



Technical Report

SAP HANA Backup and Recovery Using NetApp Storage Systems and Commvault Software

Marco Schoen, NetApp
Dr. Tristan Daude, Commvault Systems
March 2021 | TR-4711

In partnership with



Abstract

This document describes the design of a NetApp and Commvault solution for SAP HANA, which includes Commvault IntelliSnap snapshot management technology and NetApp® Snapshot™ technology. The solution is based on NetApp storage and the Commvault data protection suite.

TABLE OF CONTENTS

Overview	3
The NetApp and Commvault solution	3
Run time of Snapshot copy backups	6
Recovery time objective comparison	7
NetApp and Commvault SAP HANA backup solution	10
Solution component overview	10
Solution components	11
Supported SAP HANA releases and configurations	13
Capacity requirements for Snapshot copy backups	13
Installation and configuration overview	13
Commvault installation	13
Configuration steps	13
Data protection strategy	16
Backup retention management and housekeeping of log backups	17
Conclusion	17
Where to find additional information	18
Version history	18

LIST OF TABLES

Table 1) Data protection parameters	16
Table 2) Policies based on data protection parameters.	16

LIST OF FIGURES

Figure 1) NetApp and Commvault backup solution overview.	5
Figure 2) NetApp and Commvault backup solution overview – ANF	6
Figure 3) Customer example of Snapshot copy backup run time.	7
Figure 4) RTO for a 2TB database with file-based backups.	8
Figure 5) RTO for a 2TB database with Snapshot copy backups.	9
Figure 6) RTO comparison: file-based backup versus Snapshot copy backup.	9
Figure 7) Solution component overview.	10
Figure 8) Database and log backup configuration overview.	12
Figure 9) Commvault array management.	14
Figure 10) Commvault array management completed.	15
Figure 11) Commvault Plan example.	17

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.

Widespread adoption of public clouds, such as Microsoft Azure, Amazon Web Services (AWS), and Google Cloud Platform, has not only created attractive new cloud-based SAP offerings, but it's also fueled development of new data protection solutions as well as advanced new NetApp storage offerings such as Azure NetApp Files (ANF).

No matter which SAP landscapes are running, backup windows are always shrinking and the amount of data that must 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 that is needed to restore and recover SAP systems is a concern because downtime for SAP production and nonproduction systems must be minimized in order to reduce data loss and cost to the business.

The following points summarize the challenges with 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 that are placed on the database server, the storage system, and the storage network.
- **Shrinking backup windows.** Conventional backups can be made only when a few dialog or batch activities are in process on the SAP system. The scheduling of backups becomes more difficult when SAP systems are continuously used.
- **Rapid data growth.** Rapid data growth and shrinking backup windows require ongoing investment in the backup infrastructure. In other words, you must procure more tape drives, more backup disk space, and faster backup networks. You must also cover the ongoing expense of storing and managing these tape assets. Incremental or differential backups can help resolve these issues, but this arrangement results in a very slow, cumbersome, and complex restore process that is harder to verify. These systems usually increase recovery time objective (RTO) and recovery point objective (RPO) times in ways that are unacceptable 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 requirement to restore and recover the SAP system. Therefore, the desired 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 that is available for the upgrade process. The decision to proceed is generally based on the amount of time that is required to restore and recover the database from the previously created backup. Rather than just restoring a system to its previous state, a rapid restore provides more time to solve problems that might occur during an upgrade.
- **Backup and recovery for hybrid, in-cloud, and multicloud SAP environments.** How can RTO and RPO SLAs be met when cloud-based use cases are deployed?

The NetApp and Commvault solution

Commvault IntelliSnap uses NetApp Snapshot technology to create database backups in minutes. Because Snapshot technology does not move or copy any physical data blocks on the storage platform, the time that is needed to create a Snapshot copy is independent of the size of the database. Also, Snapshot technology does not move or copy data blocks when data in the active file system is changed, so the use of Snapshot technology has no performance effect on the live SAP system. Therefore, you can schedule the creation of Snapshot copies without having to consider peak dialog or batch activity periods. SAP and NetApp customers typically schedule multiple online Snapshot copies during the day; for example, every four hours is common. These Snapshot copies are usually kept for three to five days on the primary storage system before they are removed.

Commvault software acts as an orchestration layer on top of the Snapshot technology, so it can manage the creation and deletion of Snapshot copies, as well as NetApp SnapMirror® and

SnapVault® relationships. With the latest feature release, Commvault not only supports Snapshot copy management on premises, but also in hybrid and in-cloud environments where SAP HANA is running in Microsoft Azure virtual machines (VMs) leveraging ANF storage. The software enables you to set retention policies on Snapshot copies as if they were normal backups, so you do not need to manually delete them or to script their deletion.

Note: SnapVault and SnapMirror are currently supported only for on-premises environments by Commvault.

Commvault software also provides a GUI-based restore feature in which the end user only has to specify the system that must be restored and up to what point in time. Commvault software then uses its recovery catalog to determine which Snapshot copy must be used. Commvault software can make this determination regardless of whether or not the Snapshot copy has been vaulted by using SnapVault. When the initial Snapshot copy has been reverted to the active file system, Commvault software applies the required log backups to complete the database recovery to the specified point in time. These log backups are restored from streaming backups that Commvault catalogs along with the database Snapshot copy to enable either full or point-in-time hands-free recovery.

Snapshot copies also provide key advantages for restore and recovery operations. Commvault IntelliSnap uses NetApp SnapRestore® data recovery software. SnapRestore enables the restore of an entire database or a portion of a database to any point in time, based on the available Snapshot copies. Such restore processes finish in a few minutes, independent of the size of the database. Because several online Snapshot copies are created during the day, the time required for the recovery process is significantly reduced relative to a traditional backup approach. And because a restore operation 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 required for restore and recovery operations, is reduced to several minutes rather than the several hours that conventional single-cycle tape backups require.

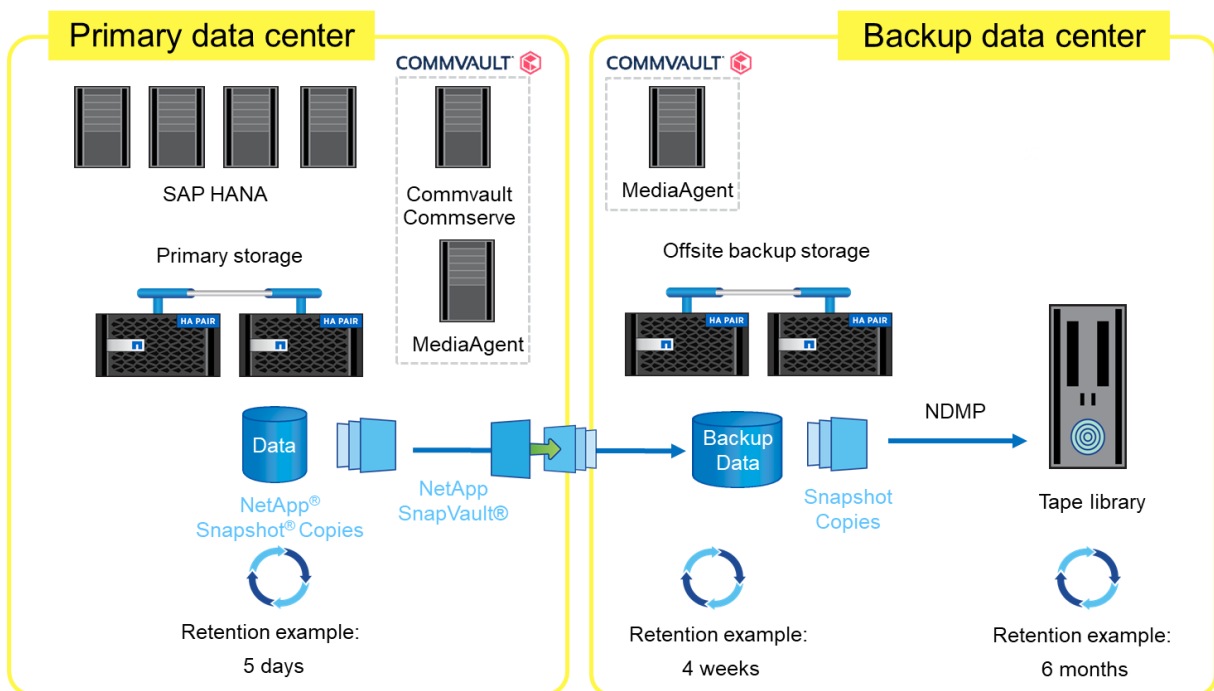
Snapshot copy backups are stored on the same disk system as the active online data. Therefore, Commvault and NetApp recommend that you use Snapshot copy backups as a supplement rather than as a replacement for backups to a secondary location. Most restore and recovery actions are handled by using SnapRestore by means of Commvault IntelliSnap technology on the primary storage system. Restores from a secondary location are necessary only if the primary storage system that contains 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 the Snapshot copies that are 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 software offers significant advantages over traditional backups. After an initial data transfer, in which all data has been transferred from the source to the destination, all subsequent backups copy only the changed blocks to the secondary storage. Therefore, the load on the primary storage system and the time that is needed for a full backup are significantly reduced. Because SnapVault software 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 that case, the tape infrastructure can be directly connected to the secondary storage, and data can be written to tape by using NDMP. Figure 1 shows an overview of the NetApp and Commvault backup solution. In addition, Commvault software offers native tape library and cloud storage container support within the Commvault platform.

Figure 1) NetApp and Commvault backup solution overview.



Architecture overview for an environment in Azure based on Azure NetApp Files

An on-premise environment can easily be extended to hybrid cloud environment or an additional in-cloud landscape can be added over time. All centrally managed through Commvault CommandCenter. Commvault IntelliSnap also supports orchestration and automation of Snapshot copy management for SAP HANA systems running on Microsoft ANF storage. For an on-premise plus Azure hybrid environment, it is required to set up at least one Commvault MediaAgent in the cloud part of the environment. This MediaAgent must be connected to multiple Azure Blob storage containers for building a Commvault cloud library so it can orchestrate the Snapshot copy backups for the cloud-based SAP HANA systems and create an Azure Blob storage copy of all ANF snapshot copies for mid- and long-term retentions. Like in an on-premises environment, you can automate the creation of Snapshot copies for SAP HANA on ANF. Commvault's Snapshot revert option (based on NetApp SnapRestore) enables rapid in-place and out-of-place SAP HANA restore operations within minutes. It also allows for the fast creation of SAP HANA database clones by leveraging NetApp FlexClone functionality.

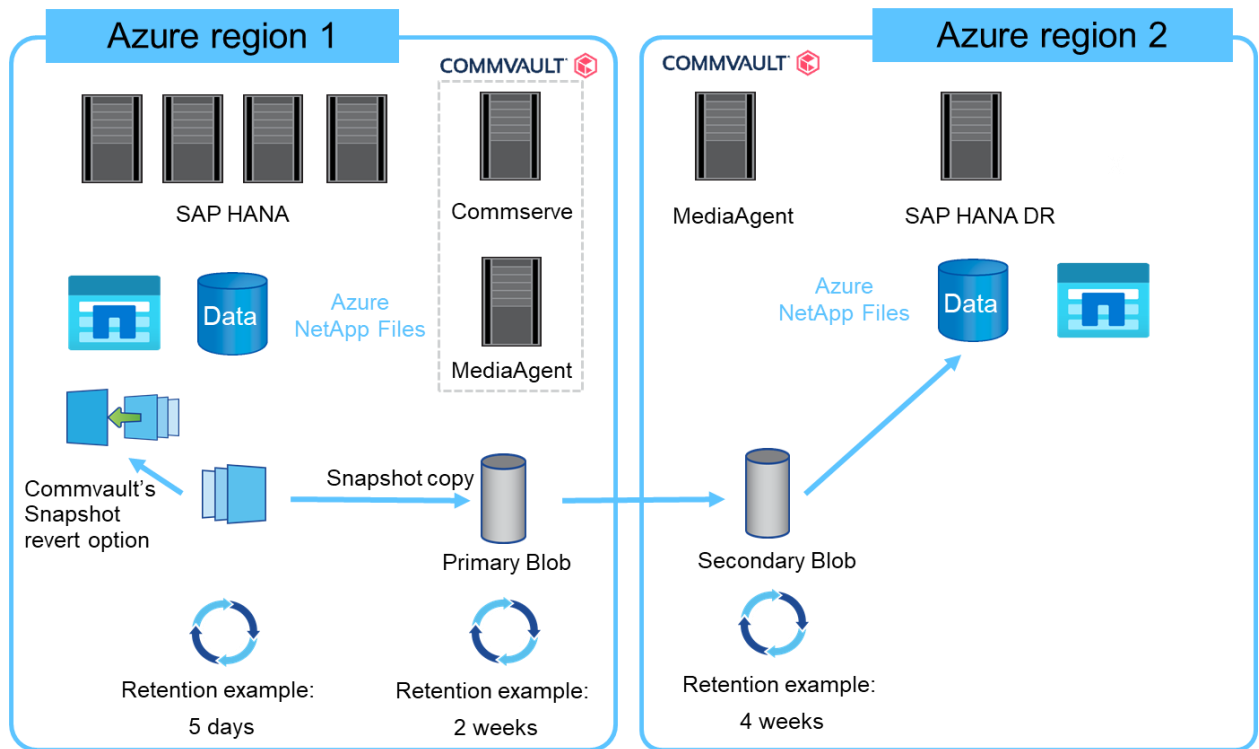
Another major Commvault benefit is that the entire SAP HANA snapshot recovery process can be streamlined and managed through Commvault's modern web-based CommandCenter GUI. This interface can help save time and enable a backup administrator to run SAP HANA recoveries or database refreshes. This also includes the creation of temporary SAP HANA database clones from Snapshot copies, which can be beneficial in supporting DevOps tasks, or if a single database table must be recovered while SAP S/4HANA production remains online. After its lifetime expires, the clone is automatically decommissioned.

The cloud-based MediaAgent can also create additional Blob storage copies for on-premises backup data to support disaster recovery use cases. In a pure in-cloud architecture, SAP HANA VMs are normally distributed across multiple Azure regions. In order to protect such environments, it is required to provision at least one MediaAgent VM in every region in conjunction with a central Commvault CommServe VM. After the ANF Snapshot copies are externalized to Azure Blob storage, Commvault can either leverage Azure geo replication (for example, through Azure GRS-RA Blob storage containers) or its own auxcopy mechanism to orchestrate backup copy replication to the other region for disaster recovery.

Figure 2 is an overview of the NetApp and Commvault backup solution for ANF.

Note: Support for SnapVault and SnapMirror is considered for a future Commvault release after 11.22

Figure 2) NetApp and Commvault backup solution overview – ANF.



Run time of Snapshot copy backups

Figure 3 shows SAP HANA Studio running SAP HANA on NetApp storage. Snapshot copies are used to back up the SAP HANA database. This figure shows that the SAP HANA database is backed up in 24 seconds by using Commvault IntelliSnap for Snapshot backup technology.

Figure 3) Customer example of Snapshot copy backup run time.

The screenshot shows the SAP HANA Studio Backup Catalog for SYSTEMDB@H23. The 'Backup Catalog' tab is active, displaying a list of backup operations. The selected backup is a Snapshot backup of size 1.66 GB, completed on 04.04.2018 at 15:17:46. The 'Backup Details' pane on the right shows the backup ID, status, type, destination, and location.

Status	Started	Duration	Size	Backup Type	Destination...
Success	18.04.2018 11:55:13	00h 00m 12s	1,73 GB	Data Backup	Snapshot
Success	18.04.2018 05:55:14	00h 00m 14s	1,73 GB	Data Backup	Snapshot
Success	17.04.2018 23:55:13	00h 00m 08s	1,73 GB	Data Backup	Snapshot
Success	17.04.2018 17:55:13	00h 00m 12s	1,73 GB	Data Backup	Snapshot
Success	17.04.2018 11:55:14	00h 00m 12s	1,73 GB	Data Backup	Snapshot
Success	17.04.2018 05:55:13	00h 00m 08s	1,73 GB	Data Backup	Snapshot
Success	16.04.2018 23:55:13	00h 00m 12s	1,72 GB	Data Backup	Snapshot
Success	16.04.2018 17:55:13	00h 00m 13s	1,72 GB	Data Backup	Snapshot
Success	16.04.2018 11:55:13	00h 00m 12s	1,72 GB	Data Backup	Snapshot
Success	16.04.2018 05:55:13	00h 00m 12s	1,72 GB	Data Backup	Snapshot
Success	15.04.2018 23:55:13	00h 00m 12s	1,72 GB	Data Backup	Snapshot
Success	15.04.2018 17:55:13	00h 00m 12s	1,72 GB	Data Backup	Snapshot
Success	11.04.2018 10:06:49	00h 00m 28s	1,67 GB	Data Backup	Snapshot
Success	04.04.2018 15:17:46	00h 00m 24s	1,66 GB	Data Backup	Snapshot
Success	04.04.2018 15:12:52	00h 02m 48s	1,66 GB	Data Backup	Backint
Success	04.04.2018 13:26:29	00h 00m 28s	1,66 GB	Data Backup	Snapshot
Success	26.03.2018 11:08:25	00h 00m 25s	1,66 GB	Data Backup	Snapshot
Success	26.03.2018 10:35:38	00h 00m 26s	1,66 GB	Data Backup	Snapshot
Success	20.03.2018 14:31:22	00h 01m 28s	1,63 GB	Data Backup	Snapshot
Success	20.03.2018 13:28:29	00h 00m 32s	1,61 GB	Data Backup	Snapshot
Success	20.03.2018 12:45:57	00h 00m 27s	1,61 GB	Data Backup	Snapshot
Success	20.03.2018 12:29:33	00h 02m 59s	1,61 GB	Data Backup	Backint
Success	02.03.2018 10:30:07	00h 00m 06s	1,48 GB	Data Backup	Snapshot

Backup Details:

- ID: 1522847866353
- Status: Successful
- Backup Type: Data Backup
- Destination Type: Snapshot
- Started: 04.04.2018 15:17:46 (Europe/Berlin)
- Finished: 04.04.2018 15:18:10 (Europe/Berlin)
- Duration: 00h 00m 24s
- Size: 1,66 GB
- Throughput: n.a.
- System ID: H23
- Comment: 306_COMPLETE_DATA_BACKUP
- Additional Information: <ok>
- Location: /hana/data/H23/mnt00001/

Recovery time objective comparison

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

Time needed to restore the 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 (MBps). For example, if the infrastructure supports a restore operation at a speed of 250MBps, it takes approximately one hour and 10 minutes to restore a database that is 1TB in size.

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

Time needed to start the 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 for a 2TB database. The start time is the same for a file-based restore and recovery operation and for a restore and recovery operation that is based on a Snapshot copy.

Time needed to recover the database

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

With file-based data backups, the backup schedule is typically once per day. A higher backup frequency is normally not possible, because the backup degrades production performance. Therefore,

in the worst case, all the logs that were written during the day must be applied during forward recovery.

Storage Snapshot copy data backups are typically scheduled with a higher frequency because they do not affect the performance of the SAP HANA database. For example, if Snapshot copy backups are scheduled every 6 hours, the recovery time would be, in the worst case, one-fourth of the recovery time for a file-based backup ($6 \text{ hours}/24 \text{ hours} = \frac{1}{4}$).

Figure 4 shows an RTO example for a 2TB database when file-based data backups are used. In this example, a backup is performed once per day. The RTO differs depending on when the restore and recovery operations were performed. If the restore and recovery operations were carried out immediately after a backup was performed, 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 operations were carried out immediately before the next backup was performed, and the maximum RTO was 4 hours and 30 minutes.

Figure 4) RTO for a 2TB database with file-based backups.




Figure 5 shows an RTO example for a 2TB database when Snapshot copy backups are used. With storage-based Snapshot copy backups, the RTO depends only on the database start time and on the forward recovery time. The restore operation 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 operations are performed. However, because of 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 2TB database with Snapshot copy backups.

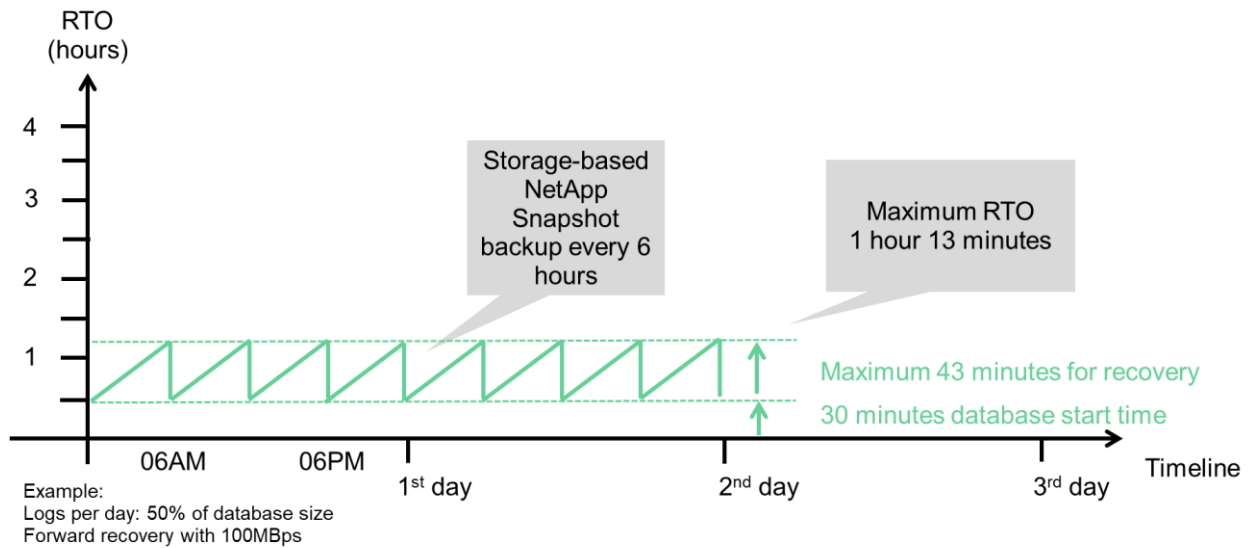
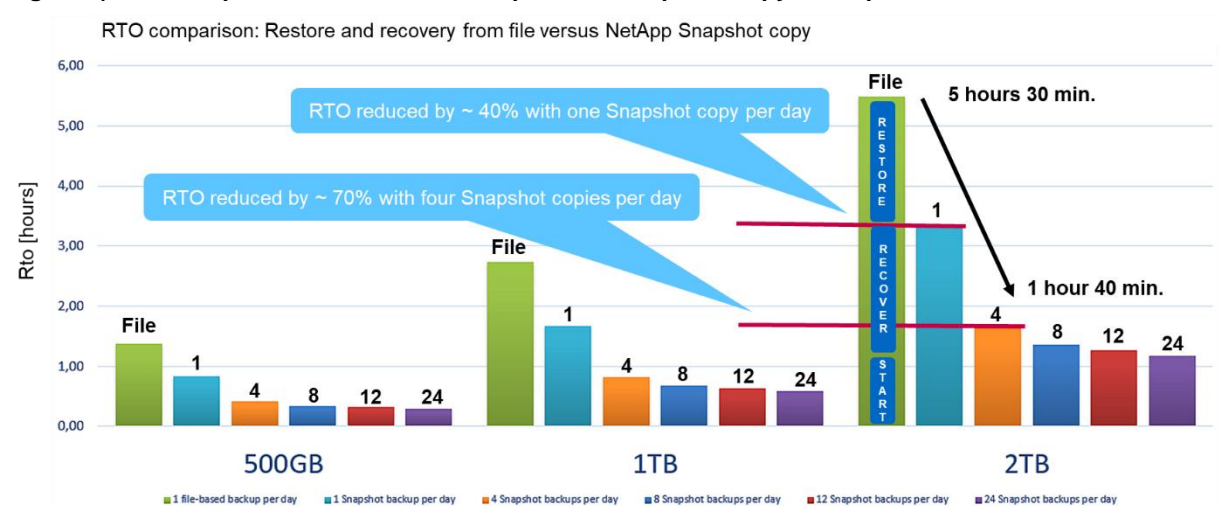


Figure 6 shows an RTO comparison of file-based and storage-based Snapshot copy backups for different database sizes and different frequencies of Snapshot copy backups. The green bar shows the file-based backup. The other bars show Snapshot copy backups with different backup frequencies.

With a single Snapshot copy data backup per day, the RTO is already reduced by 40% when compared with a file-based data backup. The reduction increases to 70% when four Snapshot copy backups are taken per day. Figure 6 also shows that the time saved starts to level off if you increase the Snapshot copy backup frequency to more than four to six Snapshot copy backups per day. Therefore, customers typically configure four to six Snapshot copy backups per day.

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



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

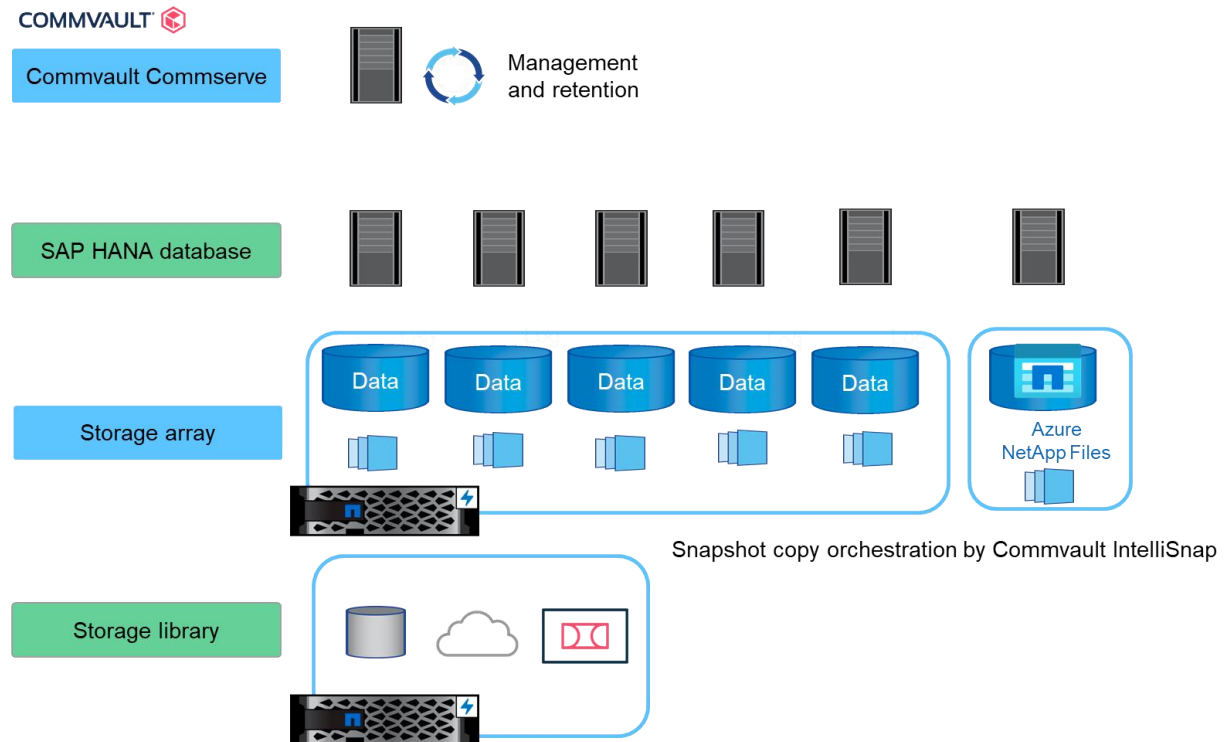
Note: The restore and recovery time is calculated based on the following assumptions. The database can be restored at 250MBps. The number of log files per day is 50% of the database size. For example, a 1TB database creates 500MB of log files per day. A recovery can be performed at 100MBps.

NetApp and Commvault SAP HANA backup solution

Solution component overview

Figure 7 shows an overview of the NetApp and Commvault components of the SAP HANA backup solution.

Figure 7) Solution component overview.



Commvault CommServe

A group of systems that use one Commvault license or account number is called a CommCell. This group is operated by a central management component called CommServe. CommServe coordinates and executes all CommCell operations, maintaining Microsoft SQL Server databases that contain all configuration, security, and operational history for the CommCell environment. There can be only one CommServe host in a CommCell environment. The CommServe software can be installed in physical, virtual, and clustered environments. For more information about the CommServe host, see [CommServe Server Overview](#).

SAP HANA database hosts

The SAP HANA Intelligent Data Agent and MediaAgent software are installed on the SAP HANA database physical hosts or in Azure SAP HANA VMs. These two components enable Commvault IntelliSnap to automate the creation of application-aware Snapshot copies. They also enable IntelliSnap to catalog Snapshot data to simplify the recovery of individual files or databases without the need for a collection of scripts and disparate Snapshot copy, backup, and recovery tools. IntelliSnap technology also supports the transfer of Snapshot copy-based backups to another on-prem NetApp storage system that is stored at a different location by using SnapVault and SnapMirror technologies. SAP HANA recoveries can be automated from each Snapshot copy either by mounting Snapshot copies or by leveraging NetApp SnapRestore technology (Commvault Snap Revert). Commvault also supports the automation of creating SAP HANA database clones from Snapshot copies.

In addition, long-term Snapshot copy backups can be moved to commodity disk, cloud, or tape to eliminate legacy backup solutions. At the end of its lifetime, the Snapshot copy metadata is removed from Commvault's internal databases as well as the actual Snapshot copies on the NetApp side. Therefore, Commvault offers full Snapshot copy lifecycle management and automation.

To protect large databases on-premises or in a hybrid or Azure cloud environment, using Snapshot copies that are managed by Commvault is an excellent choice. This option has a minimal impact on the SAP HANA database during a backup, and it greatly reduces RTO in the event of a restore and recover operation.

NetApp storage systems

NetApp AFF and FAS systems powered by NetApp ONTAP® data management software are certified for use with SAP HANA. NetApp systems provide the foundation for the backup solution, Snapshot copies, and replication that is based on SnapVault and SnapMirror technologies.

In addition, NetApp systems are built with innovative inline data reduction technologies such as inline compression, inline deduplication, and inline compaction.

ANF expands the NetApp data management capabilities offered by ONTAP into the public cloud, in this instance, Microsoft Azure.

Storage library

For on-premises environments, the joint NetApp and Commvault backup solution uses a NetApp AFF or FAS system as a storage library for the SAP HANA log backups, database integrity checks, and nondatabase file backups. In addition, this storage system is used as a SnapVault or a SnapMirror replication target of the primary storage system. As with the primary storage system, to provide space savings, innovative data reduction technologies such as compression and deduplication are also available at the secondary storage system.

In cloud environments, Commvault can leverage cloud storage containers such as Azure Blob. Multiple cloud containers can be configured to form a cloud library.

Solution components

The NetApp and Commvault backup solution for SAP HANA provides the following features:

- SAP HANA data file backup with storage-based Snapshot copies:
 - On-premises and in Microsoft Azure
 - Backup scheduling
 - Replication/DASH Copy to an off-site backup or cross-region disaster recovery location in the cloud
 - Retention management
 - GUI-based simplified restore and recovery
- SAP HANA streaming data file backup by using the SAP HANA Backint interface:
 - Database block integrity check
 - Backup scheduling
 - Retention management
 - Deduplication-aware secondary (DASH Copy) to disaster recovery location
- SAP HANA nondata volume backup with storage-based Snapshot copies:
 - Backup scheduling
 - Replication/DASH Copy to an off-site backup or cross-region disaster recovery location in cloud
 - Retention management
- SAP HANA log file backup by using the SAP HANA Backint interface:
 - Retention management

- DASH copy to disaster recovery location
- Automatic application of transaction logs during recovery operations

Database data file backups are executed by Commvault by using the SAP HANA Intelligent Data Agent and MediaAgent, which leverage an SAP HANA database backup save point. The Snapshot copies, which are created on the primary storage system, are therefore based on a consistent image of the SAP HANA database.

Commvault software enables the replication of consistent database images to an off-site backup or disaster recovery location by using SnapVault or SnapMirror technologies or as disk-based Snapshot copies through Commvault DASH Copy. Typically, different retention policies are defined for backups at primary and off-site backup storage. Commvault orchestrates and manages the retention at both sites by using NetApp functionality that is built into the storage platform.

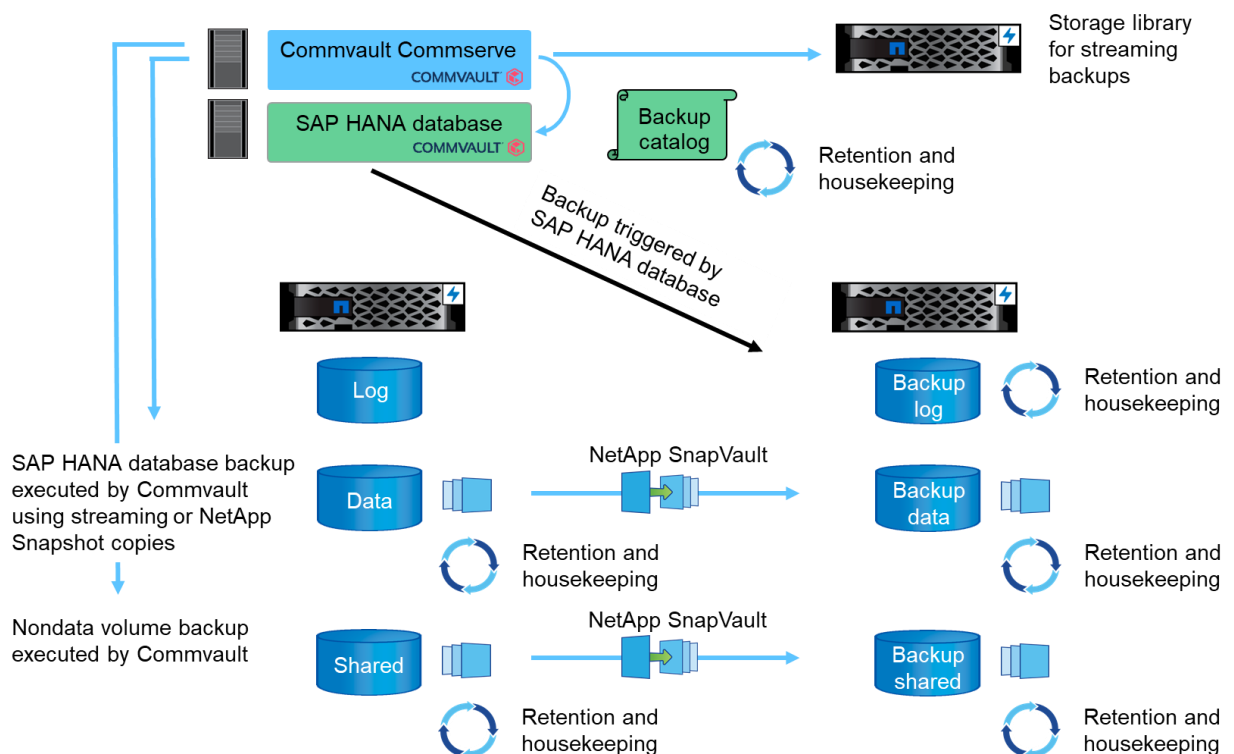
Commvault software also enables the back up of all SAP HANA data volumes by using the SAP HANA Backint interface, which is intended for streaming backups. To enable individual retention and protection policies and to collect all required information for SAP HANA disaster recovery, the backup of SAP HANA nondata volumes must be scheduled independently from the database data backups.

The SAP HANA database can automatically execute log backups. These log backups are written to Commvault through the Backint interface as streaming backups. Commvault software can mirror or DASH Copy these backups to a secondary site (and if necessary, to a tertiary site) for disaster recovery purposes. The software catalogs the backups so that roll-forward can be performed afterward; for example, restore of a data volume from a Snapshot copy backup.

SAP recommends that you combine storage-based Snapshot copy backups with a weekly streaming-based backup to execute a block integrity check. Based on your configurable retention policies, Commvault manages the housekeeping of streaming and Snapshot copy-based data file backups at the primary storage, streaming log file backups, and the SAP HANA backup catalog.

Figure 8 shows an overview of the database and log backup configuration, where the log backups are written to an NFS mount of the off-site backup storage.

Figure 8) Database and log backup configuration overview.



When executing a storage-based Snapshot copy backup of nondata volumes, Commvault software performs a streaming backup that stores the data in a storage library, which is defined within the used storage policy. This stored data can in turn be mirrored to a secondary site.

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

1. Creates an SAP HANA backup save point to create a consistent image on the persistence layer.
2. Creates a storage Snapshot copy of the data volume.
3. Registers the storage Snapshot copy backup in the SAP HANA backup catalog.
4. Releases the SAP HANA backup save point.
5. Executes a SnapVault or SnapMirror update for the data volume, if configured.
6. Deletes storage Snapshot copies from the primary storage and from SnapVault or SnapMirror based on the defined retention policy.

Note: Both the primary and secondary Snapshot copies can be assigned individual retention times based on Commvault's Storage Policy concept.

7. Whenever a backup is deleted based on the retention policy or is manually deleted, Commvault deletes all log backups that are older than the oldest data backup. Log backups are deleted from the backup device and in the Commvault backup catalog.

Note: The deletion of storage Snapshot copies at the off-site backup storage is executed by NetApp ONTAP data management software. The deletion is based on the defined retention in the ONTAP protection relationship configuration.

Supported SAP HANA releases and configurations

Commvault version 11.22 and later supports SAP HANA single-host and multiple host configurations that use NFS- or FC-attached NetApp storage systems (AFF and FAS) and ANF (NFSv4.1) by using streaming-based backups and Snapshot copies.

Commvault V11 supports the following SAP HANA releases:

- SAP HANA single container:
 - SAP HANA 1.0 SPS7 and later
 - SAP HANA 2.0 up to SPS0
- SAP HANA multitenant database container (MDC) single tenant and multiple tenants:
 - SAP HANA 2.0 SPS1 and later

Note: Storage-based Snapshot copy backups for SAP HANA MDC with more than one tenant are not supported for SAP HANA scale-out systems.

Capacity requirements for Snapshot copy 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 in addition to the block changes are written to disk. Data from typical customers shows a daily change rate between 10% and 50%.

Installation and configuration overview

Commvault installation

Perform the base Commvault software installation as instructed in the [Commvault installation documentation](#).

Configuration steps

After the base Commvault software installation process is complete, configure the following settings:

1. Set up a storage or Azure cloud library for streaming backup data like SAP HANA logs.
2. Set up NetApp storage virtual machines (SVMs) for NetApp Snapshot copy management or set up ANFs account and capacity pools.
3. Set up a Commvault storage policy/plan.
4. Install SAP HANA agent on the client systems and/or Azure VMs forming a SAP HANA instance.
5. Create new SAP HANA instance in Commvault CommandCenter and add all SAP HANA nodes to which it belongs.
6. Enable IntelliSnap in the SAP HANA instance.
7. Create a new Snapshot subclient in the SAP HANA instance, turn on IntelliSnap, and select NetApp Snap Engine for On-premises Environments.

The rest of this section provides details about each of these settings.

Set up a storage library for streaming backup data

You must set up a Commvault disk or cloud storage library for streaming backup data. A disk library is a virtual library that is associated with one or more mount paths. In this solution, a volume in the NetApp storage system is mounted by using NFS to the SAP HANA host where the Commvault MediaAgent is installed. Streaming backup data is used for SAP HANA log backup, and full data backup is used as a database consistency check once per week.

Set up NetApp SVMs for Snapshot copy management

For Commvault to communicate with a NetApp storage system, it must be able to communicate with the SVMs. For this reason, the SVMs must be defined in Commvault Array Management, as shown in Figure 9. When you use NetApp SnapMirror or SnapVault technology, you must define both the primary and secondary arrays.

Figure 9) Commvault array management.

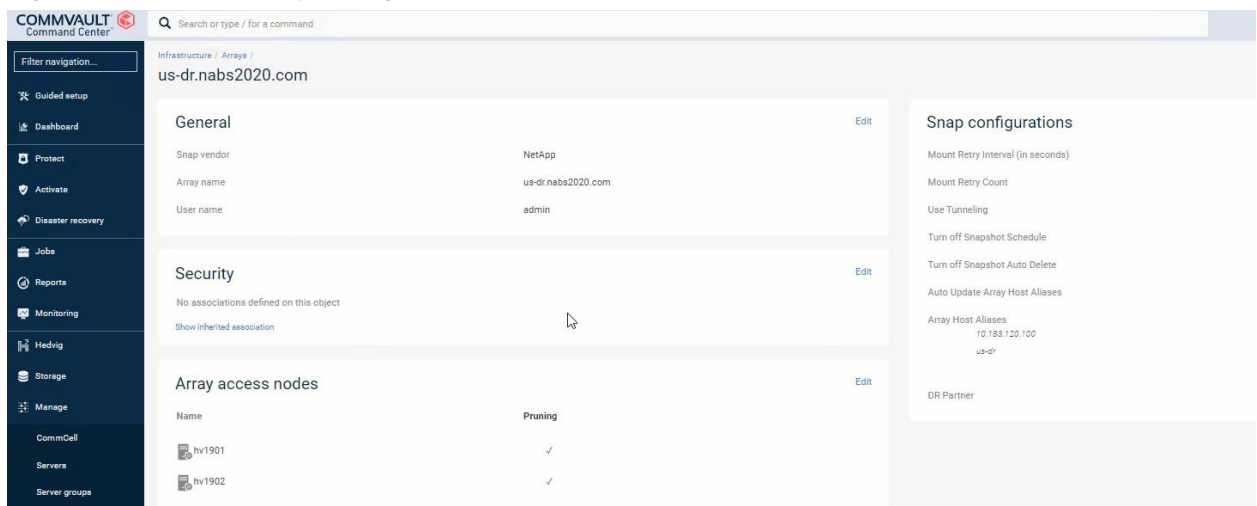
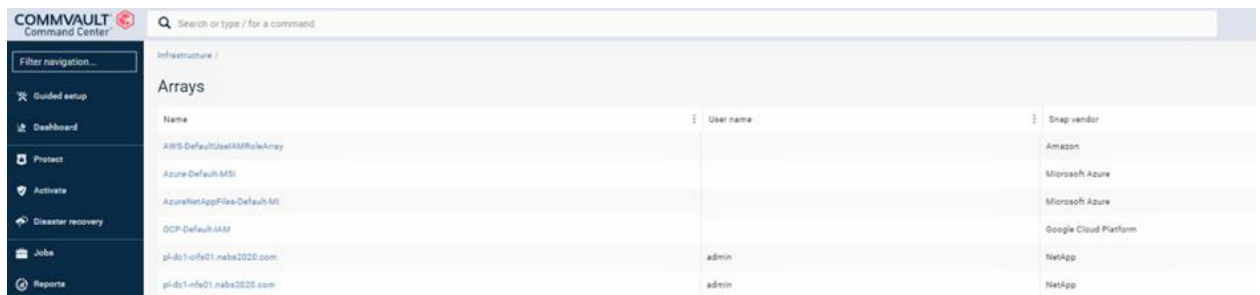


Figure 10 shows the end result after you have defined all NetApp systems in Commvault.

Figure 10) Commvault array management completed.



Name	User name	Snap vendor
AWS-DefaultIaaSAMIRoleArray		Amazon
Azure-Default-MSI		Microsoft Azure
AzureNetAppFiles-Default-MSI		Microsoft Azure
OCP-Default-IAM		Google Cloud Platform
pl-dc1-nfs01-nabx2020.com	admin	NetApp
pl-dc1-nfs01-nabx2020.com	admin	NetApp

Configure IntelliSnap for Azure NetApp files

In addition to the steps listed in the section “Configuration steps,” the following actions need to be performed:

1. On the SAP HANA VM on which IntelliSnap operations will run, do the following:
 - a. Enable managed identity access.
 - b. Set the contributor role on the level of the respective Azure resource groups.
2. Create a new Snapshot copy subclient in the SAP HANA instance, turn on IntelliSnap, and select Microsoft Azure as Snap Engine.

For more information, see the [Commvault documentation](#).

Set up a storage policy

To set up and configure a storage policy/plan, you must define a data protection strategy, as described in section, “Data protection strategy.

For Snapshot copy-based backups, use either of the following options to configure replication of Snapshot copy-based backups on NetApp storage systems:

- NetApp Active IQ Unified Manager:
 - When you use Active IQ Unified Manager, it creates and manages the SnapMirror or SnapVault relationships for you after you define the provision pools and the replication target of the source SVM. Within provision pools, aggregates are defined, which can be used for automatic provisioning of volumes—in this case, for volumes that are used as the replication target.
 - To use this option, specify the Active IQ Unified Manager server name or IP address when you create the storage policy.
- Open Systems Data Protection:
 - In the storage policy, enter the NetApp SVM host name or IP address + user name and password directly.

Note: This option requires you to manually manage the SnapMirror or SnapVault relationships.

For more information, see the [Commvault documentation](#).

Set up agents on the client systems

To set up agents on the client systems, complete the following steps:

1. Push the Commvault software to the client system or VM.
2. Create an SAP HANA instance
3. Configure the subclients.

After you complete these steps, configure the Commvault SAP on HANA agent, as described in the [Commvault documentation](#).

Data protection strategy

Before you configure a Commvault storage policy to back up SAP HANA systems, you must define the data protection strategy based on the RTO and RPO requirements of the various SAP systems.

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

The parameters that you must define are as follows:

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

Table 1 shows an example of the data protection parameters for the system types: Production, Development, and Test. For the Production and Development systems, a high backup frequency has been defined, and the backups are replicated to an off-site backup location once per day. The Test systems have lower requirements and no replication of the backups.

Table 1) Data protection parameters.

Parameters	Production systems	Development systems	Test systems
Backup frequency	Every four hours	Every four hours	Every four hours
Primary retention	Two days	Two days	Two days
Block integrity check	Once per week	Once per week	No
Replication to off-site backup site	Once per day	Once per day	No
Off-site backup retention	Two weeks	Two weeks	n/a

Table 2 shows the policies that must be configured for the data protection parameters.

Table 2) Policies based on data protection parameters.

Parameters	Policy LocalSnap	Policy LocalSnapAndSnapVault	Policy streaming
Backup type	Snapshot copy-based	Snapshot copy-based	File-based
Schedule frequency	Hourly	Daily	Weekly
Primary retention	Count = 12	Count = 2	Count = 1
SnapVault replication	No	Yes	n/a

The policy `LocalSnap` is used for the Production, Development, and Test systems to cover the local Snapshot copy backups with a retention of two days.

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

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

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

In the resource configuration, the schedule is defined for Production and Development:

- Production: Schedule every day.
- Development: Schedule every day.

The policy `Streaming` is used for the Production and Development systems to cover the weekly block integrity check by using a file-based backup.

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

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

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

With our example, for each Production and Development system, a retention of two weeks is defined at the off-site backup storage.

Note: In our example, protection policies and retention for SAP HANA database resources and nondata volume resources are not different.

Backup retention management and housekeeping of log backups

In Commvault, data retention is managed through storage policies/plans that define where a backup is stored and for how long, as shown in the example in Figure 11. This procedure can handle Snapshot copy, SnapVault, and streaming-based backups in the same way. All NetApp arrays must be entered as backup destinations and associated with an individual retention time. Commvault orchestrates the data movement between the backup destinations. The deletion of expired backups is a fully automated background process.

Figure 11) Commvault Plan example.

The screenshot displays the Commvault Command Center interface for a backup plan named "IP - Files". The left sidebar shows navigation options like Filter navigation, Backup setup, Dashboard, Protect, Activate, Disaster recovery, Jobs, Reports, Monitoring, Heklog, Storage, Manage, CommCell, Servers, Server groups, Companies, Plans, Tags, Infrastructure, Regions, License, Customization, System, Network, and Security. The main content area is divided into several sections:

- Backup content:** Lists backup content for Windows, Mac, and Unix, all associated with "All contents".
- Security:** Shows a table with columns for User/Group and Role. The "master" user is assigned the "Plan Creator Role".
- Override restrictions:** A table with columns for Storage pool, RPO, and Folders to backup, all with "Override optional" status.
- RPO:** Shows the backup frequency as "Runs every 1 Day(s) at 9:00 PM". It includes options for "Add full backup", "Backup window" (Monday through Sunday - All day), and "Full backup window" (Monday through Sunday - All day).
- Secondary copy schedule:** Set to "Automatic schedule".
- Snapshot options:** Includes "Backup Copy" (checked) and "Backup copy frequency (in HH:MM)" set to "4 hour(s)".
- Backup destinations:** A table listing backup destinations with columns for Name, Storage, Retention period, and Source.

Name	Storage	Retention period	Source
1 - Source Snapshot Snapshot primary	APF_JH2	1 Month	
2 - SnapMirror Copy Deduplicate	APF_JH2	1 Month	1 - Source Snapshot
3 - SnapVault Copy Vault/Replica	APF_JH2	90 Days	2 - SnapMirror Copy
4 - Flash Disk Copy	APF_JH2	1 Month	
5 - Object Storage Copy	StorageGRID	6 Months	4 - Flash Disk Copy

Conclusion

The NetApp and Commvault backup solution for SAP HANA is changing today's backup and recovery landscape. Commvault IntelliSnap software together with NetApp on-premises or ANF software combines simplified manageability, power, and flexibility for the SAP HANA landscape, including virtual environments. Commvault IntelliSnap for NetApp integrates with NetApp Snapshot technology in a virtually seamless way for fast and efficient backup operations in on-premises, hybrid, and in-cloud SAP environments. It also integrates on-premises with NetApp SnapVault and SnapMirror software to support content cataloging and data movement to tape-based media.

Commvault IntelliSnap for NetApp offers single-interface management for backup and recovery workflows for NetApp ONTAP and much more. Because it centralizes all these functions and offers policy-based management and granular recovery across all supported workloads, Commvault IntelliSnap for NetApp is a compelling enterprise backup and recovery solution.

Where to find additional information

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

- TR-4018: Integrating NetApp ONTAP Systems with SAP Landscape Management
<https://www.netapp.com/us/media/tr-4018.pdf>
- TR-4646: SAP HANA Disaster Recovery with Storage Replication
<https://www.netapp.com/us/media/tr-4646.pdf>
- NetApp Product Documentation
<http://docs.netapp.com>
- Commvault General Documentation
<http://documentation.commvault.com/commvault/v11/article?p=33739.htm>
- Commvault SAP HANA Documentation
https://documentation.commvault.com/11.21/essential/86678_sap_hana.html
https://documentation.commvault.com/11.21/expert/35970_getting_started_for_sap_hana_intellisnap.html
- Commvault IntelliSnap for NetApp
https://documentation.commvault.com/11.21/expert/33739_getting_started_with_netapp_storage_array.html
- Commvault IntelliSnap for HANA on Azure NetApp Files
https://documentation.commvault.com/11.21/expert/127985_azure_netapp_files_for_sap_hana.html
- SAP Support Portal
<https://support.sap.com/en/index.html>

Version history

Version	Date	Document Version History
Version 1.0	August 2018	Initial release.
Version 2.0	March 2021	General refresh ANF support

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.

Copyright Information

Copyright © 2021 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4711-0321