HITACHI
Inspire the Next

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform 5000 Series with Hitachi Storage Virtualization Operating System RF

Reference Architecture Guide

MK-SL-167-04
August 2024

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
MK-SL-167-04	▪ Engineering validation for SAP recertification	August 2024
MK-SL-167-03	▪ Added support for VSP 5200 ▪ Added support for VSP 5500 ▪ Added support for VSP 5600	October 11, 2021
MK-SL-167-02	▪ Updated CPU information in Table 2 ▪ Correction for multipath setting for rr_min_io_rq: for SLES = 128 and for RHEL = 1	April 28, 2020

Reference Architecture Guide

Use this reference architecture guide to implement a SAP HANA tailored data center integration (TDI) solution for the SAP HANA platform. This is different from the appliance model, which combines storage and server.

This solution provides the storage requirements for the maximum number of validated active production nodes in an SAP HANA deployment on the following Hitachi Virtual Storage Platform 5000 series (VSP 5000 series) models:

- VSP 5100/5100H using solid state drives (SSDs)
- VSP 5200/5200H using SSDs
- VSP 5500/5500H using SSDs
- VSP 5600/5600H using SSDs

You can implement SAP HANA as follows:

- With an SAP HANA appliance deployment, the hardware vendor pre-configures all hardware components, storage, and servers for each installation.
- With an SAP HANA tailored data center integration deployment, customize your installation by choosing hardware from any certified SAP HANA server vendor and any certified SAP HANA enterprise storage vendor. This way you may be able to leverage existing hardware to reduce the total cost of ownership (TCO).

When deploying a TDI solution for SAP HANA, 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 in your SAP HANA implementation.

Engineering validation for this solution has the following points:

- Servers from Hitachi Vantara were used. This solution is also valid for other Hitachi-supported servers.
- Testing showed that Virtual Storage Platform 5000 series meets the SAP enterprise storage certification key performance indicator (KPI) requirements for SAP HANA.

- Scalability and storage KPI testing was performed using:
  - SAP HANA Hardware Configuration Check Tool (HWCCT). See SAP Note 1943937 - Hardware Configuration Check Tool - Central Note (SAP user credentials required) for HWCCT details.
  - SAP HANA Hardware and Cloud Measurement Tool (HCMT). See SAP Note 2493172 - Hardware and Cloud Measurement Tools (SAP user credentials required) for HCMT details.
- The storage platform passed the SAP enterprise storage KPIs. While running HWCCT or HCMT in the SAP HANA systems during testing, there were no other systems or applications using the storage, with the storage fully dedicated to this testing.

Note: If you need to have mixed applications running in parallel on the same storage system, SAP HANA and non-SAP HANA, you must perform testing for the expected workload using the SAP HANA Hardware Configuration Tool or Cloud Measurement Tool to meet SAP KPIs.

The following table lists supported SAP HANA systems with VSP 5000 series storage. Scalability is extrapolated for the maximum number of nodes using engineering validation data and information.

Storage Model Name	Drive Type	Maximum HANA Systems	Minimum Parity Group	
			Data	Log
VSP 5100, VSP 5100H, VSP 5200, VSP 5200H (1 pair of nodes)	SAS SSDs	37	7 × RAID 10 (2D+2D)	7 × RAID 10 (2D+2D)
VSP 5500, VSP 5500H, VSP 5600, VSP 5600H (1 pair of nodes)	SAS SSDs	74	14 × RAID 10 (2D+2D)	14 × RAID 10 (2D+2D)
VSP 5500, VSP 5500H, VSP 5600, VSP 5600H (2 pair of nodes)	SAS SSDs	148	28 × RAID 10 (2D+2D)	28 × RAID 10 (2D+2D)

Storage Model Name	Drive Type	Maximum HANA Systems	Minimum Parity Group	
			Data	Log
VSP 5500, VSP 5500H, VSP 5600, VSP 5600H (3 pair of nodes)	SAS SSDs	222	42 × RAID 10 (2D+2D)	42 × RAID 10 (2D+2D)

Note: Since the release of SAP HANA TDI in November 2013, several versions of HWCCT have been published. To check to see if the hardware configuration of your SAP HANA TDI infrastructure meets the SAP KPIs, you must use the same version of HWCCT for HCMT used during the certification of the hardware, compute servers, and storage system for your tests. SAP Notes 1943937 and 249317 describe how to determine the right version of HWCCT or HCMT for your tests.

When implementing an SAP HANA TDI infrastructure on Hitachi Virtual Storage Platform 5000 series using SSDs, you do not have to use exactly same storage design in this reference architecture guide that was used for enterprise storage KPI testing. However, the recommendation is to follow these directions and guidelines for the setup and configuration of the storage system. See SAP HANA Tailored Data Center Integration - Frequently Asked Questions for more details.

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 reference architecture guide provides example configurations tested in the Hitachi Vantara lab using SSDs of the storage layout for SAP HANA nodes on Hitachi Virtual Storage Platform 5000 series.

The following table lists highlights of the features of Hitachi Virtual Storage Platform 5000 series.

Item			Specifications
	Hitachi Storage Virtualization Operating System RF (SVOS) version		SVOS 9.1.2 for VSP 5500/5500H/ 5100/5100H with SAS SSD SVOS 9.1.3 for VSP 5500/5500H/ 5100/5100H with NVMe SSD (does not support 2 and 3 pair nodes) SVP 9.8 for VSP 5600/5600H/ 5200/5200H
	Number of nodes		1 pair, 2 pairs, 3 pairs for VSP 5500/5500H/5600/5600H 1 pair for VSP 5100/5100H/ 5200/5200H
	Maximum number of drives per pair node	SAS SFF	768
		SAS LFF	384
		NVMe SSD	96
		FMD	192
	Drive Options	SAS SFF	600 GB, 2.4 TB
		SAS LFF	10 TB, 14 TB
		SAS SSD	960 GB, 1.9 TB, 3.8 TB, 7.6 TB, 15.3 TB, 30.6 TB
		NVMe SSD	1.9 TB, 3.8 TB, 7.6 TB, 15.3 TB, 30.6 TB
		FMD	7.0 TB, 14.0 TB
	RAID Level		RAID 10, RAID 5, RAID 6
	RAID group configuration	RAID 10	2D+2D, 4D+4D
		RAID 5	3D+1P, 7D+1P
		RAID 6	6D+2P, 14D+2P
	Maximum number of LDEVs		65280
	Maximum storage capacity per pair node		20.5 PB
	Maximum external configuration		255 PB

Item			Specifications
Memory	Cache memory capacity per pair node		2 TB for VSP 5500/5500H/ 5600/5600H 1 TB for VSP 5100/5100H/ 5200/5200H
Storage I/F	DKC to drive interface		SAS/Dual Port for SAS SSD PCIe/Dual Port for NVMe SSD
	Data transfer rate		12 Gbps for SAS SSD 8 Gbps for NVMe SSD
	Maximum number of DKBs		4/node for VSP 5500/5500H/ 5600/5600H 2/node for VSP 5100/5100H/ 5200/5200H
Device I/F	Supported channel type	Fibre Channel	32/node for VSP 5500/5500H/ 5600/5600H 16/node for VSP 5100/5100H/ 5200/5200H
		Data transfer rate	4/8/16/32 Gbps
	Maximum number of CHBs		8/node for VSP 5500/5500H/ 5600/5600H 4/node for VSP 5100/5100H/ 5200/5200H
Non-stop maint-enance	Control PCB		Supported
	Cache memory		Supported
	Cache flash memory		Supported
	Power supply, fan		Supported
	Microcode		Supported
	Flash drive/flash module drive		Supported

The following figure shows a server-to-storage Fibre Channel switch in a direct connection configuration of this solution. This example uses Hitachi Virtual Storage Platform 5000 series.

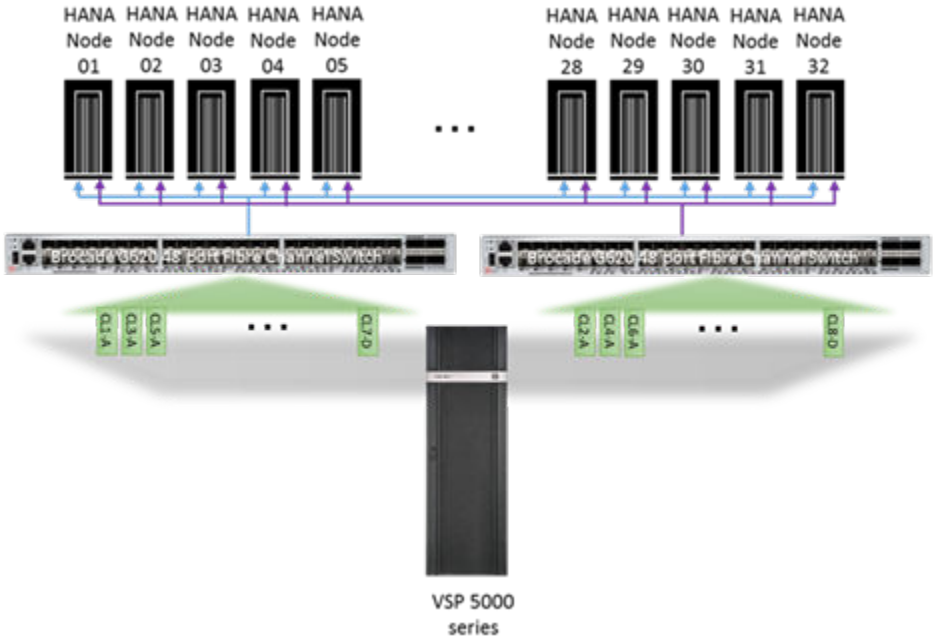

Key solution elements

These are the key hardware and software elements used for this environment.

Hardware elements

The following table lists the tested hardware for maximum scalability of active nodes on Hitachi Virtual Storage Platform 5000 series using SAS SSDs.

Hardware	Tested number of SAP HANA nodes	Storage Components	Fibre Channel Switch
Hitachi Virtual Storage Platform 5500 or VSP 5500H ▪ 1 pair node	32	CTL: 2 pairs CHB: 4 pairs (32 Gbps) MPU: 2 pairs Cache: 2 TB 48 × 960 GB SAS SSDs	48-port Brocade G620 Fibre Channel Switch: 1 pair

The following table shows the optional hardware that can be used with Virtual Storage Platform 5000 series in your environment.

Optional Items	Components
Management server, using this server	Hitachi Advanced Server DS120
1 GbE management network switch, using either one of these switches	Brocade ICX 6430-48 Port
	Cisco Nexus 3048 or 92348
10 GbE network switch, using either one of these switches	Brocade VDX 6740-48 port
	Cisco93180YC-EX/FX/FX3

Server for an SAP HANA TDI environment

Only servers certified for use with SAP HANA can be used in an SAP HANA tailored data center integration (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 list of all certified servers and enterprise storage solutions in Certified and Supported SAP HANA Hardware.

For more information on SAP HANA TDI, see SAP HANA Tailored Data Center Integration - Frequently Asked Questions.

Software elements

This environment uses the following software.

- SUSE Linux Enterprise Server for SAP Applications
  - Scalability testing was carried out using SUSE Linux Enterprise Server for SAP Applications. However, this solution also supports Red Hat Enterprise Linux.
- SAP HANA
- Hitachi Command Suite
- Hitachi Storage Virtualization Operating System RF
  - This includes Hitachi Storage Navigator and Hitachi Dynamic Provisioning.
- Hitachi Storage Advisor

Solution design

This is the detailed solution example for the SAP HANA tailored data center integration (TDI) on Hitachi Virtual Storage Platform 5500 or VSP 5500H using SSDs.

However, the same configuration is applicable to Hitachi Virtual Storage Platform 5100, VSP 5100H, VSP 5200, VSP 5200H, VSP 5600 or VSP 5600H using SSDs.

Fibre Channel architecture

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

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

The following table shows an example storage port mapping for 32 HANA nodes on VSP 5500 or VSP 5500 H. The storage port names can be different for VSP 5100, VSP 5100H, VSP 5200 and VSP 5200H (e.g. 1E, 3E, 5E and 7E instead of 1B, 3B, 5B and 7B). Make sure to use the right port names.

SAP HANA Node	HBA Port		Fibre Channel Switch		Storage Target Port		
	Port Name	Port Speed	Host	Storage	Host Group Name*	Port Speed	Port Security
HANA Node 1	Port 0	16 Gbps	SW-1-P16	SW-1-P0	CBX0-1A-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P16	SW-2-P0	CBX1-2A-1	32 Gbps	Enabled
HANA Node 2	Port 0	16 Gbps	SW-1-P17	SW-1-P0	CBX0-1A-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P17	SW-2-P0	CBX1-2A-2	32 Gbps	Enabled
HANA Node 3	Port 0	16 Gbps	SW-1-P18	SW-1-P1	CBX0-3A-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P18	SW-2-P1	CBX1-4A-1	32 Gbps	Enabled
HANA Node 4	Port 0	16 Gbps	SW-1-P19	SW-1-P1	CBX0-3A-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P19	SW-2-P1	CBX1-4A-2	32 Gbps	Enabled
HANA Node 5	Port 0	16 Gbps	SW-1-P20	SW-1-P2	CBX0-5A-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P20	SW-2-P2	CBX1-6A-1	32 Gbps	Enabled
HANA Node 6	Port 0	16 Gbps	SW-1-P21	SW-1-P2	CBX0-5A-2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch		Storage Target Port		
	Port Name	Port Speed	Host	Storage	Host Group Name*	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P22	SW-2-P2	CBX1-6A-2	32 Gbps	Enabled
HANA Node 7	Port 0	16 Gbps	SW-1-P23	SW-1-P3	CBX0-7A-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P23	SW-2-P3	CBX1-8A-1	32 Gbps	Enabled
HANA Node 8	Port 0	16 Gbps	SW-1-P24	SW-1-P3	CBX0-7A-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P24	SW-2-P3	CBX1-8A-2	32 Gbps	Enabled
HANA Node 9	Port 0	16 Gbps	SW-1-P25	SW-1-P4	CBX0-1B-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P25	SW-2-P4	CBX1-2B-1	32 Gbps	Enabled
HANA Node 10	Port 0	16 Gbps	SW-1-P26	SW-1-P4	CBX0-1B-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P26	SW-2-P4	CBX1-2B-2	32 Gbps	Enabled
HANA Node 11	Port 0	16 Gbps	SW-1-P27	SW-1-P5	CBX0-3B-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P27	SW-2-P5	CBX1-4B-1	32 Gbps	Enabled
HANA Node 12	Port 0	16 Gbps	SW-1-P28	SW-1-P5	CBX0-3B-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P28	SW-2-P5	CBX1-4B-2	32 Gbps	Enabled
HANA Node 13	Port 0	16 Gbps	SW-1-P29	SW-1-P6	CBX0-5B-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P29	SW-2-P6	CBX1-6B-1	32 Gbps	Enabled
HANA Node 14	Port 0	16 Gbps	SW-1-P30	SW-1-P6	CBX0-5B-2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch		Storage Target Port		
	Port Name	Port Speed	Host	Storage	Host Group Name*	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P30	SW-2-P6	CBX1-6B-2	32 Gbps	Enabled
HANA Node 15	Port 0	16 Gbps	SW-1-P31	SW-1-P7	CBX0-7B-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P31	SW-2-P7	CBX1-8B-1	32 Gbps	Enabled
HANA Node 16	Port 0	16 Gbps	SW-1-P32	SW-1-P7	CBX0-7B-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P32	SW-2-P7	CBX1-8B-2	32 Gbps	Enabled
HANA Node 17	Port 0	16 Gbps	SW-1-P33	SW-1-P8	CBX0-1C-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P33	SW-2-P8	CBX1-2C-1	32 Gbps	Enabled
HANA Node 18	Port 0	16 Gbps	SW-1-P34	SW-1-P8	CBX0-1C-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P34	SW-2-P8	CBX1-2C-2	32 Gbps	Enabled
HANA Node 19	Port 0	16 Gbps	SW-1-P35	SW-1-P9	CBX0-3C-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P35	SW-2-P9	CBX1-4C-1	32 Gbps	Enabled
HANA Node 20	Port 0	16 Gbps	SW-1-P36	SW-1-P9	CBX0-3C-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P36	SW-2-P9	CBX1-4C-2	32 Gbps	Enabled
HANA Node 21	Port 0	16 Gbps	SW-1-P37	SW-1-P10	CBX0-5C-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P37	SW-2-P10	CBX1-6C-1	32 Gbps	Enabled
HANA Node 22	Port 0	16 Gbps	SW-1-P38	SW-1-P10	CBX0-5C-2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch		Storage Target Port		
	Port Name	Port Speed	Host	Storage	Host Group Name*	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P38	SW-2-P10	CBX1-6C-2	32 Gbps	Enabled
HANA Node 23	Port 0	16 Gbps	SW-1-P39	SW-1-P11	CBX0-7C-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P39	SW-2-P11	CBX1-8C-1	32 Gbps	Enabled
HANA Node 24	Port 0	16 Gbps	SW-1-P40	SW-1-P11	CBX0-7C-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P40	SW-2-P11	CBX1-8C-2	32 Gbps	Enabled
HANA Node 25	Port 0	16 Gbps	SW-1-P41	SW-1-P12	CBX0-1D-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P41	SW-2-P12	CBX1-2D-1	32 Gbps	Enabled
HANA Node 26	Port 0	16 Gbps	SW-1-P42	SW-1-P12	CBX0-1D-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P42	SW-2-P12	CBX1-2D-2	32 Gbps	Enabled
HANA Node 27	Port 0	16 Gbps	SW-1-P43	SW-1-P13	CBX0-3D-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P43	SW-2-P13	CBX1-4D-1	32 Gbps	Enabled
HANA Node 28	Port 0	16 Gbps	SW-1-P44	SW-1-P13	CBX0-3D-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P44	SW-2-P13	CBX1-4D-2	32 Gbps	Enabled
HANA Node 29	Port 0	16 Gbps	SW-1-P45	SW-1-P14	CBX0-5D-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P45	SW-2-P14	CBX1-6D-1	32 Gbps	Enabled
HANA Node 30	Port 0	16 Gbps	SW-1-P46	SW-1-P14	CBX0-5D-2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch		Storage Target Port		
	Port Name	Port Speed	Host	Storage	Host Group Name*	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P46	SW-2-P14	CBX1-6D-2	32 Gbps	Enabled
HANA Node 31	Port 0	16 Gbps	SW-1-P47	SW-1-P15	CBX0-7D-1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P47	SW-2-P15	CBX1-8D-1	32 Gbps	Enabled
HANA Node 32	Port 0	16 Gbps	SW-1-P48	SW-1-P15	CBX0-7D-2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P48	SW-2-P15	CBX1-8D-2	32 Gbps	Enabled
* - Host Group Names listed indicate the controller box (for example, CBX0 or CBX1) and port name (for example, 1A, 2A, 3A, 4A, and so on.).							

Storage architecture

Each SAP HANA node requires 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 the following two dynamic provisioning pools per module on Hitachi Virtual Storage Platform 5000 series created with Hitachi Dynamic Provisioning for the storage layout. This ensures maximum utilization and optimization at a lower cost than other solutions.

- OS_SH_DT_Pool for the following:
  - OS volume
  - SAP HANA shared volume
  - SAP HANA data volume
- LG_Pool for the following:
  - SAP HANA log volume

The validated dynamic provisioning pool layout options with minimal disks and storage cache on Virtual Storage Platform 5000 series are listed in the following table.

Storage	Cache	HANA Node Number	Number of Parity Groups in OS_SH_DT_Pool	Number of Parity Groups in LG_Pool
Hitachi Virtual Storage Platform 5100 or VSP 5100H	1 TB	Up to 26	1	1
		Up to 32	2	1
Hitachi Virtual Storage Platform 5500 or VSP 5500H 1 pair node	2 TB	Up to 32	1	1
Hitachi Virtual Storage Platform 5200 or VSP 5200H	1 TB	Up to 32	1	1
Hitachi Virtual Storage Platform 5600 or VSP 5600H 1 pair node	2 TB	Up to 32	1	1

To fit 32 SAP HANA nodes, in addition to RAID 10 (2D+2P), you can use RAID 6 (6D+2P) for capacity efficiency. As an example, for the Virtual Storage Platform 5000 series listed in the previous table, you need the following:

- A minimum of 1 parity group is needed for OS_SH_DT_ Pool to fit 32 SAP HANA production nodes on one Virtual Storage Platform 5500, VSP 5500H, VSP 5600, or VSP 5600H using RAID 6 (6D+2P).
- A minimum of 2 parity groups is needed for OS_SH_DT_ Pool to fit 32 SAP HANA production nodes on one Virtual Storage Platform 5100, or VSP 5100H using RAID-6 (6D +2P).
- A minimum of 1 parity groups is needed for OS_SH_DT_ Pool to fit 32 SAP HANA production nodes on one Virtual Storage Platform 5200, or VSP 5200H using RAID-6 (6D +2P).
- A minimum of 1 parity group is needed for LG_ Pool to fit 32 SAP HANA production nodes on one Virtual Storage Platform 5000 series using RAID 6 (6D+2P).

Additional parity groups of the same type might need to be added. Drive boxes can be needed if the internal drives do not have enough capacity, depending on the following:

- The various combinations of HANA node sizes
- The number of HANA nodes to meet the capacity requirements

While it is not limited to these systems, this SAP HANA tailored data center integration solution uses the following four active SAP HANA systems as examples:

- System 1 — 384 GB
- System 2 — 768 GB
- System 3 — 1536 GB
- System 4 — 3072 GB

Provision the storage for the four SAP HANA systems listed previously and follow these guidelines:

- Determine the minimum sizes for operating system, data, log, and HANA shared using these formulas in the SAP white pager SAP HANA Storage Requirements, as follows:
  - Every HANA node requires approximately 100 GB capacity for the operating system.
  - /hana/shared size uses 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 on each HANA node
  - Log size uses these formulas:

    Systems with equal or less than 512 GB memory — size = ½ × RAM

    Systems with greater than 512 GB memory — size = 512 GB
- Provision the storage:
  - Create two dynamic provisioning pools on the storage for the three SAP HANA systems:
    - Use OS_SH_DT_Pool to provision the operating system volume, SAP HANA shared volume, and data volume.
    - Use LG_Pool to provision the log volume.
  - For SSDs, create the parity groups first, as shown in the following table, using the RAID 6 design.

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
OS_SH_DT_Pool	1	RAID 6 (6D+2P) on 960 GB SSDs	00:00:01	OS_SH_DT_1	1320 GB	MPU-010
			00:00:02	OS_SH_DT_2	1320 GB	MPU-020

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
			00:00:03	OS_SH_DT_3	1320 GB	MPU-110
			00:00:04	OS_SH_DT_4	1320 GB	MPU-120
	2	RAID 6 (6D+2P) on 960 GB SSDs	00:00:05	OS_SH_DT_5	1320 GB	MPU-010
			00:00:06	OS_SH_DT_6	1320 GB	MPU-020
			00:00:07	OS_SH_DT_7	1320 GB	MPU-110
			00:00:08	OS_SH_DT_8	1320 GB	MPU-120
	3	RAID 6 (6D+2P) on 960 GB SSDs	00:00:09	OS_SH_DT_9	1320 GB	MPU-010
			00:00:0A	OS_SH_DT_10	1320 GB	MPU-020
			00:00:0B	OS_SH_DT_11	1320 GB	MPU-110
			00:00:0C	OS_SH_DT_12	1320 GB	MPU-120
LG_Pool	4	RAID 6 (6D+2P) on 960 GB SSDs	00:00:0D	LG_1	1320 GB	MPU-010
			00:00:0E	LG_2	1320 GB	MPU-020
			00:00:0F	LG_3	1320 GB	MPU-110
			00:00:10	LG_4	1320 GB	MPU-120

- Assign all LDEVs to the dedicated pool as shown in the previous table.

- Create virtual volumes (vVols) for the operating system, SAP HANA shared, log, and data volumes. The following table lists examples for HANA systems with memory of 384 GB, 768 GB, 1536 GB, and 3072 GB.

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment	System Memory
OS_SH_DT_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-010	384 GB
	00:02:00	HANA_OS_N2	100 GB	MPU-020	768 GB
	00:03:00	HANA_OS_N3	100 GB	MPU-110	1536 GB
	00:04:00	HANA_OS_N4	100 GB	MPU-120	3072 GB
	00:01:01	HANA_SH_N1	384 GB	MPU-010	384 GB
	00:02:01	HANA_SH _N2	768 GB	MPU-020	768 GB
	00:03:01	HANA_SH _N3	1536 GB	MPU-110	1536 GB
	00:04:01	HANA_SH _N4	3072 GB	MPU-120	3072 GB
	00:01:06	HANA_DATA_N1_1	96 GB	MPU-010	384 GB
	00:01:07	HANA_DATA_N1_2	96 GB	MPU-020	
	00:01:08	HANA_DATA_N1_3	96 GB	MPU-110	
	00:01:09	HANA_DATA_N1_4	96 GB	MPU-120	
	00:02:06	HANA_DATA_N2_1	192 GB	MPU-010	768 GB
	00:02:07	HANA_DATA_N2_2	192 GB	MPU-020	
	00:02:08	HANA_DATA_N2_3	192 GB	MPU-110	
	00:02:09	HANA_DATA_N2_4	192 GB	MPU-120	
	00:03:06	HANA_DATA_N3_1	384 GB	MPU-010	1536 GB
	00:03:07	HANA_DATA_N3_2	384 GB	MPU-020	
	00:03:08	HANA_DATA_N3_3	384 GB	MPU-110	
	00:03:09	HANA_DATA_N3_4	384 GB	MPU-120	
	00:04:06	HANA_DATA_N4_1	768 GB	MPU-010	3072 GB
	00:04:07	HANA_DATA_N4_2	768 GB	MPU-020	
	00:04:08	HANA_DATA_N4_3	768 GB	MPU-110	

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment	System Memory
	00:04:09	HANA_DATA_N4_4	768 GB	MPU-120	
LG_Pool	00:01:02	HANA_LOG_N1_1	48 GB	MPU-010	384 GB
	00:01:03	HANA_LOG_N1_2	48 GB	MPU-020	
	00:01:04	HANA_LOG_N1_3	48 GB	MPU-110	
	00:01:05	HANA_LOG_N1_4	48 GB	MPU-120	
	00:02:02	HANA_LOG_N2_1	96 GB	MPU-010	768 GB
	00:02:03	HANA_LOG_N2_2	96 GB	MPU-020	
	00:02:04	HANA_LOG_N2_3	96 GB	MPU-110	
	00:02:05	HANA_LOG_N2_4	96 GB	MPU-120	
	00:03:02	HANA_LOG_N3_1	128 GB	MPU-010	1536 GB
	00:03:03	HANA_LOG_N3_2	128 GB	MPU-020	
	00:03:04	HANA_LOG_N3_3	128 GB	MPU-110	
	00:03:05	HANA_LOG_N3_4	128 GB	MPU-120	
	00:04:02	HANA_LOG_N4_1	128 GB	MPU-010	3072 GB
	00:04:03	HANA_LOG_N4_2	128 GB	MPU-020	
	00:04:04	HANA_LOG_N4_3	128 GB	MPU-110	
	00:04:05	HANA_LOG_N4_4	128 GB	MPU-120	

- While mapping the LUN path assignment for each node, add the vVols in the following order:
  1. Map the operating system volume for the specific SAP HANA platform node.
  2. Map the SAP HANA shared for the specific SAP HANA platform node.
  3. Map the log volume and data volume for the specific SAP HANA platform node.

The following table lists an example configuration of the LUN path assignments for Node 1. The LUN assignments would be the similar for all the other nodes.

LUN ID	LDEV ID	vVol Name	File System Mount Point
0000	00:01:00	HANA_OS_N1	/(root)
0001	00:01:01	HANA_SH_N1	/hana/shared

LUN ID	LDEV ID	vVol Name	File System Mount Point
0002	00:01:02	HANA_LOG_N1_1	/hana/log
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	/hana/data
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 section describes the SAP HANA volume and operating system 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 storage.

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

  Not applicable if using internal disks for the boot volume.
- SAP HANA data volume
- SAP HANA log volume
- SAP HANA shared volume

  For a scale-out solution, Hitachi Vantara uses multiple paths to the disks with two options:

  - Hitachi NAS Platform as NFS for /hana/shared. This solution is valid also for the file systems listed in SAP Note 405827 for SAP HANA.
  - Global File System 2 (GFS2) for /hana/shared. This is for Hitachi-only TDI scale-out deployments where Hitachi provides server and storage for the TDI landscape.

Hitachi Vantara recommends the multipath settings listed in the following table.

Section	Option	SUSE Linux Enterprise Server for SAP Applications	Red Hat Enterprise Linux
defaults		N/A	N/A
blacklist	devnode	^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]* ^hd[a-z] ^dcssblk[0-9]*	
devices/device	vendor	HITACHI	
	product	.*	
	user_friendly_names	No	
	path_checker	direction	
	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	

SAP HANA persistent storage volume configuration

For both operating systems, SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux, Hitachi Vantara uses an LVM-based storage layout. Once the operating system is installed and multipathing is correctly configured, you can see the assigned LUNs in the following 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 indicate 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 scale-up and scale-out installations.

- The volume groups for scale-up use vgdata, vglog, and vgshared.
- The volume groups for scale-out also include the SAP system ID as well as the node number. For example, vgHITdata001, vgHITlog001, vgHITdata002, and vgHITlog002.

The command to create the volume group takes no specific options. The following example creates the volume group for SAP HANA log in a scale-up scenario using four physical disks/ LUNs:

```
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.

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

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

Use the following table to complete the creation of logical volumes.

	Number of LUNs	lv Name	vg Name
DATA	4 - following this reference architecture, or the number of assigned LUNs	lvdata	▪ Scale-up: vgdata ▪ Scale-out: vg<SID>data<node number>
LOG	4 - following this reference architecture, or the number of assigned LUNs	lvlog	▪ Scale-up: vglog ▪ Scale-out: vg<SID>log<node number>
SHARED[1]	1 - following this reference	lvshared	▪ Scale-up: vgshared

[1] Only applicable for scale-up deployments for SAP HANA.

	Number of LUNs	lv Name	vg Name
	architecture, or the number of assigned LUNs		

Note: If you only use one LUN to create the logical volumes for data, log, or shared, the options stripes and stripesize are not needed.

Create the file system on top of the logical volume. Hitachi Vantara storage systems use the XFS file system. The following table lists the options to create and mount the file system.

	System Type	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**	N/A	vers=3,proto=tcp,hard,intr,timeo=600,retrans=2,wsize=65536, rsize=65536 (tested and verified on Hitachi NAS Platform) (These options were tested and verified on Hitachi NAS Platform)	/hana/shared/<SID>
* The nobarrier option is no longer valid from RHEL8.1 and SLES15 SP2. ** The listed options are applicable to Hitachi NAS Platform as NFS for /hana/shared. If you use GFS2 for /hana/shared, use mkfs.gfs2 to create the file system. The GFS2 file system will be managed by Linux Distributed Lock Manger (DLM).				

To create a file system, use the following command:

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

For example:

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

SAP HANA persistent storage volume configuration for scale-up deployments

This is only for scale-up systems.

For scale-up systems, you need to persist the file systems, including the mount options, in the operating system's startup file, /etc/fstab/, to mount the file systems automatically during boot operations.

To configure the HANA persistent storage volume in a scale-up system, do the following.

Procedure

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

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

   See the previous two tables for volume group and logical volume names as well as the mount options.

2. To create the mount points, use the following command:

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

   Example for scale-up:

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

3. Mount the file systems.

   To mount the file systems one at a time, use the following command:

   ```
   mount <mount point>
   ```

   Or, mount them all at once:

   ```
   mount -a
   ```

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

SAP HANA persistent storage volume configuration for scale-out deployments using Hitachi NAS Platform

This is only for scale-out systems.

To configure the SAP HANA persistent storage volume in a scale-out system, do the following.

Procedure

1. Create the mount points for SHARED, LOG, and DATA on each server. This example assumes that HIT is your SAP System ID.

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

2. Configure the SAP HANA shared file system using HNAS or GFS2.
   - Hitachi NAS Platform (HNAS)
     a. Persist the SAP HANA shared file systems in /etc/fstab by adding the following entry to /etc/fstab:

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

     b. Mount the file system:

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

   - Global File System 2 (GFS2)

     The logical volume manager (LVM) creates a single striped volume on which the GFS2 volume is created for the SAP HANA shared file system. The Linux Distributed Lock Manager manages shared file systems on a Linux computer cluster.
3. Verify that the file system is mounted using either the **df** or **mount** command.

   SAP HANA takes care of mounting the file system used for log and data with the correct options during HANA startup. For more details, see the following:
   - SAP storage connector API Fibre Channel Client (on page 25)
   - SAP HANA software installation (on page 26)

SAP storage connector API Fibre Channel Client

For a scale-out configuration, SAP HANA offers a ready-to-use storage connector client for configurations with native multipaths of Fiber Channel-attached devices. This enables host auto-failover on block storage.

The Fibre Channel storage connector, fcClient/fcLientLVM, implements the SAP Storage Connector API that provides hooks for the following:

- Database startup
- Failing-over nodes

SAP supports this solution to enable the use of high-performance Fibre Channel devices in a scale-out installation.

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

The fcClient/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 the /hana/shared/<SID>/global/hdb/custom/config.

SAP HANA software installation

After configuring the file system for the 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. Refer to SAP Note 2235581. The SAP_HANA_OS_Release_Support_Matrix.pdf file attachment outlines which combinations of OS RHEL or SLES are supported with the various SAP HANA versions.

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

Install SAP HANA software on a scale-up environment

Follow this procedure to install SAP HANA on a scale-up environment. This procedure assumes the following:

- HIT is your SAP System ID
- 10 is your SAP System Number
- saphanas.company.corp is your fully qualified hostname

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

Procedure

1. Download a supported version of SAP HANA and follow the instructions to extract the archive.
2. Install HANA by typing 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 \
--system_user_password=<password> -p <password> \
--saoadn_oasswird=<password> \
--datapath=/hana/data/HIT \
--logpath=/hana/log/HIT \
--hostname=saphanas.company.corp \
--certificates_hostmap= saphanas.company.corp=saphanas.company.corp
```

3. During the installation you need to provide various passwords.
4. Once the installation is complete, continue with Configure the SAP HANA software (on page 29).

Install SAP HANA software on a scale-out environment

Before you can start with the installation on a scale-out environment, prepare a configuration file used by the SAP Storage Connector API Fibre Channel.

This configuration file describes the communication, persistence, and storage details for the SAP HANA installation routine with similar content to the following (assuming HIT is your SAP System ID).

Create this configuration file before following the procedure:

```
[communication]
listeninterface = .global
tcp_backlog = 2048
sslminprotocolversion = tls12

[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_log lvmname = vgHITlog001-lvlog
partition_1_data lvmname = vgHITdata001-lvdata
partition_2_log lvmname = vgHITlog002-lvlog
partition_2_data lvmname = vgHITdata002-lvdata

[trace]
Ha_fcclient = info
```

See SAP HANA Fiber Channel Storage Connector Admin Guide for more details and configuration options used by this configuration file.

Use this configuration file during the installation. You can remove it afterwards. Save this file as /tmp/hana_install_HIT/global.ini. Keep the file name (global.ini) the same, even if you choose a different directory.

Follow this configuration to complete your installation. This procedure assumes the following:

- HIT is your SAP System ID
- 10 is your SAP System Number
- hana001.company.corp is your fully qualified hostname
- hana001ic is the interface listening on the HANA node interconnect network
- /tmp/hana_install_HIT/global.ini is your configuration file

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

1. Download a supported version of SAP HANA and follow the instructions to extract the archive.
2. 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 \
--system_user_password=<password> -p <password> \
--sapadm_password=<password> \
--storage_cfg=/tmp/hana_install_<number> \
--hostname=hana001ic \
--certificates_hostmap=hana001ic=hana001
```

3. During the installation you need to provide various passwords.
4. After the installation is complete, continue with Configure the SAP HANA software (on page 29).

To add a worker node to the existing installation, use the following command:

```
/hana/shared/HIT/hdblcm/hdblcm \
--action=add_hosts \
--addhosts=hana002ic:role=worker:group=default:storage_partition=2 \
--sid=HIT \
--install_hostagent \
--sapmnt=/hana/shared \
--sapadm_password=<password> -p <password>
```

Note: Make sure that your master installation is aware of the node you are adding. That means 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.

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

```
/hana/shared/HIT/hdblcm/hdblcm \
--action=add_hosts \
addhosts=hana003ic:role=standby:group=default \
--sid=HIT \
--install_hostagent \
--sapmnt=/hana/shared \
--sapadm_password=<password> -p <password>
```

Configure the SAP HANA software

The following are examples of global.ini files used during validation of this environment:

- SAP HANA 2.0 global.ini file for VSP 5200 and VSP 5600:

```
[communication]
tcp_backlog = 2048
listeninterface = .global
[fileio]
max_parallel_io_requests[DATA] = 1024
max_parallel_io_requests[LOG] = 128
max_submit_batch_size[DATA] = 1024
max_submit_batch_size[LOG] = 128
size_kernel_io_queue[DATA] = 1024
size_kernel_io_queue[LOG] = 1024
async_read_submit[DATA] = on
async_read_submit[LOG] = on
async_write_submit_blocks[DATA] = all
async_write_submit_blocks[LOG] = all
min_submit_batch_size[DATA] = 16
min_submit_batch_size[LOG] = 16
async_write_submit_active[DATA] = on
async_write_submit_active[LOG] = on
[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
partition_3_data__lvmname = vgHITdata003-lvdata
partition_3_log__lvmname = vgHITlog003-lvlog
[trace]
ha_fcclientlvm = info
```

- SAP HANA 2.0 global.ini file for VSP 5100 and VSP 5500:

```
[communication]
tcp_backlog = 2048
listeninterface = .global
[fileio]
max_parallel_io_requests[DATA] = 512
max_parallel_io_requests[LOG] = 512
max_submit_batch_size[DATA] = 384
max_submit_batch_size[LOG] = 384
size_kernel_io_queue[DATA] = 1024
size_kernel_io_queue[LOG] = 1024
async_read_submit[DATA] = on
async_read_submit[LOG] = on
async_write_submit_blocks[DATA] = all
async_write_submit_blocks[LOG] = all
min_submit_batch_size[DATA] = 16
min_submit_batch_size[LOG] = 16
async_write_submit_active[DATA] = on
async_write_submit_active[LOG] = on
[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
partition_3_data__lvmname = vgHITdata003-lvdata
partition_3_log__lvmname = vgHITlog003-lvlog
[trace]
ha_fcclientlvm = info
```

- Scale-up installation of SAP HANA 1.0 global.ini file:

```
[persistence]
basepath_datavolumes=/hana/data/HIT
basepath_logvolumes=/hana/log/HIT
Scale-out 2+1 installation of SAP HANA 1.0 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 validation of this SAP HANA tailored datacenter integration (TDI) enterprise storage configuration used the following for a 2+1 HANA scale-out system:

- SAP HANA hardware and cloud measurement tool (HCMT) for testing for the enterprise storage certification, revision hcmt-079 (HCMT is used for SAP HANA 2.0 or higher. Refer to SAP Note 2493172).
- SAP HANA Hardware Configuration Check Tool (HWCCT) for testing for the enterprise storage certification, revision hwcct-212.5. (Refer to SAP Note 1943937 for the version of SAP HANA that is supported by HWCCT)

The following table lists the hardware elements used during 2+1 scale-out test including Network-HA, FC-HA and HANA node failover.

Hardware	Quantity	Configuration	Role

Hitachi Advanced Server DS7040	3	▪ 2 compute modules
		▪ CPU for each SAP HANA node:
		  • 4 Intel Xeon Platinum 8176/8176M processors 28-core, 2.1 GHz, 165 W Or
		  • 4 Intel Xeon Platinum 8180/8180M processors 28-core, 2.5 GHz, 205 W 4 heat sinks for CPU 0/1 of both compute modules
		▪ 3 TB RAM with 48 × 64 GB DIMMs or 24 × 128 GB DIMMs per SAP HANA node
			SAP HANA server

Mellanox Technologies MT27710Family [ConnectX-4 Lx]	4	Mellanox Connectx4 dual port PCIe card for each HANA node	10 GbE SUSE High Availability Extension (HAE) or RHEL High Availability (HA) cluster communication network, SAP HANA inter-cluster network, client network
	6	Cisco SFP+ 3M twinaxial cables for each HANA node	
Emulex LPe31002M6 blade,16 Gbps Fibre Channel Dual HBA	2	Only on Module 0	Connectivity to the external storage sub-system
	2	SFP	
Management Server for Solution	1	▪ 2 Intel 4110 8-core; 85 W; 2.1 GHz ▪ 2 Samsung 32 GB DDR4 R-DIMM 2666 MHz memory ▪ 1 OCP Mezzanine card — Quanta (Intel) X557-T4 10 GbE, 4 ports ▪ 1 PCIe — Intel XXV710 SFP28 10/25 GbE ▪ 2 SATADOM 128 GB MLC	Used as the management server, and runs the following: SAP HANA Studio

Hardware	Quantity	Configuration	Role
Hitachi Virtual Storage Platform 5600 or VSP 5600H (1 pair nodes)	1	▪ Controller box (CBX): 2 (1/ Node) ▪ Controller board (CTL): 4 (1 pair/ CBX) ▪ Disk board (DKB): 4 (1 pair/ CBX) ▪ CPU:8 (1 pair/CTL) ▪ Cache flash memory (CFM): 8 (4/ CBX) ▪ Backup module and fan (BKMF): 8 (2 pair/CBX) ▪ Hitachi interconnect edge (HIE): 4 (1 pair/CBX) ▪ Used Fibre Channel ports: 8 × 16 Gbps ▪ MPU: 1 pair ▪ Maximum cache: 2048 GB (1024 GB/CBX) ▪ Used channel board (CHB): 1 pair ▪ SSD drives: 9 × 960 GB SSDs (one spare drive) ▪ Drive box (DBS2): 4 (1/DKU)	
CiscoNexus 93180YC-FX switch	4	▪ Two switches with distinct VLANs, each dedicated to cluster FS and SAP HANA inter-cluster network ▪ Two switches with one VLAN to provide uplink network to customer network infrastructure	10 GbE cluster FS and inter-cluster network 10 GbE client network
Cisco Nexus 92348 port switch	1	48 × 1 GbE ports	1 GbE management network

The following table and figure show the server components in the engineering validation.

Feature	Three HANA Nodes (2+1) Scale-out Configuration
Server	3 DS7040 servers

Feature	Three HANA Nodes (2+1) Scale-out Configuration
	Each DS7040 server as one SAP HANA node with the following: ▪ Module 0 (primary) ▪ Module 1 (non-primary)
Network ports	2 × 2-port Mellanox Connect X-4 LX PCIe card per compute module in the following locations: ▪ PCIE Slot 1 ▪ PCIE Slot 3
Fibre Channel ports	2 × 2-port LPE31002-M6 16 Gbps PCIe card on Module 0 for each HANA node in the following locations: ▪ PCIE Slot 0 ▪ PCIE Slot 2
Other interfaces	For all sizes: ▪ 1 USB 3.0 port ▪ KVM connector (VGA, COM, USB 2.0 port)

Hitachi Advanced Server DS7040 (Front)

Hitachi Advanced Server DS7040 PCIe Numbering (Back)

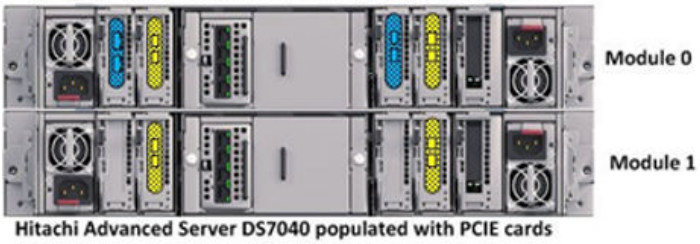

Hitachi Advanced Server DS7040 populated with PCIE cards (Back)

Mellanox Connect-4x LX 2-port 10GbE Network Card

Emulex LPE 31002 16 GB/s PCIe HBA Card

The network configuration uses four dual-port 10 GbE PCIe cards for each SAP HANA node to meet the requirements of no single point of failure (NSPOF) and 10 GbE equivalent throughput. Bond two ports from different PCIe network adapters at the operating system level using link aggregation, following the IEEE 802.3ad link aggregation standard for each of the following three networks:

- SAP HANA inter-cluster network
- SAP HANA NFS network
- SAP HANA client network

For each SAP HANA node, there are two Emulex LPe31002 M6 16 Gbps, 2-port Fibre Channel PCIe adapters. Use Fibre Channel cables to connect with the designated Hitachi Virtual Storage Platform 5000 series Fibre Channel ports to achieve no single point of failure and high performance.

The following table lists the storage configuration used for the 2+1 scale-out test.

Dynamic Provisioning Pool Name	Purpose
Microcode Level	93-03-21-40/00
Cache Memory	2048 GB
Number of Ports	6 × 16 Gbps ports used
Number of RAID Groups in Dynamic Provisioning Pool	1
Number of Dynamic Provisioning Pools	2
Number of LUs in Dynamic Provisioning Pools	4/PG
RAID Group Type	RAID 10 (2D+2D)
Number of Drives per PG	4
Drive Capacity	960 GB
Drive Type	SSD
DBSC Tray	4
Multi-Pathing Enabled	Yes

The following table shows the drive slot assignment for parity groups on testbed (2+1 scale-out test).

DB 00&01	0	1	2	3	….	9	10	11	12	13	14	….	19	20	21	22	23
DB 02&03	0	1	2	3	….	9	10	11	12	13	14	….	19	20	21	22	23
DB 04&05	0	1	2	3	….	9	10	11	12	13	14	….	19	20	21	22	23

DB 06&07	0	1	2	3	….	9	10	11	12	13	14	….	19	20	21	22	23
Parity Group, RAID 10 (2D+2D)	PG-1								PG-2								

The following table shows dynamic provisioning pool configuration for the 2+1 scale-out test.

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size	MPU Assign
OS_SH_DT_Pool	1	RAID 10 (2D+2D) on 960 GB SSDs	00:00:01	OS_SH_DT_1	440 GB	MPU-010
			00:00:02	OS_SH_DT_2	440 GB	MPU-020
			00:00:03	OS_SH_DT_3	440 GB	MPU-110
			00:00:04	OS_SH_DT_4	440 GB	MPU-120
LG_Pool	2	RAID 10 (2D+2D) on 960 GB SSDs	00:00:05	LG_1	440 GB	MPU-010
			00:00:06	LG_2	440 GB	MPU-020
			00:00:07	LG_3	440 GB	MPU-110
			00:00:08	LG_4	440 GB	MPU-120

The following table shows the drive slot assignment for parity groups on testbed (2+1 scale-out test).

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assign
OS_SH_DT_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-010
	00:02:00	HANA_OS_N2	100 GB	MPU-020
	00:03:00	HANA_OS_N3	100 GB	MPU-110
	00:01:01	HANA_SH_1	128 GB	MPU-010
	00:02:01	HANA_SH_2	128 GB	MPU-020
	00:03:01	HANA_SH_3	128 GB	MPU-110
	00:01:06	HANA_DATA_N1_1	64 GB	MPU-010

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assign
	00:01:07	HANA_DATA_N1_2	64 GB	MPU-020
	00:01:08	HANA_DATA_N1_3	64 GB	MPU-110
	00:01:09	HANA_DATA_N1_4	64 GB	MPU-120
	00:02:06	HANA_DATA_N2_1	64 GB	MPU-010
	00:02:07	HANA_DATA_N2_2	64 GB	MPU-020
	00:02:08	HANA_DATA_N2_3	64 GB	MPU-110
	00:02:09	HANA_DATA_N2_4	64 GB	MPU-120
	00:03:06	HANA_DATA_N3_1	64 GB	MPU-010
	00:03:07	HANA_DATA_N3_2	64 GB	MPU-020
	00:03:08	HANA_DATA_N3_3	64 GB	MPU-110
	00:03:09	HANA_DATA_N3_4	64 GB	MPU-120
LG_Pool	00:01:02	HANA_LOG_N1	128 GB	MPU-010
	00:01:03	HANA_LOG_N2	128 GB	MPU-020
	00:01:04	HANA_LOG_N3	128 GB	MPU-110

Product descriptions

These products are used in this reference architecture.

Hitachi Virtual Storage Platform 5000 series

This enterprise-class, flash array evolution storage, Hitachi Virtual Storage Platform 5000 series (VSP) has an innovative, scale-out design optimized for NVMe and storage class memory. It achieves the following:

- **Agility using NVMe:** Speed, massive scaling with no performance slowdowns, intelligent tiering, and efficiency.
- **Resilience:** Superior application availability and flash resilience. Your data is always available, mitigating business risk.
- **Storage simplified:** Do more with less, integrate AI (artificial intelligence) and ML (machine learning), simplify management, and save money and time with consolidation.

Hitachi Storage Virtualization Operating System RF

Hitachi Storage Virtualization Operating System RF powers the Hitachi Virtual Storage Platform (VSP) family. 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.

Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid arrays as storage capacity, consolidating resources for a higher ROI and providing a high-speed front end to slower, less-predictable arrays.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact