



Hitachi Storage Adapter for the SAP HANA Cockpit in Hitachi Unified Compute Platform for the SAP HANA Platform in a Scale-Up Configuration

Best Practice Guide

By Prasad Patkar, Morihide Nakaya, , Stephen Ma, Archana
Kuppuswamy

January 2017

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.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.

Contents

Solution Components	3
Solution Elements	4
Hardware Components.....	4
Software Components	5
Best Practices	8
Replicate Using Hitachi Thin Image	8
Replicate Using Hitachi ShadowImage.....	11
Compare Hitachi Thin Image to Hitachi ShadowImage.....	14
Recover and Restore the Operating System LUN	14
Recover and Restore the SAP HANA Storage Snapshot	16
Copy and Clone the SAP HANA Database with Hitachi ShadowImage	17
Limitations of Hitachi Storage Adapter for SAP HANA Cockpit	18
Steps to Configure fstab and Network.....	20
Steps to Register a SAP HANA System.....	21

Hitachi Storage Adapter for the SAP HANA Cockpit in Hitachi Unified Compute Platform for the SAP HANA Platform in a Scale-Up Configuration

Best Practices Guide

This describes the best practices for using Hitachi Storage Adapter for the SAP HANA cockpit in Hitachi Unified Compute Platform for the SAP HANA platform (UCP for SAP HANA) in a scale-up configuration for IT management. This solution uses the following components:

- Hitachi Compute Systems
- Hitachi storage
- A management server from Quanta Cloud Technology

The SAP HANA database keeps the bulk of its data in memory. It uses persistent storage to provide a fall-back in case of failure. However, if the persistent storage is damaged, you need snapshots of the database for recovery. The ability to restore real time data processing, such as the SAP HANA database, translates directly into revenue for your business.

When used for infrastructure management, Storage Adapter for SAP HANA cockpit uses Hitachi Thin Image (HTI) and Hitachi ShadowImage. Use this storage management software to simplify IT administrative tasks.

Hitachi Thin Image (HTI) creates a storage snapshot of the data volumes. You can restore a SAP HANA database using this replica snapshot.

Hitachi ShadowImage creates a replica of logical volume in the same storage system without host. ShadowImage enables concurrent operations such as backup and batch processing with minimum impacts on online operation continuity.

You need to use Thin Image and Hitachi ShadowImage with Storage Adapter for SAP HANA cockpit to implement these practices fully.

This solution is for use by IT administrators, database administrators, storage administrators, SAP HANA administrators, and architects implementing snapshot, recovery, and cloning solutions.

This technical paper assumes that you have familiarity with the following:

- Storage area network (SAN)-based storage systems
- General storage concepts
- General network knowledge
- Common IT storage practices
- SAP HANA
- Hitachi replication products

Note — These best practices were tested and developed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Solution Components

The IT management solution in this paper used the following for testing:

- **Servers**
 - Hitachi Compute Blade 2500 (CB 2500)
 - The rack optimized server for solutions, 2U four node, for the management server
- **Storage**
 - Hitachi Virtual Storage Platform G200 (VSP G200)

However, this solution supports using a wide range of Hitachi compute and storage systems with this storage management software.

Find the supported hardware configurations of Hitachi Storage Adapter for SAP HANA cockpit in *Hitachi Storage Adapter for the SAP HANA Cockpit User Guide* (MK-92ADPTR111-00 or later). Contact Hitachi Data Systems for a copy of this document.

Figure 1 shows the topology of this solution.

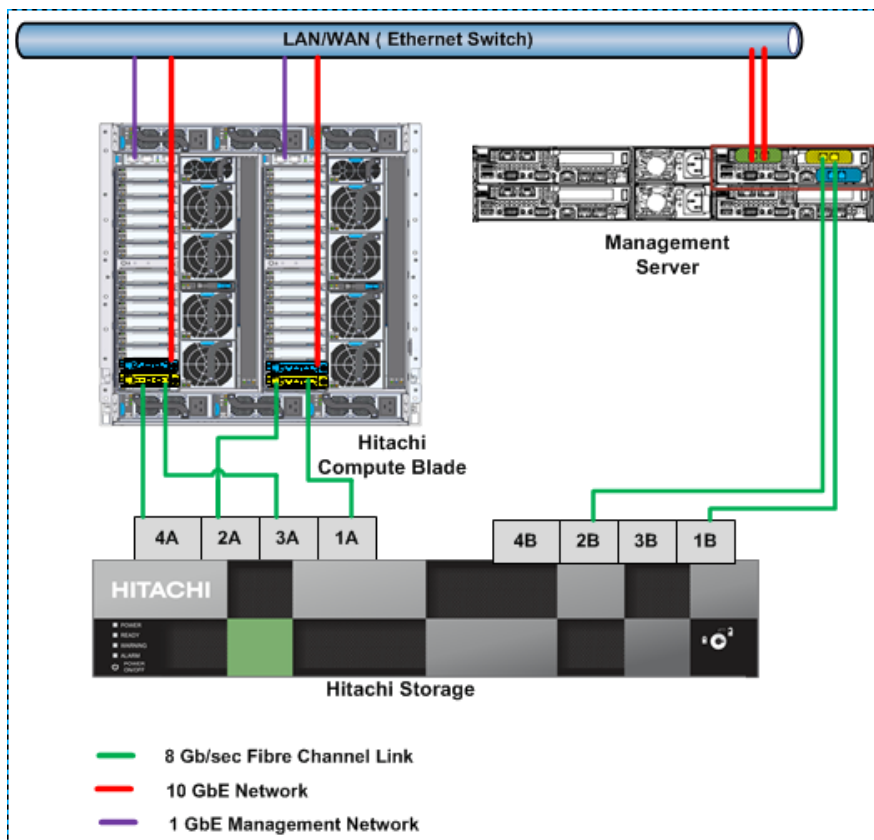


Figure 1

Solution Elements

The following are the key hardware and software components used in this infrastructure management solution.

Hardware Components

Find the supported server chassis, server blades, and storage system to use Hitachi Storage Adapter for the SAP HANA cockpit in *Hitachi Storage Adapter for the SAP HANA Cockpit User Guide* (MK-92ADPTR111-00 or later).

Table 1 lists the additional hardware you need to implement Hitachi Storage Adapter for SAP HANA cockpit to simplify IT management.

Table 1. Hardware Elements

Hardware	Quantity	Configuration (per unit)	Role
Rack optimized server for solutions, 2U four node	1	<ul style="list-style-type: none"> ▪ 2.5 inch bay chassis ▪ One server node with following components: <ul style="list-style-type: none"> ▪ 2 Intel Xeon E5-2620 v3 processors (6 cores, 2.4 GHz, 85 W) ▪ 1 heat sink for CPU0 and CPU1 ▪ 2 × 16 GB DDR4 2,133 MHz memory module ▪ 2 × 500 GB, 7200 RPM, 2.5 inch SATA HDD (6 Gb) ▪ 1 dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card ▪ 1 dual port 1 GbE Base-T Intel i350 mezzanine card ▪ 1 Emulex dual port 8 Gb/sec Fibre Channel HBA ▪ 3 server fillers 	Optional management server to run the following: <ul style="list-style-type: none"> ▪ Network time protocol ▪ Hi-Track Remote Monitoring system ▪ SAP HANA Studio
Brocade VDX 6740-48 port switch (optional)	2	<ul style="list-style-type: none"> ▪ 10 GbE ▪ 48 ports 	10 GbE client network and replication network
Brocade ICX 6430-24 port switch (optional)	1	<ul style="list-style-type: none"> ▪ 1 GbE ▪ 24 ports 	1 GbE management network

Software Components

Find the supported SAP HANA operating systems, storage microcode versions, and replication software versions for Hitachi Storage Adapter for the SAP HANA cockpit in *Hitachi Storage Adapter for the SAP HANA Cockpit User Guide* (MK-92ADPTR111-00).

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. In addition, 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.

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

SAP customers can download more information on the SAP HANA platform at the [SAP Service Marketplace](#). See the installation and upgrade guides download section for SAP In-Memory Computing (SAP In-Memory Appliance — SAP HANA). The following are available:

- [SAP HANA Master Guide](#)

This is the central starting point for the technical implementation of the SAP HANA platform. Use this for basic concepts and for planning the SAP HANA application system landscape.

- **SAP HANA Installation and Initial Configuration Guides**

Use the various installation guides to install the required SAP In-Memory Database and the other software components for the different replication technologies. Refer to the [SAP HANA Server Installation Guide](#) for an overview on how to install SAP HANA.

- [SAP HANA Technical Operations Manual](#)

This provides an end-to-end picture of the available SAP HANA appliance administration tools and the key tasks for a system administrator to perform.

- [SAP HANA Master Update Guide](#)

This explains how to update SAP HANA and its components.

- [SAP Integration and Certification Center \(SAP ICC\)](#)

This page provides information about SAP HANA appliances certified by SAP hardware partners.

SUSE Linux Enterprise Server (SLES) for SAP Applications and Red Hat Enterprise Linux (RHEL)

The following options are available as operating system:

- **SUSE Linux Enterprise Server (SLES) for SAP Applications**

Compete more effectively through improved uptime, better efficiency, and accelerated innovation using [SUSE Linux Enterprise Server](#). 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 the vast majority of SAP HANA customers.

- **Red Hat Enterprise Linux (RHEL)**

Using the stability and flexibility of [Red Hat Enterprise Linux](#), reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability on 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 SLES and RHEL 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](#)

Hitachi Storage Adapter for SAP HANA Cockpit

Hitachi Storage Adapter for SAP HANA cockpit runs as a web application. When needed, the adapter calls the web service on the management server. Then the web service communicates to the attached Hitachi storage and returns the information to the adapter.

In addition, the web service collects device-mapping information from the SAP HANA host by using SSH. The web service executes SQL queries for retrieving information from the SAP HANA database.

The web service manages all configuration information, storing it in the local disk of the management server. Then, the adapter passes configuration information to the web service.

Hitachi ShadowImage

High-speed, nondisruptive local mirroring technology of Hitachi ShadowImage, a part of [Hitachi In-System Replication Bundle](#), rapidly creates multiple copies of mission-critical information within all Hitachi storage systems.

ShadowImage keeps data RAID-protected and fully recoverable, without affecting service or performance levels. You can split replicated data volumes from the host applications and used for system backups, application testing, and data mining applications, while business continues to run at full capacity.

Hitachi Thin Image

The high-speed, non-disruptive snapshot technology of [Hitachi Thin Image](#) snapshot software rapidly creates up to one million point-in-time copies of mission-critical information within any Hitachi storage system or virtualized storage pool without impacting host service or performance levels.

Because snapshots store only the changed data, the volume of storage capacity required for each snapshot copy is substantially smaller than the source volume. As a result, Thin Image snapshot software can provide significant savings over full volume cloning methods.

Thin Image snapshot copies are fully read/write compatible with other hosts. They can be used for system backups, application testing, and data mining applications while the business continues to run at full capacity.

Hitachi Storage Adapter for SAP HANA cockpit uses Thin Image to create snapshots.

Best Practices

These are the best practices for Hitachi Storage Adapter for the SAP HANA cockpit in Hitachi Unified Compute Platform for SAP HANA in a scale-up configuration to simplify IT management with storage management software.

- “Replicate Using Hitachi Thin Image” on page 8
- “Replicate Using Hitachi ShadowImage” on page 11
- “Compare Hitachi Thin Image to Hitachi ShadowImage” on page 14
- “Recover and Restore the Operating System LUN ” on page 14
- “Recover and Restore the SAP HANA Storage Snapshot” on page 16
- “Copy and Clone the SAP HANA Database with Hitachi ShadowImage” on page 17
- “Limitations of Hitachi Storage Adapter for SAP HANA Cockpit” on page 18

Replicate Using Hitachi Thin Image

Hitachi Thin Image is used for backup to create a replica of the specified logical volume using a snapshot. Using Thin Image, you can create and restore a replica of an arbitrary point of time.

The capacity used to store a snapshot can be smaller than the storage needed to copy and store the entire volume data. This is because the snapshot data in the pool volume only stores data that is different from what is in the specified logical volume.

Hitachi Storage Adapter for SAP HANA cockpit uses Thin Image to create snapshots of the SAP HANA database.

Basic Hitachi Thin Image Components

These are the basic Hitachi Thin Image components.

- **V-VOL**

A virtual volume (V-VOL) is a volume used by the host to access a snapshot. Being defined virtually over all storage, a V-VOL does not have physical disk space.

- **Physical Disk Space**

There is physical disk space assigned in the snapshot area.

- **Thin Image Pool**

A Thin Image pool (TI pool) consists of a group of LUNs used by Hitachi Thin Image.

Steps for Replication Using Hitachi Thin Image

These are the steps to use Hitachi Storage Adapter with SAP HANA cockpit with Hitachi Thin Image for replication.

1. Estimate the Hitachi Thin Image Pool Size.

You must define the Hitachi Thin Image pool volume (TI pool) before creating Hitachi Thin Image pairs.

The maximum capacity for the snapshot created by database update is the same as the capacity of the log volume, because the log data equals the backup data (like a snapshot) of differential data. It means that the capacity for the snapshot is ensured if TI pool size is the same as the log volume size.

Total TI pool size for the data volume and shared volume are the same as log volume size, because their differential data are reflected in the log volume. The size of the TI pool and its data will not be larger than the capacity of the log volume.

You define the TI pool size considering the following points (Figure 2):

- The log volume size for all memory size models is 600 GB.
- Total TI pool size for the data volume and shared volume is the same as the log volume size.
- Decide what you want to be the threshold TI pool size to issue a warning notice about approaching the TI pool capacity.
- Allow capacity for the V-VOL management area. There is 3% of the size of the TI pool is used for management.

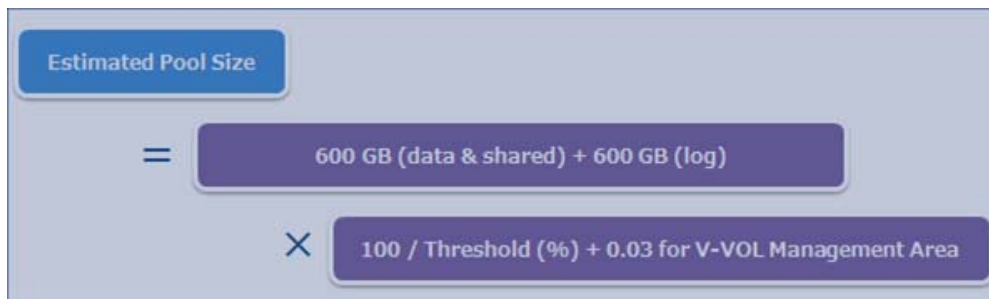


Figure 2

2. Select Host Group for mapping LUNS (Hitachi Thin Image pair).

A Hitachi Thin Image pair consists of a primary volume (P-VOL) and a pool.

The Hitachi Storage Adapter for SAP HANA cockpit supports selecting host group mapping LUNs for Hitachi Thin Image. Select all LUNs, except the LUN for the operating system, for the replication pair.

The replication runs by an assembly of pairs called the consistency group (CTG). The automatic creation of the CTG happens when you run pair creation from the user interface of Storage Adapter for SAP HANA cockpit. See Figure 3 on page 10.

The required V-VOL to assign an LDEV ID to snapshot data is created automatically by Storage Adapter for SAP HANA cockpit, as well as the CTG.

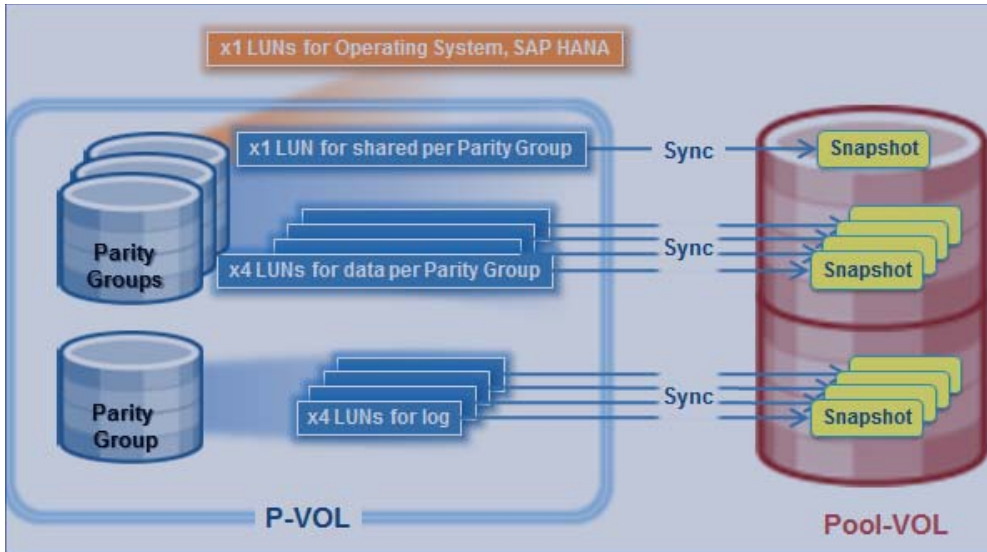


Figure 3

3. Complete savepoint execution and service shutdown before replication using Hitachi Thin Image.

P-VOL differential data is stored as snapshot data in TI pool when creating the Hitachi Thin Image pair. This replication solution requires savepoint execution to put the latest data from the memory in the host to the P-VOL before resynchronization with the TI pair.

By this step, Hitachi Thin Image can keep the snapshot up-to-date when the resynchronization procedure runs. In the event of data corruption, recover the P-VOL using this snapshot data. See Figure 4.

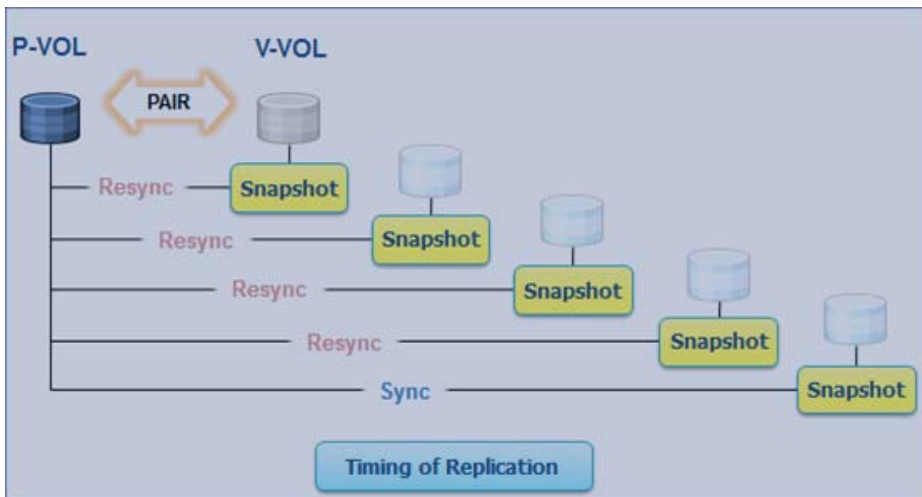


Figure 4

4. Release to open TI pool capacity.

Delete the Hitachi Thin Image pairs that are using the V-VOLs to open the capacity in TI pool. Run this process from the user interface provided by Hitachi Storage Adapter for SAP HANA cockpit.

If you delete the CTG for the Thin Image pairs, then this opens the space reserved to them in the TI pool capacity.

Replicate Using Hitachi ShadowImage

Hitachi ShadowImage creates a replica of logical volume (a pair) in the same storage system without a host. This software enables concurrent operations, such as backup and batch processing, with minimum impacts on online operation continuity. The ShadowImage pair performs all copy operations asynchronously after receiving the pair operation command.

1. Prepare either the (1) host group and parity group or (2) dynamic provisioning pool for the S-VOL.

Create the parity group or dynamic provisioning pools for the secondary volume (S-VOL) on the secondary system. Follow the same storage design for each S-VOL as that of each P-VOL.

Also, create a host group for mapping the S-VOL.

2. Select the host group to map the LUNs.

Hitachi Storage Adapter for SAP HANA cockpit supports selecting host group to map the LUNs for ShadowImage.

Select all LUNs, except the LUN for the operating system, for the replication pair. The replication runs by an assembly of pairs called the consistency group (CTG). The automatic creation of the CTG happens when you run pair creation from the user interface of Storage Adapter for SAP HANA cockpit. See Figure 5.

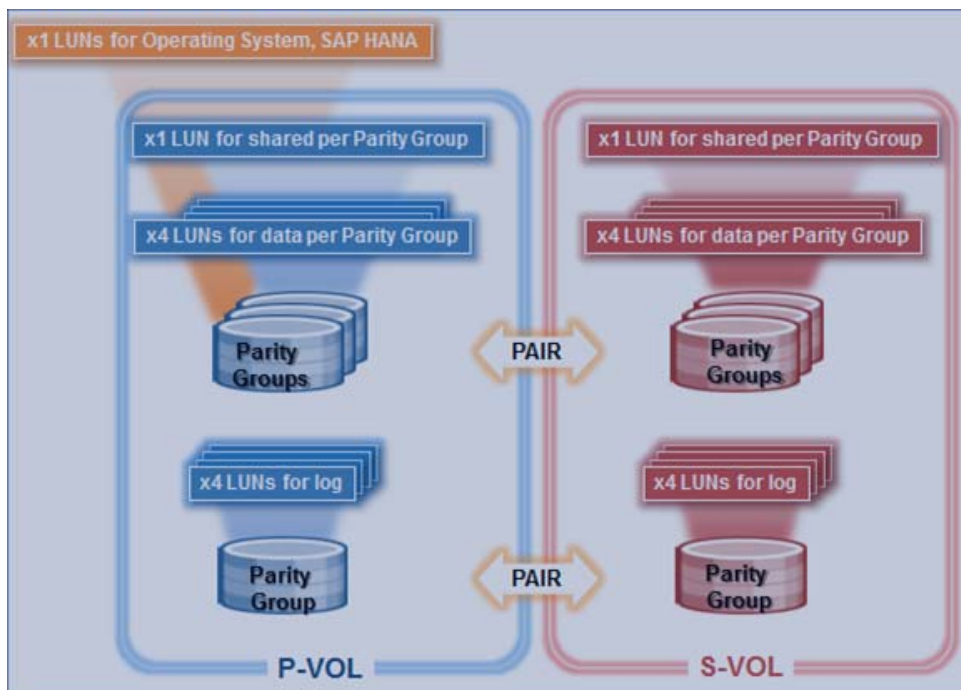


Figure 5

3. Create the pair.

Assign the logical volumes as P-VOL and S-VOL, respectively. When creating a new ShadowImage pair, all data in the P-VOL is copied to the S-VOL. See Figure 6.

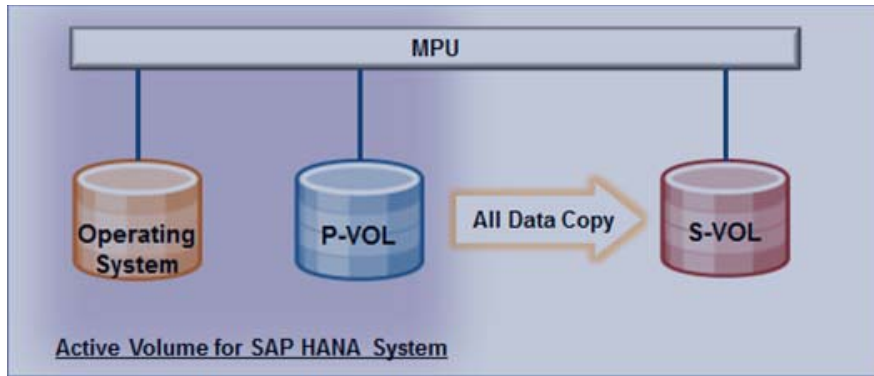


Figure 6

4. Split the pair.

During a pair split, the pair relation is kept.

If you issue a pair split command, the difference between the pair volumes at that time is copied from the P-VOL to the S-VOL so that the data between the pair volumes is fully synchronized.

Data at the time when receiving a pair split command from a host for multiple pairs belonging to the same CTG is guaranteed. See Figure 7.

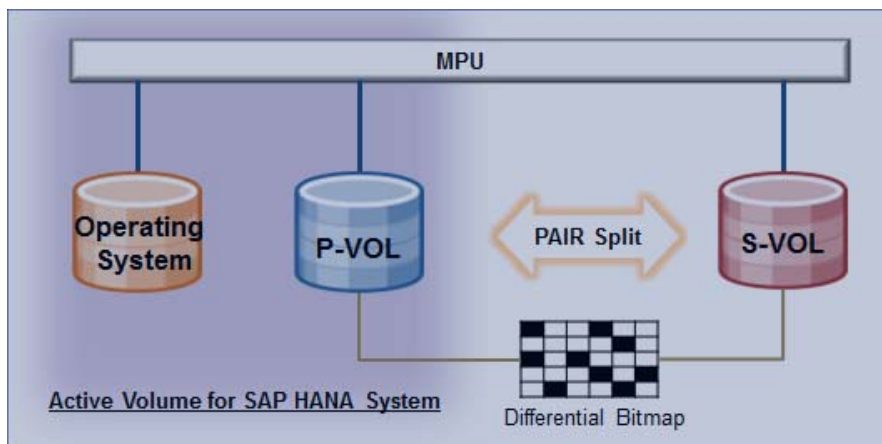


Figure 7

5. Complete savepoint execution and service shutdown before pair resync.

The recommendation is to shut down the SAP HANA service for the replication. The data in the P-VOL containing the latest available data in memory can be copied to the S-VOL by using pair resync. The consistency between P-VOL and S-VOL can be assured, because the data in the P-VOL is not updated.

When the assignment of volume for the SAP HANA system changes from P-VOL to S-VOL (or conversely), you need to do a mount and umount process.

▪ pair resync

This command copies the difference between the pair volumes after the pair split until the receipt of resync command from the P-VOL to the S-VOL.

The differential bitmap reflecting the differential data between the P-VOL and S-VOL is fixed after receiving the pair resync command. The differential data is copied from the P-VOL to the S-VOL by reference to the fixed differential bitmap. See Figure 8.

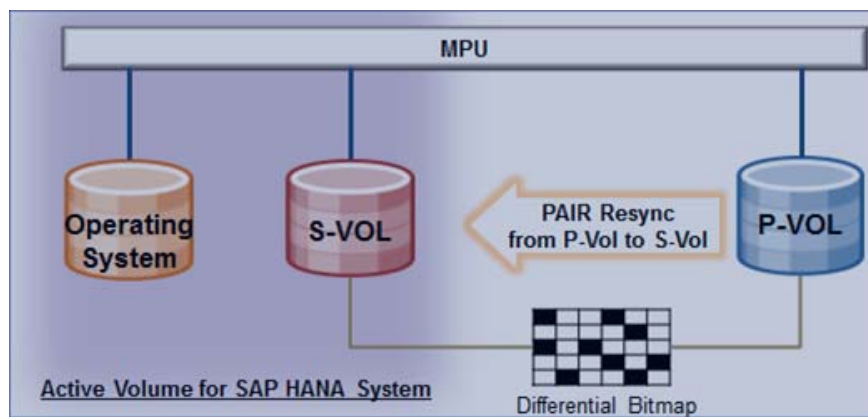


Figure 8

▪ pair revert

This copies the difference between pair volumes after pair split from S-VOL to P-VOL.

The differential data is copied from the S-VOL to the P-VOL by reference to the fixed differential bitmap. See Figure 9 on page 14.

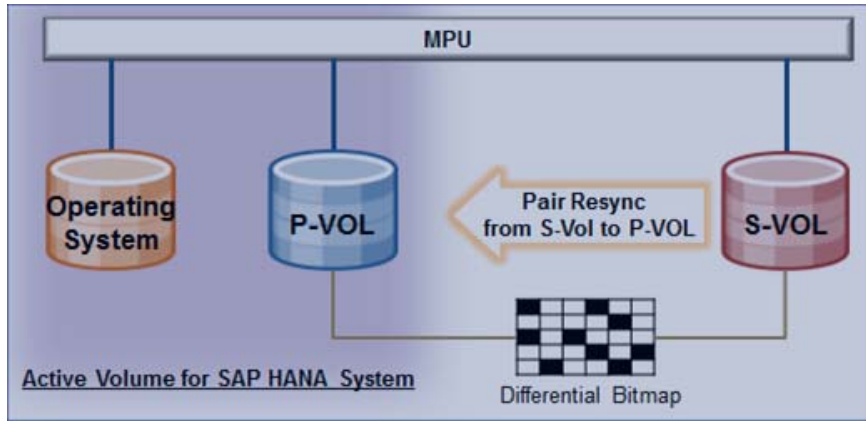


Figure 9

Compare Hitachi Thin Image to Hitachi ShadowImage

Table 2 shows the comparison of Hitachi Thin Image and Hitachi ShadowImage.

Table 2. Hitachi Thin Image vs Hitachi ShadowImage

Feature	Hitachi Thin Image	Hitachi ShadowImage
Ease of creation	Initial copy is not required. The replica can be created in a short time	Initial copy is required. The replica takes more time
Capacity efficiency of replica	More efficient because only data different from the P-VOL is copied.	Less efficient because the entire P-VOL is copied.
Use of replica	Impossible to use replica by separating it from the P-VOL	Possible to use replica by separating it from the P-VOL
P-VOL failure	When a physical failure occurs in the P-VOL, the replica cannot be guaranteed	When a physical failure occurs in P-VOL, the P-VOL can be recovered using the replica

Recover and Restore the Operating System LUN

Use Hitachi ShadowImage to create a local replica (S-VOL) of the operating system LUN (P-VOL) in the same storage. This replica can be used to restore the operating system in case of failure.

Note — You may also use Hitachi Thin Image to recover and restore the operating system LUN.

Follow these guidelines to restore and recover the operating system:

- Create a ShadowImage pair using Hitachi Storage Adapter for SAP HANA cockpit.
- Split the pair using the standalone graphical user interface and the **split** operation. You can enable Host I/O during cloning, but you have to stop the host I/O before splitting pairs.
- When you want to restore the operating system back in case of failures, shut down the operating system and revert the pair using the standalone graphical user interface and the **revert** operation.

Refer to the following for more information:

- “Steps to Configure fstab and Network” on page 20
- “Steps to Register a SAP HANA System” on page 21

Select Host Group, Parity Group, or Dynamic Provisioning Pool

Hitachi Storage Adapter for SAP HANA cockpit needs a host group for Hitachi ShadowImage. Storage Adapter for SAP HANA cockpit also needs a dynamic provisioning pool or a parity group where to create the S-VOL.

Ensure that you have sufficient capacity within the parity group or dynamic provisioning pool, since a full copy of P-VOL will be created.

Create the Pair

To create the pair between the P-VOL and the S-VOL, do the following.

- Select the operating system LUN (P-VOL) which you want to replicate with Hitachi Storage Adapter for SAP HANA cockpit. All the data in the P-VOL is copied to the S-VOL when you create a new ShadowImage pair.

After you select the LUN, Hitachi ShadowImage creates a consistency group (CTG). This automatic creation of the CTG occurs when you run pair creation from the user interface of Storage Adapter for SAP HANA cockpit.

Split the Pair

To split the pair between the P-VOL and S-VOL, do the following.

- Shut down the operating system and split pairs, use the standalone graphical user interface and the **split** operation on the management server that runs Hitachi Storage Adapter for SAP HANA cockpit.

Pair split means splitting the pair while keeping the pair relation. When issuing the pair split command, the difference between the pair volumes at that time is copied from the P-VOL to the S-VOL. This fully synchronizes the data between the pair volumes while maintaining the same CTG.

Resync the Pair

To resync pairs from the P-VOL to the S-VOL, do the following.

- Use the standalone graphical user interface and the **resync** operation on the management server that runs Hitachi Storage Adapter for SAP HANA cockpit.

This copies the difference between the pair volumes after the last pair split from the P-VOL to the S-VOL

Revert the Pair

To reverse sync, or revert, from the S-VOL to the P-VOL, do the following.

- Use the standalone graphical user interface and the **revert** operation on the management server that runs Hitachi Storage Adapter for SAP HANA cockpit.

This copies the difference between the pair volumes after the last pair split from the S-VOL to the P-VOL.

The pair splits automatically after completing the copy.

Recover and Restore the SAP HANA Storage Snapshot

Create a SAP HANA storage snapshot to recover the SAP HANA database after a failure. Hitachi Storage Adapter for SAP HANA cockpit uses Hitachi Thin Image to create a storage snapshot of the data volumes. Restore the SAP HANA database using this replica for a point in time recovery.

Follow these guidelines to use a SAP HANA storage snapshot with Hitachi Thin Image.

- Because storage snapshots miss integrity checks on block level, SAP strongly recommends that you combine storage snapshots with traditional data backups. For example, (1) perform a storage snapshot once a day and (2) perform a traditional backup once a week to the file system or to a 3rd party backup tool. If the traditional backup fails, all storage snapshots after the last successful traditional backup potentially are corrupt, as well.
- Perform a snapshot on only data volumes.
- Host I/O can be enabled during the snapshot.
- It is not possible to create SAP HANA storage snapshots with a SAP HANA multi-tenant database container. A SAP HANA storage snapshot is available only for a SAP HANA single database container, by design.
- Ensure all the pairs are in pair split status (PSUS) before trying to recover the SAP HANA database from SAP HANA Studio.

Select Host Group and Hitachi Thin Image Pool

Hitachi Storage Adapter for SAP HANA cockpit needs a host group for the SAP HANA storage snapshot. Create a host group for the data V-VOLs. Also, create a Hitachi Thin Image pool where you create the V-VOLs.

Create the Pair

To create a pair between the P-VOL and the V-VOL, do the following.

- Select all the data LUNs (P-VOLs) you want to replicate using Hitachi Storage Adapter for SAP HANA cockpit. All the differential data from the P-VOLs is copied to the V-VOLs when you create a new Hitachi Thin Image pair.

Thin Image then creates a consistency group (CTG). This automatic creation of the CTG occurs when you run pair creation from the user interface of Storage Adapter for SAP HANA cockpit.

Split the Pair

Hitachi Storage Adapter for SAP HANA cockpit splits the pair automatically after taking the snapshot. You do not split the pairs manually.

Revert the Pair

To reverse sync or revert from the V-VOL to the P-VOL, do the following.

1. Shut down the operating system.
2. Revert pairs using the standalone graphical user interface and the **revert** operation on the management server that runs Hitachi Storage Adapter for SAP HANA cockpit.

This copies the difference between pair volumes after the last pair split from the V-VOL to the P-VOL.

The pair is split automatically after completing the copy.

Copy and Clone the SAP HANA Database with Hitachi ShadowImage

Use Hitachi ShadowImage to create a local replica (S-VOLs) of the following in the same storage:

- Operating system
- HANA shared
- HANA LOG
- HANA Data

You can use this replica to copy or clone the system.

Follow these guidelines to copy or clone the SAP HANA system.

- It is not possible to copy a SAP HANA single-container system to a SAP HANA multitenant database container.
- Create a Hitachi ShadowImage pair for the operating system LUN (optional), HANA shared LUN, data LUN, and log LUN using Hitachi Storage Adapter for SAP HANA cockpit on the primary system.
 - (1) Select a host group on the target system.
 - (2) Select a dynamic provisioning pool.
 - (3) Select the operating system LUN (optional), HANA shared LUN, HANA log LUN, and data LUN.
 - (4) On the Create ShadowImage user interface of Hitachi Storage Adapter for SAP HANA cockpit, select the LUNS in this specific order:
 - i. Operating system LUN (optional)
 - ii. HANA shared LUNs
 - iii. HANA log LUNs
 - iv. HANA data LUNs
- Split the pairs using the standalone graphical user interface and the **split** operation. Host I/O can be enabled during replication, but you have to stop the host I/O before splitting pairs.

Refer to the following for more information:

- “Steps to Configure fstab and Network” on page 20.
- “Steps to Register a SAP HANA System” on page 21.

Select Host Group, Parity Group, or Dynamic Provisioning Pool

Hitachi Storage Adapter for SAP HANA cockpit needs a host group for Hitachi ShadowImage. Create a host group for the S-VOL. Also, create a dynamic provisioning pool or a parity group where you create the S-VOL.

Ensure that you have sufficient capacity within the parity group or dynamic provisioning pool, since this creates a full copy of P-VOL.

Create the Pair

To create the pair between the P-VOL and the S-VOL, do the following.

- Select the operating system (optional) LUN, HANA shared LUN, HANA log LUN, and the HANA data LUNs (P-VOLs) you want to replicate using Hitachi Storage Adapter for SAP HANA cockpit. All the data in P-VOLs is copied to S-VOLs when you create a new ShadowImage pair.

Hitachi ShadowImage then creates a consistency group (CTG). This automatic creation of the CTG occurs when you run pair creation from the user interface of Storage Adapter for SAP HANA cockpit.

Split the Pair

To split the P-VOL and the S-VOL pair, do the following.

- Stop the host I/O and split pairs, use the standalone graphical user interface and the **split** operation on the management server that runs Hitachi Storage Adapter for SAP HANA cockpit.

Pair split means splitting the pair while keeping the pair relation. When issuing the pair split command, the difference between the pair volumes at the time is copied from the P-VOL to the S-VOL. This fully synchronizes the data between the pair volumes while maintaining the same CTG.

Resync the Pair

To resync pairs from the P-VOL to the S-VOL, do the following.

- Use the standalone graphical user interface and the **resync** operation on the management server that runs Hitachi Storage Adapter for SAP HANA cockpit.

This copies the difference between the pair volumes after the last pair split from the P-VOL to the S-VOL

Limitations of Hitachi Storage Adapter for SAP HANA Cockpit

The following are the limitations for Hitachi Storage Adapter for SAP HANA cockpit.

- Hitachi Thin Image supports only Thin Image pools. You cannot create VVOLs directly on parity groups.
- SAP HANA uses multiple parity groups for HANA data and log volume (P-VOLs). Hitachi ShadowImage supports creation of S-VOLs on dynamic provisioning pools and parity groups.
 - Use a single dynamic provisioning pool or a parity group to create the pairs.
 - Only one CTG will be created for all pairs together.
 - Split operation on a single CTG ensures all pairs are consistent.
 - Although one parity group or one dynamic provisioning pool will host all the S-VOLs, it will work for quality assurance and test systems, as the same performance is not needed.
 - Use multiple dynamic provisioning pools or parity groups to create the pairs.
 - Multiple CTGs will be created for each pair of data and log volumes.
 - Split operation on all CTGs is needed at once to ensure consistency.
 - This method can be used to ensure performance as the storage design for P-VOLs and S-VOLs will have the same multiple parity groups.

- You can have one P-VOL with up to three S-VOL pairs in ShadowImage.
- The S-VOL of Pair 1 (SI_L1) can be used as the P-VOL of Pair 2 (SI_L2). Doing this creates a cascade pair. However, you cannot execute the revert (reverse resync) operation on SI_L2 in a cascade pair. The split of SI-L2 pair first requires the split of the SI-L1 pair.

Steps to Configure fstab and Network

To configure fstab entries and modify network parameters, do the following.

1. Make sure that the S-VOLs are mapped to the storage port of the system that needs to be restored or cloned.
2. Power on the system and enter the BIOS.
3. Rescan the LUNs on the BIOS and make sure that the boot priority is set up for the operating system S-VOL.
4. Boot the system from the S-VOL.
5. Once the system starts booting, enter the recovery mode to perform the following:
 - Edit /etc/fstab to replace the LUN WWNs. Make sure that you replace the WWN of P-VOL with the WWN of the V-VOL:

```
sed -i -e 's/<WWN of P-VOL>/<WWN of V-VOL>/g' /etc/fstab
```
 - Reboot the server.

This successfully restores or clones the operating system so it will be up and running.

Steps to Register a SAP HANA System

Note — Refer to [SAP HANA Help Portal](#) for additional details on registering the SAP HANA system.

Follow the steps below to register the target system using the cloned V-VOLs:

1. Change the IP and Hostname of the target system.
 - (1) Assign a new IP and hostname for the target system.
 - (2) Reboot the system.
2. Register the SAP HANA system with new SID and system number, as follows:
 - (1) Log on as the root user.
 - (2) Change the directory to the HANA Life Cycle Manager directory.

```
# cd /hana/shared/SI D/hdbl cm
```
 - (3) Register the system.

```
# ./hdbl cm --acti on=rename_system
```
 - (4) You can customize the following parameters:
 - Target hostname
 - System usage: Production, Test, Development, Custom
 - Target System ID
 - Target Instance Number
 - Location of Data volumes
 - Location of Log volumes
 - Location of Data backups
 - Location of Log Backups
 - Target System Administrator password
 - (5) Check the target system.

```
# su <SI D>adm
```

```
# HDB start
```

This clones the SAP HANA target system so it will be up and running.

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems [Global Services](#) website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate [Resources](#) website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services [Training and Certification](#) website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the [Hitachi Data Systems](#) website.

Hitachi Data Systems



Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

© Hitachi Data Systems Corporation 2016. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. Hi-Track, ShadowImage, and VSP are trademarks or registered trademarks of Hitachi Data Systems Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

AS-517-01. January 2017.