



SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter

NetApp solutions for SAP

NetApp
February 25, 2026

Table of Contents

- SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter 1
 - SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter 1
 - Backup and recovery using Amazon FSx for ONTAP 1
 - Runtime of Snapshot backup and restore operations 2
 - Recovery time objective comparison 2
 - Use cases and values of accelerated backup and cloning operations 4
- SnapCenter architecture 5
 - SnapCenter components 5
 - SnapCenter SAP HANA backup solution 6
 - Scope of this document 7
 - Data protection strategy 8
 - Example lab setup 10
- SnapCenter configuration 10
 - Overview configuration steps 11
 - SAP HANA backup user and hdbuserstore configuration 11
 - Configure storage 13
 - Add a SAP HANA host 16
 - Configure policies 18
 - Configure and protect a HANA resource 22
- SnapCenter backup operations 27
 - Create an on-demand Snapshot backup 27
 - Create an on-demand block integrity check operation 32
- Backup of non-data volumes 37
- Restore and recover 44
- Backup replication with SnapVault 52
 - Overview - Backup replication with SnapVault 52
 - Configure replication relationships on FSx for ONTAP file systems 53
 - Add a backup SVM to SnapCenter 58
 - Create a new SnapCenter policy for backup replication 59
 - Add a policy to resource protection 61
 - Create a backup with replication 62
 - Restore and recover from secondary storage 65
- Where to find additional information 66
 - Version history 67

SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter

SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter

This technical report provides best practices for SAP HANA data protection on Amazon FSx for NetApp ONTAP and NetApp SnapCenter. This document covers SnapCenter concepts, configuration recommendations, and operation workflows, including configuration, backup operations, and restore and recovery operations.

Author: Nils Bauer, NetApp

Companies today require continuous, uninterrupted availability for their SAP applications. They expect consistent performance levels in the face of ever-increasing volumes of data and the need for routine maintenance tasks, such as system backups. Performing backups of SAP databases is a critical task and can have a significant performance impact on the production SAP system.

Backup windows are shrinking while the amount of data to be backed up is increasing. Therefore, it is difficult to find a time when you can perform backups with minimal effect on business processes. The time needed to restore and recover SAP systems is a concern because downtime for SAP production and nonproduction systems must be minimized to reduce cost to the business.

Backup and recovery using Amazon FSx for ONTAP

You can use NetApp Snapshot technology to create database backups in minutes.

The time needed to create a Snapshot copy is independent of the size of the database because a Snapshot copy does not move any physical data blocks on the storage platform. In addition, the use of Snapshot technology has no performance effect on the live SAP system. Therefore, you can schedule the creation of Snapshot copies without considering peak dialog or batch activity periods. SAP and NetApp customers typically schedule multiple online Snapshot backups during the day; for example, every six hours is common. These Snapshot backups are typically kept for three to five days on the primary storage system before being removed or tiered to cheaper storage for long term retention.

Snapshot copies also provide key advantages for restore and recovery operations. NetApp SnapRestore technology enables the restoration of an entire database or, alternatively, just a portion of a database to any point in time, based on the currently available Snapshot copies. Such restore processes are finished in a few seconds, independent of the size of the database. Because several online Snapshot backups can be created during the day, the time needed for the recovery process is significantly reduced relative to a traditional once per day backup approach. Because you can perform a restore with a Snapshot copy that is at most only a few hours old (rather than up to 24 hours), fewer transaction logs must be applied during forward recovery. Therefore, the RTO is reduced to several minutes rather than the several hours required for conventional streaming backups.

Snapshot copy backups are stored on the same disk system as the active online data. Therefore, NetApp recommends using Snapshot copy backups as a supplement rather than a replacement for backups to a secondary location. Most restore and recovery actions are managed by using SnapRestore on the primary storage system. Restores from a secondary location are only necessary if the primary storage system containing the Snapshot copies is damaged. You can also use the secondary location if it is necessary to

restore a backup that is no longer available on the primary location.

A backup to a secondary location is based on Snapshot copies created on the primary storage. Therefore, the data is read directly from the primary storage system without generating load on the SAP database server. The primary storage communicates directly with the secondary storage and replicates the backup data to the destination by using the NetApp SnapVault feature.

SnapVault offers significant advantages when compared to traditional backups. After an initial data transfer, in which all data has been transferred from the source to the destination, all subsequent backups copy only move the changed blocks to the secondary storage. Therefore, the load on the primary storage system and the time needed for a full backup are significantly reduced. Because SnapVault stores only the changed blocks at the destination, any additional full database backups consume significantly less disk space.

Runtime of Snapshot backup and restore operations

The following figure shows a customer's HANA Studio using Snapshot backup operations. The image shows that the HANA database (approximately 4TB in size) is backed up in 1 minute and 20 seconds by using Snapshot backup technology and more than 4 hours with a file-based backup operation.

The largest part of the overall backup workflow runtime is the time needed to execute the HANA backup save point operation, and this step is dependent on the load on the HANA database. The storage Snapshot backup itself always finishes in a couple of seconds.

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

File-based backup: **4 hours 05 min**

(~270 MB/s throughput)

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

Snapshot backup: **1 min 20 sec**

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

Backup runtime reduced by 99%

Recovery time objective comparison

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

Time needed to restore database

With a file-based backup, the restore time depends on the size of the database and backup infrastructure, which defines the restore speed in megabytes per second. For example, if the infrastructure supports a restore operation at a speed of 250MBps, it takes approximately 4.5 hours to restore a database 4TB in size on the persistence.

With storage Snapshot copy backups, the restore time is independent of the size of the database and is always

in the range of a couple of seconds.

Time needed to start database

The database start time depends on the size of the database and the time needed to load the data into memory. In the following examples, it is assumed that the data can be loaded with 1000MBps. Loading 4TB into memory takes around 1hour and 10 minutes. The start time is the same for a file-based and Snapshot based restore and recovery operations.

Time needed to recover database

The recovery time depends on the number of logs that must be applied after the restore. This number is determined by the frequency at which data backups are taken.

With file-based data backups, the backup schedule is typically once per day. A higher backup frequency is normally not possible, because the backup degrades production performance. Therefore, in the worst case, all the logs that were written during the day must be applied during forward recovery.

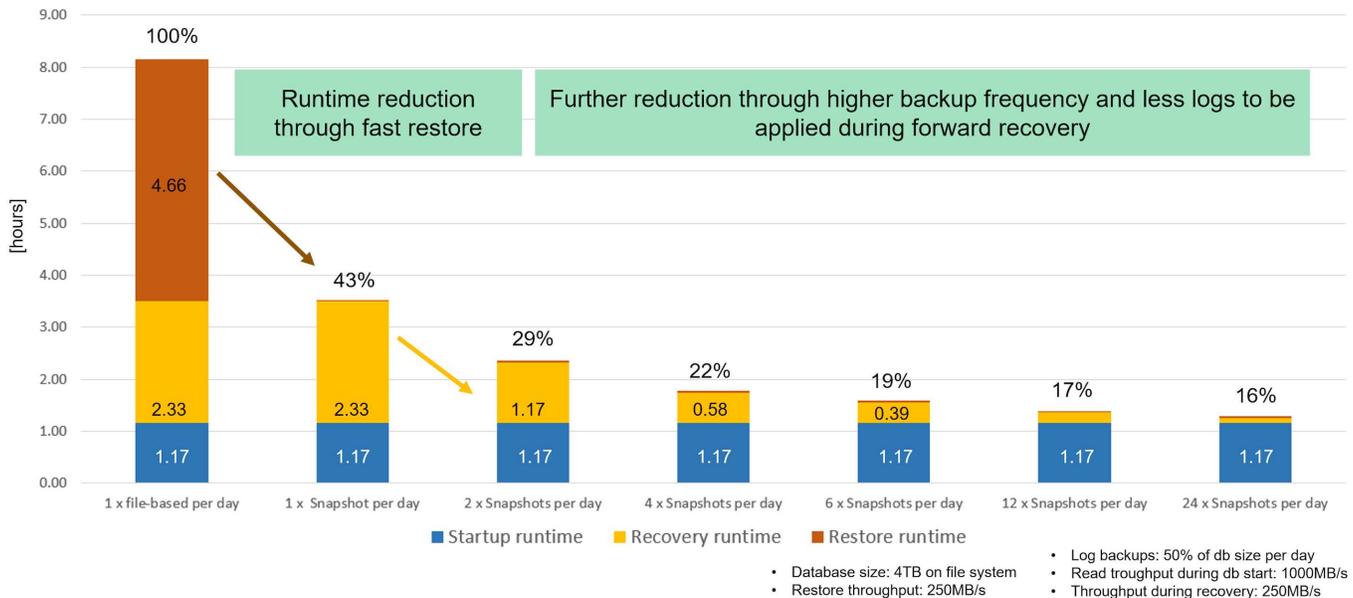
Snapshot backups are typically scheduled with a higher frequency because they do not influence the performance of the SAP HANA database. For example, if Snapshot backups are scheduled every six hours, the recovery time would be, in the worst case, one-fourth of the recovery time for a file-based backup (6 hours / 24 hours = .25).

The following figure shows a comparison of restore and recovery operations with a daily file-based backup and Snapshot backups with different schedules.

The first two bars show that even with a single Snapshot backup per day, the restore and recovery is reduced to 43% due to the speed of the restore operation from a Snapshot backup. If multiple Snapshot backups per day are created, the runtime can be reduced further because less logs need to be applied during forward recovery.

The following figure also shows that four to six Snapshot backups per day makes the most sense, because a higher frequency does not have a big influence on the overall runtime anymore.

Restore and Recovery of a 4TB HANA Database (8TB RAM)



Use cases and values of accelerated backup and cloning operations

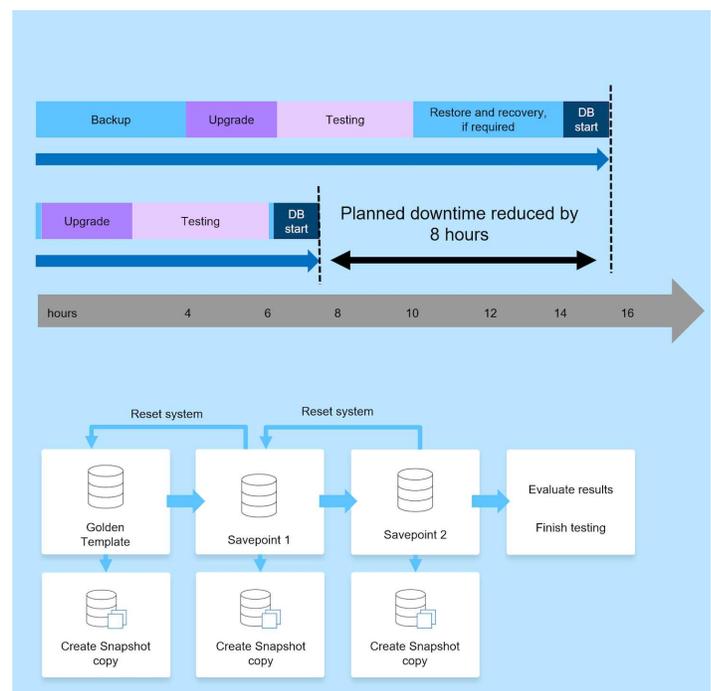
Executing backups is a critical part of any data protection strategy. Backups are scheduled on a regular basis to ensure that you can recover from system failures. This is the most obvious use case, but there are also other SAP lifecycle management tasks, where accelerating backup and recovery operations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

SAP HANA system upgrade is an example of where an on-demand backup before the upgrade and a possible restore operation if the upgrade fails has a significant impact on the overall planned downtime. With the example of a 4TB database, you can reduce the planned downtime by 8 hours by using the Snapshot-based backup and restore operations.

Another use case example would be a typical test cycle, where testing must be done over multiple iterations with different data sets or parameters. When leveraging the fast backup and restore operations, you can easily create save points within your test cycle and reset the system to any of these previous save points if a test fails or needs to be repeated. This enables testing to finish earlier or enables more testing at the same time and improves test results.

Use Cases for Backup and Recovery Operations

- Accelerate HANA system upgrade operations
 - Fast on-demand backup before HANA system upgrade
 - Fast restore operation in case of an upgrade failure
 - Reduction of planned downtime
- Accelerate test cycles
 - Fast creation of savepoints after a successful step
 - Fast reset of system to any savepoint
 - Repeat step until successful



When Snapshot backups have been implemented, they can be used to address multiple other use cases, which require copies of a HANA database. With FSx for ONTAP, you can create a new volume based on the content of any available Snapshot backup. The runtime of this operation is a few seconds, independent of the size of the volume.

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

The second use case is the creation of a repair system, which is used to address a logical corruption in the production system. In this case, an older Snapshot backup of the production system is used to start a repair system, which is an identical clone of the production system with the data before the corruption occurred. The repair system is then used to analyze the problem and export the required data before it was corrupted.

The last use case is the ability to run a disaster recover failover test without stopping the replication and therefore without influencing RTO and recovery point objective (RPO) of the disaster recovery setup. When

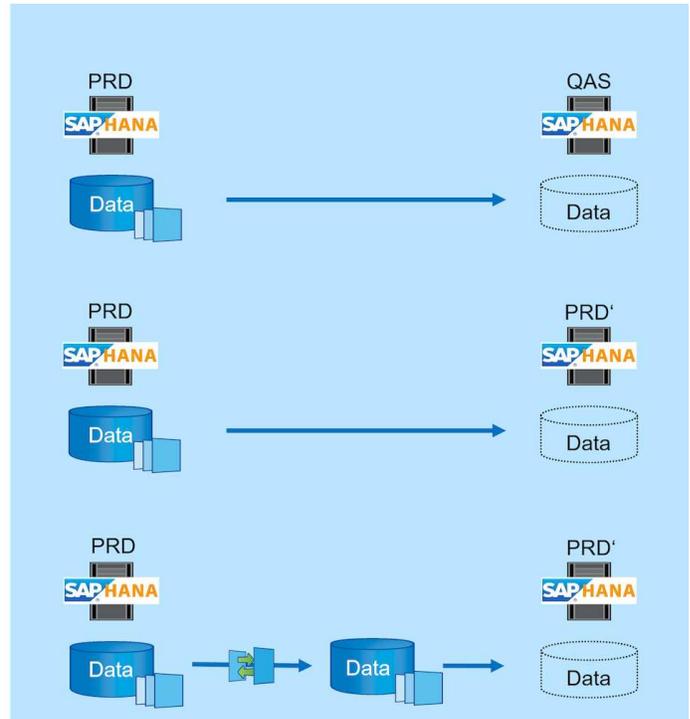
FSx for ONTAP NetApp SnapMirror replication is used to replicate the data to the disaster recovery site, the production Snapshot backups are available at the disaster recovery site as well and can then be used to create a new volume for disaster recover testing.

Use Cases for Cloning Operations

- SAP System Refresh
 - Fast creation of a new volume based on a production Snapshot backup
 - Attach volume to the test system and recover HANA database with SID change

- Repair System creation to address logical corruption
 - Fast creation of a new volume based on a production Snapshot backup
 - Attach volume to the repair system and recover HANA database w/o SID change

- Disaster Recovery testing
 - Combined with SnapMirror Replication
 - Attach storage clone from a replicated production Snapshot backup to a DR test system



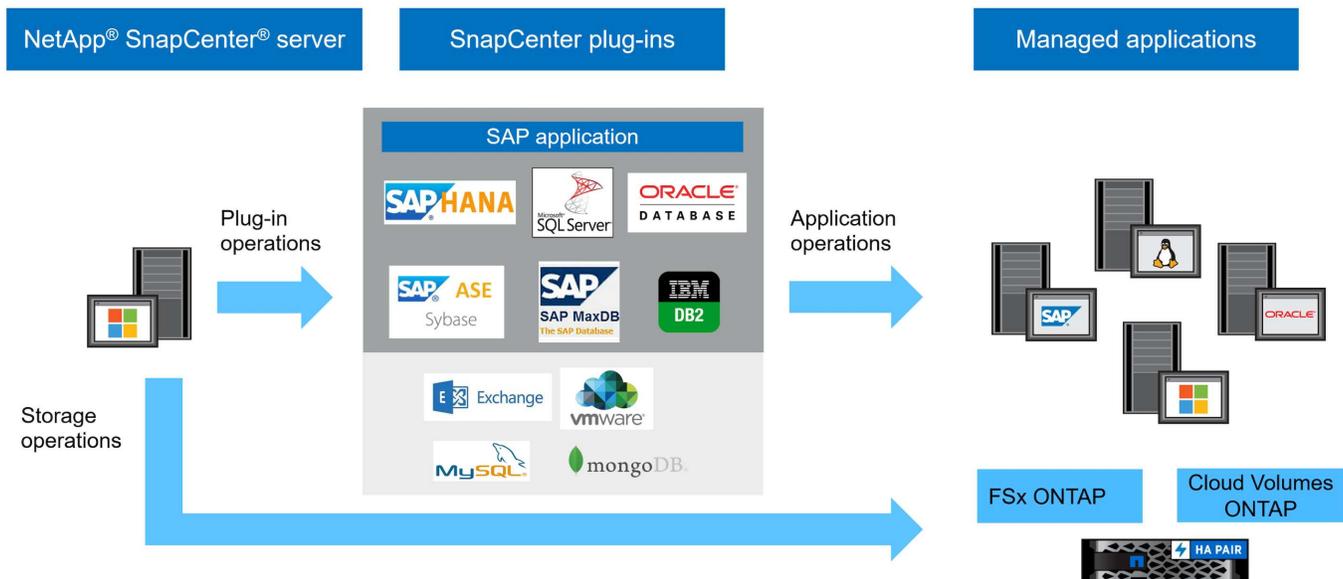
SnapCenter architecture

SnapCenter is a unified, scalable platform for application-consistent data protection. SnapCenter provides centralized control and oversight, while delegating the ability for users to manage application-specific backup, restore, and clone jobs. With SnapCenter, database and storage administrators learn a single tool to manage backup, restore, and cloning operations for a variety of applications and databases.

SnapCenter manages data across endpoints in the data fabric powered by NetApp. You can use SnapCenter to replicate data between on-premises environments; between on-premises environments and the cloud; and between private, hybrid, or public clouds.

SnapCenter components

SnapCenter includes the SnapCenter Server, the SnapCenter Plug-In Package for Windows, and the SnapCenter Plug-In Package for Linux. Each package contains plug-ins to SnapCenter for various applications and infrastructure components.



SnapCenter SAP HANA backup solution

The SnapCenter backup solution for SAP HANA covers the following areas:

- Backup operations, scheduling, and retention management
 - SAP HANA data backup with storage-based Snapshot copies
 - Non-data volume backup with storage-based Snapshot copies (for example, /hana/shared)
 - Database block integrity checks using a file-based backup
 - Replication to an off-site backup or disaster recovery location
- Housekeeping of the SAP HANA backup catalog
 - For HANA data backups (Snapshot and file-based)
 - For HANA log backups
- Restore and recovery operations
 - Automated restore and recovery
 - Single tenant restore operations for SAP HANA (MDC) systems

Database data file backups are executed by SnapCenter in combination with the plug-in for SAP HANA. The plug-in triggers the SAP HANA database backup save point so that the Snapshot copies, which are created on the primary storage system, are based on a consistent image of the SAP HANA database.

SnapCenter enables the replication of consistent database images to an off-site backup or disaster recovery location by using SnapVault or the SnapMirror feature. Typically, different retention policies are defined for backups at primary and at the off-site backup storage. SnapCenter handles the retention at primary storage, and ONTAP handles the retention at the off-site backup storage.

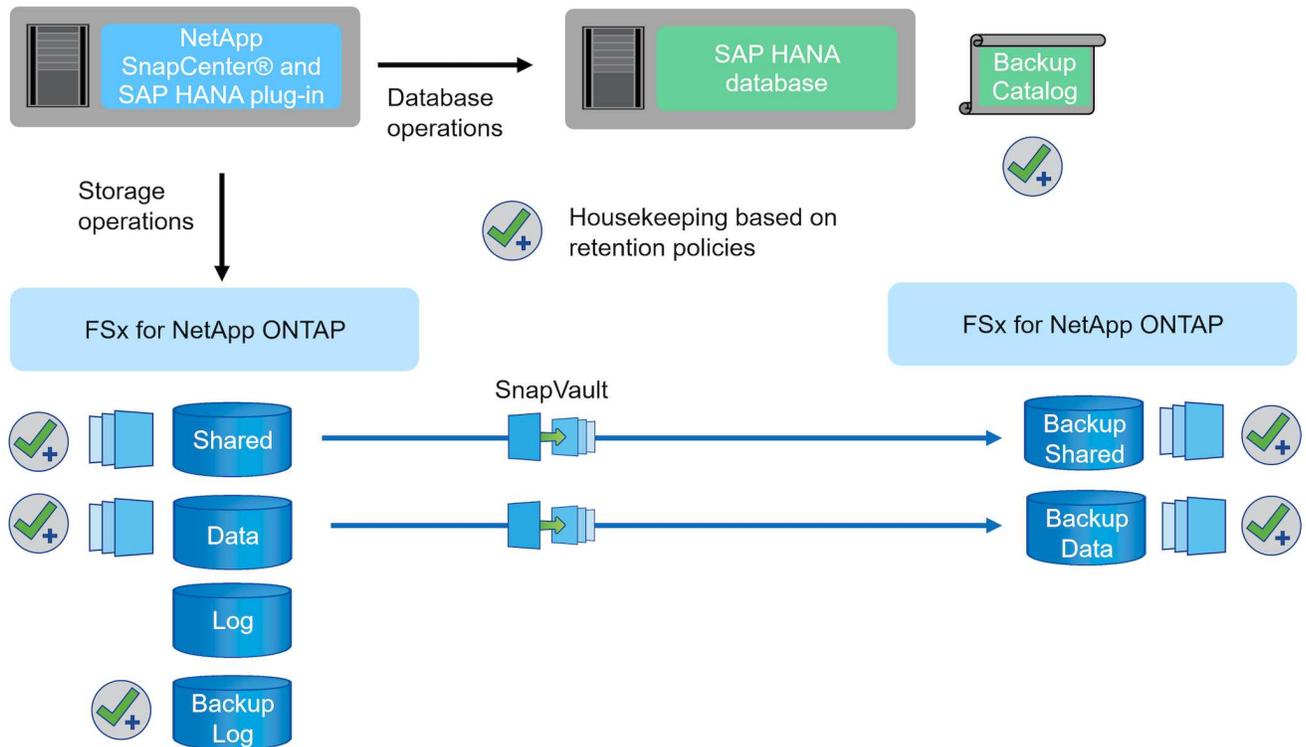
To allow a complete backup of all SAP HANA-related resources, SnapCenter also enables you to back up all non-data volumes by using the SAP HANA plug-in with storage-based Snapshot copies. You can schedule non-data volumes independently from the database data backup to enable individual retention and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

SAP recommends combining storage-based Snapshot backups with a weekly file-based backup to execute a

block integrity check. You can execute the block integrity check from within SnapCenter. Based on your configured retention policies, SnapCenter manages the housekeeping of data file backups at the primary storage, log file backups, and the SAP HANA backup catalog.

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

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



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

1. Creates an SAP HANA backup save point to create a consistent image on the persistence layer.
2. Creates a storage-based Snapshot copy of the data volume.
3. Registers the storage-based Snapshot back up in the SAP HANA backup catalog.
4. Releases the SAP HANA backup save point.
5. Executes a SnapVault or SnapMirror update for the data volume, if configured.
6. Deletes storage Snapshot copies at the primary storage based on the defined retention policies.
7. Deletes SAP HANA backup catalog entries if the backups do not exist anymore at the primary or off-site backup storage.
8. Whenever a backup has been deleted based on the retention policy or manually, SnapCenter also deletes all log backups that are older than the oldest data backup. Log backups are deleted on the file system and in the SAP HANA backup catalog.

Scope of this document

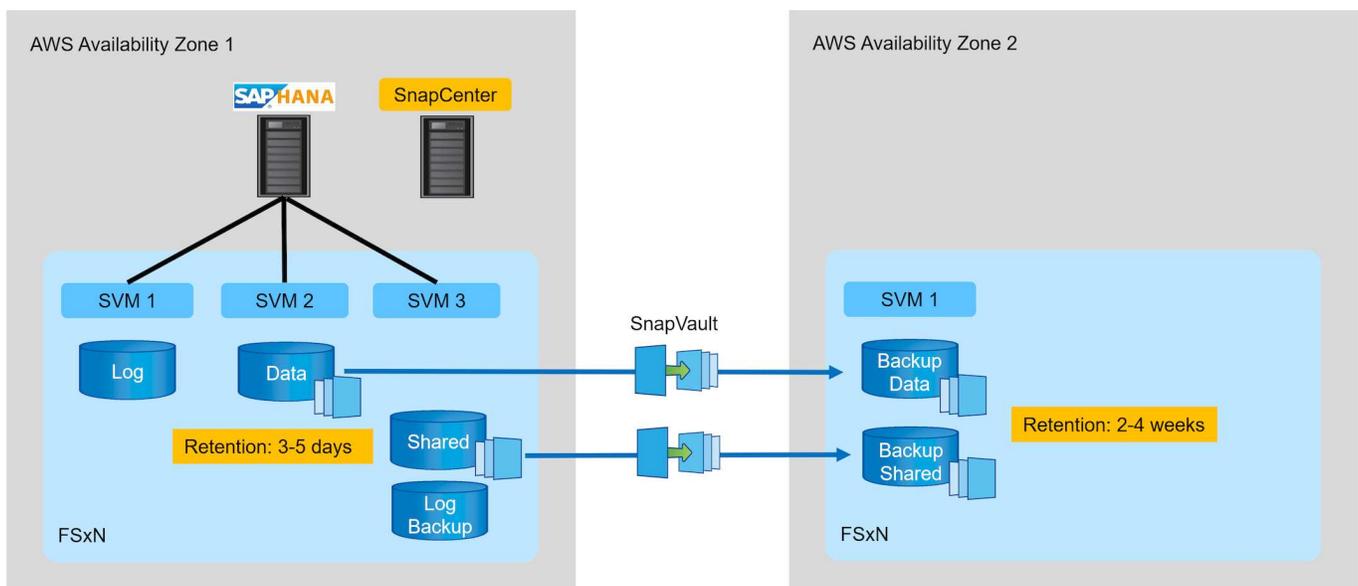
This document describes the most common SnapCenter configuration option for an SAP HANA MDC single host system with a single tenant on FSx for ONTAP. Other configuration options are possible and, in some

cases, required for specific SAP HANA systems, for example, for a multiple host system. For a detailed description about other configuration options, see [SnapCenter concepts and best practices \(netapp.com\)](https://netapp.com).

In this document, we use the Amazon Web Services (AWS) console and the FSx for ONTAP CLI to execute the required configuration steps on the storage layer. You can also use NetApp Cloud Manager to manage FSx for ONTAP, but this is out of scope for this document. For information about using NetApp Cloud Manager for FSx for ONTAP, see [Learn about Amazon FSx for ONTAP \(netapp.com\)](https://netapp.com).

Data protection strategy

The following figure shows a typical backup architecture for SAP HANA on FSx for ONTAP. The HANA system is located in the AWS availability zone 1 and is using an FSx for ONTAP file system within the same availability zone. Snapshot backup operations are executed for the data and the shared volume of the HANA database. In addition to the local Snapshot backups, which are kept for 3-5 days, backups are also replicated to an offsite storage for longer term retention. The offsite backup storage is a second FSx for ONTAP file system located in a different AWS availability zone. Backups of the HANA data and shared volume are replicated with SnapVault to the second FSx for ONTAP file system and are kept for 2-3 weeks.



Before configuring SnapCenter, the data protection strategy must be defined based on the RTO and RPO requirements of the various SAP systems.

A common approach is to define system types such as production, development, test, or sandbox systems. All SAP systems of the same system type typically have the same data protection parameters.

The following parameters must be defined:

- How often should a Snapshot backup be executed?
- How long should Snapshot copy backups be kept on the primary storage system?
- How often should a block integrity check be executed?
- Should the primary backups be replicated to an off-site backup site?
- How long should the backups be kept at the off-site backup storage?

The following table shows an example of data protection parameters for the system types: production, development, and test. For the production system, a high backup frequency has been defined, and the

backups are replicated to an off-site backup site once per day. The test systems have lower requirements and no replication of the backups.

Parameters	Production systems	Development systems	Test systems
Backup frequency	Every 6 hours	Every 6 hours	Every 6 hours
Primary retention	3 days	3 days	3 days
Block integrity check	Once per week	Once per week	No
Replication to off-site backup site	Once per day	Once per day	No
Off-site backup retention	2 weeks	2 weeks	Not applicable

The following table shows the policies that must be configured for the data protection parameters.

Parameters	Policy LocalSnap	Policy LocalSnapAndSnapVault	Policy BlockIntegrityCheck
Backup type	Snapshot based	Snapshot based	File based
Schedule frequency	Hourly	Daily	Weekly
Primary retention	Count = 12	Count = 3	Count = 1
SnapVault replication	No	Yes	Not applicable

The policy `LocalSnapshot` is used for the production, development, and test systems to cover the local Snapshot backups with a retention of two days.

In the resource protection configuration, the schedule is defined differently for the system types:

- Production: Schedule every 4 hours.
- Development: Schedule every 4 hours.
- Test: Schedule every 4 hours.

The policy `LocalSnapAndSnapVault` is used for the production and development systems to cover the daily replication to the off-site backup storage.

In the resource protection configuration, the schedule is defined for production and development:

- Production: Schedule every day.
- Development: Schedule every day. The policy `BlockIntegrityCheck` is used for the production and development systems to cover the weekly block integrity check by using a file-based backup.

In the resource protection configuration, the schedule is defined for production and development:

- Production: Schedule every week.
- Development: Schedule every week.

For each individual SAP HANA database that uses the off-site backup policy, you must configure a protection relationship on the storage layer. The protection relationship defines which volumes are replicated and the retention of backups at the off-site backup storage.

With the following example, for each production and development system, a retention of two weeks is defined at the off-site backup storage.

In this example, protection policies and retention for SAP HANA database resources and non- data volume resources are not different.

Example lab setup

The following lab setup was used as an example configuration for the rest of this document.

HANA system PFX:

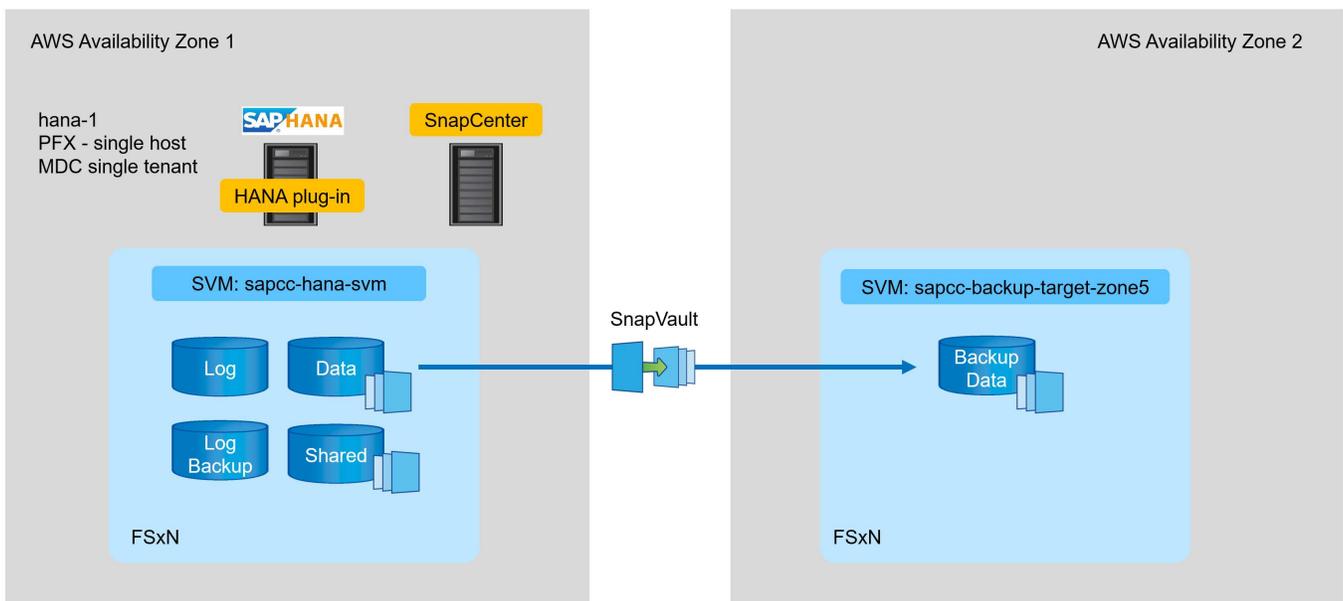
- Single host MDC system with a single tenant
- HANA 2.0 SPS 6 revision 60
- SLES for SAP 15SP3

SnapCenter:

- Version 4.6
- HANA and Linux plug-in deployed on a HANA database host

FSx for ONTAP file systems:

- Two FSx for ONTAP file systems with a single storage virtual machine (SVM)
- Each FSx for ONTAP system in a different AWS availability zone
- HANA data volume replicated to the second FSx for ONTAP file system



SnapCenter configuration

You must perform the steps in this section for base SnapCenter configuration and the protection of the HANA resource.

Overview configuration steps

You must perform the following steps for base SnapCenter configuration and the protection of the HANA resource. Each step is described in detail in the following chapters.

1. Configure SAP HANA backup user and hdbuserstore key. Used to access the HANA database with the hdbsql client.
2. Configure storage in SnapCenter. Credentials to access the FSx for ONTAP SVMs from SnapCenter
3. Configure credentials for plug-in deployment. Used to automatically deploy and install the required SnapCenter plug-ins on the HANA database host.
4. Add HANA host to SnapCenter. Deploys and installs the required SnapCenter plug-ins.
5. Configure policies. Defines the backup operation type (Snapshot, file), retentions, as well as optional Snapshot backup replication.
6. Configure HANA resource protection. Provide hdbuserstore key and attach policies and schedules to the HANA resource.

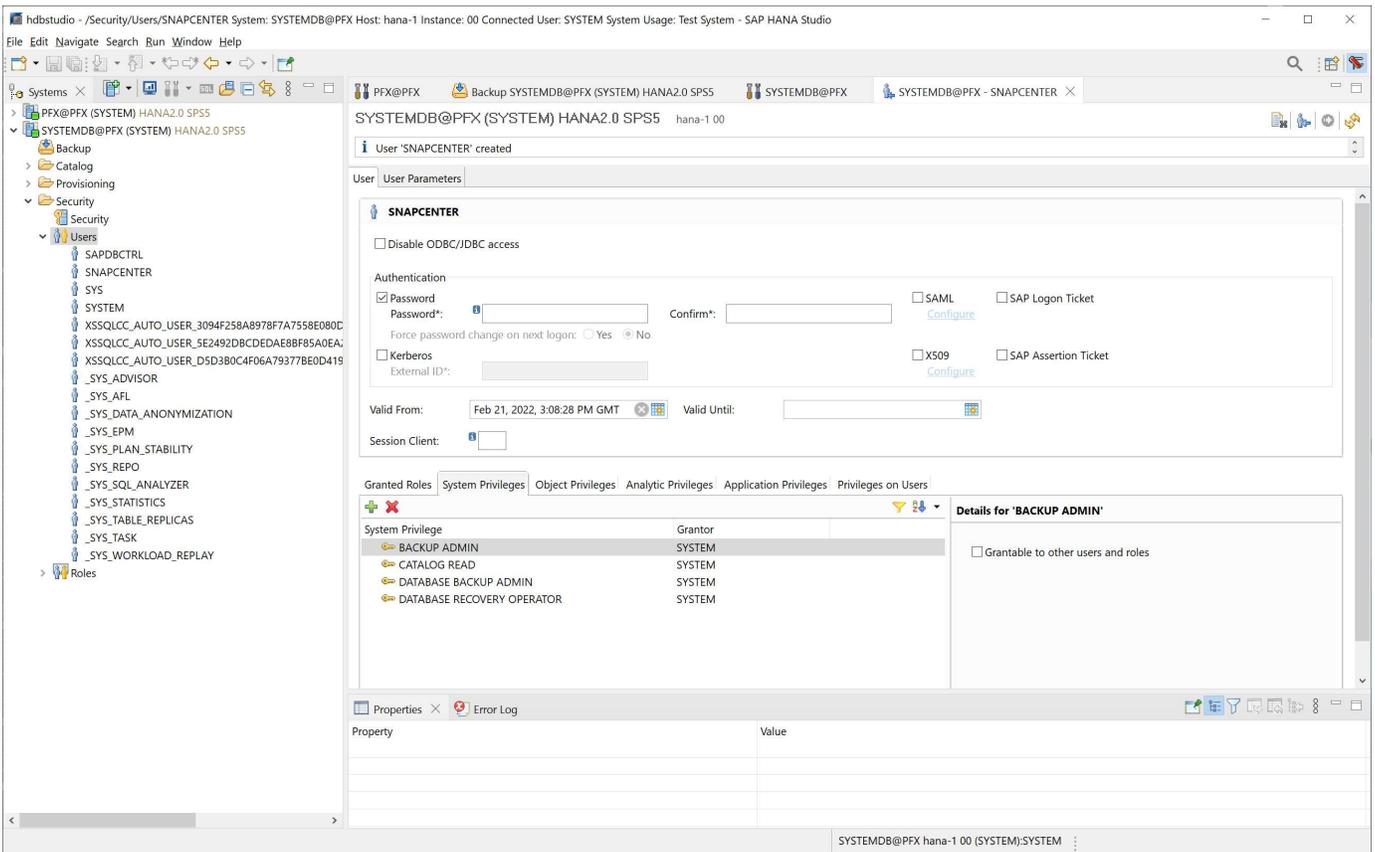
SAP HANA backup user and hdbuserstore configuration

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

The following figure shows the SAP HANA Studio through which you can create the backup user

The required privileges are changed with the HANA 2.0 SPS5 release: backup admin, catalog read, database backup admin, and database recovery operator. For earlier releases, backup admin and catalog read are sufficient.

For an SAP HANA MDC system, you must create the user in the system database because all backup commands for the system and the tenant databases are executed by using the system database.



The following command is used for the user store configuration with the <sid>adm user:

```
hdbuserstore set <key> <host>:<port> <database user> <password>
```

SnapCenter uses the <sid>adm user to communicate with the HANA database. Therefore, you must configure the user store key by using the <'sid>adm` user on the database host. Typically, the SAP HANA hdbsql client software is installed together with the database server installation. If this is not the case, you must install the hdbclient first.

In an SAP HANA MDC setup, port 3<instanceNo>13 is the standard port for SQL access to the system database and must be used in the hdbuserstore configuration.

For an SAP HANA multiple-host setup, you must configure user store keys for all hosts. SnapCenter tries to connect to the database by using each of the provided keys and can therefore operate independently of a failover of an SAP HANA service to a different host. In our lab setup, we configured a user store key for the user pfxadm for our system PFX, which is a single host HANA MDC system with a single tenant.

```
pfxadm@hana-1:/usr/sap/PFX/home> hdbuserstore set PFXKEY hana-1:30013
SNAPCENTER <password>
Operation succeed.
```

```

pfxadm@hana-1:/usr/sap/PFX/home> hdbuserstore list
DATA FILE      : /usr/sap/PFX/home/.hdb/hana-1/SSFS_HDB.DAT
KEY FILE       : /usr/sap/PFX/home/.hdb/hana-1/SSFS_HDB.KEY
ACTIVE RECORDS : 7
DELETED RECORDS : 0
KEY PFXKEY
  ENV : hana-1:30013
  USER: SNAPCENTER
KEY PFXSAPDBCTRL
  ENV : hana-1:30013
  USER: SAPDBCTRL
Operation succeed.

```

You can check the access to the HANA system database that uses the key with the `hdbsql` command.

```

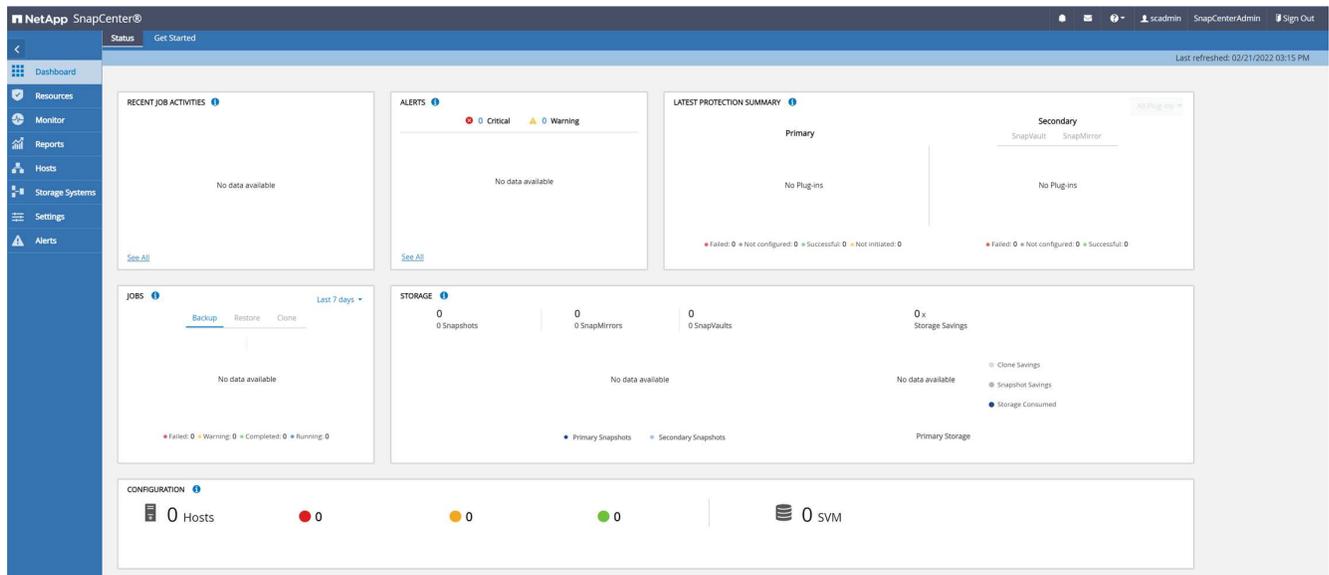
pfxadm@hana-1:/usr/sap/PFX/home> hdbsql -U PFXKEY
Welcome to the SAP HANA Database interactive terminal.
Type: \h for help with commands
      \q to quit
hdbsql SYSTEMDB=>

```

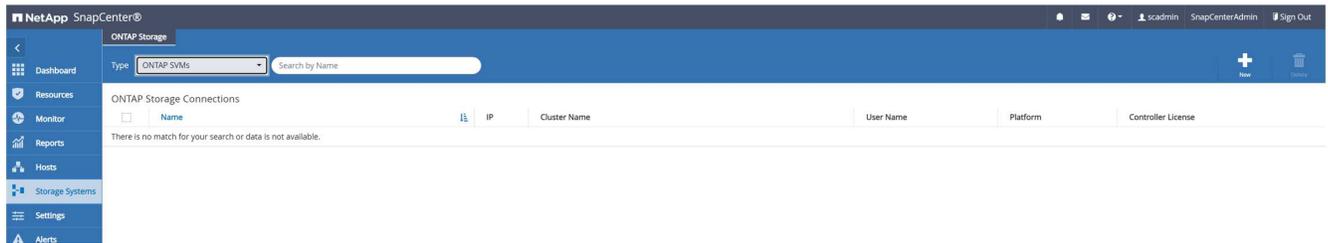
Configure storage

Follow these steps to configure storage in SnapCenter.

1. In the SnapCenter UI, select Storage Systems.

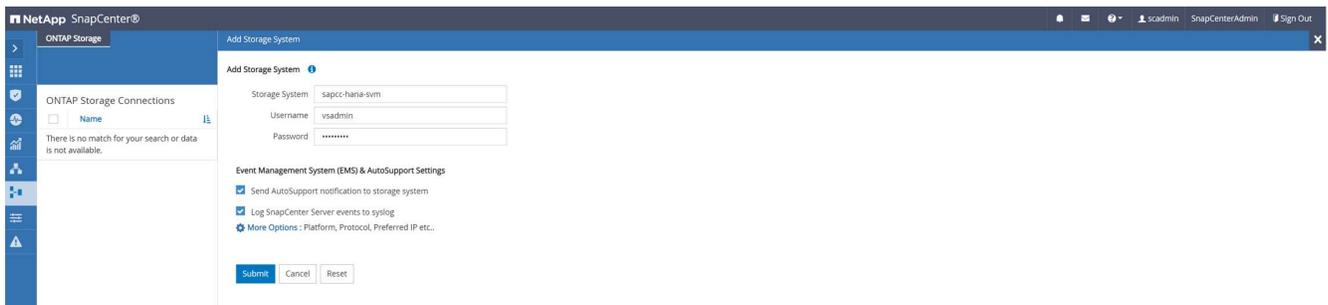


You can select the storage system type, which can be ONTAP SVMs or ONTAP Clusters. In the following example, SVM management is selected.

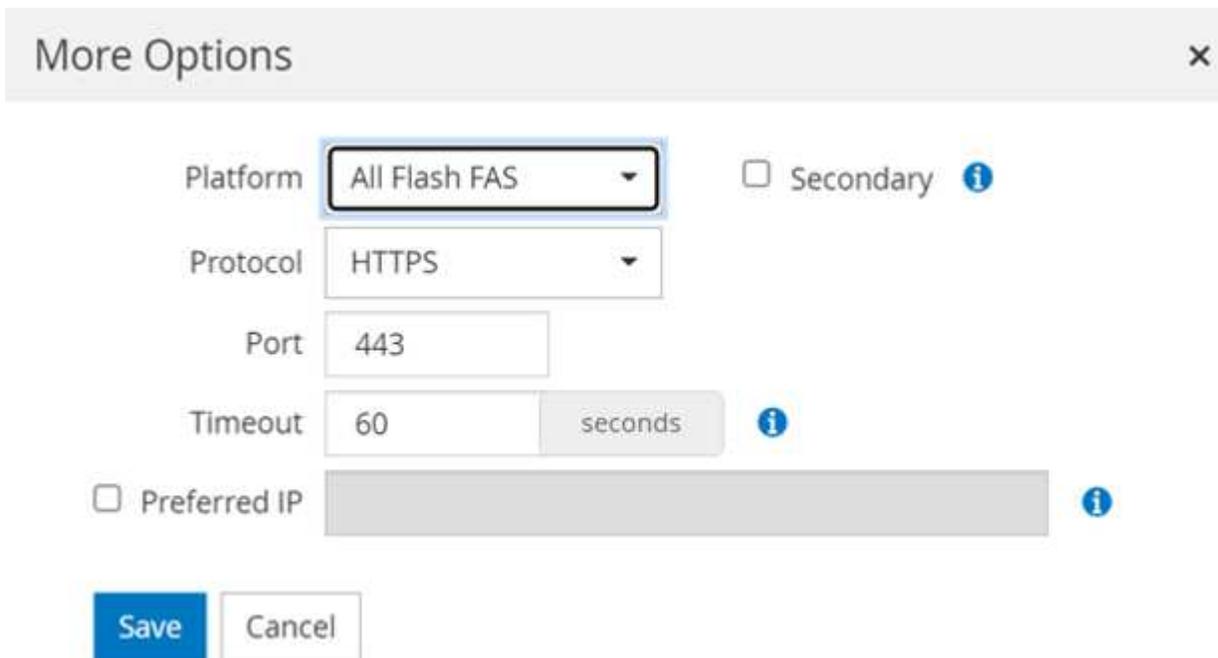


2. To add a storage system and provide the required host name and credentials, click New.

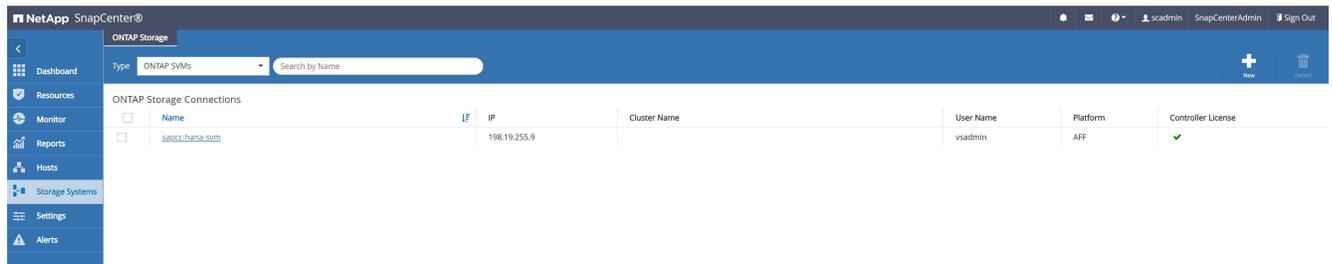
The SVM user is not required to be the vsadmin user, as shown in the following figure. Typically, a user is configured on the SVM and assigned the required permissions to execute backup and restore operations. For information about required privileges, see [SnapCenter Installation Guide](#) in the section titled “Minimum ONTAP privileges required”.



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



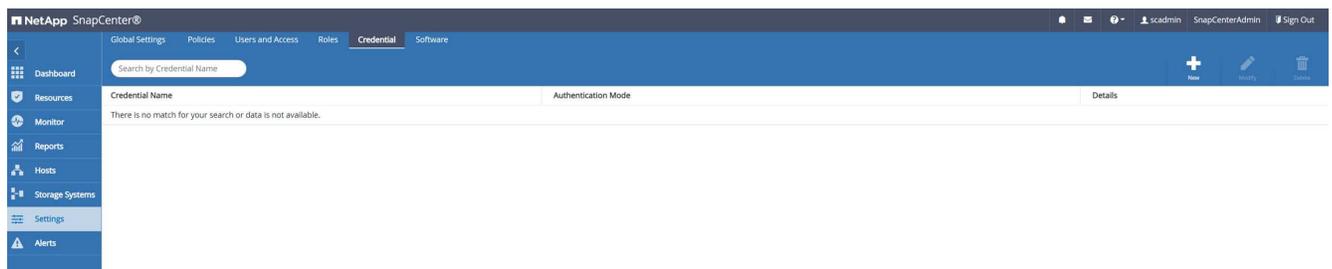
The SVM `sapcc-hana-svm` is now configured in SnapCenter.



Create credentials for plugin deployment

To enable SnapCenter to deploy the required plug-ins on the HANA hosts, you must configure user credentials.

1. Go to Settings, select Credentials, and click New.



2. In the lab setup, we configured a new user, `snapcenter`, on the HANA host that is used for the plug-in deployment. You must enable sudo privileges, as shown in the following figure.

Credential ✕

Credential Name:

Authentication Mode:

Username:

Password:

Use sudo privileges ?

```
hana-1:/etc/sudoers.d # cat /etc/sudoers.d/90-cloud-init-users
# Created by cloud-init v. 20.2-8.48.1 on Mon, 14 Feb 2022 10:36:40 +0000
# User rules for ec2-user
ec2-user ALL=(ALL) NOPASSWD:ALL
# User rules for snapcenter user
snapcenter ALL=(ALL) NOPASSWD:ALL
hana-1:/etc/sudoers.d #
```

Add a SAP HANA host

When adding an SAP HANA host, SnapCenter deploys the required plug-ins on the database host and executes auto discovery operations.

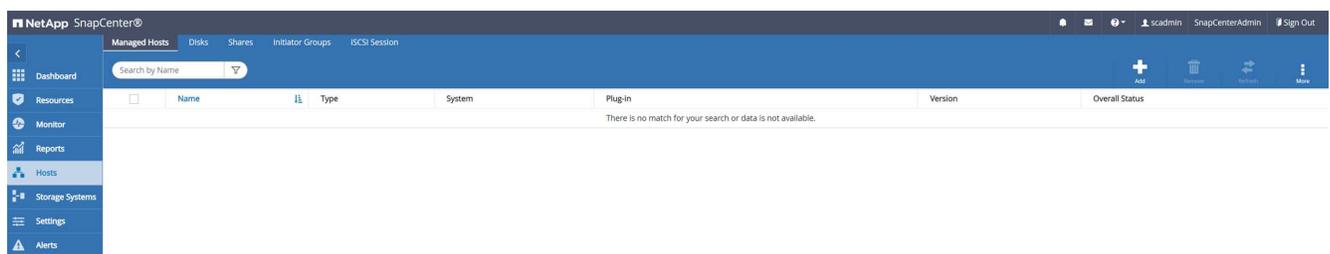
The SAP HANA plug-in requires Java 64-bit version 1.8. Java must be installed on the host before the host is added to SnapCenter.

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

OpenJDK or Oracle Java is supported with SnapCenter.

To add the SAP HANA host, follow these steps:

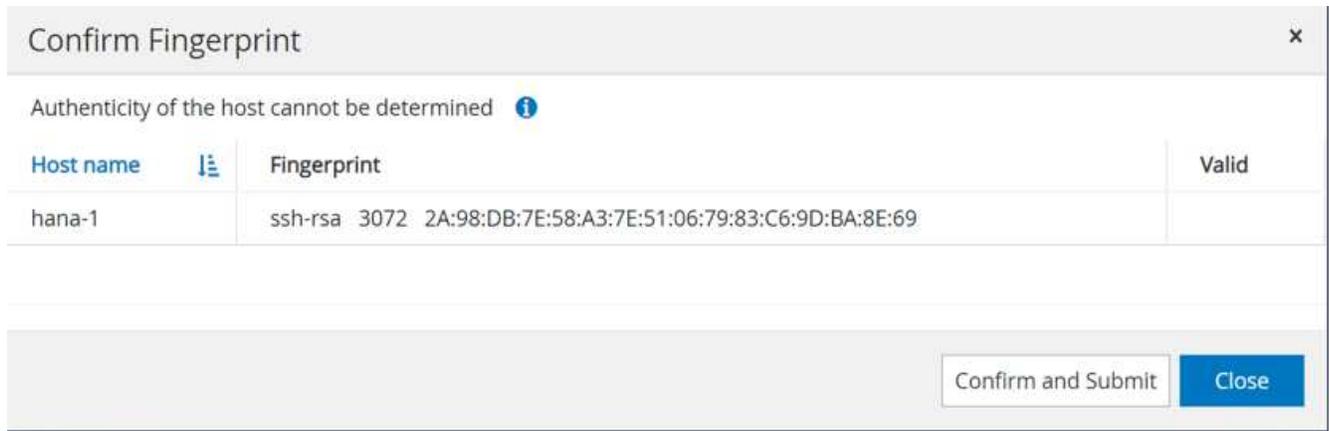
1. From the host tab, click Add.



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



3. Confirm the fingerprint.

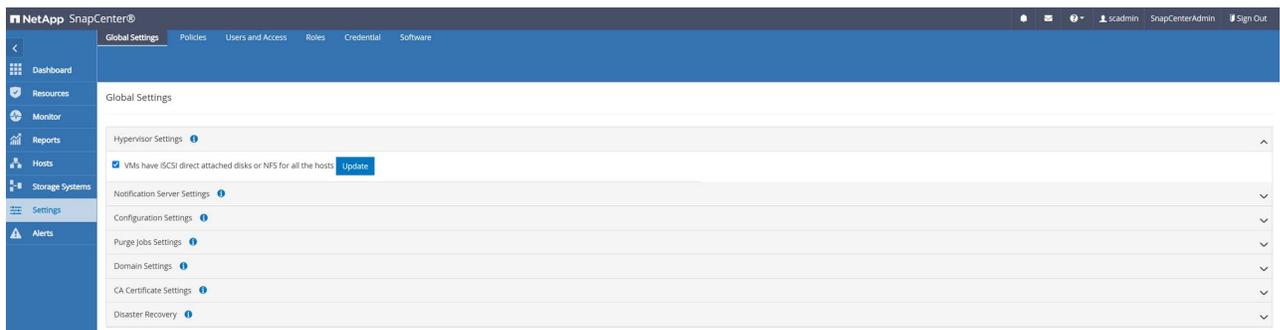


The installation of the HANA and the Linux plug-in starts automatically. When the installation is finished, the status column of the host shows Configure VMware Plug-in. SnapCenter detects if the SAP HANA plug-in is installed on a virtualized environment. This might be a VMware environment or an environment at a public cloud provider. In this case, SnapCenter displays a warning to configure the hypervisor.

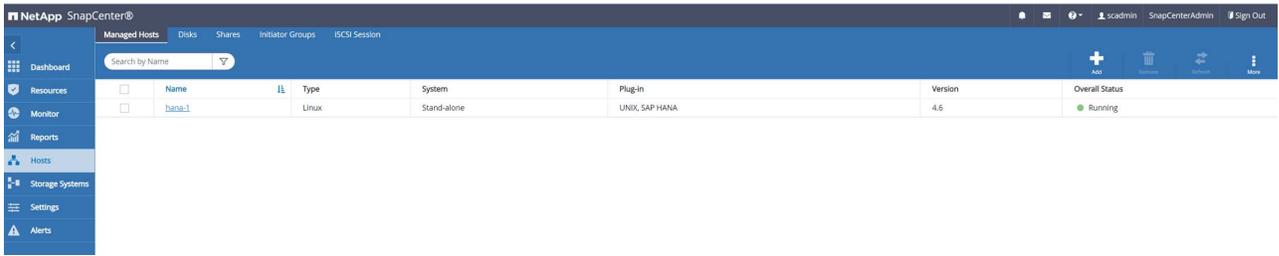
You can remove the warning message by using the following steps.



- From the Settings tab, select Global Settings.
- For the hypervisor settings, select VMs Have iSCSI Direct Attached Disks or NFS For All the Hosts and update the settings.



The screen now shows the Linux plug-in and the HANA plug-in with the status Running.



Configure policies

Policies are usually configured independently of the resource and can be used by multiple SAP HANA databases.

A typical minimum configuration consists of the following policies:

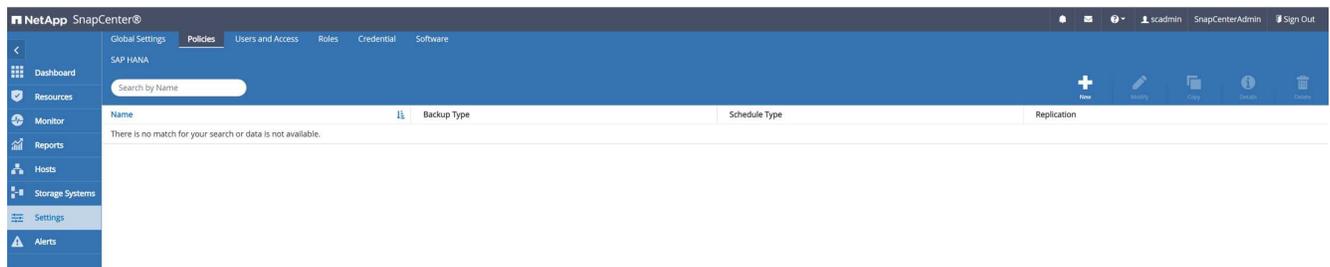
- Policy for hourly backups without replication: LocalSnap.
- Policy for weekly block integrity check using a file-based backup: BlockIntegrityCheck.

The following sections describe the configuration of these policies.

Policy for Snapshot backups

Follow these steps to configure Snapshot backup policies.

1. Go to Settings > Policies and click New.



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

New SAP HANA Backup Policy ✕

- 1 Name
- 2 Settings
- 3 Retention
- 4 Replication
- 5 Summary

Provide a policy name

Policy name

Details

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

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

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Select backup settings

Backup Type Snapshot Based File-Based i

Schedule Frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

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

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Retention settings

Hourly retention settings

Total Snapshot copies to keep i

Keep Snapshot copies for days

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

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Select secondary replication options i

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label i

Error retry count i

New SAP HANA Backup Policy ✕

1 Name	Summary	
2 Settings	Policy name	LocalSnap
3 Retention	Details	Snapshot backup at primary volume
4 Replication	Backup Type	Snapshot Based Backup
5 Summary	Schedule Type	Hourly
	Hourly backup retention	Total backup copies to retain : 7
	Replication	none

The new policy is now configured.



Policy for block integrity check

Follow these steps to configure the block integrity check policy.

1. Go to Settings > Policies and click New.
2. Enter the policy name and description. Click Next.

New SAP HANA Backup Policy ✕

1 Name	Provide a policy name	
2 Settings	Policy name	<input style="width: 100%;" type="text" value="BlockIntegrityCheck"/>
3 Retention	Details	<input style="width: 100%;" type="text" value="Check HANA DB blocks using file-based backup"/>
4 Replication		
5 Summary		

3. Set the backup type to File-Based and schedule frequency to Weekly. The schedule itself is configured later with the HANA resource protection configuration.

New SAP HANA Backup Policy ✕

1 Name

2 Settings

3 Retention

4 Summary

Select backup settings

Backup Type Snapshot Based File-Based i

Schedule Frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

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

New SAP HANA Backup Policy ✕

1 Name

2 Settings

3 Retention

4 Summary

Retention settings

Weekly retention settings

Total backup copies to keep i

Keep backup copies for days

5. On the Summary page, click Finish.

New SAP HANA Backup Policy ✕

1 Name

2 Settings

3 Retention

4 Summary

Summary

Policy name	BlockIntegrityCheck
Details	Check HANA DB blocks using file-based backup
Backup Type	File-Based Backup
Schedule Type	Weekly
Weekly backup retention	Total backup copies to retain : 1

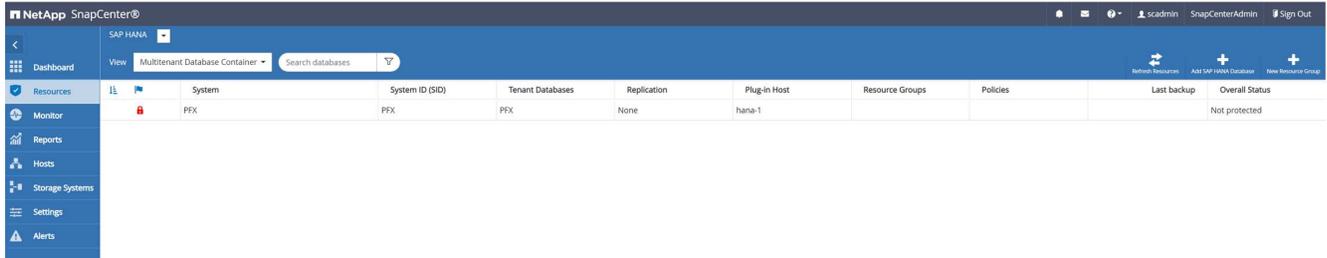
Name	Backup Type	Schedule Type	Replication
BlockIntegrityCheck	File Based Backup	Weekly	
LocalSnap	Data Backup	Hourly	

Configure and protect a HANA resource

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

1. Select and click the resource to continue the configuration.

You can also trigger the automatic discovery process manually within the Resources screen by clicking Refresh Resources.



2. Provide the userstore key for the HANA database.

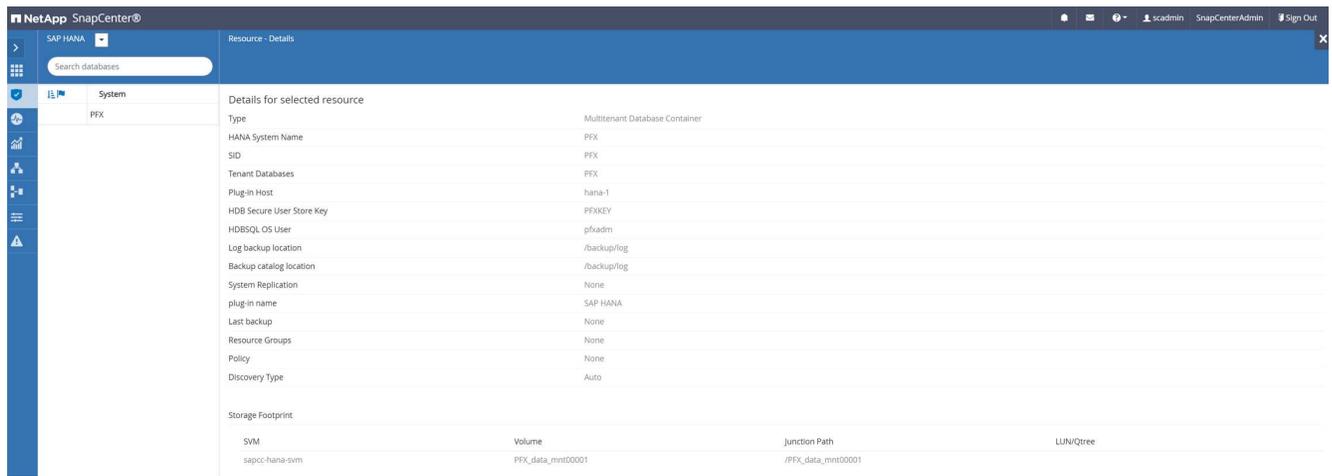
Configure Database

Plug-in host: hana-1

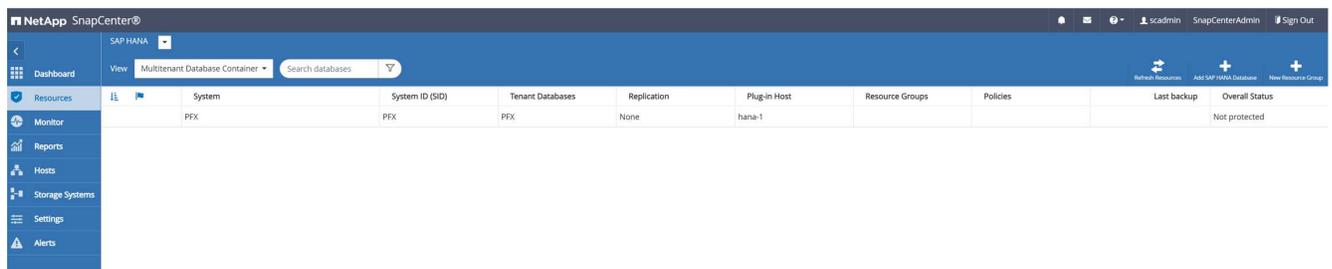
HDBSQL OS User: pfxadm

HDB Secure User Store Key:

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

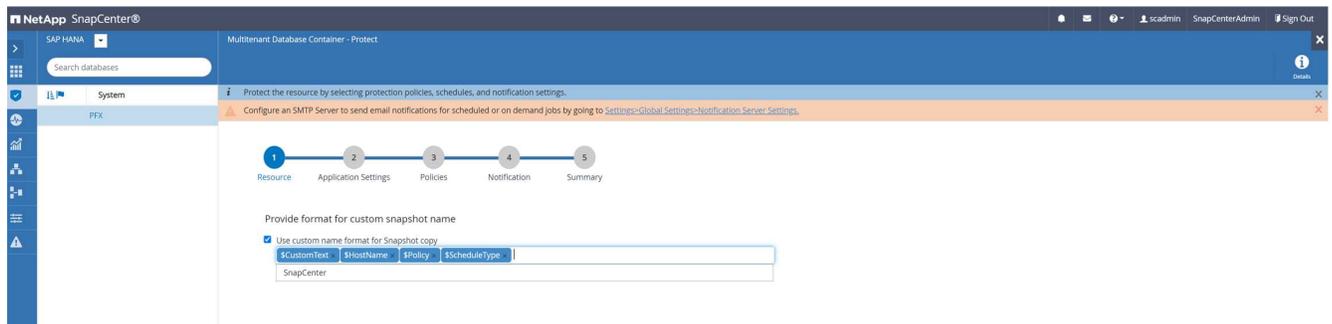


3. From the Resources tab, double click the resource to configure the resource protection.

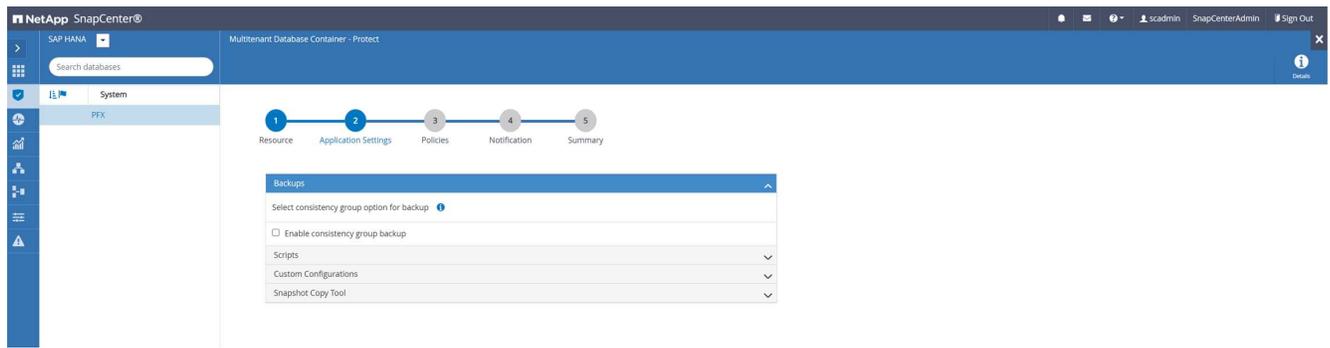


4. Configure a custom name format for the Snapshot copy.

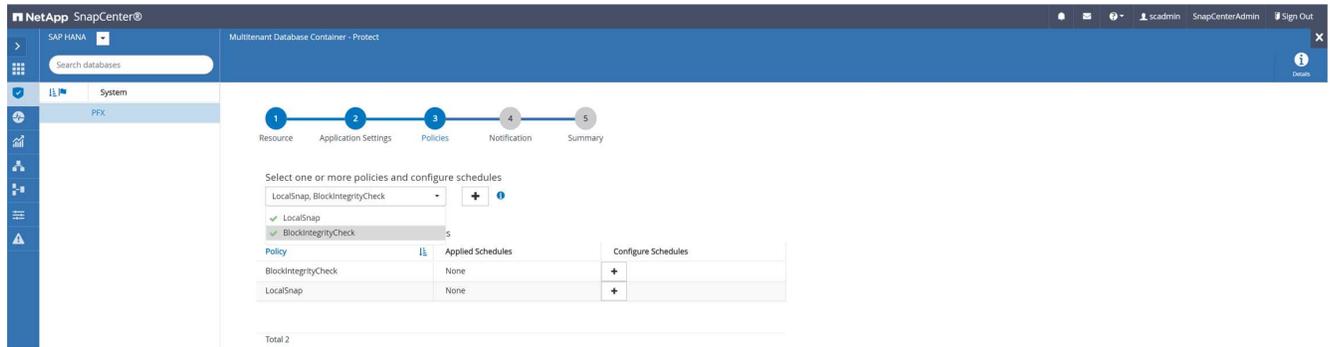
NetApp recommends using a custom Snapshot copy name to easily identify which backups have been created with which policy and schedule type. By adding the schedule type in the Snapshot copy name, you can distinguish between scheduled and on-demand backups. The `schedule` name string for on-demand backups is empty, while scheduled backups include the string `Hourly`, `Daily`, or `Weekly`.



5. No specific setting needs to be made on the Application Settings page. Click Next.



6. Select the policies to be added to the resource.



7. Define the schedule for the block integrity check policy.

In this example, it is set for once per week.

Add schedules for policy BlockIntegrityCheck



Weekly

Start date

02/22/2022 12:00 pm



Expires on

03/22/2022 12:00 pm



Days

Sunday

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday



The schedules are triggered in the SnapCenter Server time zone.



Cancel

OK

8. Define the schedule for the local Snapshot policy.

In this example, it is set for every 6 hours.

Modify schedules for policy LocalSnap



Hourly

Start date

02/22/2022 02:00 pm



Expires on

04/28/2022 11:57 am



Repeat every

6

hours

0

mins



The schedules are triggered in the SnapCenter Server time zone.



Cancel

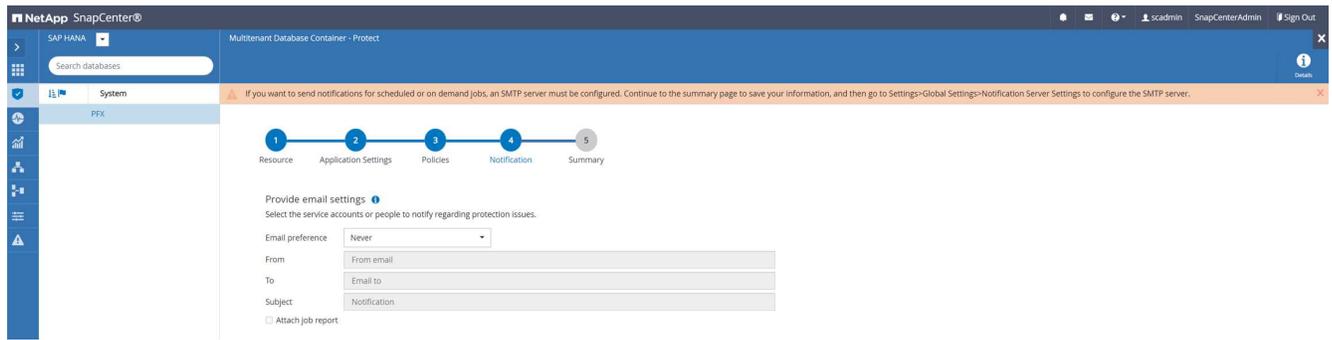
OK

The screenshot shows the NetApp SnapCenter interface for configuring a policy. The breadcrumb trail is: Resource > Application Settings > Policies > Notification > Summary. The 'Policies' step is active. A dropdown menu shows 'LocalSnap, BlockIntegrityCheck' with a plus icon and a refresh icon. Below, a table titled 'Configure schedules for selected policies' shows the configuration for two policies:

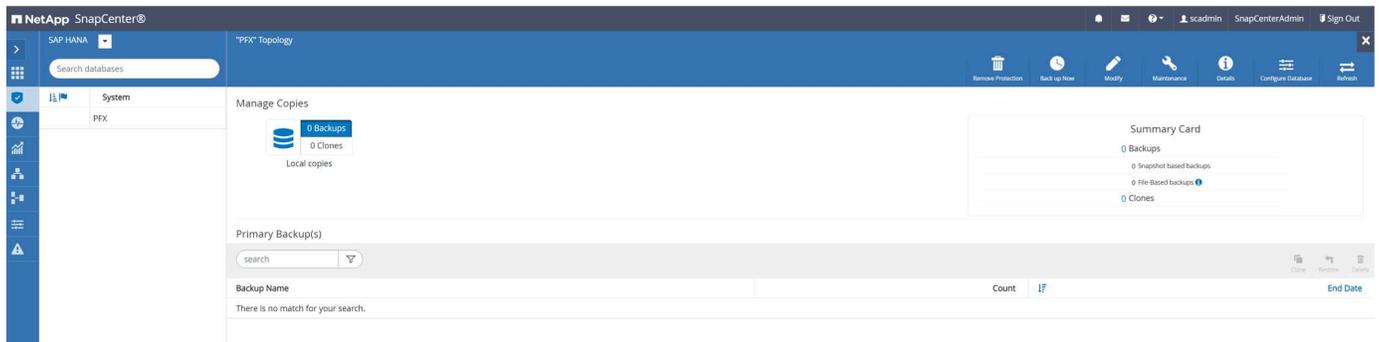
Policy	Applied Schedules	Configure Schedules
BlockIntegrityCheck	Weekly; Run on days: Sunday	
LocalSnap	Hourly; Repeat every 6 hours	

Total 2

9. Provide information about the email notification.



The HANA resource configuration is now completed, and you can execute backups.



SnapCenter backup operations

You can create an on-demand Snapshot backup and an on-demand block integrity check operation.

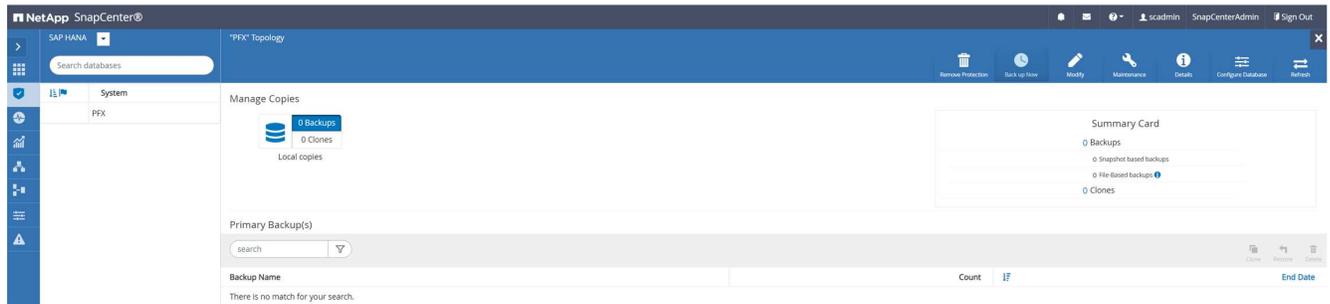
Create an on-demand Snapshot backup

Follow these steps to create on-demand Snapshot backups.

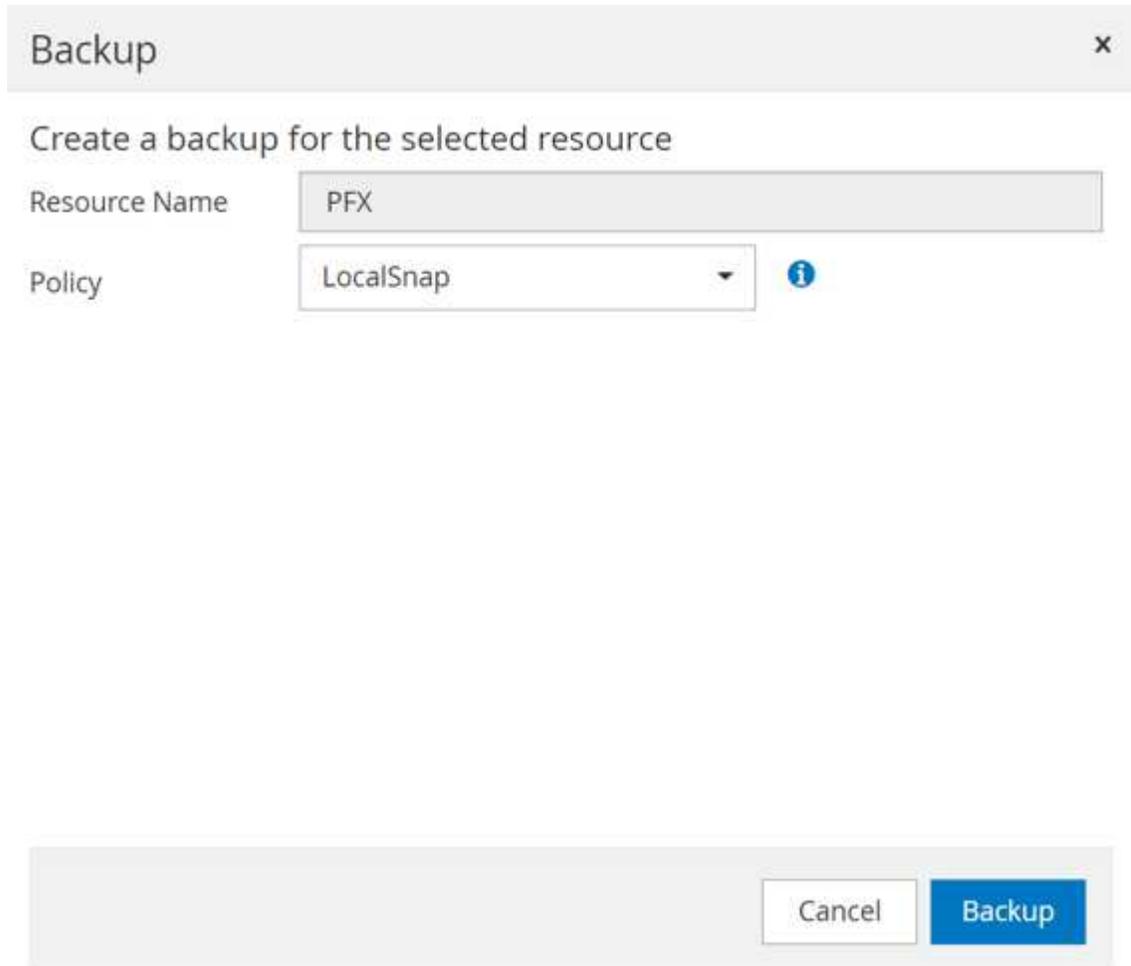
1. In the Resource view, select the resource and double-click the line to switch to the Topology view.

The Resource Topology view provides an overview of all available backups that have been created by using SnapCenter. The top area of this view displays the backup topology showing the backups on the primary storage (local copies) and, if available, on the off-site backup storage (vault copies).

2. In the top row, select the Back up Now icon to start an on-demand backup.



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



Confirmation



The policy selected for the on-demand backup is associated with a backup schedule and the on-demand backups will be retained based on the retention settings specified for the schedule type. Do you want to continue?

Yes

No

A log of the previous five jobs is shown in the Activity area at the bottom of the Topology view.

4. The job details are shown when clicking the job's activity line in the Activity area. You can open a detailed job log by clicking View Logs

Job Details

Backup of Resource Group 'hana-1_hana_MDC_PFX' with policy 'LocalSnap'

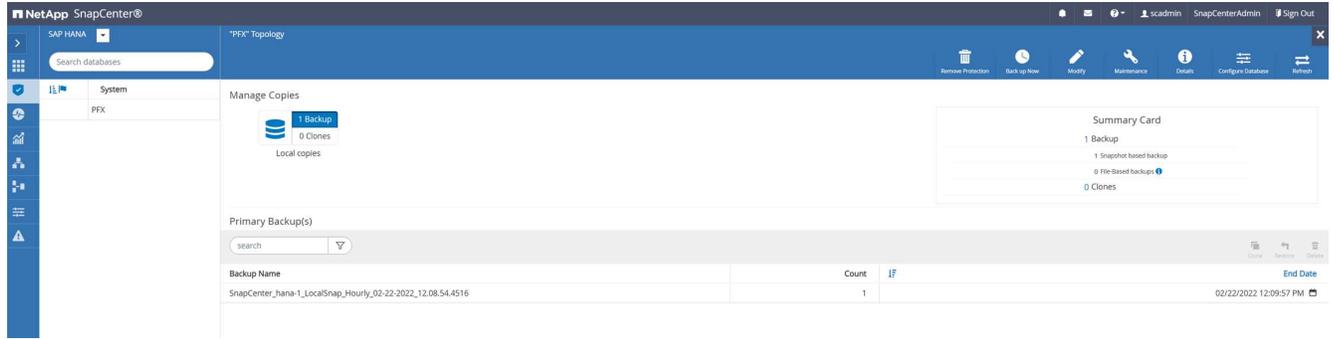
- ✓ Backup of Resource Group 'hana-1_hana_MDC_PFX' with policy 'LocalSnap'
 - ✓ hana-1
 - Backup
 - Validate Dataset Parameters
 - Validate Plugin Parameters
 - Complete Application Discovery
 - Initialize Filesystem Plugin
 - Discover Filesystem Resources
 - Validate Retention Settings
 - Quiesce Application
 - Quiesce Filesystem
 - Create Snapshot
 - UnQuiesce Filesystem
 - UnQuiesce Application
 - Get Snapshot Details
 - Get Filesystem Meta Data
 - Finalize Filesystem Plugin
 - Collect Autosupport data
 - Register Backup and Apply Retention
 - Register Snapshot attributes
 - Application Clean-Up
 - Data Collection
 - Agent Finalize Workflow

Task Name: Backup Start Time: 02/22/2022 12:08:58 PM End Time: 02/22/2022 12:10:21 PM

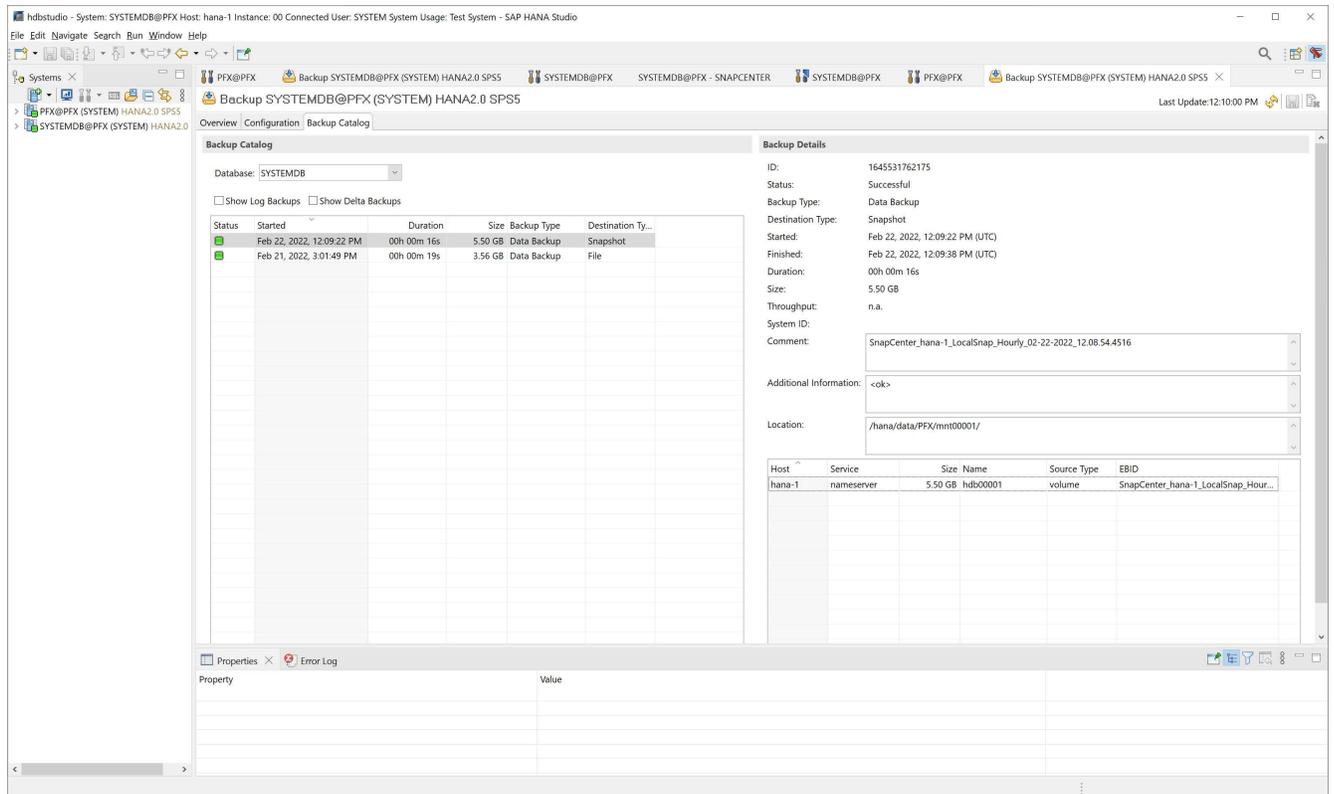
View Logs Cancel Job Close

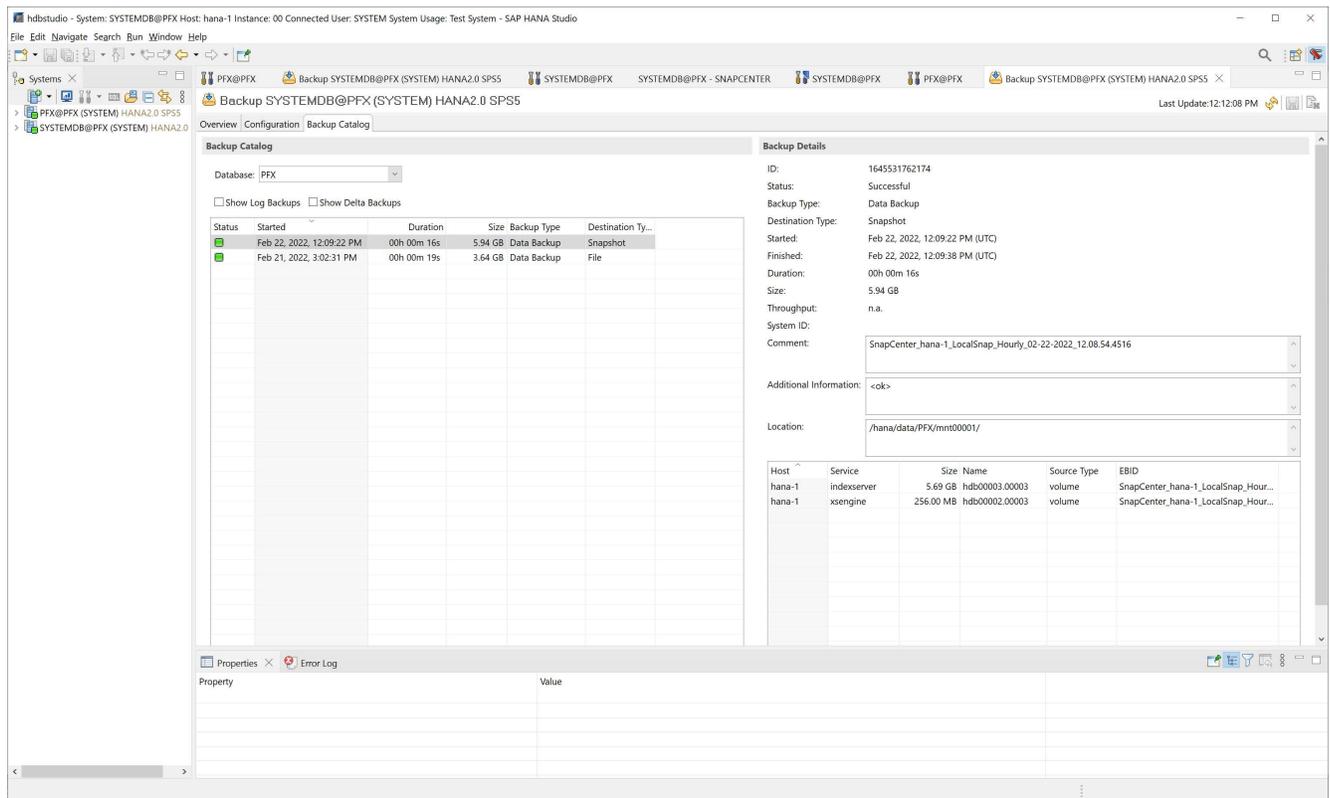
When the backup is finished, a new entry is shown in the topology view. The backup names follow the same naming convention as the Snapshot name defined in the section [“Configure and protect a HANA resource”](#).

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



In the SAP HANA backup catalog, the SnapCenter backup name is stored as a Comment field as well as External Backup ID (EBID). This is shown in the following figure for the system database and in the next figure for the tenant database PFX.





On the FSx for ONTAP file system, you can list the Snapshot backups by connecting to the console of the SVM.

```

sapcc-hana-svm::> snapshot show -volume PFX_data_mnt00001
---Blocks---
Vserver   Volume      Snapshot                                     Size Total%
Used%
-----
-----
sapcc-hana-svm
          PFX_data_mnt00001
                SnapCenter_hana-1_LocalSnap_Hourly_02-22-
2022_12.08.54.4516
                                                126.6MB      0%
2%
sapcc-hana-svm::>

```

Create an on-demand block integrity check operation

An on-demand block integrity check operation is executed in the same way as a Snapshot backup job, by selecting the policy BlockIntegrityCheck. When scheduling backups using this policy, SnapCenter creates a standard SAP HANA file backup for the system and tenant databases.

Backup



Create a backup for the selected resource

Resource Name

PFX

Policy

BlockIntegrityCheck



Cancel

Backup

Job Details



Backup of Resource Group 'hana-1_hana_MDC_PFX' with policy 'BlockIntegrityCheck'

✓ ▾ Backup of Resource Group 'hana-1_hana_MDC_PFX' with policy 'BlockIntegrityCheck'

✓ ▾ hana-1

✓ ▾ File-Based Backup

- ✓ ▶ Validate Plugin Parameters
- ✓ ▶ Start File-Based Backup
- ✓ ▶ Check File-Based Backup
- ✓ ▶ Register Backup and Apply Retention
- ✓ ▶ Data Collection

i Task Name: File-Based Backup Start Time: 02/22/2022 12:55:21 PM End Time: 02/22/2022 12:56:36 PM

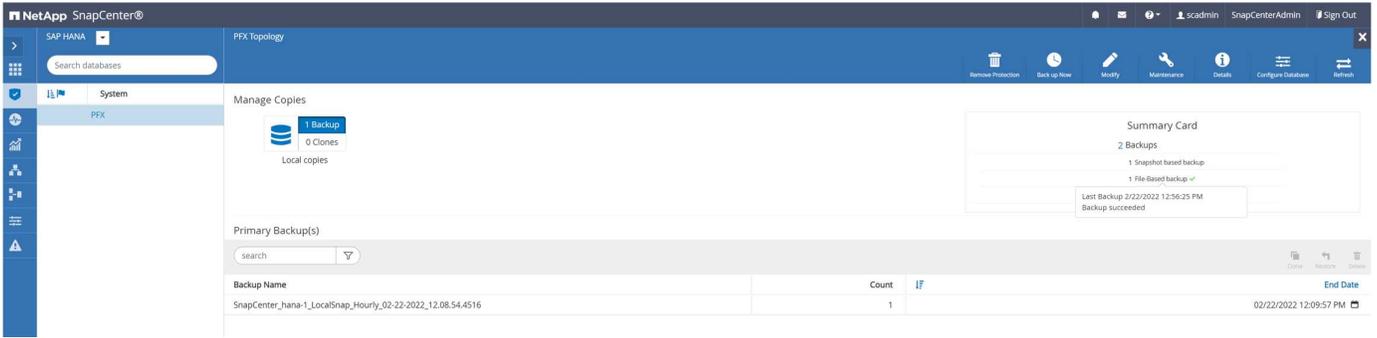
View Logs

Cancel Job

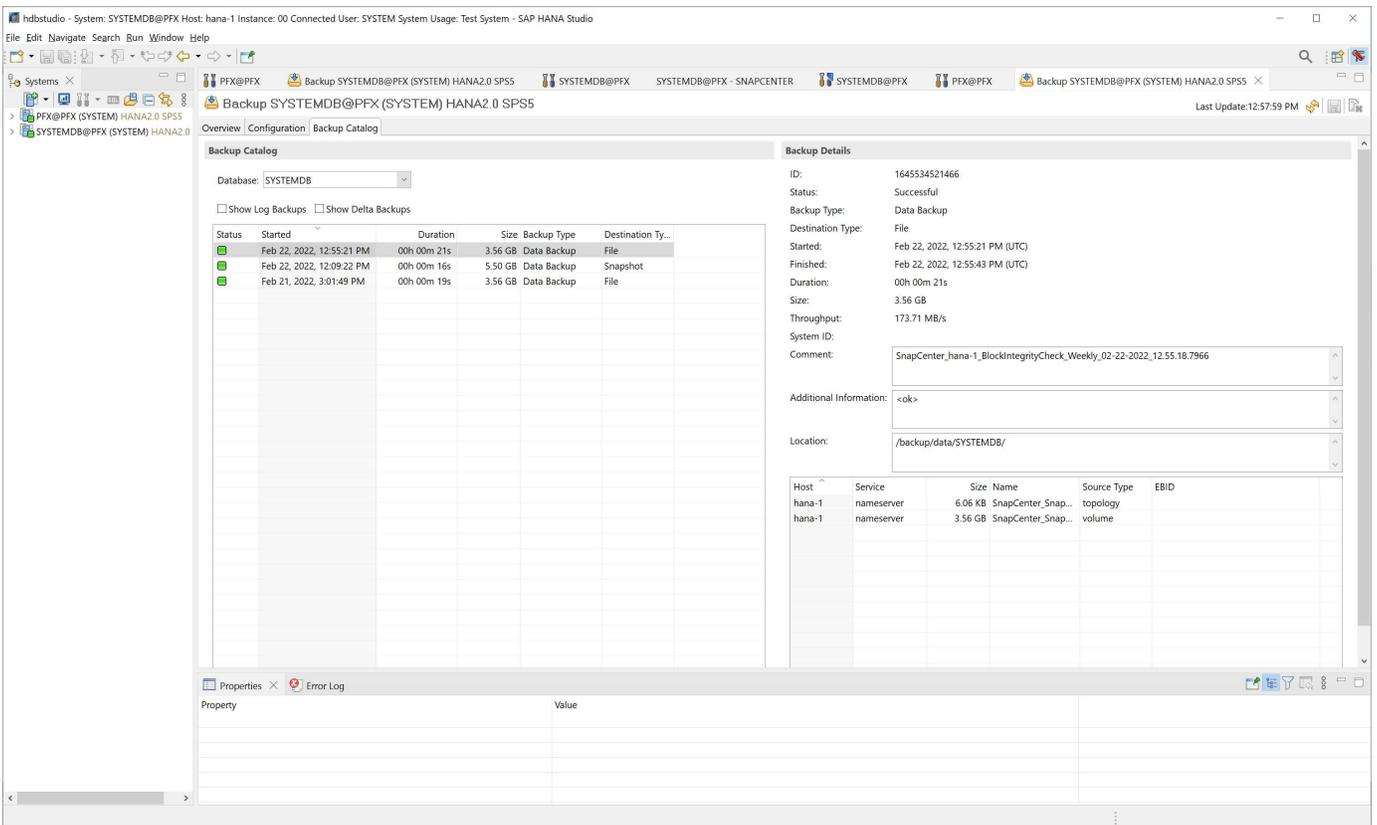
Close

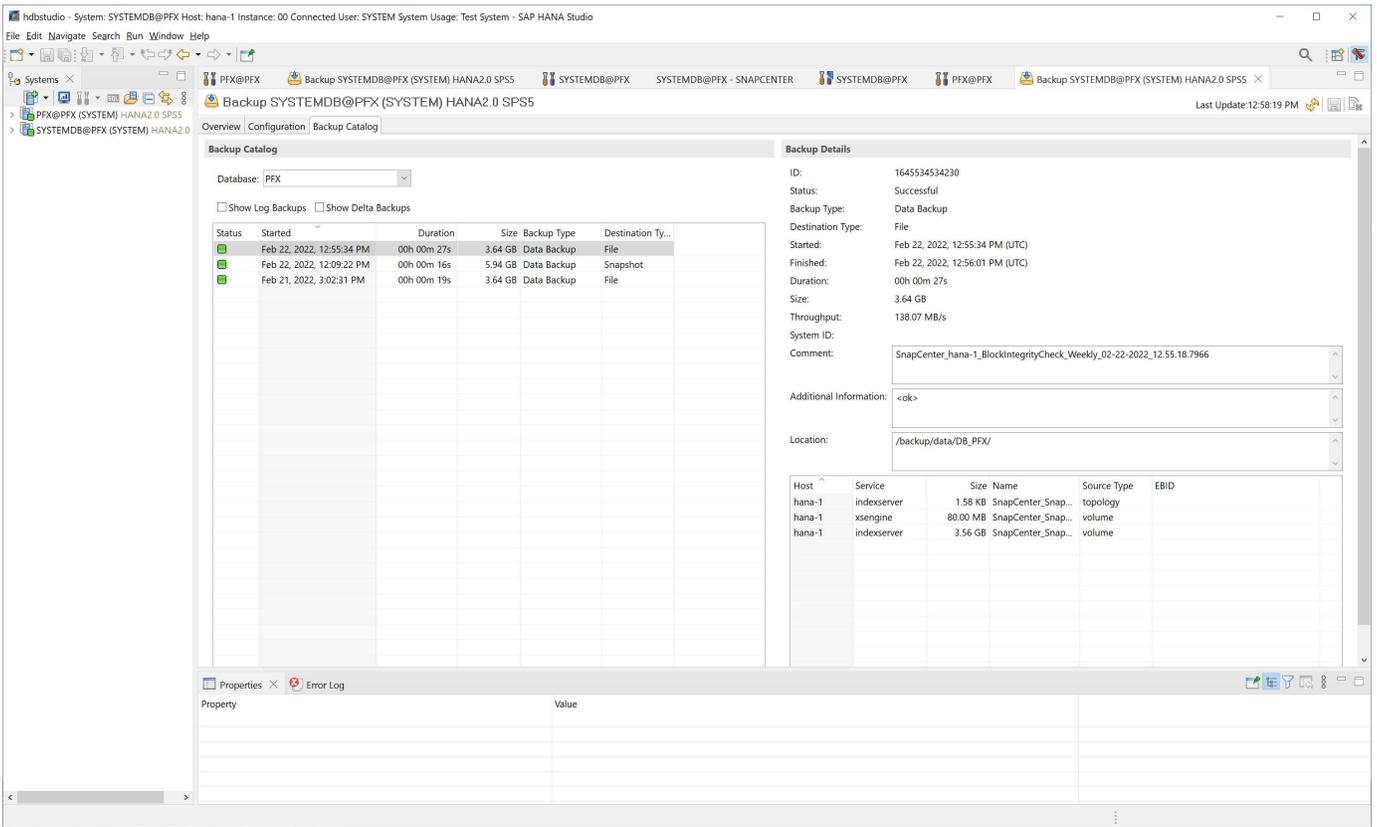
SnapCenter does not display the block integrity check in the same manner as Snapshot copy-based backups.

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



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





A successful block integrity check creates standard SAP HANA data backup files. SnapCenter uses the backup path that has been configured with the HANA database for file-based data backup operations.

```

hana-1:~ # ls -al /backup/data/*
/backup/data/DB_PFX:
total 7665384
drwxr-xr-- 2 pfxadm sapsys      4096 Feb 22 12:56 .
drwxr-xr-x 4 pfxadm sapsys      4096 Feb 21 15:02 ..
-rw-r----- 1 pfxadm sapsys    155648 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_0_1
-rw-r----- 1 pfxadm sapsys    83894272 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_2_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_3_1
-rw-r----- 1 pfxadm sapsys     155648 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_0_1
-rw-r----- 1 pfxadm sapsys    83894272 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_2_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 22 12:56
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_3_1
/backup/data/SYSTEMDB:
total 7500880
drwxr-xr-- 2 pfxadm sapsys      4096 Feb 22 12:55 .
drwxr-xr-x 4 pfxadm sapsys      4096 Feb 21 15:02 ..
-rw-r----- 1 pfxadm sapsys    159744 Feb 21 15:01
COMPLETE_DATA_BACKUP_databackup_0_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 21 15:02
COMPLETE_DATA_BACKUP_databackup_1_1
-rw-r----- 1 pfxadm sapsys     159744 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_0_1
-rw-r----- 1 pfxadm sapsys   3825213440 Feb 22 12:55
SnapCenter_SnapCenter_hana-1_BlockIntegrityCheck_Weekly_02-22-
2022_12.55.18.7966_databackup_1_1
hana-1:~ #

```

Backup of non-data volumes

The backup of non-data volumes is an integrated part of the SnapCenter and the SAP HANA plug-in.

Protecting the database data volume is sufficient to restore and recover the SAP HANA database to a given point in time, provided that the database installation resources, and the required logs are still available.

To recover from situations where other non-data files must be restored, NetApp recommends developing an

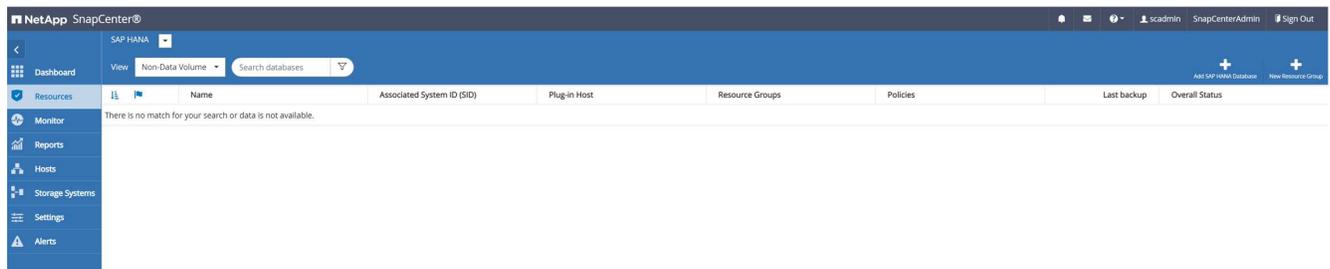
additional backup strategy for non-data volumes to augment the SAP HANA database backup. Depending on your specific requirements, the backup of non-data volumes might differ in scheduling frequency and retention settings, and you should consider how frequently non-data files are changed. For instance, the HANA volume `/hana/shared` contains executables but also SAP HANA trace files. While executables only change when the SAP HANA database is upgraded, the SAP HANA trace files might need a higher backup frequency to support analyzing problem situations with SAP HANA.

SnapCenter non-data volume backup enables Snapshot copies of all relevant volumes to be created in a few seconds with the same space efficiency as SAP HANA database backups. The difference is that there is no SQL communication with SAP HANA database required.

Configure non-data volume resources

Follow these steps to configure non-data volume resources:

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



2. In step one of the Add SAP HANA Database dialog, in the Resource Type list, select Non- data Volumes. Specify a name for the resource and the associated SID and the SAP HANA plug-in host that you want to use for the resource, then click Next.

Add SAP HANA Database

1 Name

2 Storage Footprint

3 Summary

Provide Resource Details

Resource Type	Non-data Volume
Resource Name	PFX-Shared-Volume
Associated SID	PFX
Plug-in Host	hana-1

Previous Next

3. Add the SVM and the storage volume as storage footprint, then click Next.

Add SAP HANA Database x

1 Name

2 Storage Footprint

3 Summary

Provide Storage Footprint Details

Storage Type ONTAP

Add Storage Footprint x

Storage System

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

Volume name	LUNs or Qtrees
<input type="text" value="PFX_shared"/>	<input type="text" value="Default is 'None' or type to find"/>

4. To save the settings, in the summary step, click Finish.

Add SAP HANA Database

- 1 Name
- 2 Storage Footprint
- 3 Summary

Summary

Resource Type	Non-data Volume
Resource Name	PFX-Shared-Volume
Associated SID	PFX
Plug-in Host	hana-1

Storage Footprint

Storage System	Volume	LUN/Qtree
sapcc-hana-svm	PFX_shared	

Previous
Finish

The new non-data volume is now added to SnapCenter. Double click the new resource to execute the resource protection.

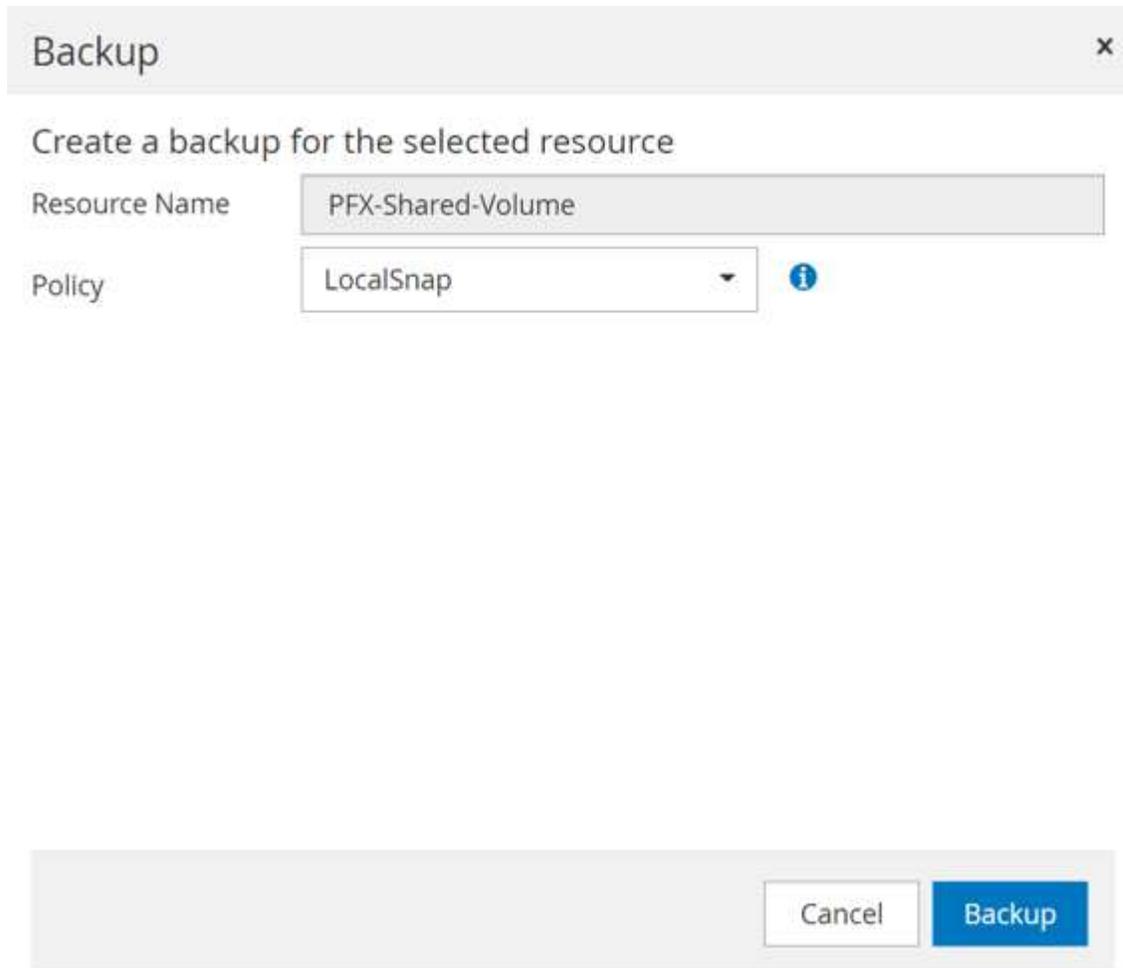
Name	Associated System ID (SID)	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
PFX-Shared-Volume	PFX	hana-1				Not protected

The resource protection is done in the same way as described before with a HANA database resource.

5. You can now execute a backup by clicking on Backup Now.



6. Select the policy and start the backup operation.



The SnapCenter job log shows the individual workflow steps.

Job Details



Backup of Resource Group 'hana-1_hana_NonDataVolume_PFX_PFX-Shared-Volume' with policy 'LocalSnap'

✓ ▾ Backup of Resource Group 'hana-1_hana_NonDataVolume_PFX_PFX-Shared-Volume' with policy 'LocalSnap'

✓ ▾ hana-1

✓ ▾ Backup

- ✓ ▶ Validate Dataset Parameters
- ✓ ▶ Validate Plugin Parameters
- ✓ ▶ Validate Retention Settings
- ✓ ▶ Create Snapshot
- ✓ ▶ Get Snapshot Details
- ✓ ▶ Collect Autosupport data
- ✓ ▶ Register Backup and Apply Retention
- ✓ ▶ Register Snapshot attributes
- ✓ ▶ Data Collection
- ✓ ▶ Agent Finalize Workflow

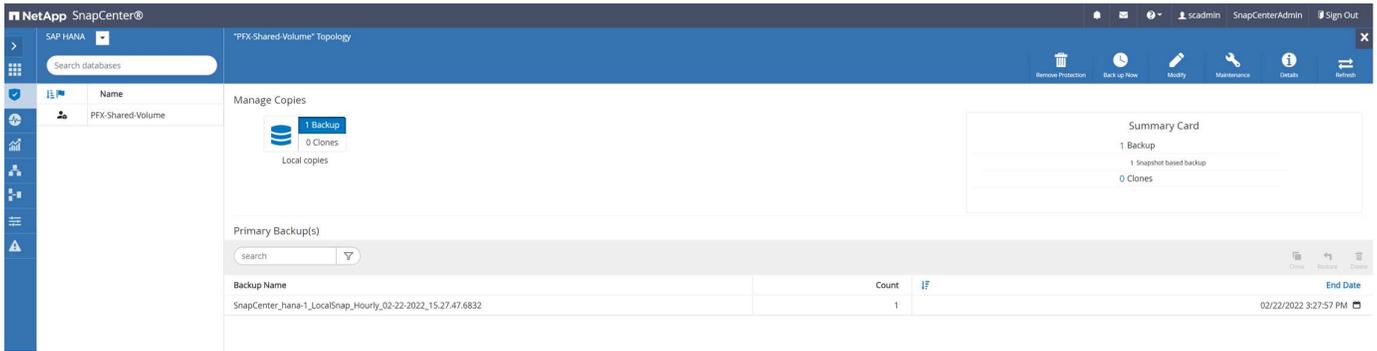
i Task Name: Backup Start Time: 02/22/2022 3:27:48 PM End Time:

View Logs

Cancel Job

Close

The new backup is now visible in the resource view of the non- data volume resource.



Restore and recover

With SnapCenter, automated restore and recovery operations are supported for HANA single host MDC systems with a single tenant. For multiple-host systems or MDC systems with multiple tenants, SnapCenter only executes the restore operation and you must perform the recovery manually.

You can execute an automated restore and recovery operation with the following steps:

1. Select the backup to be used for the restore operation.
2. Select the restore type. Select Complete Restore with Volume Revert or without Volume Revert.
3. Select the recovery type from the following options:
 - To most recent state
 - Point in time
 - To specific data backup
 - No recovery

The selected recovery type is used for the recovery of the system and the tenant database.

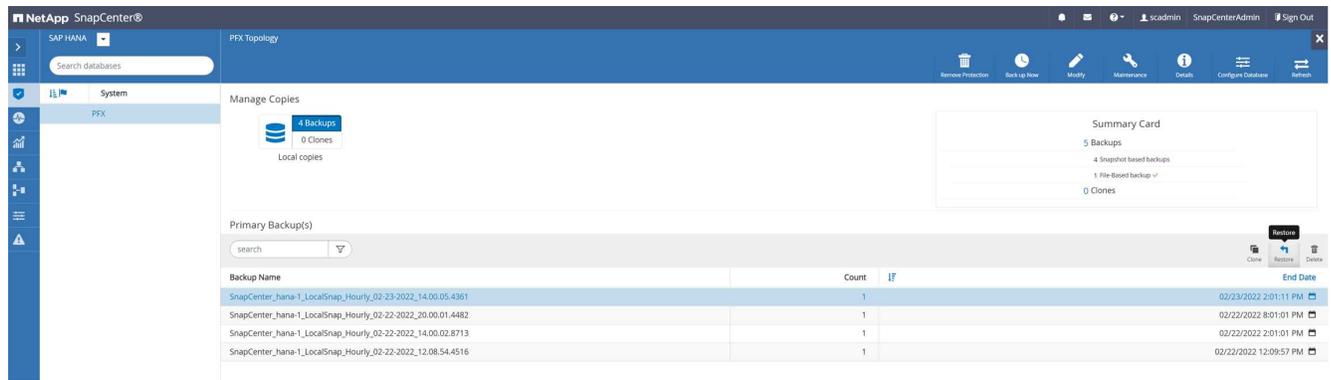
Next, SnapCenter performs the following operations:

1. It stops the HANA database.
2. It restores the database. Depending on the selected restore type, different operations are executed.
 - If Volume Revert is selected, then SnapCenter unmounts the volume, restores the volume by using volume-based SnapRestore on the storage layer, and mounts the volume.
 - If Volume Revert is not selected, then SnapCenter restores all files by using single file SnapRestore operations on the storage layer.
3. It recovers the database:
 - a. By recovering the system database
 - b. recovering the tenant database
 - c. starting the HANA database

If No Recovery is selected, SnapCenter exits, and you must perform the restore operation for the system and the tenant database manually.

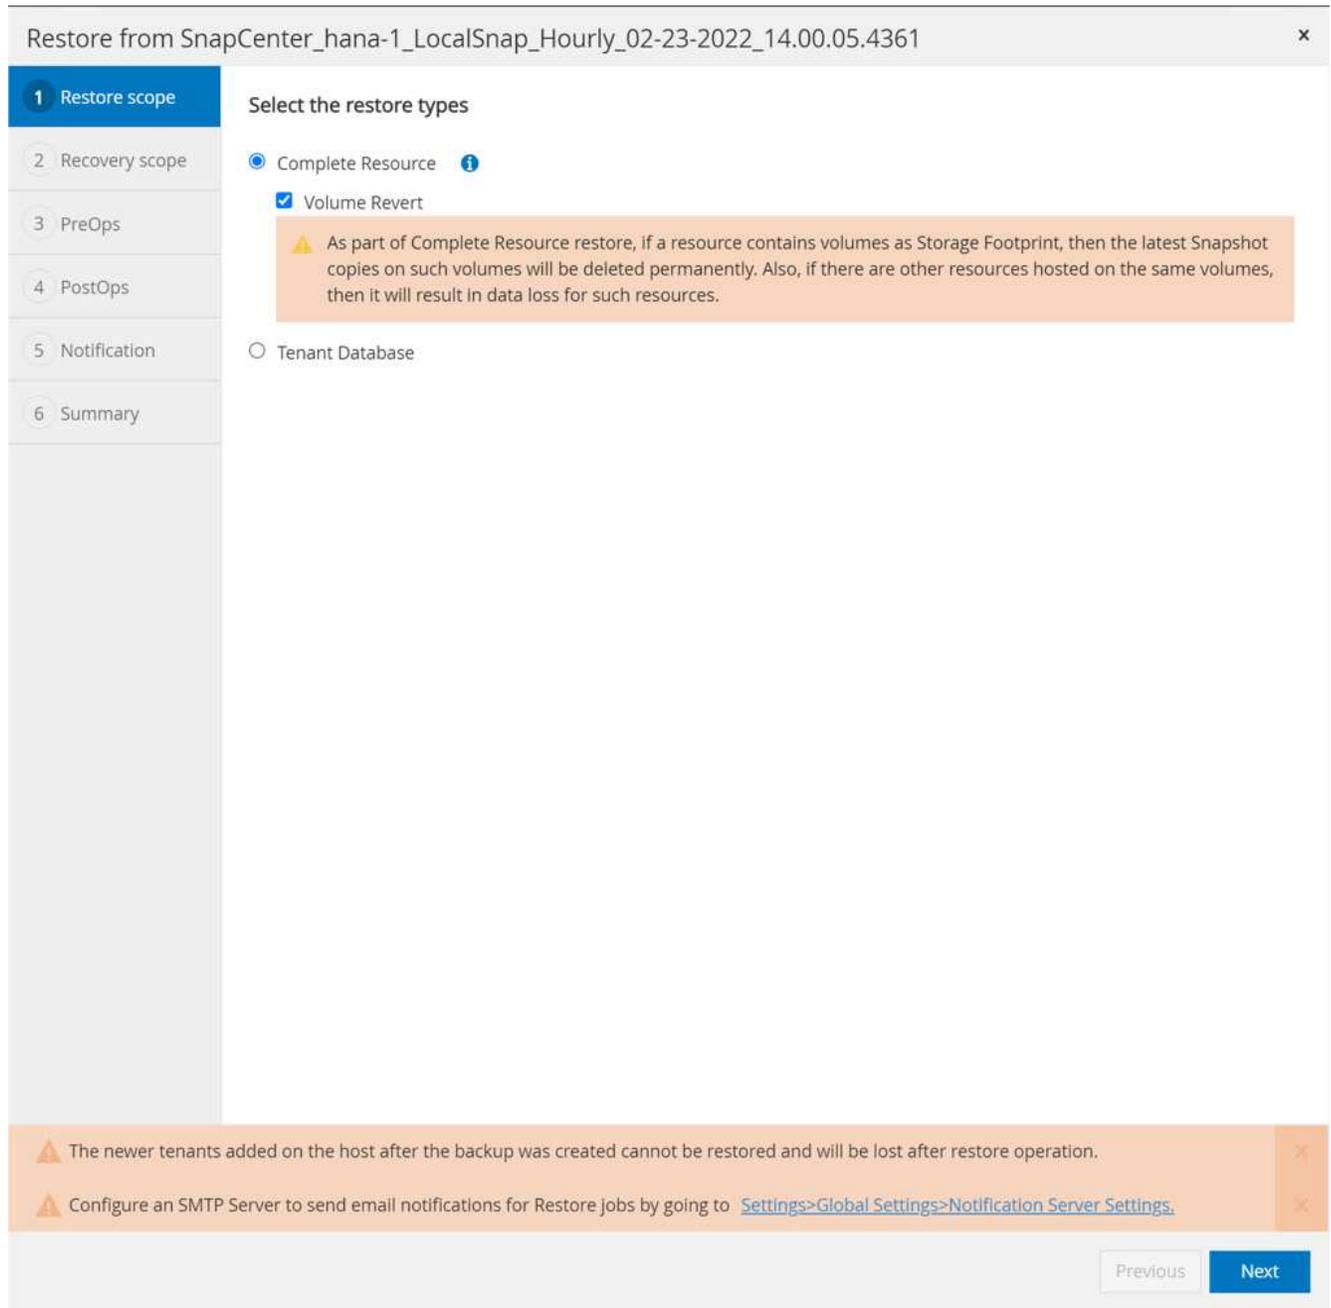
To perform a manual restore operation, follow these steps:

1. Select a backup in SnapCenter to be used for the restore operation.



2. Select the restore scope and type.

The standard scenario for HANA MDC single tenant systems is to use complete resource with volume revert. For a HANA MDC system with multiple tenants, you might want to restore only a single tenant. For more information about the single tenant restore, see [Restore and recovery \(netapp.com\)](https://netapp.com).



3. Select Recovery Scope and provide the location for log backup and catalog backup.

SnapCenter uses the default path or the changed paths in the HANA global.ini file to pre-populate the log and catalog backup locations.

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

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

Recover database files using

- Recover to most recent state ?
- Recover to point in time ?
- Recover to specified data backup ?
- No recovery ?

Specify log backup locations ?

Add

Specify backup catalog location ?

⚠ Recovery options are applicable to both system database and tenant database. ×

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

Previous Next

4. Enter the optional pre-restore commands.

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

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

Enter optional commands to run before performing a restore operation ?

Pre restore command

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

Previous Next

5. Enter the optional post-restore commands.

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

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

Enter optional commands to run after performing a restore operation ⓘ

Post restore command

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

Previous Next

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

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

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

Summary

Backup Name	SnapCenter_hana-1_LocalSnap_Hourly_02-23-2022_14.00.05.4361
Backup date	02/23/2022 2:01:11 PM
Restore scope	Complete Resource with Volume Revert
Recovery scope	Recover to most recent state
Log backup locations	/backup/log
Backup catalog location	/backup/log
Pre restore command	
Post restore command	
Send email	No

⚠ If you want to send notifications for Restore Jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous
Finish

SnapCenter executes the restore and recovery operation. This example shows the job details of the restore and recovery job.

Job Details



Restore 'hana-1\hana\MDC\PFX'

- ✓ ▾ Restore 'hana-1\hana\MDC\PFX'
- ✓ ▾ hana-1
 - ✓ ▾ Restore
 - ✓ ▶ Validate Plugin Parameters
 - ✓ ▾ Pre Restore Application
 - ✓ ▶ Stopping HANA instance
 - ✓ ▶ Filesystem Pre Restore
 - ✓ ▾ Restore Filesystem
 - ✓ ▶ Filesystem Post Restore
 - ✓ ▾ Recover Application
 - ✓ ▶ Recovering system database
 - ✓ ▶ Checking HDB services status
 - ✓ ▶ Recovering tenant database 'PFX'
 - ✓ ▶ Starting HANA instance
 - ✓ ▶ Clear Catalog on Server
 - ✓ ▶ Application Clean-Up
 - ✓ ▶ Data Collection
 - ✓ ▶ Agent Finalize Workflow

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

View Logs

Cancel Job

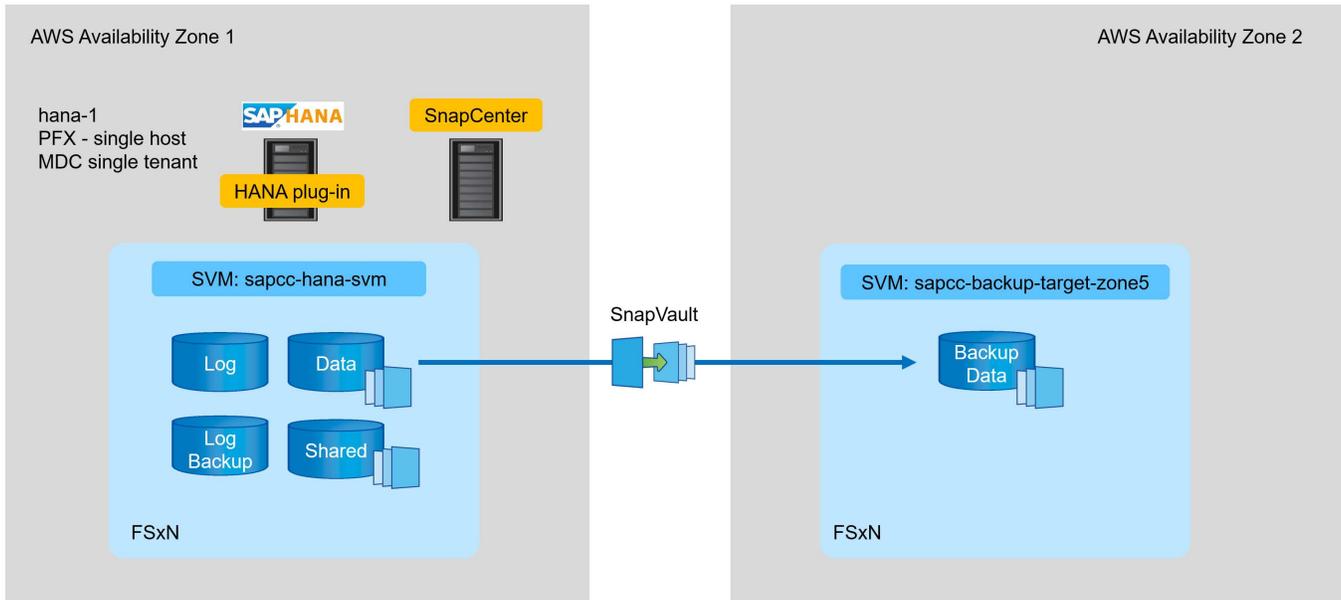
Close

Backup replication with SnapVault

Overview - Backup replication with SnapVault

In our lab setup, we use a second FSX for ONTAP file system in a second AWS availability zone to showcase the backup replication for the HANA data volume.

As discussed in chapter “Data protection strategy”, the replication target must be a second FSx for ONTAP file system in another availability zone to be protected from a failure of the primary FSx for ONTAP file system. Also, the HANA shared volume should be replicated to the secondary FSx for ONTAP file system.



Overview of configuration steps

There are a couple of configuration steps that you must execute on the FSx for ONTAP layer. You can do this either with NetApp Cloud Manager or the FSx for ONTAP command line.

1. Peer FSx for ONTAP file systems. FSx for ONTAP file systems must be peered to allow replication between each other.
2. Peer SVMs. SVMs must be peered to allow replication between each other.
3. Create a target volume. Create a volume at the target SVM with volume type `DP`. Type `DP` is required to be used as a replication target volume.
4. Create a SnapMirror policy. This is used to create a policy for replication with type `vault`.
 - a. Add a rule to policy. The rule contains the SnapMirror label and the retention for backups at the secondary site. You must configure the same SnapMirror label later in the SnapCenter policy so that SnapCenter creates Snapshot backups at the source volume containing this label.
5. Create a SnapMirror relationship. Defines the replication relationship between the source and target volume and attaches a policy.
6. Initialize SnapMirror. This starts the initial replication in which the complete source data is transferred to the target volume.

When volume replication configuration is complete, you must configure the backup replication in SnapCenter

as follows:

1. Add the target SVM to SnapCenter.
2. Create a new SnapCenter policy for Snapshot backup and SnapVault replication.
3. Add the policy to HANA resource protection.
4. You can now execute backups with the new policy.

The following chapters describe the individual steps in more detail.

Configure replication relationships on FSx for ONTAP file systems

You can find additional information about SnapMirror configuration options in the ONTAP documentation at [SnapMirror replication workflow \(netapp.com\)](https://netapp.com).

- Source FSx for ONTAP file system: FsxId00fa9e3c784b6abbb
- Source SVM: sapcc-hana-svm
- Target FSx for ONTAP file system: FsxId05f7f00af49dc7a3e
- Target SVM: sapcc-backup-target-zone5

Peer FSx for ONTAP file systems

```
FsxId00fa9e3c784b6abbb::> network interface show -role intercluster
      Logical      Status      Network      Current      Current
Is
Vserver      Interface  Admin/Oper  Address/Mask      Node      Port
Home
-----
----
FsxId00fa9e3c784b6abbb
      inter_1      up/up      10.1.1.57/24
FsxId00fa9e3c784b6abbb-01
true
      inter_2      up/up      10.1.2.7/24
FsxId00fa9e3c784b6abbb-02
true
      e0e
      e0e
2 entries were displayed.
```

```

FsxId05f7f00af49dc7a3e::> network interface show -role intercluster
          Logical      Status      Network      Current      Current
Is
Vserver   Interface  Admin/Oper  Address/Mask  Node         Port
Home
-----
----
FsxId05f7f00af49dc7a3e
          inter_1      up/up      10.1.2.144/24
FsxId05f7f00af49dc7a3e-01
                                     e0e

true
          inter_2      up/up      10.1.2.69/24
FsxId05f7f00af49dc7a3e-02
                                     e0e

true
2 entries were displayed.

```

```

FsxId05f7f00af49dc7a3e::> cluster peer create -address-family ipv4 -peer
-addr 10.1.1.57, 10.1.2.7
Notice: Use a generated passphrase or choose a passphrase of 8 or more
characters. To ensure the authenticity of the peering relationship, use a
phrase or sequence of characters that would be hard to guess.
Enter the passphrase:
Confirm the passphrase:
Notice: Now use the same passphrase in the "cluster peer create" command
in the other cluster.

```



peer-addr are cluster IPs of the destination cluster.

```

FsxId00fa9e3c784b6abbb::> cluster peer create -address-family ipv4 -peer
-addr 10.1.2.144, 10.1.2.69
Notice: Use a generated passphrase or choose a passphrase of 8 or more
characters. To ensure the authenticity of the peering relationship, use a
phrase or sequence of characters that would be hard to guess.
Enter the passphrase:
Confirm the passphrase:
FsxId00fa9e3c784b6abbb::>
FsxId00fa9e3c784b6abbb::> cluster peer show
Peer Cluster Name          Cluster Serial Number Availability
Authentication
-----
FsxId05f7f00af49dc7a3e    1-80-000011             Available      ok

```

Peer SVMs

```

FsxId05f7f00af49dc7a3e::> vserver peer create -vserver sapcc-backup-
target-zone5 -peer-vserver sapcc-hana-svm -peer-cluster
FsxId00fa9e3c784b6abbb -applications snapmirror
Info: [Job 41] 'vserver peer create' job queued

```

```

FsxId00fa9e3c784b6abbb::> vserver peer accept -vserver sapcc-hana-svm
-peer-vserver sapcc-backup-target-zone5
Info: [Job 960] 'vserver peer accept' job queued

```

```

FsxId05f7f00af49dc7a3e::> vserver peer show
Peer          Peer          Peering
Remote
Vserver      Vserver      State      Peer Cluster      Applications
Vserver
-----
sapcc-backup-target-zone5
peer-source-cluster
peered      FsxId00fa9e3c784b6abbb
snapmirror
sapcc-hana-svm

```

Create a target volume

You must create the target volume with the type DP to flag it as a replication target.

```
FsxId05f7f00af49dc7a3e::> volume create -vserver sapcc-backup-target-zone5
-volume PFX_data_mnt00001 -aggregate aggr1 -size 100GB -state online
-policy default -type DP -autosize-mode grow_shrink -snapshot-policy none
-foreground true -tiering-policy all -anti-ransomware-state disabled
[Job 42] Job succeeded: Successful
```

Create a SnapMirror policy

The SnapMirror policy and the added rule define the retention and the Snapmirror label to identify Snapshots that should be replicated. When creating the SnapCenter policy later, you must use the same label.

```
FsxId05f7f00af49dc7a3e::> snapmirror policy create -policy snapcenter-
policy -tries 8 -transfer-priority normal -ignore-atime false -restart
always -type vault -vserver sapcc-backup-target-zone5
```

```
FsxId05f7f00af49dc7a3e::> snapmirror policy add-rule -vserver sapcc-
backup-target-zone5 -policy snapcenter-policy -snapmirror-label
snapcenter -keep 14
```

```
FsxId00fa9e3c784b6abbb::> snapmirror policy showVserver Policy
Policy Number      Transfer
Name      Name              Type    Of Rules  Tries  Priority  Comment
-----  -
FsxId00fa9e3c784b6abbb
          snapcenter-policy  vault          1      8   normal   -
          SnapMirror Label: snapcenter                                Keep:      14
                                                                Total Keep: 14
```

Create SnapMirror relationship

Now the relation between the source and target volume is defined as well as the type XDP and the policy we created earlier.

```
FsxId05f7f00af49dc7a3e::> snapmirror create -source-path sapcc-hana-
svm:PFX_data_mnt00001 -destination-path sapcc-backup-target-
zone5:PFX_data_mnt00001 -vserver sapcc-backup-target-zone5 -throttle
unlimited -identity-preserve false -type XDP -policy snapcenter-policy
Operation succeeded: snapmirror create for the relationship with
destination "sapcc-backup-target-zone5:PFX_data_mnt00001".
```

Initialize SnapMirror

With this command, the initial replication starts. This is a full transfer of all data from the source volume to the target volume.

```
FsxId05f7f00af49dc7a3e::> snapmirror initialize -destination-path sapcc-
backup-target-zone5:PFX_data_mnt00001 -source-path sapcc-hana-
svm:PFX_data_mnt00001
Operation is queued: snapmirror initialize of destination "sapcc-backup-
target-zone5:PFX_data_mnt00001".
```

You can check the status of the replication with the `snapmirror show` command.

```
FsxId05f7f00af49dc7a3e::> snapmirror show

Progress
Source          Destination Mirror Relationship Total
Last
Path           Type Path           State Status           Progress Healthy
Updated
-----
-----
sapcc-hana-svm:PFX_data_mnt00001
                XDP  sapcc-backup-target-zone5:PFX_data_mnt00001
                Uninitialized
                Transferring  1009MB  true
02/24 12:34:28
```

```
FsxId05f7f00af49dc7a3e::> snapmirror show
```

Progress

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

sapcc-hana-svm:PFX_data_mnt00001	XDP	sapcc-backup-target-zone5:PFX_data_mnt00001	Snapmirrored	Idle	-	true -

Add a backup SVM to SnapCenter

To add a backup SVM to SnapCenter, follow these steps:

1. Configure the SVM where the SnapVault target volume is located in SnapCenter.



2. On the More Options window, select All Flash FAS as the platform and select Secondary.

More Options ✕

Platform All Flash FAS Secondary i

Protocol HTTPS

Port 443

Timeout 60 seconds i

Preferred IP i

Save Cancel

The SVM is now available in SnapCenter.

NetApp SnapCenter®

ONTAP Storage

Type: ONTAP SVMs Search by Name

<input type="checkbox"/>	Name	IP	Cluster Name	User Name	Platform	Controller License
<input type="checkbox"/>	sapcc-backup-target-zone5	10.1.2.31		vsadmin	AFF	Not applicable
<input type="checkbox"/>	sapcc-hana-svm	198.19.255.9		vsadmin	AFF	✓

Create a new SnapCenter policy for backup replication

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

1. Provide a name for the policy.

NetApp SnapCenter®

Global Settings Policies Users and Access Roles Credential Software

SAP HANA

Search by Name

Name	Backup Type	Schedule Type	Replication
BlockIntegrityCheck	File Based Backup	Weekly	
LocalSnap	Data Backup	Hourly	

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

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Provide a policy name

Policy name

Details

3. Select the retention for the Snapshot backups.

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Select backup settings

Backup Type Snapshot Based File-Based i

Schedule Frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

This is the retention for the daily Snapshot backups taken at the primary storage. The retention for secondary backups at the SnapVault target has already been configured previously using the add rule command at the ONTAP level. See “Configure replication relationships on FSx for ONTAP file systems” (xref).

New SAP HANA Backup Policy x

1 Name

2 Settings

3 Retention

4 Replication

5 Summary

Retention settings

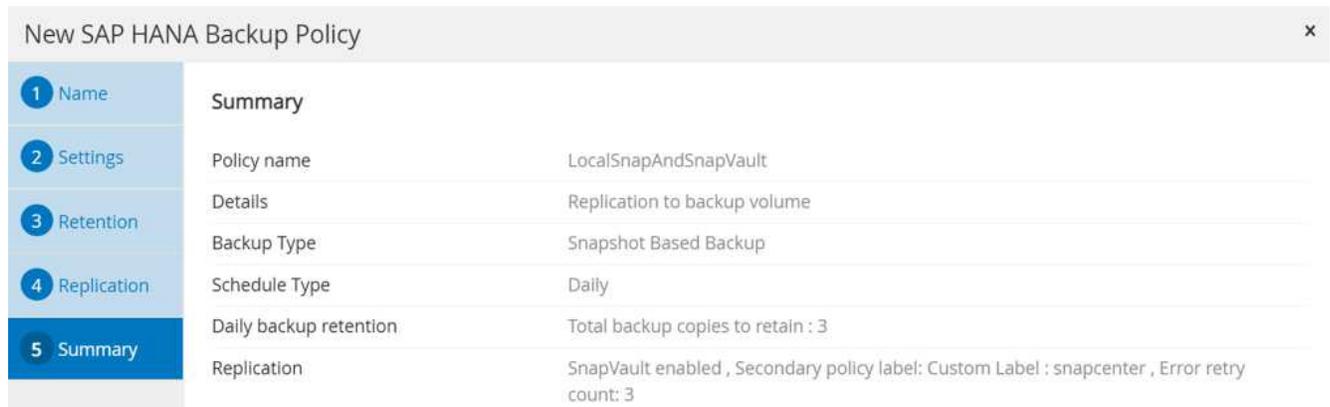
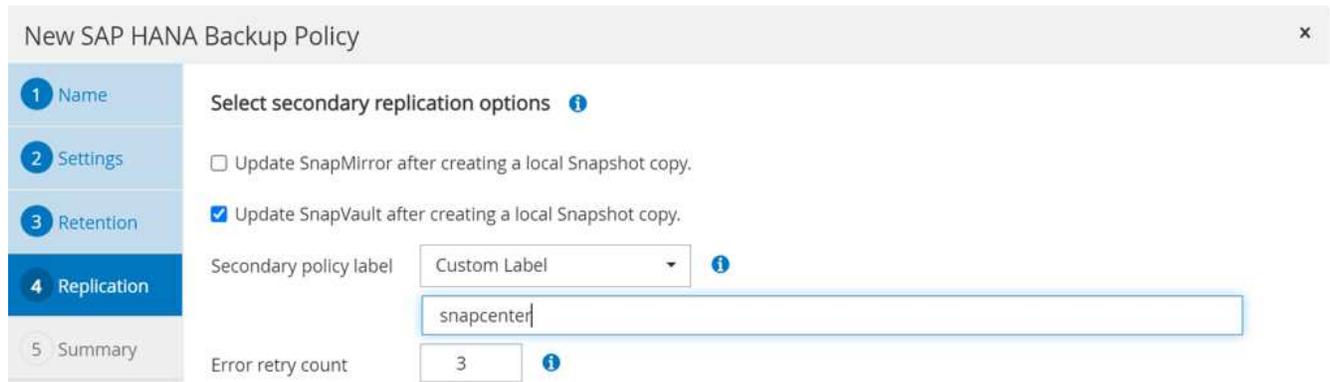
Daily retention settings

Total Snapshot copies to keep

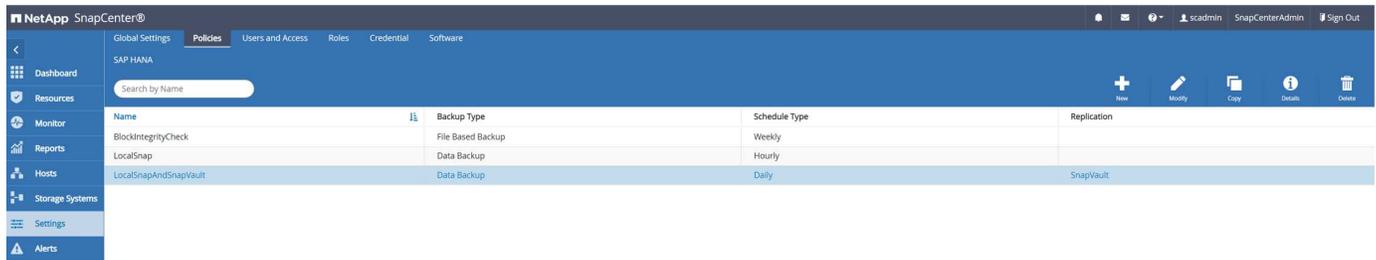
Keep Snapshot copies for days

4. Select the Update SnapVault field and provide a custom label.

This label must match the SnapMirror label provided in the add rule command at ONTAP level.

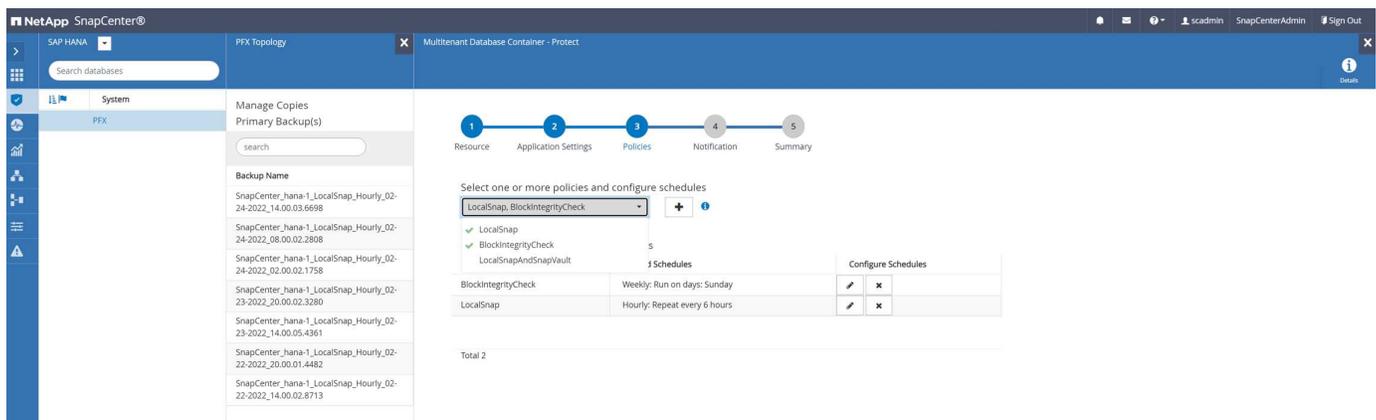


The new SnapCenter policy is now configured.

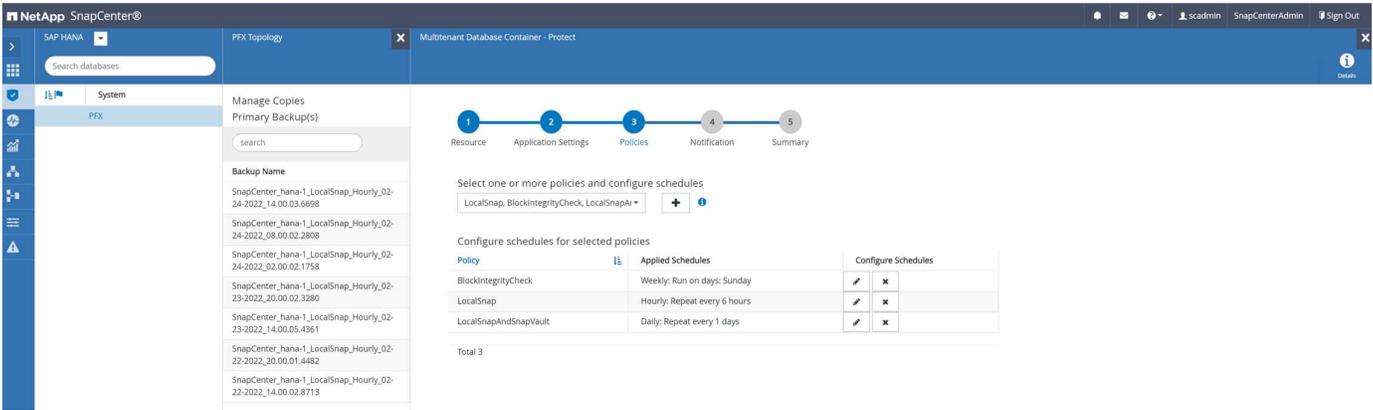


Add a policy to resource protection

You must add the new policy to the HANA resource protection configuration, as shown in the following figure.



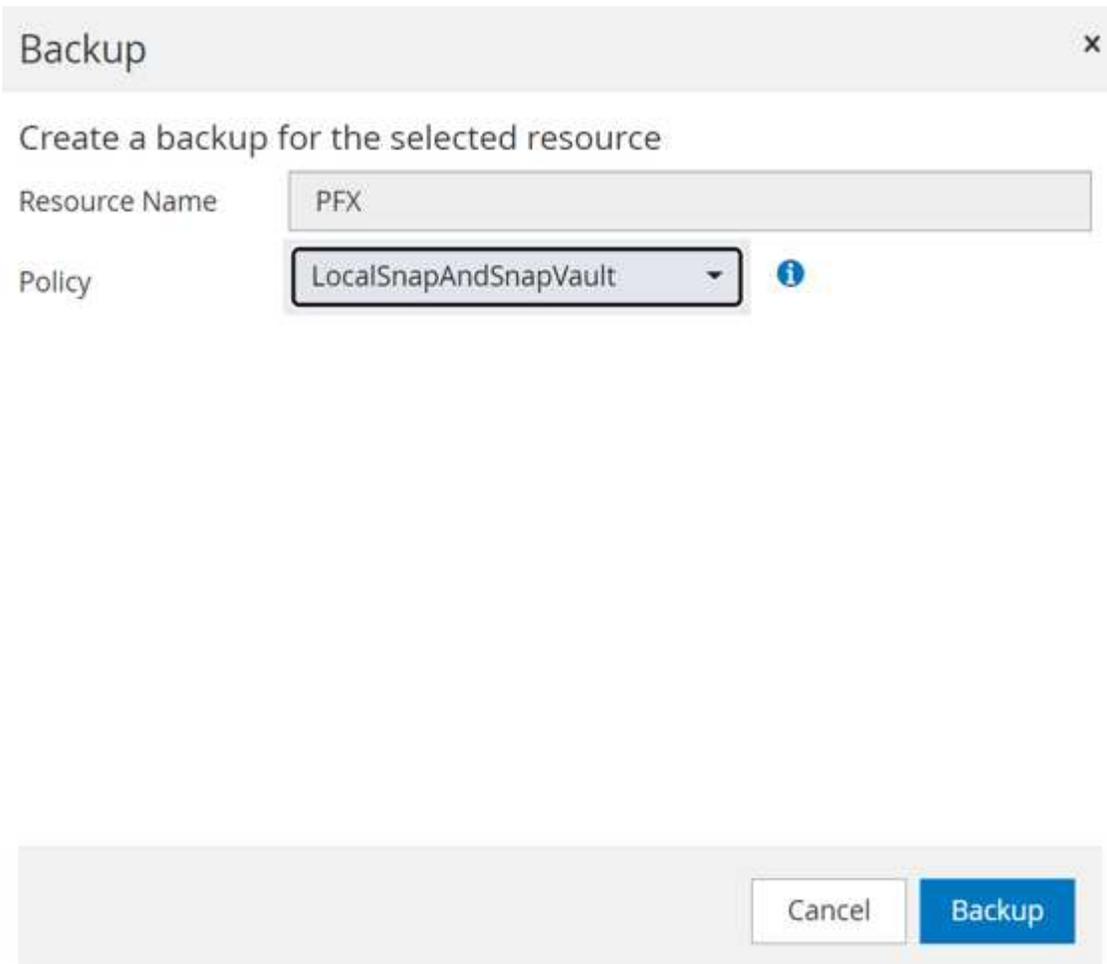
A daily schedule is defined in our setup.



Create a backup with replication

A backup is created in the same way as with a local Snapshot copy.

To create a backup with replication, select the policy that includes the backup replication and click Backup.



Within the SnapCenter job log, you can see the Secondary Update step, which initiates a SnapVault update operation. Replication changed blocks from the source volume to the target volume.

Job Details

Backup of Resource Group 'hana-1_hana_MDC_PFX' with policy 'LocalSnapAndSnapVault'

- Backup of Resource Group 'hana-1_hana_MDC_PFX' with policy 'LocalSnapAndSnapVault'
 - hana-1
 - Backup
 - Validate Dataset Parameters
 - Validate Plugin Parameters
 - Complete Application Discovery
 - Initialize Filesystem Plugin
 - Discover Filesystem Resources
 - Validate Retention Settings
 - Quiesce Application
 - Quiesce Filesystem
 - Create Snapshot
 - UnQuiesce Filesystem
 - UnQuiesce Application
 - Get Snapshot Details
 - Get Filesystem Meta Data
 - Finalize Filesystem Plugin
 - Collect Autosupport data
 - Secondary Update**
 - Register Backup and Apply Retention
 - Register Snapshot attributes
 - Application Clean-Up
 - Data Collection
 - Agent Finalize Workflow
 - (Job 49) SnapVault update

Task Name: Secondary Update Start Time: 02/24/2022 3:14:37 PM End Time: 02/24/2022 3:14:46 PM

View Logs Cancel Job Close

On the FSx for ONTAP file system, a Snapshot on the source volume is created using the SnapMirror label,

snapcenter, as configured in the SnapCenter policy.

```
FsxId00fa9e3c784b6abbb::> snapshot show -vserver sapcc-hana-svm -volume
PFX_data_mnt00001 -fields snapmirror-label
vserver          volume          snapshot
snapmirror-label
-----
-----
-----
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_03-31-
2022_13.10.26.5482 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_03-31-
2022_14.00.05.2023 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-05-
2022_08.00.06.3380 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-05-
2022_14.00.01.6482 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-14-
2022_20.00.05.0316 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-28-
2022_08.00.06.3629 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-1_LocalSnap_Hourly_04-28-
2022_14.00.01.7275 -
sapcc-hana-svm PFX_data_mnt00001 SnapCenter_hana-
1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853

snapcenter
8 entries were displayed.
```

At the target volume, a Snapshot copy with the same name is created.

```
FsxId05f7f00af49dc7a3e::> snapshot show -vserver sapcc-backup-target-zone5
-volume PFX_data_mnt00001 -fields snapmirror-label
vserver          volume          snapshot
snapmirror-label
-----
-----
-----
sapcc-backup-target-zone5 PFX_data_mnt00001 SnapCenter_hana-
1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853 snapcenter
FsxId05f7f00af49dc7a3e::>
```

The new Snapshot backup is also listed in the HANA backup catalog.

Status	Started	Duration	Size	Backup Type	Destination Ty...
	Apr 28, 2022, 4:22:06 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
	Apr 28, 2022, 2:00:26 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
	Apr 28, 2022, 8:00:35 AM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
	Apr 15, 2022, 5:00:44 PM	00h 06m 59s	5.50 GB	Data Backup	Snapshot
	Apr 14, 2022, 8:00:32 PM	00h 00m 16s	5.50 GB	Data Backup	Snapshot
	Apr 5, 2022, 2:00:29 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
	Apr 5, 2022, 8:00:39 AM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
	Mar 31, 2022, 2:00:29 PM	00h 00m 15s	5.50 GB	Data Backup	Snapshot
	Mar 31, 2022, 1:10:57 PM	00h 00m 16s	5.50 GB	Data Backup	Snapshot
	Feb 22, 2022, 12:55:21 PM	00h 00m 21s	3.56 GB	Data Backup	File

ID:	1651162926424
Status:	Successful
Backup Type:	Data Backup
Destination Type:	Snapshot
Started:	Apr 28, 2022, 4:22:06 PM (UTC)
Finished:	Apr 28, 2022, 4:22:21 PM (UTC)
Duration:	00h 00m 15s
Size:	5.50 GB
Throughput:	n.a.
System ID:	
Comment:	SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853
Additional Information:	<ok>
Location:	/hana/data/PFX/mnt00001/

Host	Service	Size	Name	Source Type	EBID
hana-1	nameserver	5.50 GB	hdb00001	volume	SnapCent...

In SnapCenter, you can list the replicated backups by clicking Vault Copies in the topology view.

The screenshot shows the NetApp SnapCenter interface. The 'Manage Copies' section displays a topology diagram with 'Local copies' (8 Backups, 0 Clones) and 'Vault copies' (1 Backup, 0 Clones). Below this, the 'Secondary Vault Backup(s)' section shows a table with one entry:

Backup Name	Count	IF	End Date
SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853	1		04/28/2022 4:22:40 PM

Restore and recover from secondary storage

To restore and recover from secondary storage, follow these steps:

To retrieve the list of all the backups on the secondary storage, in the SnapCenter Topology view, click Vault Copies, then select a backup and click Restore.

This screenshot is similar to the previous one, but the 'Restore' button is now visible in the bottom right corner of the 'Secondary Vault Backup(s)' table, indicating the user has selected a backup to restore.

The restore dialog shows the secondary locations.

Restore from SnapCenter_hana-1_LocalSnapAndSnapVault_Daily_04-28-2022_16.21.41.5853 ×

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

Select the restore types

Complete Resource ?

Tenant Database

Choose archive location

sapcc-hana-svm:PFX_data_mnt00001 sapcc-backup-target-zone5:PFX_data_mnt00

⚠ The newer tenants added on the host after the backup was created cannot be restored and will be lost after restore operation. ×

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

Previous Next

Further restore and recovery steps are identical to those previously covered for a Snapshot backup at the primary storage.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- FSx for NetApp ONTAP user guide — What is Amazon FSx for NetApp ONTAP?

<https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/what-is-fsx-ontap.html>

- SnapCenter resources page

<https://www.netapp.com/us/documentation/snapcenter-software.aspx>

- SnapCenter Software documentation

<https://docs.netapp.com/us-en/snapcenter/index.html>

- Automating SAP HANA System Copy and Clone Operations with SnapCenter

[Automating SAP HANA System Copy and Clone Operations with SnapCenter](#)

- SAP HANA System Replication — Backup and Recovery with SnapCenter

[Backup and Recovery with SnapCenter](#)

Version history

Version	Date	Document version history
Version 1.0	May 2022	Initial release.

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.