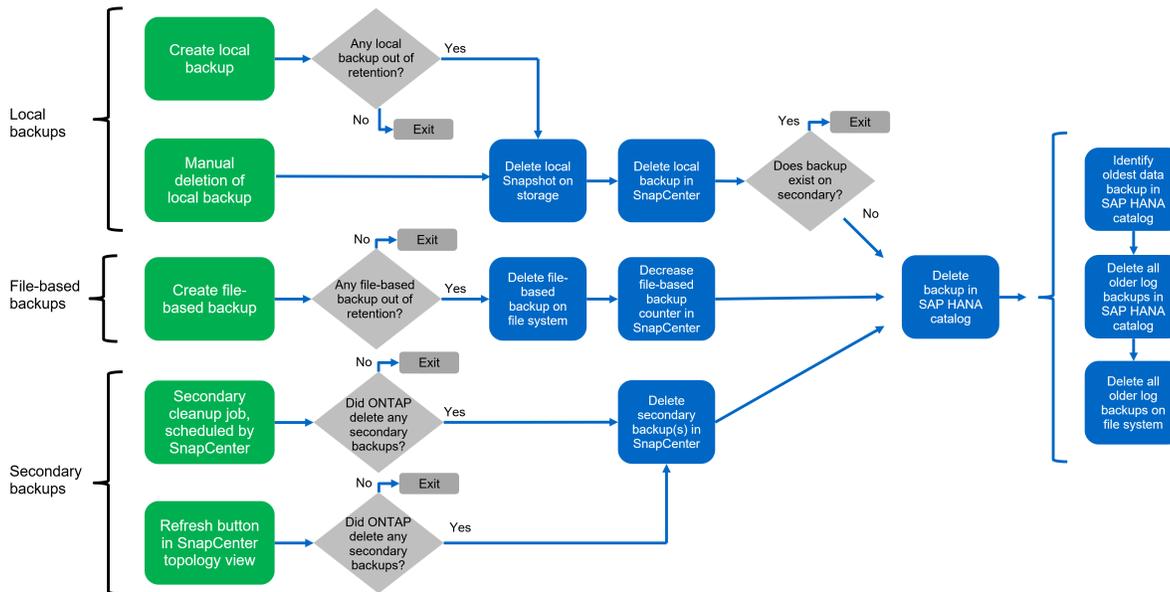


Backup retention management

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:

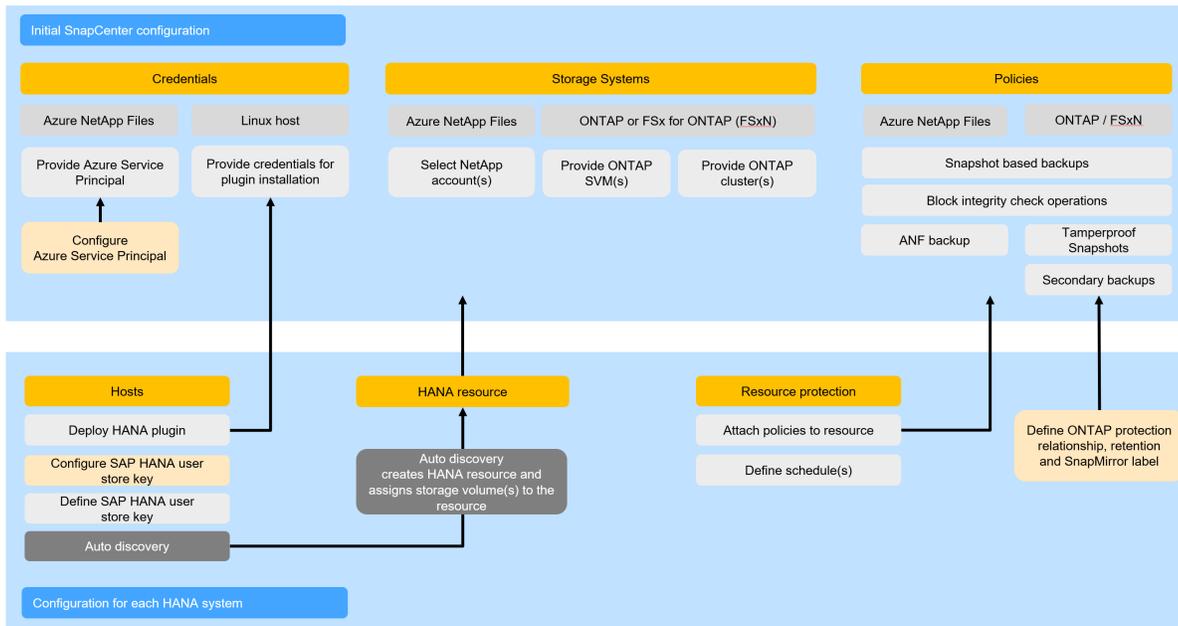
- The initial configuration
 - Credential, storage and policy configurations.
These settings or resources are typically consumed by multiple HANA systems.
- The HANA resource-specific configuration
 - Host, HANA and resource protection configuration must be done for each HANA system individually.

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.

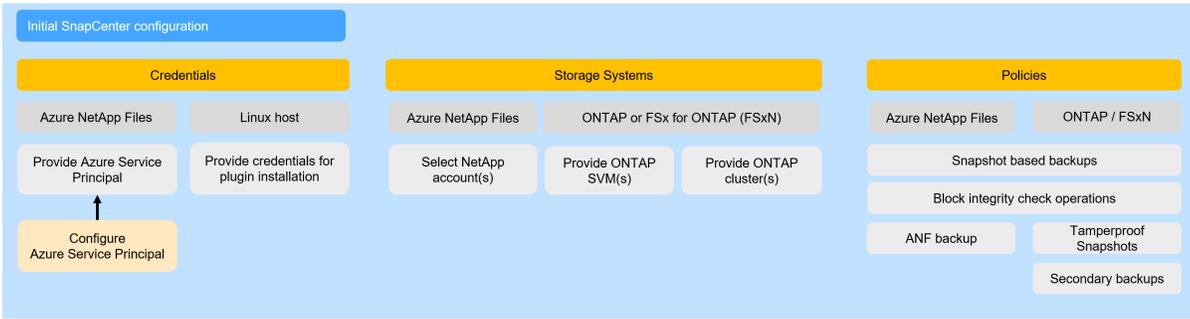
The SnapCenter initial configuration includes the following steps:

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.

Credential Name	Authentication Mode	Details
InstallOnBareMetal	Linux	User:root
InstallPluginOnLinux	Linux	User:root
InstallPluginOnWindows	Windows	User:svcacct/admin
SCV-sapcc	Linux	User: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 shows the following configuration:

- Credential Name: InstallPluginOnLinux
- Authentication Mode: Linux
- Authentication Type: Password Based, SSH Key Based
- Username: scuser
- Password: [masked]
- Use sudo privileges

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/ac
tion",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/breakFileLocks/acti
on",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/getGroupIdListForLd
apUser/action",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/backups/write",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/backups/restoreFile
s/action",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/read",

```

```
"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.



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

Authentication Mode

Azure Details

Tenant ID

Client ID

Client Secret Key

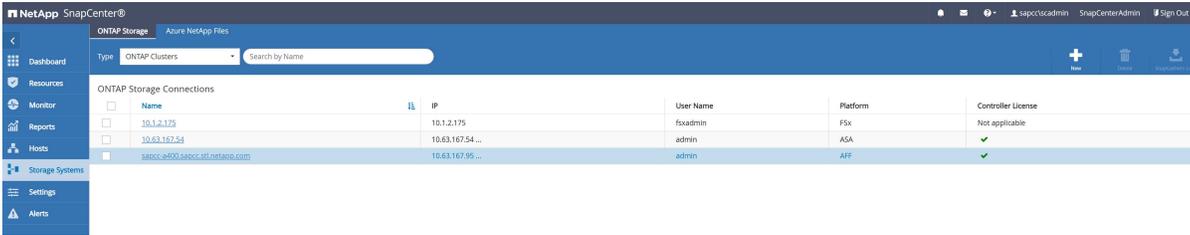
Cancel
OK

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.

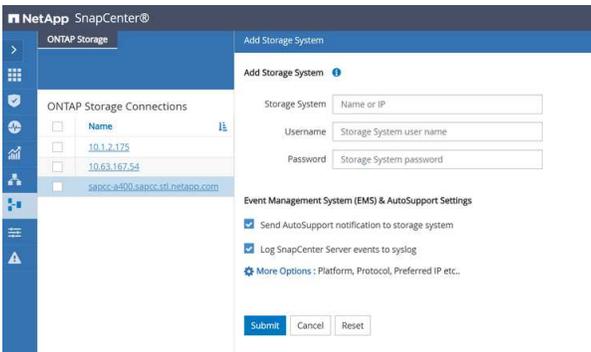
In our lab setup, we added the storage clusters to SnapCenter. ONTAP clusters are configured in the Storage systems section by selecting the ONTAP storage tab and the ONTAP Cluster type. A new cluster is added by clicking the + icon.



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.

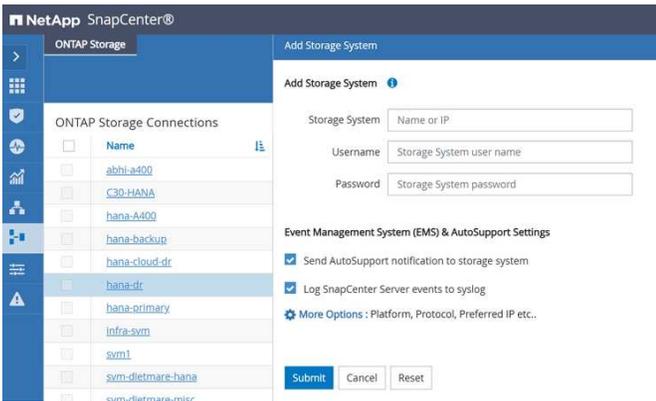
In the following screen you need to provide the credentials for a cluster user.



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](#).

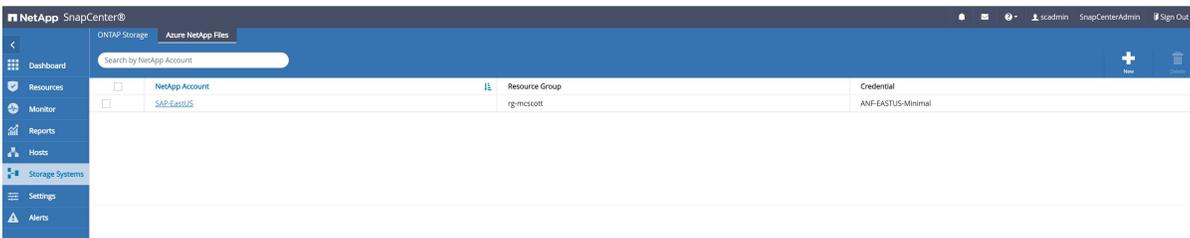


The DNS name for the SVM must match with the SVM name configured at the ONTAP system.

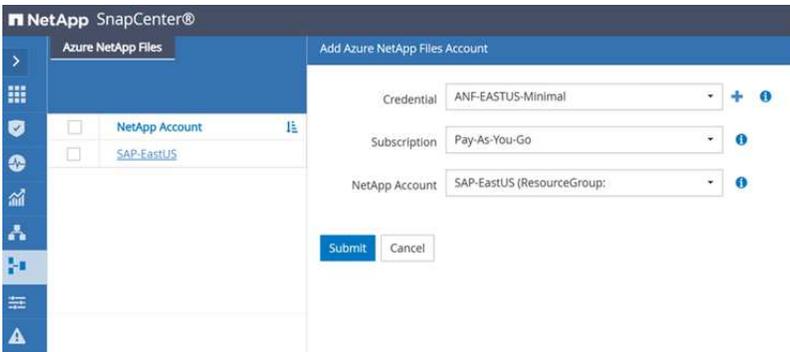


Azure NetApp Files

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		
LocalSnapAndANFBackup	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		SnapVault
LocalSnapKeep2	Data Backup	Hourly	Copies to keep: 2 copies		
LocalSnap-OnDemand	Data Backup	On demand	Copies to keep: 2 copies		
LocalSnapTamperProof	Data Backup	Daily	Copies to keep: 4 copies		
Policy4CBA	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.

Modify SAP HANA Backup Policy

- Name**: Provide a policy name. Policy name: LocalSnapAndSnapVault. Details: [Empty field]
- Policy type
- Snapshot And Replication
- Summary

Select ONTAP storage type and Snapshot policy scope.

Modify SAP HANA Backup Policy

- Name
- Policy type**: Choose storage type. ONTAP/FSx/Cloud volumes ONTAP Azure NetApp Files
- Snapshot And Replication
- Summary

Choose policy scope: Snapshot Based File-Based

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 screenshot shows the 'Modify SAP HANA Backup Policy' dialog box with the 'Snapshot And Replication' tab selected. The 'Choose schedule frequency' section has 'Daily' selected. Under 'Snapshot settings', 'Copies to keep' is set to 2, 'Retain copies for' is 14 days, and both primary and secondary snapshot copy locking periods are set to 7 days. The 'Policy label' is 'daily'. Under 'Select secondary replication options', 'Update SnapVault after creating a local Snapshot copy' is checked, and the 'Error retry count' is 3.

The summary screen shows the configured parameters.

The screenshot shows the 'Summary' tab of the 'Modify SAP HANA Backup Policy' dialog box. The summary table is as follows:

Summary	
Policy name	LocalSnapAndSnapVault
Details	
Backup Type And Replication	Snapshot Based Backup
Schedule Type	Daily
Daily backup retention	Total copies to retain : 2
Replication	SnapVault: enabled , Secondary policy label: daily , Error retry count: 3

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.

Modify SAP HANA Backup Policy

1 Name

Provide a policy name

Policy name: LocalSnapAndANFBackup

Details: [Empty field]

Select Azure NetApp Files storage type and Snapshot policy scope.

Modify SAP HANA Backup Policy

2 Policy type

Choose storage type

ONTAP/FSx/Cloud volumes ONTAP Azure NetApp Files

Choose policy scope

Snapshot Based File-Based

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.

Modify SAP HANA Backup Policy

3 Snapshot and backup

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: 2 copies
 Retain copies for: 14 days

Backup settings

Enable backup

Data backup retention settings

Copies to keep: 7 copies
 Retain copies for: 14 days

Previous Next

The summary screen shows the configured parameters.

Summary	
Policy name	LocalSnapAndANFBackup
Details	
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 retention	Total copies to retain : 7

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.

Provide a policy name

Policy name:

Details:

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

Choose storage type

ONTAP/FSx/Cloud volumes ONTAP Azure NetApp Files

Choose policy scope

Snapshot Based 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.

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: 1 copies

Retain copies for: 14 days

The summary screen shows the configured parameters.

Modify SAP HANA Backup Policy

1 Name

2 Policy type

3 Snapshot and Replication

4 Summary

Summary

Policy name	BlockIntegrityCheck
Details	
Backup Type And Replication	File-Based Backup
Schedule Type	Weekly
Weekly backup retention	Total copies to retain : 1

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

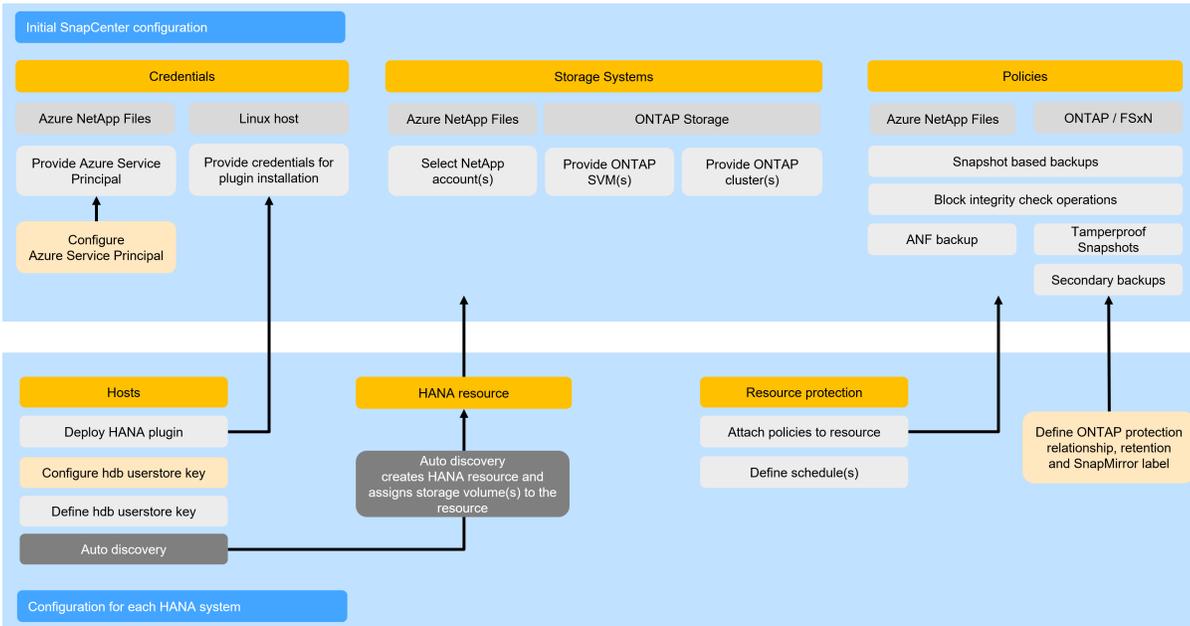
Configure individual SAP HANA databases in SnapCenter by creating backup users and user store keys, setting up storage replication for secondary backups, deploying the HANA plug-in for auto discovery, and configuring resource protection with policies and schedules.

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.

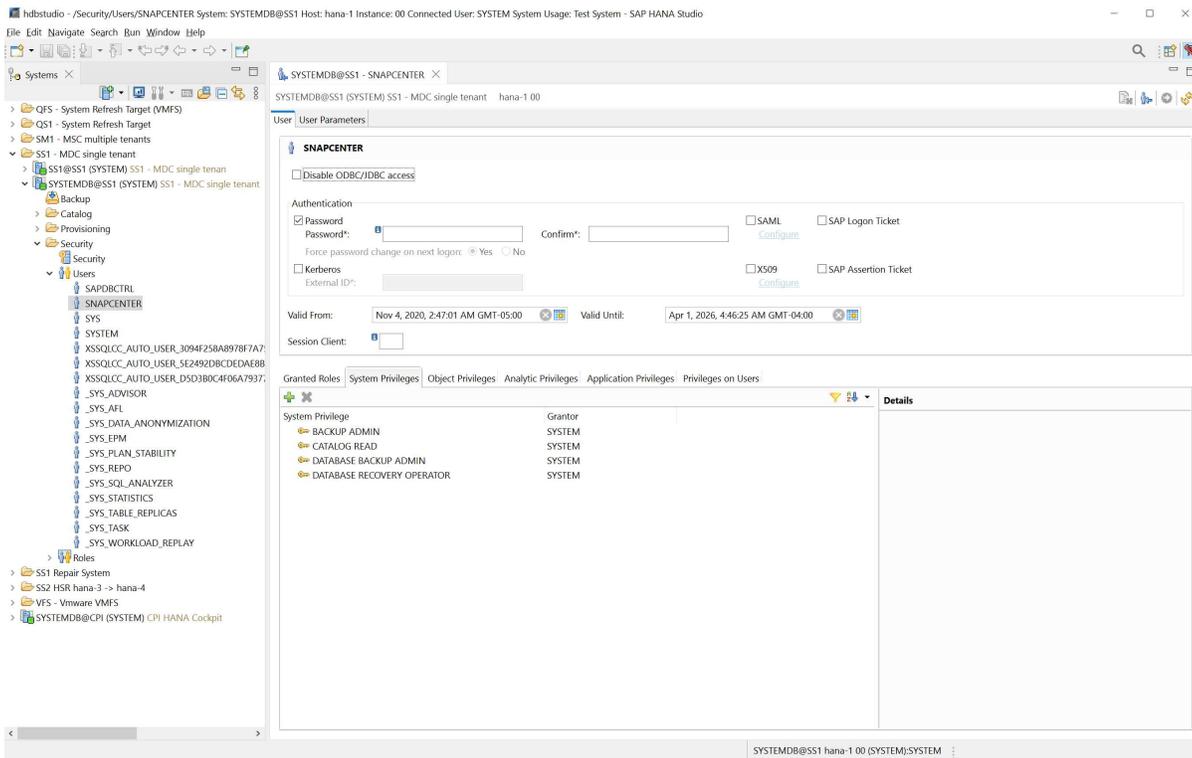
The following figure shows the SAP HANA Studio through which the backup user, in this example SNAPCENTER can be created.



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.

```

smladm@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.
smladm@hana-2:/usr/sap/SM1/HDB12>

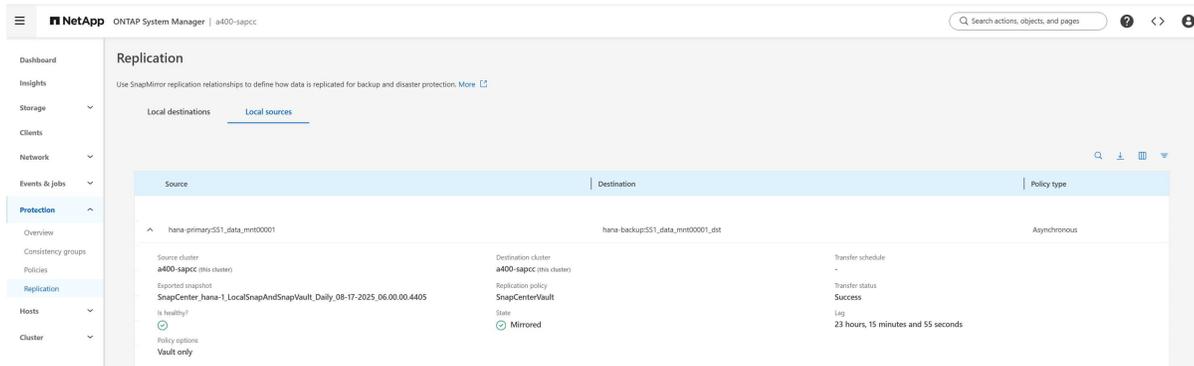
```

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 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.



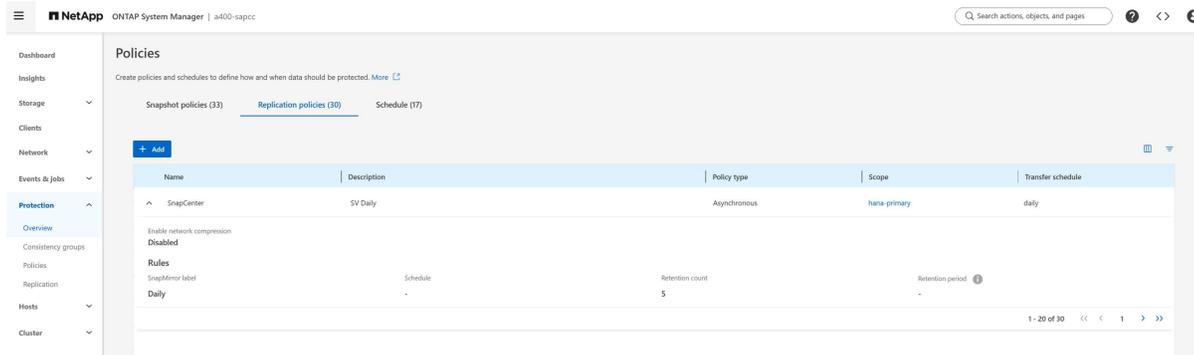
The SnapMirror label in the replication policy must match the label defined in the SnapCenter policy configuration.



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.

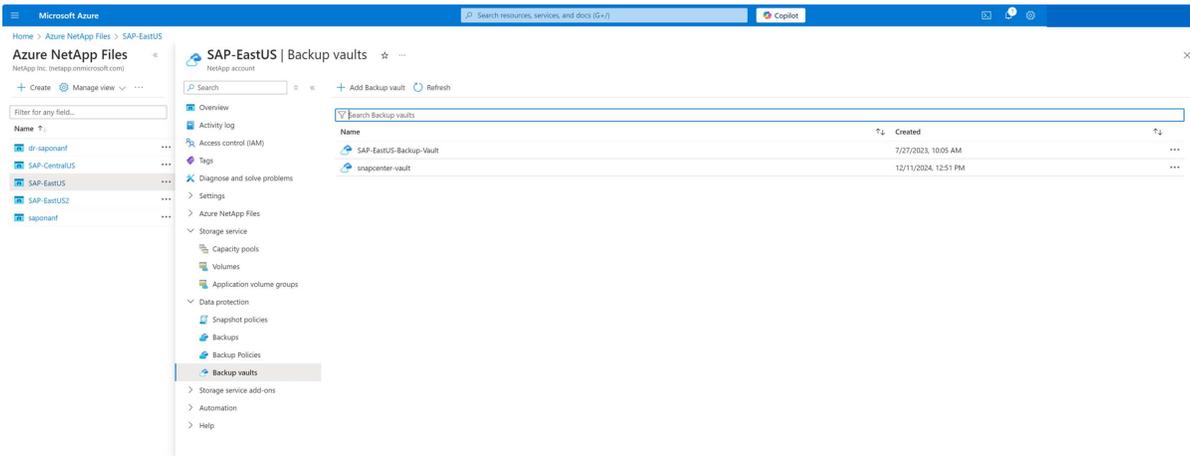


The retention for backups at the secondary backup storage is defined in the policy and controlled by ONTAP.



ANF backup configuration

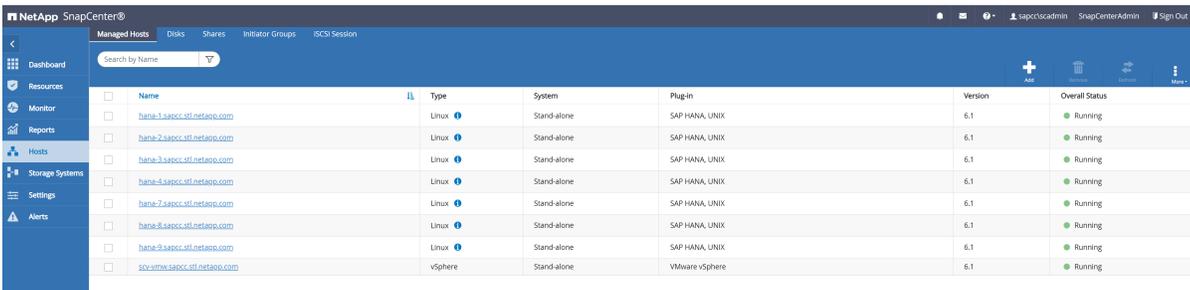
For ANF backup no specific preparation is required. As soon as the first backup with enabled ANF backup is executed an Azure backup vault with the name snapcenter-vault is created by SnapCenter. This backup vault is then used by all following ANF backup operations executed by SnapCenter.



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](#).

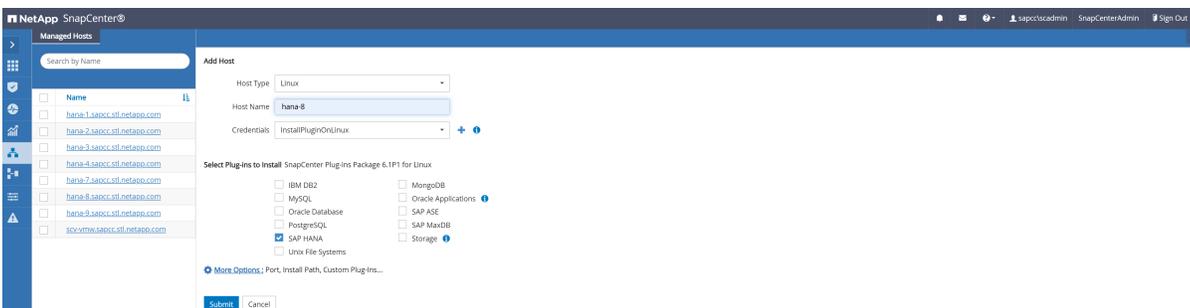
The HANA plug-in deployment is done by clicking the Add button in the Hosts section of the SnapCenter UI.



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.

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.stl.netapp.com				Not protected
Q51	Q51	Q51 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	08/18/2025 12:01:04 PM	Backup succeeded
SS1	SS1	SS1	None	Online	hana-1.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap LocalSnapAndSnap/vault 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 Replication	BlockIntegrityCheck LocalSnap	08/18/2025 11:57:30 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Secondary)	Online	hana-4.sapcc.stl.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapKeep2	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:

Buttons: 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.

Details for selected resource			
Type	Multitenant Database Container		
HANA System Name	SS1		
SID	SS1		
Tenant Databases	SS1		
Plug-in Host	hana-1.sapcc.stl.netapp.com		
HDB Secure User Store Key	SS1KEY		
HDBSQL OS User	ss1adm		
Log backup location	/mnt/rog-backup		
Backup catalog location	/mnt/rog-backup		
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, LocalSnapAndSnap/vault, BlockIntegrityCheck, LocalSnap-OnDemand		
Discovery Type	Auto		
Storage Footprint			
SVM	Volume	Junction Path	LUN/Qtree
hana-primary	SS1_data_mnt00001	/SS1_data_mnt00001	

Resource protection configuration

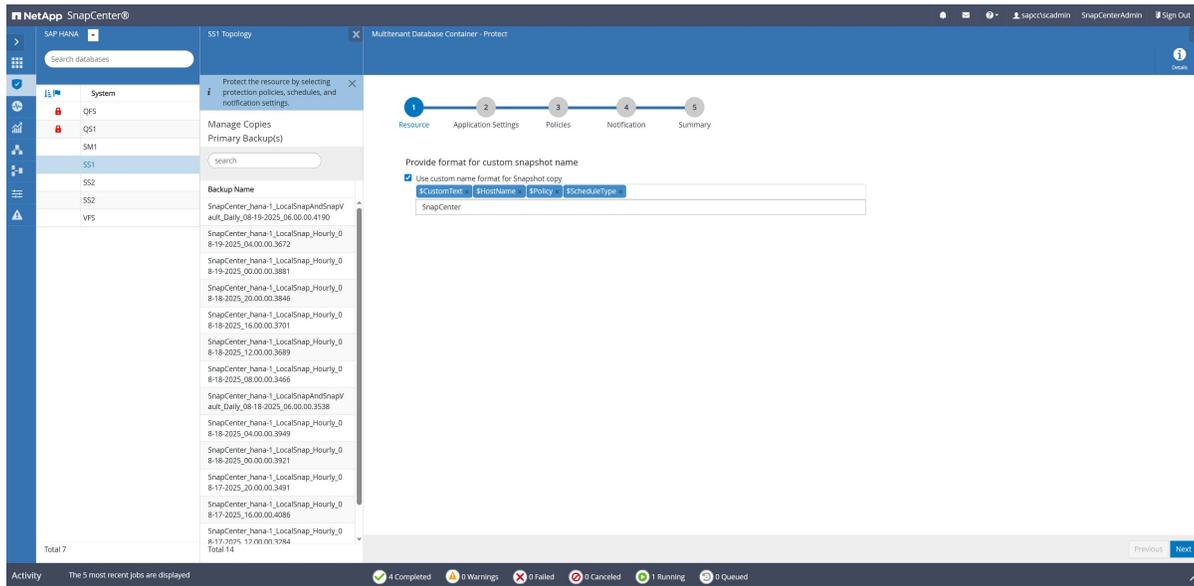
The resource protection configuration screen is opened by clicking on a resource after the auto discovery process has finished. The screenshots in this documentation show the protection configuration of an existing resource.

Configure a custom name format for the Snapshot. NetApp recommends using a custom Snapshot name to

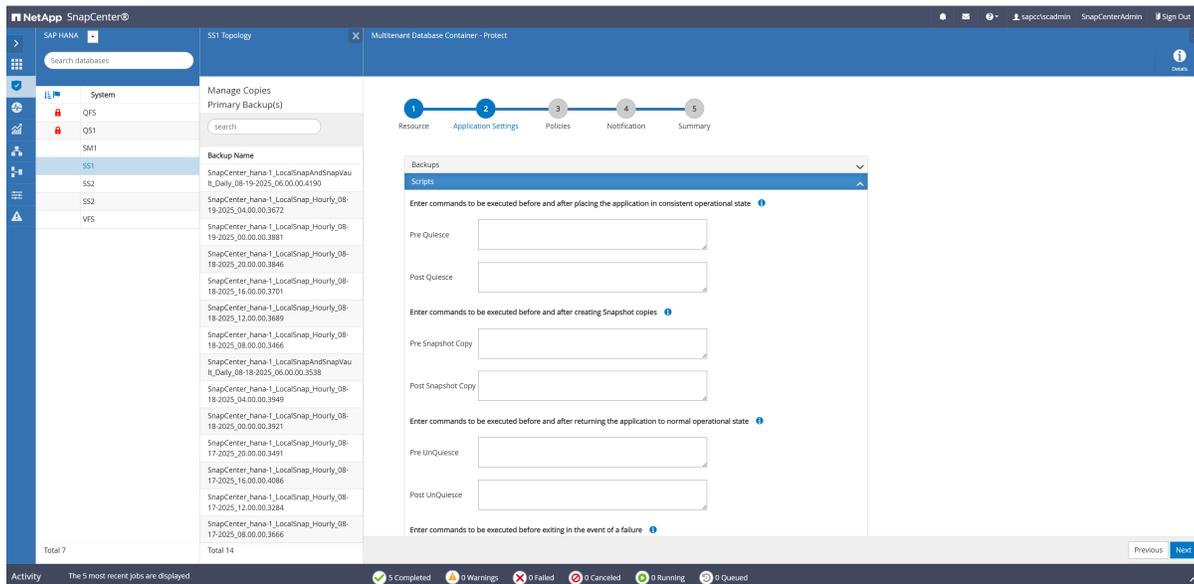
easily identify which backups have been created with which policy and schedule type.

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.

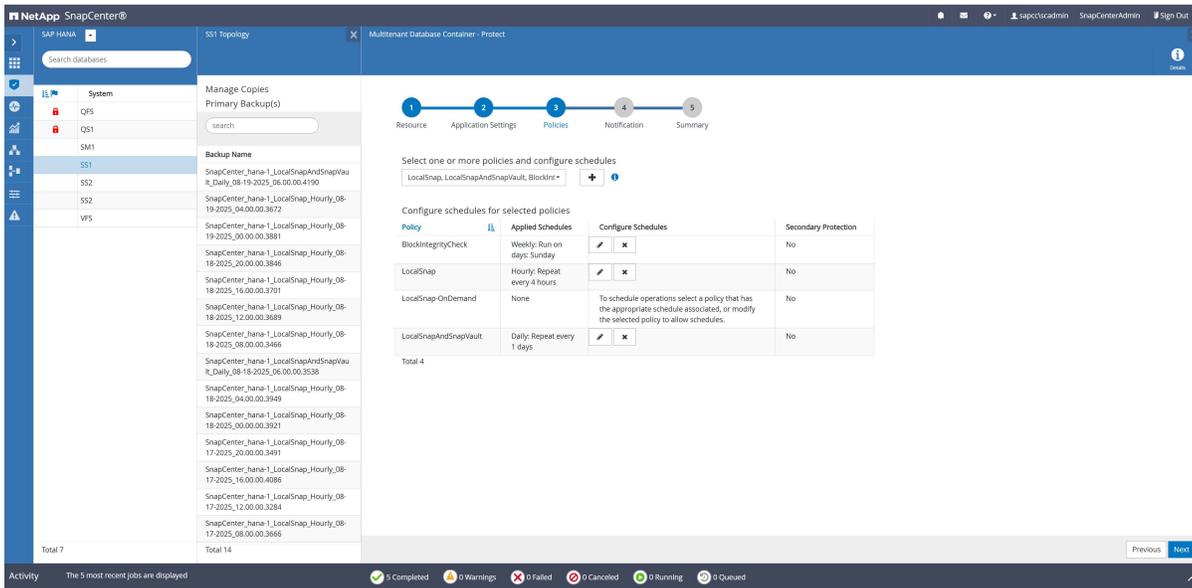


Now policies are attached to the resource and schedules are defined.

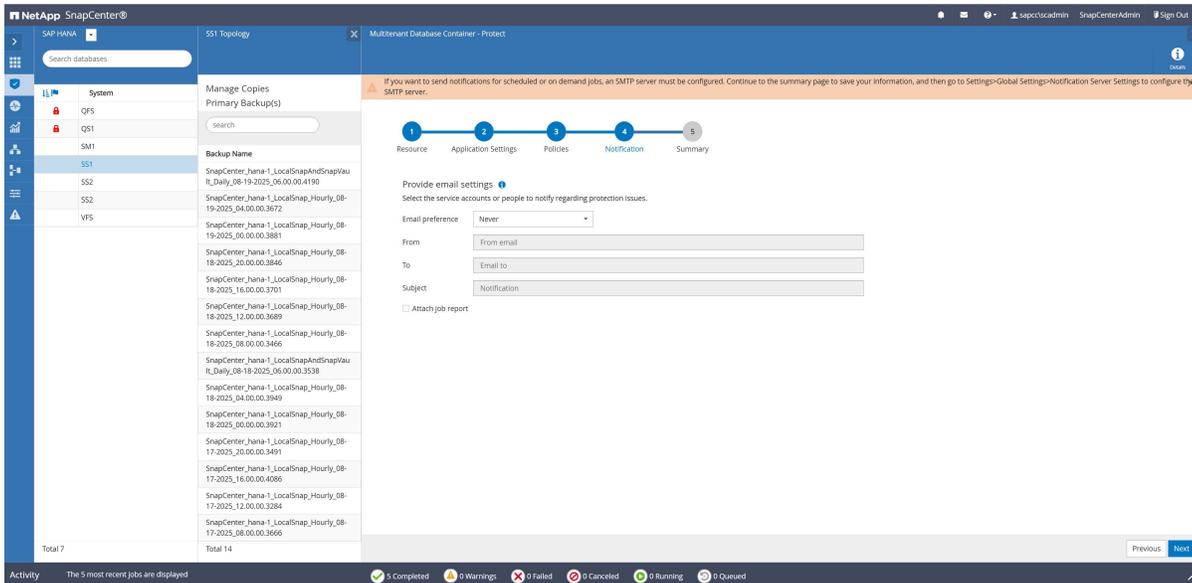
In this example we have configured

- 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.

Resource Group	System ID (SID)	Tenant Databases	Replication	System State	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
Single Container	QF5	QF5	None	Offline	hana-9.sapcc.stl.netapp.com				Not protected
Resource Group									
Q51	Q51	Q51 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 LocalSnapEndSnapVault LocalSnap-OnDemand	08/21/2025 6:01:04 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Primary)	Online	hana-3.sapcc.stl.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapKeep2	08/21/2025 6:57:25 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Secondary)	Online	hana-4.sapcc.stl.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapKeep2	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 LocalSnap-OnDemand	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.

Add SAP HANA Database ✕

1 Name

Provide Resource Details

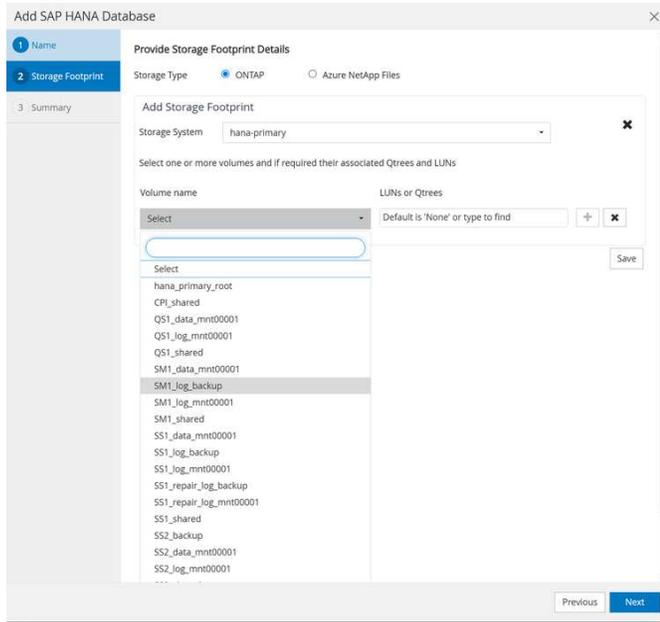
Resource Type:

Resource Name:

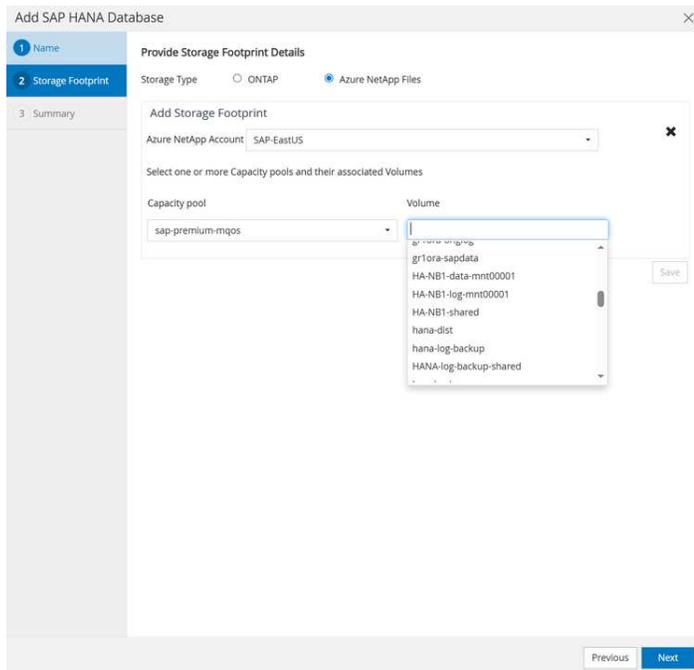
Associated SID:

Plug-in Host:

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:

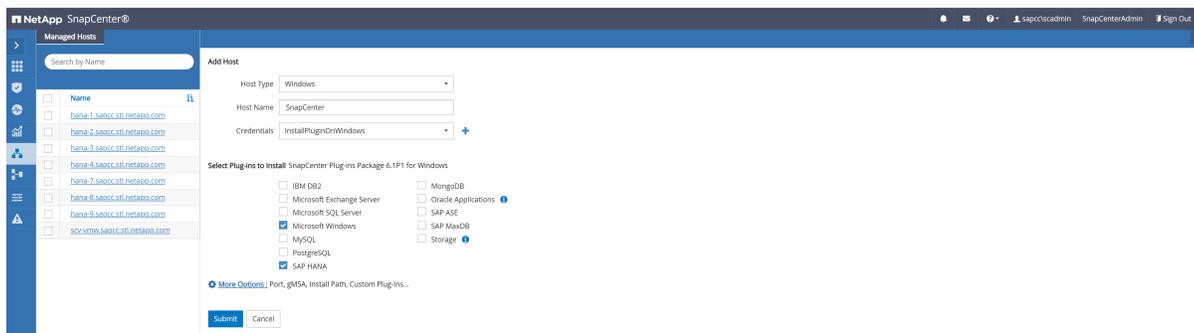
- SnapCenter HANA plug-in deployment
- SAP HANA hdbsql client installation and configuration
- SAP HANA user store configuration for each HANA system which is protected by the central plug-in host

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.

```

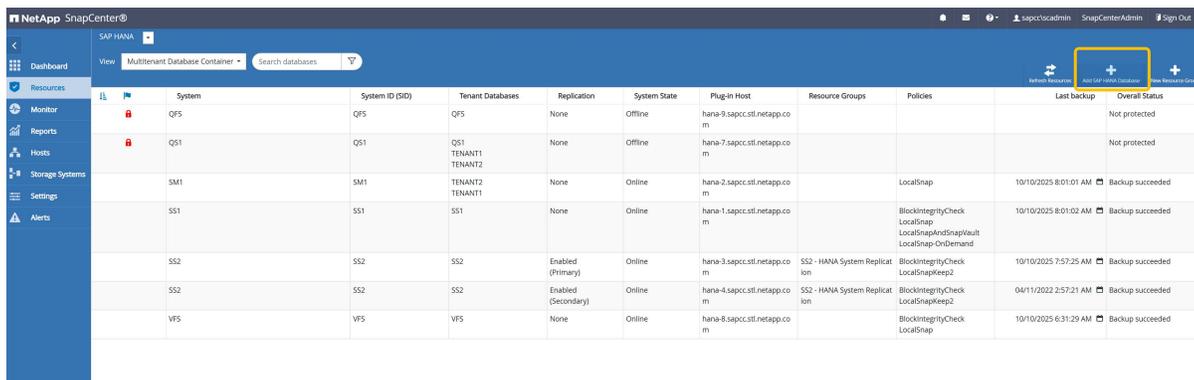
hdbuserstore.exe -u SYSTEM set MS1KEYHOST1 hana-4:30013 SNAPCENTER
password
hdbuserstore.exe -u SYSTEM set MS1KEYHOST2 hana-5:30013 SNAPCENTER
password
hdbuserstore.exe -u SYSTEM set MS1KEYHOST3 hana-6:30013 SNAPCENTER
password
C:\Program Files\sap\hdbclient>hdbuserstore.exe -u SYSTEM list
DATA FILE : 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\hdbclient>

```

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.stl.netapp.com				Not protected
Q51	Q51	Q51 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	10/10/2025 8:01:01 AM	Backup succeeded
SS1	SS1	SS1	None	Online	hana-1.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap LocalSnapAndSnapVault LocalSnap-OnDemand	10/10/2025 8:01:02 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Primary)	Online	hana-3.sapcc.stl.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapkeep2	10/10/2025 7:57:25 AM	Backup succeeded
SS2	SS2	SS2	Enabled (Secondary)	Online	hana-4.sapcc.stl.netapp.com	SS2 - HANA System Replication	BlockIntegrityCheck LocalSnapkeep2	04/11/2022 2:57:21 AM	Backup succeeded
VFS	VFS	VFS	None	Online	hana-8.sapcc.stl.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.

The screenshot shows the 'Add SAP HANA Database' wizard at the 'Provide Resource Details' step. The left sidebar has three steps: '1 Name' (selected), '2 Storage Footprint', and '3 Summary'. The main area contains the following fields:

- 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 field)

Information icons are present next to the SID, Plug-in Host, HDB Secure User Store Keys, and HDBSQL OS User fields. At the bottom right, there are 'Previous' and 'Next' buttons.

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.

The screenshot shows the 'Add SAP HANA Database' wizard at the 'Provide Storage Footprint Details' step. The left sidebar has three steps: '1 Name', '2 Storage Footprint' (selected), and '3 Summary'. The main area contains the following elements:

- Storage Type: ONTAP, Azure NetApp Files
- Add Storage Footprint dialog box:
 - Storage System: hana-primary
 - Select one or more volumes and if required their associated Qtrees and LUNs
 - Volume name: Select (dropdown)
 - LUNs or Qtrees: Default is 'None' or type to find (input field)
 - Save button
- Volume list (from the dropdown):
 - 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

At the bottom right, there are 'Previous' and 'Next' buttons.

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)

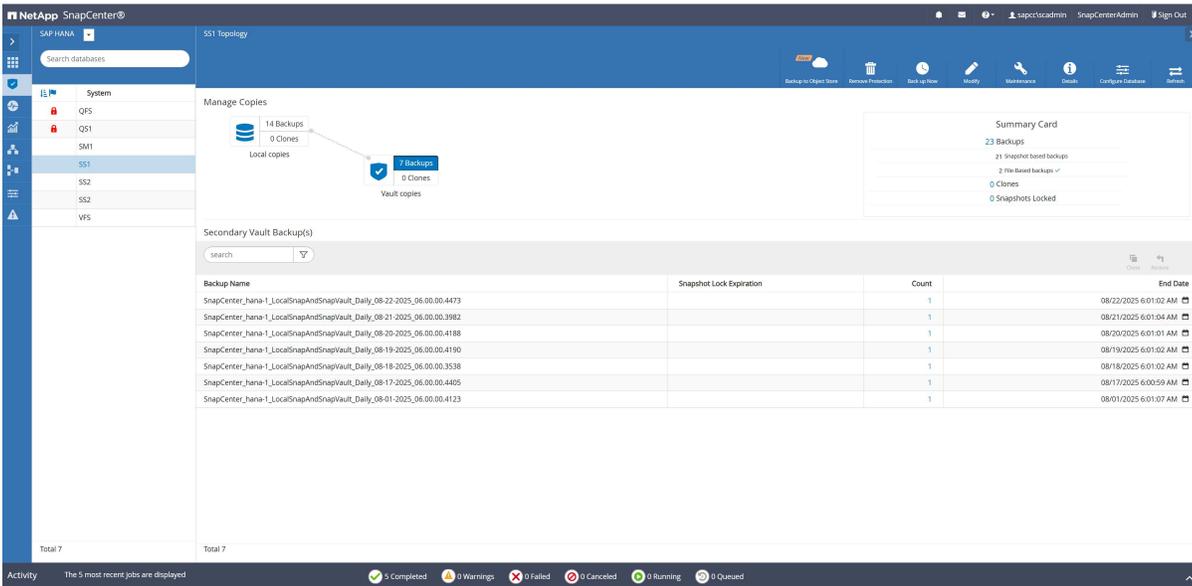
SAP HANA Snapshot backups in SnapCenter

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.

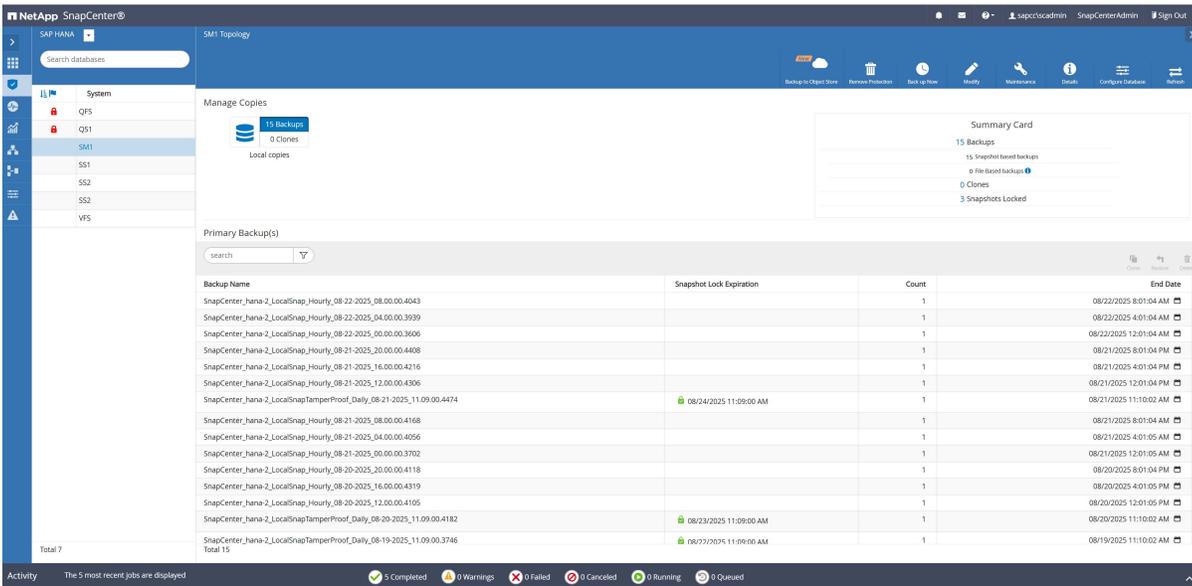
The screenshot displays the SnapCenter interface for a SAP HANA system. The 'SS1 Topology' section shows a diagram with 'Local copies' (14 Backups) and 'Vault copies' (7 Backups). A 'Summary Card' indicates 23 Backups, 21 Snapshot based backups, 3 File based backups, 0 Clones, and 0 Snapshots Locked. Below this is a table of 'Primary Backup(s)' with columns for Backup Name, Snapshot Lock Expiration, Count, and End Date. The most recent backup is highlighted with a blue box.

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_LocalSnapAndSnapVault_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 4:01:03 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-22-2025_00.00.00.3445		1	08/22/2025 12:01:02 AM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_20.00.00.3933		1	08/21/2025 8:01:01 PM
SnapCenter_hana-1_LocalSnap_Hourly_08-21-2025_16.00.00.4185		1	08/21/2025 8:01:02 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_LocalSnapAndSnapVault_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.



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



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.

Backup SYSTEMDB@SS1 (SYSTEM) SS1 - MDC single tenant

Database: SYSTEMDB

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 19s	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, 12:00:26 PM	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, 12:00:26 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, 4:00:25 AM	00h 00m 19s	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, 6:00:24 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

Backup Details:

ID: 1755864026172
 Status: Successful
 Backup Type: Data Backup
 Destination Type: Snapshot
 Started: Aug 22, 2025, 8:00:26 AM (America/Indianapolis)
 Finished: Aug 22, 2025, 8:00:44 AM (America/Indianapolis)
 Duration: 00h 00m 18s
 Size: 4.75 GB
 Throughput: n.a.
 System ID:
 Comment: SnapCenter:hana-1_LocalSnap_Hourly_08-22-2025_08.00.00.3884
 Additional Information: <ok>
 Location: /hana/data/SS1/mnt00001/

Host	Service	Size	Name	Sou...	EBID
hana-1	nameserver	4.75 GB	hdb00001	vol...	SnapCenter:hana-1_LocalSnap_Hourly...

Backup SYSTEMDB@SS1 (SYSTEM) SS1 - MDC single tenant

Database: SS1

Status	Started	Duration	Size	Backup Type	Destination Ty...
Success	Aug 22, 2025, 8:00:26 AM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 22, 2025, 6:00:26 AM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 22, 2025, 4:00:26 AM	00h 00m 19s	4.88 GB	Data Backup	Snapshot
Success	Aug 22, 2025, 12:00:26 A...	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 8:00:25 PM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 4:00:25 PM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 12:00:26 PM	00h 00m 19s	4.88 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 8:00:26 AM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 6:00:26 AM	00h 00m 19s	4.88 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 4:00:25 AM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 21, 2025, 12:00:26 A...	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 8:00:26 PM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 6:00:25 PM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 12:00:26 PM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 20, 2025, 6:00:25 AM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 19, 2025, 6:00:26 AM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 18, 2025, 6:00:26 AM	00h 00m 18s	4.88 GB	Data Backup	Snapshot
Success	Aug 18, 2025, 6:00:24 AM	00h 00m 17s	4.88 GB	Data Backup	Snapshot
Success	Aug 17, 2025, 6:00:24 AM	00h 00m 13s	3.53 GB	Data Backup	File
Success	Aug 17, 2025, 5:32:10 AM	00h 00m 13s	3.53 GB	Data Backup	File
Success	Aug 1, 2025, 6:00:27 AM	00h 00m 19s	4.88 GB	Data Backup	Snapshot

Backup Details:

ID: 1755864026173
 Status: Successful
 Backup Type: Data Backup
 Destination Type: Snapshot
 Started: Aug 22, 2025, 8:00:26 AM (America/Indianapolis)
 Finished: Aug 22, 2025, 8:00:44 AM (America/Indianapolis)
 Duration: 00h 00m 18s
 Size: 4.88 GB
 Throughput: n.a.
 System ID:
 Comment: SnapCenter:hana-1_LocalSnap_Hourly_08-22-2025_08.00.00.3884
 Additional Information: <ok>
 Location: /hana/data/SS1/mnt00001/

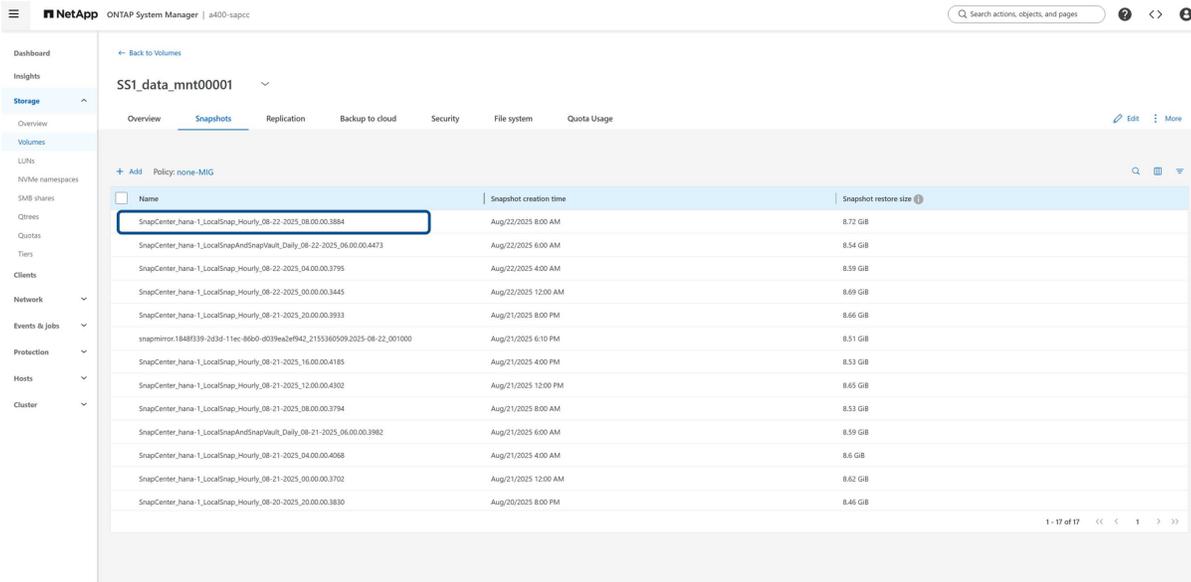
Host	Service	Size	Name	Sou...	EBID
hana-1	xsengine	320.00 MB	hdb00002...	vol...	SnapCenter:hana-1_LocalSnap_Hourly...
hana-1	indexserver	4.56 GB	hdb00003...	vol...	SnapCenter:hana-1_LocalSnap_Hourly...



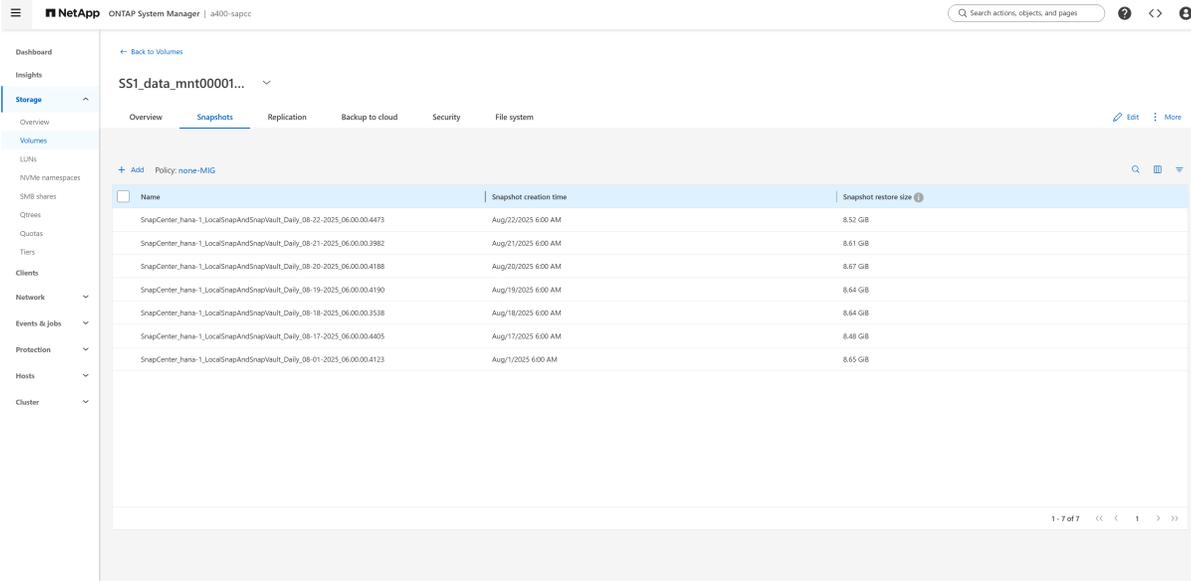
SnapCenter is only aware of its own backups. Additional backups created, for example, with SAP HANA Studio, are visible in the SAP HANA catalog but not in SnapCenter. Also Snapshots created directly on the storage system will not be visible in SnapCenter,

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 SAP HANA Studio in the previous images and has the same naming convention.

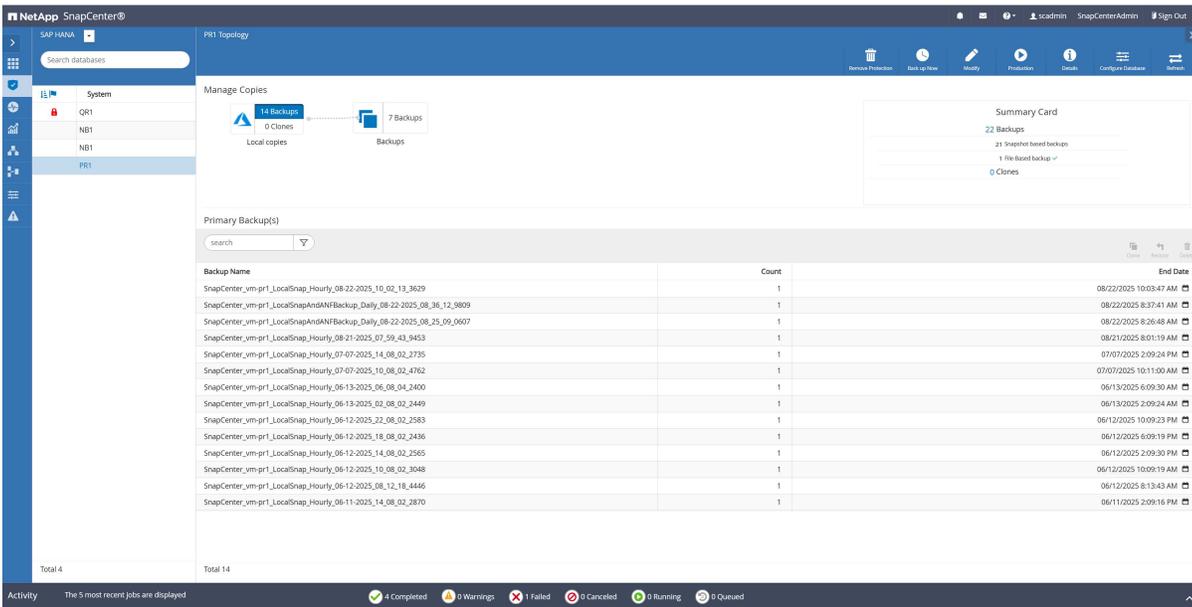


The following screenshot shows the available backups for the replication target volume hana_SS1_data_mnt00001_dest at the secondary storage system.

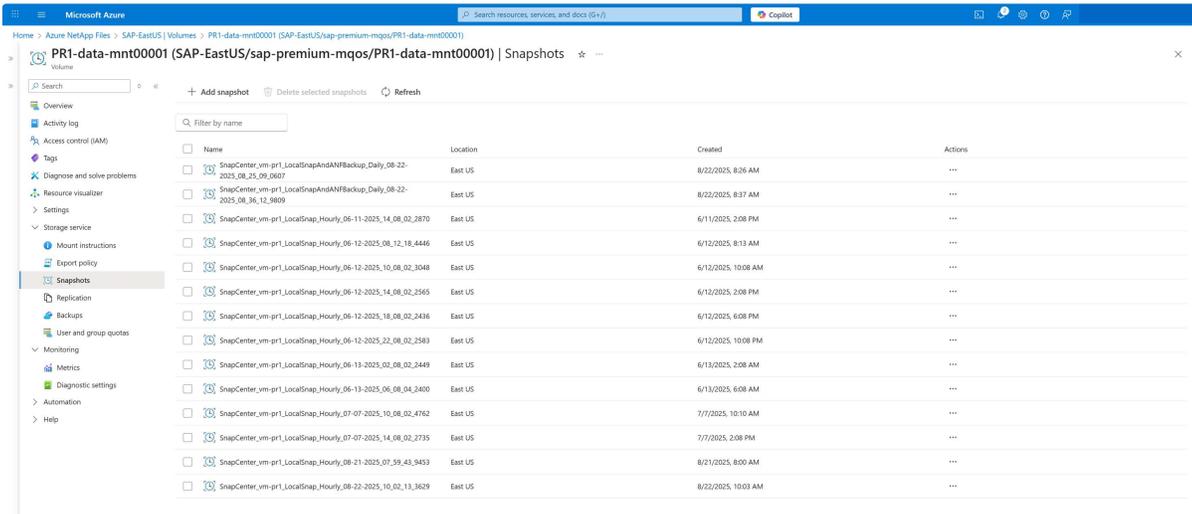


SAP HANA Snapshot backups with ANF

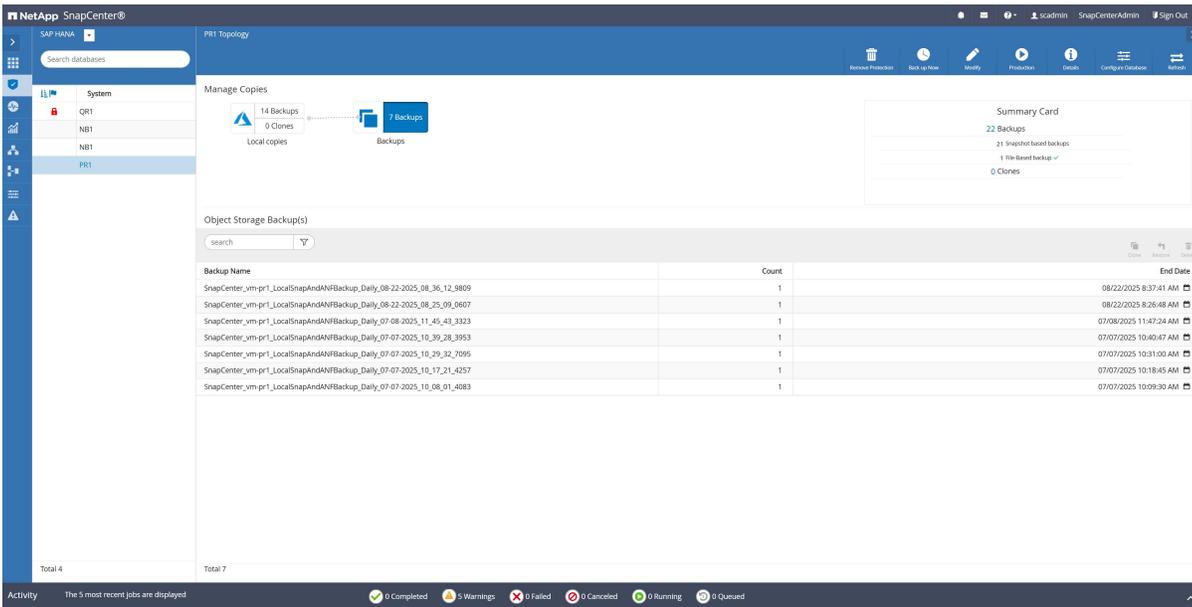
The following screenshot shows the topology view of a HANA system using Azure NetApp Files. For this HANA system local Snapshot backups as well as backup replication using ANF backup has been configured.



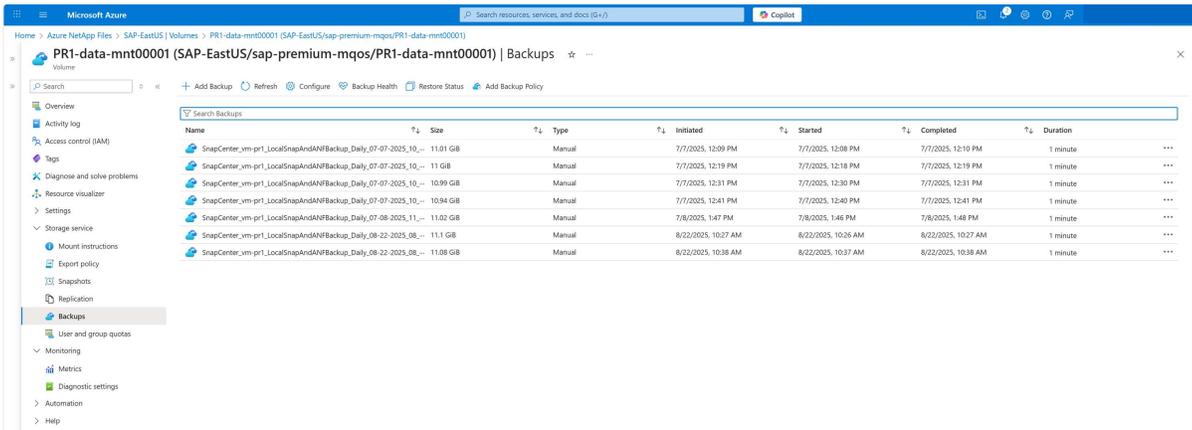
Snapshot backups on the ANF volume can be listed using the Azure portal.



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

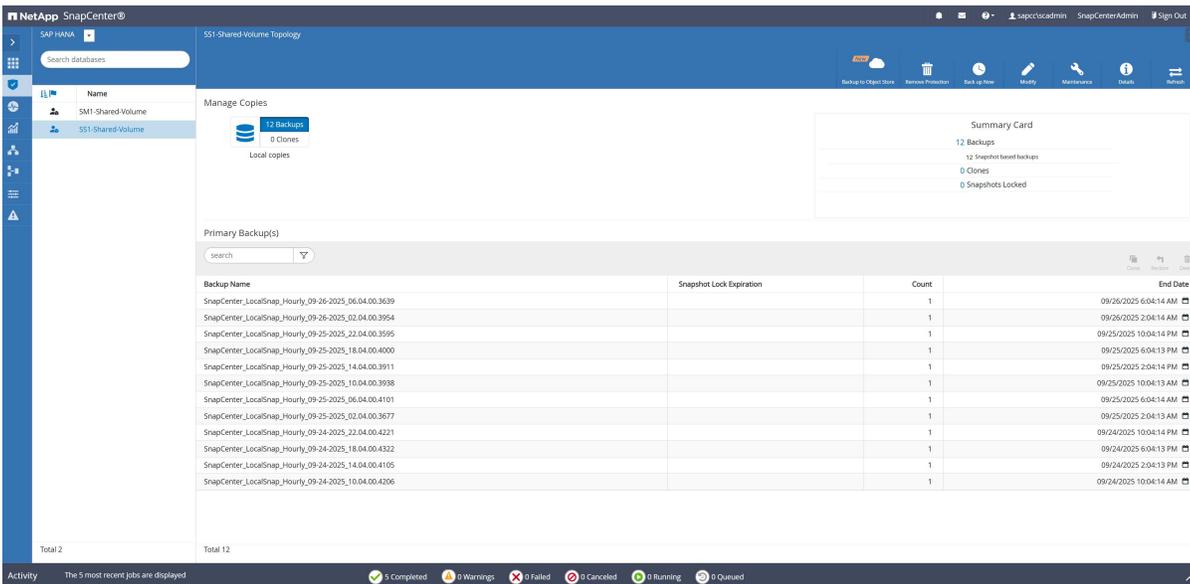


ANF backups can also be listed in the Azure portal.



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 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

Job Details

Backup of Resource Group 'hana-1_sapcc_stl_ne.....na_MDC_SS1' with policy 'LocalSnap'

- ✓ Backup of Resource Group 'hana-1_sapcc_stl_netapp_com_hana_MDC_SS1' with policy 'LocalSnap'
- ✓ hana-1.sapcc.stl.netapp.com
 - ✓ Backup
 - ✓ ▶ Validate Dataset Parameters
 - ✓ ▶ Validate Plugin Parameters
 - ✓ ▶ Complete Application Discovery
 - ✓ ▶ Initialize Filesystem Plugin
 - ✓ ▶ Discover Filesystem Resources
 - ✓ ▶ 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
 - ✓ ▶ Finalize Filesystem Plugin
 - ✓ ▶ Collect Autosupport data
 - ✓ ▶ Register Backup and Apply Retention
 - ✓ ▶ Register Snapshot attributes
 - ✓ ▶ Application Clean-Up
 - ✓ ▶ Data Collection
 - ✓ ▶ Agent Finalize Workflow

Auto discovery operations
HANA database snapshot, storage Snapshot
Retention management

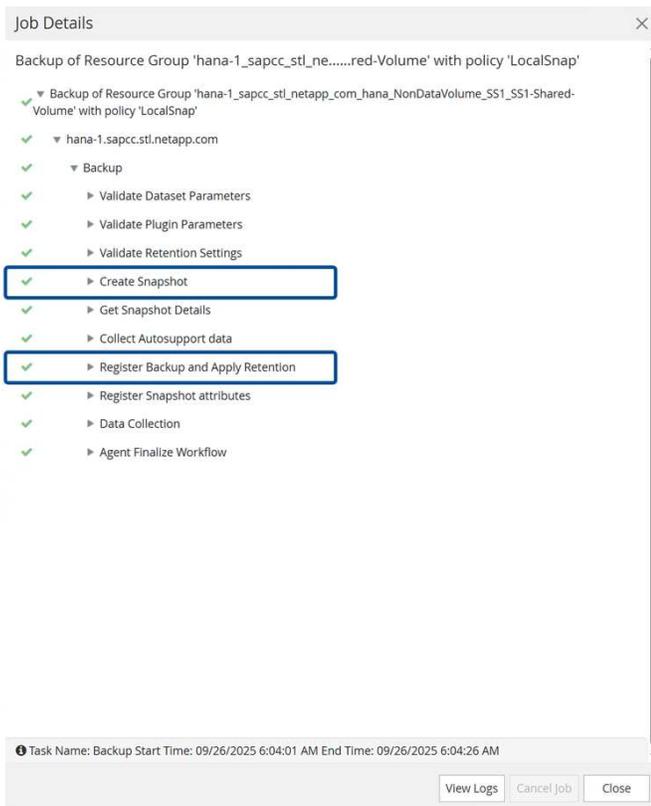
▶ (Job 156547) Uncataloging Backup(s) SnapCenter_hana-1_LocalSnap_Hourly_08-22-2025_08.00.00.3884

Task Name: Backup Start Time: 08/24/2025 8:00:01 AM End Time: 08/24/2025 8:01:30 AM

View Logs Cancel Job Close

Backup workflow for non-data volumes

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

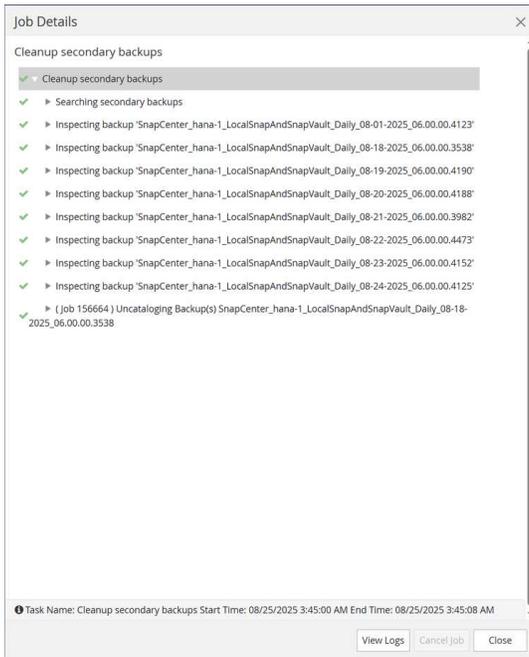


Cleanup of secondary backups

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.....na_MDC_SMI' with policy 'LocalSnap'	08/25/2025 4:00:00 AM	8/25/2025 4:01:35 AM	SAPCC'scadmin
156669	✓	Backup of Resource Group 'hana-1_sapcc_stl_ne.....na_MDC_SS1' with policy 'LocalSnap'	08/25/2025 4:00:00 AM	8/25/2025 4:01:32 AM	SAPCC'scadmin
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	SAPCC'scadmin
156663	✓	Cleanup secondary backups	08/25/2025 3:45:00 AM	08/25/2025 3:45:08 AM	SAPCC'scadmin
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	SAPCC'scadmin
156657	✓	Backup of Resource Group 'hana-8_sapcc_stl_ne.....na_MDC_VFS' with policy 'LocalSnap'	08/25/2025 2:29:00 AM	08/25/2025 2:30:43 AM	SAPCC'scadmin
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	SAPCC'scadmin
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	SAPCC'scadmin



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
ReplaceJobIfExists : 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.

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 5:04:00 AM	Hourly	SS1 - Shared Volume	Backup	0h 0m 26s		Completed	Active
08/30/2025 6:39:00 AM	Weekly	SS2 - HANA System Replication	Backup	0h 1m 33s		Completed	Active
08/30/2025 6:39:00 AM	Weekly	SS2 - HANA System Replication	Backup	0h 1m 33s		Completed	Active
08/25/2025 4:56:00 AM	Hourly	SS2 - HANA System Replication	Backup	0h 1m 54s		Completed	Active
08/25/2025 4:56:00 AM	Hourly	SS2 - HANA System Replication	Backup	0h 1m 54s		Completed	Active
08/25/2025 6:06:00 AM	Hourly	SM1 - Shared Volume	Backup	0h 0m 29s		Completed	Active
08/25/2025 8:00:00 AM	Hourly	SM1	Backup	0h 1m 35s		Completed	Active
08/25/2025 11:09:00 AM	Daily	SM1	Backup			Warning	Active
08/25/2025 7:06:00 AM	Daily	SnapCenterDataCollection	Predefined			Completed	Active
08/31/2025 2:59:00 AM	Weekly	SnapCenter_RefreshBackupSnapLock	Predefined	0h 1m 3s		Completed	Active
08/26/2025 12:00:00 AM	Daily	SnapCenter_ComputeStorageSavings	Predefined	0h 0m 2s		Completed	Active
08/25/2025 11:59:00 PM	Daily	SnapCenter_AlertCleanup	Predefined			Completed	Active
08/25/2025 11:59:00 PM	Daily	SnapCenter_AuditDiskSpaceCheck	Predefined	0h 0m 0s		Completed	Active
08/31/2025 4:59:00 AM	Weekly	SnapCenter_AuditIntegrityCheck	Predefined			Completed	Active
08/31/2025 12:59:00 AM	Weekly	SnapCenter_RefreshSyncSnapMirrorBackups	Predefined	0h 0m 0s		Completed	Active
08/31/2025 4:30:00 PM	Weekly	SnapCenterJobRetention	Purgopolis			Completed	Active
08/25/2025 11:59:00 PM	Daily	SnapCenter_StateJobCleanup	Predefined			Completed	Active
08/25/2025 3:46:00 AM	Daily	SnapCenter_RemoveSecondaryBackup	Predefined	0h 0m 36s		Completed	Active
08/25/2025 7:00:00 AM	Hourly	SS1	Backup	0h 1m 33s		Completed	Active
08/25/2025 5:00:00 AM	Daily	SS1	Backup	0h 2m 13s		Completed	Active
08/31/2025 5:32:00 AM	Weekly	SS1	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.

The screenshot shows the 'SS1 Topology' view in SnapCenter. On the left, a sidebar lists resources: QFS, QS1, SM1, SS1 (selected), SS2, SS2, and VFS. The main area is divided into 'Manage Copies' and 'Secondary Vault Backups(s)'. 'Manage Copies' shows 14 Backups (0 Clones) for Local copies and 7 Backups (0 Clones) for Vault copies. A 'Summary Card' on the right indicates 23 Backups, 21 Snapshot based backups, 2 File-based backups, 0 Clones, and 0 Snapshots Locked. Below, a table lists 'Secondary Vault Backups(s)' with columns for Backup Name, Snapshot Lock Expiration, Count, and End Date.

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

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

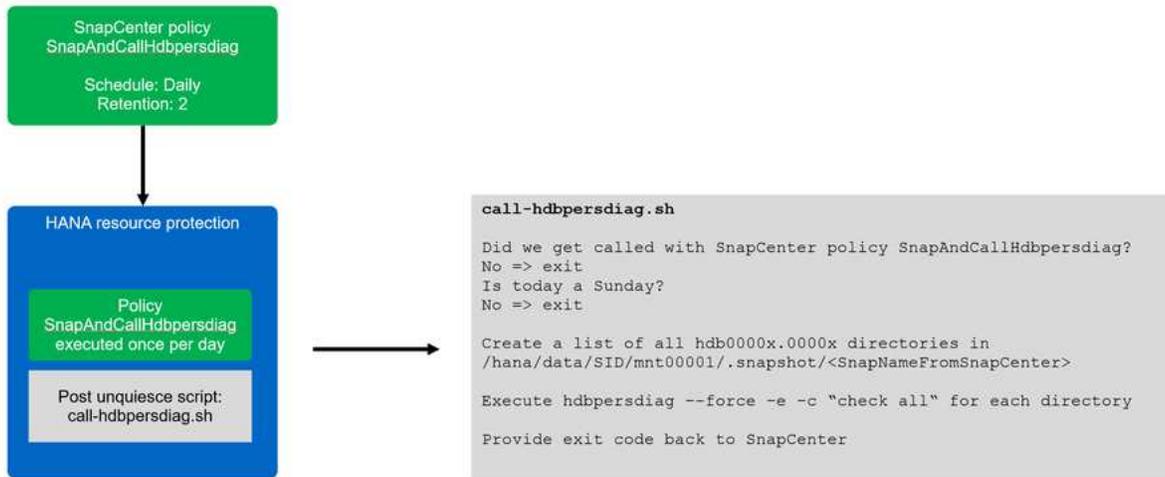
Within SnapCenter a dedicated policy for hdbpersdiag operations is created with a daily schedule and a retention of two. We don't use the weekly schedule, since we would then have at least 2 Snapshot backups (minimum retention=2), where one of them would be up to two weeks old.

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 Provide a policy name

Policy name

Details

2 Policy type

3 Snapshot And Replication

4 Summary

Modify SAP HANA Backup Policy

1 Name

2 Policy type

Choose storage type

ONTAP/FSx/Cloud volumes ONTAP Azure NetApp Files

Choose policy scope

Snapshot Based File-Based

3 Snapshot And Replication

4 Summary

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

Retain copies for: days

Primary snapshot copy locking period: days

Secondary snapshot copy locking period: days

Policy label:

Select secondary replication options

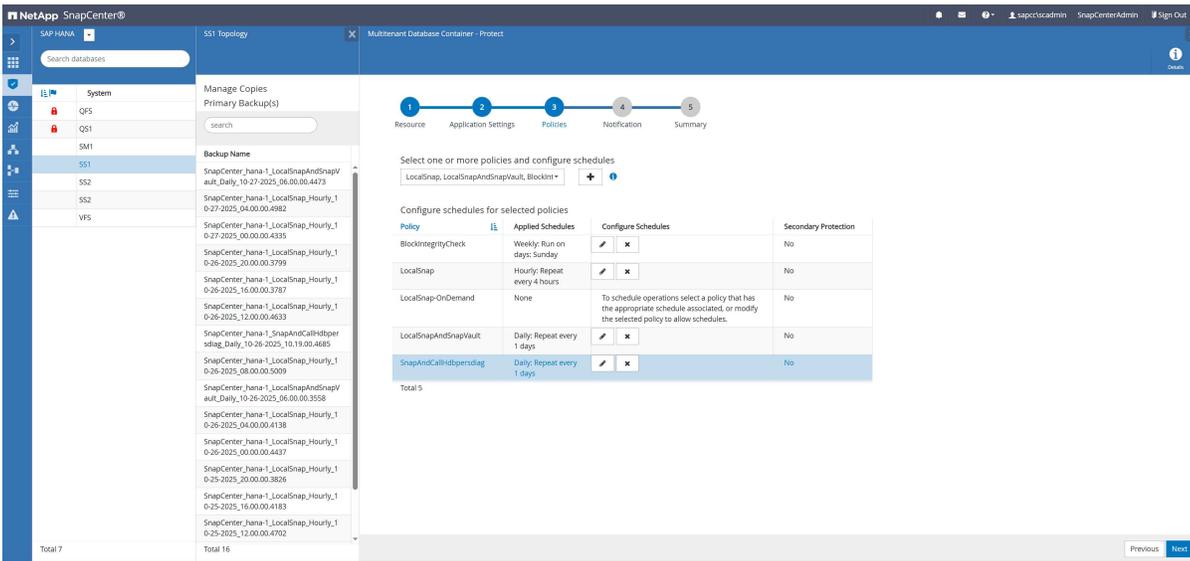
Update SnapMirror after creating a local Snapshot copy.
 Update SnapVault after creating a local Snapshot copy.

Error retry count:

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

The screenshot shows the NetApp SnapCenter interface for configuring a resource protection policy. The left sidebar lists databases including QFS, QFS1, SM1, S51, S52, and VFS. The central pane displays a list of backup jobs with columns for Backup Name and Total. The right pane shows the configuration steps: 1. Resource, 2. Application Settings, 3. Policies, 4. Notification, and 5. Summary. The 'Policies' step is active, showing fields for Post Quiesce, Pre Snapshot Copy, Post Snapshot Copy, Pre UnQuiesce, Post UnQuiesce, Pre Exit, and Custom Configurations. A status bar at the bottom indicates 4 Completed, 0 Warnings, 0 Failed, 0 Canceled, 1 Running, and 0 Queued jobs.

resource.



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 'libhdblivercache'
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
succesful 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 'libhdblivercache'

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
successeful for volume /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: Executing hdbpersdiag in:
/hana/data/SS1/mnt00001/.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 'libhdblivercache'
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/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
successeful for volume /hana/data/SS1/mnt00001/.snapshot/SnapCenter_hana-
1_SnapAndCallHdbpersdiag_Daily_10-24-2025_05.57.37.0274/hdb00003.00003.
20251024060048###hana-1###call-hdbpersdiag.sh: Current policy is
```

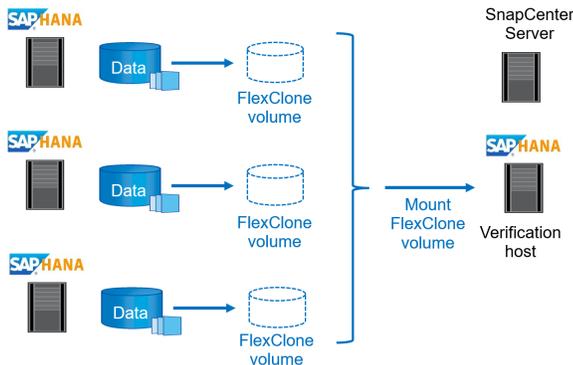
```
LocalSnapAndSnapVault, consistency check is only done with Policy
SnapAndCallHdbpersdiag
20251024080048###hana-1###call-hdbpersdiag.sh: Current policy is
LocalSnap, consistency check is only done with Policy SnapAndHdbpersdiag
```

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](#)



1.) SnapCenter clone create operation

- Clone server = Verification host
- SnapCenter creates the FlexClone and mounts it to the verification host
- Post Clone Command Script
 - Create a list of hdb0000x directories
 - Execute hdbpersdiag -c 'check all' for each directory

2.) SnapCenter clone delete operation

- SnapCenter unmounts the FlexClone from the verification host
- SnapCenter deletes the FlexClone

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-flexclone.sh
```



The example script call-hdbpersdiag-flexclone.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.

Backup Name	Snapshot Lock Expiration	Count	End Date
SnapCenter_hana-3_LocalSnapKeep2_Hourly_10-29-2025_10:56:27.1217		1	10/29/2025 10:57:26 AM
SnapCenter_hana-3_LocalSnapKeep2_Hourly_10-29-2025_09:56:27.3665		1	10/29/2025 9:57:25 AM

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.

Clone From Backup ✕

1 Location

2 Scripts

3 Notification

4 Summary

Select the host to create the clone

Plug-in host: ⓘ

Target Clone SID: ⓘ

NFS Export IP Address: ⓘ

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

Clone From Backup
✕

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

2 Scripts
Enter optional commands to run before performing a clone operation ⓘ

3 Notification

4 Summary

Pre clone command

Enter optional commands to run after performing a clone operation ⓘ

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
```

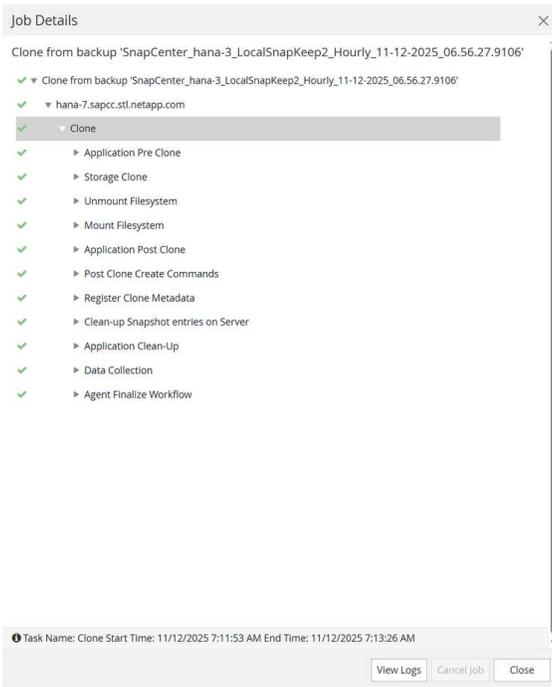
```

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 'libhdblivercache'
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.

```

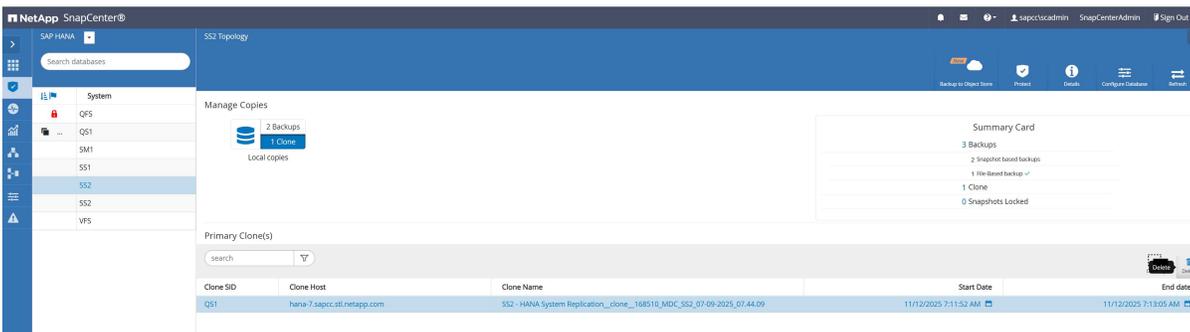


The script calls hdbpersdiag with the “-e” command line option which is required for data volume encryption. If HANA data volume encryption is not used the parameter must be removed. When the post clone script is finished the SnapCenter job is finished as well.



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.stl.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.stl.netapp.com]
Found backup name [SnapCenter_hana-3_LocalSnapKeep2_Hourly_11-21-
2025_07.56.27.5547]
Creating clone from backup [hana-
3.sapcc.stl.netapp.com/SS2/SnapCenter_hana-3_LocalSnapKeep2_Hourly_11-21-
2025_07.56.27.5547]: [hana-7.sapcc.stl.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 'libhdblivercache'

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

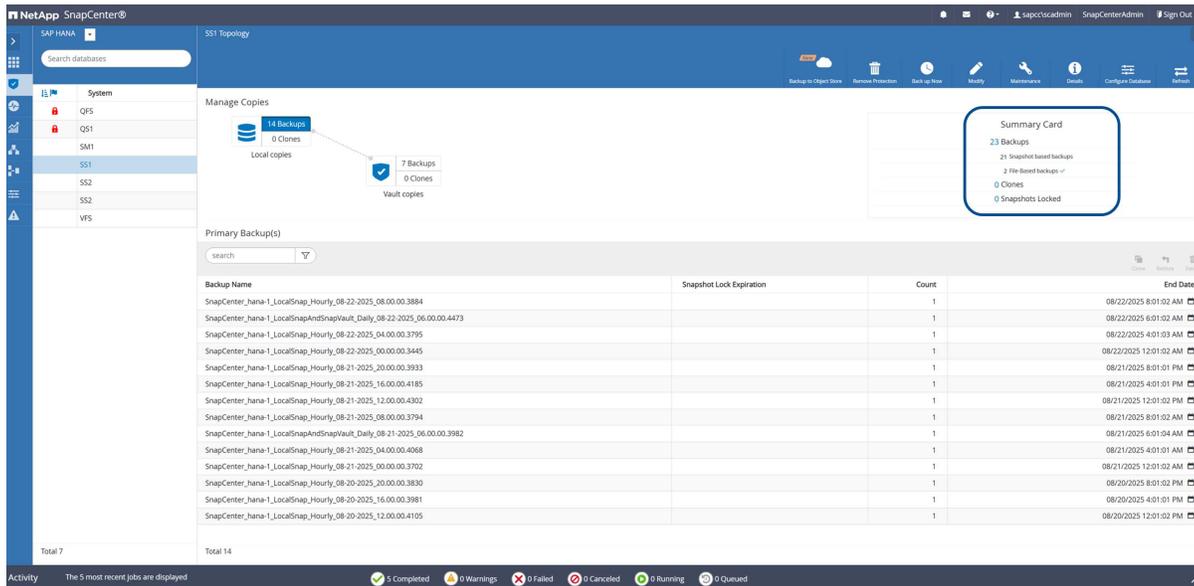
```
operation successful for volume /hana/data/QS1/mnt00001/hdb00003.00003.
hana-7:/mnt/sapcc-share/hdbpersdiag #
```

File-based backup

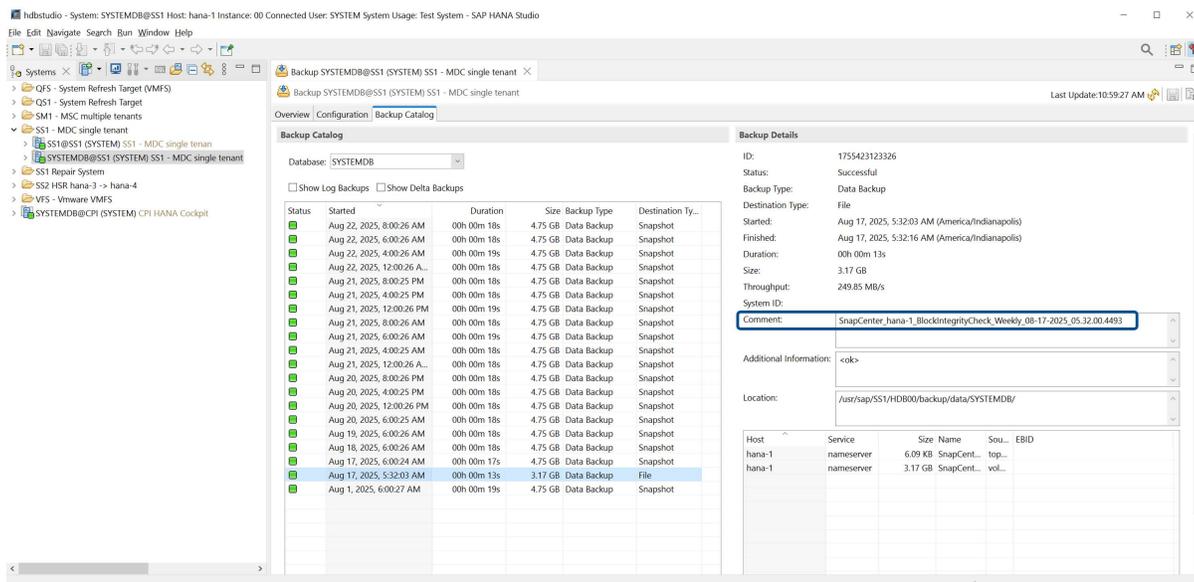
SnapCenter supports the execution of a block integrity check by using a policy in which file-based backup is selected as the backup type.

When scheduling backups using this policy, SnapCenter creates a standard SAP HANA file backup for the system and all tenant databases.

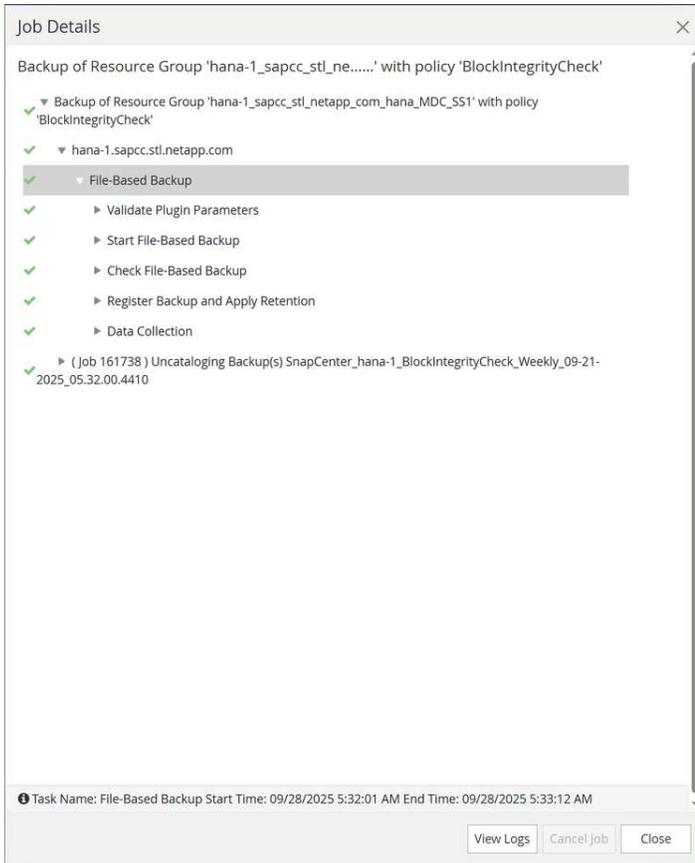
SnapCenter does not display the block integrity check in the same manner as Snapshot copy-based backups. Instead, the summary card shows the number of file-based backups and the status of the previous backup.



The SAP HANA backup catalog shows entries for both the system and the tenant databases. The following figure shows a SnapCenter block integrity check in the backup catalog of the system database.



A successful block integrity check creates standard SAP HANA data backup files.



SnapCenter uses the backup path that has been configured in the HANA database for file-based data backup operations.

```

hana-1:/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.



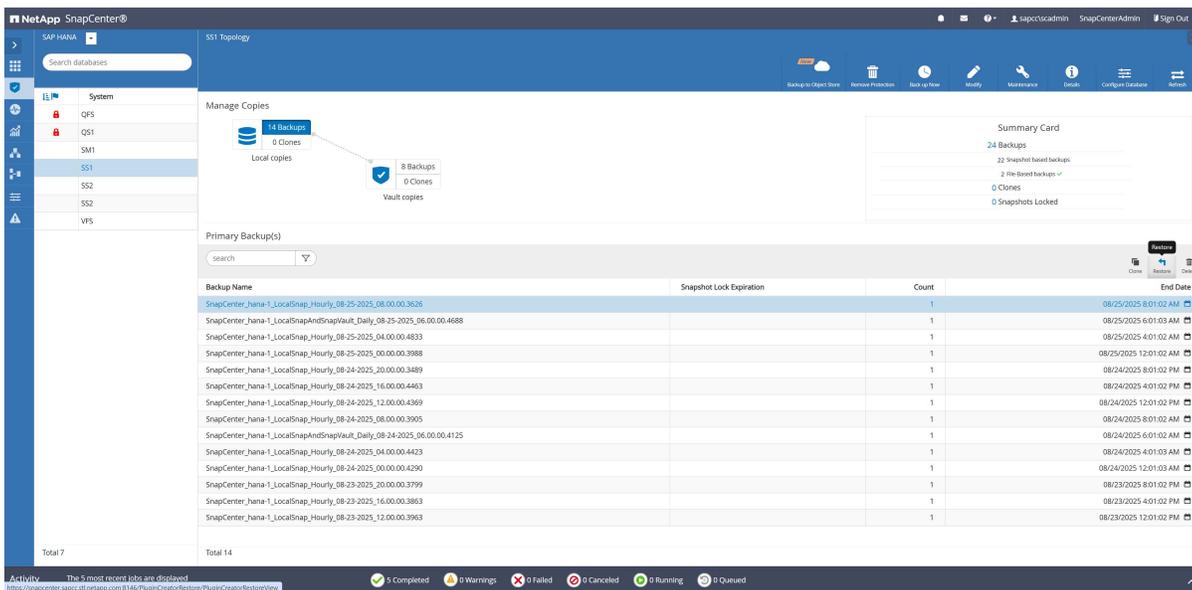
With Azure NetApp Files, restore operations are supported for primary ANF volumes or ANF backups. For primary ANF volumes a volume revert, for ANF backups an in-place restore with single file restore will be executed. In both cases application volume group configurations will be maintained.



If volume encryption is enabled and SAP local secure store (LSS) is used, a recovery with SnapCenter is supported as long as the root key backup password in the LSS has not been changed since the backup has been taken. If the password got changed and a restore and recovery is done using an older Snapshot with a different password, the recovery needs to be done manually and the old password needs to be provided with the recovery statement: "RECOVER DATA USING SNAPSHOT CLEAR LOG ENCRYPTION ROOT KEYS BACKUP PASSWORD 'old-password'"

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.



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.

Restore from SnapCenter_hana-1_LocalSnap_Hourly_08-26-2025_04.00.00.3631

1 Restore scope

2 Recovery scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Select the restore types

Complete Resource ⓘ

Volume Revert

Tenant Database

Previous Next

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.

Restore from SnapCenter_hana-8_LocalSnap_Hourly_08-26-2025_02.29.00.4210

1 Restore scope

2 Recovery scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Select the restore types

Complete Resource ⓘ

Tenant Database

Previous Next

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.

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

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

Enter optional commands to run before performing a restore operation ⓘ

Pre restore command

Previous Next

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

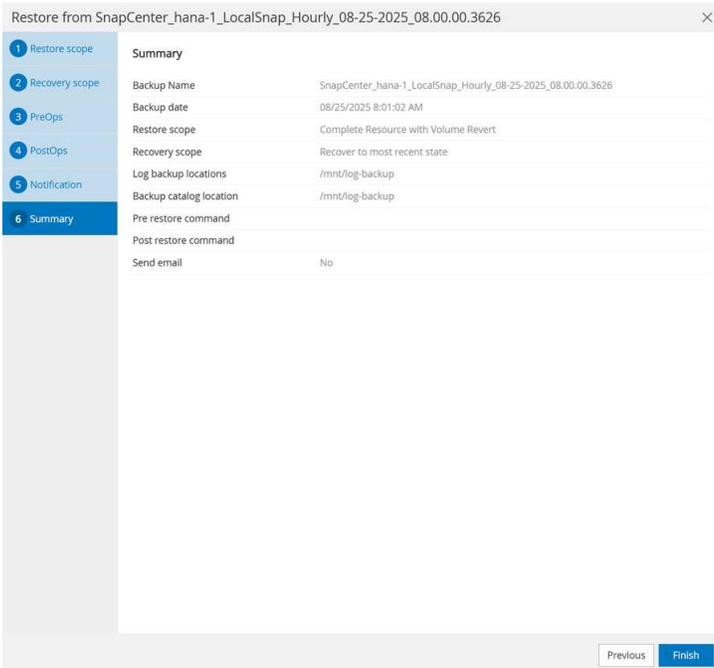
- 1 Restore scope
- 2 Recovery scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Enter optional commands to run after performing a restore operation ⓘ

Post restore command

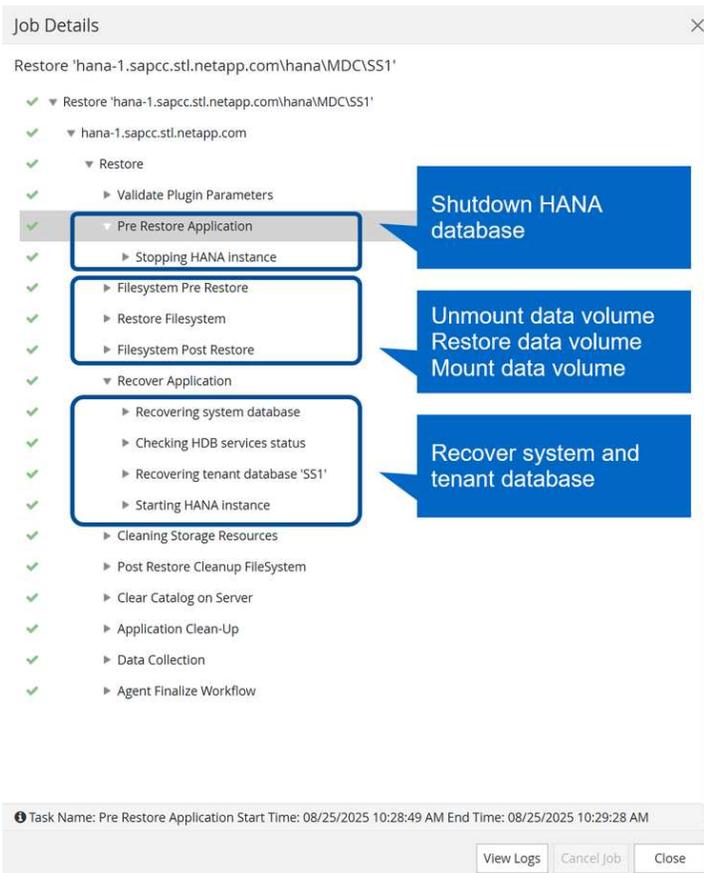
Previous Next

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



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.

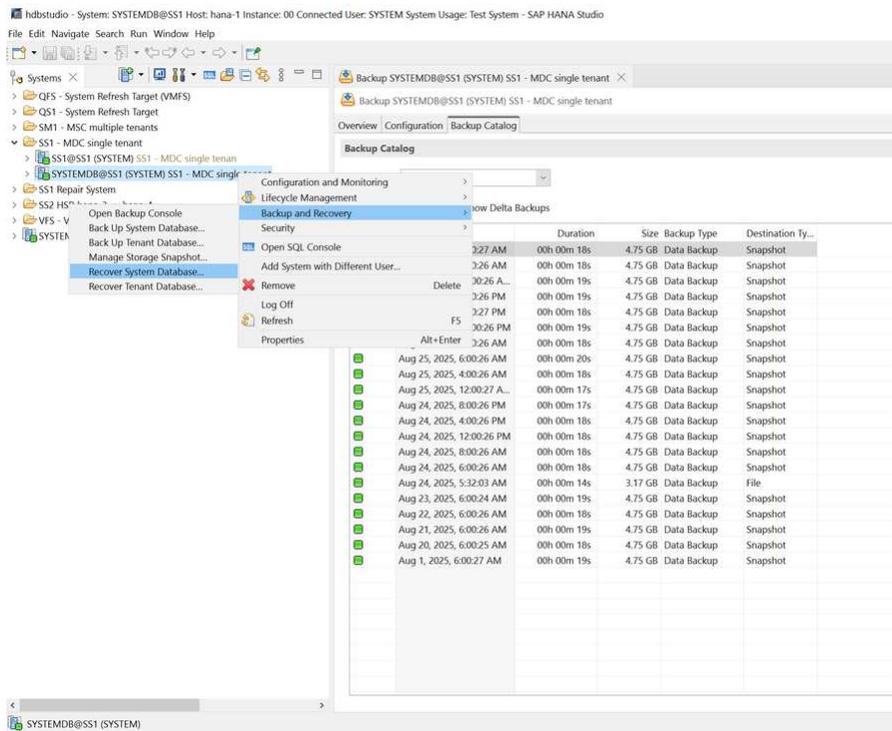
- b. Select the recovery type and the log backup location.
 - c. A list of data backups displays. Because the data volume has already been restored, the tenant backup is indicated as available (in green).
 - d. Select this backup and start the recovery process. After the recovery process is finished,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



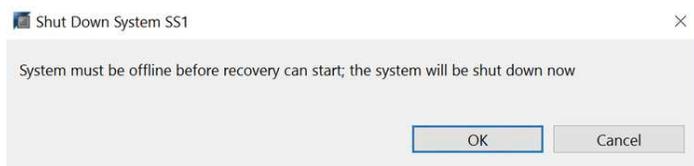
A manual recovery with SAP HANA Cockpit is done with the same steps.

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

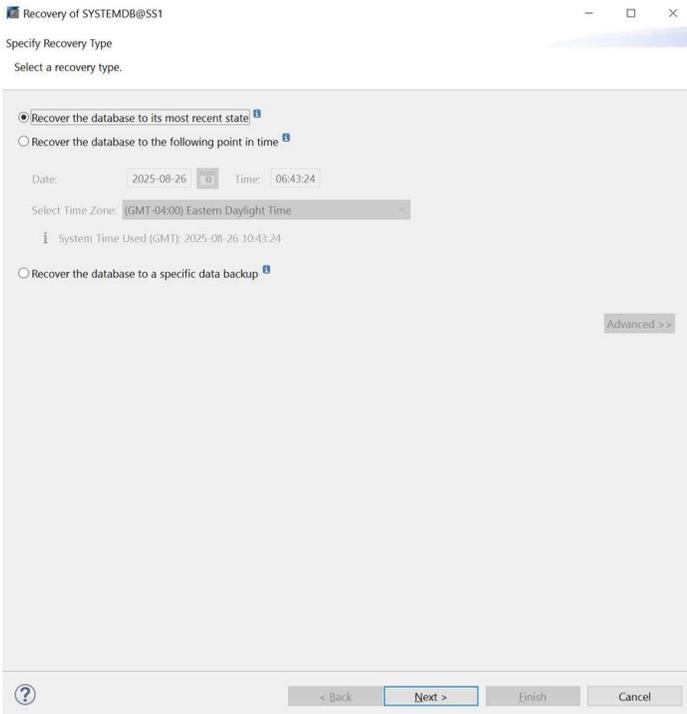
In HANA Studio select Backup and Recovery and Recover System Database.



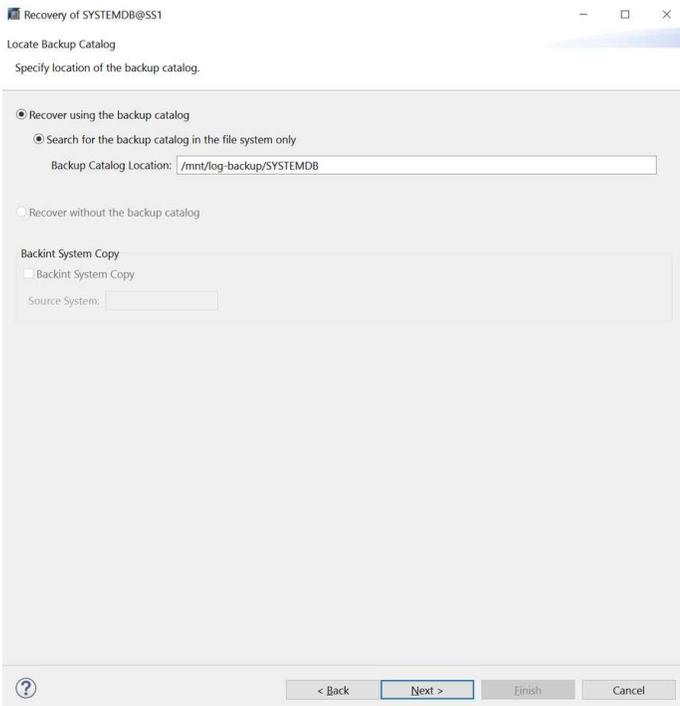
Confirm shutdown operation; only required if the HANA system is still running.



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

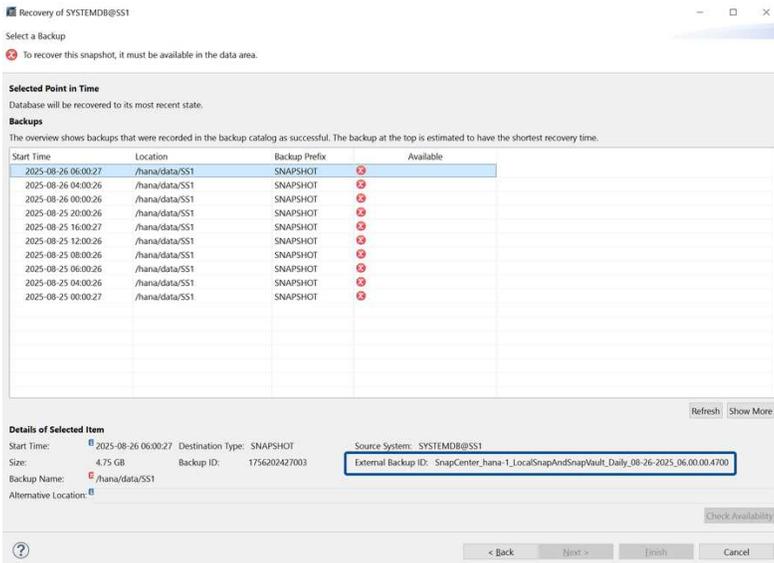


Provide backup catalog location.

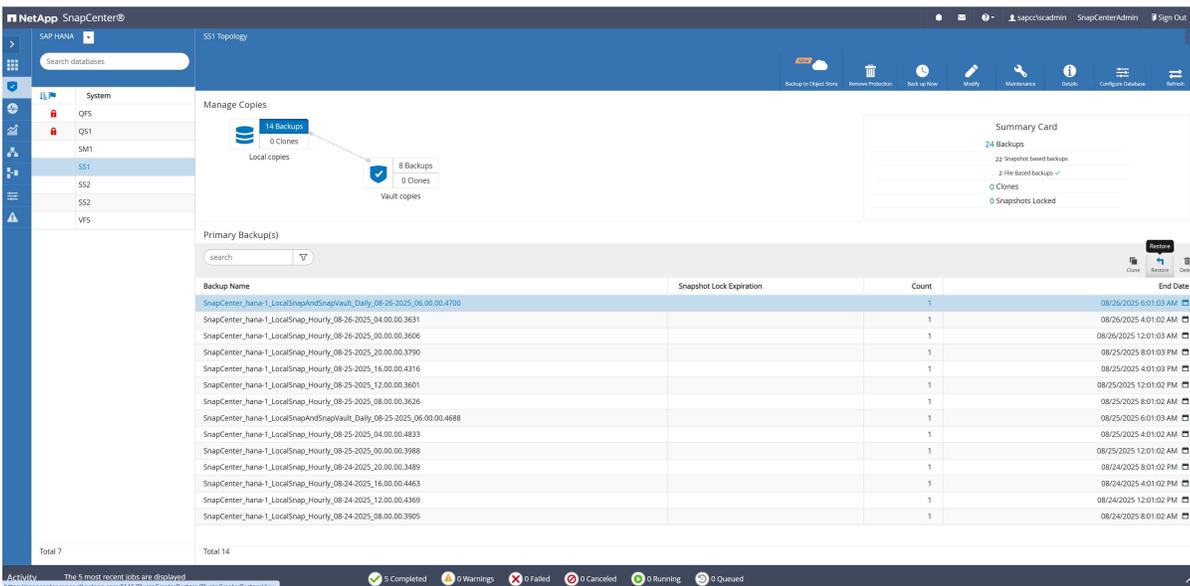


HANA Studio lists the most recent backups stored in the HANA backup catalog.

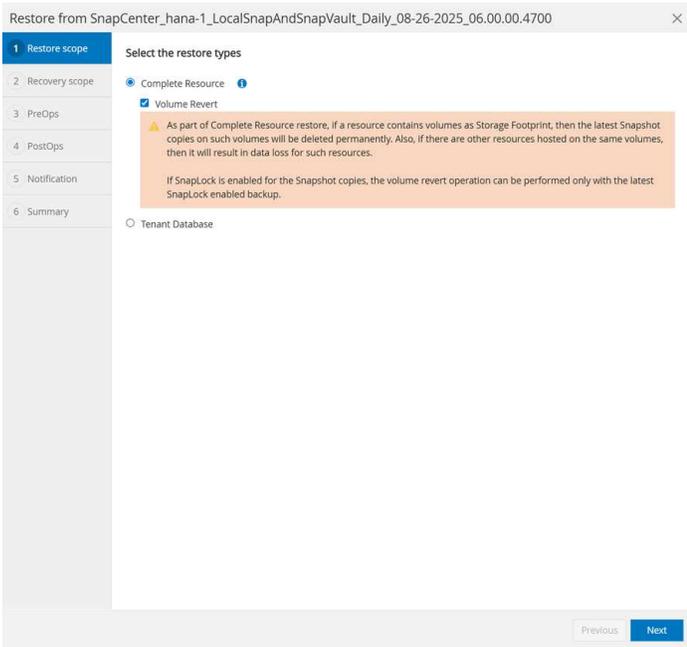
A list of available backups is shown based on the content of the backup catalog. Choose the required backup and note the external backup ID: in this example, the most recent backup.



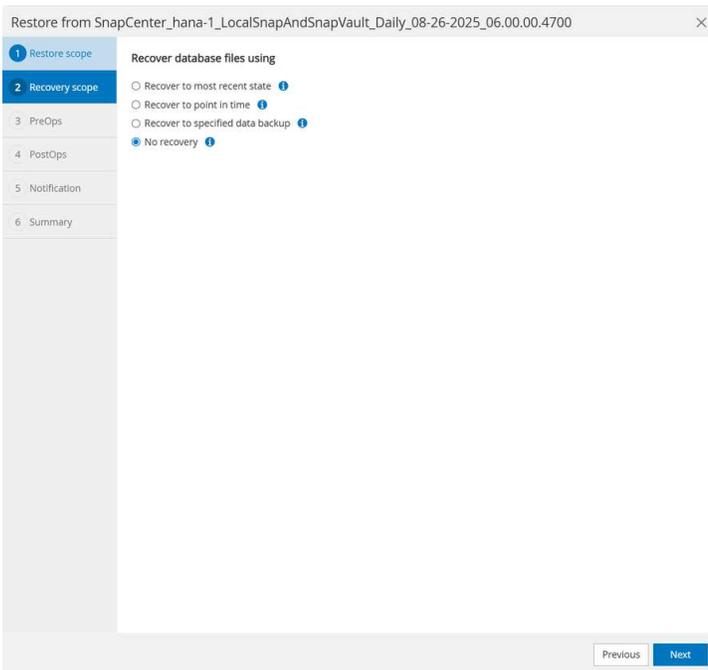
From the SnapCenter GUI, select the resource topology view and select the backup that should be restored, in this example, the most recent primary backup. Click the Restore icon to start the restore.



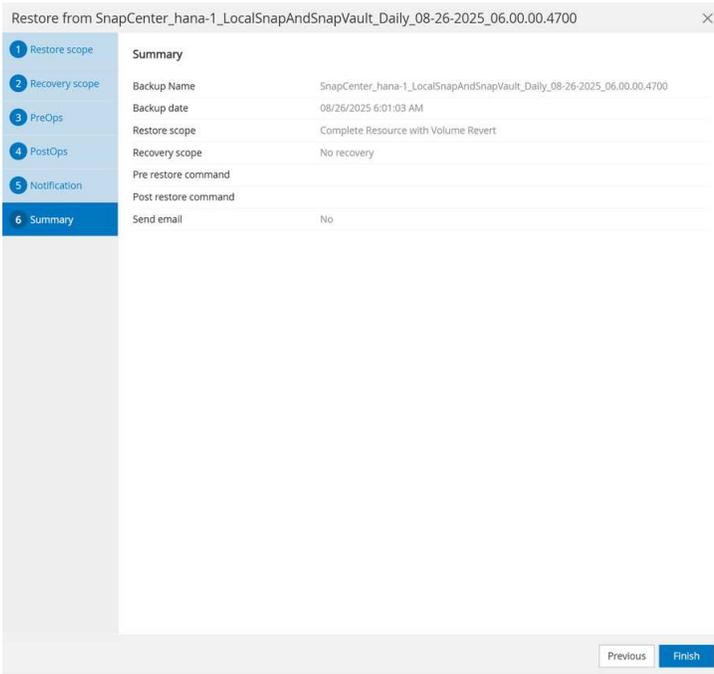
The SnapCenter restore wizard starts. Select the restore type Complete Resource and Volume revert to use a volume-based restore.



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

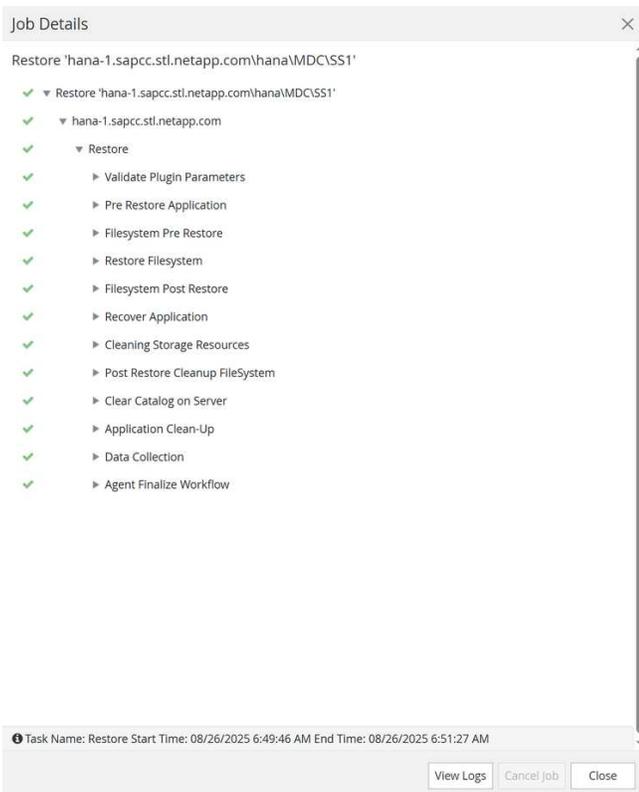


Click on Finish to start the restore operation.



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

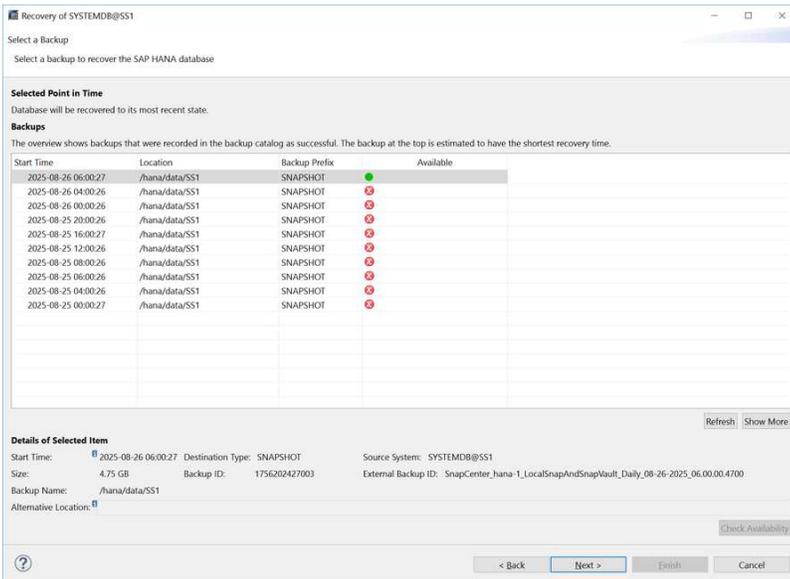


When the Snapshot got restored by SnapCenter a snapshot_databackup_0_1 file is available in the system and tenant database subdirectory of the HANA data volume. This file got created by the HANA database during the HANA database Snapshot creation. HANA deletes the file when the backup operation is finished, 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 #
```

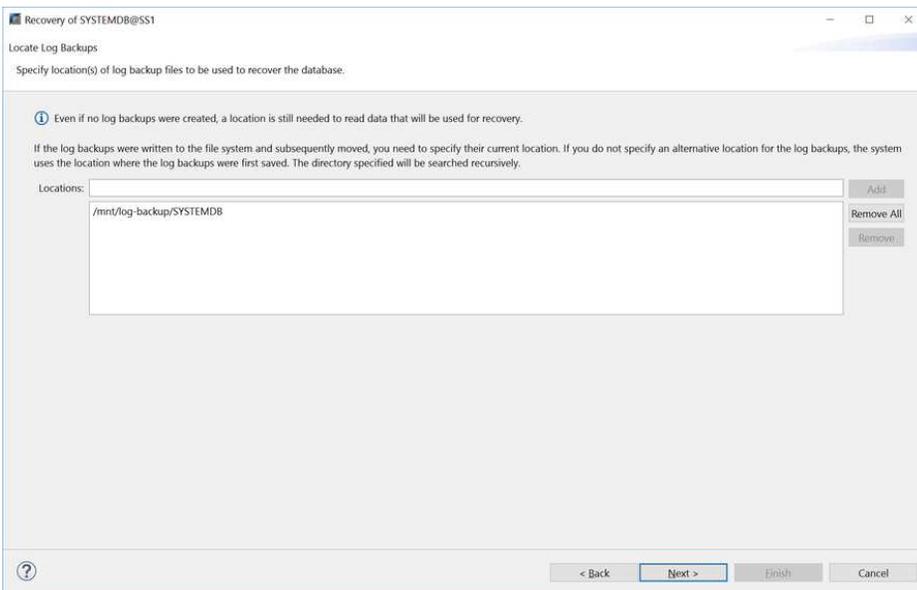
Go to SAP HANA Studio and click Refresh to update the list of available backups. The backup that was restored with SnapCenter is now shown with a green icon in the list of backups. Select the backup and click Next.



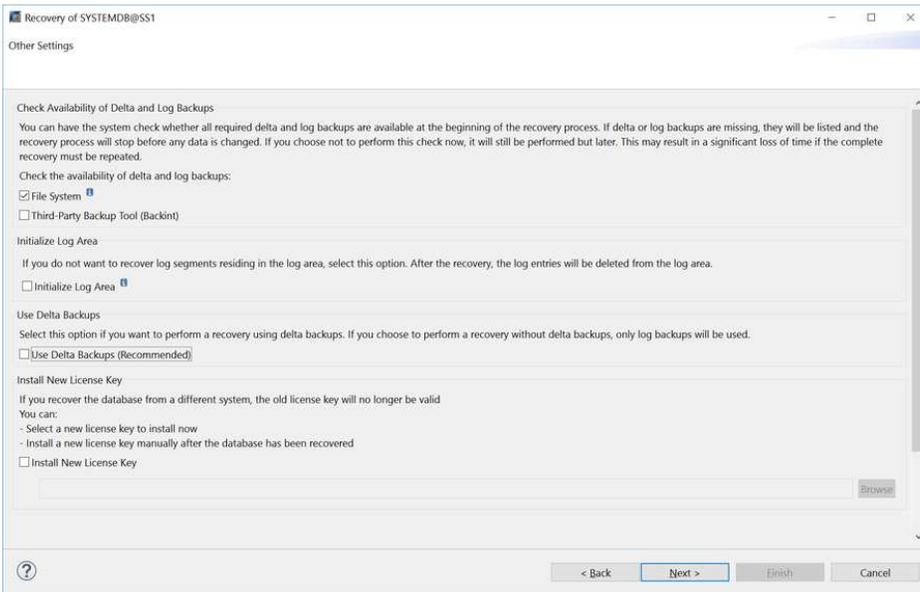
Provide the location of the log backups. Click Next.



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.

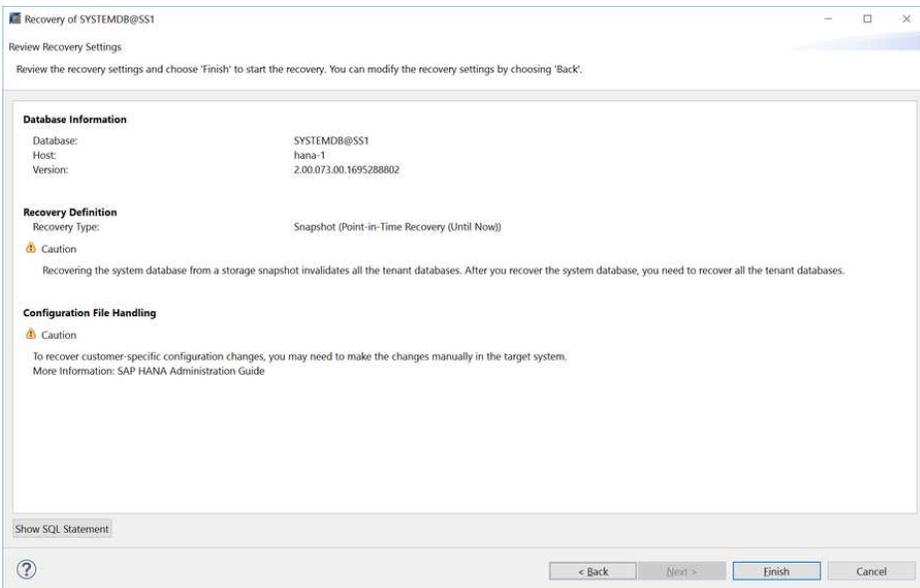


Select other settings as required. Make sure Use Delta Backups is not selected. Click Next.

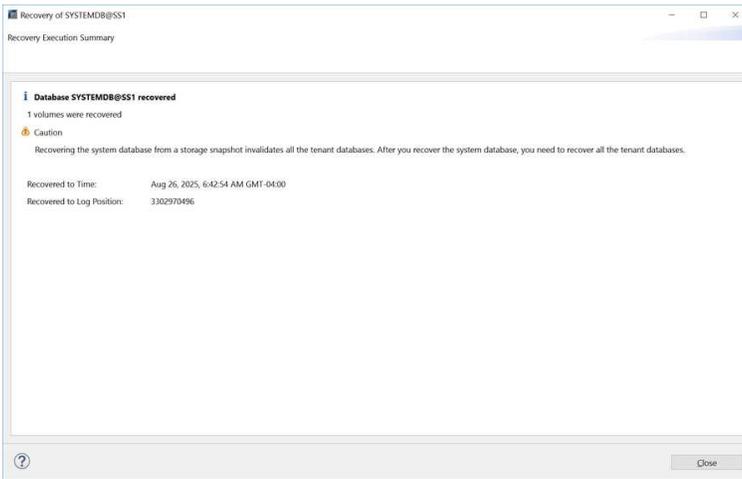


Review the recovery settings and click Finish.

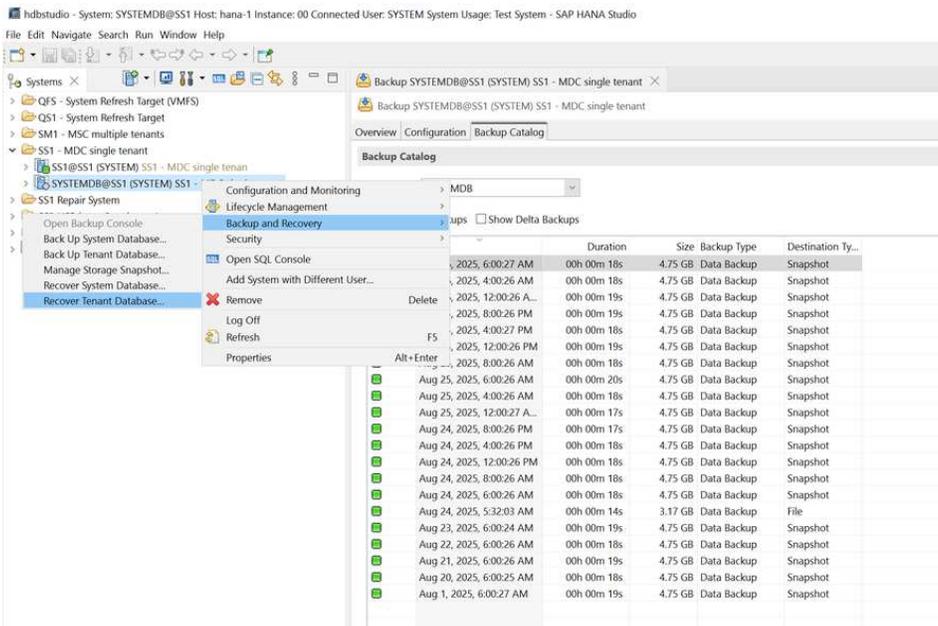
By clicking on show SQL statement, HANA Studio shows the SQL command which is executed for the recovery operation.



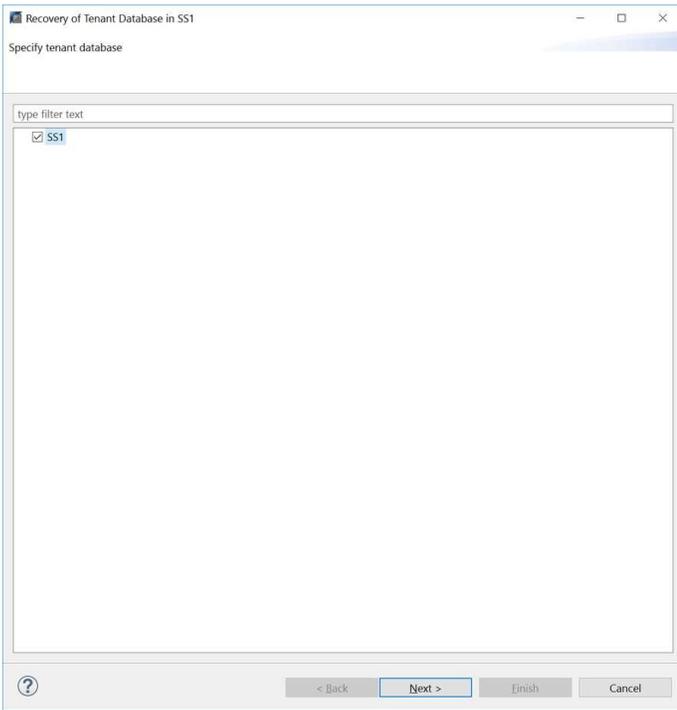
The recovery process starts. Wait until the recovery of the system database is completed.



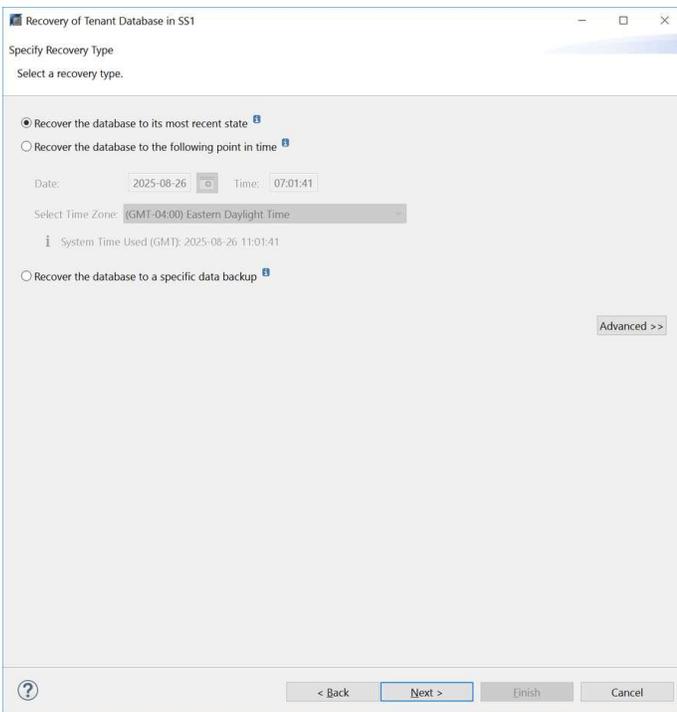
In SAP HANA Studio, select the entry for the system database and start Backup Recovery - Recover Tenant Database.



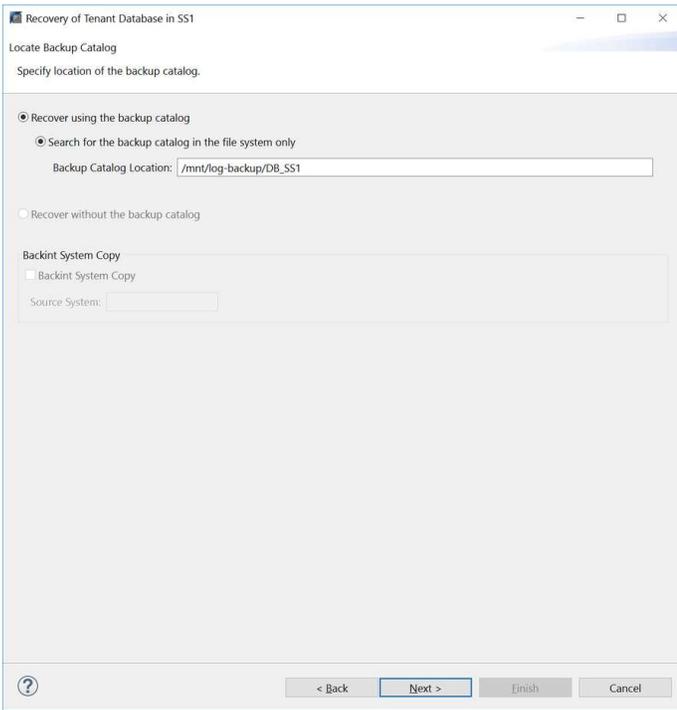
Select the tenant to recover and click Next.



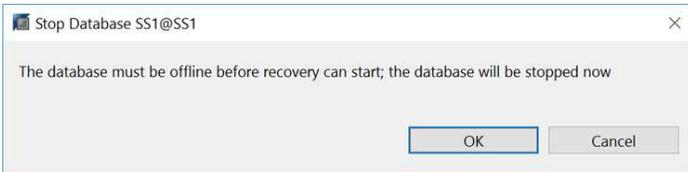
Specify the recovery type and click Next.



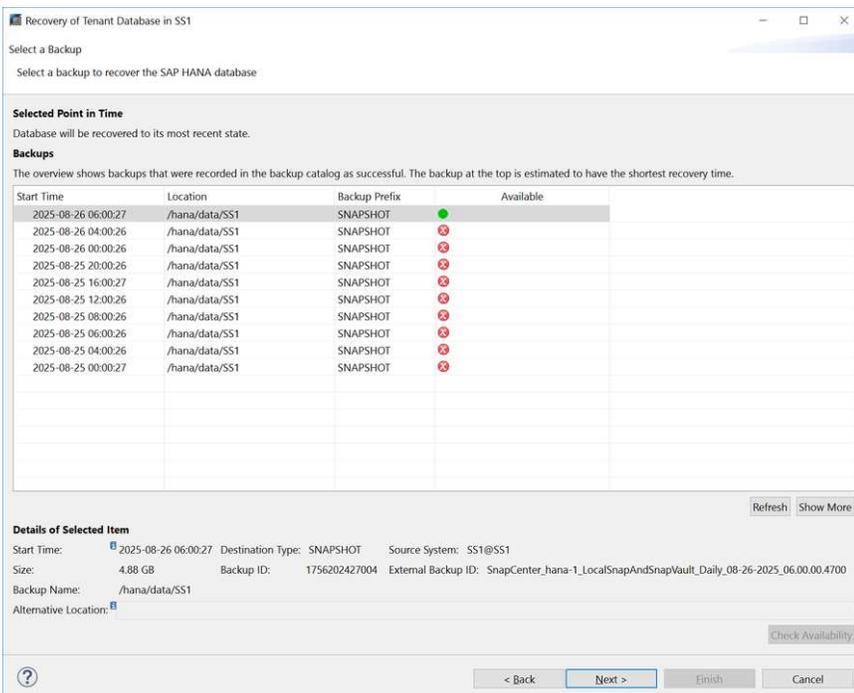
Confirm the backup catalog location and click Next.



Confirm that the shutdown of the tenant database.



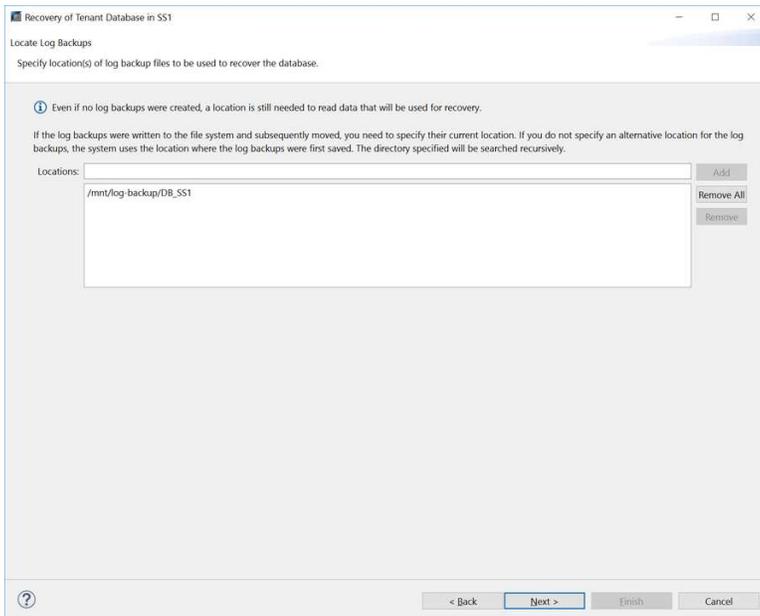
Because the restore of the data volume has been done before the recovery of the system database, the tenant backup is immediately available. Select the backup highlighted in green and click Next.



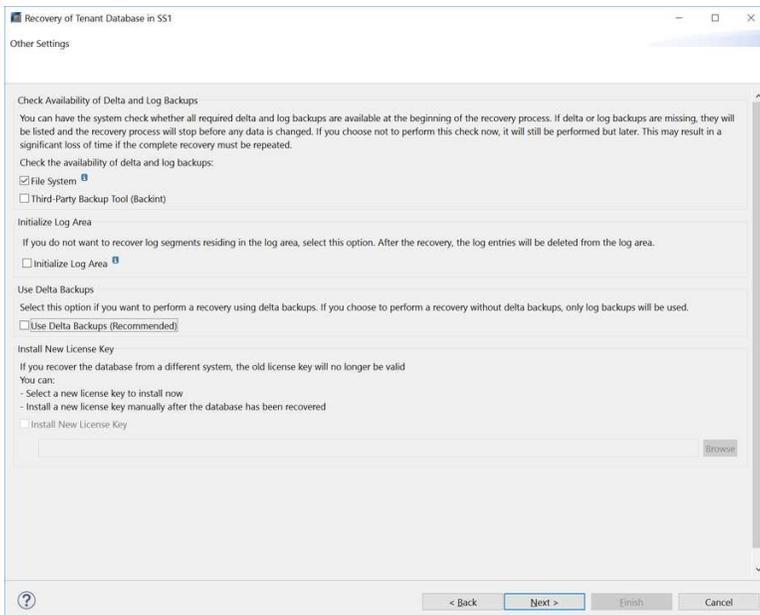
Provide the location of the log backups. Click Next.



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.

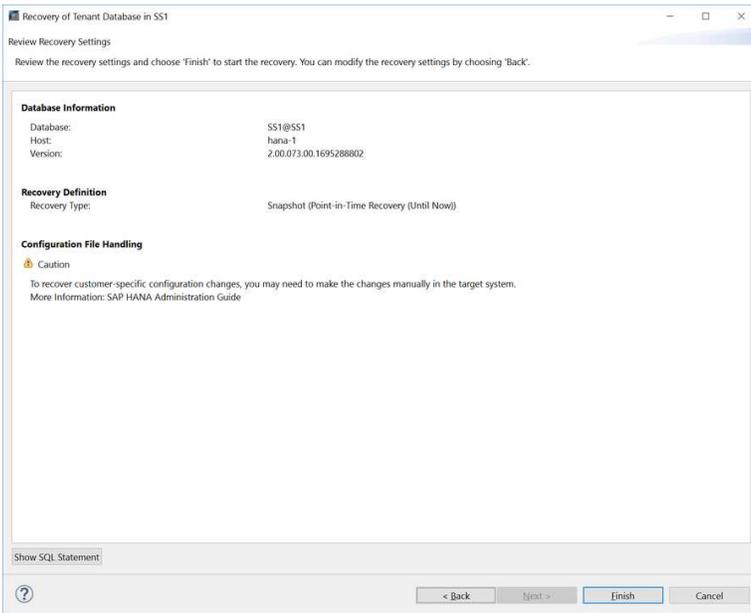


Select other settings as required. Make sure Use Delta Backups is not selected. Click Next.

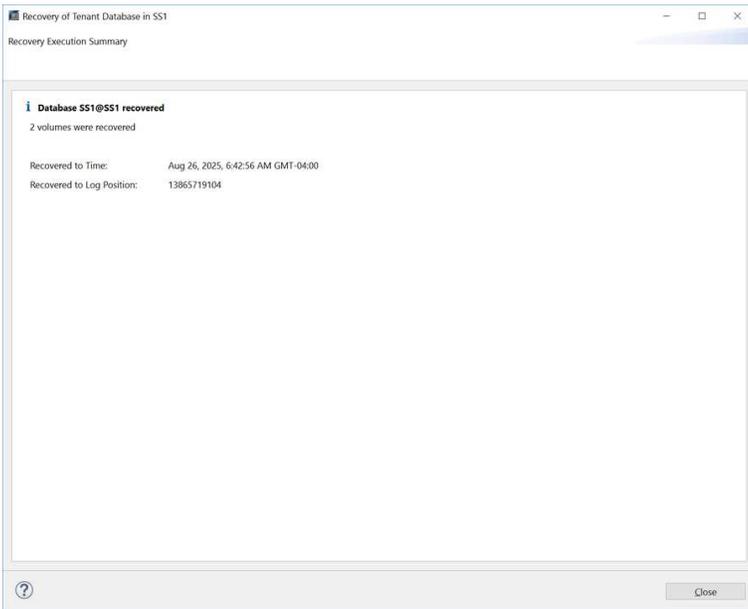


Review the recovery settings and click Finish.

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.

```
hdbsql 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

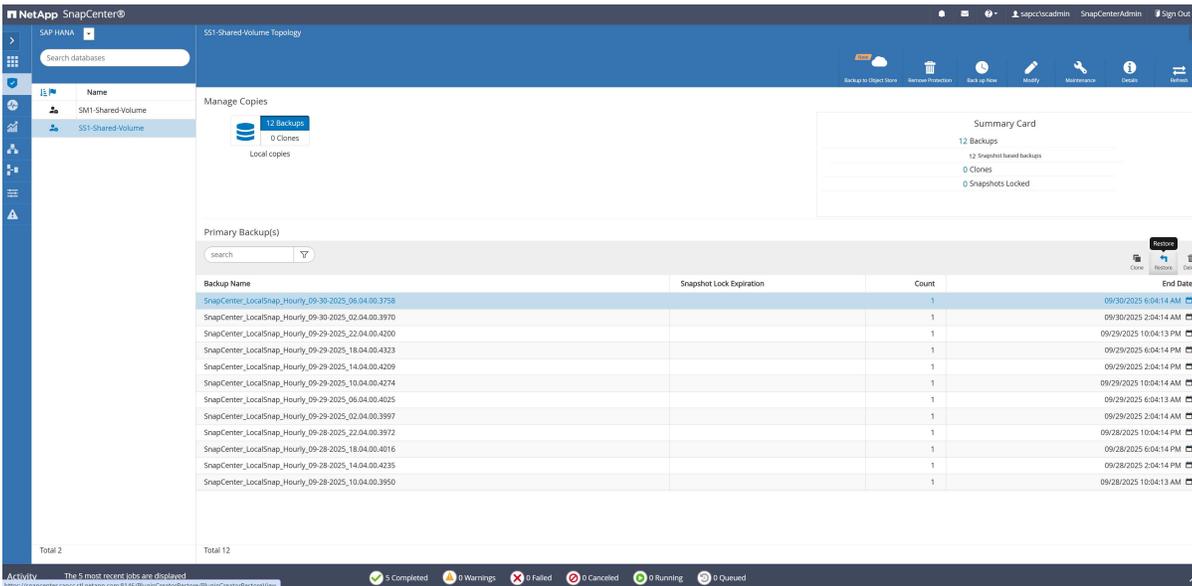
A single tenant restore and recovery operation with SnapCenter is very similar to the workflow described in the previous topic "[Manual recovery with HANA Studio](#)".

To restore and recover an SAP HANA MDC single-tenant system using SAP HANA Studio and SnapCenter, complete the following steps:

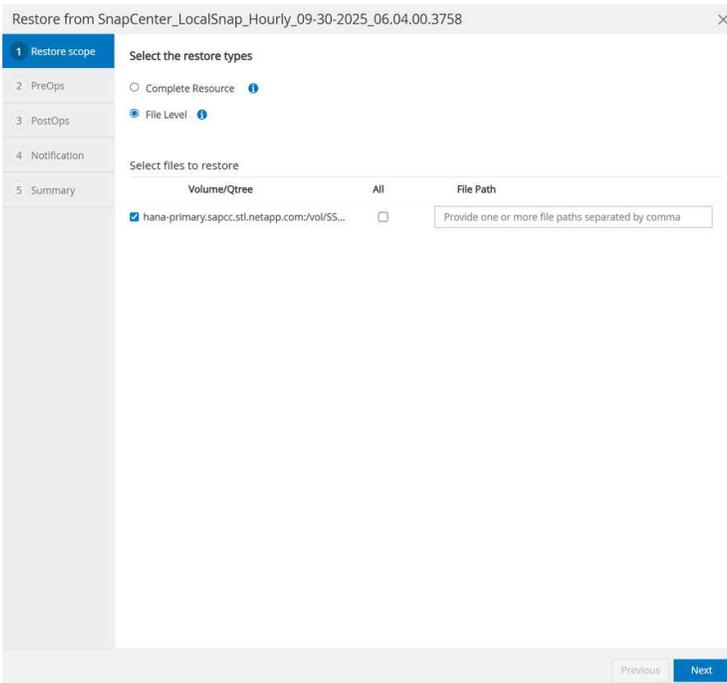
1. Prepare the restore and recovery process with SAP HANA Studio:
 - a. Select Recover Tenant Database and confirm shutdown of the tenant database.
 - b. Select the recovery type and provide the backup catalog location.
 - c. The list of data backups is shown. Select Backup to see the external backup 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-SmConfigSettings -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 #
```

Copyright information

Copyright © 2026 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.