Technical Report

SAP HANA Backup and Recovery by Using Snap Creator

Nils Bauer, Bernd Herth, NetApp
March 2016 | TR-4313

Abstract

The document describes the installation and configuration of the NetApp® backup and recovery solution for SAP HANA. The solution is based on the NetApp Snap Creator® framework and the Snap Creator plug-in for SAP HANA. This solution is supported with the certified Cisco SAP HANA multinode appliance in combination with NetApp storage. This solution is also supported with single-node and multinode SAP HANA systems in tailored data center integration (TDI) projects.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Overview

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

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

The following points summarize the challenges facing SAP backup and recovery:

- **Performance effects on production SAP systems.** Typically, traditional copy-based backups create a significant performance drain on production SAP systems because of the heavy loads placed on the database server, the storage system, and the storage network.
- **Shrinking backup windows.** Conventional backups can only be made when few dialog or batch activities are in process on the SAP system. The scheduling of backups becomes more difficult when SAP systems are in use around the clock.
- **Rapid data growth.** Rapid data growth and shrinking backup windows require ongoing investment in backup infrastructure. In other words, you must procure more tape drives, newer tape drive technology, and faster storage networks. You must also cover the ongoing expense of storing and managing these tape assets. Incremental or differential backups can address these issues, but this arrangement results in a very slow, cumbersome, and complex restore process that is harder to verify. Such systems usually increase RTO and RPO times in ways that are not acceptable to the business.
- **Increasing cost of downtime.** Unplanned downtime of an SAP system typically affects business finances. A significant part of any unplanned downtime is consumed by the need to restore and recover the SAP system. Therefore, the desired recovery time objective (RTO) dictates the design of the backup and recovery architecture.
- **Backup and recovery time for SAP upgrade projects.** The project plan for an SAP upgrade includes at least three backups of the SAP database. These backups significantly reduce the time available for the upgrade process. The decision to proceed is generally based on the amount of time required to restore and recover the database from the previously created backup. Rather than just restoring a system back to its previous state, a rapid restore provides more time to solve problems that might occur during an upgrade.

1.1 The NetApp Solution

NetApp Snapshot® technology can be used to create database backups within 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, again because the NetApp Snapshot technology does not move or copy data blocks when the Snapshot copy is created or when data in the active file system is changed. Therefore, the creation of Snapshot copies can be scheduled without considering peak dialog or batch activity periods. SAP and NetApp customers typically schedule multiple online Snapshot backups during the day; for example, every four hours is common. These Snapshot backups are typically kept for three to five days on the primary storage system before being removed.

Snapshot copies also provide key advantages for restore and recovery operations. NetApp SnapRestore® data recovery software enables the restore of an entire database or, alternatively, a portion of a database to any point in time, based on the available Snapshot copies. Such restore processes are finished in a few minutes, independent of the size of the database. Because several online Snapshot backups are

created during the day, the time needed for the recovery process is significantly reduced relative to a traditional backup approach. Because a restore can be performed with a Snapshot copy that is only a few hours old (rather than up to 24 hours), fewer transaction logs must be applied. Therefore the mean time to recover, which is the time needed for restore and recovery operations, is reduced to several minutes rather than the multiple hours required for conventional single-cycle tape backups.

Snapshot backups are stored on the same disk system as the active online data. Therefore, NetApp recommends using Snapshot backups as a supplement rather than a replacement for backups to a secondary location. Most restore and recovery actions are handled by using SnapRestore on the primary storage system. Restores from a secondary location are only necessary if the primary storage system holding the Snapshot copies is damaged. The secondary location can also be used if it is necessary to restore a backup that is no longer available from a Snapshot copy: a month-end backup, for example.

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

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

Backing up data to tape as a long-term backup might still be required. This backup 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 data can be written to tape by using the Network Data Management Protocol (NDMP). Figure 1 shows an overview of the backup solution.

Figure 1) Backup solution overview.
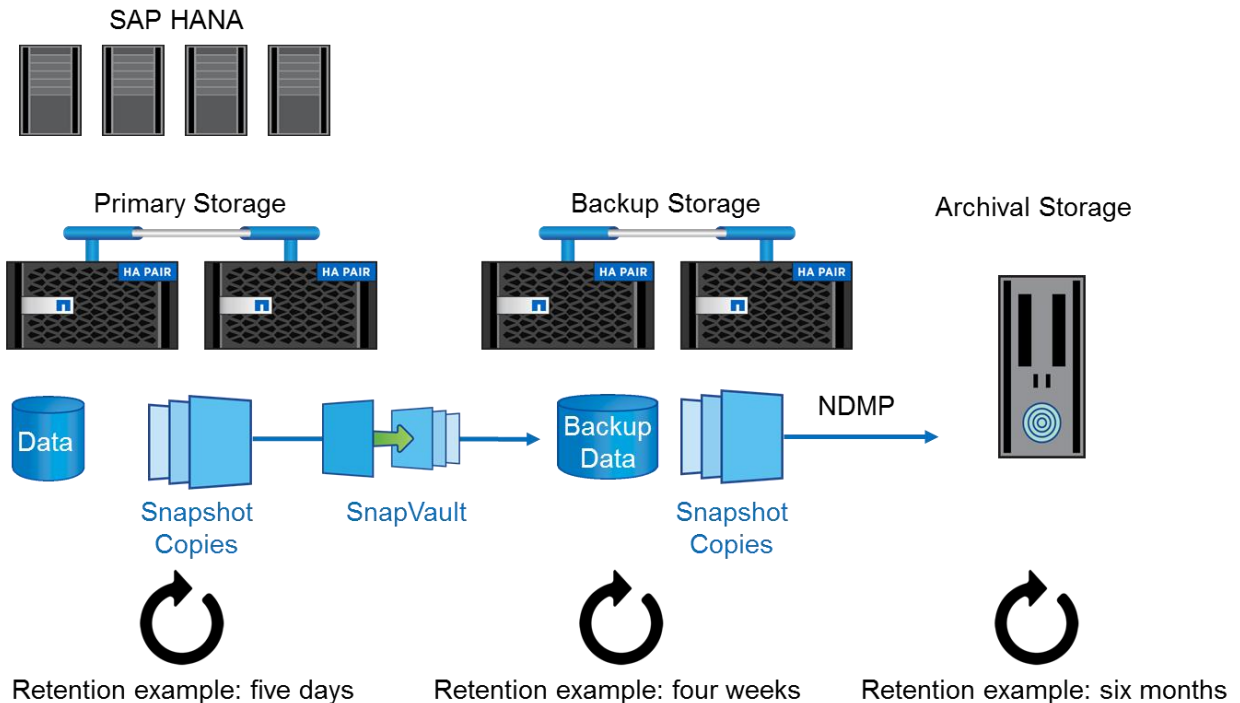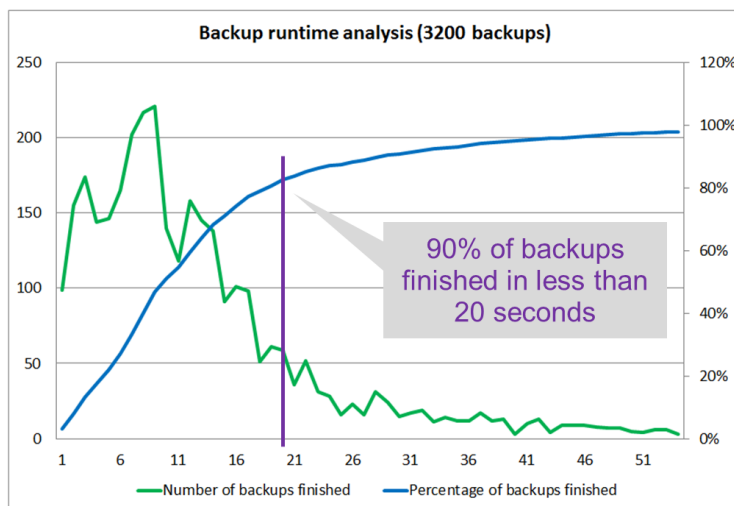
1.2 Runtime of Snapshot Backups

Figure 2 shows a screenshot of a customer's HANA Studio running SAP HANA on NetApp storage. The customer is using the NetApp Snap Creator storage management framework to manage backups of the HANA database. The screenshot shows that the HANA database (approximately 620GB in size) is backed up in 11 seconds by using Snapshot backup technology.

Figure 2) Customer example of Snapshot backup runtime.
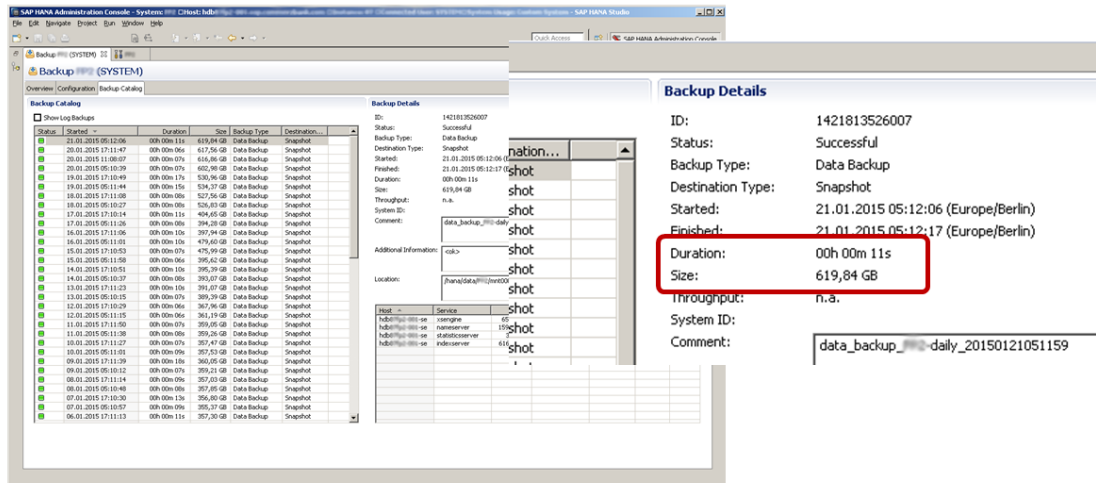
Further analysis of more than 3,000 backup runs demonstrated that more than 90% of the backups were finished in less than 20 seconds. All of the backups were finished in less than a minute.

Figure 3) Backup runtime analysis.

1.3 Recovery Time Objective Comparison

This section provides an RTO comparison of file-based and storage-based Snapshot backups. The RTO is defined by the sum of the time needed to restore the database and the time needed to start and recover the database.

Time Needed to Restore Database

With a file-based backup, the restore time depends on the size of the database and backup infrastructure, which defines the restore speed in megabytes per second. For example, if the infrastructure supports a restore operation at a speed of 250MBps, it takes approximately one hour and ten minutes to restore a database 1TB in size.

With storage Snapshot copy backups, the restore time is independent of the size of the database and is always in the range of a couple of seconds when the restore can be performed from primary storage. A restore from secondary storage is only required in the case of a disaster when the primary storage is no longer available.

Time Needed to Start Database

The database start time depends on the size of the row and column store. For the column store, the start time also depends on how much data is preloaded during the database start. In the following examples, we assume that the start time is 30 minutes. The start time is the same for a file-based restore and recovery and a restore and recovery based on Snapshot.

Time Needed to Recover Database

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

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

Storage Snapshot data backups are typically scheduled with a higher frequency because they do not influence the performance of the SAP HANA database. For example, if Snapshot 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 = ¼).

Figure 4 shows an RTO example for a 1TB database when file-based data backups are used. In this example, a backup is taken once per day. The RTO differs depending on when the restore and recovery were performed. If the restore and recovery were performed immediately after a backup was taken, the RTO is primarily based on the restore time, which is 1 hour and 10 minutes in the example. The recovery time increased to 2 hours and 50 minutes when restore and recovery were performed immediately before the next backup was taken, and the maximum RTO was 4 hours and 30 minutes.

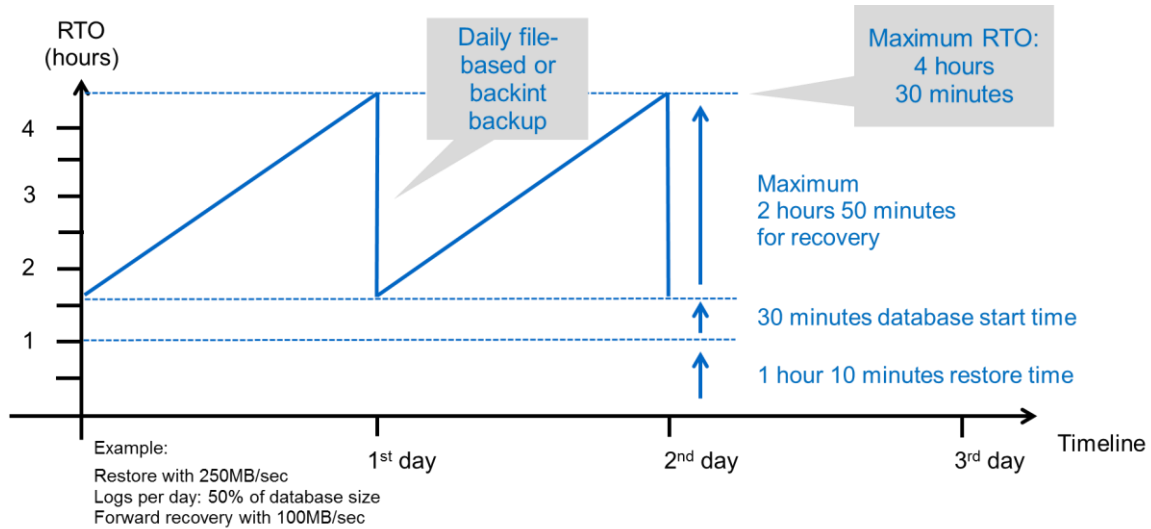**Figure 4) RTO for a 1TB database with file-based backups.**

Figure 5 shows an RTO example for a 1TB database when Snapshot backups are used. With storage-based Snapshot backups, the RTO only depends on the database start time and the forward recovery time because the restore is completed in a few seconds, independent of the size of the database. The forward recovery time also increases depending on when the restore and recovery are done, but due to the higher frequency of backups (every 6 hours in this example), the forward recovery time is 43 minutes at most. In this example, the maximum RTO is 1 hour and 13 minutes.

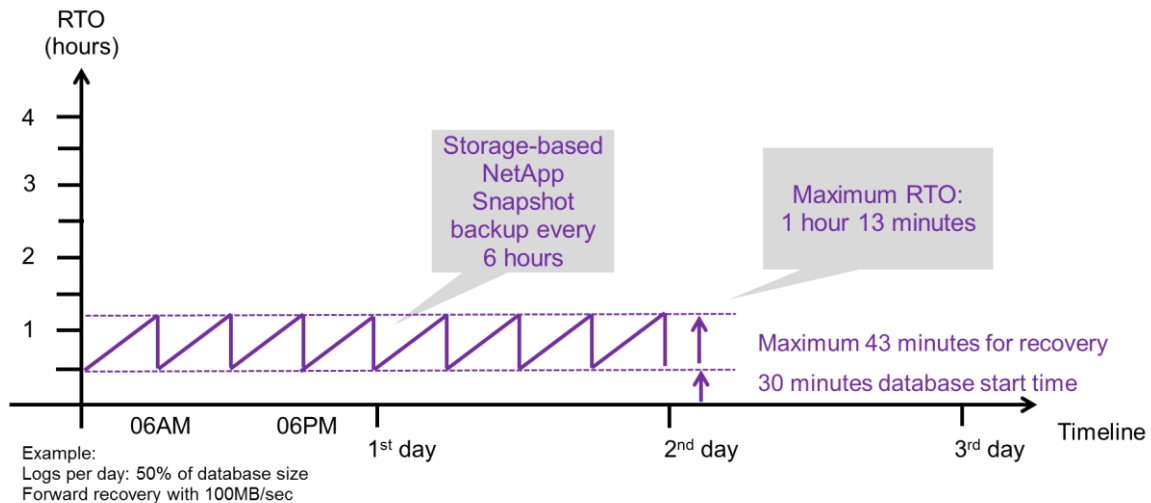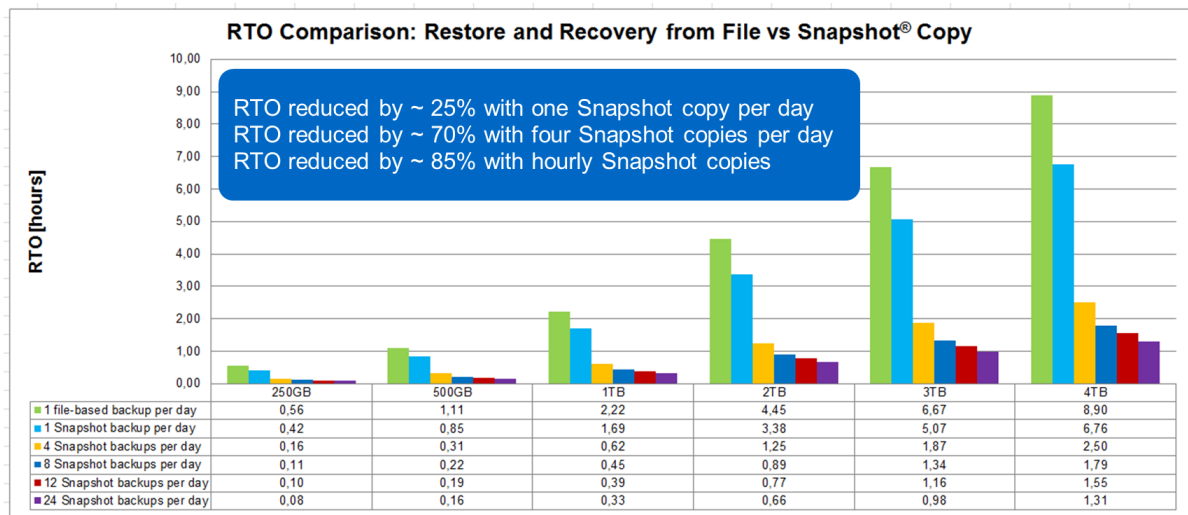**Figure 5) RTO for a 1TB database with Snapshot backups.**

Figure 6 shows an RTO comparison of file-based and storage-based Snapshot data backups for different database sizes and different frequencies of Snapshot backups. The green bar shows the file-based backup. The other bars show Snapshot backups with different backup frequencies.

With a single Snapshot data backup per day, the RTO is already reduced by 25% compared to a file-based data backup. The reduction increases to 70% when 4 Snapshot data backups are taken per day and goes up to 85% when 24 Snapshot data backups are taken per day.

Figure 6) RTO comparison: file-based backup versus Snapshot backup.

RTO Comparison: Restore and Recovery from File vs Snapshot® Copy

RTO reduced by ~ 25% with one Snapshot copy per day
RTO reduced by ~ 70% with four Snapshot copies per day
RTO reduced by ~ 85% with hourly Snapshot copies

	250GB	500GB	1TB	2TB	3TB	4TB
1 file-based backup per day	0,56	1,11	2,22	4,45	6,67	8,90
1 Snapshot backup per day	0,42	0,85	1,69	3,38	5,07	6,76
4 Snapshot backups per day	0,16	0,31	0,62	1,25	1,87	2,50
8 Snapshot backups per day	0,11	0,22	0,45	0,89	1,34	1,79
12 Snapshot backups per day	0,10	0,19	0,39	0,77	1,16	1,55
24 Snapshot backups per day	0,08	0,16	0,33	0,66	0,98	1,31

Assumptions: restore from file with 250MB/sec; log files per day: 50% of database size;
forward recovery with 100MB/sec

Note: The graph shows the HANA server RAM size. The database size in memory is calculated to be half of the server RAM size.

Note: The restore and recovery time is calculated based on the following assumptions. The database can be restored at 250MBps. The number of log files per day is 50% of the database size. For example, a 1TB database creates 500MB of log files per day. A recovery can be performed at 100MBps.

2 Snap Creator Backup Solution Components

Snap Creator can be used to back up SAP HANA database files as well as nondatabase files such as the /hana/shared file system. The backup of nondatabase files can either be included in the Snap Creator configuration for database backups, or it can be put in a separate Snap Creator configuration. The configuration and workflows described in this document cover database backup, but not the backup of nondatabase files.

The Snap Creator backup solution for SAP HANA covers the following areas:

- SAP HANA data file backup with storage-based Snapshot copies
- Replication of data file backups to a secondary off-site backup location
- SAP HANA log file backup with the HANA database log backup functionality
- Database block integrity checks with a file-based backup
- Housekeeping of data file backups, log file backups, and the SAP HANA backup catalog

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

Snap Creator enables the replication of consistent database images to a secondary storage location by using SnapVault. Typically, different retention policies are defined for backups in primary storage and backups in secondary storage. Snap Creator handles retention at primary storage as well as at secondary storage.

Log backup is executed automatically by the SAP HANA database. Depending on the recovery point objectives, there are several options for the storage location of the log backups:

- The log backup is written to a storage system that synchronously mirrors the data to a second location with NetApp MetroCluster™ high-availability (HA) and disaster recovery storage software.
- The log backup destination can be configured on the same primary storage system and then replicated asynchronously to a secondary storage with SnapMirror®.
- The log backup destination can be configured on the same secondary storage in which the database backups are replicated with SnapVault. With this configuration, the secondary storage has availability requirements similar to those of the primary storage so that log backups can always be written to the secondary storage.

SAP recommends combining storage-based Snapshot backups with a weekly file-based backup to execute a block integrity check. The block integrity check can be executed from within the Snap Creator GUI or CLI. Based on your configurable retention policies, Snap Creator manages the housekeeping of log file backups, the SAP HANA backup catalog, and data file backups at the primary and secondary storage location.

Figure 7 shows an overview of the database and log backup configuration.

Figure 7) Database and log backup.
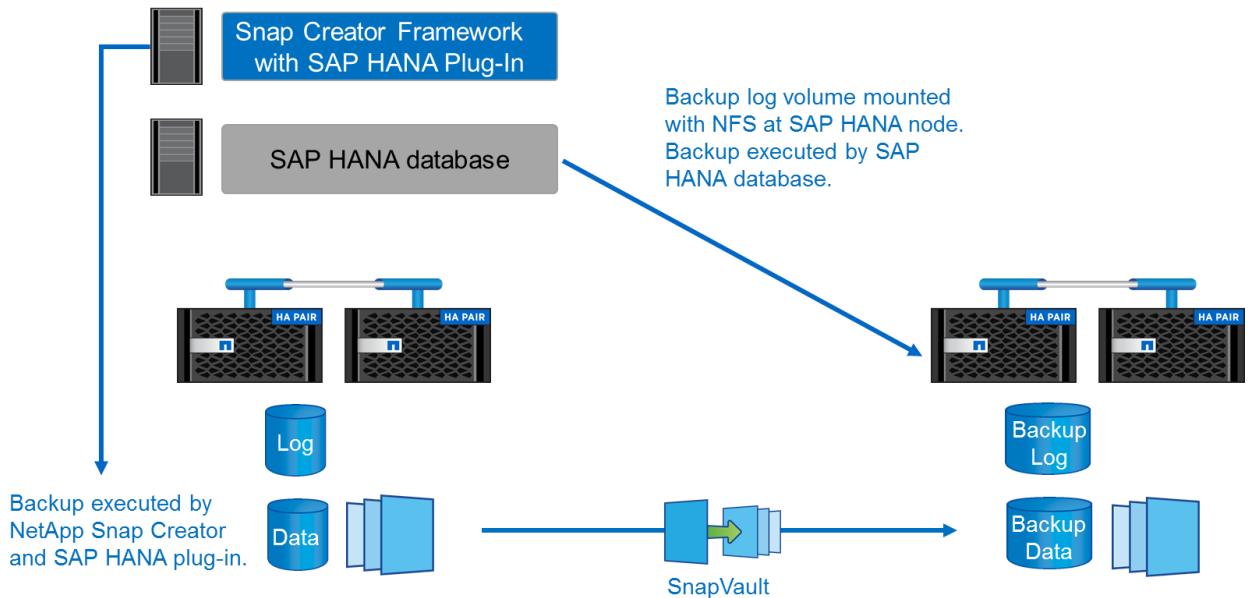
Backup schedules and retention policies must be defined based on customer requirements. Table 1 shows an example configuration of different schedules and retention policies.

Table 1) Example: backup schedules and retention policies.

	Executed By	Primary Storage	Secondary Storage
Database backups	Snap Creator Schedule 1: every four hours Schedule 2: once per day	Retention: 6 Retention: 3 - 6 hourly Snapshot copies - 3 daily Snapshot copies	Retention: 6 Retention: 28 (4 weeks) - 6 hourly Snapshot copies - 28 daily Snapshot copies

	Executed By	Primary Storage	Secondary Storage
Log backups	SAP HANA database tools Schedule: every 15 minutes	N/A	Retention: 28 days (4 weeks)
Block integrity check	Scheduled by Snap Creator Executed by SAP HANA database Schedule: once per week	Retention: 1 The backup is overwritten when the next block integrity check is executed	

With this example, six hourly backups and three daily backups are kept at the primary storage. Database backups are retained for four weeks at the secondary storage.

2.1 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.

2.2 Capacity Requirements for Snapshot Backups

You must consider the higher block change rate on the storage layer relative to the change rate with traditional databases. Due to the 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 estimate for the change rate is 20% to 50% per day.

2.3 Snap Creator Plug-In for SAP HANA

The NetApp Snap Creator framework in combination with an application-specific plug-in for SAP HANA is used to implement a backup solution for SAP HANA based on storage-based Snapshot backups.

Snap Creator and the SAP HANA plug-in are supported with Data ONTAP® operating in 7-Mode and with clustered Data ONTAP with the SAP HANA database nodes attached to the storage controllers with either NFS or Fibre Channel. An SAP HANA database can run in a single-node or multinode configuration. 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 by 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 Server commands for the following tasks:

- Provision of database consistency to prepare a storage-based Snapshot backup
- Management of log file backup retention on the file system level
- Management of the SAP HANA backup catalog for data file and log file backups
- Execution of a file-based backup for a block integrity check

Figure 8 shows an overview of the communication paths of Snap Creator with the storage and the SAP HANA database.

Figure 8) Snap Creator and SAP HANA plug-in.
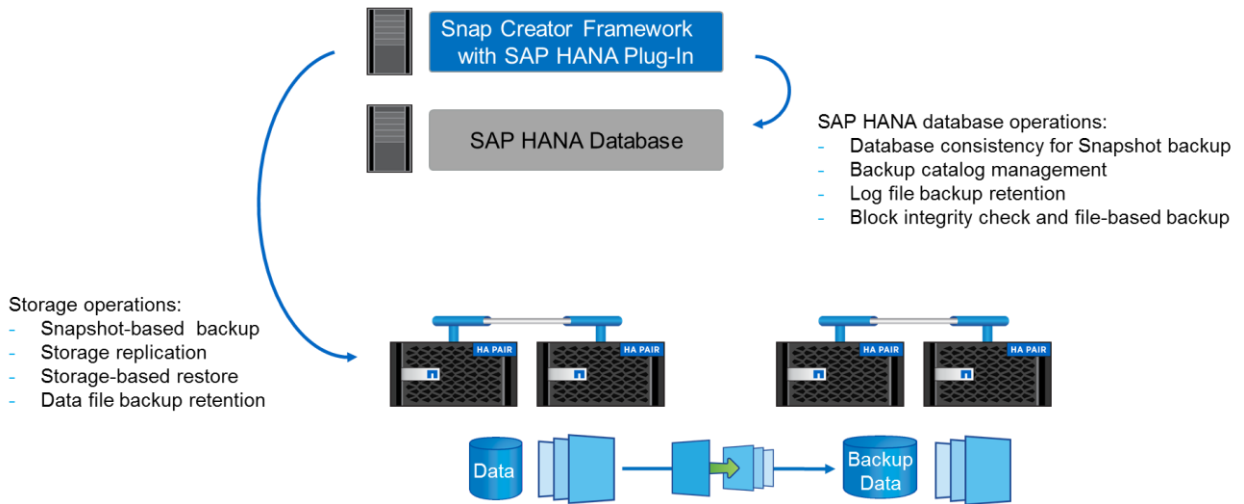
When executing a backup, Snap Creator performs the following tasks:

1. Creation of an SAP HANA backup save point to create a consistent image on the persistence layer.
2. Creation of a storage Snapshot copy of the data volume.
3. Registration of the storage Snapshot backup within the SAP HANA backup catalog.
4. Deletion of the SAP HANA backup save point.
5. Execution of a SnapVault update for the data volume.
6. Deletion of storage Snapshot copies at the primary and/or secondary storage based on the defined retention policies for backups at these locations.
7. Deletion of SAP HANA backup catalog entries if the backups do not exist anymore at the primary storage and the secondary storage.
8. Deletion of all log backups that are older than the oldest data backup on file system and within the SAP HANA backup catalog.

3 Overview of Installation and Configuration Steps

This document covers the required installation and configuration steps for NetApp storage with clustered Data ONTAP as well as for systems with Data ONTAP operating in 7-Mode. Indeed, all Snap Creator and SAP HANA Studio operations are the same with clustered Data ONTAP and Data ONTAP operating in 7-Mode.

The initial SnapVault configuration on the storage systems and all SnapVault commands that must be executed directly on the storage are different. The differences are highlighted and described in this document.

To perform installation, complete the following steps:

1. Install the SAP HANA hdbsql client software on the management server.
2. Install the Snap Creator framework on the management server.

To perform configuration, complete the following steps:

1. Create the backup user and SAP HANA user store configuration.
2. Create volumes for backup replication at the secondary storage controller.
3. Configure SnapVault relationships for database volumes.

4. Configure Snap Creator.

5. Configure log backups.

The differences between clustered Data ONTAP and Data ONTAP operating in 7-Mode are highlighted in the following sections:

- 5.2, "Initial Configuration for Data Protection to Secondary Storage"
- 8.1, "Restore and Recovery from Primary Storage by Using Volume-Based Restore"
- 8.4, "Resume SnapVault Relation After a Restore"
- 8.5, "Restore After Primary Storage Failure"

3.1 Lab Setup Used with Clustered Data ONTAP

Figure 9 shows the setup used with clustered Data ONTAP. This setup is based on a single-node SAP HANA configuration with the storage virtual machines (SVMs) and volume names shown in Figure 9.

Figure 9) Setup used with clustered Data ONTAP.
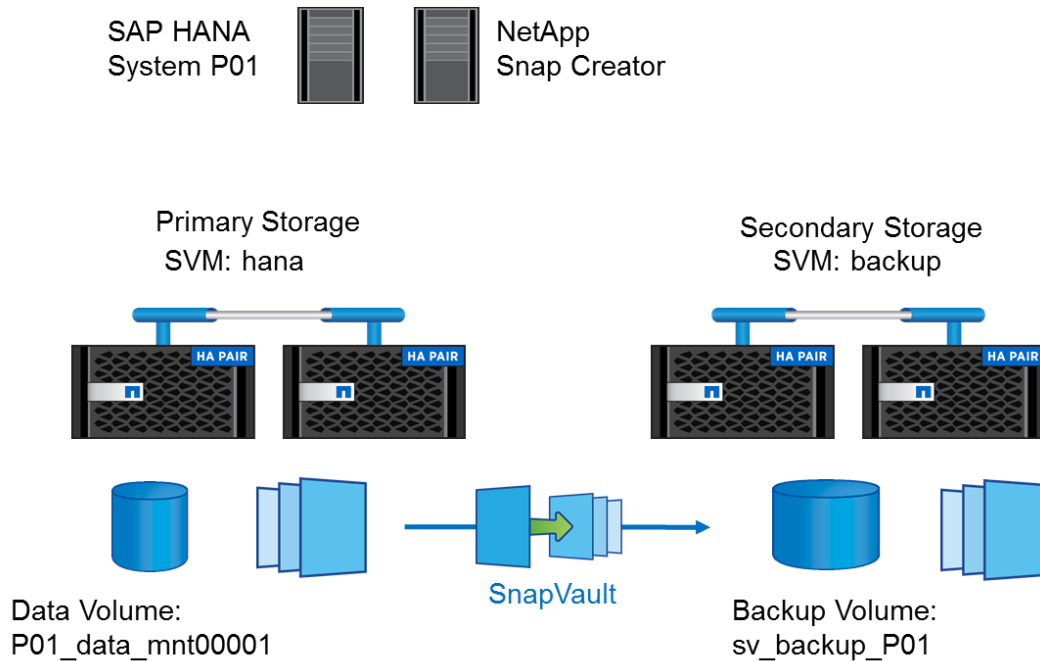
3.2 Lab Setup with Data ONTAP Operating in 7-Mode

The setup with DATA ONTAP operating in 7-Mode is based on a Cisco and NetApp SAP HANA multinode appliance.

With this example setup, the SAP HANA database runs on a three-plus-one database node configuration. All of the Snap Creator software components, the server, the agent, and the plug-in are installed on the management host, which is part of the Cisco appliance.

One NetApp 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 from another NetApp HA controller pair is used as the secondary storage. Each data volume is replicated to a dedicated backup volume on the secondary storage.

Figure 10 shows the data volumes on the primary storage and the replication path to the secondary storage.

Figure 10) Setup used with Data ONTAP operating in 7-Mode.

SAP HANA 3 plus 1 multinode setup
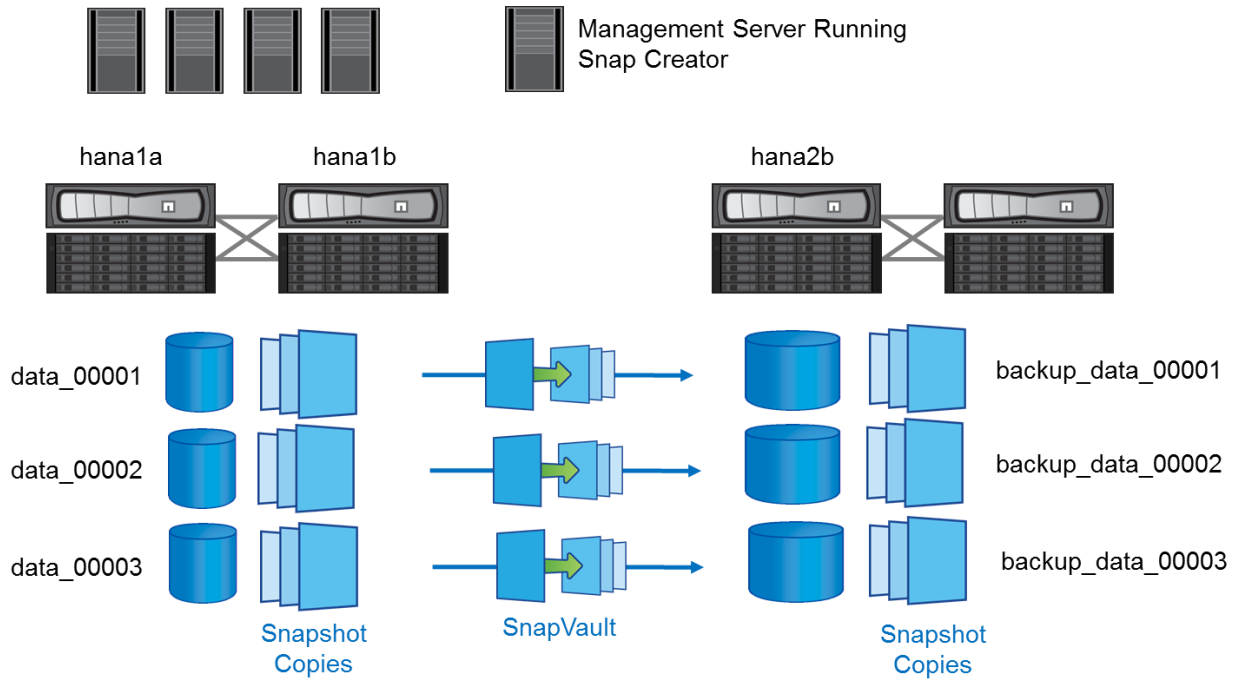
4 Installation

4.1 Installation of SAP HANA hdbsql Client Software

The Snap Creator SAP HANA plug-in uses hdbsql commands to execute commands in the SAP HANA database. The SAP HANA hdbsql client software must be installed on the host on which the Snap Creator software is installed.

```
stlrx300s8-1:/software/HANA-SPS11/51050506/DATA_UNITS/HDB_CLIENT_LINUX_X86_64 # ./hdbinst
SAP HANA Database Client installation kit detected.


SAP HANA Lifecycle Management - Client Installation 1.00.110.00.1447753075
**************************************************************************


Enter Installation Path [/usr/sap/hdbclient]:
Checking installation...
Preparing package 'Python Runtime'...
Preparing package 'Product Manifest'...
Preparing package 'SQLDBC'...
Preparing package 'REPOTOOLS'...
Preparing package 'Python DB API'...
Preparing package 'ODBC'...
Preparing package 'JDBC'...
Preparing package 'HALM Client'...
Preparing package 'Client Installer'...
Installing SAP HANA Database Client to /usr/sap/hdbclient...
Installing package 'Python Runtime'...
Installing package 'Product Manifest'...
Installing package 'SQLDBC'...
Installing package 'REPOTOOLS'...
```

```
Installing package 'Python DB API'...
Installing package 'ODBC'...
Installing package 'JDBC'...
Installing package 'HALM Client'...
Installing package 'Client Installer'...
Installation done
Log file written to '/var/tmp/hdb_client_2016-02-08_04.12.13_32307/hdbinst_client.log' on host '
stlrx300s8-1'.
stlrx300s8-1:/software/HANA-SPS11/51050506/DATA_UNITS/HDB_CLIENT_LINUX_X86_64
```

4.2   Snap Creator Framework

The following description is based on Snap Creator version 4.3.0. For additional information, see the
Snap Creator Installation Guide.

1.   Create the installation directory.

```
stlrx300s8-1:/ # mkdir -p /opt/NetApp/SnapCreator
```

2.   Download the Snap Creator software and copy it to `/opt/NetApp/SnapCreator`.

3.   Unzip and untar the file.

```
stlrx300s8-1://opt/NetApp/SnapCreator # gzip -d Linux64.tar.gz
stlrx300s8-1:/opt/NetApp/SnapCreator # tar -xvf SC4.3.0-Linux64.tar
scAgent4.3.0/
scAgent4.3.0/bin/
scAgent4.3.0/etc/
scAgent4.3.0/lib/
scAgent4.3.0/logs/
scAgent4.3.0/plugins/
scAgent4.3.0/plugins/examples/
scAgent4.3.0/plugins/examples/filesystem/
scAgent4.3.0/plugins/examples/native/
scAgent4.3.0/plugins/examples/native/bat/
scAgent4.3.0/plugins/examples/native/c++/
scAgent4.3.0/plugins/examples/native/filesystem/
scAgent4.3.0/plugins/examples/native/java/
scAgent4.3.0/plugins/examples/native/powershell/
scAgent4.3.0/plugins/examples/native/python/
scAgent4.3.0/plugins/examples/native/shell/
scAgent4.3.0/plugins/examples/perl-style/
scAgent4.3.0/plugins/java/
scAgent4.3.0/plugins/java/hana/….


Truncated
….
```

4.   Run the Snap Creator server setup.

```
stlrx300s8-1:/opt/NetApp/SnapCreator/scServer4.3.0 # ./snapcreator --setup

Welcome to the NetApp Snap Creator Framework 4.3.0!
### Installation options ###
01. NetApp Snap Creator Framework 4.3.0 Server
02. NetApp Snap Creator Framework 4.3.0 Remote CLI
Select install option (enter a number or "q" to quit): 01
END USER LICENSE AGREEMENT

This end user license agreement ("Agreement") is a contractual agreement between you ("You" or
"Your") and NetApp ("NetApp"), and provides the terms under which NetApp licenses its i)
software, including where relevant, backup and recovery, disaster recovery, storage efficiency
and management software, operating systems, protocols, updates and upgrades ("Software"), and ii)
technical documentation describing the Software ("Documentation") to You, whether supplied by
NetApp, Your NetApp distributor, reseller or partner. Any support is provided under a separate
agreement.
```

```
1. Acceptance. By downloading, installing, copying, accessing or using the Software and

…….. truncated

NetApp and You and prevails over any conflicting or additional terms in any quote, purchase
order, acknowledgment, or similar communication between the parties.

Version: 28 February 2013

Do you accept the End User License Agreement (y|n): y

Enter controller serial number (Recommended): 12345678

Enter Snap Creator server port [8443]:

Enable job monitor (Y|N): y

Enter job monitor size, how many jobs to allow [100]:
Enter scServer Administrator Username: admin
Enter password for admin:
Confirm password for admin:

INFO: Updated NetApp Snap Creator Framework 4.3.0
/opt/NetApp/SnapCreator/scServer4.3.0/engine/etc/snapcreator.properties
INFO: Updated NetApp Snap Creator Framework 4.3.0
/opt/NetApp/SnapCreator/scServer4.3.0/bin/scServer

INFO: To start scServer please do the following:

/opt/NetApp/SnapCreator/scServer4.3.0/bin/scServer start

INFO: To access NetApp Snap Creator Framework 4.3.0 GUI goto https://hostname:8443
```

5. Start the Snap Creator server.

```
stlrx300s8-1:/opt/NetApp/SnapCreator/scServer4.3.0/bin # ./scServer start
Starting scServer:
Checking Status of scServer:
Running
stlrx300s8-1:/opt/NetApp/SnapCreator/scServer4.3.0/bin #
```

6. Run the Snap Creator agent setup.

```
stlrx300s8-1:/opt/NetApp/SnapCreator/scAgent4.3.0 # ./snapcreator --setup

Welcome to the NetApp Snap Creator Framework 4.3.0!
### Installation options ###
01. NetApp Snap Creator Framework 4.3.0 Agent
Select install option (enter a number or "q" to quit): 01
END USER LICENSE AGREEMENT

This end user license agreement ("Agreement") is a contractual agreement between you ("You" or
"Your") and NetApp ("NetApp"), and provides the terms under which NetApp licenses its i)

…. Truncated

order, acknowledgment, or similar communication between the parties.




Version: 28 February 2013

Do you accept the End User License Agreement (y|n): y

Enter Agent Port [9090]:
INFO: Updated NetApp Snap Creator Framework 4.3.0
/opt/NetApp/SnapCreator/scAgent4.3.0/etc/agent.properties

INFO: To start scAgent please do the following:
```

```
/opt/NetApp/SnapCreator/scAgent4.3.0/bin/scAgent start
```

7. Start the agent.

```
stlrx300s8-1:/opt/NetApp/SnapCreator/scAgent4.3.3/bin # ./scAgent start
Starting scAgent:
    Watchdog: Running
       Agent: Running
stlrx300s8-1:/opt/NetApp/SnapCreator/scAgent4.3.3/bin #
```

5 Configure Data Backups

5.1 SAP HANA Backup User and Hdbuserstore Configuration

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

Figure 11 shows a screenshot of the SAP HANA Studio on which the backup user can be created. The backup user must have the privileges BACKUP ADMIN and CATALOG READ.

Figure 11) Database user for SAP HANA backups.
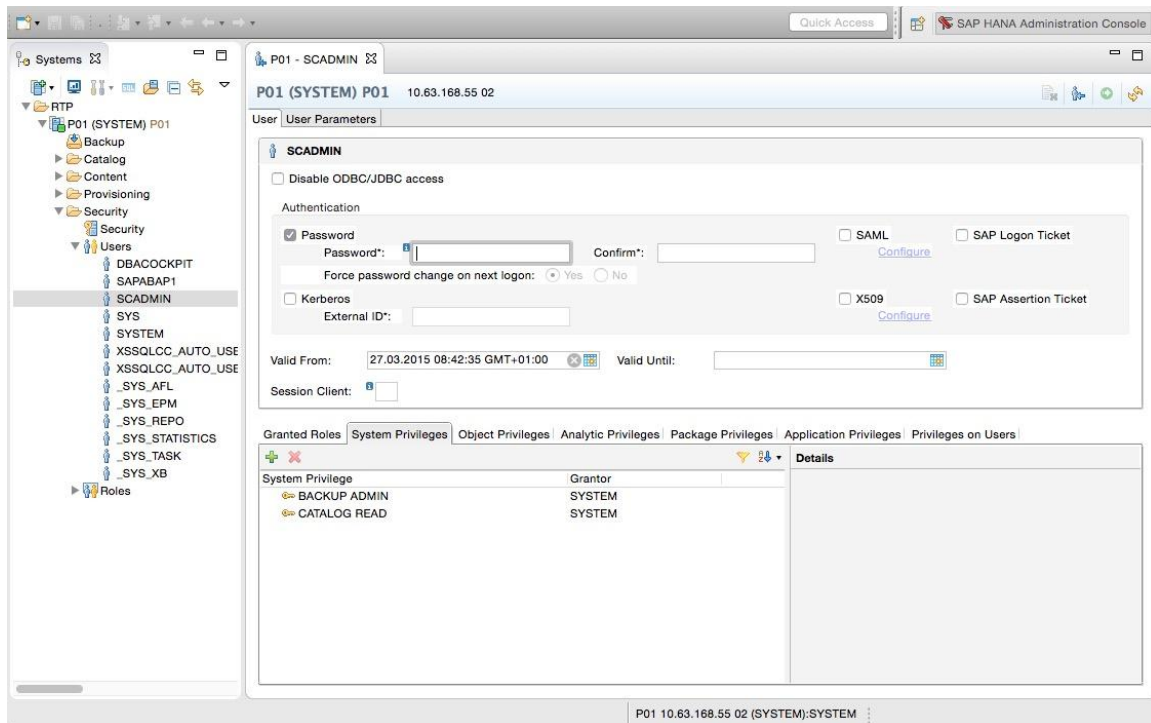
At the administration host on which Snap Creator was installed, a user store key is configured for all database hosts that belong to the SAP HANA database. The user store key is configured with the OS root user.

```
hdbuserstore set <key> <host><3[instance]15> <user> <password>
```

In the example configuration, only one key must be configured for the single-node HANA database host.

```
stlrx300s8-1:/usr/sap/hdbclient # ./hdbuserstore set SCADMINKEY stlrx300s8-3:30215 SCADMIN
Password
```

```
stlrx300s8-1:/usr/sap/hdbclient # ./hdbuserstore LIST
DATA FILE       : /root/.hdb/stlrx300s8-1/SSFS_HDB.DAT
KEY FILE        : /root/.hdb/stlrx300s8-1/SSFS_HDB.KEY

KEY SCADMINKEY
  ENV : stlrx300s8-3:30215
  USER: SCADMIN
stlrx300s8-1:/usr/sap/hdbclient #
```

5.2   Initial Configuration for Data Protection to Secondary Storage

Create Volumes at Secondary Storage

All volumes to be backed up must be created at the secondary storage controller or SVM. In our example, the volume `sv_backup_P01` is created at the secondary storage.

SnapVault Configuration in Clustered Data ONTAP

Create SVM peering between the primary SVMs and the backup SVMs. If the backup SVM resides on a different cluster, then cluster peering must be performed as well. Consult the [Data ONTAP documentation](#) for details about cluster and SVM peering.

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

1.   Create the SnapMirror policy.

```
stl-cmode-INC10364933::> snapmirror policy create -vserver backup -policy HANA_backup_policy
```

2.   Add a rule to SnapMirror policy.

```
stl-cmode-INC10364933::> snapmirror policy add-rule -vserver backup -policy HANA_backup_policy -
snapmirror-label P01_HANA -keep 999
```

> **Note:**   SnapMirror policy configuration requires a SnapMirror label and a SnapMirror retention policy. Although a SnapMirror label is set with the SnapMirror configuration, Snap Creator is configured so that the SnapMirror label is not used, and the retention of secondary backups is managed by Snap Creator.

3.   Create the SnapMirror relationship.

```
stl-cmode-INC10364933::> snapmirror create -source-path hana:P01_data_mnt00001 -destination-path
backup:sv_backup_P01 -type XDP -policy HANA_backup_policy
Operation succeeded: snapmirror create the relationship with destination backup:sv_backup_P01.
```

4.   Initialize the SnapMirror relationship.

```
stl-cmode-INC10364933::> snapmirror initialize -destination-path backup:sv_backup_P01 -type XDP
```

> **Note:**   In an SAP HANA multinode setup, there are multiple data volumes, one for each database node. Execute the `snapmirror create` and `snapmirror initialize` commands for each volume.

Configuring SnapVault Relationships in Data ONTAP Operating in 7-Mode

1.   SnapVault and NDMP must be enabled 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. Access to the secondary storage controller must be configured on all primary storage controllers.

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

> **Note:** NetApp recommends using a dedicated network for replication traffic. The host name of this interface at the secondary storage controller must be configured. Rather than `hana2b`, consider using the host name `hana2b-rep`.

3. Configure access for all primary storage controllers on the secondary storage controller.

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

> **Note:** NetApp recommends using a dedicated network for replication traffic. The host name of this interface at the primary storage controllers must be configured. Rather than `hana1b` and `hana1a`, consider using the host names `hana1a-rep` and `hana1b-rep`.

4. To start the SnapVault relationship, run the following commands at the secondary storage:

```
hana2b> snapvault start -S hana1a:/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 hana1a:/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 hana1b:/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>
```

> **Note:** NetApp recommends using a dedicated network for replication traffic. The host name of this interface at the primary storage controllers must be configured. Rather than `hana1b` and `hana1a`, consider using the host names `hana1a-rep` and `hana1b-rep`.

5.3 Snap Creator Configuration

To configure Snap Creator, complete the following steps:

1. Connect to the Snap Creator GUI.
2. Log in with the user and password that were configured during the installation. Click Sign In.

NetApp Snap Creator Framework

Server	10.63.168.51
Port	8443
User Name	admin
Password	••••••••

Sign In

3. Enter a profile name and click OK. In our example, P01 is the system identifier (SID) of the database.
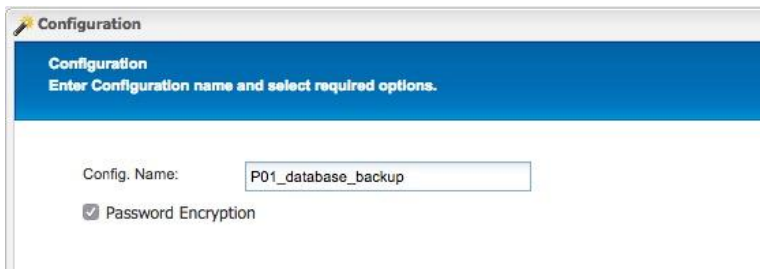
New Profile

Enter new profile name: HANA_profile_P01

OK   Cancel

4. Enter the configuration name and click Next.
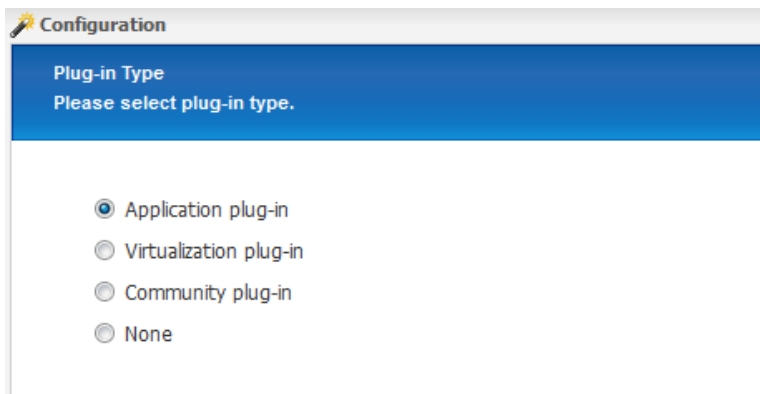
Configuration

Configuration
Enter Configuration name and select required options.

Config. Name: P01_database_backup

☑ Password Encryption

5. Select Application Plug-In.

Configuration

Plug-in Type
Please select plug-in type.

⦿ Application plug-in
○ Virtualization plug-in
○ Community plug-in
○ None

6. Select SAP HANA as the application plug-in and click Next.
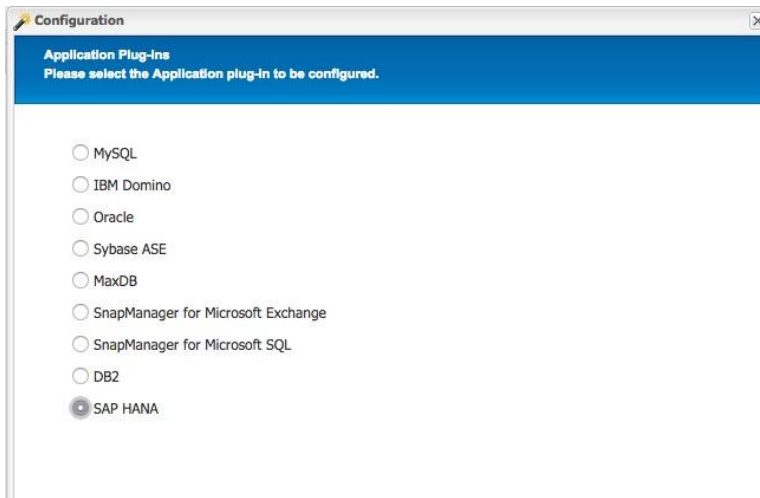
7. Enter the database SID, the hdbuserstore keys for each SAP HANA node, the path to the hdbsql executable, and the OSDB user. You can also enable automated LOG cleanup. Click Next.
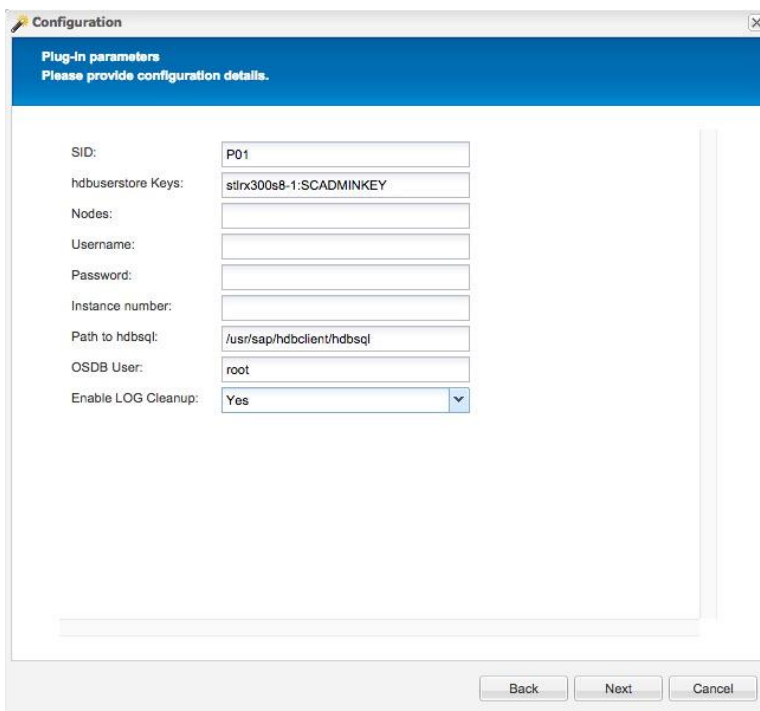
8. Enter the file backup location and the file backup prefix and select if you want to enable file backup. Click Next.

Configuration

File-Based Backup Configuration Details
Provide File-Based Backup Details

File-Backup Location:	/mnt/filebased/P01
File-Backup prefix:	BACKUP-P01
Enable File-Backup:	☐

9. Enter the temporary file backup location for the block integrity check operation and select Enable DB Integrity Check. Click Next.

Configuration

Integrity Check Configuration Details
Provide Integrity Check Details

Temporary File-Backup Location:	/mnt/filebased/TMP
Enable DB Integrity Check:	☐

10. Enter the agent configuration parameter. Click Next.

Configuration ⊠

Agent Configuration
Enter agent configuration details

IP/DNS:	10.63.168.51
Port:	9090
Timeout (secs):	300

Test agent connection

11. Enter the storage connection settings and click Next.

Configuration

Storage Connection Settings
Please Provide Storage Connection Settings

Use OnCommand Proxy:	☐
Transport:	HTTPS ▾
Controller/Vserver Port:	443

12. Enter the storage login credentials and click Next.
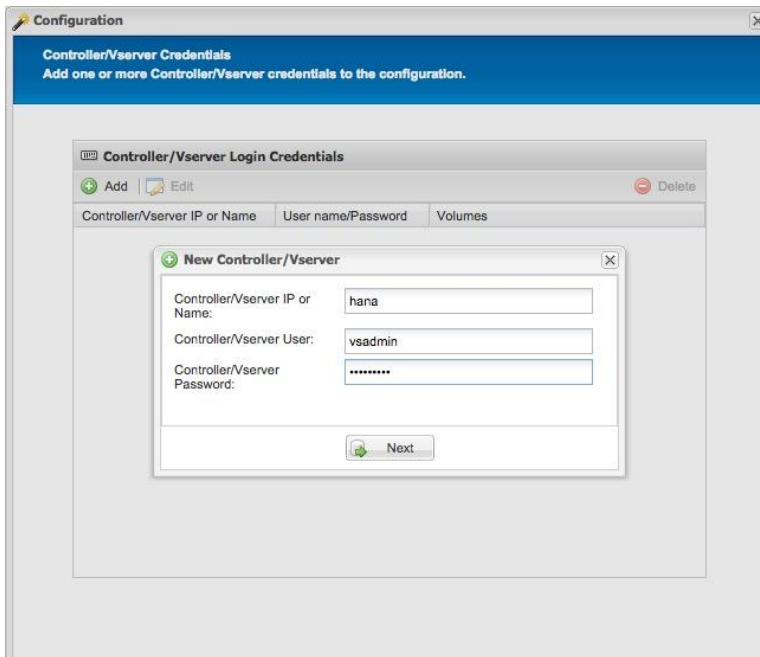
13. Select the data volumes that are stored on this SVM and click Save.
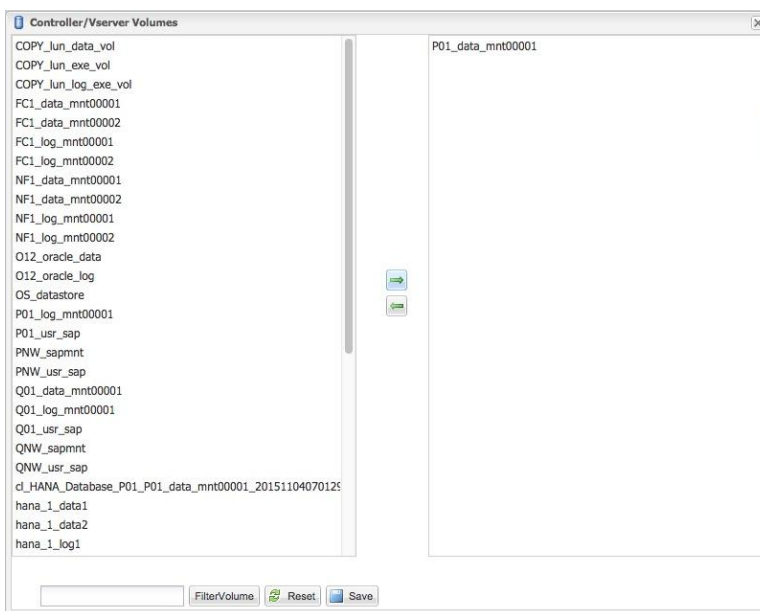
14. The next screen shows the SVM and volumes for our example configuration. Click Next.

> **Note:** In an SAP HANA multinode setup, there are typically volumes on different controllers. Click Add to add another controller.

15. Enter the Snapshot policy and retention configuration. The retention of 3 daily and 12 hourly Snapshot copies is presented as an example. Your system can be configured differently depending on your requirements.

> **Note:** A Snapshot copy label is not required and must not be configured to make sure that Snap Creator is able to control backup retention at the secondary storage. If you were to set a label, clustered Data ONTAP would also control backup retention and would delete backups based on the retention defined in the SnapMirror relationship.

> **Note:** You must select the naming convention Timestamp. The naming convention Recent is not supported with the SAP HANA plug-in because the timestamp of the Snapshot copy is also used for SAP HANA backup catalog entries.

16. No changes are needed. Click Next.
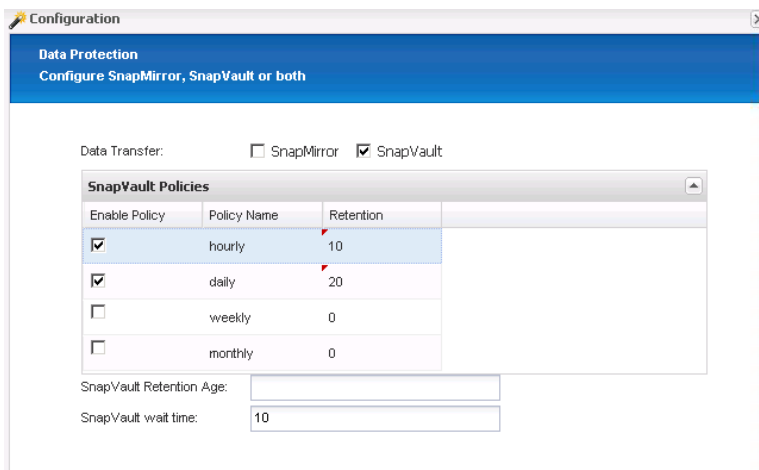
17. Select SnapVault and configure the SnapVault retention policies and the SnapVault wait time.

18. Click Add.

19. Select a source SVM from the list and click Next.
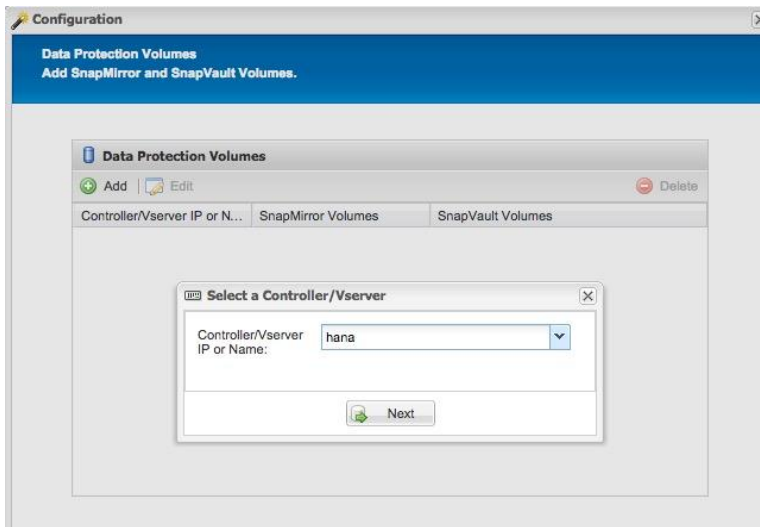
20. Select all volumes and click Save.
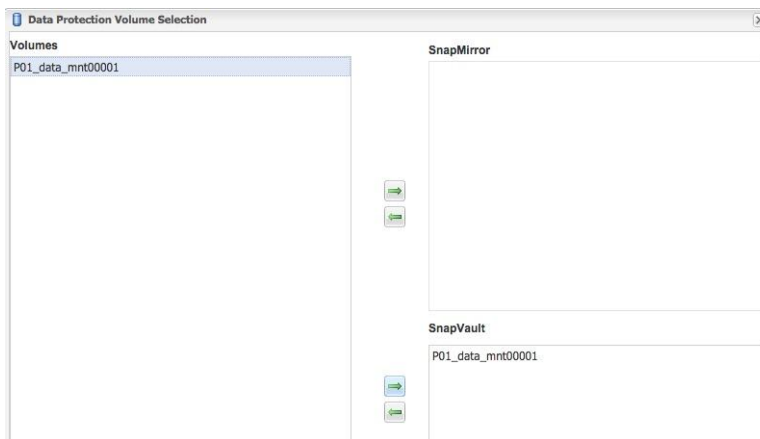
21. The next screen shows all volumes that are protected in our example configuration. Click Next.
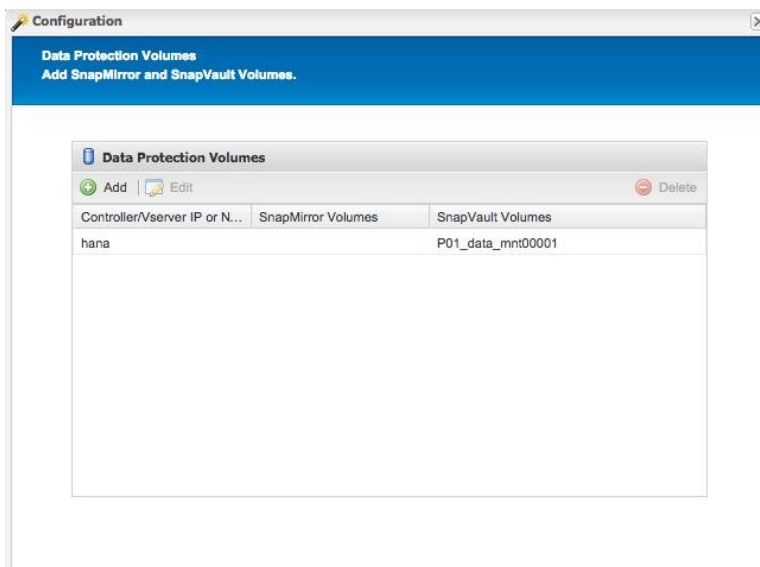
22. Enter the credentials for the target SVM and click Next. In this example, the `vsadmin` user is used to access the SVM. Typically, a dedicated backup user is configured on the storage system and then used with Snap Creator.

23. Click Next.

24. After the configuration is completed, click Finish.
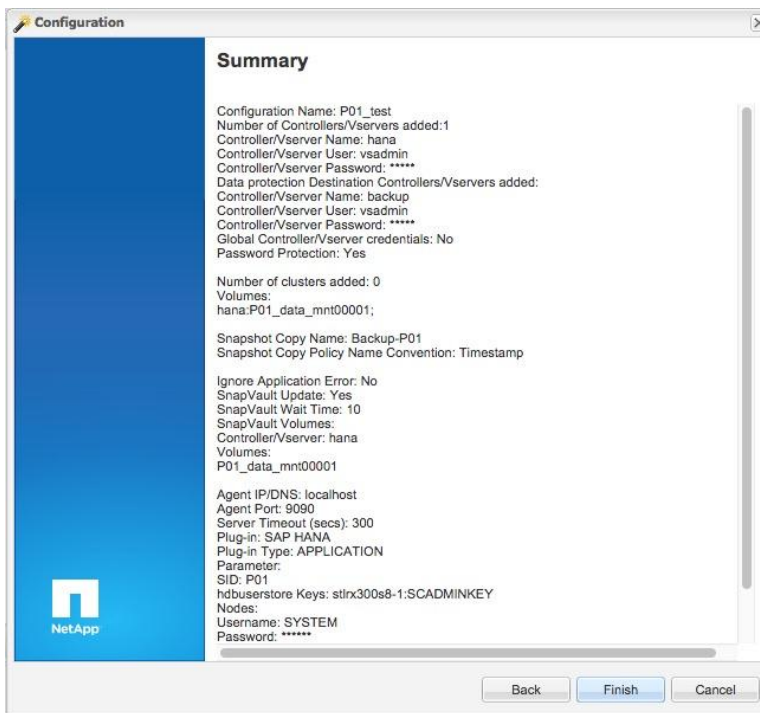
25. Go to the SnapVault Settings tab.
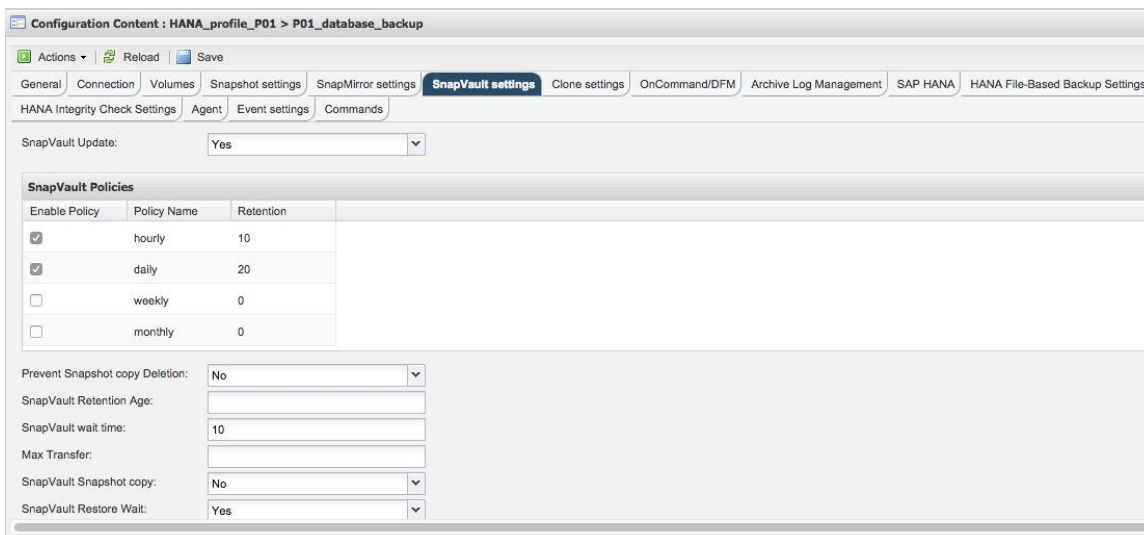
26. Set the retention policies for backups at the secondary SVM and set SnapVault Restore Wait to Yes. Click Save.

Configuration of Dedicated Replication and Management Interfaces

With Data ONTAP operating in 7-Mode, NetApp recommends a dedicated network for replication traffic. You must 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=hana1a:hana1a-rep/hana2b;hana1b:hana1b-rep/hana2b
MANAGEMENT_INTERFACES=hana2b:hana2b-mgmt
```

5.4 Additional Configuration Steps for Fibre Channel Environments

Additional configuration steps are required in environments in which the SAP HANA systems are connected with FCP to the storage controllers.

When a global synchronized backup save point is triggered by Snap Creator within SAP HANA, SAP HANA writes the file `/hana/data/SID/mnt00001/hdb00001/snapshot_databackup_0_1` as a last step. 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 a situation in which the backup is restored. Due to metadata caching with the XFS file system on the Linux host, the file is not immediately visible at the storage layer. The standard XFS configuration for metadata caching is 30 seconds.

Within Snap Creator, you must configure a postapplication quiesce command that waits until the XFS metadata cache is flushed to the disk layer.

The actual configuration of the metadata caching can be checked by using the following command:

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

In the configuration file at `/opt/NetApp/SnapCreator/scServer4.3.0/engine/configs`, the command `/bin/sleep` must be added in the Post Commands section, as is shown in the following example. NetApp recommends using a wait time that is twice the value of the `fs.xfs.xfssyncd_centisecs` parameter. Because the default value is 30 seconds, set the sleep command to 60 seconds.

```
#############################################################################################
#######################
#     Post Commands
#############################################################################################
#######################
POST_NTAP_DATA_TRANSFER_CMD01=
POST_APP_QUIESCE_CMD01=/bin/sleep 60
POST_CLONE_CREATE_CMD01=
```

At the Snap Creator host, the Sleep command must be added to the configuration file `/opt/NetApp/SnapCreator/scAgent4.3.0/etc/allowed_commands.config`.

```
command: /bin/sleep 60
```

6 Configure Log Backups

Store log backups on a storage system different from the primary storage. The storage system that is used for the data backup can also be used for the log backup. Configure a volume at the secondary storage to hold the log backups. Make sure that automatic Snapshot copies are switched off for this volume.

The volume is mounted at each database node: for example, the `/etc/fstab` entry at the database nodes.

```
backup:/log_backup/P01 /mnt/log_backup/P01 nfs
rw,bg,vers=3,hard,timeo=600,rsize=65536,wsize=65536,actimeo=0,noatime   0   0
```

Within SAP HANA Studio, the log backup destination is configured as is shown in Figure 12.

Figure 12) Configuration of log backup destination.
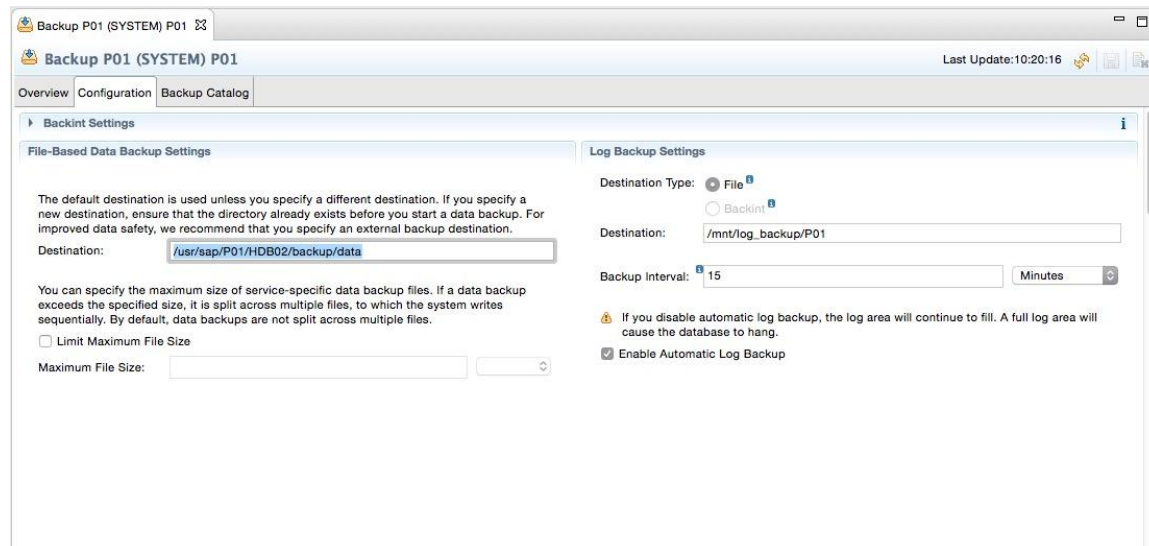
6.1 Housekeeping of Log Backups

The SAP HANA concept for log backup housekeeping is based on a function within HANA Studio or a SQL Server statement that allows the deletion of all backups that are older than a selected backup. Snap Creator handles the housekeeping of data backups (Snapshot copies) by deleting Snapshot copies on the primary or secondary storage according to a defined retention policy and by deleting the corresponding entries within the HANA catalog. Log backups that are older than the latest data backup are not required anymore, so they can also be deleted.

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

1. Snap Creator reads the backup catalog and gets the backup ID of the oldest successful data or Snapshot backup.

2. Snap Creator deletes all backups that are older than the oldest backup.

    Note: Snap Creator only handles housekeeping for backups based on Snapshot copies. If additional file-based backups are created, the user must make sure 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 log backups are not deleted before this file-based backup is deleted.

Automated log backup housekeeping can be enabled or disabled within the SAP HANA tab of the database configuration.

Figure 13) Enable log cleanup.
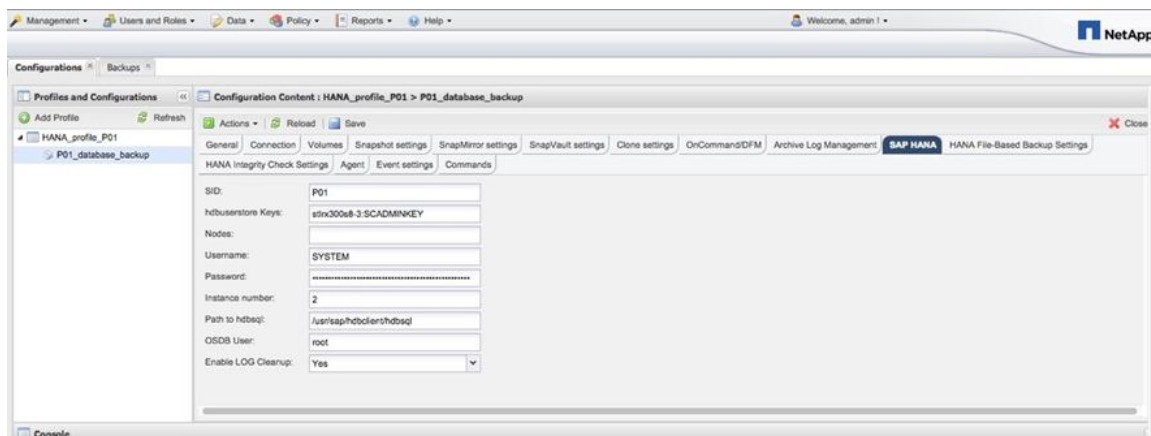
7   Execute Database Backups

A database backup can be performed by using either the Snap Creator GUI or the command line. The built-in scheduler in the Snap Creator GUI or the command line in combination with an external scheduler can be used to schedule the backups.

7.1   Backup Workflow

When Snap Creator is backing up the database, it performs the following steps:

1. It creates a global synchronized backup save point to create a consistent database image on the persistence layer.
2. It creates storage Snapshot copies for all data volumes of the database. In our example of a single-node HANA database, there is only one data volume. With an SAP HANA multinode database, there are multiple data volumes.
3. It registers the storage Snapshot backup within the SAP HANA backup catalog.
4. It deletes the SAP HANA backup save point.
5. It starts a SnapVault update for all data volumes (if configured and enabled).
6. It checks the SnapVault status and waits until it is finished or a configurable timeout has occurred.
7. It deletes storage Snapshot copies and deletes backups in the SAP HANA backup catalog based on the defined retention policy for backups at the primary and secondary storage.
8. It deletes all log backups that are older than the oldest data backup on the file system and within the SAP HANA backup catalog. This step is only executed if Log Backup Cleanup is enabled.

7.2   Database Backup with Snap Creator GUI

1. Click the `P01_database_backup` configuration. From the Actions menu, select Backup.
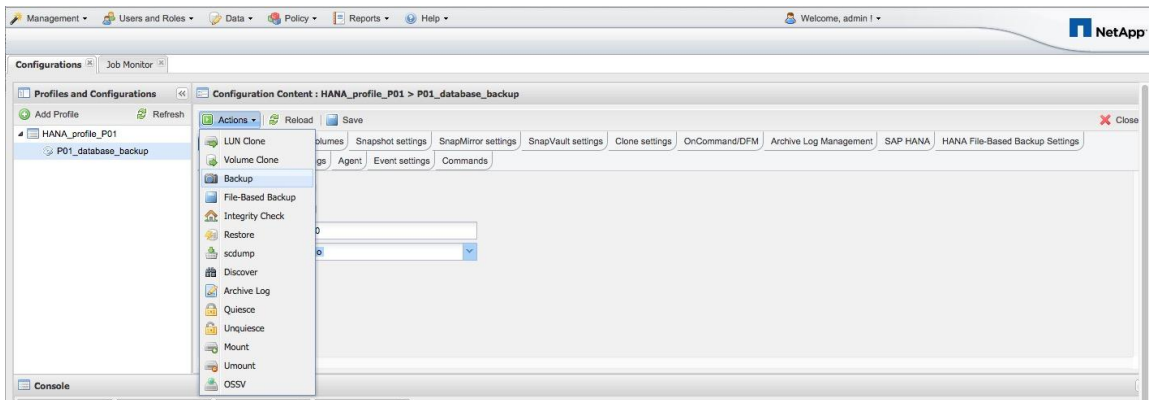
2. Select the backup policy and click OK.
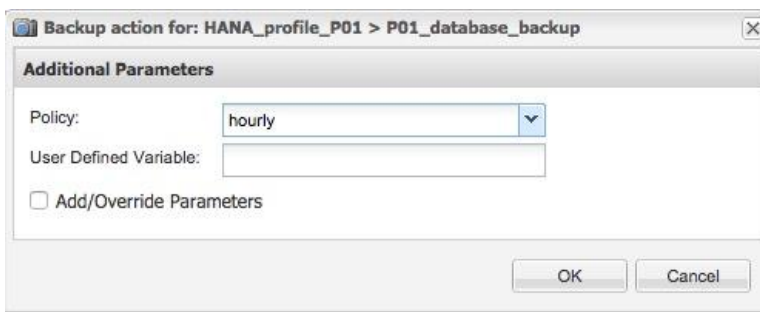
3. The action is started.
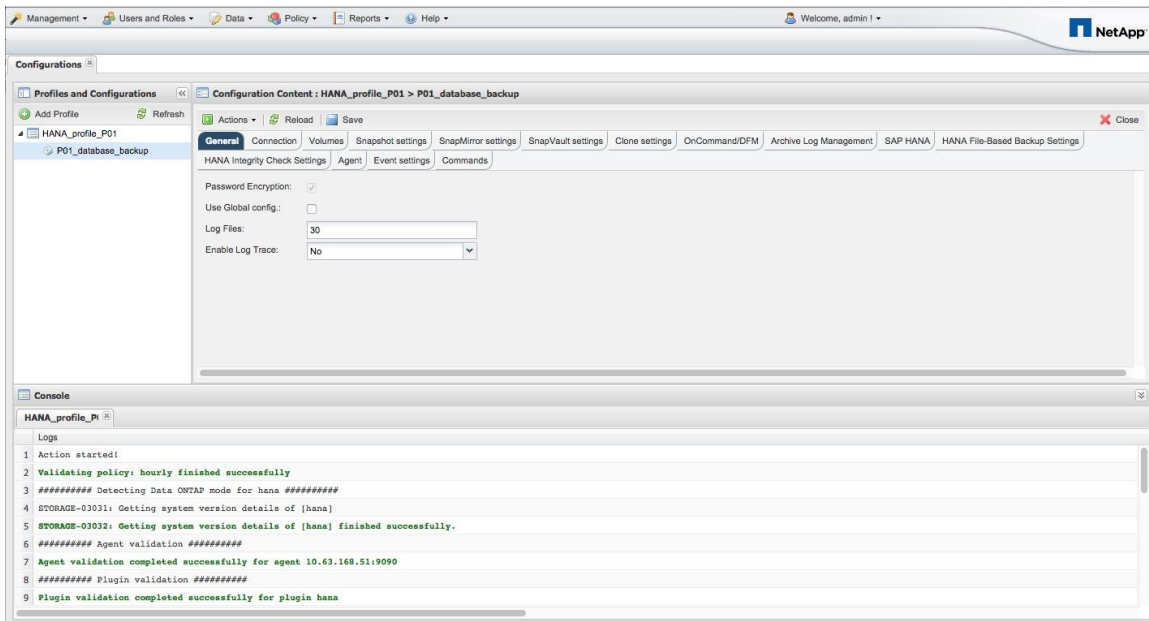
4. Snap Creator triggers the SnapVault update, and Snap Creator waits until the data is replicated to the secondary storage. The wait time is set during configuration and can be adapted in the SnapVault Settings tab. Snap Creator triggers the SnapVault updates in parallel for each volume on the same storage controller and in sequence for each storage controller.

7.3 Database Backup with Snap Creator CLI

To back up the database, run the following command:

```
 stlrx300s8-1:~ # /opt/NetApp/SnapCreator/scServer4.3.0/snapcreator --server localhost --port
8443 --user admin --passwd Netapp123 --profile HANA_profile_P01 --config P01_database_backup --
action backup --policy daily --verbose

 [Tue Feb 9 04:46:42 2016] INFO: Validating policy: daily finished successfully

########## Detecting Data ONTAP mode for hana ##########
[Tue Feb  9 04:46:47 2016] INFO: STORAGE-03031: Getting system version details of [hana]
[Tue Feb  9 04:46:47 2016] INFO: STORAGE-03032: Getting system version details of [hana] finished
successfully.

########## Agent validation ##########
[Tue Feb 9 04:46:43 2016] INFO: Agent validation completed successfully for agent
10.63.168.51:9090

########## Plugin validation ##########
[Tue Feb 9 04:46:43 2016] INFO: Plugin validation completed successfully for plugin hana
[Tue Feb 9 04:46:43 2016] INFO: Application auto discovery is not enabled skipping .

########## Running storage discovery ##########

… Truncated
```
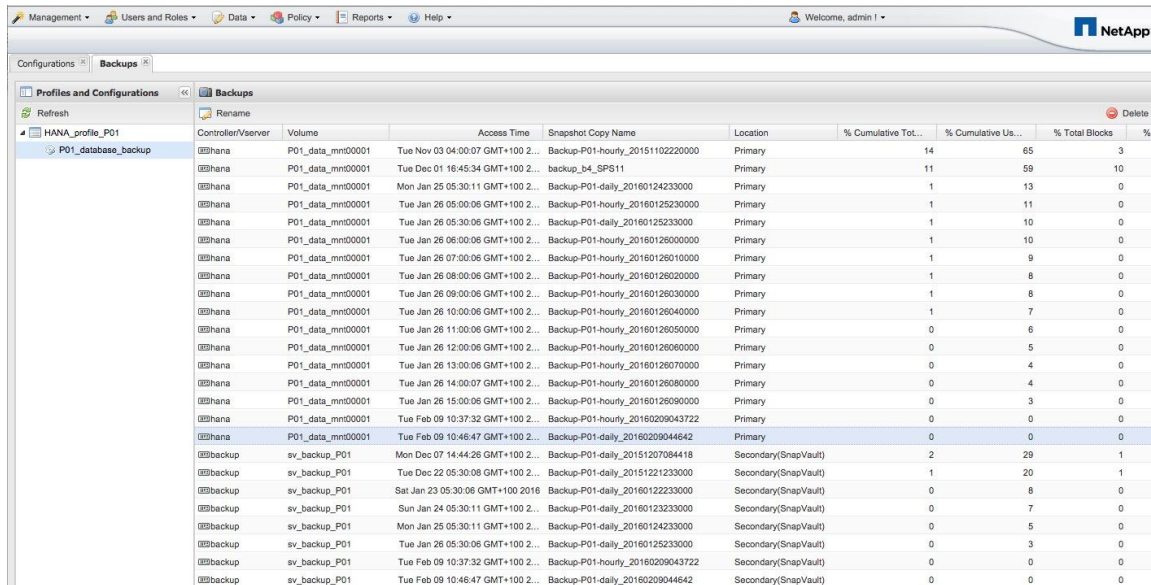
7.4 Backups in SAP HANA Studio

Figure 14 shows a list of backups within the Snap Creator GUI. The highlighted backup shows a Snapshot copy named `Backup-P01-daily_20160209044642`. This backup includes the Snapshot copy for the data volume of the SAP HANA system. This backup is also available at the secondary storage.

Figure 14) List of backups within Snap Creator.

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 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 is shown in Figure 15.

Figure 15) Backup catalog within SAP HANA Studio.
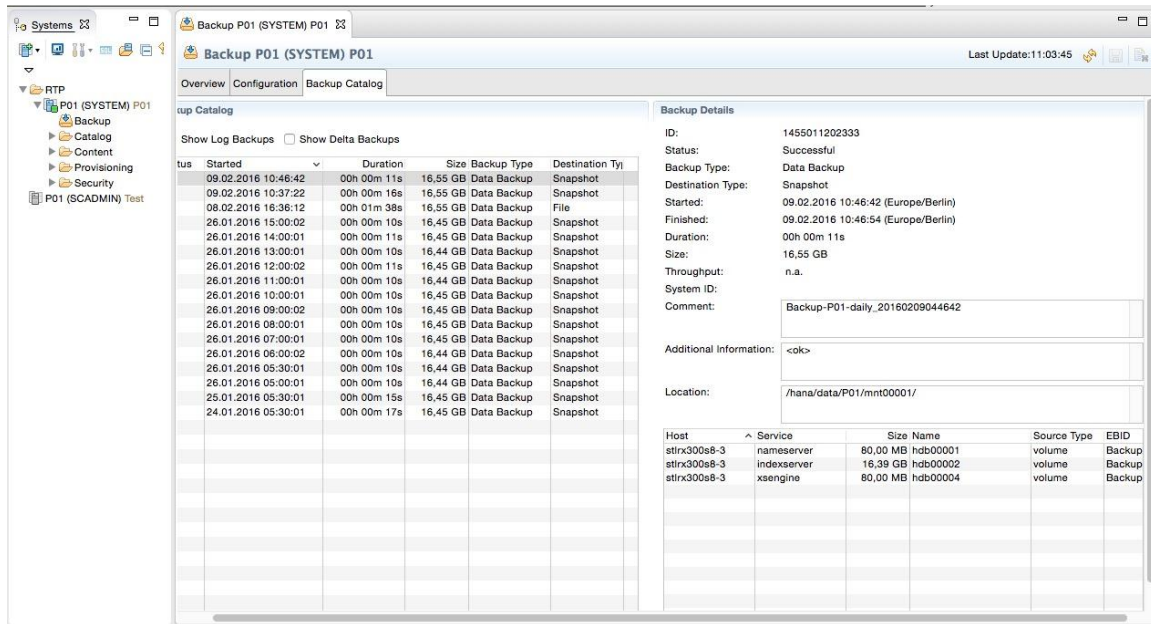
With every backup run, Snap Creator deletes Snapshot backups at the primary storage and at the secondary storage based on the retention policies defined for the different schedules (hourly, daily, and so on).

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

8 Restore and Recover

To restore and recover an SAP HANA database by using SAP HANA Studio and Snap Creator, complete the following steps:

1. To prepare the restore and recovery process with SAP HANA Studio, complete the following steps:
   a. Select Recover System and confirm shutdown of the SAP HANA system.
   b. Select the recovery type and the log backup location.
   c. The list of data backups is shown. Select Backup to see the external backup ID.
2. To perform the restore process with Snap Creator, complete the following steps:
   a. Select Restore from primary or secondary storage (volume or file based).
   b. Select the Snapshot backup that matches the external backup ID from SAP HANA Studio.
   c. Start the restore process.

   Note: If a volume-based restore from primary storage is chosen, the data volumes must be unmounted from all SAP HANA database hosts before the restore and mounted again when the restore process is finished.

   Note: In an SAP HANA multinode setup with Fibre Channel, the unmount and mount operation is executed by the SAP HANA name server as part of the shutdown and startup process of the database.

Note: The SnapVault relation must be deactivated if a volume-based restore from primary storage is chosen with clustered Data ONTAP and if you must restore a Snapshot copy that is older than the Snapshot copy currently used as the base Snapshot copy for SnapVault replication.

3. To run the recovery process with SAP HANA Studio, complete the following steps:

   a. Click Refresh from the backup list and select the available backup for recovery (green item).

   b. Start the recovery process. When the recovery process is finished, the SAP HANA system is started.

8.1 Restore and Recovery from Primary Storage by Using Volume-Based Restore

To restore and recover from primary storage, complete the following steps:

1. Within SAP HANA Studio, select the Recover option for the SAP HANA system.
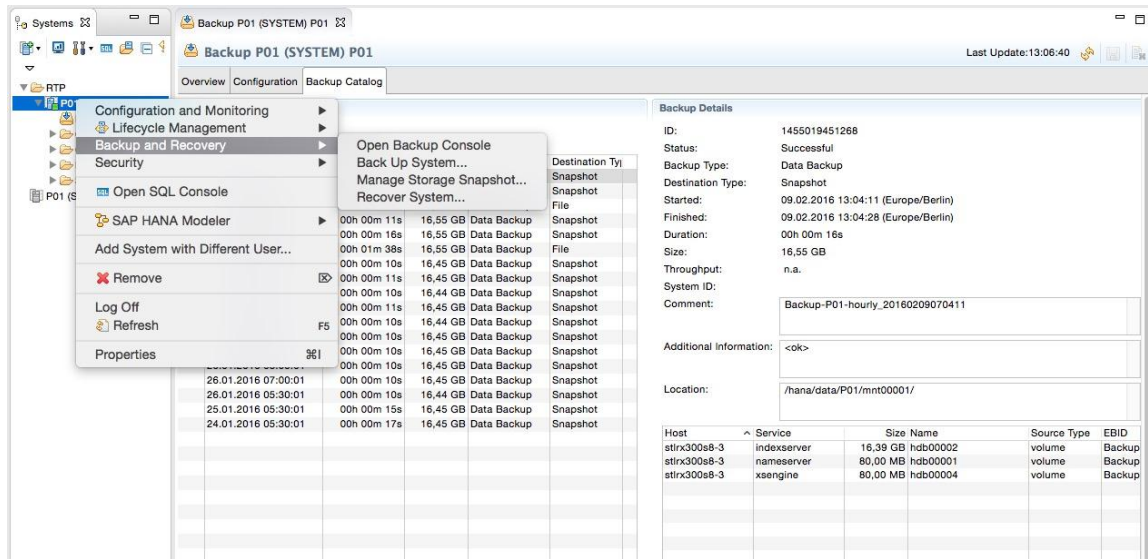
2. Click OK to shut down the SAP HANA database.
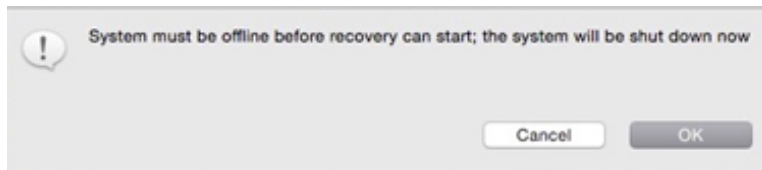
3. Select the recovery type and click Next.

4. Provide log backup locations and click Next.

5. A list of available backups is shown based on the content of the backup catalog. Choose the required backup and write down the external backup ID.

6. Deactivate the SnapVault relationship with the following commands on the backup cluster console:

   Note: This step is required only with clustered Data ONTAP and if you need to restore a Snapshot copy that is older than the Snapshot copy currently used as the base for SnapVault replication.

   a. Quiesce the SnapMirror relationship.

```
stl-cmode-INC10364933::> snapmirror quiesce -destination-path backup:sv_backup_P01
Operation succeeded: snapmirror quiesce for destination "backup:sv_backup_P01".
```

   b. Delete the SnapMirror relationship.

```
stl-cmode-INC10364933::> snapmirror delete -destination-path backup:sv_backup_P01
Operation succeeded: snapmirror delete for the relationship with destination
"backup:sv_backup_P01".
```

   c. Release the SnapMirror relationship.

```
stl-cmode-INC10364933::> snapmirror release -destination-path backup:sv_backup_P01

Warning: Snapshot copies on source volume "hana:P01_data_mnt00001" generated by SnapMirror for
the purpose of mirroring to destination
        volume "backup:sv_backup_P01" will be deleted. Once these Snapshot copies are deleted,
it will likely not be possible to
        re-establish a mirroring relationship between these two volumes.
Do you want to continue? {y|n}: y
[Job 8214] Job succeeded: SnapMirror Release Succeeded
```

7. On each database node, unmount all data volumes. In our example, only one volume must be unmounted on the single database node.

```
stlrx300s8-3:/ # umount /hana/data/P01/mnt00001
```

   Note: In an SAP HANA multinode setup with Fibre Channel, the unmount operation is executed by the SAP HANA name server as a part of the shutdown process.

8. From the Snap Creator GUI, select the configuration of the SAP HANA system. Go to Actions and select Restore.
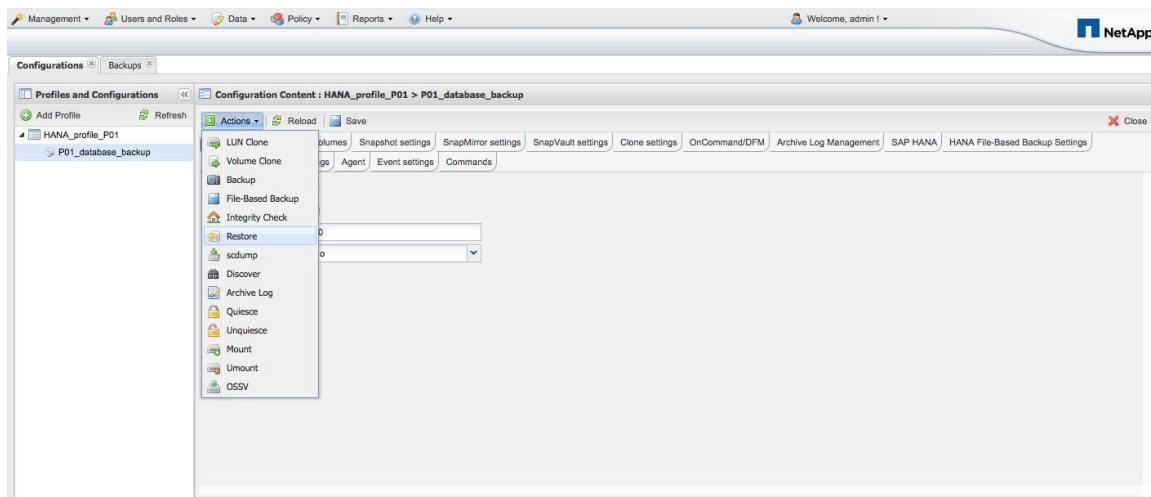
9. Click Next.
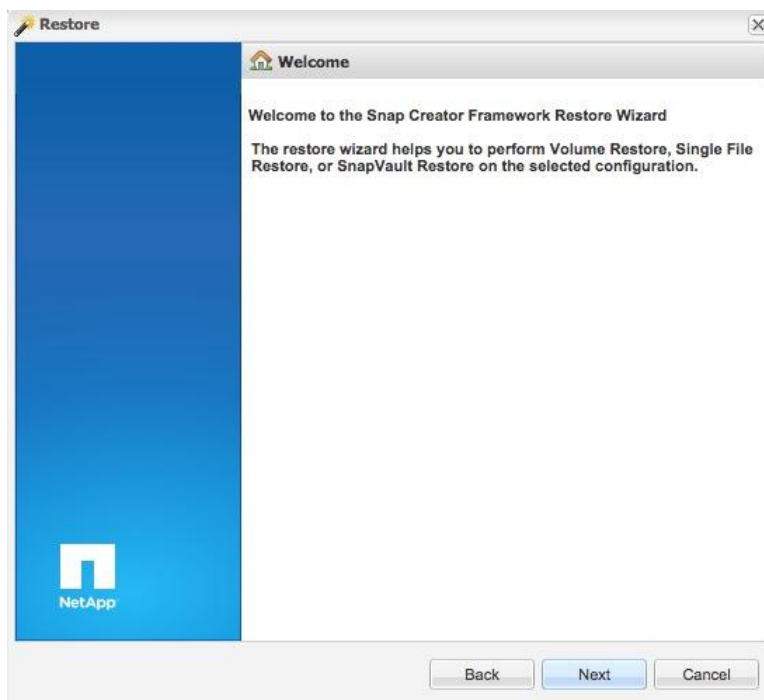
10. Select to restore from the primary storage.

11. Select the SVM, the volume name, and the Snapshot name. The Snapshot name correlates with the backup ID that has been selected within SAP HANA Studio.
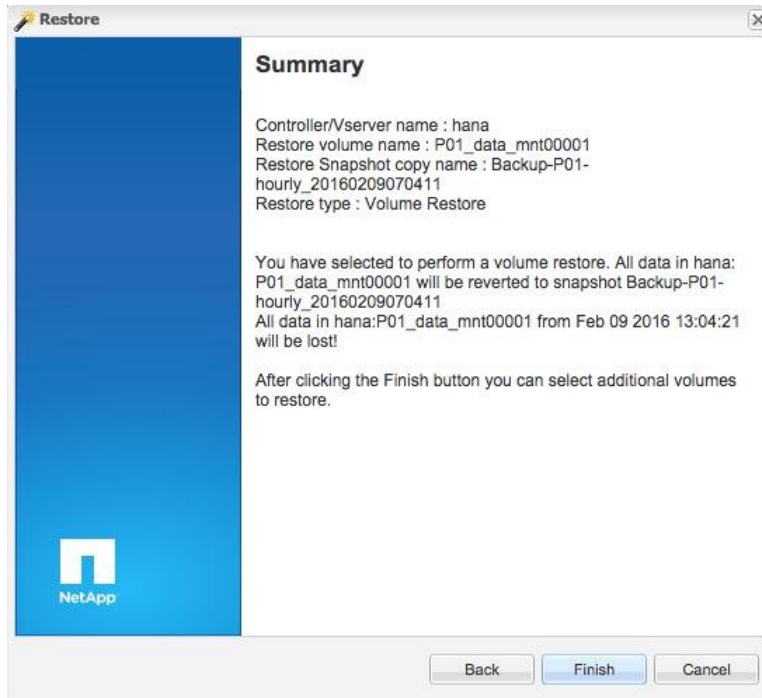

12. Click Finish.

Restore — Summary

Controller/Vserver name : hana
Restore volume name : P01_data_mnt00001
Restore Snapshot copy name : Backup-P01-hourly_20160209070411
Restore type : Volume Restore

You have selected to perform a volume restore. All data in hana: P01_data_mnt00001 will be reverted to snapshot Backup-P01-hourly_20160209070411
All data in hana:P01_data_mnt00001 from Feb 09 2016 13:04:21 will be lost!

After clicking the Finish button you can select additional volumes to restore.

Back   Finish   Cancel

13. Select Yes to add more restore items; in our example, there is only one volume.

   Note: In an SAP HANA multinode setup, there are multiple volumes that must be selected for the restore process.

Warning

Would you like to add more restore items on this configuration? Select 'No' to start the entire restore operation or 'Yes' to add more restore items.

Yes   No

14. When all volumes are selected, click OK to start the restore process.
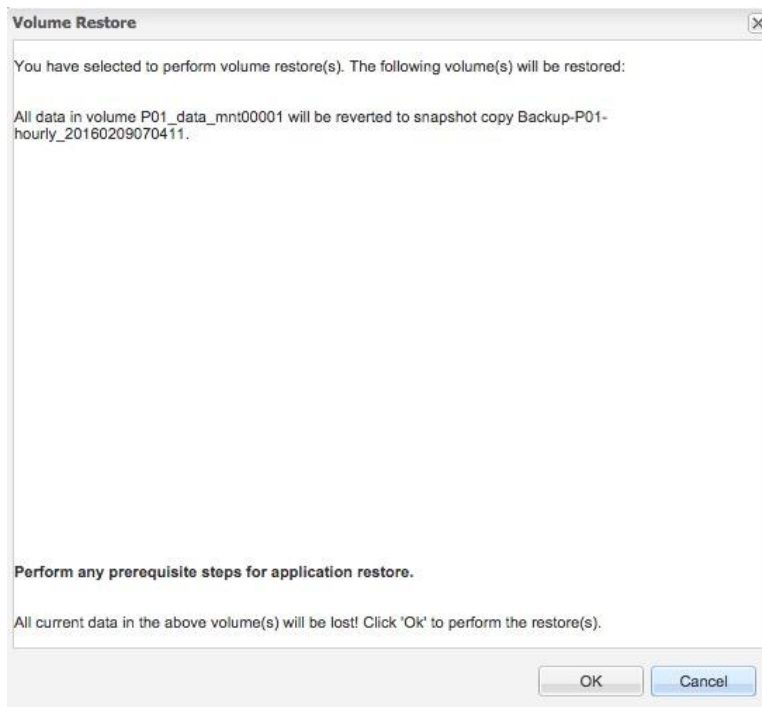
15. Wait until the restore process completes.
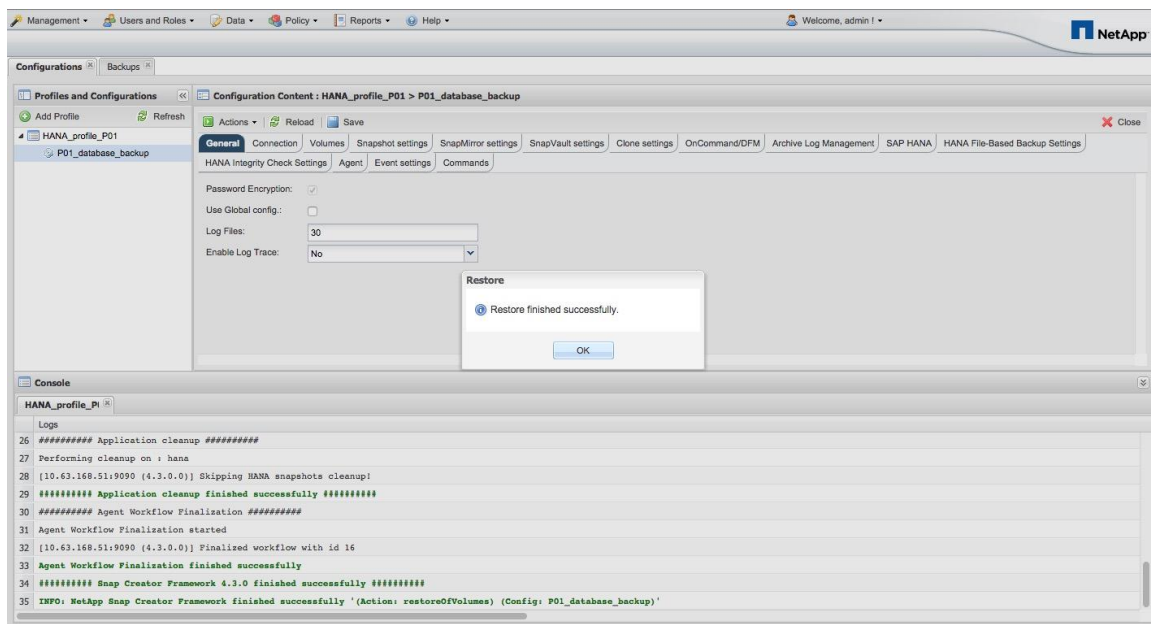
16. On each database node, mount all data volumes. With our example, only one volume must be remounted at the database node.

```
mount /hana/data/P01/mnt00001
```

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

Select a Backup
Select a backup to recover the SAP HANA database

Selected Point in Time
Database will be recovered to its most recent state.

Backups
The overview shows backups that were recorded in the backup catalog as successful. The backup at the top is estimated to have the shortest recovery time.

Start Time	Location	Backup Prefix	Avail
2016-02-09 13:04:11	/hana/data/P01	SNAPSHOT	✦
2016-02-09 12:59:48	/hana/data/P01	SNAPSHOT	⊗
2016-02-09 12:56:48	/mnt/filebased/P01/	SnapCreator_BACKUP-P01...	◇
2016-02-09 10:46:42	/hana/data/P01	SNAPSHOT	⊗
2016-02-09 10:37:22	/hana/data/P01	SNAPSHOT	⊗
2016-02-08 16:36:12	/mnt/filebased/TMP/	SnapCreator_IC	◇
2016-01-26 15:00:02	/hana/data/P01	SNAPSHOT	⊗
2016-01-26 14:00:01	/hana/data/P01	SNAPSHOT	⊗
2016-01-26 13:00:01	/hana/data/P01	SNAPSHOT	⊗
2016-01-26 12:00:02	/hana/data/P01	SNAPSHOT	⊗

[Refresh] [Show More]

Details of Selected Item

Start Time:	🛈 2016-02-09 13:04:11	Destination Type:	SNAPSHOT	Source System:	P01
Size:	16,55 GB	Backup ID:	1455019451268	External Backup ID:	Backup-P01-hourly_20160209070411
Backup Name:	/hana/data/P01				

Alternative Location: 🛈 [   ]

[Check Availability]

[?]   [< Back] [Next >] [Cancel] [Finish]

---

18. Select other settings as required. Deselect Use Delta Backups (Recommended). Click Next.

Other Settings

Check Availability of Delta and Log Backups

You can have the system check whether all required delta and log backups are available at the beginning of the recovery process. If delta or log backups are missing, they will be listed and the recovery process will stop before any data is changed. If you choose not to perform this check now, it will still be performed but later. This may result in a significant loss of time if the complete recovery must be repeated.

Check the availability of delta and log backups:
☑ File System 🛈
☐ Third-Party Backup Tool (Backint)

Initialize Log Area

If you do not want to recover log segments residing in the log area, select this option. After the recovery, the log entries will be deleted from the log area.

☐ Initialize Log Area 🛈

Use Delta Backups

Select this option if you want to perform a recovery using delta backups. If you choose to perform a recovery without delta backups, only log backups will be used.

☐ Use Delta Backups (Recommended)

Install New License Key

If you recover the database from a different system, the old license key will no longer be valid
You can:
- Select a new license key to install now
- Install a new license key manually after the database has been recovered
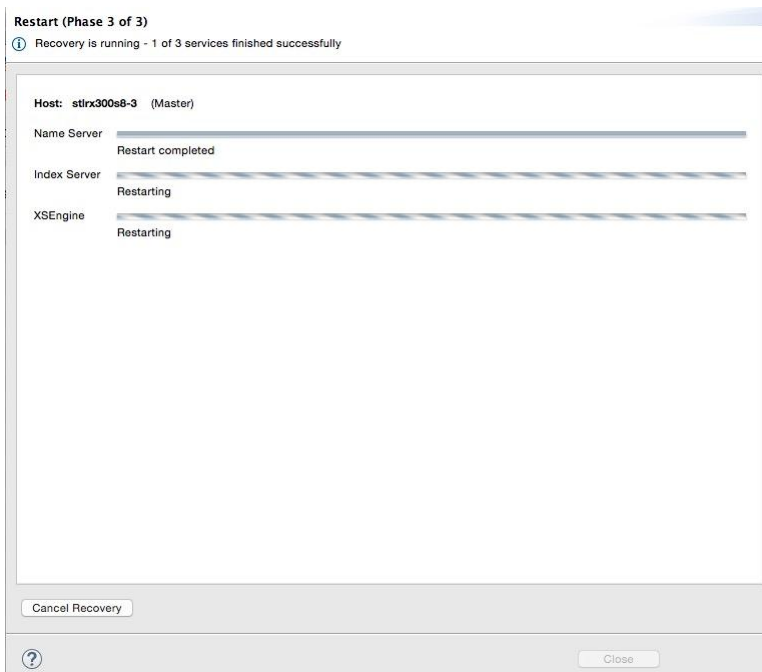
☐ Install New License Key

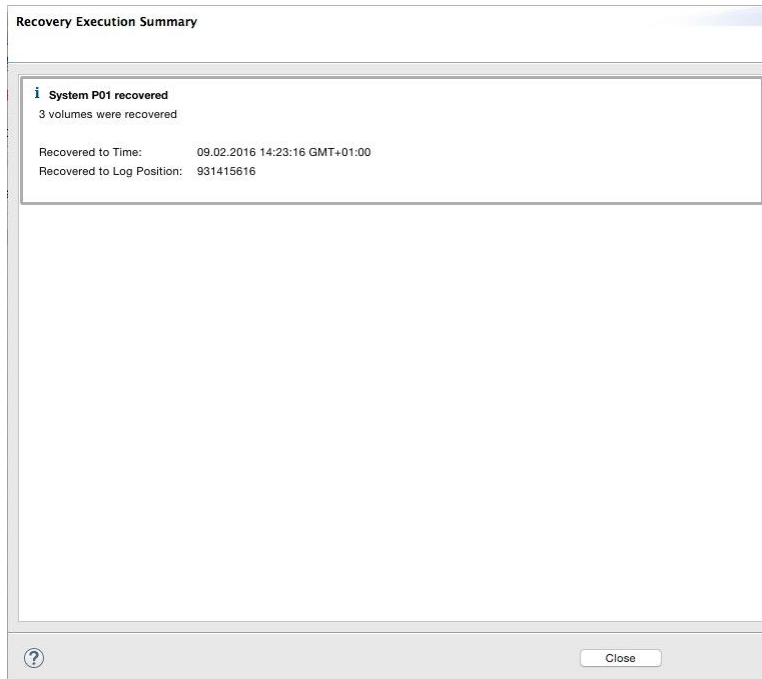[   ] [Browse]

[?]   [< Back] [Next >] [Cancel] [Finish]

---

19. Click Finish.

Review Recovery Settings

Review the recovery settings and choose 'Finish' to start the recovery. You can modify the recovery settings by choosing 'Back'.

System Information

System:	P01
Host:	10.63.168.55
Version:	1.00.110.00.1447753075

Recovery Definition

Recovery Type: Snapshot (Point-in-Time Recovery (Until Now))

Configuration File Handling

⚠ Caution

If you want to recover customer-specific configuration changes, you may need to make the changes manually in the target system.
If you are performing a recovery to a different system:
Note that the target system and the source system must have the same configuration. In particular, the number of database services with their own persistency must be the same in both systems.
More Information: SAP HANA Administration Guide

[Show SQL Statement]

⑦   [< Back] [Next >] [Cancel] [**Finish**]

20. The recovery process starts.

Restart (Phase 3 of 3)

ⓘ Recovery is running - 1 of 3 services finished successfully

Host: **stlrx300s8-3** (Master)

Name Server
Restart completed

Index Server
Restarting

XSEngine
Restarting

[Cancel Recovery]

⑦   [Close]

21. Recovery completes.

Recovery Execution Summary

i **System P01 recovered**
3 volumes were recovered

Recovered to Time: 09.02.2016 14:23:16 GMT+01:00
Recovered to Log Position: 931415616

Close

22. Resume any SnapVault relationships, if needed. See section 8.4, "Resume SnapVault Relation After a Restore."

8.2   Restore and Recovery from Secondary Storage

To restore and recover from secondary storage, complete the following steps:
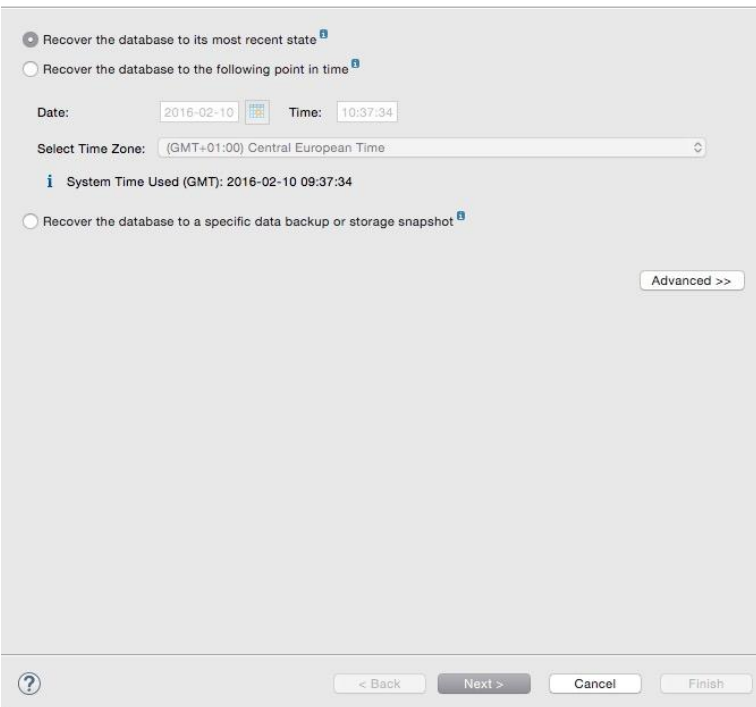
1.   In SAP HANA Studio, select the Recover option for the SAP HANA system.
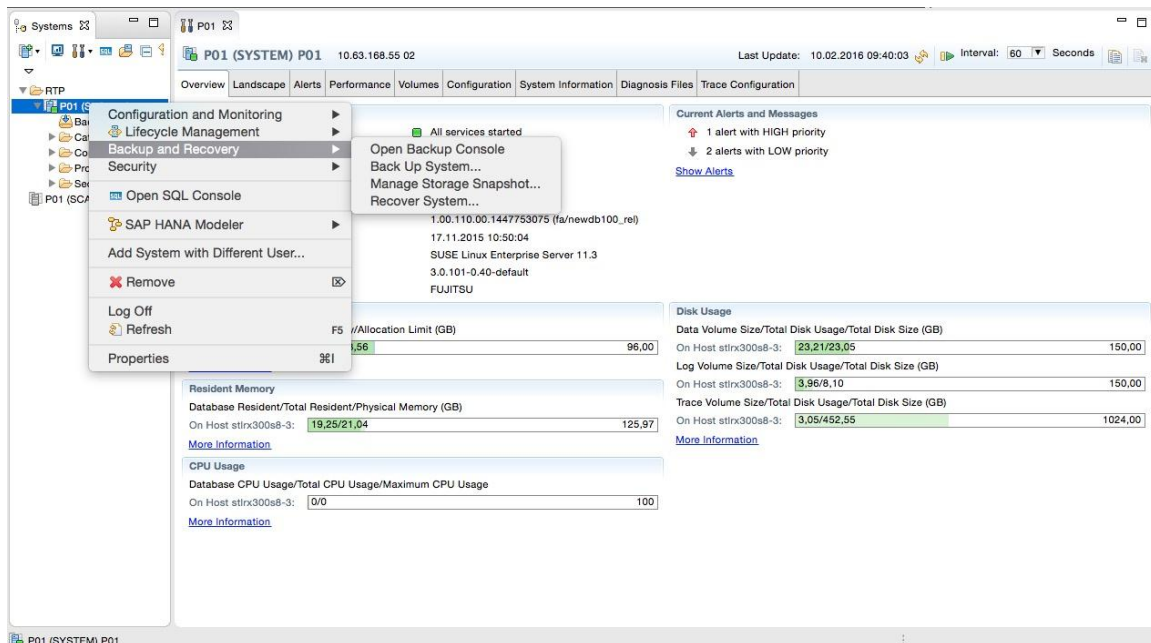
2.   The SAP HANA system is shut down.

3.   Select the recovery type and click Next.

4. Provide the log backup locations and click Next.

SAP HANA Backup and Recovery by Using Snap Creator

5. A list of available backups is shown based on the content of the backup catalog. Select the required backup and write down the external backup ID.

6. On each database node, unmount all data volumes. In this example, only one volume must be unmounted from the database node.

```
umount /hana/data/P01/mnt00001
```

Note: In an SAP HANA multinode setup with Fibre Channel, the umount operation is executed by the SAP HANA name server as part of the shutdown process.

7. In the Snap Creator GUI, select the SAP HANA system, go to Actions, and select Restore.

8. Click Next.

9. Select Secondary.

10. Enter the required information. The Snapshot name correlates with the backup ID that was selected in SAP HANA Studio. The source and destination path is `<SVM name>:<volume_name>`.

11. Select Finish.

12. Click No.

Warning

Would you like to add more restore items on this configuration? Select 'No' to start the entire restore operation or 'Yes' to add more restore items.
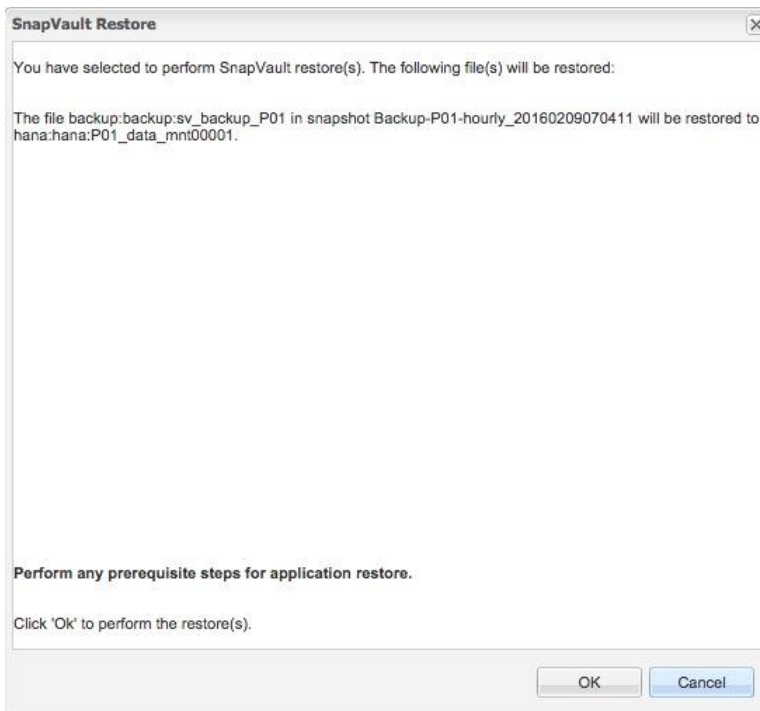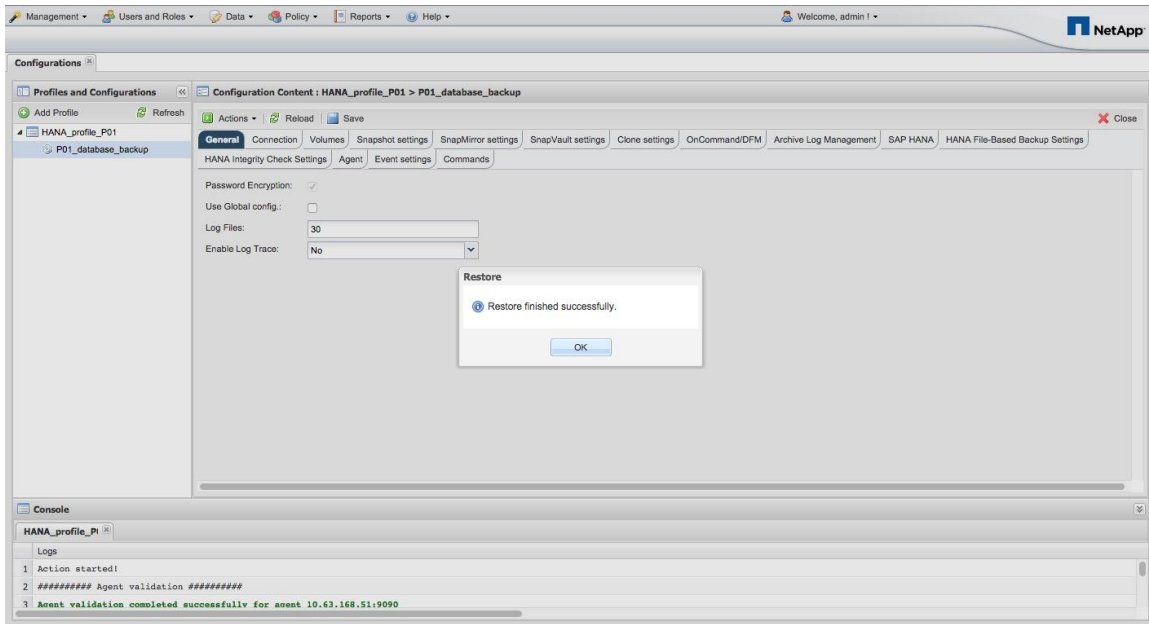
Yes   No

13. Select Yes to add more restore items; in our example, there is only one volume.

    Note: In an SAP HANA multinode setup, there are multiple volumes that must be selected for the restore process.

14. When all volumes are selected, click OK to start the restore process.

SnapVault Restore

You have selected to perform SnapVault restore(s). The following file(s) will be restored:

The file backup:backup:sv_backup_P01 in snapshot Backup-P01-hourly_20160209070411 will be restored to hana:hana:P01_data_mnt00001.

Perform any prerequisite steps for application restore.

Click 'Ok' to perform the restore(s).

OK   Cancel

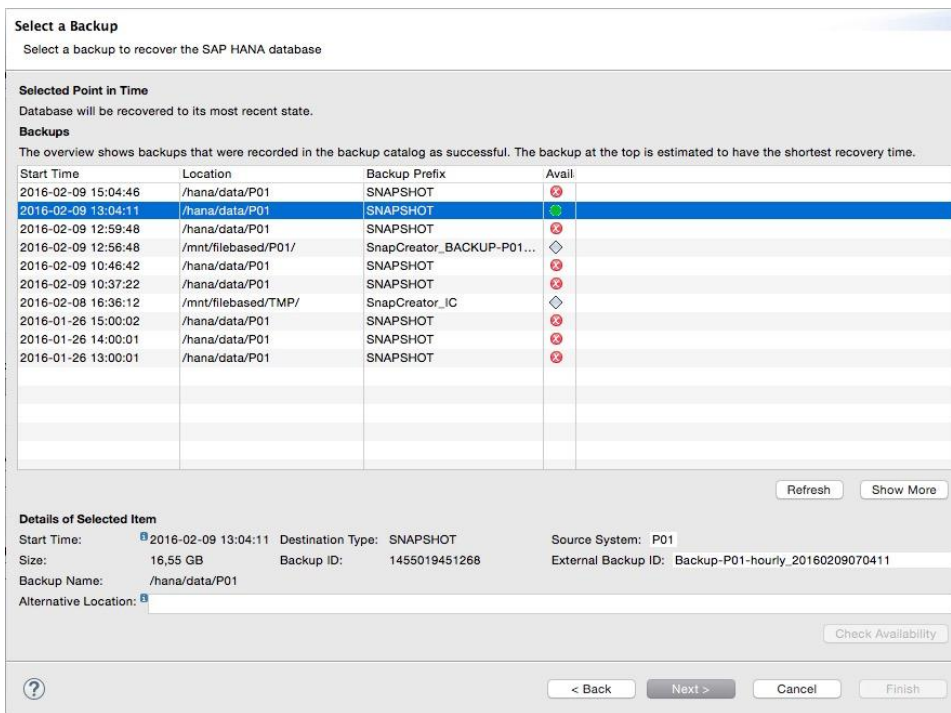15. Wait until the restore process completes.

16. Mount the data volume.

```
mount /hana/data/P01/mnt00001
```

> **Note:** In an SAP HANA multinode setup with Fibre Channel, the mount operation is executed by the SAP HANA name server as a part of the recovery process.

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

18. 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.

19. Select other settings as required and deselect Use Delta Backups. Click Next.

20. Click Finish.

21. The recovery process starts.

Recovery Progress Information
ⓘ Prepare Recovery - Start Master Name Server

Host: stlrx300s8-3
🕐 Recovery of host 'stlrx300s8-3'

22. The recovery process completes.

Recovery Execution Summary

ⓘ **System P01 recovered**
3 volumes were recovered

Recovered to Time: 10.02.2016 10:34:56 GMT+01:00
Recovered to Log Position: 938866944

⑦   Close

8.3   Restore and Recovery from Primary Storage by Using Single File Restore

This workflow is similar to the volume restore workflow described in section 8.1, "Restore and Recovery from Primary Storage by Using Volume-Based Restore." However, instead of restoring the complete volume, all single files within this volume are restored. One of the benefits of this method is that all existing Snapshot copies are preserved, even if they are older than the most current Snapshot copy that was used as the source for the restore.

Note: In a file-based restore, the mount and unmount operation of the target volume is not required. Also, a restore to a Snapshot copy older than the Snapshot copy currently used as the base for SnapVault replication does not require the deletion of the SnapVault relationship.

This description starts with the restore process (step 11) from section 8.1. The initial steps to start the recovery in SAP HANA Studio are identical to the method described in the previous section.

1.   Select the SVM, the volume name, and the Snapshot name. The Snapshot name correlates with the backup ID that has been selected within SAP HANA Studio. Select Single File Restore and click Next.

2. Select the first folder, `hdb00001`, and expand the folder list. This selects all files within the folder. Click Next.

3. As the target, select the same destination path in the destination volume and click Next.

   Note: If all data, including the directory structure, has been deleted, the directory structure must be set up before the restore process is started.
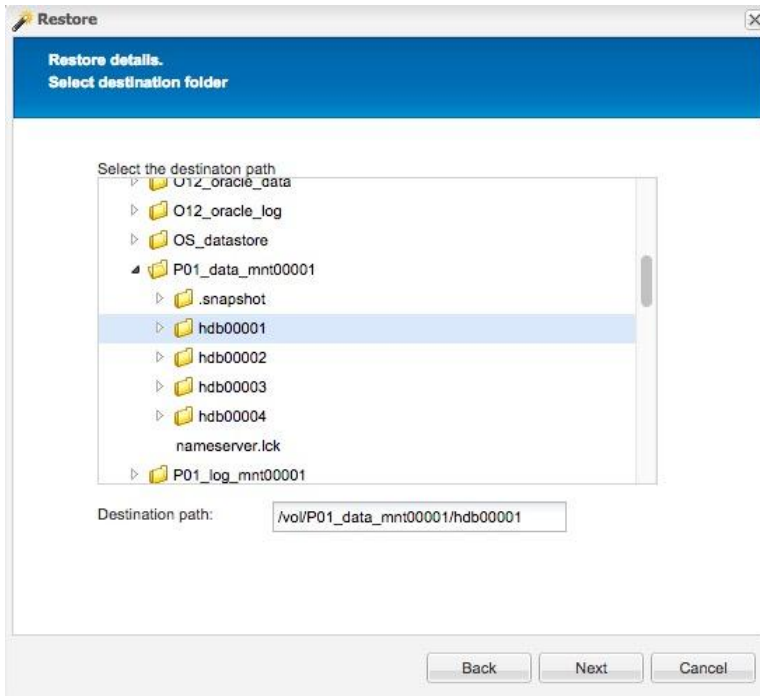
4. The summary page shows the matching source and destination files. Click Finish.
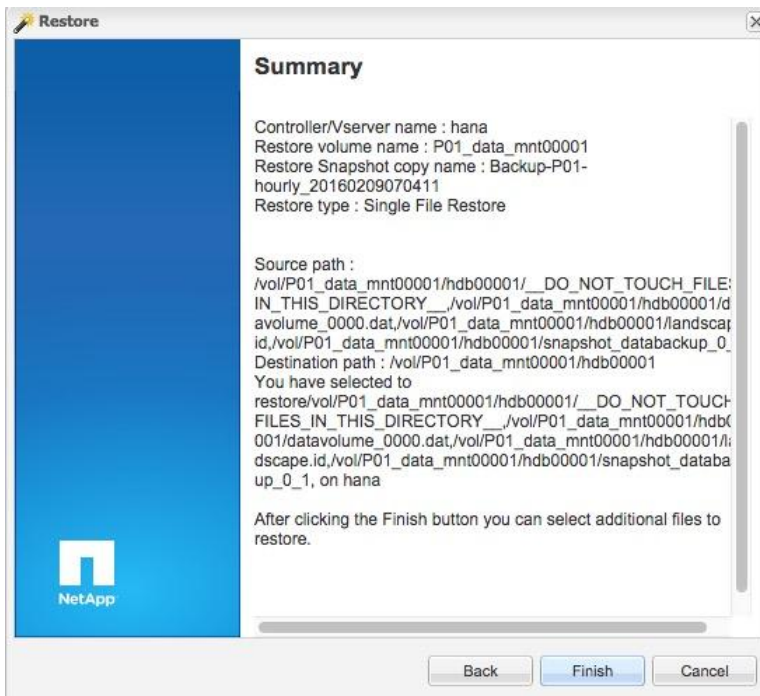
5. Click Yes to add additional files for the restore.
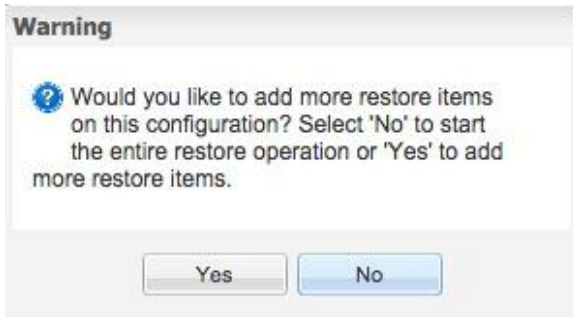
6. Repeat step 2 to step 5 for each of the folders within the HANA data volume. Finally, add another cycle (step 2 to step 5) for the single top-level file `nameserver.lck`. This time select No in step 5.

7. On the final dialog page, check if all files are selected and if they are restored into the correct folder. Click OK to start the restore.
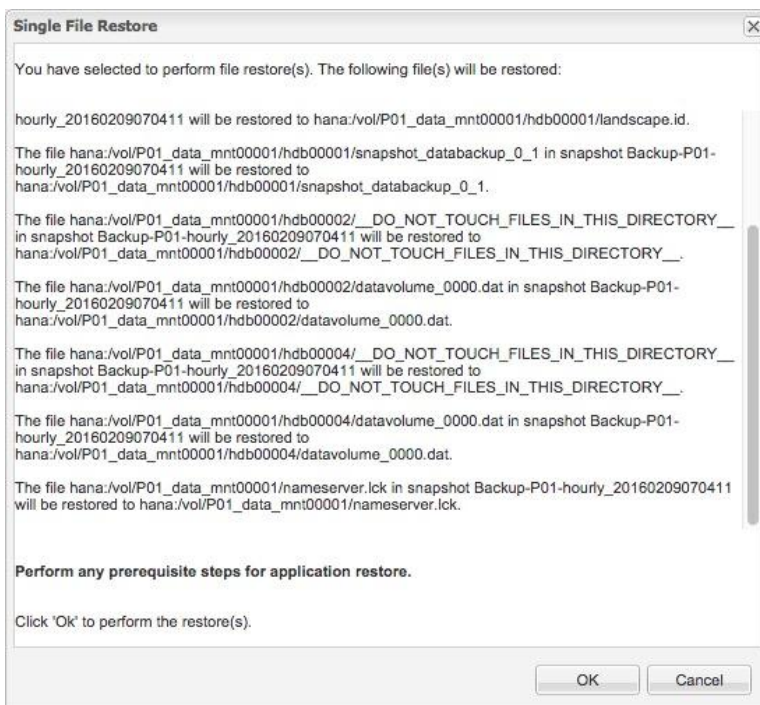
8. Wait until the restore is finished.
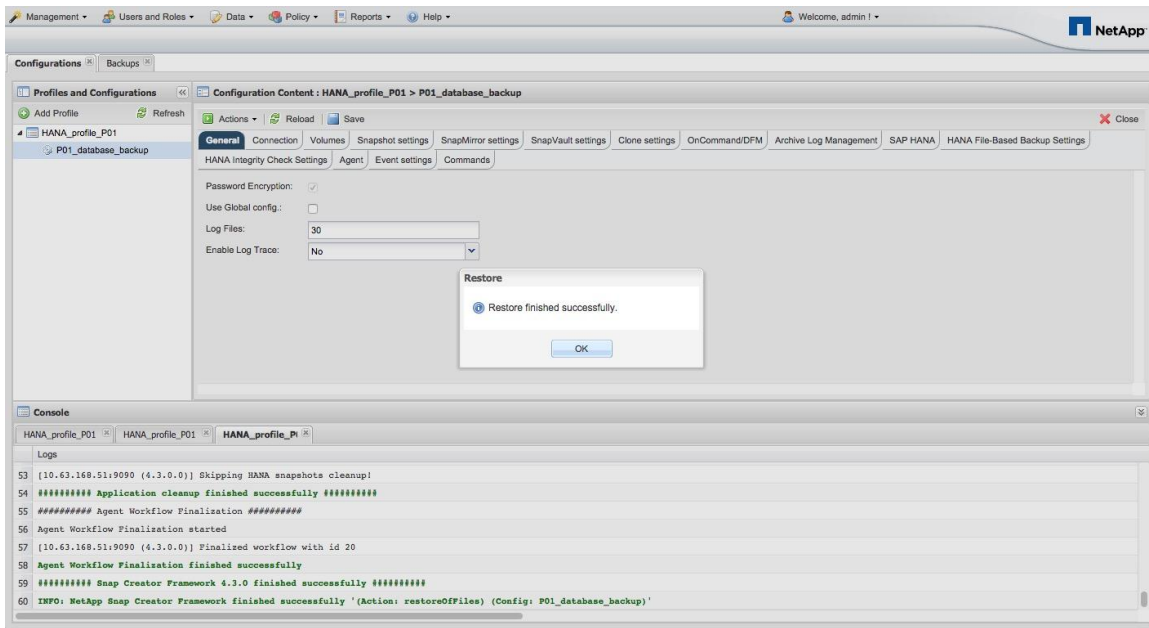
9. Continue the recovery process with step 17 from section 8.1.

8.4 Resume SnapVault Relation After a Restore

The SnapVault relationship is deleted with any restore that has been performed by using a Snapshot copy that is older than the Snapshot copy currently used as the base for SnapVault replication. After the restore and recovery process is finished, the SnapVault relation must be resumed so that backups can be performed again with Snap Creator. Otherwise, Snap Creator issues an error message because it cannot find the SnapVault relations on the primary storage systems.

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

Resume SnapVault Relation in Clustered Data ONTAP

With clustered Data ONTAP, run the following commands to resume the SnapVault relationship.

1. Recreate and resynchronize the SnapVault relationship.

```
stl-cmode-INC10364933::> snapmirror create -source-path hana:P01_data_mnt00001 -destination-path
backup:sv_backup_P01 -type XDP
Operation succeeded: snapmirror create for the relationship with destination
"backup:sv_backup_P01".
```

```
stl-cmode-INC10364933::> snapmirror resync -destination-path backup:sv_backup_P01 -type XDP

Warning: All data newer than Snapshot copy Backup-P01-hourly_20160209070411 on volume
backup:sv_backup_P01 will be deleted.
Do you want to continue? {y|n}: y
Operation is queued: initiate snapmirror resync to destination "backup:sv_backup_P01".
```

2. To restart the SnapVault transfer, a manual Snapshot copy and SnapVault update are required.

```
stl-cmode-INC10364933::> snapshot create -vserver hana -volume P01_data_mnt00001 -snapshot temp -
snapmirror-label HANA_P01
```

```
stl-cmode-INC10364933::> snapmirror update -destination-path backup:sv_backup_P01
Operation is queued: snapmirror update of destination "backup:sv_backup_P01".
```

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

```
hana::> snapmirror list-destinations -source-path hana1a:hana_data

stl-cmode-INC10364933::> snapmirror list-destinations -source-path hana:P01_data_mnt00001
                                                   Progress
Source            Destination        Transfer Last           Relationship
Path        Type  Path        Status Progress Updated        Id
----------- ----- ------------ ------- --------- ------------ ---------------
hana:P01_data_mnt00001
            XDP   backup:sv_backup_P01
                               Transferring
                                        378.6MB  02/11 12:37:43
                                                               d2c56330-d0bb-11e5-846e-
00a09851f8c6
```

Resume SnapVault Relation in Data ONTAP 7-Mode

1. To resume the SnapVault relation, run the following command:

```
snapvault start -r -S <Source Controller>:<Source Volume> <Backup Controller>:<Backup Volume>
```

2. This process must be performed for all volumes belonging to the SAP HANA database.

```
hana2b> snapvault start -r -S hana1a:/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 hana1a:/vol/data_00001/mnt00001 was successful.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hana2b> snapvault start -r -S hana1b:/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 hana1b:/vol/data_00002/mnt00002 was successful.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hana2b> snapvault start -r -S hana1a:/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 hana1a:/vol/data_00003/mnt00003 was successful.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

3. When data transfer is finished, backups can be scheduled again by using Snap Creator.

8.5 Restore 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 cannot handle the restore because there is no SnapVault relation on the primary storage systems.

Restore After Primary Storage Failure in Clustered Data ONTAP

If the primary volume is completely lost, complete the following steps:

1. Create a primary volume with the data protection type.

```
stl-cmode-INC10364933::> volume create -vserver hana -volume P01_data_mnt00001 -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.

```
stl-cmode-INC10364933::> snapmirror restore -destination-path hana: P01_data_mnt00001 -source-
path backup:sv_backup_P01 -source-snapshot sc-backup-daily_20140505121000
[Job 6746] Job is queued: snapmirror restore from source "backup:sv_backup_P01" for the snapshot
sc-backup-daily_20140505121000.
```

3. If other volumes and file systems, such as `/hana/shared`, have also been backed up with Snap Creator, then the previous commands must be repeated for those volumes.

4. When the restore process is finished, the recovery can be performed by using SAP HANA Studio, as is described in section 8.2, "Restore and Recovery from Secondary Storage."

Restore After Primary Storage Failure in Data ONTAP Operating in 7-Mode

To restore after a primary storage failure in Data ONTAP operating in 7-Mode, complete the following steps:

1. The restore must be performed directly on the secondary storage system by using the following command:

```
snapvault restore –s <Snapshot Name> -S <Backup Controller>:<Backup Volume> <Source
Controller>:<Source Volume>
```

2. This process must be performed for all volumes belonging to the SAP HANA database.

```
hana1a>  snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00001/mnt00001 hana1a:/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 [hana1a:vdisk.qtreePreserveComplete:info]: Qtree preserve is complete
for /vol/data_00001/mnt00001.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hana1a>  snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00003/mnt00003 hana1a:/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 [hana1a:vdisk.qtreePreserveComplete:info]: Qtree preserve is complete
for /vol/data_00003/mnt00003.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hana1b>  snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00002/mnt00002 hana1b:/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 [hana1b:vdisk.qtreePreserveComplete:info]: Qtree preserve is complete
for /vol/data_00002/mnt00002.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

3. If other volumes and file systems, such as /hana/shared, have also been backed up with Snap Creator, the preceding commands must be repeated for those volumes.

4. When the restore process is finished, the recovery can be performed by using SAP HANA Studio, as is described in section 8.2, "Restore and Recovery from Secondary Storage."

9 Block Integrity Check

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 GUI or CLI.

The block integrity check must be enabled in the Snap Creator configuration (see Figure 16).

**Figure 16) Integrity check configuration.**
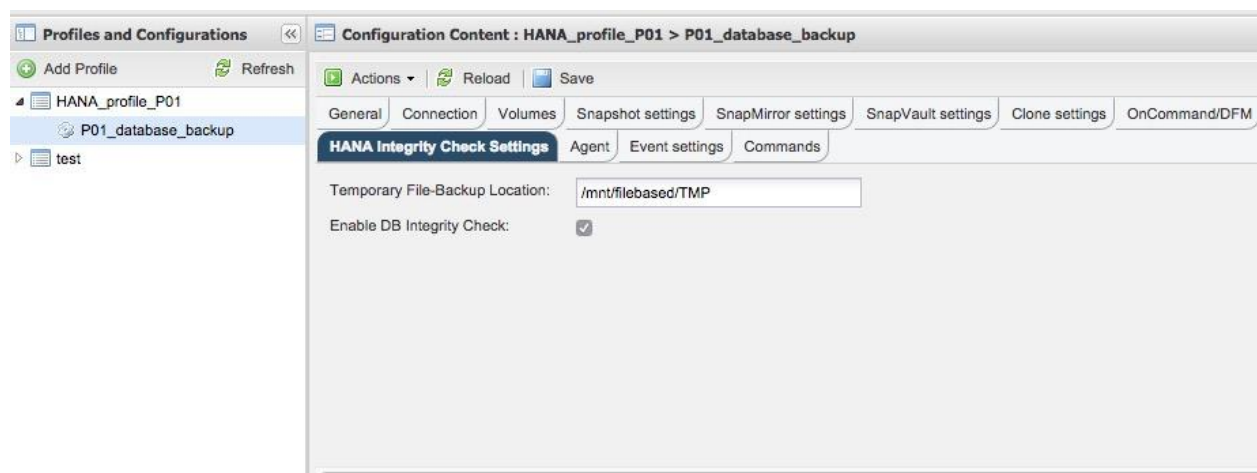
The integrity check writes a standard file-based HANA backup into the specified folder. Any restart of the integrity check overwrites the existing backup files and deletes the previous integrity check backup from the HANA catalog.

9.1 Integrity Check Workflow

1. Click the P01_database_backup configuration. From the Actions menu, select Integrity Check.
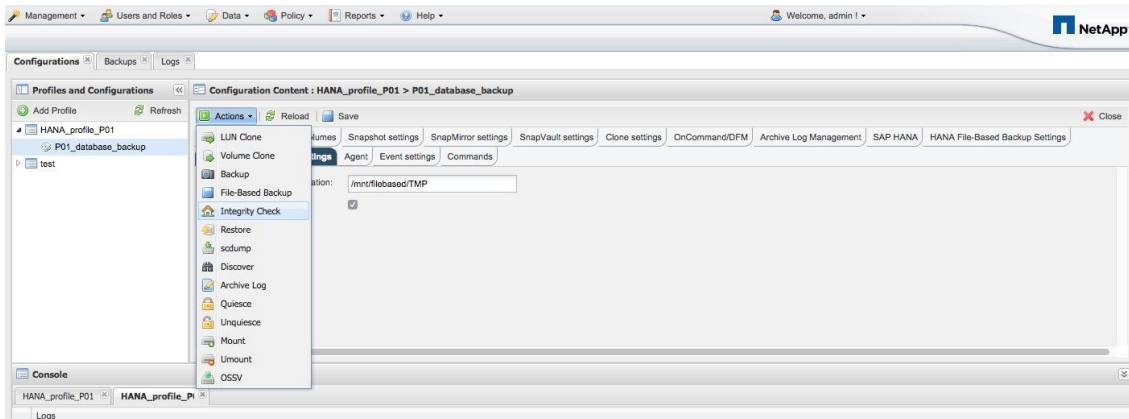
2. In the Policy field, select None and click OK.
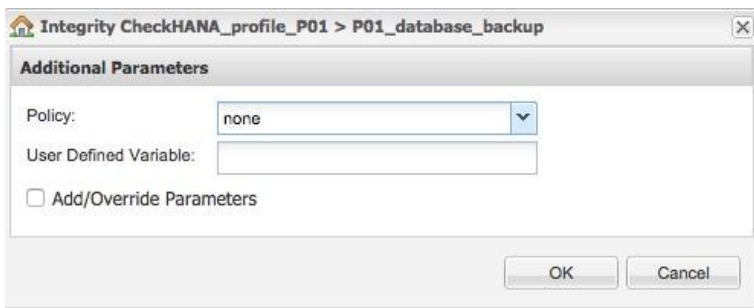
3. The integrity check creates a standard HANA file-based backup in the folder specified.

```
stlrx300s8-1:/mnt/filebased/TMP # ll
total 17435612
-rw-r----- 1 p01adm sapsys      155648 Feb 11 13:37 SnapCreator_IC_databackup_0_1
-rw-r----- 1 p01adm sapsys     83894272 Feb 11 13:37 SnapCreator_IC_databackup_1_1
-rw-r----- 1 p01adm sapsys 17616084992 Feb 11 13:39 SnapCreator_IC_databackup_2_1
-rw-r----- 1 p01adm sapsys     83894272 Feb 11 13:37 SnapCreator_IC_databackup_4_1
stlrx300s8-1:/mnt/filebased/TMP #
```

10 File-Based Backup

File-based backups can be used to transfer an existing HANA database into a different environment or to use a file-based backup storage solution such as NetApp AltaVault™ cloud-integrated storage to store backups in the cloud.

Note: With a file-based backup, retention is not managed with Snap Creator. File-based backups must be deleted manually by using SAP HANA Studio.
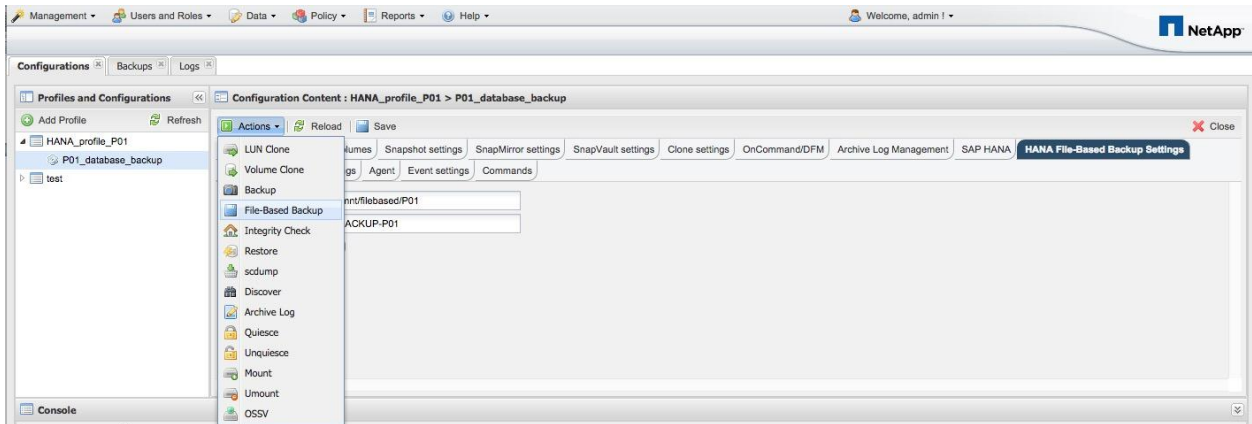
Configure file-based backup in Snap Creator, as is depicted in Figure 17) .

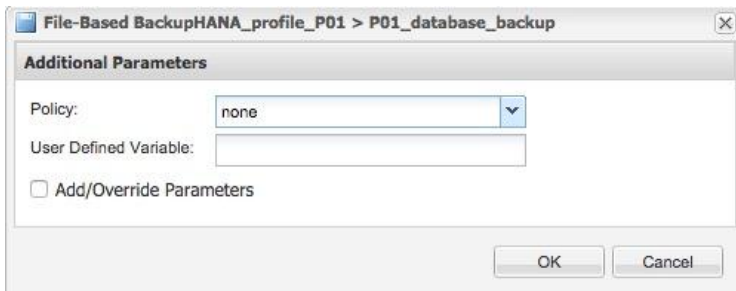**Figure 17) File-based backup configuration.**
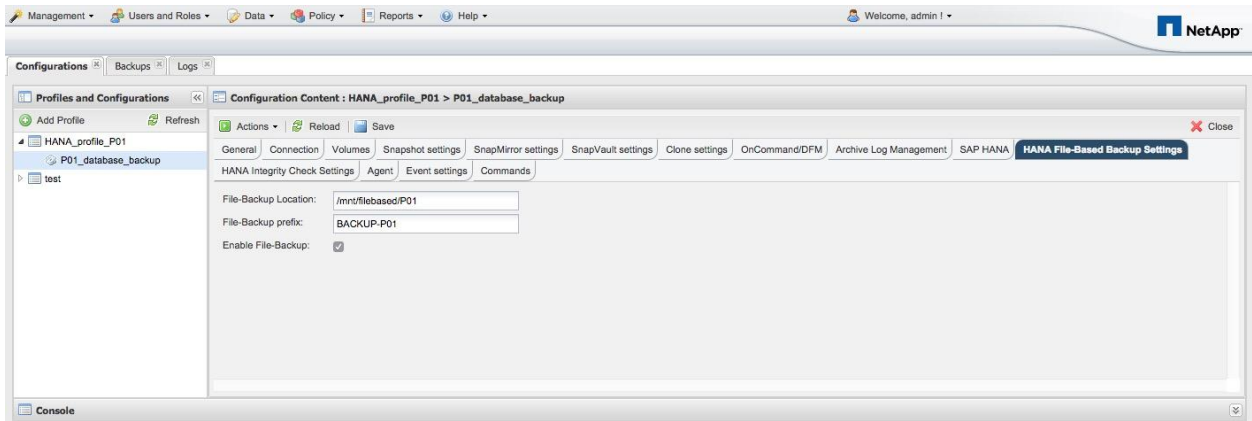
10.1 File-Based Backup Workflow

1. Select the `P01_database_backup` configuration. From the Actions menu, select File-Based Backup.

2. In the Policy field, select None and then click Next.

3. The following files are created.

```
stlrx300s8-1:/mnt/filebased/P01 # ll
total 34871224
-rw-r----- 1 p01adm sapsys      155648 Feb 11 13:55 SnapCreator_BACKUP-P01_2016-02-11-
08:55:28.739_databackup_0_1
-rw-r----- 1 p01adm sapsys    83894272 Feb 11 13:55 SnapCreator_BACKUP-P01_2016-02-11-
08:55:28.739_databackup_1_1
-rw-r----- 1 p01adm sapsys 17632862208 Feb 11 13:57 SnapCreator_BACKUP-P01_2016-02-11-
08:55:28.739_databackup_2_1
```

```
-rw-r----- 1 p01adm sapsys    83894272 Feb 11 13:55 SnapCreator_BACKUP-P01_2016-02-11-
08:55:28.739_databackup_4_1
```

11 Upgrade Snap Creator

The latest upgrade to Snap Creator is described in the Snap Creator documentation. The following section contains a short summary of the steps required to activate the new features in Snap Creator 4.3.

11.1 Upgrade Summary

The following steps describe the upgrade process for Snap Creator running on Linux. This example assumes that the agent and server are running on the same host. This example depicts an upgrade from Snap Creator 4.1.1 to Snap Creator 4.3.0 and assumes that all installations are installed into the path `/opt/Netapp/SnapCreator`, as recommended in section 4.2, "Snap Creator Framework."

1. First, stop the Snap Creator server and Snap Creator agent.

```
stlrx300s8-1:/opt/NetApp/SnapCreator # cd scServer4.1.1/bin/
stlrx300s8-1:/opt/NetApp/SnapCreator/scServer4.1.1/bin # ./scServer stop
Shutting down scServer:
scServer stopped
stlrx300s8-1:/opt/NetApp/SnapCreator/scServer4.1.1/bin #

stlrx300s8-1:/opt/NetApp/SnapCreator/scServer4.1.1/bin # cd ../../scAgent4.1.1/bin/
stlrx300s8-1:/opt/NetApp/SnapCreator/scAgent4.1.1/bin # ./scAgent stop
Shutting down scAgent:
    Watchdog: Stopped
       Agent: Stopped
stlrx300s8-1:/opt/NetApp/SnapCreator/scAgent4.1.1/bin #
```

2. Install the Snap Creator server and agent as described in section 4.2, but do not start the new server and agent yet. The new version is installed into a different path.

3. Delete the newly installed, empty Snap Creator internal database folder
   `/opt/Netapp/SnapCreator/scServer4.3.0/engine/snapcreator`.

```
stlrx300s8-1:/ # cd /opt/Netapp/SnapCreator/scServer4.3.0/engine/
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine # ll
total 100
drwxr-xr-x 1 root root    60 Feb 16 04:23 configs
-rw-r--r-- 1 root root  2343 Feb 16 04:26 derby.log
drwxr-xr-x 1 root root   410 Feb 16 03:51 etc
drwxr-xr-x 1 root root  1216 Feb 16 03:51 lib
drwxr-xr-x 1 root root  1890 Feb 16 04:26 logs
-rw------- 1 root root  1121 Feb 16 04:26 nohup.out
drwxr-xr-x 1 root root   142 Feb 16 04:26 snapcreator
-rwxr-xr-x 1 root root 93002 Jan 18 06:15 snapcreator.jar
drwxr-xr-x 1 root root     0 Feb 16 04:11 tmp
drwxr-xr-x 1 root root    36 Feb 16 03:51 wars
drwxr-xr-x 1 root root   160 Feb 16 04:07 work
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine # rm -r snapcreator/
```

4. Copy the old database files to the new installation.

```
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine # cp -ar
../../scServer4.1.1/engine/snapcreator .
```

5. Copy the old configuration files to the new installation.

```
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine # cd config
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine/config # cp –ar
../../../scServer4.1.1/engine/configs/* .
```

6. Copy the old log file to the new installation.

```
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine/config # cd ../logs
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine/logs # cp -ar
../../../scServer4.1.1/engine/logs/* .
```

7. To perform the upgrade, run the following command:

```
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine/logs # cd ..
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine # java -jar snapcreator.jar -upgrade
```

8. After this upgrade, you can start the Snap Creator service.

```
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine/ # cd ../bin
stlrx300s8-1:/opt/Netapp/SnapCreator/scServer4.3.0/engine/bin # ./scServer start
```

After verifying the functionality of the new version of Snap Creator, you can delete the old installation folder.

11.2 Activate New SAP HANA Plug-In Features

Configure the block integrity check as described in section 9, "Block Integrity Check."

Configure the file-based backup as described in section 10, "File-Based Backup."

With version 4.3, Snap Creator implements automated management of log backup files. In the previous version, customers deployed a special script to enable Snap Creator to clean up old log backups. In order to switch the configuration to the new function, complete the following steps:

1. Remove the call of the script in the corresponding Snap Creator config file:

   `/opt/NetApp/SnapCreator/scServer4.1.1c/engine/configs`. The corresponding line is `POST_NTAP_CMD01`.

```
###############################################################################################
######################
#      Post Commands                                                                          #
###############################################################################################
######################
POST_NTAP_DATA_TRANSFER_CMD01=
POST_RESTORE_CMD01=
POST_APP_QUIESCE_CMD01=
POST_NTAP_CMD01=/mnt/hwval/backup-script/delete-log-files.sh
POST_CLONE_CREATE_CMD01=
POST_APP_UNQUIESCE_CMD01=
```

2. Enable log cleanup, as is described in step 7 of section 5.3, "Snap Creator Configuration."

Version History

Version	Date	Document Version History
Version 1.0	June 2014	Initial version
Version 1.1	August 2014	Minor updates for Snap Creator release 4.1.1
Version 1.2	December 2014	Minor updates to include FCP support
Version 1.3	August 2015	• Added section on RTO comparison • Updated log backup and log backup housekeeping section
Version 2.0	February 2016	• Updates to cover new features of Snap Creator 4.3 • Reorganized document based on lab setup with HANA SPS11 and NetApp clustered Data ONTAP

SAP HANA Backup and Recovery by Using Snap Creator

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp®
www.netapp.com