



SAP HANA Plug-in Operations

Snap Creator Framework

NetApp

January 20, 2026

This PDF was generated from https://docs.netapp.com/us-en/snap-creator-framework/sap-hana-ops/concept_considerations_for_backing_up_sap_hana_systems.html on January 20, 2026. Always check docs.netapp.com for the latest.

Table of Contents

SAP HANA Plug-in Operations Guide	1
SAP HANA backup and restore solution overview	1
Considerations for backing up SAP HANA systems	1
The NetApp solution	2
Backup solution components	3
SAP HANA plug-in overview	4
Requirements	6
Installing and configuring required software components	6
Setup assumptions in this guide	7
Setup used with clustered Data ONTAP	8
Configuring data backups	9
Configuring the backup user and hdbuserstore	9
Configuring SnapVault relationships	11
Starting the SnapVault relationships	12
Configuring the Snap Creator Framework and SAP HANA database backup	14
Configuring SAP HANA for SAN environments	30
Configuring log backups	31
Housekeeping of log backups	32
Modifying the housekeeping of log backups	32
Executing database backups	33
Overview of database backups	33
Backing up the database with the Snap Creator GUI	33
Backing up the database with Snap Creator command line	35
Reviewing available backups in SAP HANA Studio	36
SAP HANA File-Based Backup and Database Integrity Checks	38
Modifying configuration for File-Based Backup	38
Modifying configuration for Database Integrity Checks	38
Scheduling file-based backup	39
Scheduling database integrity checks	40
Performing File-Based Backup from the Snap Creator GUI	40
Performing File-Based Backup from Snap Creator command line	41
Performing Database Integrity Checks from Snap Creator GUI	41
Performing Database Integrity Checks from Snap Creator command line	42
Restoring and recovering SAP HANA databases	42
Restoring and recovering databases from primary storage	43
Restoring and recovering databases from secondary storage	56
Resuming a SnapVault relationship after a restore	66
Restoring databases after primary storage failure	69
SAP HANA plug-in parameters	71
Troubleshooting	73
Where to go next	75

SAP HANA Plug-in Operations Guide

You can configure and use the SAP HANA plug-in for Snap Creator 4.3.3 to back up and restore SAP HANA databases.

SAP HANA backup and restore solution overview

Corporations today require their SAP applications to be available 24 hours a day, seven days a week. Consistent levels of performance are expected regardless of increasing data volumes and routine maintenance tasks such as system backups.

Running SAP database backups can have a significant performance effect on a production SAP system. Because backup windows are shrinking and the amount of data that needs to be backed up is increasing, it is difficult to define a point in time when backups can be performed with minimal effect on business processes. The time needed to restore and recover SAP systems is of particular concern because the downtime must be minimized.

Considerations for backing up SAP HANA systems

SAP HANA administrators must deliver a reliable level of service, minimizing downtime or performance degradation due to backups.

To deliver this level of service, SAP HANA administrators contend with challenges in the following areas:

- Performance effect on production SAP systems

Backups typically have a significant performance impact on the production SAP system because there is a heavy load on the database server, the storage system, and the storage network during backups.

- Shrinking backup windows

Backups can be created only during times with low I/O or batch activities occurring on the SAP system. It is very difficult to define a backup window when the SAP system is active all the time.

- Rapid data growth

Rapid data growth together with shrinking backup windows result in ongoing investments in the backup infrastructure: more tape drives, new tape drive technology, faster storage networks. Growing databases also result in more tape media or disk space for backups. Incremental backups can address these issues, but result in a very slow restore process, which is usually not acceptable.

- Increasing cost of downtime

Unplanned downtime of an SAP system always has a financial effect on the business. A significant part of the unplanned downtime is the time that is required to restore and recover the SAP system in case of a failure. The backup and recovery architecture must be designed based on an acceptable recovery time objective (RTO).

- Backup and recovery time

Backup and recovery time are included in SAP upgrade projects. The project plan for a SAP upgrade always includes at least three backups of the SAP database. The time required to perform these backups

reduces the total available time for the upgrade process. The decision whether to backup and recover is generally based on the amount of time required to restore and recover the database from the backup that was created previously. The option to restore very quickly provides more time to solve problems that might occur during the upgrade rather than just restore the system back to its previous state.

The NetApp solution

A database backup can be created in minutes by using NetApp Snapshot technology. The time needed to create a Snapshot copy is independent of the size of the database because a Snapshot copy does not move any data blocks.


The use of Snapshot technology also has no performance effect on the production SAP system. Therefore, the creation of Snapshot copies can be scheduled without having to consider peak activity periods. SAP and NetApp customers typically schedule several online Snapshot backups during the day. For example, backups might occur every four hours. These Snapshot backups are typically kept for three to five days on the primary storage system.

Snapshot copies also provide key advantages for the restore and recovery operation. NetApp SnapRestore functionality allows restoring the entire database or parts of the database to the point in time when any available Snapshot copy was created. This restore process is done in a few minutes, independently of the size of the database. The time needed for the recovery process is also dramatically reduced, because several Snapshot copies have been created during the day, and fewer logs need to be applied.

Snapshot backups are stored on the same disk system as the active online data. Therefore NetApp recommends using Snapshot backups as a supplement, not a replacement for backups to a secondary location such as disk or tape. Although backups to a secondary location are still necessary, there is only a slight probability that these backups will be needed for restore and recovery. Most restore and recovery actions are handled by using SnapRestore on the primary storage system. Restores from a secondary location are only necessary if the primary storage system holding the Snapshot copies is damaged or if it is necessary to restore a backup that is no longer available from a Snapshot copy. For example, you might need to restore a backup from two weeks ago.

A backup to a secondary location is always based on Snapshot copies created on the primary storage. Therefore, the data is read directly from the primary storage system without generating load on the SAP database server. The primary storage communicates directly with the secondary storage and sends the backup data to the destination using the SnapVault disk-to-disk backup. The NetApp SnapVault functionality offers significant advantages compared to traditional backups. After an initial data transfer, in which all the data has to be transferred from the source to the destination, all subsequent backups copy only the changed blocks to the secondary storage. This significantly reduces the load on the primary storage system and the time needed for a full backup. A full database backup requires less disk space because SnapVault stores only the changed blocks at the destination.

Backing up data to tape as a long-term backup might still be required. This could be, for example, a weekly backup that is kept for a year. In this case, the tape infrastructure can be directly connected to the secondary storage, and the data could be written to tape by using the Network Data Management Protocol (NDMP).




Backup solution components

The Snap Creator backup solution for SAP HANA consists of SAP HANA data file backup using storage-based Snapshot copies, replication of data file backups to a secondary offsite backup location, SAP HANA log file backup using the HANA database log backup functionality, database block integrity check using a file-based backup, and housekeeping of data file, log file backups, and the SAP HANA backup catalog.

Database backups are executed by Snap Creator in conjunction with a plug-in for SAP HANA. The plug-in ensures database consistency so that the Snapshot copies that are created on the primary storage system are based on a consistent image of the SAP HANA database.

Snap Creator allows you to replicate the consistent database images to a secondary storage using SnapVault. Typically, there will be different retention policies defined for the backups at the primary storage and the backups at the secondary storage. Snap Creator handles the retention at the primary storage as well as the secondary storage.

The log backup is executed automatically by the SAP HANA database tools. The log backup destination should not be on the same storage system where the log volume of the database is located. Configuring the log backup destination on the same secondary storage where the database backups get replicated with SnapVault is recommended. With this configuration, the secondary storage has similar availability requirements as the primary storage so that it is certain that the log backups can always be written to the secondary storage.



The backup schedules and retention policies must be defined based on customer requirements. The following table shows an example configuration of the different schedules and retention policies.

	Executed by Snap Creator	Primary storage	Secondary storage
Database backups	Schedule 1: every 4 hours	Retention: 6 (=> 6 hourly Snapshot copies)	Retention: 6 (=> 6 hourly Snapshot copies)
Schedule 2: once per day	Retention: 3 (=> 3 daily Snapshot copies)	Retention: 28 (4 weeks) (=> 28 daily Snapshot copies)	Log backups
SAP HANA database tools schedule: every 15 minutes	NA	Retention: 28 days (4 weeks)	Block integrity check

With this example, six hourly and three daily backups are kept at the primary storage. At the secondary storage, the database backups are kept for four weeks. To be able to recover any of the data backups, you must set the same retention for the log backups.

SAP HANA plug-in overview

The SAP HANA plug-in works in conjunction with the Snap Creator Framework to provide a backup solution for SAP HANA databases that rely on a NetApp storage back end. The Snapshot backups created by Snap Creator are registered in the HANA Catalog and are visible in HANA Studio.

Snap Creator Framework supports two types of SAP HANA databases: single containers and multitenant database containers (MDC) single tenant database.

Snap Creator and the SAP HANA plug-in are supported with Data ONTAP operating in 7-Mode and clustered Data ONTAP with the SAP HANA database nodes attached to the storage controllers using either NFS or Fibre Channel. The required interfaces to the SAP HANA database are available for Service Pack Stack (SPS) 7 and later.


The Snap Creator Framework communicates with the storage systems to create Snapshot copies and to replicate the data to a secondary storage using SnapVault. Snap Creator is also used to restore the data either with SnapRestore at the primary storage or with SnapVault restore from the secondary storage.

The Snap Creator plug-in for SAP HANA uses the SAP HANA hdbsql client to execute SQL commands in order to provide database consistency and to manage the SAP HANA backup catalog. The SAP HANA plug-in is supported for both SAP Certified Hardware Appliances and Tailored Datacenter Integration (TDI) programs.

The Snap Creator plug-in for SAP HANA uses the SAP HANA hdbsql client to execute SQL commands for the following tasks:

- Provide database consistency to prepare a storage-based Snapshot backup
- Manage log file backup retention on file system level
- Manage the SAP HANA backup catalog for data file and log file backups
- Execute a file-based backup for block integrity check

The following illustration shows an overview of the communication paths of Snap Creator with the storage and the SAP HANA database.



Snap Creator performs the following steps to back up the database:

1. Creates an SAP HANA database Snapshot copy to obtain a consistent image on the persistence layer.
2. Creates a storage Snapshot copy of the data volume(s).
3. Registers the storage Snapshot backup within the SAP HANA backup catalog.
4. Deletes the SAP HANA Snapshot copy.
5. Executes a SnapVault update for the data volume.
6. Deletes the storage Snapshot copies at the primary and/or secondary storage, based on the defined retention policies for backups at the primary and secondary storage.
7. Deletes the SAP HANA backup catalog entries if the backups do not exist anymore at the primary and the secondary storage.
8. Deletes all log backups that are older than the oldest data backup on the file system and within the SAP HANA backup catalog.

Requirements

The SAP HANA plug-in enables you to create backups and perform point-in-time recovery of HANA databases.

Support for the SAP HANA plug-in is as follows:

- Host operating system: SUSE Linux Enterprise Server (SLES), 32-bit and 64-bit
- Clustered Data ONTAP or Data ONTAP operating in 7-Mode
- At least one SAP HANA database node attached via NFS
- SAP HANA running Service Pack Stack (SPS) 7 or later



For the latest information about support or to view compatibility matrices, see the [NetApp Interoperability Matrix Tool](#).

Required licenses

The primary storage controllers must have a SnapRestore and SnapVault license installed. The secondary storage must have a SnapVault license installed.

No license is required for Snap Creator and the Snap Creator SAP HANA plug-in.

Capacity requirements for Snapshot backups

A higher block change rate on the storage layer has to be considered compared to the change rate with traditional databases. Due to the table merge process of the column store, much more data than just the block changes is written to disk. Until more customer data is available, the current estimation for the change rate is 20% to 50% per day.

Installing and configuring required software components

For the SAP HANA backup and restore solution using the Snap Creator Framework and the SAP HANA plug-in, you need to install Snap Creator software components and the SAP HANA hdbsql client software.

You do not need to install the plug-in separately. It is installed with the Agent.

1. Install the Snap Creator Server on a host that shares network connectivity with the host where you install the Agent.
2. Install the Snap Creator Agent on a host that shares network connectivity with the Snap Creator Server host.
 - In a single SAP HANA node environment, install the Agent on the database host. Alternately, install the Agent on another host that has network connectivity to the database host and the Snap Creator Server host.
 - In a multinode SAP HANA environment, you should not install the Agent on the database host; the Agent needs to be installed on a separate host that has network connectivity to the database host and the Snap Creator Server host.
3. Install the SAP HANA hdbsql client software on the host where you installed the Snap Creator Agent.

Configure the user store keys for the SAP HANA nodes that you manage through this host.

```
mgmtsrv01:/sapcd/HANA_SP5/DATA_UNITS/HDB_CLIENT_LINUXINTEL # ./hdbinst

SAP HANA Database Client installation kit detected.

SAP HANA Database Installation Manager - Client Installation
1.00.46.371989
*****
***

Enter Installation Path [/usr/sap/hdbcclient32]:
Checking installation...
Installing and configuring required software components | 13
Preparing package "Product Manifest"...
Preparing package "SQLDBC"...
Preparing package "ODBC"...
Preparing package "JDBC"...
Preparing package "Client Installer"...
Installing SAP HANA Database Client to /usr/sap/hdbcclient32...
Installing package 'Product Manifest' ...
Installing package 'SQLDBC' ...
Installing package 'ODBC' ...
Installing package 'JDBC' ...
Installing package 'Client Installer' ...
Installation done
Log file written to '/var/tmp/hdb_client_2013-07-
05_11.38.17/hdbinst_client.log'
mgmtsrv01:/sapcd/HANA_SP5/DATA_UNITS/HDB_CLIENT_LINUXINTEL #
```

Related information

[Snap Creator Framework Installation Guide](#)

Setup assumptions in this guide

Though a typical Snap Creator installation assumes that the Server is installed on one host and the Agent is installed on a different host, the setup used in this guide is based on an SAP HANA multinode appliance.


In this configuration, the SAP HANA database runs on a 3+1 database node configuration and all Snap Creator software components—Server, Agent, and plug-in—are installed on the same host.

The NetApp storage systems used in this setup are running Data ONTAP operating in 7-Mode. One high-availability (HA) controller pair is used on the storage layer. The data and log volumes of the three SAP HANA database nodes are distributed to both storage controllers. With the example setup, one storage controller of another HA controller pair is used as the secondary storage. Each data volume is replicated to a dedicated

backup volume on the secondary storage. The size of the backup volumes depend on the number of backups that will be kept at the secondary storage.

All Snap Creator and SAP HANA Studio operations described here are the same with storage systems running clustered Data ONTAP. However, the initial SnapVault configuration on the storage systems and all SnapVault commands that need to be executed directly on the storage are different with clustered Data ONTAP. The differences are highlighted and described in this guide.


The following figure shows the data volumes on the primary storage and the replication path to the secondary storage:



i All volumes that need to be backed up must be created on the secondary storage controller. In this example, the volumes backup_data_00001, backup_data_00002, and backup_data_00003 are created on the secondary storage controller.

Setup used with clustered Data ONTAP

The following figure shows the setup that has been used with clustered Data ONTAP. The setup is based on a single-node SAP HANA configuration with the storage virtual machines (SVMs) and volume names shown in the following illustration.



The way you prepare, start, resume, and restore SnapVault operations is different in clustered Data ONTAP and Data ONTAP operating in 7-Mode. These differences are called out in the corresponding sections of this guide.

Configuring data backups

After you install the required software components, follow these steps to complete the configuration:


1. Configure a dedicated database user and the SAP HANA userstore.
2. Prepare SnapVault replication on all storage controllers.
3. Create volumes at secondary storage controller.
4. Initialize the SnapVault relationships for database volumes.
5. Configure Snap Creator.

Configuring the backup user and hdbuserstore

You should configure a dedicated database user within the HANA database to run the backup operations with Snap Creator. In a second step, you should configure a SAP HANA userstore key for this backup user. This userstore key is used within the configuration of the Snap Creator SAP HANA plug-in.

The backup user must have the following privileges:

- BACKUP ADMIN
- CATALOG READ



1. At the administration host, the host where Snap Creator got installed, a userstore key is configured for all database hosts that belong to the SAP HANA database. The userstore key is configured with the OS root user: `hdbuserstore set keyhost 3[instance]15 userpassword`
2. Configure a key for all four database nodes.

```

mgmtsrv01:/usr/sap/hdbclient32 # ./hdbuserstore set SCADMIN08
cishanar08:34215 SCADMIN Password
mgmtsrv01:/usr/sap/hdbclient32 # ./hdbuserstore set SCADMIN09
cishanar09:34215 SCADMIN Password
mgmtsrv01:/usr/sap/hdbclient32 # ./hdbuserstore set SCADMIN10
cishanar10:34215 SCADMIN password
mgmtsrv01:/usr/sap/hdbclient32 # ./hdbuserstore set SCADMIN11
cishanar11:34215 SCADMIN Password
mgmtsrv01:/usr/sap/hdbclient32 # ./hdbuserstore LIST
DATA FILE      : /root/.hdb/mgmtsrv01/SSFS_HDB.DAT

KEY SCADMIN08
  ENV : cishanar08:34215
  USER: SCADMIN
KEY SCADMIN09
  ENV : cishanar09:34215
  USER: SCADMIN
KEY SCADMIN10
  ENV : cishanar10:34215
  USER: SCADMIN
KEY SCADMIN11
  ENV : cishanar11:34215
  USER: SCADMIN
mgmtsrv01:/usr/sap/hdbclient32

```

Configuring SnapVault relationships

When you configure SnapVault relationships, the primary storage controllers must have a valid SnapRestore and SnapVault license installed. The secondary storage must have a valid SnapVault license installed.

1. Enable SnapVault and NDMP on the primary and the secondary storage controllers.

```

hana1a> options snapvault.enable on
hana1a> ndmp on
hana1a>
hana1b> options snapvault.enable on
hana1b> ndmpd on
hana1b
hana2b> options snapvault.enable on
hana2b> ndmpd on
hana2b>

```

2. On all primary storage controllers, configure the access to the secondary storage controller.

```
hana1a> options snapvault.access host=hana2b
hana1a>
hana1b> options snapvault.access host=hana2b
hana1b>
```



Using a dedicated network for replication traffic is recommended. In such cases, the host name of this interface at the secondary storage controller needs to be configured. Instead of hana2b, the host name could be hana2brep.

3. On the secondary storage controller, configure the access for all primary storage controllers.

```
hana2b> options snapvault.access host=hana1a,hana1b
hana2b>
```



Using a dedicated network for replication traffic is recommended. In such cases, the host name of this interface at the primary storage controllers needs to be configured. Instead of hana1b and hana1a the host name could be hana1arep and hana1brep.

Starting the SnapVault relationships

You need to start the SnapVault relationship with Data ONTAP operating in 7-Mode and clustered Data ONTAP.

Starting the SnapVault relationships with Data ONTAP operating in 7-Mode

You can start a SnapVault relationship with commands executed on the secondary storage system.

1. For storage systems running Data ONTAP operating in 7-Mode, you start the SnapVault relationships by running the following command:

```
hana2b> snapvault start -S hanala:/vol/data_00001/mnt00001
/vol/backup_data_00001/mnt00001
Snapvault configuration for the qtree has been set.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
hana2b>
hana2b> snapvault start -S hanala:/vol/data_00003/mnt00003
/vol/backup_data_00003/mnt00003
Snapvault configuration for the qtree has been set.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
hana2b>
hana2b> snapvault start -S hanalb:/vol/data_00002/mnt00002
/vol/backup_data_00002/mnt00002
Snapvault configuration for the qtree has been set.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
hana2b>
```



It is recommended that you use a dedicated network for replication traffic. In that case, configure the host name of this interface at the primary storage controllers. Instead of hana1b and hana1a, the host name could be hana1arep and hana1brep.

Starting the SnapVault relationships with clustered Data ONTAP

You need to define a SnapMirror policy before you start a SnapVault relationship.

1. For storage systems running clustered Data ONTAP, you start the SnapVault relationships by running the following command.

```

hana::> snapmirror policy create -vserver hana2b -policy SV_HANA
hana::> snapmirror policy add-rule -vserver hana2b -policy SV_HANA
-snapmirror-label daily -keep 20
hana::> snapmirror policy add-rule -vserver hana2b -policy SV_HANA
-snapmirror-label hourly -keep 10

hana::> snapmirror policy show -vserver hana2b -policy SV_HANA

        Vserver: hana2b
        SnapMirror Policy Name: SV_HANA
        Policy Owner: vserver-admin
        Tries Limit: 8
        Transfer Priority: normal
        Ignore accesstime Enabled: false
        Transfer Restartability: always
        Comment: -
        Total Number of Rules: 2
        Total Keep: 8
        Rules: Snapmirror-label   Keep  Preserve  Warn
        -----
        daily           20   false      0
        hourly          10   false      0

```

The policy must contain rules for all retention classes (labels) that are used in the Snap Creator configuration. The above commands show how to create a dedicated SnapMirror policy SV_HANA

2. To create and start the SnapVault relationship on the cluster console of the backup cluster, run the following commands.

```

hana::> snapmirror create -source-path hanala:hana_data -destination
-path
hana2b:backup_hana_data -type XDP -policy SV_HANA
Operation succeeded: snapmirror create the relationship with destination
hana2b:backup_hana_data.


hana::> snapmirror initialize -destination-path hana2b:backup_hana_data
-type XDP

```

Configuring the Snap Creator Framework and SAP HANA database backup

You must configure the Snap Creator Framework and the SAP HANA database backup.

1. Connect to the Snap Creator graphical user interface (GUI): <https://host:8443/ui/>.
2. Log in using the user name and password that were configured during the installation. Click **Sign in**.




3. Enter a profile name and click **OK**.



For example, “ANA” is the SID of the database.


4. Enter the configuration name, and click **Next**.



5. Select **Application plug-in** as the plug-in type, and click **Next**.



6. Select **SAP HANA** as the application plug-in, and click **Next**.



7. Enter the following configuration details:

- Select **Yes** from the drop-down menu to use the configuration with a multitenant database. For a single container database select **No**.
- If Multitenant Database Container is set to **No**, you must provide the database SID.
- If Multitenant Database Container is set to **Yes**, you must add the hdbuserstore keys for each SAP

HANA node.

- d. Add the name of the tenant database.
- e. Add the HANA nodes on which the hdbsql statement must be executed.
- f. Enter the HANA node instance number.
- g. Provide the path to the hdbsql executable file.
- h. Add the OSDB user.
- i. Select **Yes** from the drop-down list to Enable LOG Cleanup.


NOTE:

- Parameter `HANA_SID` is available only if the value for parameter `HANA_MULTITENANT_DATABASE` is set to `N`
- For multitenant database containers (MDC) with a “Single Tenant” resource type, the SAP HANA Snapshot copies work with UserStore Key based authentication. If the `HANA_MULTITENANT_DATABASE` parameter is set to `Y`, then the `HANA_USERSTORE_KEYS` parameter must be set to the appropriate value.
- Similar to non-multitenant database containers, the file-based backup and integrity check feature is supported


- j. Click **Next**.

Multitenant Database Container (MDC) - Single Tenant:	
SID:	<input type="text" value="H66"/>
hdbuserstore Keys:	<input type="text"/>
Tenant Database Name:	<input type="text"/>
Nodes:	<input type="text" value="10.235.220.66"/>
Username:	<input type="text" value="SYSTEM"/>
Password:	<input type="password" value="*****"/>
Instance number:	<input type="text" value="66"/>
Path to hdbsql:	<input type="text" value="/usr/sap/H66/HDB66/exe/hdbsql"/>
OSDB User:	<input type="text"/>
Enable LOG Cleanup:	<input type="text" value="Yes"/>

8. Enable the File-Based Backup operation:
 - a. Set the File-Backup Location.
 - b. Specify the file-backup prefix.
 - c. Select the **Enable File-Backup** checkbox.
 - d. Click **Next**.



9. Enable the Database Integrity Check operation:
 - a. Set the temporary File-Backup location.
 - b. Select the **Enable DB Integrity Check** checkbox.
 - c. Click **Next**.



10. Enter the details for the agent configuration parameter, and click **Next**.

IP/DNS:	localhost
Port:	9090
Timeout (secs):	300

11. Enter the storage connection settings, and click **Next**.

Storage Connection Settings
Please Provide Storage Connection Settings

Use OnCommand Proxy:

Transport:

Controller/Vserver Port:

12. Enter the storage login credentials, and click **Next**.

Controller/Vserver Credentials
Add one or more Controller/Vserver credentials to the configuration.

Controller/Vserver Login Credentials

Controller/Vserver IP or Name	User name/Password	Volumes
New Controller/Vserver		
Controller/Vserver IP or Name:	<input type="text" value="hana1a"/>	
Controller/Vserver User:	<input type="text" value="root"/>	
Controller/Vserver Password:	<input type="password" value="*****"/>	
<input type="button" value="Next"/>		


13. Select the data volumes that are stored on this storage controller, and click **Save**.

Controller/Vserver Volumes	
SLES	
SLES4SAP	
SLES4SAP_K3074	
cishanar08_3080	
cishanar08_30807	
cishanar08_PTF	
cishanar08_SLES4SAP	
cishanar09	
cishanar09_3080	
cishanar09_PTF	
cishanar09_SLES4SAP	
cishanar10	
cishanar10_3080	
cishanar10_PTF	
cishanar10_SLES4SAP	
cishanar11	
cishanar11_3080	
cishanar11_PTF	
cishanar11_SLES4SAP	
log_00002	
log_00004	
osmaster	
osmaster_30807	
osmaster_PTF_S745	
osmaster_PTF_S819	
sapcd	
tftpboot	
vol0	
<input type="button" value="Save"/>	


14. Click **Add** to add another storage controller.

Controller/Vserver Credentials		
Add one or more Controller/Vserver credentials to the configuration.		
Controller/Vserver Login Credentials		
<input type="button" value="Add"/>	<input type="button" value="Edit"/>	<input type="button" value="Delete"/>
Controller/Vserver IP or Name	User name/Password	Volumes
hanaf1a	root/***	data_00001 data_00003

15. Enter the storage login credentials, and click **Next**.



16. Select the data volumes that are stored on the second storage controller that you created, and click **Save**.



17. The Controller/Vserver Credentials window displays the storage controllers and volumes that you added.

Click Next.

Controller/Vserver Credentials
Add one or more Controller/Vserver credentials to the configuration.

Controller/Vserver Login Credentials

Controller/Vserver IP or Name	User name/Password	Volumes
hana1a	root/****	data_00001 data_00003
hana1b	root/****	data_00002

18. Enter the Snapshot policy and retention configuration.

The retention of three daily and eight hourly Snapshot copies is just an example and could be configured differently depending on the customer requirements.



Select **Timestamp** as the naming convention. The use of the naming convention **Recent** is not supported with the SAP HANA plug-in, because the timestamp of the Snapshot copy is also used for the SAP HANA backup catalog entries.

Configuration

Snapshot Details
Provide Snapshot copy related information.

Snapshot copy Name:

Snapshot copy Label:

Policy Type: Use Policy Use Policy Object

Snapshot copy Policies

Enable Policy	Policy Name	Retention
<input checked="" type="checkbox"/>	hourly	12
<input checked="" type="checkbox"/>	daily	3
<input type="checkbox"/>	weekly	0
<input type="checkbox"/>	monthly	0

Prevent Snapshot copy Deletion:

Policy Retention Age:

Naming Convention: Recent Timestamp

19. No changes required. Click **Next**.

Snapshot Details Continued
Provide Snapshot copy related information.

Consistency Group:	<input type="checkbox"/>
Consistency Timeout:	MEDIUM
SnapDrive Discovery:	No
Consistency Group WAFL Sync:	No
Snapshot copy Delete by age only:	N
Snapshot copy Dependency Ignore:	No
Restore Auto Detect:	No
Ignore Application Errors:	No
Snapshot Copy Disable:	No

20. Select **SnapVault**, and configure the SnapVault retention policies and the SnapVault wait time.

Configuration

Data Protection
Configure SnapMirror, SnapVault or both

Data Transfer: SnapMirror SnapVault

SnapVault Policies		
Enable Policy	Policy Name	Retention
<input checked="" type="checkbox"/>	hourly	10
<input checked="" type="checkbox"/>	daily	20
<input type="checkbox"/>	weekly	0
<input type="checkbox"/>	monthly	0

SnapVault Retention Age:

SnapVault wait time:

21. Click **Add**.

Data Protection Volumes
Add SnapMirror and SnapVault Volumes.

Data Protection Volumes		
Add	Edit	Delete
Controller/Vserver IP or Name	SnapMirror Volumes	SnapVault Volumes

22. Select a source storage controller from the list, and click **Next**.

Data Protection Volumes
Add SnapMirror and SnapVault Volumes.

Data Protection Volumes

Add Edit Delete


Controller/Vserver IP or Name	SnapMirror Volumes	SnapVault Volumes

Select a Controller/Vserver


Controller/Vserver IP or Name:

Next


23. Select all the volumes that are stored on the source storage controller, and click **Save**.



24. Click **Add**, and select the second source storage controller from the list, and then click **Next**.




25. Select all the volumes that are stored on the second source storage controller, and click **Save**.




26. The Data Protection Volumes window displays all the volumes that should be protected in the configuration that you created. Click **Next**.

Controller/Server IP or Name	SnapMirror Volumes	SnapVault Volumes
hans1a	data_00001 data_00003	data_00002
hans1b		data_00002


27. Enter the credentials for the target storage controllers, and click **Next**. In this example, the “root” user credentials are used to access the storage system. Typically, a dedicated backup user is configured on the storage system and is then used with Snap Creator.



28. Click **Next**.



29. Click **Finish** to complete the configuration.



30. Click the **SnapVault settings** tab.
31. Select **Yes** from the drop-down list of the **SnapVault Restore Wait** option, and click **Save**.

It is recommended that you use a dedicated network for replication traffic. If you decide to do so, you should include this interface in the Snap Creator configuration file as a secondary interface.

You can also configure dedicated management interfaces so that Snap Creator can access the source or the target storage system by using a network interface that is not bound to the storage controller's host name.

```
mgmtsrv01:/opt/NetApp/Snap_Creator_Framework_411/scServer4.1.1c/engine/configs/HANA_profile_ANA
# vi ANA_database_backup.conf

#####
# Connection Options
#####
#PORT=443
SECONDARY_INTERFACES=hanal1:hanal1-rep/hana2b;hanal2:hanal2-rep/hana2b
MANAGEMENT_INTERFACES=hana2b:hana2b-mgmt
```

Configuring SAP HANA for SAN environments

After you configure the data backups, you will need to add a new command to the Snap

Creator configuration file in environments where a SAP HANA system is connected using Fibre Channel storage area network (SAN) to the storage controller(s).

When a global synchronized backup savepoint is triggered by Snap Creator within SAP HANA, the last step occurs when SAP HANA writes the /hana/data/SID/mnt00001/hdb00001/snapshot_databackup_0_1 file. This file is part of the data volume on the storage and is therefore part of the storage Snapshot copy. This file is mandatory when performing a recovery in case the backup is restored. Due to metadata caching with the 'X' File System (XFS) on the Linux host, the file is not immediately visible at the storage layer. The standard XFS configuration for metadata caching is 30 seconds.

Within Snap Creator, you need to add a post-application quiesce command, which waits until the XFS metadata cache is flushed to the disk layer.

You can check the configuration of the metadata caching by using the following command:

```
stlrx300s8-2:/ # sysctl -A | grep xfssyncd_centisecs
fs.xfs.xfssyncd_centisecs = 3000
```

1. In the configuration file (install_path/scServerversion_number/engine/configs), add the /bin/sleep command to the Post Commands section as shown in the following example:

```
#####
#      Post Commands      #####
POST_NTAP_DATA_TRANSFER_CMD01=
POST_APP QUIESCE_CMD01=/bin/sleep 60
POST_CLONE_CREATE_CMD01=
```



You should allow a wait time that is twice the value of the fs.xfs.xfssyncd_centisecs parameter. For example, with the default value 30 seconds, the sleep command should be configured with 60 seconds.

Configuring log backups


Log backups should be stored on a different storage system than the primary storage. The storage system that is used for the data backup can also be used for the log backup.

At the secondary storage, a volume needs to be configured to hold the log backups. Ensure that automatic Snapshot copies are switched off for this volume.

1. Mount the volume at each database node, either by running the mount command or editing the file system table (fstab) file.

```
hana2b:/vol/backup_log_ANA /mnt/backup_log_ANA nfs
rw,bg,vers=3,hard,timeo=600,rsize=65536,wszie=65536,actimeo=0,noatime
0 0
```

Within SAP HANA Studio, the log backup destination is configured as shown in the following figure.



Housekeeping of log backups

Housekeeping of log backups in SAP HANA is based on a function within the HANA Studio or based on an SQL statement that allows deleting all backups that are older than a selected backup.

Snap Creator handles the housekeeping of data backups (Snapshot copies) by deleting the Snapshot copies on the primary or secondary storage and by deleting the corresponding entries within the HANA catalog, based on a defined retention policy.

The log backups that are older than the latest data backup are deleted because they are not required.

Snap Creator handles the housekeeping of log file backups on file system level and within the SAP HANA backup catalog. As part of each Snapshot backup with Snap Creator, the following steps are executed:


- Read backup catalog and obtain the backup ID of the oldest successful data or Snapshot backup.
- Delete all backups that are older than the oldest backup.

i Snap Creator only handles housekeeping for backups based on Snapshot copies. If additional file-based backups are created, you must ensure that the file-based backups are deleted from the backup catalog and file system. If such a data backup is not deleted manually from the backup catalog, it can become the oldest data backup, and the log backup housekeeping operation will fail.

Modifying the housekeeping of log backups

You can modify the parameters that are configured for the housekeeping of log backups if you want to disable the log cleanup operation.

1. Select the SAP HANA profile that you want to modify.
2. Select the configuration you want to modify, and click **SAP HANA Settings**.
3. Edit the Enable LOG cleanup parameter, and click **Save**.



Executing database backups

You can back up your SAP HANA database by using the Snap Creator GUI or the command line. To schedule backups, you can use the scheduler within the GUI, or you can use the command line in combination with an external scheduler like cron.

Overview of database backups

When Snap Creator is backing up the database, the following steps are executed.

1. Create a global synchronized backup save point (SAP HANA Snapshot copy) to obtain a consistent image on the persistence layer.
2. Create storage Snapshot copies for all data volumes.


In the example, there are three data volumes, which are distributed to both storage controllers, hana1a and hana1b.

3. Register the storage Snapshot backup within the SAP HANA backup catalog.
4. Delete the SAP HANA Snapshot copy.
5. Start SnapVault update for all data volumes.
6. Check SnapVault status and wait until finished or configurable timeout.
7. Delete storage Snapshot copies and delete backups in the SAP HANA backup catalog based on the defined retention policy for backups at the primary and secondary storage.
8. Delete all log backups, which are older than the oldest data backup on the file system and within the SAP HANA backup catalog.


Backing up the database with the Snap Creator GUI

You can back up a database with the Snap Creator GUI.


1. Select the **HANA_database_backup configuration** and then select **Actions > Backup**.



2. Select the backup policy and click **OK**.



The backup starts. Snap Creator triggers the “SnapVault update,” and Snap Creator waits until the data is replicated to the secondary storage. The wait time has been configured during the configuration and can be adapted in the SnapVault settings tab. Snap Creator triggers the SnapVault updates in parallel for each volume on the same storage controller, but in sequence for each storage controller.



Backing up the database with Snap Creator command line

You can also back up the database by using the Snap Creator command line.

1. To back up the database, run the following command.

```

mgmtsrv01:~ #
/opt/NetApp/Snap_Creator_Framework_411/scServer4.1.1/snapcreator
--server
localhost --port 8443 --user scadmin --passwd scadmin --profile
HANA_profile_ANA --config
ANA_database_backup --action backup --policy daily --verbose
[Wed Mar 5 14:17:08 2014] INFO: Validating policy: daily finished
successfully

##### Detecting Data ONTAP mode for hanala #####
##### Detecting Data ONTAP mode for hanalb #####
[Wed Mar 5 14:17:13 2014] INFO: STORAGE-03031: Getting system version
details of [hana2b]
[Wed Mar 5 14:17:13 2014] INFO: STORAGE-03032: Getting system version
details of [hana2b] finished successfully.
[Wed Mar 5 14:17:13 2014] INFO: STORAGE-03031: Getting system version
details of [hanala]
[Wed Mar 5 14:17:13 2014] INFO: STORAGE-03032: Getting system version
details of [hanala] finished successfully.
[Wed Mar 5 14:17:13 2014] INFO: STORAGE-03031: Getting system version
details of [hanalb]
[Wed Mar 5 14:17:13 2014] INFO: STORAGE-03032: Getting system version
details of [hanalb] finished successfully.

.....
Truncated
...

```

Reviewing available backups in SAP HANA Studio

You can see the list of storage Snapshot backups in the SAP HANA Studio.

The highlighted backup in the following figure shows a Snapshot copy named “Backup-ANA_hourly_20140320103943.” This backup includes Snapshot copies for all three data volumes of the SAP HANA system. The backup is also available at the secondary storage.

The Snapshot copy name is used by Snap Creator as a backup ID when Snap Creator registers the storage Snapshot copy in the SAP HANA backup catalog. Within the SAP HANA Studio, the storage Snapshot backup is visible in the backup catalog. The external backup ID (EBID) has the same value as the Snapshot copy name as shown in the following figure.

Backup Catalog							Backup Details	
	Status	Started	Duration	Size	Backup Type	Destination Type	ID	Status
1	Mar 20, 2014 2:39:47 PM	00h 00m 12s	4.96 GB	Data Backup	Snapshot		12953003000000000000	Successful
2	Mar 20, 2014 3:39:47 AM	00h 00m 12s	4.99 GB	Data Backup	Snapshot		12953003000000000001	Successful
3	Mar 20, 2014 12:39:47 AM	00h 00m 14s	1.00 GB	Data Backup	Snapshot		12953003000000000002	Successful
4	Mar 19, 2014 3:39:47 PM	00h 00m 12s	4.97 GB	Data Backup	Snapshot		12953003000000000003	Successful
5	Mar 19, 2014 10:39:47 PM	00h 00m 12s	4.94 GB	Data Backup	Snapshot		12953003000000000004	Successful
6	Mar 19, 2014 9:39:47 PM	00h 00m 12s	4.90 GB	Data Backup	Snapshot		12953003000000000005	Successful
7	Mar 19, 2014 8:39:47 PM	00h 00m 12s	4.89 GB	Data Backup	Snapshot		12953003000000000006	Successful
8	Mar 19, 2014 7:39:47 PM	00h 00m 12s	4.87 GB	Data Backup	Snapshot		12953003000000000007	Successful
9	Mar 19, 2014 6:39:47 PM	00h 00m 12s	4.94 GB	Data Backup	Snapshot		12953003000000000008	Successful
10	Mar 19, 2014 5:39:47 PM	00h 00m 12s	4.92 GB	Data Backup	Snapshot		12953003000000000009	Successful
11	Mar 19, 2014 4:39:47 PM	00h 00m 12s	4.88 GB	Data Backup	Snapshot		12953003000000000010	Successful
12	Mar 19, 2014 3:40:00 PM	00h 00m 08s	4.86 GB	Data Backup	Snapshot		12953003000000000011	Successful
13	Mar 19, 2014 2:39:47 PM	00h 00m 30s	4.94 GB	Data Backup	Snapshot		12953003000000000012	Successful
14	Mar 19, 2014 12:00:04 PM	00h 00m 12s	4.85 GB	Data Backup	Snapshot		12953003000000000013	Successful
15	Mar 18, 2014 12:00:03 PM	00h 00m 12s	4.79 GB	Data Backup	Snapshot		12953003000000000014	Successful
16	Mar 17, 2014 12:00:03 PM	00h 00m 12s	4.55 GB	Data Backup	Snapshot		12953003000000000015	Successful
17	Feb 26, 2014 1:15:35 AM	00h 00m 54s	3.67 GB	Data Backup	Snapshot		12953003000000000016	Successful

With every backup run, Snap Creator deletes Snapshot backups at the primary and at the secondary storage based on the retention policies defined for the different schedules (hourly, daily, and so on).

Snap Creator also deletes the backups within the SAP HANA backup catalog if the backup does not exist at the primary or secondary storage. The SAP HANA backup catalog always shows the complete list of backups that are available at the primary and/or the secondary storage.

SAP HANA File-Based Backup and Database Integrity Checks

SAP recommends combining storage-based Snapshot backups with a weekly file-based backup to execute a block integrity check. The block integrity check can be executed from within the Snap Creator graphical user interface (GUI) or command line interface (CLI).

The File-Based Data Backup operation is used when the backup copies of files are to be retained. The Database Integrity Checks operation is used when backup copies have to be discarded.

You can configure either one or both of the operations. During on demand backup, you can choose either one of the operations.

Modifying configuration for File-Based Backup

You can modify the parameters that are configured for File-Based Backup. The subsequent scheduled or on-demand File-Based Backup operation reflects the updated information.

1. Click on the SAP HANA profile.
2. Select the configuration that you want to modify, and click **HANA File Based Backup Settings**.



3. Edit the information, and click **Save**.

Modifying configuration for Database Integrity Checks

You can modify the parameters that are configured for Database Integrity Checks. The subsequent scheduled or on-demand Integrity Check operation reflects the updated information.

1. Click on the SAP HANA profile.
2. Select the configuration that you want to modify, and click **HANA Integrity Check Settings**.




3. Edit the information, and click **Save**.

Scheduling file-based backup

For SAP HANA configurations, you can schedule additional operations such as file-based backup and database integrity checks. You can schedule the file-based backup operation to occur at specific intervals.

1. From the main menu of the Snap Creator GUI, select **Management > Schedules**, and click **Create**.
2. In the New Job window, enter the details for the job.

The file-based backup policy is set to “none” by default.




Scheduling database integrity checks

For SAP HANA configurations, you can schedule additional operations such as file-based backup and database integrity checks. You can schedule the database integrity checks operation to occur at specific intervals.

1. From the main menu of the Snap Creator GUI, select **Management > Schedules**, and click **Create**.
2. In the New Job window, enter the details for the job.

The integrity check policy is set to “none” by default.




Performing File-Based Backup from the Snap Creator GUI


You can perform File-Based Backup from the Snap Creator graphical user interface (GUI).

You must have enabled the File-Based Backup parameter in the HANA File-Based Backup Settings tab.

1. Select the HANA_database_backup configuration.
2. Select **Actions > File-Based Backup**.



3. Set the Policy option to **None**, and click **OK**.



Performing File-Based Backup from Snap Creator command line

You can perform File-Based Backup using the Snap Creator command line.

1. To perform File-Based Backup, run the following command:


```
./snapcreator --server localhost --port 8443 --user sc --passwd sc
--profile hana_testing --config HANA_Test --action fileBasedBackup
--policy none --verbose
```

Performing Database Integrity Checks from Snap Creator GUI


You can perform Database Integrity Checks from the Snap Creator graphical user interface (GUI).

You must have enabled the DB Integrity Check parameter in the HANA Integrity Check Settings tab.

1. Select the HANA_database_integrity_check configuration.
2. Select **Actions > Integrity Check**.



3. Set the Policy option to **None**, and click **OK**.



Performing Database Integrity Checks from Snap Creator command line

You can perform Database Integrity Checks using the Snap Creator command line.

1. To perform Database Integrity Checks, run the following command:

```
./snapcreator --server localhost --port 8443 --user sc --passwd sc
--profile hana_testing --config HANA_Test --action integrityCheck
--policy none --verbose
```

Restoring and recovering SAP HANA databases

You use SAP HANA Studio and Snap Creator to restore and recover SAP HANA databases.

1. Within SAP HANA Studio:
 - a. Select Recover for the SAP HANA system.
 - b. SAP HANA system is shut down.

- c. Select the recovery type.
- d. Provide log backup locations.
- e. List of data backups is shown
- f. Select backup to see the external backup ID.

2. For a storage system running clustered Data ONTAP only:

- a. Only required if any other backup than the latest has been used for the restore.
- b. Only required for "Volume SnapRestore" from primary storage.
- c. Deactivate SnapVault relationships

3. Within Snap Creator:

- a. Select "Restore" for the SAP HANA system.
- b. Select restore from primary or secondary storage, depending on the availability of the backup at the primary storage.
- c. Select storage controller, volume name, and Snapshot copy name. Snapshot copy name correlates with the backup ID earlier.
- d. For multinode SAP HANA systems, multiple volumes need to be restored:
 - i. Choose **Add more restore items**.
 - ii. Select storage controller, volume name, and Snapshot copy name.
 - iii. Repeat this process for all required volumes.
- e. For multitenant database containers (MDC) single tenant database systems, both the SYSTEM and TENANT databases are restored.
- f. Restore process is started
- g. Restore finished for all volumes.

4. At the database nodes, unmount and mount all data volumes to clean "Stale NFS Handles."

5. Within SAP HANA Studio:

- a. Select **Refresh** on backup list.
- b. Select available backup for recovery (green item).
- c. Start recovery process.
- d. For multitenant database containers (MDC) single tenant database systems, start the recovery process first for the SYSTEM database, and then for the TENANT database.
- e. The SAP HANA system is started.

6. (Optional) Resume SnapVault relationships for all restored volumes.



At the storage systems, this step is only required if a backup other than the latest one has been used for the restore.


Restoring and recovering databases from primary storage

You can restore and recover the database from the primary storage.



You cannot restore file-based backup copies from Snap Creator.

1. Within SAP HANA Studio, select **Recover** for the SAP HANA system.




The SAP HANA system shuts down.

2. Select the recovery type and click **Next**.




3. Provide the log backup locations and click **Next**.



The list of available backups you see is based on the content of the backup catalog.

4. Select the required backup and record the external backup ID.



5. Deactivate the SnapVault relationship.



This step is only required with clustered Data ONTAP.

If you need to restore a Snapshot copy that is older than the Snapshot copy currently used as the base Snapshot copy for SnapVault, you must first deactivate the SnapVault relationship in clustered Data ONTAP. To do that, execute the following commands on the backup cluster console:

```


hana::> snapmirror quiesce -destination-path hana2b:backup_hana_data
Operation succeeded: snapmirror quiesce for destination
hana2b:backup_hana_data.

hana::> snapmirror delete -destination-path hana2b:backup_hana_data
Operation succeeded: snapmirror delete the relationship with destination
hana2b:backup_hana_data.

hana::> snapmirror release -destination-path hana2b:backup_hana_data
[Job 6551] Job succeeded: SnapMirror Release Succeeded


```

6. In the Snap Creator GUI, select the SAP HANA system, then select **Actions > Restore**.




The Welcome to the Snap Creator Framework Restore Wizard screen appears.

7. Click **Next**.




8. Select **Primary** and click **Next**.




9. Select restore from primary storage.
10. Select the storage controller, the volume name, and the Snapshot name.


The Snapshot name correlates with the backup ID that has been selected within SAP HANA Studio.



11. Click **Finish**.




12. Click **Yes** to add more restore items.




13. Select the storage controller, the additional volume name, and the Snapshot name.


The Snapshot name correlates with the backup ID that has been selected within SAP HANA Studio.



14. Repeat steps 10 through 13 until all required volumes are added; in our example, data_00001, data_00002, and data_00003 need to be selected for the restore process.
15. When all volumes are selected, click **OK** to start the restore process.



The restore process is started.




Wait until the restore process is finished.

16. On each database node, remount all data volumes to clean Stale NFS Handles.

In the example, all three volumes need to be remounted at each database node.


```
mount -o remount /hana/data/ANA/mnt00001
mount -o remount /hana/data/ANA/mnt00002
mount -o remount /hana/data/ANA/mnt00003
```

17. Go to SAP HANA Studio and click **Refresh** to update the list of available backups.




The backup that has been restored with Snap Creator is shown with a green icon in the list of backups.


18. Select the backup and click **Next**.




19. Select other settings as required and click **Next**.




20. Click **Finish**.



The recovery process begins.



21. After recovery is finished, resume the SnapVault relationships, if required.



Restoring and recovering databases from secondary storage

You can restore and recover the database from the secondary storage.

1. Within SAP HANA Studio, select **Recover** for the SAP HANA system.

The screenshot shows the SAP HANA Administration Console with the following details:

- Backup Catalog View:** The main area displays a table of backup jobs. The columns are: Duration, Size, Backup Type, and Destination Type. The last row is selected, showing the following details:


2:47	00h 00m 12s	5.07 GB	Data Backup	Snapshot
------	-------------	---------	-------------	----------
- Context Menu:** A context menu is open over the selected backup job, listing options such as:
 - Add System
 - Add System Archive List
 - Add Additional User
 - Administration
 - System Replication...
 - Open Security
 - Backup...
 - Storage Snapshot...
 - Transport Management
 - Lifecycle Management
 - Open Memory Overview
 - Open Resource Utilization
 - Delete
 - Refresh
 - Log Off
 - SQL Console
 - Find Table
 - Properties
 - Ad-Hoc
- Backup Details:** A panel on the right shows detailed information for the selected backup job:

ID	1396548004221
Status	Successful
Backup Type	Data Backup
Destination Type	Snapshot
Started	Apr 3, 2014 11:00:04 AM (America/Los_Angeles)
Finished	Apr 3, 2014 11:00:16 AM (America/Los_Angeles)
Duration	00h 00m 12s
Size	5.07 GB
Throughput	n/a
Comment	Backup-ANA-daily_20140403200000
Additional Information	none
Location	Hana/data/ANA/mtd/00001/
- File List:** A table below the details shows a list of backup files:


Host	Service	Size	EBID
osthana08	namaserver	76.89 MB	Backup-ANA-daily_20140403200000
osthana08	indexserver	260.36 MB	Backup-ANA-daily_20140403200000
osthana08	indexserver	257.16 MB	Backup-ANA-daily_20140403200000
osthana08	indexserver	256.28 MB	Backup-ANA-daily_20140403200000
osthana08	statisticsserver	1.96 GB	Backup-ANA-daily_20140403200000
osthana08	esengine	45.03 MB	Backup-ANA-daily_20140403200000
osthana10	indexserver	325.27 MB	Backup-ANA-daily_20140403200000
osthana10	indexserver	324.27 MB	Backup-ANA-daily_20140403200000
osthana10	indexserver	325.09 MB	Backup-ANA-daily_20140403200000
osthana11	indexserver	324.83 MB	Backup-ANA-daily_20140403200000
osthana11	indexserver	322.45 MB	Backup-ANA-daily_20140403200000
osthana11	indexserver	324.41 MB	Backup-ANA-daily_20140403200000

The SAP HANA system will be shut down.

2. Select the recovery type and click **Next**.




3. Provide log backup locations and click **Next**.




The list of available backups appear based on the content of the backup catalog.


4. Select the required backup and write down external backup ID.




5. Go to the Snap Creator GUI.
6. Select the SAP HANA system, and then click **Actions > Restore**.



The Welcome screen appears.




7. Click **Next**.
8. Select **Secondary** and click **Next**.




9. Enter the required information. The Snapshot name correlates with the backup ID that has been selected in


SAP HANA Studio.




10. Select **Finish**.



11. Click **Yes** to add more items to restore.



12. Provide the required information for all volumes that need to be restored. In the setup data_00001, data_00002, and data_00003 need to be selected for the restore process.



13. When all volumes are selected, select **OK** to start the restore process.


Wait until the restore process is finished.

14. On each database node remount all data volumes to clean “Stale NFS Handles.”

In the example, all three volumes need to be remounted at each database node.


```
mount -o remount /hana/data/ANA/mnt00001
mount -o remount /hana/data/ANA/mnt00002
mount -o remount /hana/data/ANA/mnt00003
```

15. Go to SAP HANA Studio and click **Refresh** to update the backup list.




16. The backup that has been restored with Snap Creator is shown with a green icon in the list of backups. Select the backup and click **Next**.


17. Select other settings as required and click **Next**.



18. Click **Finish**.



The recovery process begins.



19. After the recovery process is finished, resume the SnapVault Relationships, if required.



Resuming a SnapVault relationship after a restore

Any restore that is not done using the latest Snapshot backup will delete the SnapVault relationship at the primary storage systems.

After the restore and recovery process is finished, the SnapVault relationship has to be resumed so that backups can be executed again with Snap Creator. Otherwise, Snap Creator will issue an error message, because it can't find the SnapVault relationship anymore at the primary storage systems.

The data transfer that is required will be based on a delta transfer, if there is still a common Snapshot copy between the source volume and the destination volume.

Resuming a SnapVault relationship with Data ONTAP operating in 7-Mode

If you restore using a Snapshot backup other than the latest one, you need to resume the SnapVault relationship so that Snap Creator can continue to run backups.

1. Resume the SnapVault relationship with Data ONTAP operating in 7-Mode by entering the following command. `snapshot start -r -S source_controller:source_volume backup_controller:backup_volume`

Perform this step for all volumes belonging to the SAP HANA database.

```
hana2b> snapshot start -r -S hanala:/vol/data_00001/mnt00001
hana2b:/vol/backup_data_00001/mnt00001
The resync base snapshot will be: Backup-ANA-SV_daily_20140406200000
Resync may alter the data in this qtree.
Are you sure you want to resync the qtree? y
Mon Apr  7 14:08:21 CEST [hana2b:replication.dst.resync.success:notice]:
SnapVault resync of
/vol/backup_data_00001/mnt00001 to hanala:/vol/data_00001/mnt00001 was
successful.
Transfer started.
Monitor progress with 'snapshot status' or the snapmirror log.
```

```
hana2b> snapshot start -r -S hanalb:/vol/data_00002/mnt00002
hana2b:/vol/backup_data_00002/mnt00002
The resync base snapshot will be: Backup-ANA-SV_daily_20140406200000
Resync may alter the data in this qtree.
Are you sure you want to resync the qtree? y
Mon Apr  7 14:09:49 CEST [hana2b:replication.dst.resync.success:notice]:
SnapVault resync of
/vol/backup_data_00002/mnt00002 to hanalb:/vol/data_00002/mnt00002 was
successful.
Transfer started.
Monitor progress with 'snapshot status' or the snapmirror log.
```

```

hana2b> snapvault start -r -S hanala:/vol/data_00003/mnt00003
hana2b:/vol/backup_data_00003/mnt00003
The resync base snapshot will be: Backup-ANA-SV_daily_20140406200000
Resync may alter the data in this qtree.
Are you sure you want to resync the qtree? y
Mon Apr  7 14:10:25 CEST [hana2b:replication.dst.resync.success:notice]:
SnapVault resync of
/vol/backup_data_00003/mnt00003 to hanala:/vol/data_00003/mnt00003 was
successful.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.

```

When the data transfer is finished, you can again schedule backups by using Snap Creator.

Resuming a SnapVault relationship with clustered Data ONTAP

If you restore using a Snapshot backup other than the latest one, you need to resume the SnapVault relationship so that Snap Creator can continue to run backups.

1. Re-create and resynchronize the SnapVault relationship.

```

hana::> snapmirror create -source-path hanala:hana_data -destination
-path
hana2b:backup_hana_data -type XDP
Operation succeeded: snapmirror create the relationship with destination
hana2b:backup_hana_data.

hana::> snapmirror resync -destination-path hana2b:backup_hana_data
-type XDP

Warning: All data newer than Snapshot copy sc-backup-
daily_20140430121000 on volume
hana2b:backup_hana_data will be deleted.
Do you want to continue? {y|n}: y
[Job 6554] Job is queued: initiate snapmirror resync to destination
"hana2b:backup_hana_data".
[Job 6554] Job succeeded: SnapMirror Resync Transfer Queued

```

2. To actually restart the SnapVault transfer, a manual Snapshot copy is required.

```

hana::> snapshot create -vserver hanala -volume hana_data -snapshot
sv_resync

hana::> snapshot modify -vserver hanala -volume hana_data -snapshot
sv_resync -snapmirror-label daily

hana::> snapmirror update -destination-path hana2b:backup_hana_data
Operation is queued: snapmirror update of destination
hana2b:backup_hana_data.

```

3. Verify that the SnapVault relationship appears in the destination list.

```

hana::> snapmirror list-destinations -source-path hanala:hana_data
                                         Progress
                                         Source          Destination          Transfer      Last
                                         Relationship
                                         Path      Type    Path      Status Progress      Updated      Id
-----  -----
-----  -----
hanala:hana_data
          XDP    hana2b:backup_hana_data
          Transferring
          38.46KB  04/30 18:15:54
          9137fb83-
cba9-11e3-85d7-123478563412

```

Restoring databases after primary storage failure

After a primary storage failure, or when all Snapshot copies are deleted from the volumes at the primary storage, Snap Creator will not be able to handle the restore, because there will no longer be a SnapVault relationship on the primary storage systems.

Restoring databases after a primary storage failure with Data ONTAP operating in 7-Mode

You can restore an SAP HANA database after a primary storage system running Data ONTAP operating in 7-Mode fails.

1. In this case, the restore has to be executed directly on the secondary storage system by using the following command: `snapshot restore --snapshot_name -S backup_controller:backup_volumesource_controller:source_volume`

Perform this step for all volumes belonging to the SAP HANA database.

```
hanala> snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00001/mnt00001 hanala:/vol/data_00001/mnt00001
Restore will overwrite existing data in /vol/data_00001/mnt00001.
Are you sure you want to continue? y
Thu Apr 10 11:55:55 CEST [hanala:vdisk.qtreePreserveComplete:info]:
Qtree preserve is complete for /vol/data_00001/mnt00001.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hanala> snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00003/mnt00003 hanala:/vol/data_00003/mnt00003
Restore will overwrite existing data in /vol/data_00003/mnt00003.
Are you sure you want to continue? y
Thu Apr 10 11:58:18 CEST [hanala:vdisk.qtreePreserveComplete:info]:
Qtree preserve is complete for /vol/data_00003/mnt00003.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hanalb> snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00002/mnt00002 hanalb:/vol/data_00002/mnt00002
Restore will overwrite existing data in /vol/data_00002/mnt00002.
Are you sure you want to continue? y
Thu Apr 10 12:01:29 CEST [hanalb:vdisk.qtreePreserveComplete:info]:
Qtree preserve is complete for /vol/data_00002/mnt00002.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

When the restore process is finished, you use SAP HANA to perform the recovery.

Restoring databases after a primary storage failure with clustered Data ONTAP

You can restore an SAP HANA database after a primary storage system running clustered Data ONTAP fails.

Assuming the primary volume is lost completely, you need to create a new primary volume and then restore from the backup volume.

1. Create a primary volume with type data protection.

```

hana::> volume create -vserver hanala -volume hana_data -aggregate
aggr_sas_101 -size 300G -state online -type DP -policy default -autosize
-mode grow_shrink -space-guarantee none
-snapshot-policy none -foreground true
[Job 6744] Job is queued: Create hana_data.
[Job 6744] Job succeeded: Successful

```

2. Restore all data from the backup volume.

```

hana::> snapmirror restore -destination-path hanala:hana_data -source
-path hana2b:backup_hana_data -source-snapshot sc-backup-
daily_20140505121000
[Job 6746] Job is queued: snapmirror restore from source
"hana2b:backup_hana_data" for the
snapshot sc-backup-daily_20140505121000.

hana::> job show -id 6746
Owning
Job ID Name Vserver Node State
----- -----
6746 SnapMirror restore hana hana01 Running
Description: snapmirror restore from source
"hana2b:backup_hana_data" for the snapshot sc-backup-
daily_20140505121000

```

When the restore process is finished, you use SAP HANA to perform the recovery.

SAP HANA plug-in parameters

The following table lists the SAP HANA plug-in parameters, provides the parameter settings, and describes the parameters.

Parameter	Setting	Description
HANA_SID	Example: ABC	HANA database SID.
HANA_NODES	Example: node1, node2, node3	Comma-separated list of HANA nodes on which the hdbsql statements can be executed.
HANA_USER_NAME	Example: backupUser	HANA database user name. The minimum privilege required for this user is BACKUP ADMIN privilege.

Parameter	Setting	Description
HANA_PASSWORD	Example: hfasfh87r83r	HANA database password.
HANA_INSTANCE	Example: 42	HANA node instance number.
HANA_HDBSQL_CMD	Example: /usr/sa p/hdbclient/ hdbsql	Path to the HANA hdbsql command. If this parameter is not set, hdbsql on the search path is used. The default is hdbsql.
HANA_OSDB_USER	Example: user1	The operating system user for executing hdbsql (usually sidadm) must have the hdbsql binary in the search path and the permission to execute it.
HANA_USERSTORE_KEYS	Example: node1:key1, node 2:key2, node3:ke y3	Comma-separated list of HANA userstore keys and node pairs using which the hdbsql statements can be executed.
HANA_FILE_BACKUP_ENABLE	“Y” or “N”	Determines whether Snap Creator should enable file-based backup for the SAP HANA plug-in. This setting is useful when you want to perform the SAP HANA file-based backup operation.
HANA_FILE_BACKUP_PATH	Example:/hana/data/SCN/mnt0000 1	(Optional) Path to the directory where database file backup can be stored. If this parameter is not set, use default.
HANA_FILE_BACKUP_PREFIX	Example: SnapCreator_<HANA_FILE_BACK UP_PREFIX>_<CURRENT_TIME STAMP>	(Optional) Adds a prefix to the backup file name. Default: SnapCreator_<CURRENT_TIME STAMP>
HANA_INTEGRITY_CHECK_ENA BLE	“Y” or “N”	Determines whether Snap Creator should enable Integrity Check for the SAP HANA plug-in. This setting is usual when you want to perform the SAP HANA Integrity Check operation.
HANA_TEMP_FILE_BACKUP_PA TH	Example:/temp	(Optional) Path where the temporary database file for Integrity Check can be stored. If not sure, use default.

Parameter	Setting	Description
HANA_LOG_CLEANUP_ENABLE	“Y” or “N”	Enables Log Catalog cleanup.

Troubleshooting

The troubleshooting section provides information about the error codes, error messages, and includes the description or resolution to solve the issue.

The following table lists the SAP HANA plug-in error messages.

Error code	Error message	Description/Resolution
hdb-00001	Unable to find an accessible HANA node for executing hdbsql commands using the provided configuration parameters. Verify and update HANA settings in the configuration and try again.	Verify that HANA nodes are running and reachable, and the instance number provided is correct.
hdb-00002	Creating database snapshot for [\$sid] failed.	Check if a HANA database snapshot is already created on the database. If already created, delete the HANA database snapshot or run unquiesce operation. If not already created, check the logs for other error messages and details.
hdb-00003	Deleting database snapshot for [\$sid] failed.	Check if a HANA database snapshot is already deleted. If yes, this error can be ignored. If no, check SAP HANA plug-in parameters and make sure that nodes are reachable and instance number provided is correct.
hdb-00004	Connection to [\$hana_node] node with instance [\$instance] failed as the connection was refused.	The HANA node with instance displayed in the message are not reachable. This can be just a warning as the plug-in will attempt to run hdbsql commands on other nodes. Check the logs to see if the operation was successful.
hdb-00005	Database [\$sid] already has a snapshot!	HANA database snapshot already exists on the database. Delete the HANA database snapshot or run unquiesce operation to resolve this issue.

Error code	Error message	Description/Resolution
hdb-00006	Unable to resolve hostname [\$hana_node].	The HANA node hostname cannot be resolved. Check your DNS server or etc hosts entries.
hdb-00007	Invalid username or password. Verify the credentials and try again.	The user name and password provided for HANA database is incorrect. Correct the entries in the configuration file and try again.
hdb-00008	Running command [\$hdbsql_cmd] on [\$hana_node] failed.	Plug-in failed to execute hdbsql command on all HANA nodes provided in the configuration. Verify the HANA nodes and instance parameters and ensure at least one HANA node is up and reachable.
hdb-00009	Unable to find HANA [\$info].	The SAP HANA plug-in SCDUMP operation was unable to retrieve a particular information from the HANA databases. Verify the HANA nodes and instance parameters and make sure at least one HANA node is up and reachable.
hdb-00010	Collection of OS information failed.	The collection of OS information failed in the Windows environment; the SAP HANA plug-in is not supported on Windows. Use an SLES operating system instead.
hdb-00011	Collection of OS information failed.	Snap Creator was unable to collect OS information for the SCDUMP operation. Check your agent configuration file and correct the settings.
hdb-00012	Collection of SnapDrive information failed.	The SAP HANA plug-in is only supported in an NFS environment. Your configuration for HANA database has SnapDrive enabled; set SNAPDRIVE=Nin the configuration file.
hdb-00013	The HANA_NODES parameter is not set. Check HANA settings in the configuration file.	HANA nodes (HANA_NODES) parameter is required for the SAP HANA plug-in. Set the parameter and try again.

Error code	Error message	Description/Resolution
hdb-00014	Unable to find an accessible HANA node for executing hdbsql commands using the provided configuration parameters. Verify and update HANA settings in the configuration and try again.	Verify that HANA nodes are running and reachable, and the instance number provided is correct.
hdb-00015	The HANA_INSTANCE parameter is not set. Check HANA settings in the configuration file.	HANA instance (HANA_INSTANCE) parameter is required for the SAP HANA plug-in. Set the parameter and try again.
hdb-00016	The HANA_PASSWORD parameter is not set. Check HANA settings in the configuration file.	HANA password (HANA_PASSWORD) parameter is required for the SAP HANA plug-in. Set the parameter and try again.
hdb-00017	Path to hdbsql, value of parameter HANA_HDBSQL_CMD is invalid!	<p>One of the following has occurred:</p> <ul style="list-style-type: none"> • You have not provided the hdbsql path • The hdbsql path provided is incorrect. <p>Ensure you have the HANA hdbsql client installed on the management host where Snap Creator Agent is installed, and provide the correct path of the hdbsql binary in HANA parameters; then, try again.</p>

Where to go next

You can find more information about Snap Creator, including release-specific information, on the NetApp Support Site.

- [Snap Creator Framework 4.3.3 Installation Guide](#)

Describes how to install the Snap Creator Server and Agent. The Agent installation includes the SAP Hana plug-in.

- [Snap Creator Framework 4.3.3 Administration Guide](#)

Describes how to administer the Snap Creator Framework after installation is complete.

- [Snap Creator Framework 4.3.3 Release Notes](#)

Describes new features, important cautions, known problems, and limitations for the Snap Creator Framework 4.1.1 product.

- [Snap Creator Framework Discussions](#)

Connect with peers, ask questions, exchange ideas, find resources, and share Snap Creator best practices.

- [NetApp Video: SnapCreatorTV](#)

View videos demonstrating key Snap Creator technologies.

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.