Hitachi Data Protection Suite powered by Commvault with Hitachi Content Platform for SAP HANA

Reference Architecture Guide

MK-SL-265-00
February 2023

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@HitachiVantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision history

Changes	Date
Initial release	February 24, 2023

Reference architecture guide

Continued growth of data in the computing industry has led to an increased demand on storage capacity and optimization of business-critical processes, such as backup and recovery. Furthering the complexity of managing growth are the critical needs surrounding data security, risk mitigation, loss prevention, and compliance with industry standards. Hitachi Data Protection Suite (HDPS), powered by Commvault, can synchronize and protect your data and applications across the globe, making it easier for you to meet your recovery time objectives (RTO) and recovery point objectives (RPO).

You can easily integrate HDPS with Hitachi Content Platform (HCP), Hitachi Global-active Device, Hitachi Unified Compute Platform (UCP), and Hitachi Virtual Storage Platform (VSP) to gain even more control over your distributed IT infrastructure.

This reference architecture guide describes the backup, recovery, and auxiliary copy of SAP HANA Appliance and SAP HANA Tailored Data Center Integration (TDI) databases in a scale-up configuration using HDPS, powered by Commvault, HCP, and Hitachi Content Platform for cloud scale (HCP for cloud scale) for long-term data retention.

The SAP HANA application is an in-memory database that keeps the bulk of its data in memory and uses persistent storage to provide a backup in case of a failure. However, if the persistent storage itself is damaged, for example because of drive failures or database corruption, then additional backups are needed to protect the data against these types of loss.

The loss of business-critical system resources and services, such as the SAP HANA platform, can translate directly into lost revenue. Therefore, it is critical to perform SAP HANA database backups to a secondary storage media, such as external storage or cloud-based storage, and to be able to restore the SAP HANA database from the backup in case of a failure.

Backup solution

Continuous use of SAP enterprise systems produces massive amounts of data. These large databases increase storage costs, and impact database performance, the backup window, and migration downtime.

Hitachi Vantara and Commvault offer a validated solution that brings together HDPS, powered by Commvault, and HCP for cloud scale. Combining these technologies results in a highly scalable, flexible, and resilient cloud-ready data protection platform. This platform can meet the most stringent performance requirements – with a single interface that breaks down those silos and natively manages SAP HANA, with the same SLAs, from one location.

HDPS and HCP for cloud scale offer better return on investment by reducing downtime, preventing data loss, more efficiently storing data, and easily scaling as your data volumes grow. A highly secure framework limits access to backup data while AI-based anomaly detection alerts and remediates when an attack is suspected. Reserve primary storage resources for near-term capacity and tier enterprise data backups using policies, to HCP for long-term, cost efficient data retention.

HCP for cloud scale is a software-defined object storage solution that is based on a massively parallel microservice architecture and is compatible with the Amazon Simple Storage Service (S3) application programming interface (API). HCP for cloud scale is especially well suited to service applications requiring high bandwidth and compatibility with Amazon S3 APIs. It can federate S3-compatible storage from virtually any private or public source, and present the combined capacity in a single, centrally managed, global namespace.

Primary backups are tiered to HCP for cloud scale based on aging or capacity policies and are further made immutable, encrypted, and self-healing by setting policy-driven data management capabilities. This helps you achieve the highest speed, lowest cost, and fully secured recovery data that is managed in a coordinated fashion between the HDPS and the HCP for cloud scale platform.

For more information, see *Safeguard Your Enterprise Data With Hitachi Data Protection Suite* at https://www.hitachivantara.com/en-us/pdf/solution-profile/data-protection-suite-powered-by-commvault-solution-profile.pdf.

Use this document to understand the example reference architecture for archival, backup, and recovery of SAP HANA in a scale-up configuration, and to help with deployment of the configuration.

This technical paper assumes familiarity with the following:

- Storage area network-based storage systems
- General storage and backup concepts
- General network knowledge
- SAP HANA
- SAP HANA platform OS
- Common IT best practices

Note: These procedures were developed in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

Solution overview

This SAP HANA backup and recovery solution using SAP Backint and IntelliSnap for SAP HANA Appliance uses the following components:

- Hitachi Advanced Server DS220 for HANA Appliance
- Hitachi Advanced Server DS220 for HANA TDI
- Hitachi Virtual Storage Platform (VSP) F350 for HANA TDI
- Hitachi Virtual Storage Platform (VSP) G700 for HDPS and HCP
- Hitachi Content Platform for cloud scale
- Hitachi Data Protection Suite, powered by Commvault
- SAP High-Performance Analytic Appliance (HANA)

The following illustration shows the high-level design of this solution.

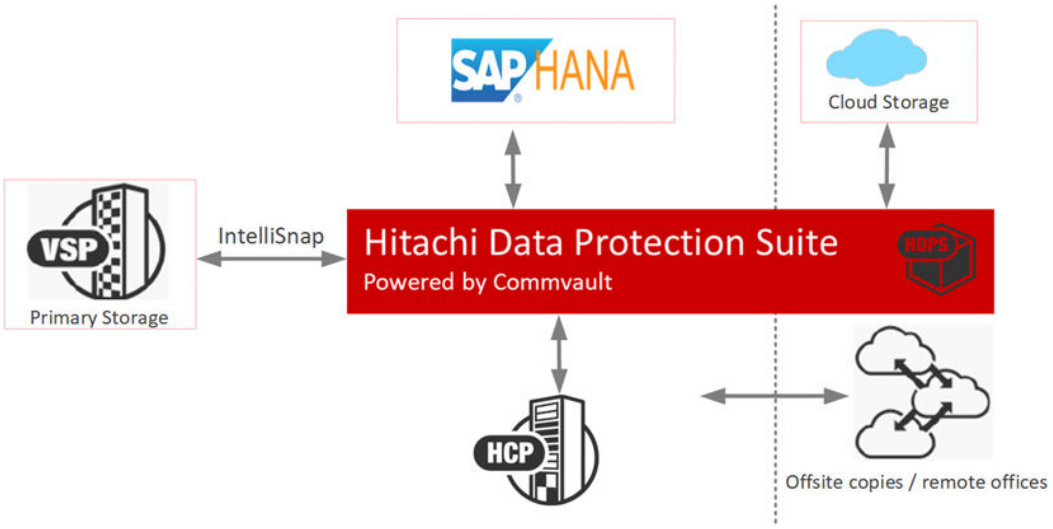

Key solution elements

The key hardware and software components used in this reference architecture are described in the following tables. For detailed component information, see Product descriptions (on page 30)

Hardware elements

The following table lists the hardware configuration used in this solution.

Hardware	Quantity	Configuration	Role	Implementation Type
Hitachi Advanced Server DS220	1	▪ CPU – 2 × Intel(R) Xeon(R) Plaintum 8176 CPU @ 2.10GHz ▪ Memory - 768 GB ▪ RAID Controller - Broadcom QS-3916 Type B Mega RAID (4G DDR) storage controller card ▪ Disks - 6 × 1.92 TB Intel S4510 SATA SSD ▪ Network – 2 × Intel E810-XXVDA 2-port 10/25 GbE network cards	SAP HANA Appliance server	HANA Database with internal disks
Hitachi Advanced Server DS220	1	▪ CPU – 2 × Intel(R) Xeon(R) Platinum 8176 CPU @ 2.10GHz ▪ Memory - 768 GB ▪ Network – 2 × Intel E810-XXVDA 2-port 10/25 GbE network cards ▪ Fibre Channel Card - 2-port Broadcom Emulex LPE 31002 PCIe 16 Gbps HBA card ▪ LUNS allocated from Hitachi VSP F350	SAP HANA TDI Server	HANA Database with external storage
Hitachi Advanced Server DS220	1	▪ CPU – 2 × Intel(R) Xeon(R) Platinum 8276 CPU @ 2.20GHz ▪ Memory - 768 GB ▪ Network – 2 × Intel E810-XXVDA 2-port 10/25 GbE network cards	vSphere 7.3 U3	HDPS and HCP VMs

Hardware	Quantity	Configuration	Role	Implementation Type
		▪ Fibre Channel Card - 2-port Broadcom Emulex LPE 31002 PCIe 16 Gbps HBA card ▪ LUNS allocated from Hitachi VSP F700		
Hitachi Virtual Storage Platform F350	1	▪ CTL: 1 pair ▪ 16/32 Gbps 4-port ▪ CHB: 1 pair ▪ MPU: 1 pair ▪ Cache: 128 GB	HANA Volumes	Storage for HANA TDI
Hitachi Virtual Storage Platform F700	1	▪ CTL: 1 pair ▪ 16/32 Gbps 4-port ▪ CHB: 1 pair ▪ MPU: 1 pair ▪ Cache: 512 GB	ESXi Datastores for HDPS VMs	Use for Performance Tier storage for HDPS
Cisco Nexus 92348 switch	1	48 × 1 GbE ports	Optional switch for management network.	All (optional)
Cisco Nexus 93180YC-FX	2	48 × 10/25 GbE ports	Optional switches for the client network or additional backup network.	All (optional)
Brocade 6505	1	24 × 16 Gbps Fibre Channel ports	Fibre Channel Switch for Storage Network	All (optional)

Software elements

The following table lists the software used in this solution.

Software	
Operating systems	▪ SUSE Linux Enterprise Server (SLES) Linux Enterprise Server for SAP Applications 15 SP3 ▪ RedHat Enterprise Linux for SAP 8.4 ▪ Windows Server 2019 for HDPS components (CommServe and Media Agent)
Database	▪ SAP HANA 2.0 SPS06
DB Management	▪ SAP HANA Studio
Benchmark Suite	▪ BW4H data
Hypervisor	▪ VMware vSphere 7.0 U3
Object storage	▪ Hitachi Content Platform 9.4.0.267 ▪ Hitachi Content Platform for cloud scale 2.4.1.2
Data Protection	▪ Hitachi Data Protection Suite powered by Commvault 11.28.32
Agent	▪ Commvault Intelligent Data Agent for HANA
Storage Management	▪ Command control interface
Storage Snapshots	▪ Hitachi Thin Image (HTI)
Storage Replication	▪ Hitachi True Copy (HTC)

Note: For the lab testing, we used VMware vSphere for HDPS components (CommServe and Media Agent) as VMware VMs and performance tier storage attached as datastores. However, in a production environment, it is recommended to use dedicated Hitachi Advanced Servers for HDPS components.

System and hardware requirements

For small and medium environments, it is recommended to deploy the HDPS software in either an all-in-one configuration or a server and storage (distributed) configuration. For hardware requirements for large and extra-large environments, see component sizing topics at https://documentation.commvault.com/2022e/expert/1644_commcell_sizing.html.

All-in-one configuration

- Windows Server 2019 operating system
- 16 CPU cores
- 32 GB RAM
- 2 TB of working space for the CommServe database, the deduplication database (DDB), index, and job results (SSD disks are recommended)

Server and storage (distributed) configuration

- Server 1 with CommServe, Web Server, Web Console, Workflow Engine installed
  - Windows Server 2019 operating system
  - 8 CPU cores
  - 32 GB RAM
  - 500 GB for SQL database (SSD is recommended)
- Server 2 with MediaAgent and access nodes installed
  - Windows Server 2019 operating system or Linux operating system, such as Red Hat and Ubuntu
  - 8 CPU cores
  - 16 GB RAM
  - 1.5 TB of working space for DDB, index, and job results (SSD is recommended)
  - To scale horizontally, several access nodes can be added. The access nodes are co-located with the data either in the cloud or on-premises.

Note: Many hyperlinks for access to Commvault documentation pages require expert login credentials.

Scaling out backup infrastructure horizontally or vertically

When you add new regions to an existing configuration, scale up your backup infrastructure based on the increased workload on your backup environment. We ideally recommend starting with a small configuration, and then scale up to medium, large, and extra-large configurations before you scale out and add new nodes. For more information about new regions, see https://documentation.commvault.com/2022e/essential/107547_regions.html.

Solution design

The following detailed design for HDPS powered by Commvault for SAP HANA using IntelliSnap and SAP Backint methods, HCP, and HCP for cloud scale includes the following:

- SAP HANA Appliance and TDI installation
- HDPS server and storage (distributed) configuration
  - Server 1 with CommServe, Web Server, Web Console, and Workflow Engine installation
  - Server 2 with MediaAgent and Access Nodes installation
- Storage library creation for primary and secondary copy
- Storage system and command control interface configuration for IntelliSnap method

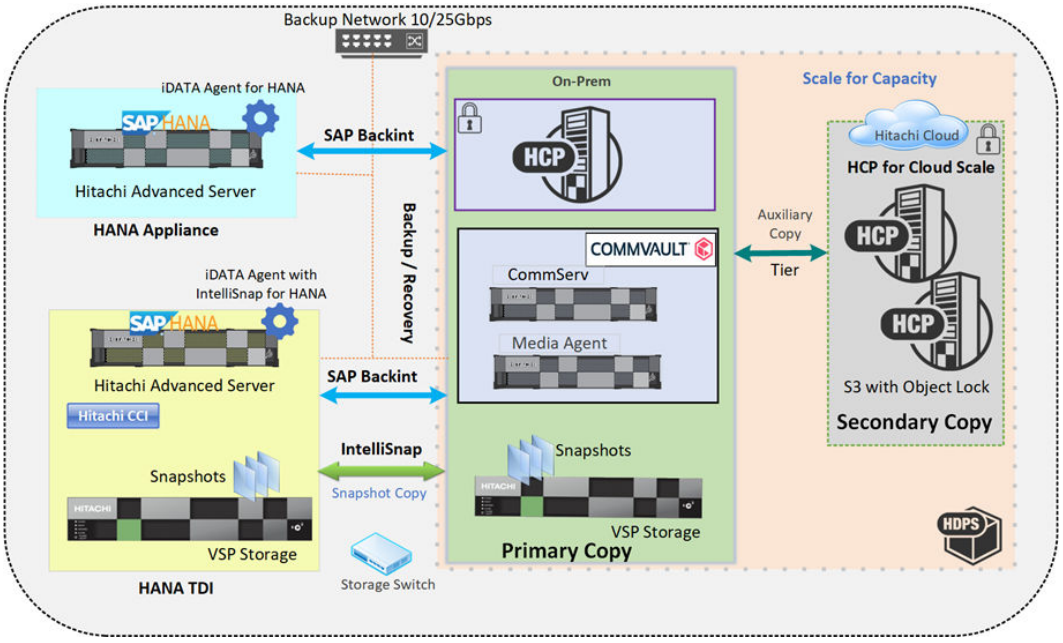

The following is a list of HDPS components and the configuration used for SAP HANA database backup:

- SAP HANA 2.0 SPS06 was installed on the SLES 15 SP3 and RHEL 8.4 operating systems
- BW/4HANA data was loaded into the database before backup and recovery were performed; there were 910GB data loaded in SAP HANA Appliance (SLES) and approximately 550 GB data loaded in SAP HANA TDI (RHEL) servers.
- SAP HANA Studio was used to perform backup and recovery operations. Commvault CommCell Browser and Command Center can also be used for SAP HANA backup and restore with more options and preferred tools.
- HDPS components CommServe and Media Agent were installed on Windows Server 2019. Both components were installed on separate servers.

  See Installation of Commvault Software at https://documentation.commvault.com/2022e/expert/1716_installations.html for installation of CommServe, and dedicated Media Agent at https://documentation.commvault.com/2022e/expert/1798_mediaagent_installations.html for configurations.

  Note: There should be a dedicated Media Agent Server in the environment or data center. Then, if one Media Agent exists in the primary data center there should be a separate Media Agent for disaster recovery (DR) datacenter or cloud environments.
- While adding client HANA servers in CommServe, SAP HANA iData Agent, and Media Agent will install automatically.
- Use SAP HANA iData Agent to communicate between SAP HANA Database, CommServe, and Media Agent and modify global.ini so that SAP Backint can use Commvault for data and log backup and recovery. This will create a default parameter file /opt/hds/iDataAgent/param for SAP HANA Database. For HANA TDI, install IntelliSnap software also for snapshot-based backup. For more information, see https://documentation.commvault.com/2022e/expert/22305_sap_hana.html.
- Set up the command control interface with a raw command device for communicating between CommServe and Hitachi VSP storage to take IntelliSnap backups (snapshot-based backup on VSP Storage). For more information, seehttps://knowledge.hitachivantara.com/Documents/Management_Software/Command_Control_Interface.
- Set up a storage pool for Hitachi Thin Image in the VSP storage systems (used for HANA TDI) for creating volume snapshots of HANA Data Volumes using IntelliSnap. For more information, see https://www.hitachivantara.com/en-us/pdf/datasheet/thin-image-snapshot-software-datasheet.pdf.
- Set up Hitachi TrueCopy for replication of Hitachi Thin Images from HANA TDI/HANA TDI Storage to Performance Tier Storage. For more information, see https://www.hitachivantara.com/en-us/pdf/datasheet/truecopy-remote-replication-software-datasheet.pdf.
- Set up an S3-compatible bucket for SAP HANA in HCP for cloud scale backup as secondary copy storage using Commvault Auxiliary Copy for data tiering. For more information, see https://www.hitachivantara.com/en-us/pdf/white-paper/hcp-for-cloud-

scale-whitepaper.pdf and https://documentation.commvault.com/2022e/expert/11490_additional_copies_of_backup_data.html.

- Configure Hitachi Content Platform VM Access Storage Node to use for testing of on-premises primary object storage. For more information, see https://www.hitachivantara.com/en-us/pdf/white-paper/content-platform-architecture-fundamentals-whitepaper.pdf.

Note: The IntelliSnap backup method supports only the HANA Data volume of HANA TDI servers, which exist on Hitachi VSP storage systems. For HANA Log backup, SAP Backint is used.

Storage configuration in HDPS

You can manage various storage components such as MediaAgents, storage targets, and the storage architecture. A MediaAgent is a data transmission manager in the CommCell environment that manages data movement and data storage targets, including tape, disk, and cloud. The tape, disk, and cloud storage provide scalable storage solutions where multiple cross-platform MediaAgents can access the shared storage. Distributed storage can be configured as a disk library. You can configure both deduplicated and non-deduplicated storage for tape, disk, and cloud storage types. However, hyperscale supports only deduplicated storage. The storage can be scaled-out and you can set the retention criteria for the backed-up data.

The following figure shows different storage types.

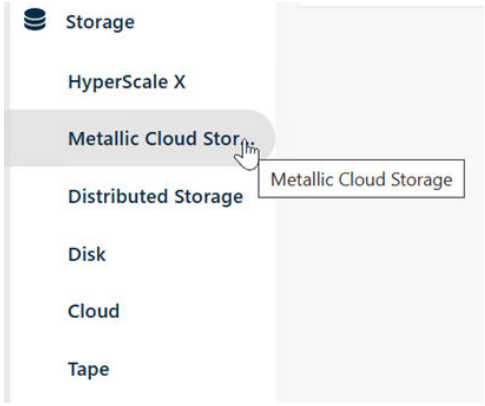

Disk storage

Disk storage provides a scalable storage solution where multiple cross-platform MediaAgents can access shared storage. You can increase the storage capacity by adding more backup locations and can configure both deduplicated and non-deduplicated disk storage.

Use the following procedure to add Disk Storage using Command Center.

Procedure

1. From the navigation pane, go to **Storage** > **Disk**.
2. In the upper-right corner of the page, click **Add**.

3. In the **Name** field, type a name for the disk storage.
4. In the **Backup location** section, click **Add**.
   a. From the MediaAgent list, select the MediaAgent that will write the data to the disk storage.
   b. To set the disk access path, use one of the following options:
      - To use a local disk as the disk access path, click **Local**.
      - To use a network drive as the disk access path, click **Network** and then enter the username and password needed to access the backup location.
   c. To set a login credential, do one of the following:
      - To use a pre-defined credential, move the **Use saved credentials** toggle key to the right, and then select the pre-defined credential from the **Name** list.
      - To add a new credential, click the + sign to create new credentials to access a network drive.

        See *Creating a Credential to Access a Network Drive* at https://documentation.commvault.com/2022e/essential/116877_creating_new_users_for_accessing_disk_storage.html for more information.
   d. In the **Backup location** field, type or browse the full path name to the storage location.
   e. Click **Add**.

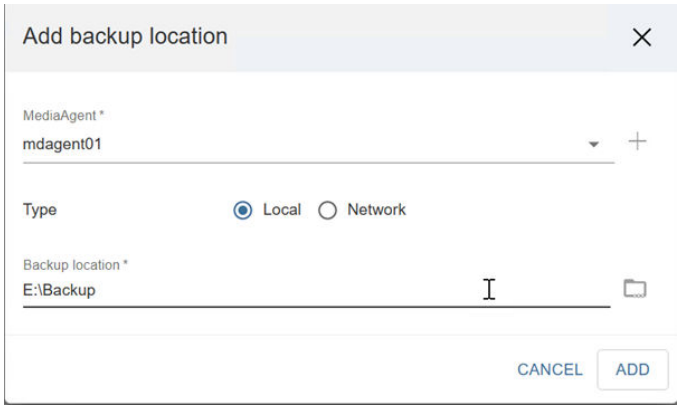

5. To enable deduplication on the storage, complete the following steps:
   a. Move the **Use deduplication** toggle key to the right.
   b. In the **Deduplication DB location** section, click **Add**.
   c. From the **MediaAgent** list, select the MediaAgent that will write the data to the DB location.
   d. In the **Deduplication DB location** field, type the full path name or browse and select the path to the DB location.
   e. Click **Add**.

   For more information about configuring additional DDB partitions, see *Configuring Additional Partitions for a Deduplication Database* at https://documentation.commvault.com/2022e/expert/12455_configuring_additional_partitions_for_deduplication_database.html.

6. Click **Save**.

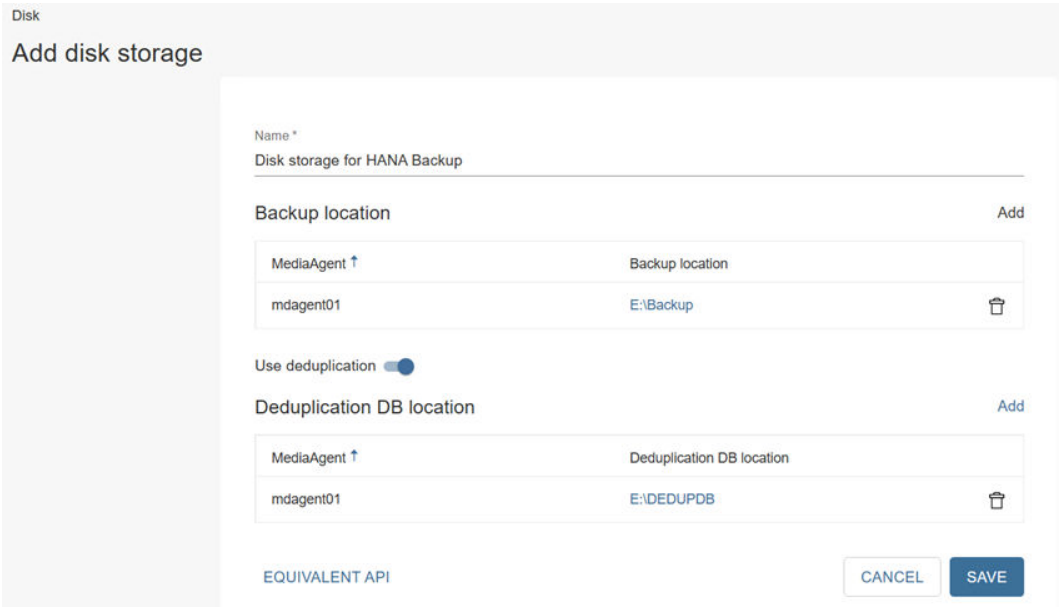

Cloud storage

Cloud storage provides a scalable and easy to configure storage solution where multiple cross-platform MediaAgents can access shared storage. The storage capacity can be scaled-out on demand by simply adding more backup locations. The added MediaAgents or the backup locations are automatically shared using the network.

To verify that the storage products are supported by HDPS software, see Cloud Storage Products at https://documentation.commvault.com/2022e/essential/9236_supported_cloud_storage_products.html .

Use the following procedure to configure cloud storage.

Procedure

1. From the navigation pane, click **Storage** > **Cloud**.
2. In the upper right of the page, click **Add** to add cloud storage.
3. Enter the details needed to add the cloud storage.

   For more information, see *Hitachi Content Platform* https://documentation.commvault.com/2022e/expert/117242_hitachi_content_platform.html or *Hitachi Content Platform for Cloud Scale* https://documentation.commvault.com/2022e/expert/91295_hitachi_content_platform_for_cloud_scale.html.

   a. From the **Type** list, select the cloud vendor.
   b. In the **Name** field, type the name of the storage.
   c. From the **MediaAgent** list, select the MediaAgent that will write the data to the cloud storage.
   d. In the **Service host** field, type the name of the cloud server host.
   e. From the **Credentials** drop down, choose a pre-defined credential.

      Optionally, click the **+** sign to create a new credential to access a cloud library.

f. In the **Bucket** or **Namespace** field, type the name of an existing bucket (or namespace) that you want to use.

g. Slide the Use deduplication toggle key to the right to enable deduplication on the storage.

h. Click **Add** to add a Deduplication DB location.

- From the **MediaAgent** list, select the MediaAgent associated with the deduplication database.
- In the **Deduplication DB location** field, type a path to the deduplication database in the selected MediaAgent.

Optionally, click the folder button to select a pre-existing Deduplication DB location.

i. Select the deduplication database and then click **Save**.

4. Click **Add** to exit the **Add Deduplication DB location** dialog box.

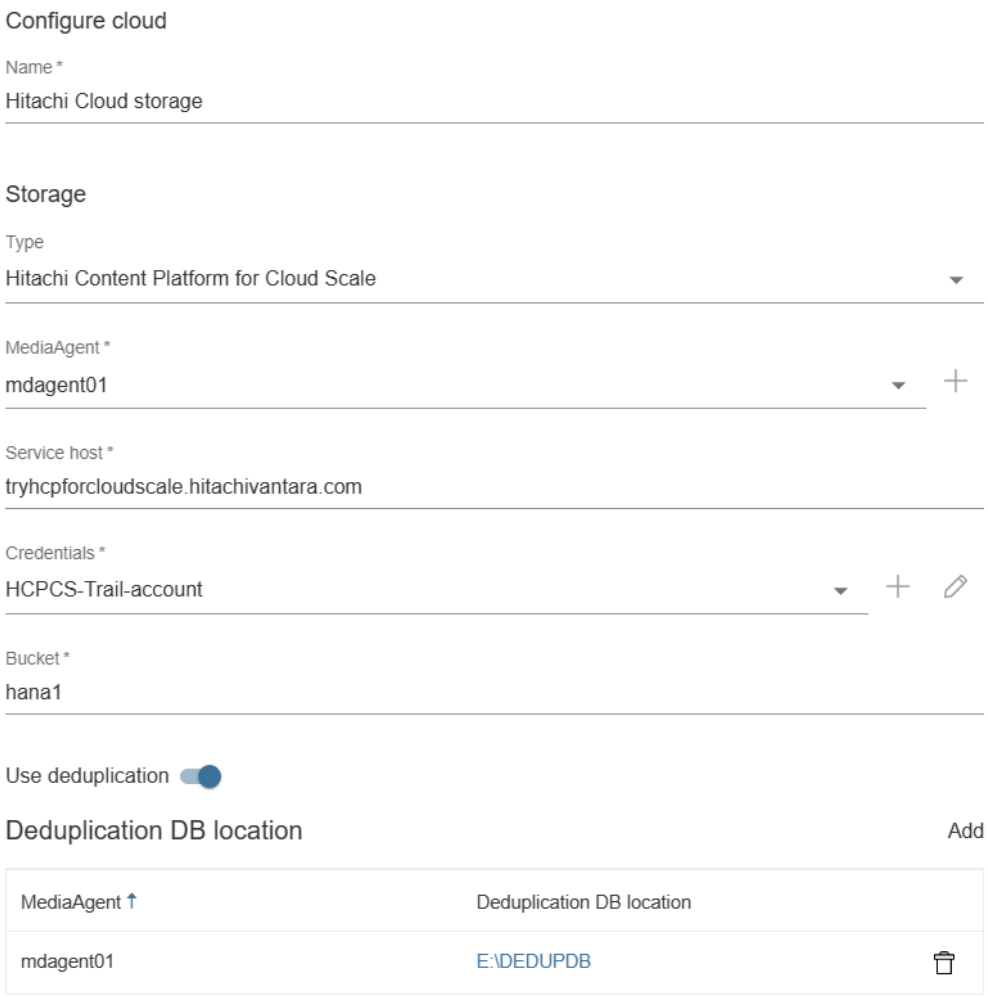

Server backup plan

A server backup plan is a combination of a storage policy, a schedule policy, and a subclient policy.

You can create a server plan to specify the following:

- The storage pool that stores the backup data
- The data retention period
- The backup frequency to use to automatically back up the data
- IntelliSnap options

Use the following procedure to configure a server backup plan.

Procedure

1. From the navigation pane, go to **Manage** > **Plans**.
2. Click **Create plan**, and then select **Server backup**.
3. Complete the fields on the **Create server backup plan** page.

Sub-page	Steps
General	1. Choose whether to create a new plan or use a base plan.
	2. In **Plan name**, enter a name for the server plan.
Backup destinations	1. Click **Add copy**.
	2. In **Name**, enter a name for the backup destination.
	3. From the **Storage list**, select the storage to use for the backups.
	4. To specify the start time for backups on a specific date and later, move the **Backups On and After toggle** key to the right.
	5. For **Retention rules**, enter the amount of time to retain the backups.
	6. To specify additional backups, such as weekly full backups, move the **Extended retention rules** toggle key to the right, and then add rules.
	7. Click **Save**.
	8. To add additional regions to the backup destination, move the **Multi-region** toggle key to the right, and then select the regions.
	You can also add more regions and associate the storage copy to the region.
Recovery point objective	1. For **Backup frequency**, click **Add**.

Sub-page	Steps
Backup schedule dates and times that you select in this section are applied to the time zones of the servers that are associated to the plan. In other words, backup jobs for different servers located in different time zones run at the same time but according to different time zones. However, you can change the time zone to suit your requirements. For example, you can change the time zone to match the CommServe time zone so that all associated servers (even in different time zones) are backed up at the same time according to the CommServe time zone. A default incremental backup is scheduled every day at 9 PM for all agents. You can add additional backup schedules. For example, you can add an additional backup schedule only for databases. Or, you can choose to delete the default incremental backup schedule. Server plans without associated backup schedules can be used for on-demand backups.	a. From the **Backup type** list, select the backup type, and from the **for** list, select **All agents or Databases**. b. Specify the frequency of the backup. c. To specify the client time zone or to add exceptions, move the **Advanced** toggle key to the right. From the time zone list, select the time zone according to which backups must run. To add exceptions when backups must not run, click **Exceptions** and add exceptions for any day or week of a month. These exceptions will repeat every month. 2. For **Backup window**, specify when you want incremental and differential backups to run. 3. For **Full backup window**, specify when you want full backups to run.
Backup content	1. To back up only some content, in **Content to back up**, enter the content to back up.
	By default, all content is backed up.
	2. To exclude folders or files from the backup, in **Exclude - files/folders/patterns**, enter the content to exclude.
	3. Specify whether to include the system state in backups:
	To include the system state in all backups, move the **Back up system state** toggle key to the right.
	To include the system state only in full backups, select the **Back up system state** check box and the **Only with full backup** check box.

Sub-page	Steps
Options & restrictions	1. Under **Snapshot options**, do the following:
	a. Specify how to retain snapshots:
	To specify a number of jobs to retain on a snapshot copy, select **Number of snap recovery points**, and then enter the number of jobs to retain.
	To specify a retention period, select **Retention period**, and then enter the amount of time to retain the jobs.
	b. If you don't want to create backup copies, move the **Enable backup copy** toggle key to the left to turn it off.
	c. For **Backup copy frequency**, enter how often to run backup copy jobs.
	2. Under **Database options**, do the following:
	a. For **Log backup RPO**, enter how often to run log backups.
	b. To automatically back up the database logs to the MediaAgent cache, do the following:
	i. Move the **Use disk cache for log backups** toggle key to the right.
	ii. For **Commit**, enter how often to commit the logs to the CommServe computer.
	Disk caching of database logs applies to the following: Informix, Microsoft SQL Server on Windows, Oracle, Oracle RAC, and SAP HANA.
	3. Under **Override restrictions**, to allow the derived plans that use this server plan as the base plan to override the settings, move the **Allow plan to be overridden** toggle key to the right, and then select one of the following options for Storage pool, RPO, and Backup content:
	Override required: Plans derived from this base plan do not inherit the base plan values.
	Override optional: Plans derived from this base plan can either inherit the base plan values or have different values.

Sub-page	Steps
	Override not allowed: Plans derived from this base plan must inherit the base plan values.

4. Click **Submit** to save the plan.

Configure the command control interface engine on the storage system software

Before adding a Hitachi Vantara storage system in Array Management, configure the system software.

For the SAP HANA agent, the following applies:

- The Virtual Server Agent client computer can be a physical server or a VM, with Hitachi Vantara devices exposed through physical Raw Device Mapping (RDM).
- For storage area network (SAN) mode, the Virtual Server Agent client computer must be a physical server.
- Verify that all Hitachi Vantara prerequisites are installed and configured on the Virtual Server Agent client computer.
- Expose Hitachi Vantara Command Devices to the Virtual Server Agent's proxy computer.

Before you begin

- Before changing any settings or mapping or unmapping a Command Device, shut down all Hitachi Open Remote Copy Manager (HORCM) instances that Commvault software created.
- Configure enough destination ShadowImage secondary volume (S-VOL) devices to meet your retention requirements. If you are using virtual storage machines (VSMs), the ShadowImage devices must belong to the same VSM as the production primary volume (P-VOL).

Procedure

1. For each host that is connected to the storage system, complete the following steps to configure a Command Device:
   a. On the Command Device interface, select the **User Authentication** attribute.
   b. On the Command Device interface, clear the **Command Device Security** and **Device Group Definition** attributes.
2. For Thin Image operations on VSP storage systems and on the Hitachi Unified Storage Virtual Machine (HUS VM) series, do the following:
   a. Create Thin Image pools. If you are using microcode version 80-05-44 and later, then any Hitachi Dynamic Provisioning (HDP) pool can be used to create TI snapshots.
   b. Unmap the virtual volumes (V-VOLs) that the TI snapshot backup will use.

c. Create V-VOLs that match the exact block size of the P-VOL devices. For V-VOLs to be created automatically by the IntelliSnap feature, select the snapshot configuration property Create VVOLs for Thin Image in Array Management.

   For more information, see *Create VVOLs for Thin Image* at https://documentation.commvault.com/2022e/expert/60141_snapshot_configuration_properties_for_hitachi_vantara_storage_array.html.

d. For the Hitachi VSM storage system, use volumes that belong to the VSM resource group.

3. For ShadowImage operations on the VSP series and on the HUS VM series, create S-VOL devices that match the exact block size of the P-VOL devices.

   Do not map existing S-VOL devices unless you also select the Use preexisting clones property on the Snap Configuration tab.

Configure the Hitachi Vantara storage system using Array Management

For IntelliSnap backups from Commvault, add the Hitachi Vantara storage system and its configuration information to the CommServe database.

Before you begin

- Verify that the client and proxy computers are configured to have access to the storage system.
- If you are configuring the replication feature, then you must add and configure both local and remote Hitachi Vantara systems.
- Determine the serial number of the storage system for the Name field. Hitachi Storage Navigator displays a 5-digit serial number, but the actual serial number has 6 digits. As a workaround, when you add a storage system in Array Management, you must add a prefix number.

For this model	Add this prefix
VSP G1000 (EOSL June 2022)	3
VSP G800	4
VSP 5500	5

Procedure

1. To verify the serial number of the storage system, type the following on the command line:`BaseFolder\etc> raidcom.exe get resource -In`

   Where `n` is a Hitachi Open Remote Copy Manager (HORCM) instance number.

2. On the ribbon in the CommCell Console, click the **Storage** tab, and then click **Array Management**.

3. Click **Add**.
4. On the **General** tab, specify the following information:
   a. From the **Snap Vendor** list, select Hitachi Vantara.
   b. In the **Name** field, enter the serial number of the storage system.
   c. In the **Control Host** field, enter the IP address or the host name of the Device Manager server.

      Tip: The **Control Host** field does not apply to CCI-based engines for Available Array or for Associated Array. For these engines, enter **NA**.

   d. In the **Credentials** area, click **Change**.
   e. Type the credentials of a user who has Modify permission and View permission for the Device Manager, and then click **OK**.

      Note: If you are using a VSM, then these permissions do not apply.

5. To configure additional MediaAgents to perform operations on the storage system, click the **Array Access Nodes** tab, and then select configuration options.

   See *Configuring Array Access Nodes* at https://documentation.commvault.com/2022e/expert/36943_configuring_array_access_nodes_for_storage_array.html.
6. On the **Snap Configuration** tab, configure the snapshot configuration properties for your environment or use the default values.

   See *Snap Configuration* athttps://documentation.commvault.com/2022e/expert/60141_snapshot_configuration_properties_for_hitachi_vantara_storage_array.html.
7. Click **OK**.

   The following figure shows Hitachi Vantara storage Array Management.

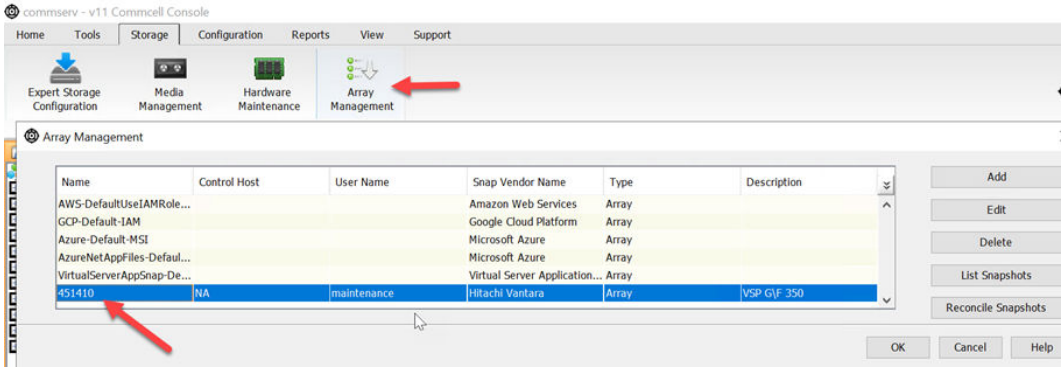

The following figure shows Array Properties.

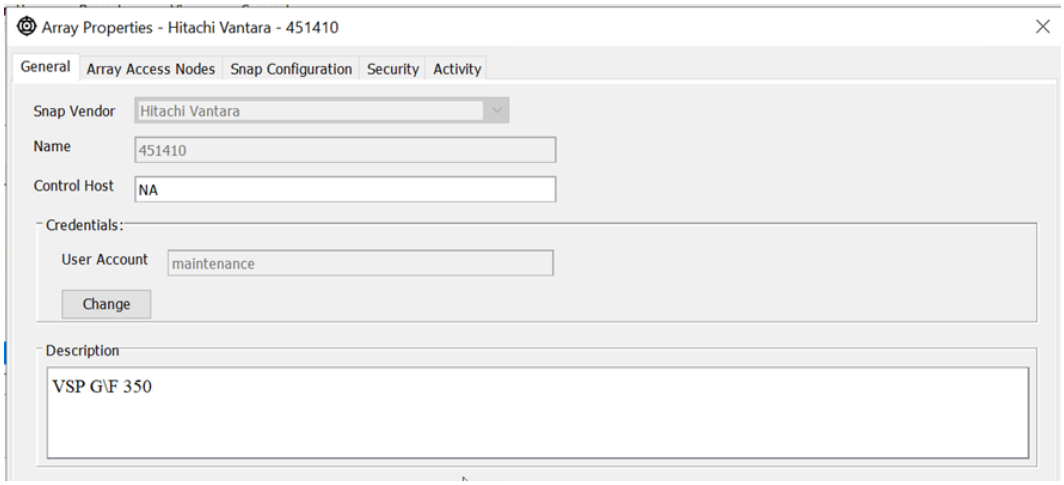

Enable IntelliSnap and select a snapshot engine for database solutions

To back up a subclient with IntelliSnap, you must enable IntelliSnap and select a snapshot engine.

Procedure

1. From the navigation pane, go to **Protect** > **Databases**.
2. Click an instance name.
3. In the **Subclient** section, click a subclient.
4. In the **Snapshot Engine** section, click **Edit**.
5. Select **Enable snapshot engine**.
6. Select a snapshot engine from the **Engine name** list, and then click **Save**.

   The following figure shows snapshot management settings.

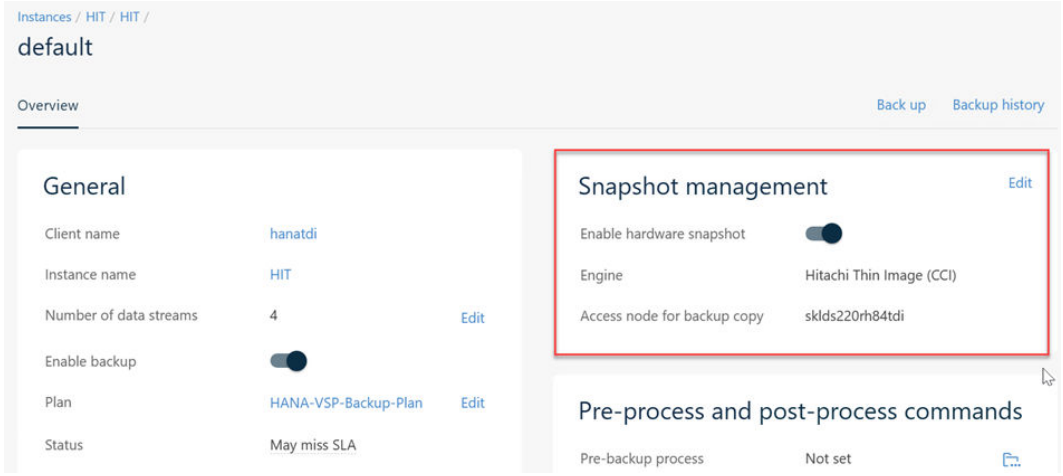

The following figure shows the IntelliSnap backup method.

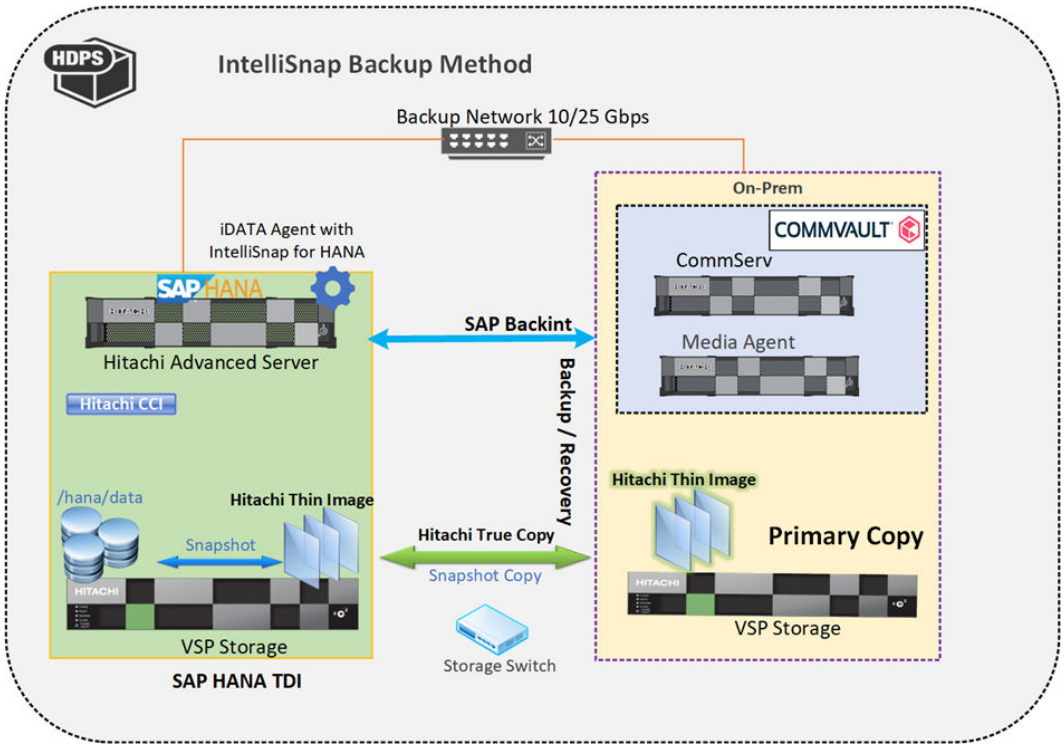

Engineering validation

After all components were configured, backup and recovery of SAP HANA database to/from HDPS was performed using Commvault CommCell Browser, Command Center, and SAP HANA Studio.

SAP HANA backup and restore

You can use the Commvault software to back up and restore SAP HANA. SAP HANA is also supported in a multi-tenant environment.

Backups

Data you can back up:

- All of the database files using SAP Backint and IntelliSnap on VSP storage systems
- All of the database files using SAP Backint on HCP (S3-based object storage)
- The HANA log files
- The catalog files

Data that cannot be backed up:

- Application profiles and binaries that are associated with the SAP HANA installation

Types of backups:

- Full backups
- Incremental backups
- Differential backups

When you can perform backups:

- On a schedule: The server plan that you assign manages scheduled backups.
- On demand: You can perform on-demand backups at any time.

The following figure shows the HANA Studio Backup console.

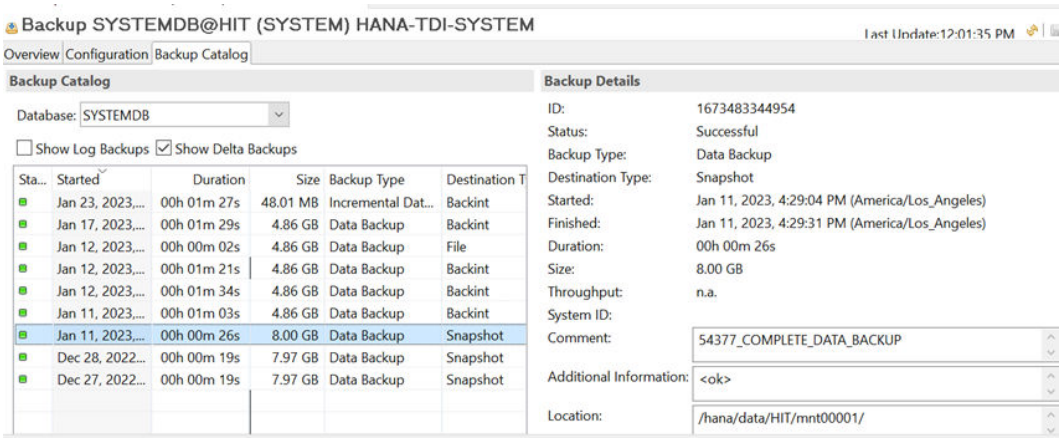

The following figure shows the Command Center backup history.

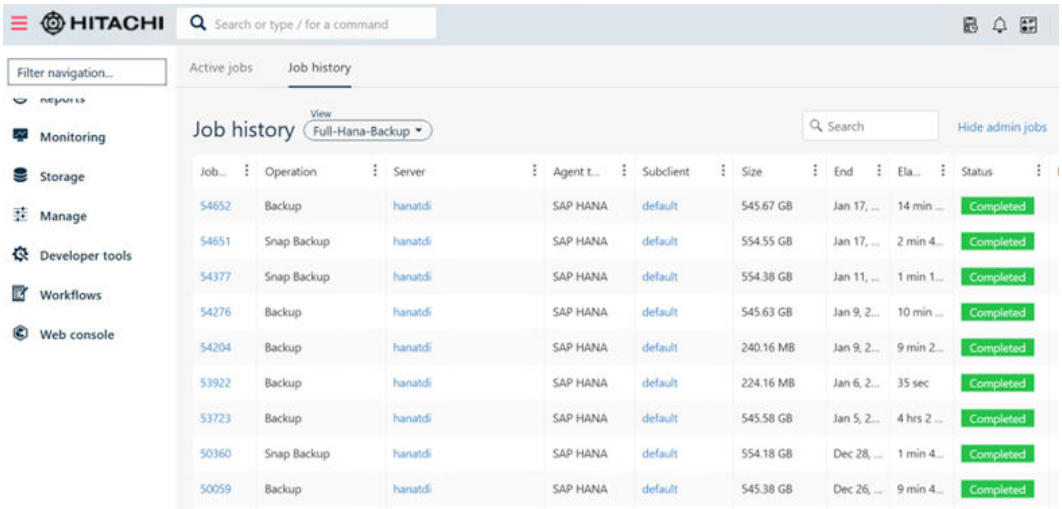

Backup copies

Data you can copy:

- All backups can be copied from Primary Storage to Secondary Storage using Auxiliary Copy.
- Tiering and recovery to and from HCP for cloud scale.
- Copy Thin Image data using Hitachi TrueCopy®

When you can perform backup copies:

- On a schedule: Create a copy policy associated with the server backup plan and run by default every 30 minutes.
- On demand: You can perform on-demand backup copies at any time.

The following figure shows an Auxiliary Copy list.

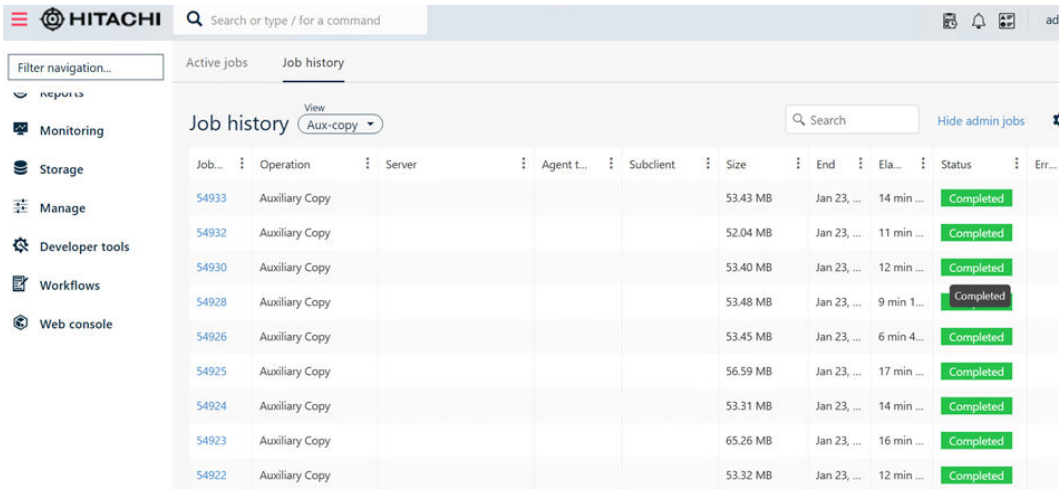

The following figure shows HCP for cloud scale data.

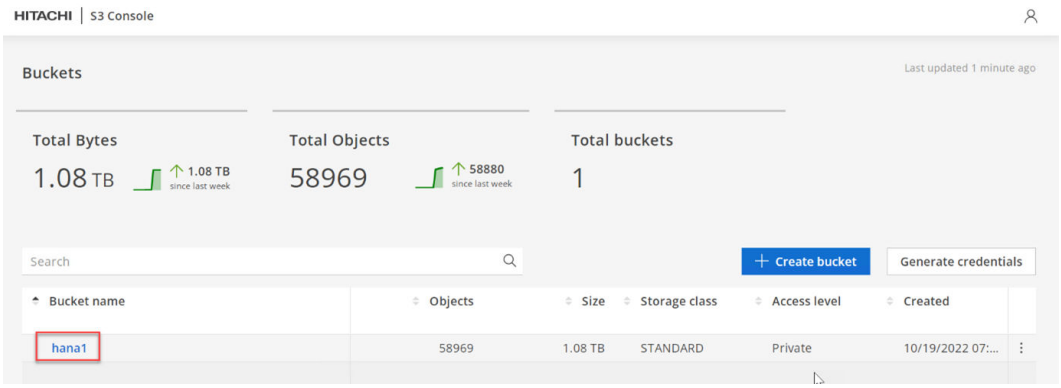

Restores

Data you can restore:

- Full database: This restore operation includes the database and the log files.
- Data only using IntelliSnap.

Backups you can use for restores:

- The most recent backup.
- A backup from a specific date (point-in-time).
- A backup identified by a backup prefix or an internal backup job ID.

Destinations you can restore to:

- The current location (in place).
- A different server (out of place).
- A clone.

The following figure shows a restore list.

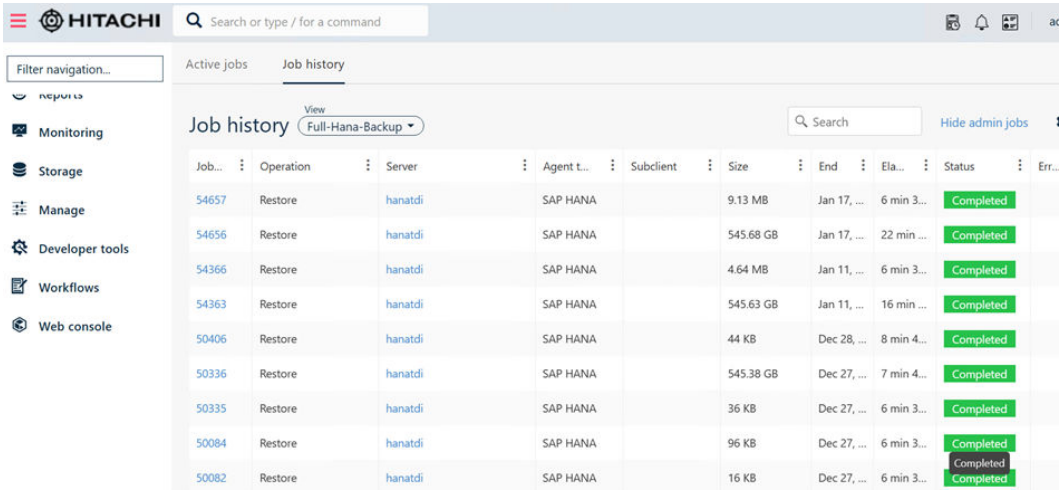

Test cases

The following test cases were performed during functional testing of HDPS for SAP HANA in the Hitachi test lab.

HANA Backup and Aux Copy test cases:

- Test Case 1: Take Full Backup of HANA Appliance server on Primary Storage-Disks (VSP) using SAP Backint
- Test Case 2: Take Full Backup of HANA TDI server on Primary Storage-Disks (VSP) using SAP Backint
- Test Case 3: Take Full Backup of HANA TDI server on Primary Storage-Disks (VSP) using Commvault IntelliSnap
- Test Case 4: Primary Storage-Disks (VSP) to Secondary Storage- HCP (Object Storage) replication
- Test Case 5: Take Full Backup of HANA-HC server on Primary Storage-HCP (Object Storage) using SAP Backint
- Test Case 6: Take Full Backup of HANA TDI server on Primary Storage-HCP (Object Storage) using SAP Backint

HANA restore test cases:

- Test Case 1: Restore of HANA-HC server from Full Backup through the Backint method
- Test Case 2: Restore of HANA TDI server from Full Backup through the Backint method
- Test Case 3: Restore of HANA TDI server through the Commvault IntelliSnap method

The following figure shows test case results.

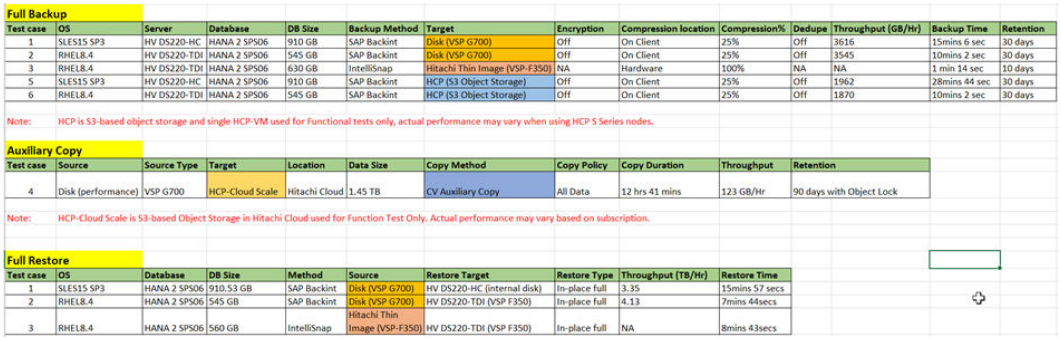

Full Backup

Test case	OS	Server	Database	DB Size	Backup Method	Target	Encryption	Compression location	Compression%	Dedupe	Throughput (GB/Hr)	Backup Time	Retention
1	SLES15 SP3	HV DS220-HC	HANA 2 SPS06	910 GB	SAP Backint	Disk (VSP G700)	Off	On Client	25%	Off	3616	15mins 6 sec	30 days
2	RHEL8.4	HV DS220-TDI	HANA 2 SPS06	545 GB	SAP Backint	Disk (VSP G700)	Off	On Client	25%	Off	3545	10mins 2 sec	30 days
3	RHEL8.4	HV DS220-TDI	HANA 2 SPS06	630 GB	IntelliSnap	Hitachi Thin Image (VSP-F350)	NA	Hardware	100%	NA	NA	1 min 14 sec	10 days
5	SLES15 SP3	HV DS220-HC	HANA 2 SPS06	910 GB	SAP Backint	HCP (S3 Object Storage)	Off	On Client	25%	Off	1962	28mins 44 sec	30 days
6	RHEL8.4	HV DS220-TDI	HANA 2 SPS06	545 GB	SAP Backint	HCP (S3 Object Storage)	Off	On Client	25%	Off	1870	10mins 2 sec	30 days

Note: HCP is S3-based object storage and single HCP-VM used for Functional tests only, actual performance may vary when using HCP S Series nodes.

Auxiliary Copy

Test case	Source	Source Type	Target	Location	Data Size	Copy Method	Copy Policy	Copy Duration	Throughput	Retention
4	Disk (performance)	VSP G700	HCP-Cloud Scale	Hitachi Cloud	1.45 TB	CV Auxiliary Copy	All Data	12 hrs 41 mins	123 GB/Hr	90 days with Object Lock

Note: HCP-Cloud Scale is S3-based Object Storage in Hitachi Cloud used for Function Test Only. Actual performance may vary based on subscription.

Full Restore

Test case	OS	Database	DB Size	Method	Source	Restore Target	Restore Type	Throughput (TB/Hr)	Restore Time
1	SLES15 SP3	HANA 2 SPS06	910.53 GB	SAP Backint	Disk (VSP G700)	HV DS220-HC (internal disk)	In-place full	3.35	15mins 57 secs
2	RHEL8.4	HANA 2 SPS06	545 GB	SAP Backint	Disk (VSP G700)	HV DS220-TDI (VSP F350)	In-place full	4.13	7mins 44secs
3	RHEL8.4	HANA 2 SPS06	560 GB	IntelliSnap	Hitachi Thin Image (VSP-F350)	HV DS220-TDI (VSP F350)	In-place full	NA	8mins 43secs

Sample global.ini file

This is a sample *global.ini* file that was configured and used for validating the HANA backup environment as an appliance solution or TDI solution with SUSE Linux Enterprise Server for SAP Applications 15 SP3 and later or Red Hat Linux Enterprise Server for SAP 8.4 and later versions:

```
[backup]
parallel_data_backup_backint_channels = 4
parallel_data_backup_backint_size_threshold = 32
data_backup_parameter_file = /usr/sap/HIT/SYS/global/hdb/opt/hdbconfig/param
log_backup_parameter_file = /usr/sap/HIT/SYS/global/hdb/opt/hdbconfig/param
log_backup_using_backint = true
catalog_backup_parameter_file = /usr/sap/HIT/SYS/global/hdb/opt/hdbconfig/param
catalog_backup_using_backint = true
backint_response_timeout = 1800

[communication]
tcp_backlog = 2048
sslminprotocolversion = tls12

[fileio]
async_read_submit[log] = on
async_write_submit_active[log] = auto
async_write_submit_blocks[log] = all
min_submit_batch_size[log] = 16
max_submit_batch_size[log] = 64
max_parallel_io_requests[log] = 64
size_kernel_io_queue[log] = 512
async_read_submit[data] = on
async_write_submit_active[data] = auto
async_write_submit_blocks[data] = all
min_submit_batch_size[data] = 16
max_submit_batch_size[data] = 64
max_parallel_io_requests[data] = 64
size_kernel_io_queue[data] = 512

[ldap]
sslminprotocolversion = tls12

[multidb]
mode = multidb
database_isolation = low
singletenant = yes

[persistence]
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
```

Product descriptions

The following products are part of this solution.

Hitachi Advanced Server DS220

With a combination of two Intel Xeon Scalable processors and high storage capacity in a 2U rack-space package, Hitachi Advanced Server DS220 delivers the storage and I/O to meet the needs of converged solutions and high-performance applications in the data center.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 24 slots for high-speed DDR4 memory, allowing up to 3 TB of memory per node when 128 GB DIMMs are used. This server supports up to 12 large form factor storage devices and an additional 2 small form factor storage devices.

This server has three storage configuration options:

- 12 large form factor storage devices and an additional 2 small form factor storage devices in the back of the chassis
- 16 SAS or SATA drives, 8 NVMe drives, and an additional 2 small form factor storage devices in the back of the chassis
- 24 SFF devices and an additional 2 SFF storage devices in the back of the chassis

Hitachi Content Platform for cloud scale

Hitachi Content Platform for cloud scale (HCP for cloud scale) is a software-defined object storage solution that is based on a massively parallel microservice architecture, and is compatible with the Amazon S3 application programming interface (API). HCP for cloud scale is well suited to service applications requiring high bandwidth and compatibility with Amazon S3 APIs.

Hitachi Content Platform

Hitachi Content Platform (HCP) is a secure, simple and intelligent web-scale object storage platform that delivers superior scale, performance, security, efficiency and interoperability. It allows any organization to deliver unique, feature-rich, private, hybrid, multicloud, or public cloud storage services at a cost comparable to public cloud. The rich feature set and extensive ecosystem surrounding the platform allow organizations to improve efficiencies and optimize costs. They can choose to move data to on-premises storage tiers, off-site to a choice of public cloud providers or to a combination of both.

Hitachi Data Protection Suite

With Hitachi Data Protection Suite (HDPS), powered by Commvault, Hitachi Vantara and Commvault together deliver a unified, modern offering that facilitates the backup, recovery and management of enterprise and application such as SAP HANA with the industry-leading object storage solution, Hitachi Content Platform (HCP). HDPS and HCP offer the reliability required by the world's largest organizations, while featuring the simplicity, cost-effectiveness and modern capabilities that are needed to remain agile and competitive. HCP seamlessly extends the secured and guaranteed management of long-term data retention at petabyte scale.

Hitachi Virtual Storage Platform F Series family

Use Hitachi Virtual Storage Platform F series family storage for a flash-powered cloud platform for your mission critical applications. This storage meets demanding performance and uptime business needs. Extremely scalable, its 4.8 million random read IOPS allows you to consolidate more applications for more cost savings.

Hitachi Virtual Storage Platform F series family delivers superior all-flash performance for business-critical applications, with continuous data availability.

SAP HANA

SAP HANA converges database and application platform capabilities in-memory to transform transactions, analytics, text analysis, predictive and spatial processing so businesses can operate in real-time. This combines database, data processing, and application platform capabilities in a single in-memory platform. Also, the platform provides libraries for predictive, planning, text processing, spatial, and business analytics — all on the same architecture. This architecture comes from leading hardware partners of SAP, including Hitachi Vantara. For more information, see https://www.sap.com/products/hana.html.

By eliminating the divide between transactions and analytics, SAP HANA allows you to answer any business question anywhere in real time.

As a SAP customer, you can download more information, including the following:

- SAP HANA Master Guide

  This is the central starting point for the technical implementation of SAP HANA. Use this guide for basic concepts and for planning.

- SAP HANA Server Installation and Update Guide

  This guide provides an overview of how to install and update a SAP HANA system with the SAP HANA lifecycle management tools.

- SAP HANA Administration Guide

  This guide explains how to configure, manage, maintain, and optimize your SAP HANA installation using SAP HANA administration tools.

SAP HANA hardware directory provides information about SAP HANA appliances certified by SAP hardware partners.

Operating system options for SAP HANA

SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux for SAP HANA are available operating systems when running SAP HANA.

- SUSE Linux Enterprise Server (SLES) for SAP Applications

  Compete more effectively through improved uptime, better efficiency, and accelerated innovation using SUSE Linux Enterprise Server for SAP Applications. This is a versatile server operating system for efficiently deploying highly available enterprise-class IT services in mixed IT environments with performance and reduced risk.

  SUSE Linux Enterprise Server was the first Linux operating system to be certified for use with SAP HANA. It remains the operating system of choice for most SAP HANA customers.

- Red Hat Enterprise Linux (RHEL) for SAP HANA

  Using the stability and flexibility of Red Hat Enterprise Linux for SAP HANA, reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability and military-grade security. Use Enterprise Linux to tailor your infrastructure as markets shift and technologies evolve.

Changing the configuration settings is only supported along the guidelines of SAP and the operating system distributor and may otherwise cause significant performance problems. The following SAP Notes for SUSE Linux Enterprise Server and Red Hat Enterprise Linux are a good starting point for information on this topic:

- 1944799 - SAP HANA Guidelines for SLES Operating System Installation
- 2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL) Operating System

For more details, see "Updating and Patching the Operating System" by searching in the "View SAP HANA document" from Technical Information and Best Practices.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact