HITACHI
Inspire the Next

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform E990 and E1090 with Hitachi Storage Virtualization Operating System RF

Reference Architecture Guide

MK-SL-197-03
April 2024

Legal Notices

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-197-03	Add HANA ES VSSO	April 2024
MK-SL-197-02	Added support for Virtual Storage Platform E1090	February 2022
MK-SL-197-01	▪ Added a note to Table 1. ▪ Updated node numbers in Table 6 to 32.	December 2020

Revision	Changes	Date
	▪ Corrected the number of SAP HANA production nodes to 32 in the two bullets after Table 6. ▪ Added a note to Table 12.	

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 a SAP HANA deployment on the following Hitachi Virtual Storage Platform (VSP) models:

- VSP E990, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- VSP E1090, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP E1090 storage systems

You can implement SAP HANA as follows:

- With a SAP HANA appliance deployment, the hardware vendor pre-configures all hardware components—storage and server.
- With Hitachi systems for SAP HANA tailored data center (TDI) deployment, you customize each installation by using any supported servers and certified enterprise storage vendor. You may be able to leverage your existing hardware to reduce your total cost of ownership (TCO) for a high-performance SAP HANA infrastructure.

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 for this Hitachi system for SAP HANA TDI has the following notes:

- Server from Hitachi were used. This solution is also valid for other Hitachi-supported servers.
- Testing showed that Hitachi Virtual Storage Platform E990 and E1090 with NVMe SSDs meet the SAP enterprise storage certification key performance indicator (KPI) requirements for SAP HANA:
- Scalability and storage KPI testing was performed using 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 following table lists Hitachi supported maximum number of SAP HANA systems connected in parallel to a single Virtual Storage Platform to pass the SAP enterprise storage KPIs. While running HCMT in the SAP HANA systems during testing, there are 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 and Cloud Measurement Tool (HCMT) to meet SAP KPIs.

Table 1 Maximum supported number of SAP HANA systems

Storage	Drive Type	Maximum HANA Systems	Minimal Parity Group	
			Data	Log
VSP E990	NVMe SSDs	50*	7 × RAID 10 (2D+2D)	7 × RAID 10 (2D+2D)
VSP E1090	NVMe SSDs	50*	7 × RAID 10 (2D+2D)	7 × RAID 10 (2D+2D)
VSSO for 10 × VSP E1090	NVMe SSDs	500	70 × RAID 10 (2D+2D)	70 × RAID 10 (2D+2D)
*Extrapolated scalability for maximum number of nodes by using the data and information of the engineering validation performed for VSP E990 with HCMT version 047.				

When implementing a SAP HANA TDI infrastructure on Hitachi Virtual Storage Platform using NVMe SSDs, you do not have to use exactly the 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: Since the release of SAP HANA TDI in November 2013, several versions of HCMT have been published. To verify that the hardware configuration of your SAP HANA TDI infrastructure meets the SAP KPIs, you must use the same version of HCMT used during the certification of the hardware, compute servers, and storage system for your tests. SAP Note 2493172 describes how to determine the right version of HCMT 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 reference architecture guide provides example configurations tested in the Hitachi Vantara lab using NVMe SSDs of the storage layout for SAP HANA nodes with variable sizes of main memory consolidated on the following for storage:

- Hitachi Virtual Storage Platform E990 using NVMe SSDs
- Hitachi Virtual Storage Platform E1090 using NVMe SSDs
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP E1090 storage systems

The following table lists Virtual Storage Platform E990 and E1090 features.

Table 2 Hitachi Virtual Storage Platform E990 and E1090 features

Item			Specifications
	Type		VSP E990 and VSP E1090
System	SVOS version		Hitachi Storage Virtualization Operating System v9.2 (for E990), v9.8 (for E1090)
	Max. drives	NVMe SSD	96
	Drive Options	NVMe SSD	1.9 TB, 3.8TB, 7.6 TB, 15.3 TB
	DBN (Drive Box for NVMe)		4
	RAID Level		RAID 10, RAID 5, RAID 6
	RAID group configuration	RAID 10	2D+2D
		RAID 5	3D+1P, 4D+1P, 6D+1P, 7D+1P
		RAID 6	6D+2P, 12D+2P, 14D+2P
	Maximum number of LDEVs		65280
	Maximum storage capacity		1.3 PB
	Maximum external configuration		255 PB
Memory	Cache memory capacity		1024 GB
Device I/F	Supported channel type	Fibre Channel	32
		Data transfer rate	8 Gbps, 16 Gbps, 32 Gbps
	Max. number of CHB		8
Non-stop maintenance	Control PCB		Supported
	Cache memory		Supported

Item		Specifications
	Type	**VSP E990 and VSP E1090**
	Cache flash memory	Supported
	Power supply, fan	Supported
	Microcode	Supported
	Flash drive	Supported
	Flash module drive	Supported

The following figure shows a server-to-storage Fibre Channel switch indirect connection configuration for this solution. This example uses a single Virtual Storage Platform E990 with 32 SAP HANA systems and NVMe SSDs (you can replace VSP E990 with VSP E1090).

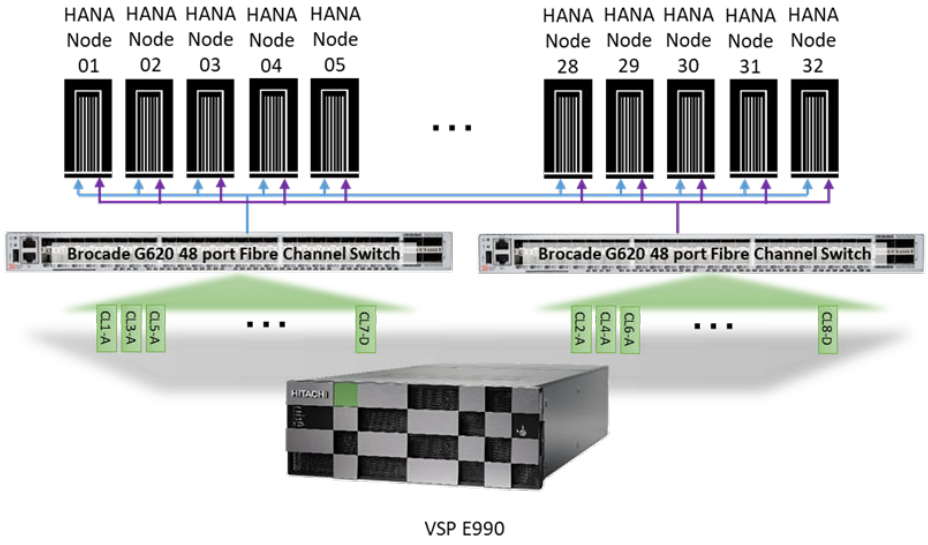

VSP E990

The following figure shows an architecture of a single VSSO storage image with clustered VSP E1090 storage systems to run multiple SAP HANA instances.

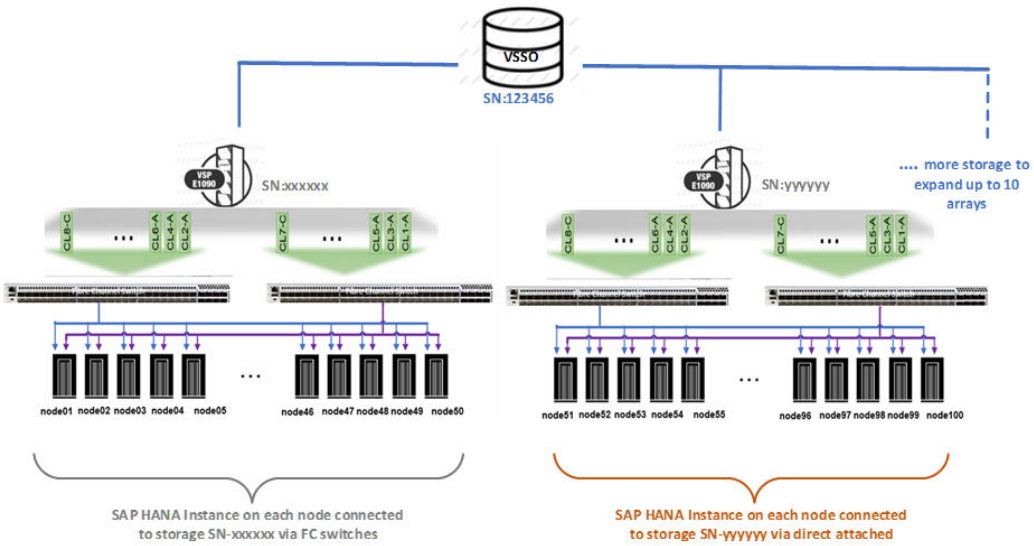

This loosely clustered VSSO provides a flexible and scalable storage solution for SAP HANA Enterprise Storage with the following:

- Achieve greater HANA node scalability by adding up the total number of HANA nodes for multiple VSP storages systems.
- No performance impact when adding more storage in a loosely clustered VSSO.
- Manage all storage from a single management tool - Hitachi OpsCenter.
- Physical connections between SAP HANA servers and storage systems are the same as non-cluster storage for SAN (via Fibre Channel switches) and direct attached storage (DAS).

Key solution elements

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

Hardware elements

The following table lists the hardware used to test the scalability of maximum active nodes on Hitachi Virtual Storage Platform E990 using NVMe SSDs.

Table 3 Hardware elements

	Tested Number of SAP HANA Nodes	Storage Components	Fibre Channel Switch
VSP E990 (NVMe SSD)	32	▪ CTL: 1 pair ▪ 32 Gbps 4-port CHB: 4 pairs	48-port Brocade G620 Fibre Channel Switch: 1 pair

	Tested Number of SAP HANA Nodes	Storage Components	Fibre Channel Switch
		▪ MPU: 1 pair ▪ Cache: 1024 GB ▪ NVMe SSDs: 48 × 1.9 TB ▪ DKBN: 8 ▪ DBN: 4	

Note: View a list of certified configurations in the Certified and Supported SAP HANA Hardware Directory.

The following table lists the optional hardware that can be used with any of these storage systems in your environment.

Table 4 Optional hardware for use with any storage system

Optional Item	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
	Or Cisco Nexus 92348 or Cisco Nexus 3048
10 GbE Network Switch, using either one of these switches	Brocade VDX 6740-48 port
	Or Cisco 93180YC-FX/EX

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 for SAP HANA.
- SAP HANA
- Hitachi Storage Virtualization Operating System
  - This includes Hitachi Device Manager - Storage Navigator and Hitachi Dynamic Provisioning.
- Hitachi Storage Advisor

Solution design

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

- Hitachi Virtual Storage Platform E990 using NVMe SSDs
- Hitachi Virtual Storage Platform E1090 using NVMe SSDs
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP E1090 storage systems

Fibre Channel architecture

For each SAP HANA node, there are two Fibre Channel ports on a VSP E990 or VSP E1090 storage system. The two Fibre Channel cables connect HBA ports on the node to the designated storage target ports to achieve the following:

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

The following table lists storage port mapping for Fibre Channel switches with hosts for the validated number of SAP HANA systems using NVMe SSDs.

Table 5 Storage port mapping for validated SAP HANA nodes using NVMe SSDs

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port		
	Port Name	Port Speed	Host	Storage	Host Group Name	Port Speed	Port Security
Node1	Port 0	16 Gbps	SW-1-P0	SW-1-P32	1A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P0	SW-2-P32	2A-Host Group 1	32 Gbps	Enabled
Node2	Port 0	16 Gbps	SW-1-P1	SW-1-P32	1A-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P1	SW-2-P32	2A-Host Group 2	32 Gbps	Enabled
Node3	Port 0	16 Gbps	SW-1-P2	SW-1-P33	3A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P2	SW-2-P33	4A-Host Group 1	32 Gbps	Enabled
Node4	Port 0	16 Gbps	SW-1-P3	SW-1-P33	3A-Host Group 2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port		
	Port Name	**Port Speed**	**Host**	**Storage**	**Host Group Name**	**Port Speed**	**Port Security**
	Port 1	16 Gbps	SW-2-P3	SW-2-P33	4A-Host Group 2	32 Gbps	Enabled
Node5	Port 0	16 Gbps	SW-1-P4	SW-1-P34	5A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P4	SW-2-P34	6A-Host Group 1	32 Gbps	Enabled
Node6	Port 0	16 Gbps	SW-1-P5	SW-1-P34	5A-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P5	SW-2-P34	6A-Host Group 2	32 Gbps	Enabled
Node7	Port 0	16 Gbps	SW-1-P6	SW-1-P35	7A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P6	SW-2-P35	8A-Host Group 1	32 Gbps	Enabled
Node8	Port 0	16 Gbps	SW-1-P7	SW-1-P35	7A-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P7	SW-2-P35	8A-Host Group 2	32 Gbps	Enabled
Node9	Port 0	16 Gbps	SW-1-P8	SW-1-P36	1B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P8	SW-2-P36	2B-Host Group 1	32 Gbps	Enabled
Node10	Port 0	16 Gbps	SW-1-P9	SW-1-P36	1B-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P9	SW-2-P36	2B-Host Group 2	32 Gbps	Enabled
Node11	Port 0	16 Gbps	SW-1-P10	SW-1-P37	3B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P10	SW-2-P37	4B-Host Group 1	32 Gbps	Enabled
Node12	Port 0	16 Gbps	SW-1-P11	SW-1-P37	3B-Host Group 2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port		
	Port Name	**Port Speed**	**Host**	**Storage**	**Host Group Name**	**Port Speed**	**Port Security**
	Port 1	16 Gbps	SW-2-P11	SW-2-P37	4B-Host Group 2	32 Gbps	Enabled
Node13	Port 0	16 Gbps	SW-1-P12	SW-1-P38	5B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P12	SW-2-P38	6B-Host Group 1	32 Gbps	Enabled
Node14	Port 0	16 Gbps	SW-1-P13	SW-1-P38	5B-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P13	SW-2-P38	6B-Host Group 2	32 Gbps	Enabled
Node15	Port 0	16 Gbps	SW-1-P14	SW-1-P39	7B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P14	SW-2-P39	8B-Host Group 1	32 Gbps	Enabled
Node16	Port 0	16 Gbps	SW-1-P15	SW-1-P39	7B-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P15	SW-2-P39	8B-Host Group 2	32 Gbps	Enabled
Node17	Port 0	16 Gbps	SW-1-P16	SW-1-P40	1C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P16	SW-2-P40	2C-Host Group 1	32 Gbps	Enabled
Node18	Port 0	16 Gbps	SW-1-P17	SW-1-P40	1C-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P17	SW-2-P40	2C-Host Group 2	32 Gbps	Enabled
Node19	Port 0	16 Gbps	SW-1-P18	SW-1-P41	3C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P18	SW-2-P41	4C-Host Group 1	32 Gbps	Enabled
Node20	Port 0	16 Gbps	SW-1-P19	SW-1-P41	3C-Host Group 2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port		
	Port Name	**Port Speed**	**Host**	**Storage**	**Host Group Name**	**Port Speed**	**Port Security**
	Port 1	16 Gbps	SW-2-P19	SW-2-P41	4C-Host Group 2	32 Gbps	Enabled
Node21	Port 0	16 Gbps	SW-1-P20	SW-1-P42	5C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P20	SW-2-P42	6C-Host Group 1	32 Gbps	Enabled
Node22	Port 0	16 Gbps	SW-1-P21	SW-1-P42	5C-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P21	SW-2-P42	6C-Host Group 2	32 Gbps	Enabled
Node23	Port 0	16 Gbps	SW-1-P22	SW-1-P43	7C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P22	SW-2-P43	8C-Host Group 1	32 Gbps	Enabled
Node24	Port 0	16 Gbps	SW-1-P23	SW-1-P43	7C-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P23	SW-2-P43	8C-Host Group 2	32 Gbps	Enabled
Node25	Port 0	16 Gbps	SW-1-P24	SW-1-P44	1D-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P24	SW-2-P44	2D-Host Group 1	32 Gbps	Enabled
Node26	Port 0	16 Gbps	SW-1-P25	SW-1-P44	1D-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P25	SW-2-P44	2D-Host Group 2	32 Gbps	Enabled
Node27	Port 0	16 Gbps	SW-1-P26	SW-1-P45	3D-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P26	SW-2-P45	4D-Host Group 1	32 Gbps	Enabled
Node28	Port 0	16 Gbps	SW-1-P27	SW-1-P45	3D-Host Group 2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port		
	Port Name	**Port Speed**	**Host**	**Storage**	**Host Group Name**	**Port Speed**	**Port Security**
	Port 1	16 Gbps	SW-2-P27	SW-2-P45	4D-Host Group 2	32 Gbps	Enabled
Node29	Port 0	16 Gbps	SW-1-P28	SW-1-P46	5D-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P28	SW-2-P46	6D-Host Group 1	32 Gbps	Enabled
Node30	Port 0	16 Gbps	SW-1-P29	SW-1-P46	5D-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P29	SW-2-P46	6D-Host Group 2	32 Gbps	Enabled
Node31	Port 0	16 Gbps	SW-1-P30	SW-1-P47	7D-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P30	SW-2-P47	8D-Host Group 1	32 Gbps	Enabled
Node32	Port 0	16 Gbps	SW-1-P31	SW-1-P47	7D-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P31	SW-2-P47	8D-Host Group 2	32 Gbps	Enabled

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 the following two 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.

- 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 Hitachi Virtual Storage Platform E990 and E1090 storage is listed in the following table.

Table 6 Dynamic Provisioning Pools with disks and storage cache

Storage	Cache	Validated Node Number	Number of Parity Groups in OS_SH_Data_Pool	Number of Parity Groups in LG_Pool
			RAID 6 (6D+2P)	RAID 6 (6D +2P)
VSP E990 (with NVMe SSD)	1024 GB	32	1	1
VSP E1090 (with NVMe SSD)	1024 GB	32	1	1

To fit 32 SAP HANA nodes, in addition to RAID 10 (2D+2D), you can use RAID 6 (6D+2P) for capacity efficiency. As an example, for Hitachi Virtual Storage Platform E990 or E1090 listed in the previous table, you need the following:

- A minimum of one parity group is needed for OS_SH_DT_ Pool to fit 32 SAP HANA production nodes on one VSP E990 or E1090 using RAID 6 (6D+2P).
- A minimum of one parity group is needed for LG_ Pool to fit 32 SAP HANA production nodes on one VSP E990 or E1090 using RAID 6 (6D+2P).

Additional parity groups of the same type might need to be added. Drive boxes may be needed if the internal drives on storage are not sufficient, depending on the following:

- The various combinations of node sizes
- The number of 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:

- 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 the following formulas:

    Single node (scale-up) — Size = MIN (1 × RAM; 1 TB)

    Multi-node (scale-out) — Size = 1 × RAM of every 4 worker nodes
  - Data size requires at least 1 × RAM of each HANA node
  - Log size uses the following formulas:

    Systems with equal or less than 512 GB memory — Size = 1/2 × RAM

    Systems with greater than 512 GB memory — Size = 512 GB

Note: Hitachi recommends these minimum sizes, as calculated using formulas. The SAP quicksizer and sizing reports might call for smaller sizes.

- Provision the storage:
  - Create two dynamic provisioning pools for the three SAP HANA systems on storage:
    - 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 NVMe SSDs, create the parity groups first, as the example listed in the following table for Hitachi Virtual Storage Platform E990 or E1090 using the RAID 6 storage design.

Table 7 Dynamic Provisioning Pool with RAID 6 (6D+2P) for 16 nodes on VSP E990 or VSP E1090

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 6 (6D +2P) on 1.9 TB SSD	00:00:01	OS_SH_DT_DPVOL_1	2640 GB	MPU-10
			00:00:02	OS_SH_DT_DPVOL_2	2640 GB	MPU-20
			00:00:03	OS_SH_DT_DPVOL_3	2640 GB	MPU-10

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
			00:00:04	OS_SH_DT_DPVOL_4	2640 GB	MPU-20
Log_Pool	4	RAID 6 (6D +2P) on 1.9 TB SSD	00:00:13	LG_DPVOL_1	2640 GB	MPU-10
			00:00:14	LG_DPVOL_2	2640 GB	MPU-20
			00:00:15	LG_DPVOL_3	2640 GB	MPU-10
			00:00:16	LG_DPVOL_4	2640 GB	MPU-20

- Assign all LDEVs to the dedicated pool as listed in the examples in the previous table for VSP E990 or VSP E1090.
- Create virtual volumes (vVols) for the operating system, SAP HANA shared, log, and data volumes. The following table lists examples for HANA systems with 384 GB, 768 GB, 1536 GB, and 3072 GB of memory.

Table 8 vVols for SAP HANA nodes for four memory sizes of SAP HANA systems

Dynamic Provisioning Pool	vVOl	vVol Name	vVol Size	MPU Assignment	System Memory
OS_SH_Data_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-10	384 GB
	00:02:00	HANA_OS_N2	100 GB	MPU-20	768 GB
	00:03:00	HANA_OS_N3	100 GB	MPU-10	1536 GB
	00:04:00	HANA_OS_N4	100 GB	MPU-20	3072 GB
	00:01:01	HANA_SH_N1	384 GB	MPU-10	384 GB
	00:02:01	HANA_SH_N2	768 GB	MPU-20	768 GB
	00:03:01	HANA_SH_N3	1536 GB	MPU-10	1536 GB
	00:04:01	HANA_SH_N4	3072 GB	MPU-20	3072 GB
	00:01:06	HANA_DATA_N1_1	96 GB	MPU-10	384 GB

Dynamic Provisioning Pool	vVOl	vVol Name	vVol Size	MPU Assignment	System Memory
	00:01:07	HANA_DATA_N1_2	96 GB	MPU-20	
	00:01:08	HANA_DATA_N1_3	96 GB	MPU-10	
	00:01:09	HANA_DATA_N1_4	96 GB	MPU-20	
	00:02:06	HANA_DATA_N2_1	192 GB	MPU-10	768 GB
	00:02:07	HANA_DATA_N2_2	192 GB	MPU-20	
	00:02:08	HANA_DATA_N2_3	192 GB	MPU-10	
	00:02:09	HANA_DATA_N2_4	192 GB	MPU-20	
	00:03:06	HANA_DATA_N3_1	384 GB	MPU-10	1536 GB
	00:03:07	HANA_DATA_N3_2	384 GB	MPU-20	
	00:03:08	HANA_DATA_N3_3	384 GB	MPU-10	
	00:03:09	HANA_DATA_N3_4	384 GB	MPU-20	
	00:04:06	HANA_DATA_N4_1	768 GB	MPU-10	3072 GB
	00:04:07	HANA_DATA_N4_2	768 GB	MPU-20	
	00:04:08	HANA_DATA_N4_3	768 GB	MPU-10	
	00:04:09	HANA_DATA_N4_4	768 GB	MPU-20	
Log_Pool	00:01:02	HANA_LOG_N1_1	48 GB	MPU-10	384 GB
	00:01:03	HANA_LOG_N1_2	48 GB	MPU-20	
	00:01:04	HANA_LOG_N1_3	48 GB	MPU-10	
	00:01:05	HANA_LOG_N1_4	48 GB	MPU-20	
	00:02:02	HANA_LOG_N2_1	96 GB	MPU-10	768 GB
	00:02:03	HANA_LOG_N2_2	96 GB	MPU-20	
	00:02:04	HANA_LOG_N2_3	96 GB	MPU-10	
	00:02:05	HANA_LOG_N2_4	96 GB	MPU-20	
	00:03:02	HANA_LOG_N3_1	128 GB	MPU-10	1536 GB
	00:03:03	HANA_LOG_N3_2	128 GB	MPU-20	
	00:03:04	HANA_LOG_N3_3	128 GB	MPU-10	
	00:03:05	HANA_LOG_N3_4	128 GB	MPU-20	

Dynamic Provisioning Pool	vVOl	vVol Name	vVol Size	MPU Assignment	System Memory
	00:04:02	HANA_LOG_N4_1	128 GB	MPU-10	3072 GB
	00:04:03	HANA_LOG_N4_2	128 GB	MPU-20	
	00:04:04	HANA_LOG_N4_3	128 GB	MPU-10	
	00:04:05	HANA_LOG_N4_4	128 GB	MPU-20	

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

The following table lists an example configuration of the LUN path assignment for SAP HANA Node 1. Configure the LUN assignment similarly for all other nodes.

Table 9 Example LUN path assignment for the SAP HANA configuration on Node 1

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

Virtual Storage Scale-out configuration

Virtual Storage Scale-out (VSSO) uses Hitachi Virtual Storage Machine (VSM) as a software-defined storage to create a virtualized storage infrastructure for multiple VSP E1090 storage systems. This does not change physical connections between SAP HANA server and each storage system (either via DAS or SAN).

The following gives guidance to implement VSSO:

- Make physical connections between each SAP HANA server and a storage system (via DAS or SAN), the same as you would non-cluster storage.
- Storage provisioning:
  - Infrastructure level provisioning (creating RAID groups, Dynamic Provisioning Pools, DP-VOL, etc) are configured from each Hitachi VSP E1090 storage system.
  - VSSO with multiple VSP E1090 storage systems is set up using Hitachi Ops Center (a centralized management interface for Hitachi storage environments to manage multiple storage systems from a single console).
  - VSSO is configured using a new single storage serial number. A new storage resource group is generated for a VSSO to allow you to logically group together multiple storage resources (such as pools, volumes and ports) within VSSO.
  - Logical provisioning (expand VSSO with more storage, create Virtual Volumes, add host groups, assign LUN path, etc.) is configured using Hitachi Ops Center.
  - For more information see the VSSO setup video at: Introduction to Virtual Storage Scale Out – YouTube.

SAP HANA configuration

This 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 operating system 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 uses multiple paths to the disks with two options:
  - Hitachi NAS Platform as NFS for /hana/shared. This solution is also valid for the other 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 recommends the multipath settings in the following table.

Table 10 Recommended multipath settings

Section	Option	SUSE Linux Enterprise Server for SAP Applications	Red Hat Enterprise Linux for SAP HANA
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	tur	
	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	

HANA persistent storage volume configuration

For both operating systems, SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux for SAP HANA, Hitachi uses an LVM-based storage layout. When the operating system is installed and multipathing configured correctly, 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 used during the LUN assignment. In the previous example, 000101 maps to LDEV ID: 00:01:01.

For all the LUNs except the one hosting the operating system, you need to initialize the LUNs for use by LVM, running the `pvcreate` command as follows which is part of the lvm2 rpm package:

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

After you have prepared all the LUNs, you need to configure the volume groups (VGs) 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 4 physical disks/ LUNs:

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

To create additional volume groups use the same syntax, exchanging the volume group name as well as the physical disks or LUNs.

When the volume groups are created, you need to create a logical volume (LV) on top. The general syntax is the following:

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

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

Table 11 Details for creating 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 (Only applicable for scale-up SAP	1 - following this reference architecture, or	lvshared	▪ Scale-up: vgshared

	Number of LUNs	LV Name	VG Name
HANA deployments)	the number of assigned LUNs		

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.

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

Table 12 File system create and mount options

	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>
	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 (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, the 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

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

Procedure

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

2. Create the mount points using 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.
   Mount them one at a time using the following command:

   ```
   mount <mount point>
   ```

   Or mount them all at once using the following command:

   ```
   mount -a
   ```

   Verify the mounted file systems using either the **`df`** or **`mount`** command.

SAP HANA persistent storage volume configuration for scale-out deployments

This is only for scale-out systems.

To configure the SAP HANA persistent storage volume, 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 HANA shared file systems 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 mounts the file system used for LOG and DATA with the correct options during SAP HANA startup. For additional details, see the following:
   - SAP Storage Connector API Fibre Channel Client
   - SAP HANA Software Installation

SAP storage connector API Fibre Channel client

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

The Fibre Channel storage connector, fcClient/fcCLientLVM, implements SAP's 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 `/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 SAP HANA on servers. See SAP Note 2235581, the attachment SAP_HANA_OS_Release_Support_Matrix.pdf file outlines which combinations of OS RHEL Minor Releases or SLES Support Packages 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.

SAP HANA software installation 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> \
--sapadm_password=<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 SAP HANA Software Configuration.

SAP HANA software installation in 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 content similar to the following (assuming HIT is your SAP System ID).

Create this configuration file before following the procedure:

```
[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__lvmname = vgHITlog001-lvlog
partition_1_data__lvmname = vgHITdata001-lvdata
partition_2_log__lvmname = vgHITlog002-lvlog
partition_2_data__lvmname = vgHITdata002-lvdata
partition_3_log__lvmname = vgHITlog003-lvlog
partition_3_data__lvmname = vgHITdata003-lvdata

[trace]
ha_fcclient = info
```

See the 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. It can be removed 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=/var/tmp/hana_install_<number>
--hostname=hana001ic \
--certificates_hostmap=hana001ic=hana001
```

3. During the installation you need to provide various passwords.
4. Once the installation is complete, continue with SAP HANA software configuration (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 \
--system_user_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 SAP HANA failover. Use the following command:

```
/hana/shared/HIT/hdblcm/hdblcm –action=add_hosts \
--sid=HIT \
--install_hostagent \
--sapmnt=/hana/shared \
--addhosts=hana003ic:role=standby:group=default
```

SAP HANA software configuration

The following are examples of *global.ini* files used during engineering validation:

- SAP HANA 2.0 *global.ini* file for VSP E1090 with NVMe SSDs:

```
[Communication]
Tcp_backlog = 2048
Listeninterface = .global

[fileio]
max_parallel_io_requests[data] = 1024
max_submit_batch_size[data] = 1024
size_kernel_io_queue[data] = 1024
async_read_submit[data] = on
async_write_submit_blocks[data] = all
min_submit_batch_size[data] = 16
async_write_submit_active[data] = on
max_parallel_io_requests[log] = 128
max_submit_batch_size[log] = 128
size_kernel_io_queue[log] = 1024
async_read_submit[log] = on
async_write_submit_blocks[log] = all
min_submit_batch_size[log] = 16
async_write_submit_active[log] = on

[multidb]
Mode = multidb
Database_isolation = low
Singletenant = yes

[persistence]
basepath_datavolumes = /hana/data/HIQ
basepath_logvolumes = /hana/log/HIQ

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*__prtype = 5
partition_1_data__lvmname = vgHIQdata001-lvdata
partition_1_log__lvmname = vgHIQlog001-lvlog
partition_2_data__lvmname = vgHIQdata002-lvdata
partition_2_log__lvmname = vgHIQlog002-lvlog
partition_3_data__lvmname = vgHIQdata003-lvdata
partition_3_log__lvmname = vgHIQlog003-lvlog

[trace]
ha_fcclientlvm = info
```

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

- SAP HANA 2.0 *global.ini* file for VSP E990:

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

[fileio]
max_parallel_io_requests[data] = 512
max_submit_batch_size[data] = 384
size_kernel_io_queue[data] = 512
async_read_submit[data] = on
async_write_submit_blocks[data] = all
min_submit_batch_size[data] = 16
async_write_submit_active[data] = on
max_parallel_io_requests[log] = 512
max_submit_batch_size[log] = 384
size_kernel_io_queue[log] = 512
async_read_submit[log] = on
async_write_submit_blocks[log] = all
min_submit_batch_size[log] = 16
async_write_submit_active[log] = on

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

[persistence]
basepath_datavolumes = /hana/data/HIQ
basepath_logvolumes = /hana/log/HIQ

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*__prtype = 5
partition_1_data__lvmname = vgHIQdata001-lvdata
partition_1_log__lvmname = vgHIQlog001-lvlog
partition_2_data__lvmname = vgHIQdata002-lvdata
partition_2_log__lvmname = vgHIQlog002-lvlog
partition_3_data__lvmname = vgHIQdata003-lvdata
partition_3_log__lvmname = vgHIQlog003-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 SAP HANA scale-out system:

- SAP HANA hardware and cloud measurement tool (HCMT) for testing for the enterprise storage certification, revision hcmt-057.

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

Table 13 Hardware elements in the engineering validation test

Hardware	Quantity	Configuration	Role
Hitachi Compute Blade chassis	1	▪ 6 blades on chassis ▪ 2 management modules ▪ 10 cooling fan modules ▪ 6 power supply modules ▪ 6 Hitachi 16 Gbps 2-port Fibre Channel PCIe cards (2 per SAP HANA node) ▪ 12 × 2-port 10GBASE-SR PCIe Cards (4 per node)	Server blade chassis
Hitachi Compute server blade	6	▪ 4 Intel Xeon CPU E7-8880 v3 @ 2.30 GHz ▪ 768 GB RAM per blade ▪ 1 pass-through mezzanine card on Mezzanine Slot 2	SAP HANA server: ▪ (2 blades per HANA node)
Management Server for Solution	1	Rack Mount server node with following components: ▪ 2 Intel Xeon E5-2620 v3 processor (6C, 2.4 GHz, 85 W) ▪ 1 heatsink CPU0 and CPU1 ▪ 2 × 16 GB DDR4, 2,133 MHz memory module ▪ 2 HDD SATA 500 GB, 7200 RPM, 2.5 inch (6 Gb) ▪ 1 dual port 1 GigE Base-T i350 Mezzanine card ▪ 1 dual port 10 GigE Intel 82599ES SFP+ OCP Mezzanine Card	Used as the management server, and runs the following: ▪ SAP HANA Studio
Hitachi Virtual Storage Platform E1090	1	▪ Controller board (CTL): 1 pair ▪ Disk board for NVMe (DKBN): 8 (4/CTL)	

Hardware	Quantity	Configuration	Role
		▪ CPU: 4 (1 pair/CTL) ▪ Cache flash memory (CFM): 4 (2/CTL) ▪ Backup module and fan (BKMF): 8 (4 pair/CTL) ▪ Battery: 4 ▪ MPU: 1 pair ▪ Maximum cache: 1024 GB (512 GB/CTL) ▪ Used 32 Gbps 4-port channel board (CHB): 1 pair ▪ Used Fibre Channel ports: 6 (six ports used by three SAP HANA nodes (2/node); additional ports needed if using HNAS for /hana/shared file system) ▪ NVMe SSD drives: 9 × 1.9 TB (one spare drive) ▪ Drive box (DBN): 4	
Hitachi NAS Platform 4060 file system module	2	For every NAS Platform server: ▪ 2 cluster ports ▪ 2 × 10 GbE ports ▪ 2 × 8 Gbps Fibre Channel ports ▪ 2 Ethernet ports ▪ 1 Virtual system management unit	Provides an NFS shared file system for SAP HANA binaries and cluster-wide configuration files.
Brocade VDX 6740-48 port switch	4	▪ Two switches with distinct VLANs, each dedicated to NFS and SAP HANA inter-cluster network ▪ Two switches with one VLAN to provide uplink network to customer network infrastructure	10 GbE NFS and inter-cluster network 10 GbE client network
Brocade ICX 6430-48 port switch	1	▪ 1 GbE ▪ 48 ports	1 GbE management network

Hardware	Quantity	Configuration	Role
Brocade ICX6430-24 port switch	1	▪ 1 GbE ▪ 24 ports	Hitachi NAS Platform 1 GbE private network

- The following table and figure show the server components used during the 2+1 scale-out test.

Table 14 Server components in the engineering validation

Feature	Three HANA Nodes (2+1) Scale-out Configuration
Server blades	6 Hitachi Compute Server blades
Blade location	▪ Blade 1 (primary for HANA Node 1) ▪ Blade 5 (primary for HANA Node 2) ▪ Blade 9 (primary for HANA Node 3)
Network ports	4 × 2-port 10GBASE-SR LAN PCIe adapter on four I/O board modules for each HANA node in the following locations: ▪ HANA Node 1: IOBD 01B and 04B; IOBD 02A and 03A ▪ HANA Node 2: IOBD 05B and 08B; IOBD 06A and 07A ▪ HANA Node 3: IOBD 09B and 12B; IOBD 10A and 11A
Fibre Channel ports	2 Hitachi 16 Gbps, 2-port Fibre Channel PCIe adapters on two I/O board modules for each HANA node in the following locations: ▪ HANA Node 1: IOBD 01A and 04A ▪ HANA Node 2: IOBD 05A and 08A ▪ HANA Node 3: IOBD 09A and 12A
Other interfaces	For all sizes: ▪ 1 USB 3.0 port ▪ KVM connector (VGA, COM, USB 2.0 port)

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 networks:

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

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

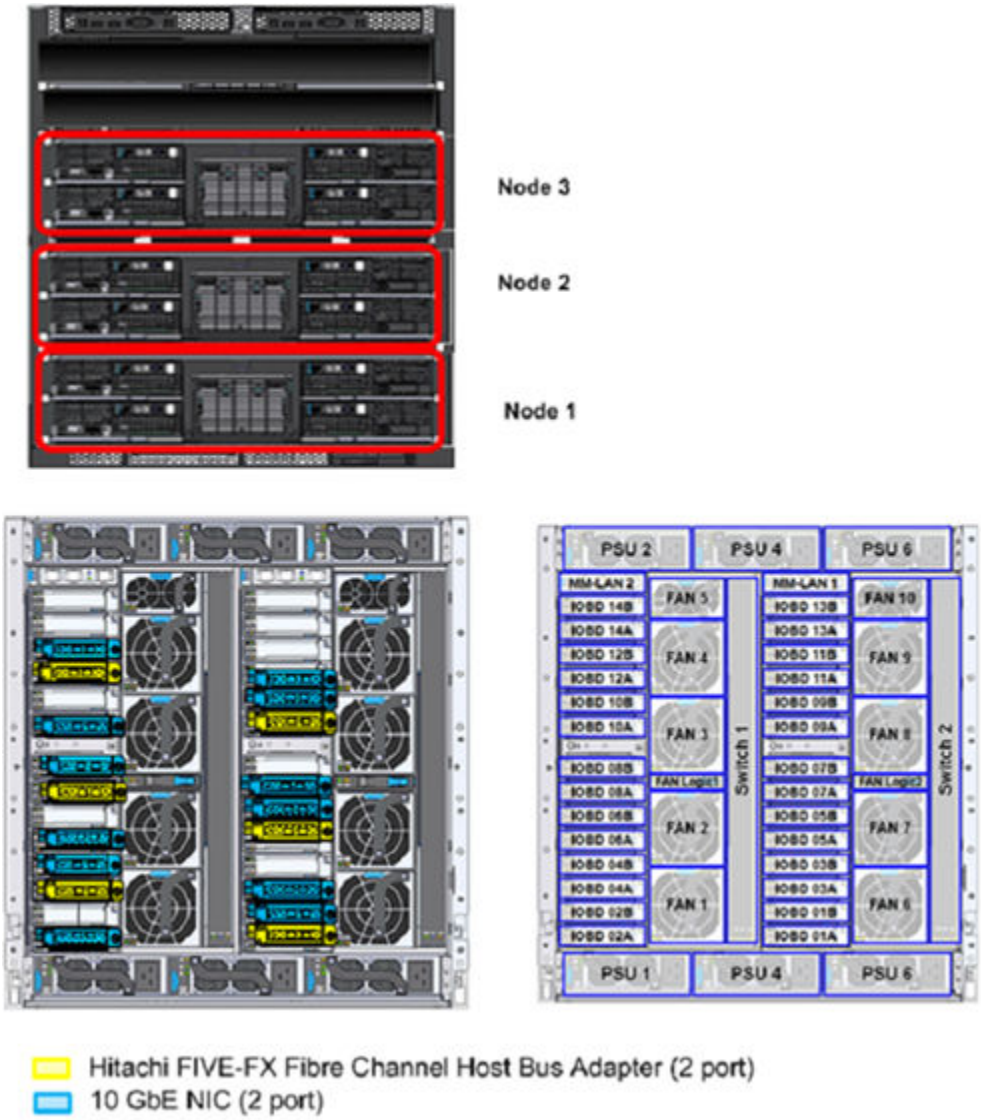

- The following four tables show the storage configuration and provisioning used for the 2+1 scale-out test.

Table 15 Storage configuration for engineering validation

Dynamic Provisioning Pool Name	Purpose
Microcode Level	93-06-00-80/13
Cache Memory	1024 GB
Number of Ports	6 × 32 Gbps ports used
Number of RAID Groups in Dynamic Provisioning Pool	1
Number of Dynamic Provisioning Pools	2

Dynamic Provisioning Pool Name	Purpose
Number of LUs in Dynamic Provisioning Pools	4/PG (Table 17)
Size of each Virtual LU	Table 18
RAID Group Type	RAID 10 (2D+2D)
Number of Drives per PG	4
Drive Capacity	1.9 TB
Drive Type	NVMe SSD
DBN Tray	4
Multi-Pathing Enabled	Yes

Table 16 Drive Box (DBN) assignment for parity groups on testbed (2+1 scale-out test)

DB 00	0	1	2	3	4	5	….	11	12	13	14	….	19	20	21	22	23
DB 01	0	1	2	3	4	5	….	11	12	13	14	….	19	20	21	22	23
DB 02	0	1	2	3	4	5	….	11	12	13	14	….	19	20	21	22	23
DB 03	0	1	2	3	4	5	….	11	12	13	14	….	19	20	21	22	23
Parity Group, RAID 10 (2D +2D)						PG-1		PG-2									

The following table shows the Dynamic Provisioning pool configuration during the 2+1 scale-out test.

Table 17 Dynamic Provisioning Pool configuration (2+1 scale-out test)

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 10 (2D+2D) on 1.9 TB SSDs	00:00:01	OS_SH_DT_1	880 GB	MPU-10
			00:00:02	OS_SH_DT_2	880 GB	MPU-20

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	880 GB	MPU-10
			00:00:04	OS_SH_DT_4	880 GB	MPU-20
LG_Pool	2	RAID 10 (2D+2D) on 1.9 TB SSDs	00:00:05	LG_1	880 GB	MPU-10
			00:00:06	LG_2	880 GB	MPU-20
			00:00:07	LG_3	880 GB	MPU-10
			00:00:08	LG_4	880 GB	MPU-20

The following table shows the vVol assignments during the 2+1 scale-out test.

Table 18 Virtual Volume assignments (2+1 scale-out test)

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment
OS_SH_DT_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-10
	00:02:00	HANA_OS_N2	100 GB	MPU-20
	00:03:00	HANA_OS_N3	100 GB	MPU-10
	00:01:01	HANA_SH_1	384 GB	MPU-10
	00:02:01	HANA_SH_2	384 GB	MPU-20
	00:03:01	HANA_SH_3	384 GB	MPU-10
	00:03:01	HANA_SH_4	384 GB	MPU-20
	00:01:06	HANA_DATA_N1_1	64 GB	MPU-10
	00:01:07	HANA_DATA_N1_2	64 GB	MPU-20
	00:01:08	HANA_DATA_N1_3	64 GB	MPU-10
	00:01:09	HANA_DATA_N1_4	64 GB	MPU-20
	00:02:06	HANA_DATA_N2_1	64 GB	MPU-10
	00:02:07	HANA_DATA_N2_2	64 GB	MPU-20

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment
	00:02:08	HANA_DATA_N2_3	64 GB	MPU-10
	00:02:09	HANA_DATA_N2_4	64 GB	MPU-20
	00:03:06	HANA_DATA_N3_1	64 GB	MPU-10
	00:03:07	HANA_DATA_N3_2	64 GB	MPU-20
	00:03:08	HANA_DATA_N3_3	64 GB	MPU-10
	00:03:09	HANA_DATA_N3_4	64 GB	MPU-20
LG_Pool	00:01:02	HANA_LOG_N1_1	32 GB	MPU-10
	00:01:03	HANA_LOG_N1_2	32 GB	MPU-20
	00:01:04	HANA_LOG_N1_3	32 GB	MPU-10
	00:01:05	HANA_LOG_N1_4	32 GB	MPU-20
	00:02:02	HANA_LOG_N2_1	32 GB	MPU-10
	00:02:03	HANA_LOG_N2_2	32 GB	MPU-20
	00:02:04	HANA_LOG_N2_3	32 GB	MPU-10
	00:02:05	HANA_LOG_N2_4	32 GB	MPU-20
	00:03:02	HANA_LOG_N3_1	32 GB	MPU-10
	00:03:03	HANA_LOG_N3_2	32 GB	MPU-20
	00:03:04	HANA_LOG_N3_3	32 GB	MPU-10
	00:03:05	HANA_LOG_N4_4	32 GB	MPU-20

Product Descriptions

These products are used in this reference architecture.

Hitachi Virtual Storage Platform E series family

The Hitachi Virtual Storage Platform E series family provides agile and automated storage built upon the innovative technologies found in our high-end enterprise systems. The expansion of the VSP E Series portfolio includes 2 new all NVMe flash models that deliver super charged, ultra-low latency performance for the business-critical applications that small and midsized businesses rely on.

- Improve IT agility: “Faster-to-market” for IT projects with proven high-performance infrastructure. Brings “enterprise-class” features and benefits to customers of all sizes whose business is outpacing their existing infrastructure and supports modern business processes like DevOPs.
- Financial elasticity that aligns costs to business goals, growth, and use: Customers can “have it their way” with purchase, lease, or cloud-like consumption models.
- Improved workforce efficiency: a better digital experience which boosts customer satisfaction (with both internal LOBs and end-users) and increases business productivity and profitability.
- VSSO Scale-out capability: allows expanding the storage infrastructure seamlessly with the following benefits:
  - Clustering of resources with increased availability.
  - Greater scalability with capacity expansion to meet expanding needs.
  - Data protection and failover support.
  - Array life cycling, resource migration.
  - Workload balancing, re-distribution and fencing

SAP HANA Tailored Data Center Integration

SAP increases flexibility and provides an alternative to SAP HANA Appliances with SAP HANA tailored data center integration (TDI) in currently 5 phases. This includes many kinds of virtualization technology. Understanding the possibilities and requirements of an SAP HANA TDI environment is crucial. One of these requirements is that all virtualized environments for SAP HANA are considered as SAP HANA TDI by SAP. SAP provides documentation around SAP HANA TDI environments that explain the 5 phases of SAP HANA TDI as well as hardware and software requirements for the whole stack:

- SAP Help Portal: SAP HANA Platform (frequently asked questions)
- SAP HANA Storage Requirements
- SAP HANA Network Requirements

Taking all this into account, Hitachi offers a solution for the whole SAP HANA TDI stack from hardware infrastructure to software deployment.

Servers for SAP HANA TDI environments

Only compute servers listed in the official SAP HANA Hardware Directory are supported. 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, see *SAP HANA Tailored Data Center Integration - Frequently Asked Questions* at https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html.

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