

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F600, G600, F800, and G800 with Solid State Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

By Yingping Niu and Dietmar Ebert

May 2020

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

Revision	Changes	Date
AS-610-00	Initial release	August 2017
MK-SL-044-00	Previously AS-610-00, updated as follows: <ul style="list-style-type: none">▪ Added multipath configuration▪ How to set up volume groups / logical volumes on LVM-based systems▪ How to set up XFS file systems▪ Recommendations for shared filesystem (e. g. HNAS, and point to documentation and what kind of configurations need to be done)▪ Updated HANA parameters▪ How to install SAP HANA (scale-up, scale-out - LVM-based, scale-out - not LVM-based (if applicable))▪ Reference documentation	March 19, 2018
MK-SL-044-01	Correct multipath settings	May 11, 2020

Table of Contents

Solution Overview	2
Key Solution Elements	5
Hardware Elements	5
Software Elements	6
Solution Design	6
Fibre Channel Architecture	7
Storage Architecture	7
SAP HANA Configuration	13
SAP HANA Software Installation	18
Engineering Validation	23

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F600, G600, F800, and G800 with Solid State Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

This reference architecture guide uses the [SAP HANA tailored data center integration](#) (TDI) approach to implement the SAP HANA platform, rather than using the appliance model. This reference architecture guide provides the storage requirements for the maximum number of validated active SAP HANA production nodes on the following:

- Hitachi Virtual Storage Platform F600
- Virtual Storage Platform G600 using solid state drives (SSD)
- Virtual Storage Platform F800
- Virtual Storage Platform G800 using SSD

With a SAP HANA system, all hardware components are pre-configured by the hardware vendor. With a SAP HANA tailored data center integration deployment, each system is a customized solution and you can choose hardware from any certified SAP HANA server vendor, along with any certified SAP HANA enterprise storage vendor, to implement a SAP HANA platform. You can leverage existing hardware to reduce the total cost of ownership (TCO).

When deploying a SAP HANA TDI solution, SAP only allows using homogeneous compute server hardware from a single hardware partner in a single implementation.

If a certificate provided by SAP is for a specific operating system, you can only use that operating system for SAP HANA in that implementation.

Engineering validation at Hitachi systems for this approach for SAP HANA has the following points:

- Server blades from Hitachi Vantara were used during validation.
- Testing showed that the storage design of Virtual Storage Platform F600, G600, F800, and G800 using SSDs meets the SAP enterprise storage certification key performance indicator (KPI) requirements for a SAP HANA system.
- Scalability and storage KPI testing was performed using [SAP HANA Hardware Configuration Check Tool](#) (HWCCT). Refer to [SAP Note 1943937 - Hardware Configuration Check Tool - Central Note](#) (SAP user credentials required) for details about HWCCT.
- Up to 28 SAP HANA systems were connected in parallel to a single VSP G600/F600 or VSP G800/F800 using SSDs to pass the SAP enterprise storage KPIs. These were configured as RAID10 (2D+2D). There were no other systems or applications using the storage while running HWCCT revision 112 on the 28 SAP HANA systems in parallel for testing.

Note – If you need to have mixed applications running in parallel on the same storage, SAP HANA with non-SAP HANA, you must perform testing for the expected workload using the SAP HANA Hardware Configuration Tool (HWCCT) to meet SAP KPIs.

When implementing a SAP HANA TDI infrastructure on Hitachi Virtual Storage Platform F600, G600, F800, or F800 using SSDs, you do not have to use exactly the same storage design that Hitachi Vantara used for enterprise storage KPI testing, as described in this reference architecture guide. However, it is recommended to follow the directions and guidelines for the setup and configuration of the storage system. Refer to [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#) for more details.

Since the release of SAP HANA TDI in November 2013, several versions of HWCCT have been published. To check if the hardware configuration of your SAP HANA TDI infrastructure meets SAP KPIs, you must use the same version of HWCCT used during the certification of the hardware, compute servers, and storage system for your tests. [SAP Note 1943937](#) describes how to determine the right version of HWCCT for your tests.

Note – Testing of this configuration was performed 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 matches your production environment before your production implementation of this solution.

Solution Overview

This document provides an example configuration of the storage layout for SAP HANA nodes with variable sizes of main memory consolidated on Hitachi Virtual Storage Platform F600 or G600 using solid state drives (SSDs), tested in the Hitachi Vantara lab. The same configuration of the storage layout is applicable for Hitachi Virtual Storage Platform F800 or G800 using solid stable drives.

This configuration uses one of the following for storage:

- **Hitachi Virtual Storage Platform F600 or F800** – This delivers superior all-flash performance for business-critical applications, with continuous data availability.
- **Hitachi Virtual Storage Platform G600 or G800** – This model scales for all data types, flexibly adapting for performance, capacity, and multi-vendor storage.

Table 1, “Hitachi Virtual Storage Platform F600, G600, F800, and G800 Feature Comparison,” on page 3 highlights the features of the available storage used in this TDI architecture.

TABLE 1. HITACHI VIRTUAL STORAGE PLATFORM F600, G600, F800, AND G800 FEATURE COMPARISON


Item	Type		Specifications			
			VSP G600	VSP F600	VSP G800	VSP F800
System	SVOS Version		SVOS 7.x			
	Performance (IOPS)		800,000		1,400,000	
	Maximum number of drives	SFF	720		1,440	
		LFF	720		1,440	
		SSD	720		1,440	
		FMD	288		576	
	Drive Options	SSF	480 GB, 960 GB, 1.9 TB, 3.8 TB		960 GB, 1.9 TB, 3.8 TB	
		LFF/SSD	4 TB, 6 TB, 10 TB			
		FMD	3.5 TB, 7 TB, 14 TB			
	RAID level and group configuration	RAID-10	2D+2D, 4D+4D			
		RAID-5	3D+1P, 4D+1P, 6D+1P, 7D+1P			
		RAID-6	6D+2P, 12D+2P, 14D+2P			
	Maximum number of LDEVs		48K		64K	
Maximum storage capacity		7,049 TB		14,098 TB		
Maximum external configuration		16 PB		64 PB		
Memory	Cache memory capacity and speed		256 GB		512 GB	
			DDR3		DDR3	
Storage Interface	DKC to drive interface		SAS/Dual Port			
	Data transfer rate		12 Gbps			
	Maximum number of DKB		24		48	
Device Interface	Supported channel type	Fibre channel	64		80	
		Data transfer rate	8/16/32 Gb/s			
	Maximum number of Channel Boards		10		16	

TABLE 1. HITACHI VIRTUAL STORAGE PLATFORM F600, G600, F800, AND G800 FEATURE COMPARISON (CONTINUED)

Item	Type	Specifications			
		VSP G600	VSP F600	VSP G800	VSP F800
Non-stop maintenance	Control PCB	Supported			
	Cache memory	Supported			
	Cache flash memory	Supported			
	Power supply, fan	Supported			
	Microcode	Supported			
	Flash drive / flash module drive	Supported			

Figure 1 shows the server to storage configuration of this solution using Hitachi Virtual Storage Platform F600 with 28 SAP HANA systems. The configuration for Virtual Storage Platform G600, G800, or F800 is similar.

Figure 1



Key Solution Elements

These are the key hardware and software elements used for scalability testing.

Hardware Elements

Table 2 describes the hardware required to test the scalability of 28 active nodes on Hitachi Virtual Storage Platform F600, G600, F800, or G800 using solid state drives (SSD).

TABLE 2. HARDWARE ELEMENTS

Hardware	Quantity	Configuration	Role
One of the following Hitachi Virtual Storage Platform models: <ul style="list-style-type: none">VSP F600 or G600 using solid state drivesVSP F800 or G800 using solid state drives	1	<ul style="list-style-type: none">CTL: 1 pairDKB: 1 pairCHB pairs: 7 × 8 Gb/sMPU: 2 pairsCache: 256 GB	Block storage for SAP HANA platform nodes
Certified server for SAP HANA (Note 1)	28	<ul style="list-style-type: none">Rack servers or server blade chassis2-socket server	SAP HANA Nodes
Brocade ICX 6430-48 port switch (optional)	1	<ul style="list-style-type: none">1 GbE48 ports	1 GbE management network
Brocade VDX 6740-48 port switch (optional)	1	<ul style="list-style-type: none">10 GbE48 ports	10 GbE connectivity

Note 1. See [Certified and Supported SAP HANA Hardware](#).

Hitachi Virtual Storage Platform G Series Family

The [Hitachi Virtual Storage Platform G series](#) family enables the seamless automation of the data center. It has a broad range of efficiency technologies that deliver maximum value while making ongoing costs more predictable. You can focus on strategic projects and to consolidate more workloads while using a wide range of media choices.

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.

Server for SAP HANA

Only servers certified for use with SAP HANA can be used in a SAP HANA TDI environment. This follows the exact same bill of materials as the certified SAP HANA appliance server, but without the storage or local disks.

Find a description of all certified servers and enterprise storage solutions in [Certified and Supported SAP HANA Hardware](#). For more information on SAP HANA TDI, consult [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#).

Software Elements

This is the software used to deploy the test configuration.

- SUSE Linux Enterprise Server for SAP Applications
- SAP HANA
- Hitachi Storage Navigator
- Hitachi Command Suite
- Hitachi Storage Virtualization Operating System RF
 - Includes Hitachi Dynamic Provisioning
- Microcode for Hitachi Virtual Storage Platform G600, F600, G800, or F800

Note – Scalability testing was carried out using SUSE Linux Enterprise Server for SAP Applications. However, the solution also supports using Red Hat Enterprise Linux.

Hitachi Storage Virtualization Operating System RF

[Hitachi Storage Virtualization Operating System RF](#) (SVOS RF) spans and integrates multiple platforms. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Storage Virtualization Operating System RF uses Hitachi Dynamic Provisioning (HDP) to provide wide striping and thin provisioning. Dynamic Provisioning provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) without initially allocating any physical space. Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs).

Solution Design

This is the detailed solution example for the SAP HANA tailored data center integration (TDI) on either of the following:

- Hitachi Virtual Storage Platform F600 or G600 using SSD
- Hitachi Virtual Storage Platform F800 or G800 using SSD

Fibre Channel Architecture

For each SAP HANA node, there are two dedicated Fibre Channel ports on the storage. The two Fibre Channel cables directly connect HBA ports on the node to the designated storage target port to achieve the following:

- No single point of failure for high availability
- Expected throughput of data and log volumes of SAP HANA

Table 3 shows the storage port mapping for 4 nodes, as an example.

TABLE 3. STORAGE PORT MAPPING FOR 4 NODES

SAP HANA Node	Fibre Channel Port	Storage Ports
Node001	Port 0	1A
	Port 1	2A
Node002	Port 0	3A
	Port 1	4A
Node003	Port 0	5A
	Port 1	6A
Node004	Port 0	7A
	Port 1	8A

Storage Architecture

Each SAP HANA node needs the following storage layout:

- Operating system (OS) volume
- SAP HANA shared volume
- SAP HANA log volume
- SAP HANA data volume

This SAP HANA TDI setup utilizes dynamic provisioning pools created with Hitachi Dynamic Provisioning for the storage layout. This ensures maximum utilization and optimization at a lower cost than other solutions.

The dynamic provisioning pool layout options with minimal disks and storage cache for 1 up to 28 nodes are listed in Table 4, "Dynamic Provisioning Pools with Disks and Storage Cache ," on page 8.

TABLE 4. DYNAMIC PROVISIONING POOLS WITH DISKS AND STORAGE CACHE

Dynamic Provisioning Pool Name	Purpose	Parity Group RAID Level and Disks	Number of Nodes	Storage Cache
OS_SH_Data_Pool	<ul style="list-style-type: none"> ▪ OS Volume 	1 × RAID-10 (2D+2D) on 1.9 TB SSD	up to 17	256 GB
	<ul style="list-style-type: none"> ▪ SAP HANA Shared Volume 	2 × RAID-10 (2D+2D) on 1.9 TB SSD	up to 24	
	<ul style="list-style-type: none"> ▪ SAP HANA Data Volume 	3 × RAID-10 (2D+2D) on 1.9 TB SSD	up to 28	
Log_Pool	<ul style="list-style-type: none"> ▪ SAP HANA Log Volume 	1 × RAID-10 (2D+2D) on 1.9 TB SSD	up to 21	
		2 × RAID-10 (2D+2D) on 1.9 TB SSD	up to 24	
		3 × RAID-10 (2D+2D) on 1.9 TB SSD	up to 27	
		4 × RAID-10 (2D+2D) on 1.9 TB SSD	up to 28	

A minimum of three parity groups is needed for OS_SH_Data_Pool, and a minimum of four parity groups is needed for Log_Pool to fit 28 SAP HANA production nodes on a Hitachi Virtual Storage Platform G600/F600 or Virtual Storage Platform G800/F800 using RAID10 (2D+2D). Additional parity groups of the same type may need to be added, depending on the various combination of node sizes as well as the number of nodes to meet the capacity requirements.

Provision the storage for the 28 SAP HANA systems. The following example is an example.

- Create two dynamic provisioning pools for the 28 SAP HANA systems on storage:
 - Use **OS_SH_Data_Pool** to provision the operating system volume, SAP HANA shared volume, and Data volume.
 - Use **Log_Pool** to provision the Log volume.
- Assign all LDEVs to the dedicated pool, as shown in Table 5, “Dynamic Provisioning Pool Provisioning with RAID-10 (2D+2D) for 28 Nodes,” on page 9.

TABLE 5. DYNAMIC PROVISIONING POOL PROVISIONING WITH RAID-10 (2D+2D) FOR 28 NODES

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
OS_SH_Data_Pool	1	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:01	OS_SH_DATA_DPVOL_1	878 GB	MPU-10
			00:00:02	OS_SH_DATA_DPVOL_2	878 GB	MPU-11
			00:00:03	OS_SH_DATA_DPVOL_3	878 GB	MPU-20
			00:00:04	OS_SH_DATA_DPVOL_4	878 GB	MPU-21
	2	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:05	OS_SH_DATA_DPVOL_5	878 GB	MPU-10
			00:00:06	OS_SH_DATA_DPVOL_6	878 GB	MPU-11
			00:00:07	OS_SH_DATA_DPVOL_7	878 GB	MPU-20
			00:00:08	OS_SH_DATA_DPVOL_8	878 GB	MPU-21
	3	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:09	OS_SH_DATA_DPVOL_9	878 GB	MPU-10
			00:00:10	OS_SH_DATA_DPVOL_10	878 GB	MPU-11
			00:00:11	OS_SH_DATA_DPVOL_11	878 GB	MPU-20
			00:00:12	OS_SH_DATA_DPVOL_12	878 GB	MPU-21
Log_Pool	4	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:13	LOG_DPVOL_1	878 GB	MPU-10
			00:00:14	LOG_DPVOL_2	878 GB	MPU-11
			00:00:15	LOG_DPVOL_3	878 GB	MPU-20
			00:00:16	LOG_DPVOL_4	878 GB	MPU-21
	5	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:17	LOG_DPVOL_5	878 GB	MPU-10
			00:00:18	LOG_DPVOL_6	878 GB	MPU-11
			00:00:19	LOG_DPVOL_7	878 GB	MPU-20
			00:00:20	LOG_DPVOL_8	878 GB	MPU-21
	6	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:21	LOG_DPVOL_9	878 GB	MPU-10
			00:00:22	LOG_DPVOL_10	878 GB	MPU-11
			00:00:23	LOG_DPVOL_11	878 GB	MPU-20
			00:00:24	LOG_DPVOL_12	878 GB	MPU-21
	7	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:25	LOG_DPVOL_13	878 GB	MPU-10
			00:00:26	LOG_DPVOL_14	878 GB	MPU-11
			00:00:27	LOG_DPVOL_15	878 GB	MPU-20
			00:00:28	LOG_DPVOL_16	878 GB	MPU-21

- Provision virtual volumes for each of the nodes, as follows:
 - Determine the minimum sizes for operating system, data, log, and HANA shared using these formulas in SAP white paper [SAP HANA Storage Requirements](#) as following:
 - Every HANA node requires approximately 100 GB capacity for the operating system
 - /hana/shared size uses the following formulas:
 - Single node (scale-up):** Size = MIN (1 × RAM; 1 TB)
 - Multinode (scale-out):** Size = 1 × RAM_of_worker per 4 worker nodes
 - Data size requires at least 1 × RAM of each HANA node
 - Log size uses the following formulas:
 - Systems with equal to or less than 512 GB memory:** Size = 1/2 × RAM
 - Systems with greater than 512 GB memory:** Size = 512 GB
 - Create virtual volumes (vVOLs) for the operating system, SAP HANA shared, Log, and Data volumes. Table 6, “VVOLs for SAP HANA Nodes for 4 Different Memory Size HANA Systems,” on page 11 shows examples of four different memory sizes of HANA systems including 512 GB, 1 TB, 2 TB and 4 TB.

TABLE 6. VVOLS FOR SAP HANA NODES FOR 4 DIFFERENT MEMORY SIZE HANA SYSTEMS

Dynamic Provisioning Pool	vVOL ID	vVOL Name	vVOL Size	MPU Assignment	System Memory
OS_SH_Data_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-10	512 GB
	00:02:00	HANA_OS_N2	100 GB	MPU-11	1 TB
	00:03:00	HANA_OS_N3	100 GB	MPU-20	2 TB
	00:04:00	HANA_OS_N4	100 GB	MPU-21	4 TB
	00:01:01	HANA_SH_N1	512 GB	MPU-10	512 GB
	00:02:01	HANA_SH_N2	1024 GB	MPU-11	1 TB
	00:03:01	HANA_SH_N3	1024 GB	MPU-20	2 TB
	00:04:01	HANA_SH_N4	1024 GB	MPU-21	4 TB
	00:01:06	HANA_DATA_N1_1	128 GB	MPU-10	512 GB
	00:01:07	HANA_DATA_N1_2	128 GB	MPU-11	
	00:01:08	HANA_DATA_N1_3	128 GB	MPU-20	
	00:01:09	HANA_DATA_N1_4	128 GB	MPU-21	
	00:02:06	HANA_DATA_N2_1	256 GB	MPU-10	1 TB
	00:02:07	HANA_DATA_N2_2	256 GB	MPU-11	
	00:02:08	HANA_DATA_N2_3	256 GB	MPU-20	
	00:02:09	HANA_DATA_N2_4	256 GB	MPU-21	
	00:03:06	HANA_DATA_N3_1	512 GB	MPU-10	2 TB
	00:03:07	HANA_DATA_N3_2	512 GB	MPU-11	
	00:03:08	HANA_DATA_N3_3	512 GB	MPU-20	
	00:03:09	HANA_DATA_N3_4	512 GB	MPU-21	
	00:04:06	HANA_DATA_N4_1	1024 GB	MPU-10	4 TB
	00:04:07	HANA_DATA_N4_2	1024 GB	MPU-11	
	00:04:08	HANA_DATA_N4_3	1024 GB	MPU-20	
	00:04:09	HANA_DATA_N4_4	1024 GB	MPU-21	

TABLE 6. VVOLS FOR SAP HANA NODES FOR 4 DIFFERENT MEMORY SIZE HANA SYSTEMS (CONTINUED)

Dynamic Provisioning Pool	vVOL ID	vVOL Name	vVOL Size	MPU Assignment	System Memory
Log_Pool	00:01:02	HANA_LOG_N1_1	64 GB	MPU-10	512 TB
	00:01:03	HANA_LOG_N1_2	64 GB	MPU-11	
	00:01:04	HANA_LOG_N1_3	64 GB	MPU-20	
	00:01:05	HANA_LOG_N1_4	64 GB	MPU-21	
	00:02:02	HANA_LOG_N2_1	128 GB	MPU-10	1 TB
	00:02:03	HANA_LOG_N2_2	128 GB	MPU-11	
	00:02:04	HANA_LOG_N2_3	128 GB	MPU-20	
	00:02:05	HANA_LOG_N2_4	128 GB	MPU-21	
	00:03:02	HANA_LOG_N3_1	128 GB	MPU-10	2 TB
	00:03:03	HANA_LOG_N3_2	128 GB	MPU-11	
	00:03:04	HANA_LOG_N3_3	128 GB	MPU-20	
	00:03:05	HANA_LOG_N3_4	128 GB	MPU-21	
	00:04:02	HANA_LOG_N4_1	128 GB	MPU-10	4 TB
	00:04:03	HANA_LOG_N4_2	128 GB	MPU-11	
	00:04:04	HANA_LOG_N4_3	128 GB	MPU-20	
	00:04:05	HANA_LOG_N4_4	128 GB	MPU-21	

- While mapping the LUN path assignment for each node, add vVOLs in the following order:

- (1) The operating volume
- (2) The SAP HANA shared volume
- (3) The log volume
- (4) The data volume

Table 7 shows an example configuration of the LUN path assignment for Node001. Configure the LUN assignment similarly for all other nodes.

TABLE 7. EXAMPLE LUN PATH ASSIGNMENT FOR THE SAP HANA CONFIGURATION ON NODE001

LUN ID	LDEV ID	LDEV Name
0000	00:01:00	HANA_OS_N1
0001	00:01:01	HANA_SH_N1
0002	00:01:02	HANA_LOG_N1_1
0003	00:01:03	HANA_LOG_N1_2
0004	00:01:04	HANA_LOG_N1_3
0005	00:01:05	HANA_LOG_N1_4
0006	00:01:06	HANA_DATA_N1_1
0007	00:01:07	HANA_DATA_N1_2
0008	00:01:08	HANA_DATA_N1_3
0009	00:01:09	HANA_DATA_N1_4

SAP HANA Configuration

This describes the SAP HANA volume configuration.

Multipath Configuration

This reference architecture uses Device-mapper Multipath, a native component of the Linux operating system.

Using Device-mapper Multipath allows the configuration of multiple I/O paths between the server blades and storages.

Each node has two I/O paths connected with the storage. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available unless both paths fail.

Device-mapper Multipath is used for the following I/O paths:

- SAP HANA server boot volume
- SAP HANA data volume
- SAP HANA log volume
- SAP HANA shared volume

Hitachi recommends the multipath settings listed in Table 8.

TABLE 8. MULTIPATH SETTINGS

Section	Option	SUSE Linux Enterprise Server	Red Hat Enterprise Linux
defaults		N/A	N/A
blacklist	devnode	$\wedge(\text{ram} \text{raw} \text{loop} \text{fd} \text{md} \text{dm-} \text{sr} \text{scd} \text{st})[0-9]^*$ $\wedge\text{hd}[\text{a-z}]$ $\wedge\text{dcssblk}[0-9]^*$	
devices/device	vendor	HITACHI	
	product	.*	
	user_friendly_names	no	
	path_checker	directio	
	path_grouping_policy	multibus	
	path_selector	queue-length 0	
	uid_attribute	ID_SERIAL	
	failback	immediate	
	rr_weight	uniform	
	rr_min_io_rq	128	1
	features	0	
	no_path_retry	5	
	prio	weightedpath	
prio_args	hbtl 0:0:0:.* 1 1:0:0:.* 2		

HANA Persistent Storage Volume Configuration

For both operating systems, SUSE Linux Enterprise Server and Red Hat Enterprise Linux, Hitachi uses an LVM-based storage layout. Once installing the operating system and multipathing is configured correctly, you can see the assigned LUNs in these directories:

- /dev/mapper
- /dev/disk/by-id

For example:

- /dev/mapper/360060e801227fc00504027fc00000101
- /dev/disk/by-id/scsi-360060e801227fc00504027fc00000101

The last 6 digits of this number indicates the LDEV ID you have used during the LUN assignment. In the example above, 000101 maps to LDEV ID: 00:01:01.

For all the LUNs besides the one hosting the operating system, you need to initialize the LUNs for use by LVM, running the **pvcreate** command, which is part of the lvm2 rpm package. For example:

```
pvcreate -ff -y /dev/mapper/360060e801227fc00504027fc00000101
```

After you have prepared all the LUNs, you need to configure the volume groups using the **vgcreate** command. The names for the volume group differ between a scale-up and scale-out installation. While for scale-up, Hitachi uses vgdata, vglog, and vgshared, the volume groups for scale-out additionally includes the SAP system ID as well as the node number (for example, vgHITdata001, vgHITlog001, vgHITdata002, and vgHITlog002).

The command for creating the volume group takes no specific options. The following example creates the volume group for HANA log in a scale-up scenario using 4 physical disks per LUN.

```
vgcreate vglog /dev/mapper/360060e801227fc00504027fc0000010[2,3,4,5]
```

For creating other volume groups, use the same syntax, exchanging the volume group name as well as the physical disks or LUNs.

Once creating the volume groups, you need to create a logical volume on top. The general syntax is the following:

```
lvcreate --yes --extents=100%VG --stripes <# luns> --stripesize 1024 --name <lv name> <volume group>
```

Use Table 9 to complete the creation of the logical volumes.

TABLE 9. DETAILS FOR CREATING LOGICAL VOLUMES

	Number of LUNs	Logical Volume Name	Volume Group Name
DATA	4 - as per this reference architecture, or number of assigned LUNs	lvdata	<ul style="list-style-type: none"> ▪ Scale-Up: vgdata ▪ Scale-Out: vg<SID>data<node number>
LOG	4 - as per this reference architecture, or number of assigned LUNs	lvlog	<ul style="list-style-type: none"> ▪ Scale-Up: vglog ▪ Scale-Out: vg<SID>log<node number>
SHARED ^a	1 - as per this reference architecture, or number of assigned LUNs	lvshared	<ul style="list-style-type: none"> ▪ Scale-Up: vgshared

a. Only applicable for HANA Scale-Up deployments

Note – If you only use 1 LUN to create the Logical Volumes for DATA, LOG, or SHARED the options --stripes and --stripesize are not needed.

Finally, create the filesystem on top of the logical volume. Hitachi Vantara recommends using the XFS filesystem. In Table 10, lists the filesystem related options to create and mount the filesystem.

TABLE 10. FILESYSTEM CREATE AND MOUNT OPTIONS

		Create options	Mount options	Mount Point
DATA	Scale-up	-F	inode64, nobarrier	/hana/data
	Scale-out		N/A	/hana/data/<SID>
LOG	Scale-up	-F	inode64, nobarrier	/hana/log
	Scale-out		N/A	/hana/log/<SID>
SHARED	Scale-up	-F	inode64, nobarrier	/hana/shared
	Scale-out		vers=3,proto=tcp,hard,intr,timeo=600,retrans=2,wsiz=65536,rsiz=65536 ^a	/hana/shared/<SID>

a. These options were tested and verified on a Hitachi NAS Platform server.

To create a filesystem, use the following command:

```
mkfs.xfs <create options> /dev/mapper/<vg name>--<lv name>
```

For example:

```
mkfs.xfs -F /dev/mapper/vglog-lvlog
```

HANA Persistent Storage Volume Configuration for Scale-Up and Scale-Out Deployments

For scale-up and scale-out systems, you need to persist the filesystems, including the mount options in one of the operating system's startup files, /etc/fstab/, in order to mount the filesystems automatically during boot operations.

If you are using Hitachi NAS Platform with a scale-out system, see "HANA Persistent Storage Volume Configuration for Scale-Out Deployments Using Hitachi NAS Platform" on page 17.

To persist the filesystems so that they mount automatically during boot operations, do the following.

1. Add the following entry for each filesystem to /etc/fstab:

```
/dev/mapper/<vg name>--<lv name> <mount point> xfs <mount options> 0 0
```

Refer to Table 9, "Details for Creating Logical Volumes," on page 15 and Table 10, "Filesystem Create and Mount Options," on page 16 for volume group and logical volume names, as well as for the mount options.

2. Create the mount points using the following:

```
mkdir -p -m 755 <mount point>
```

- This is an example for a scale-up configuration:

```
mkdir -p -m 755 /hana/{shared,log,data}
```

- This is an example for a scale-out configuration, assuming *HIT* is your SAP system ID:

```
mkdir -p -m 755 /hana/{shared,log,data}/HIT
```

3. Mount the filesystems,

- To mount the filesystems one by one, use this:

```
mount <mount point>
```

- To mount the filesystems all at once, use this:

```
mount -a
```

You can check the mounted filesystem using either the **df** or the **mount** command.

HANA Persistent Storage Volume Configuration for Scale-Out Deployments Using Hitachi NAS Platform

To persist the filesystems so that they mount automatically during boot operations with scale-out systems using Hitachi NAS Platform, do the following.

1. To create the mount points for SHARED, LOG, and DATA on each server, assuming *HIT* is your SAP System ID, use this command:

```
mkdir -p -m 755 /hana/{shared,log,data}/HIT
```

2. To persist the HANA shared filesystems in */etc/fstab*, add the following entry to */etc/fstab*:

```
<ip of HNAS>:<share> <mount point> nfs <mount options> 0 0
```

3. Mount the filesystem:

```
mount /hana/shared/HIT
```

Verify that the filesystem is mounted using either **df** or **mount** command.

HANA takes care of mounting the filesystem with the correct options during the HANA startup itself. Find more details in "SAP HANA Software Installation on a Scale-Up Configuration" on page 18 and "SAP HANA Software Installation on a Scale-Out Configuration" on page 19.

SAP Storage Connector API Fibre Channel Client

The SAP HANA Storage Connector API Fibre Channel Client, **fcClient**, defines a setup interface functions call during normal SAP HANA cluster operation and failover handling.

The scale-out configuration for SAP HANA uses the **fcClientLVM** implementation of the API, which supports the use of the logical volume manager. SAP supports this solution to enable the use of high-performance Fibre Channel devices in a scale-out installation.

The **fcClientLVM** implementation uses standard Linux commands, such as **multipath** and **sg_persist**. Install and configure these commands.

The **fcClientLVM** implementation is responsible for mounting the SAP HANA volumes. It also implements a proper fencing mechanism during a host failover by means of SCSI-3 persistent reservations for SAP HANA failover.

Configuration of the SAP Storage Connector API is contained within the SAP *global.ini* file in */hana/shared/<SID>/global/hdb/custom/config*.

SAP HANA Software Installation

After configuring the file system for SAP HANA data volume, log volume, and HANA shared volume, install the current version of SAP HANA supported by SAP on the SAP HANA server or servers.

By default, the SAP HANA database as well as the SAP HANA client need to be installed.

SAP HANA Software Installation on a Scale-Up Configuration

Follow this procedure to install SAP HANA on a scale-up configuration. These steps assume the following:

- Your SAP system ID: *HIT*
- Your SAP system number: *10*
- Your fully qualified hostname: *saphanas.company.corp*

To install SAP HANA on a scale-up configuration, do the following.

1. Download a supported version of SAP HANA, and follow the instructions to extract the archive.
2. Install HANA by running the following command:

```
<path to HDB_LCM_LINUX_X86_64>/hdblcm --action install \  
--components=server,client \  
--install_hostagent \  
--number 10 \  
--sapmnt=/hana/shared \  
--sid=HIT \  
--datapath=/hana/data/HIT \  
--logpath=/hana/log/HIT \  
--hostname=saphanas.company.corp \  
--certificates_hostmap= saphanas.company.corp=saphanas.company.corp
```

During the installation, you need to provide various passwords.

3. Once the installation is complete, continue with “SAP HANA Software Configuration” on page 21.

SAP HANA Software Installation on a Scale-Out Configuration

As part of the scale-out installation, you need to prepare a configuration file used by the SAP Storage Connector API Fibre Channel called *global.ini*.

This configuration file describes the communication, persistence, and storage details for the HANA installation routine. Assuming your SAP system ID is HIT, it will have similar content similar to the following:

```
[communication]  
listeninterface = .global  
[persistence]  
basepath_datavolumes = /hana/data/HIT  
basepath_logvolumes = /hana/log/HIT  
[storage]  
ha_provider = hdb_ha.fcClientLVM  
partition_*_*__prtype = 5  
partition_1_log__lvname = vgHITlog001-lvlog  
partition_1_data__lvname = vgHITdata001-lvdata  
partition_2_log__lvname = vgHITlog002-lvlog  
partition_2_data__lvname = vgHITdata002-lvdata  
...
```

Refer to [SAP HANA Fiber Channel Storage Connector Admin Guide](#) for more details and options used by the configuration file.

This file is used only during the installation. You can remove and save it afterwards to another directory, such as `/tmp/hana_install_HIT/global.ini`. Keep the file name *global.ini*.

This procedure assumes the following:

- Your SAP system ID: *HIT*
- Your SAP system number: *10*
- Your fully qualified hostname: *hana001.company.corp*
- You interface listening on the HANA node interconnect network: *hana001ic*
- Your configuration file: `/tmp/hana_install_HIT/global.ini`

To install SAP HANA on a scale-out configuration, do the following.

1. Prepare and save the configuration file.
2. Download a supported version of SAP HANA, and then follow the instructions to extract the archive.

3. Install SAP HANA by running the following command:

```
<path to HDB_LCM_LINUX_X86_64>/hdblcm --action install \  
--components=server,client \  
--install_hostagent \  
--number 10 \  
--sapmnt=/hana/shared \  
--sid=HIT \  
--storage_cfg=/tmp/hana_install_HIT \  
--hostname=hana001ic \  
--certificates_hostmap=hana001ic=hana001
```

During the installation, you need to provide various passwords.

4. Once the installation is complete, continue with “SAP HANA Software Configuration” on page 21.
5. To add a worker node to the installation, use the following command:

```
/hana/shared/HIT/global/hdb/install/bin/hdbaddhost \  
--sid=HIT \  
--install_hostagent \  
--sapmnt=/hana/shared \  
--role=worker \  
--hostname=hana002ic \  
--storage_partition=2
```

Note – Make sure that your master installation is aware of the node you are adding. This means that your configuration needs to include the partition information. The master node owns partition ID 1. The first additional node can be added as node 2, and so forth.

6. To add a standby node to the installation, you do not need to provide the partition information. A standby node gets the needed partition information during a HANA failover. Use the following command:

```
/hana/shared/HIT/global/hdb/install/bin/hdbaddhost \  
--sid=HIT \  
--install_hostagent \  
--sapmnt=/hana/shared \  
--role=standby \  
--hostname=hana003ic
```

SAP HANA Software Configuration

To achieve optimized performance, use the following suggested SAP HANA database parameters:

- "async_read_submit=on"
- "async_write_submit_blocks=all"
- "max_parallel_io_requests=256"

These parameters were used during validation of the environment.

Follow SAP Note [2267798](#) to set up these parameters using the hdbparam tool for SAP HANA 1.0.

Follow SAP Note [2399079](#) to set up these parameters defined in the global.ini for file SAP HANA 2.0.

The following are examples of *global.ini* files used by the installation:

- SAP HANA 2.0 scale-up installation *global.ini* file

```
[communication]  
tcp_backlog = 2048  
  
[fileio]  
max_parallel_io_requests = 256  
async_read_submit = on  
async_write_submit_blocks = all  
  
[multidb]  
mode = multidb  
database_isolation = low  
singletenant = yes  
  
[persistence]  
basepath_datavolumes = /hana/data/HIT  
basepath_logvolumes = /hana/log/HIT
```

- SAP HANA 2.0 scale-out 2+1 installation *global.ini* file

```
[communication]  
tcp_backlog = 2048  
listeninterface = .global  
  
[fileio]  
max_parallel_io_requests = 256  
async_read_submit = on  
async_write_submit_blocks = all  
  
[multidb]  
mode = multidb  
database_isolation = low  
singletenant = yes  
  
[persistence]  
basepath_datavolumes = /hana/data/HIT  
basepath_logvolumes = /hana/log/HIT  
  
[storage]  
ha_provider = hdb_ha.fcClientLVM
```

```
partition_*_*_prtype = 5
partition_1_data__lvmname = vgHITdata001-lvdata
partition_1_log__lvmname = vgHITlog001-lvlog
partition_2_data__lvmname = vgHITdata002-lvdata
partition_2_log__lvmname = vgHITlog002-lvlog
[trace]
ha_fcclientlvm = info
```

- SAP HANA 1.0 scale-up installation global.ini file

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

- SAP HANA 1.0 scale-out 2+1 installation global.ini file

```
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*_prtype = 5
partition_1_data__lvmname = vgHITdata001-lvdata
partition_1_log__lvmname = vgHITlog001-lvlog
partition_2_data__lvmname = vgHITdata002-lvdata
partition_2_log__lvmname = vgHITlog002-lvlog
[trace]
ha_fcclientlvm = info
```

Engineering Validation

The test methodology for validating this SAP HANA tailored datacenter integration (TDI) enterprise storage configuration used SAP HANA Hardware Configuration Check Tool (HWCCT) for testing Hitachi Virtual Storage Platform F600 and G600 for the enterprise storage certification, revision hwcct-112

For More Information

Hitachi Vantara 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 [Services](#) website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the [Resources](#) website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara [Training and Certification](#) website.

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

Hitachi Vantara



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

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

© Hitachi Vantara LLC, 2020. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. VSP is a trademark or registered trademark of Hitachi Vantara LLC. 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 Vantara.

MK-SL-044-01. May 2020.