SAP HANA Tailored Data Center Integration on Hitachi VSP One Block 20 with Hitachi Storage Virtualization Operating System RF

Reference Architecture Guide

MK-SL-300-02
October 2025

Feedback

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

Revision history

Changes	Date
Update scalability information with SVOS 10.4.1.	October 2025
Add scalability certificate for VSP B26 and VSP B24, and VSSO for VSP One Block 20.	April 2025
Initial release.	July 2024

Reference Architecture Guide

Use this reference architecture guide to implement an 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. See SAP HANA Tailored Data Center Integration-Frequently Asked Questions for more information.

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 One Block 20 (VSP One Block 20) models:

- VSP One Block 24, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- VSP One Block 26, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- VSP One Block 28, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP One Block 24, VSP One Block 26, or VSP One Block 28 storage systems.

With an SAP HANA appliance deployment, the hardware vendor pre-configures all hardware components—storage and server. With Hitachi systems for SAP HANA TDI deployments, you customize each installation by using any supported servers and certified enterprise storage vendor. You might 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.

Note the following for the engineering validation of this Hitachi system for SAP HANA TDI:

- DS7040 servers from Hitachi were used. This solution is also valid for other servers on Hitachi Supported Servers for SAP on Linux that are also supported for SAP HANA.
- Testing showed that VSP One Block 24, VSP One Block 26, and VSP One Block 28 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 Tools (HCMT) - Replacement of HWCCT Tool. See SAP Note 2493172 (SAP user credentials required) for HCMT details.

The following table shows 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 storage KPI testing using HCMT, the storage was fully dedicated to the testing and not shared by other systems or applications.

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.

Storage	Drive Type	Maximum SAP HANA Systems	Minimum Dynamic Drive Protection (DDP) Data Volume and Log Volume	HCMT Version
VSP One Block 24	NVMe SSDs	11	1 × DDP RAID 6 (6D+2P)	HCMT-085
VSSO for 10 × Block 24	NVMe SSDs	110	10 × DDP RAID 6 (6D+2P)	N/A
VSP One Block 26	NVMe SSDs	17	1 × DDP RAID 6 (14D+2P)	HCMT-085
VSSO for 10 × Block 26	NVMe SSDs	170	10 × DDP RAID 6 (14D +2P)	N/A
VSP One Block 28	NVMe SSDs	32	1 × DDP RAID 6 (14D+2P)	HCMT-085
VSSO for 10 × Block 28	NVMe SSDs	320	10 × DDP RAID 6 (14D +2P)	N/A

When implementing an SAP HANA TDI infrastructure on Virtual Storage Platform using NVMe SSDs, you do not have to use the exact 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 factors 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 systems:

- VSP One Block 24, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- VSP One Block 26, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- VSP One Block 28, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP One Block storage systems.

The following table lists Hitachi Virtual Storage Platform One Block 20 features.

Item			Specifications
System	SVOS Version		Hitachi Storage Virtualization Operating System v10.4.1
	Max. Drives	NVMe SSD	72
	Drive Options	NVMe SSD	1.9TB, 3.8 TB, 7.6 TB, 15 TB, 30 TB, and 60 TB
	RAID Level		DDP RAID 6
	Drive Protection Supported		Dynamic Drive Protection (DDP)
			Dual Parity Protection Scheme (6D+2P, 14D+2P)
	Maximum number of LDEVs		16384 (VVSP One Block 24), 32768 (VSP One Block 26), 49152 (VSP One Block 28)
	Maximum storage capacity		2.1 PB
	Maximum external configuration		64 PB (VSP One Block 24), 128 PB (VSP One Block 26), 192 PB (VSP One Block 28)
	Compression Accelerator Module (CAM)		0 (VSP One Block 24), 2 (VSP One Block 26), 4 (VSP One Block 28)
Memory	Cache memory capacity		256 GB (VSP One Block 24), 768 GB (VSP One Block 26), 1024 GB (VSP One Block 28)
Device I/F	Supported channel type	Fibre Channel	32
		Data transfer rate	16 Gbps, 32 Gbps, 64 Gbps
	Maximum number of CHBs		6
Non-stop maintenance	Control PCB		Supported
	Cache memory		Supported
	Cache flash memory		Supported

Item		Specifications
	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 VSP One Block 28 with 24 SAP HANA systems and NVMe SSDs.

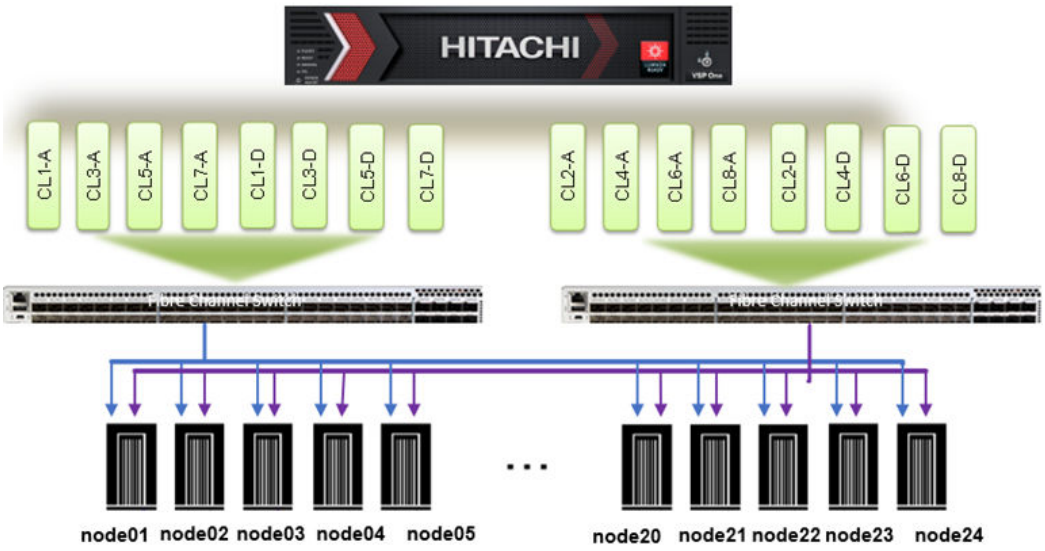

The following figure shows a server-to-storage Fibre Channel switch indirect connection configuration for this solution. This example uses VSP One Block 28 with 24 SAP HANA systems and NVMe SSDs.

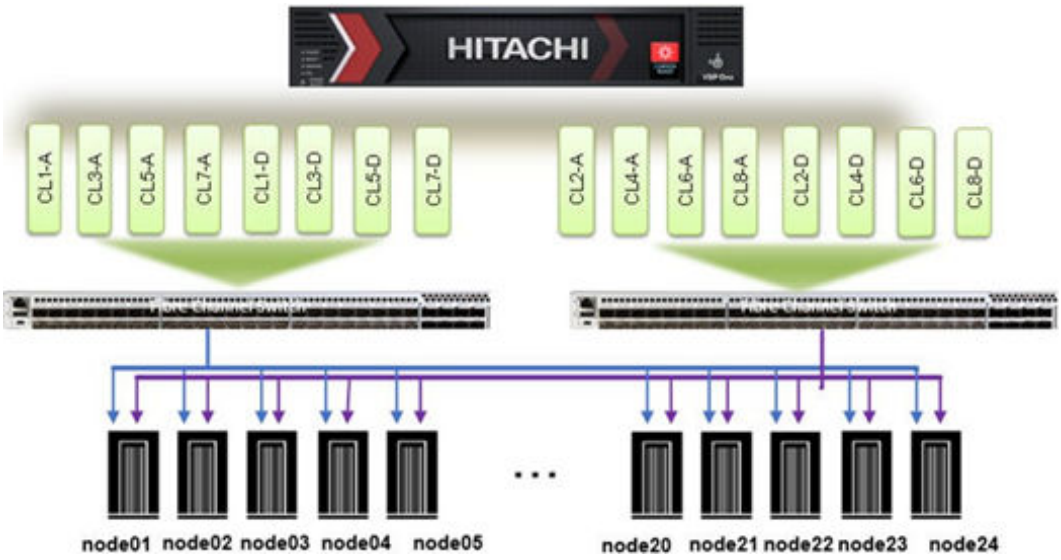

The following figure shows an architecture of a single VSSO storage image with a clustered VSP One Block 28 to run multiple SAP HANA instances. The same for clustered VSP One Block 24, VSP One Block 26, and VSP One Block 28 storage systems.

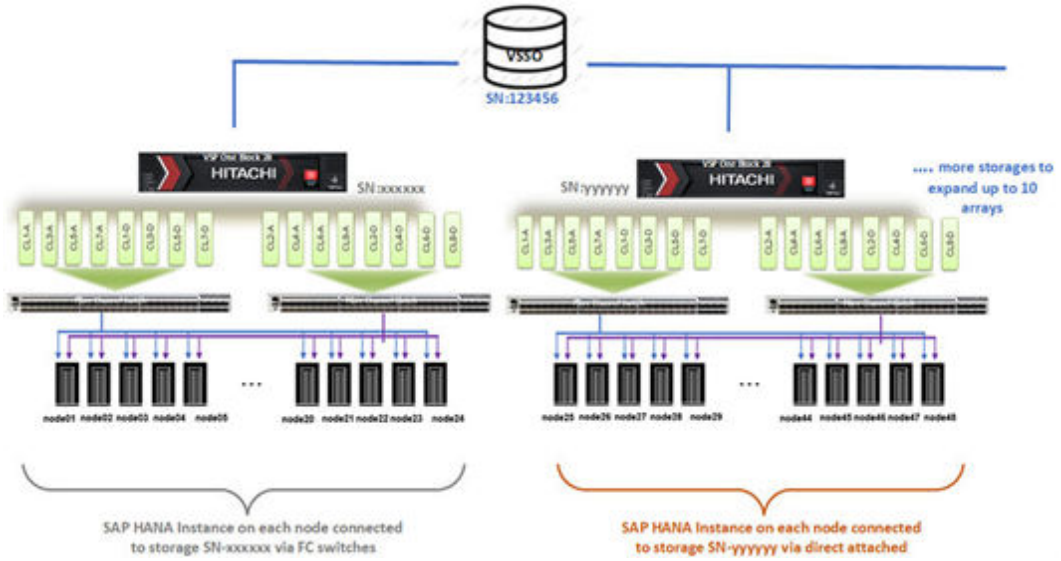

This loosely clustered VSSO provides a flexible and scalable storage solution for SAP HANA Enterprise Storage with the following:

- Achieves greater HANA node scalability by adding up to the total number of HANA nodes for multiple VSP One Block 20 storage systems.
- No performance impact when adding more storage devices in a loosely clustered VSSO.
- Manages all storage from a single management tool – Hitachi Ops Center.
- Physical connections between SAP HANA hosts and storage are the same as non-cluster storage for SAN (using Fibre Channel switches) and 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 in the Hitachi Vantara lab environment to test the maximum active nodes on VSP One Block 24, VSP One Block 26, and VSP One Block 28 using NVMe SSDs that passed SAP enterprise storage KPIs.

	Tested Number of SAP HANA Nodes	Storage Components	Fibre Channel Switch
VSP One Block 24 (NVMe SSD)	11	▪ CTL: 1 pair ▪ 32/64 Gbps 4-port CHB: 3 pairs ▪ MPU: 1 pair ▪ Cache: 256 GB ▪ NVMe SSDs: 9 × 1.92 TB	48-port Brocade G720 Fibre Channel Switch: 1 pair
VSP One Block 26 (NVMe SSD)	17	▪ CTL: 1 pair ▪ CAM: 1 / CTL ▪ 32/64 Gbps 4-port CHB: 3 pairs ▪ MPU: 1 pair ▪ Cache: 768 GB ▪ NVMe SSDs: 16 × 1.92 TB	48-port Brocade G720 Fibre Channel Switch: 1 pair
VSP One Block 28 (NVMe SSD)	32	▪ CTL: 1 pair ▪ CAM: 2 / CTL ▪ 32/64 Gbps 4-port CHB: 3 pairs ▪ MPU: 1 pair ▪ Cache: 1024 GB ▪ NVMe SSDs: 24 × 1.92 TB	48-port Brocade G720 Fibre Channel Switch: 1 pair

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

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

Optional Items	Components
Management server	Hitachi Advanced Server HA810 G3
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
	Cisco 93180YC-EX/FX/FX3

Software elements

This environment uses the following software:

- SUSE Linux Enterprise Server for SAP Applications
  - Engineering validation 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
  - Engineering validation testing used SAP HANA 2.0 SPS07.

Solution design

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

- VSP One Block 24 using NVMe SSDs
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP One Block 24 storage systems
- VSP One Block 26 using NVMe SSDs
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP One Block 26 storage systems
- VSP One Block 28 using NVMe SSDs
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP One Block 28 storage systems

The configuration uses Dynamic Drive Protection (DDP) RAID 6. A DDP group is configured with any number of devices to distribute data and spares to each drive with no spare drive required.

Fibre Channel architecture

For each SAP HANA node, there are two dedicated Fibre Channel ports on the storage. 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 shows VSP One Block 28 storage (64 Gbps) port mapping and Fibre Channel switches with hosts for the validated number of SAP HANA systems using NVMe SSDs as an example.

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP One Block 28	Port Speed	Port Security
Node1	Port 0	16 Gbps	SW-1-P0	SW-1-P32	1A-Host Group 1	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P0	SW-2-P32	2A-Host Group 1	64 Gbps	Enabled
Node2	Port 0	16 Gbps	SW-1-P1	SW-1-P32	1A-Host Group 2	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P1	SW-2-P32	2A-Host Group 2	64 Gbps	Enabled
Node3	Port 0	16 Gbps	SW-1-P2	SW-1-P32	1A-Host Group 3	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P2	SW-2-P32	2A-Host Group 3	64 Gbps	Enabled
Node4	Port 0	16 Gbps	SW-1-P3	SW-1-P32	1A-Host Group 4	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P3	SW-2-P32	2A-Host Group 4	64 Gbps	Enabled
Node5	Port 0	16 Gbps	SW-1-P4	SW-1-P33	3A-Host Group 1	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P4	SW-2-P33	4A-Host Group 1	64 Gbps	Enabled
Node6	Port 0	16 Gbps	SW-1-P5	SW-1-P33	3A-Host Group 2	64 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP One Block 28	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P5	SW-2-P33	4A-Host Group 2	64 Gbps	Enabled
Node7	Port 0	16 Gbps	SW-1-P6	SW-1-P33	3A-Host Group 3	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P6	SW-2-P33	4A-Host Group 3	64 Gbps	Enabled
Node8	Port 0	16 Gbps	SW-1-P7	SW-1-P33	3A-Host Group 4	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P7	SW-2-P33	4A-Host Group 4	64 Gbps	Enabled
Node9	Port 0	16 Gbps	SW-1-P8	SW-1-P34	5A-Host Group 1	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P8	SW-2-P34	6A-Host Group 1	64 Gbps	Enabled
Node10	Port 0	16 Gbps	SW-1-P9	SW-1-P34	5A-Host Group 2	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P9	SW-2-P34	6A-Host Group 2	64 Gbps	Enabled
Node11	Port 0	16 Gbps	SW-1-P10	SW-1-P34	5A-Host Group 3	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P10	SW-2-P34	6A-Host Group 3	64 Gbps	Enabled
Node12	Port 0	16 Gbps	SW-1-P11	SW-1-P34	5A-Host Group 4	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P11	SW-2-P34	6A-Host Group 4	64 Gbps	Enabled
Node13	Port 0	16 Gbps	SW-1-P12	SW-1-P38	7A-Host Group 1	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P12	SW-2-P38	8A-Host Group 1	64 Gbps	Enabled
Node14	Port 0	16 Gbps	SW-1-P13	SW-1-P38	7A-Host Group 2	64 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP One Block 28	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P13	SW-2-P38	8A-Host Group 2	64 Gbps	Enabled
Node15	Port 0	16 Gbps	SW-1-P14	SW-1-P39	7A-Host Group 3	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P14	SW-2-P39	8A-Host Group 3	64 Gbps	Enabled
Node16	Port 0	16 Gbps	SW-1-P15	SW-1-P39	7A-Host Group 4	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P15	SW-2-P39	8A-Host Group 4	64 Gbps	Enabled
Node17	Port 0	16 Gbps	SW-1-P16	SW-1-P40	1D-Host Group 1	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P16	SW-2-P40	2D-Host Group 1	64 Gbps	Enabled
Node18	Port 0	16 Gbps	SW-1-P17	SW-1-P40	1D-Host Group 2	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P17	SW-2-P40	2D-Host Group 2	64 Gbps	Enabled
Node19	Port 0	16 Gbps	SW-1-P18	SW-1-P41	1D-Host Group 3	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P18	SW-2-P41	2D-Host Group 3	64 Gbps	Enabled
Node20	Port 0	16 Gbps	SW-1-P19	SW-1-P41	1D-Host Group 4	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P19	SW-2-P41	2D-Host Group 4	64 Gbps	Enabled
Node21	Port 0	16 Gbps	SW-1-P20	SW-1-P42	3D-Host Group 1	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P20	SW-2-P42	4D-Host Group 1	64 Gbps	Enabled
Node22	Port 0	16 Gbps	SW-1-P21	SW-1-P42	3D-Host Group 2	64 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP One Block 28	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P21	SW-2-P42	4D-Host Group 2	64 Gbps	Enabled
Node23	Port 0	16 Gbps	SW-1-P22	SW-1-P43	3D-Host Group 3	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P22	SW-2-P43	4D-Host Group 3	64 Gbps	Enabled
Node24	Port 0	16 Gbps	SW-1-P23	SW-1-P43	3D-Host Group 4	64 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P23	SW-2-P43	4D-Host Group 4	64 Gbps	Enabled

Storage architecture

Each SAP HANA node needs the following storage layout:

- Operating system volume
- SAP HANA shared volume
- SAP HANA log volume
- SAP HANA data volume

This SAP HANA TDI setup uses the following dynamic drive protection (DDP) pool created with DDP group(s) for the storage layout. This ensures maximum utilization and optimization at a lower cost than other solutions.

One DDP pool is needed for the following:

- OS volume
- SAP HANA shared volume
- SAP HANA data volume
- SAP HANA log volume

The validated dynamic provisioning pool layout options with minimal disks and storage cache on VSP One Block 24, VSP One Block 26, and VSP One Block 28 storage are listed in the following table.

Storage	Cache	Validated Number of Nodes	Number of Drives in Pool
			DDP RAID 6
VSP One Block 24 (with NVMe SSD)	256 GB	11	9 for 1 × DDP group The group is configured as RAID 6 (6D + 2P)
VSP One Block 26 (with NVMe SSD)	768 GB	17	18 for 1 × DDP group The group is configured as RAID 6 (14D + 2P)
VSP One Block 28 (with NVMe SSD)	1024 GB	32	24 for 1 × DDP group Each group is configured as RAID 6 (14D + 2P)

In this example, for VSP One Block 28, you need the following:

- A minimum of 18 drives is needed for the DDP pools to fit 24 SAP HANA production nodes on one VSP One Block 28 using one DDP RAID 6 (14D+2P) with additional drives to support distributed spare capacity per DDP.

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

- The various combinations of node sizes
- The number of nodes needed 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 as follows:

- Create one dynamic drive protection (DDP) pool for SAP HANA systems on the storage:
  - Use the DDP pool (`Pool_for_HANA` as an example) created to provision the operating system volume, SAP HANA shared volume, data volume, and log volume.
- Refer to the example in the following table by using the DDP RAID 6 storage design.

Dynamic Drive Protection Pool	Parity Group ID	DDRAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
Pool_for_HANA	0	DDP RAID 6 (14D+2P) on 16 × 1.92 TB NVMe SSD	0	1319 GB	MPU-10
			1	1319 GB	MPU-20

Dynamic Drive Protection Pool	Parity Group ID	DDRAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
			2	1319 GB	MPU-10
			3	1319 GB	MPU-20
			4	1319 GB	MPU-10
			5	1319 GB	MPU-20
			6	1319 GB	MPU-10
			7	1319 GB	MPU-20
			8	1319 GB	MPU-10
			9	1319 GB	MPU-20
Pool_for_HANA	0	DDP RAID 6 (14D+2P) on 16 × 7.6 TB NVMe SSD	10	1319 GB	MPU-10
			11	1319 GB	MPU-20
			12	1319 GB	MPU-10
			13	1319 GB	MPU-20
			14	1319 GB	MPU-10
			15	1319 GB	MPU-20
			16	1319 GB	MPU-10
			17	1319 GB	MPU-20
Note: LDEV sizes are defined automatically by the storage.					

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

DDP Pool	LDEV ID	vVol Name	vVol Size	MPU Assignment	System Memory
Pool_for_HANA	100	HANA_OS_N1	100 GB	MPU-10	384 GB
	110	HANA_OS_N2	100 GB	MPU-20	768 GB
	120	HANA_OS_N3	100 GB	MPU-10	1536 GB
	130	HANA_OS_N4	100 GB	MPU-20	3072 GB

DDP Pool	LDEV ID	vVol Name	vVol Size	MPU Assignment	System Memory
	101	HANA_SH_N1	384 GB	MPU-10	384 GB
	111	HANA_SH_N2	768 GB	MPU-20	768 GB
	121	HANA_SH_N3	1024 GB	MPU-10	1536 GB
	131	HANA_SH_N4	1024 GB	MPU-20	3072 GB
	106	HANA_DATA_N1_1	96 GB	MPU-10	384 GB
	107	HANA_DATA_N1_2	96 GB	MPU-20	
	108	HANA_DATA_N1_3	96 GB	MPU-10	
	109	HANA_DATA_N1_4	96 GB	MPU-20	
	116	HANA_DATA_N2_1	192 GB	MPU-10	768 GB
	117	HANA_DATA_N2_2	192 GB	MPU-20	
	118	HANA_DATA_N2_3	192 GB	MPU-10	
	119	HANA_DATA_N2_4	192 GB	MPU-20	
	126	HANA_DATA_N3_1	384 GB	MPU-10	1536 GB
	127	HANA_DATA_N3_2	384 GB	MPU-20	
	128	HANA_DATA_N3_3	384 GB	MPU-10	
	129	HANA_DATA_N3_4	384 GB	MPU-20	
	136	HANA_DATA_N4_1	768 GB	MPU-10	3072 GB
	137	HANA_DATA_N4_2	768 GB	MPU-20	
	138	HANA_DATA_N4_3	768 GB	MPU-10	
	139	HANA_DATA_N4_4	768 GB	MPU-20	
Pool_for_HANA	102	HANA_LOG_N1_1	48 GB	MPU-10	384 GB
	103	HANA_LOG_N1_2	48 GB	MPU-20	
	104	HANA_LOG_N1_3	48 GB	MPU-10	
	105	HANA_LOG_N1_4	48 GB	MPU-20	
	112	HANA_LOG_N2_1	96 GB	MPU-10	768 GB
	113	HANA_LOG_N2_2	96 GB	MPU-20	
	114	HANA_LOG_N2_3	96 GB	MPU-10	

DDP Pool	LDEV ID	vVol Name	vVol Size	MPU Assignment	System Memory
	115	HANA_LOG_N2_4	96 GB	MPU-20	
	122	HANA_LOG_N3_1	128 GB	MPU-10	1536 GB
	123	HANA_LOG_N3_2	128 GB	MPU-20	
	124	HANA_LOG_N3_3	128 GB	MPU-10	
	125	HANA_LOG_N3_4	128 GB	MPU-20	
	132	HANA_LOG_N4_1	128 GB	MPU-10	3072 GB
	133	HANA_LOG_N4_2	128 GB	MPU-20	
	134	HANA_LOG_N4_3	128 GB	MPU-10	
	135	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 shows an example configuration of the LUN path assignment for Node 1. Configure the LUN assignment similarly for all other nodes.

LUN ID	LDEV ID	LDEV Name
0000	100	HANA_OS_N1
0001	101	HANA_SH_N1
0002	102	HANA_LOG_N1_1
0003	103	HANA_LOG_N1_2
0004	104	HANA_LOG_N1_3
0005	105	HANA_LOG_N1_4
0006	106	HANA_DATA_N1_1
0007	107	HANA_DATA_N1_2
0008	108	HANA_DATA_N1_3
0009	109	HANA_DATA_N1_4

Virtual Storage Scale-out (VSSO) configuration

VSSO uses Virtual Storage Machine (VSM) as a software-defined storage to create a virtualized storage infrastructure for multiple VSP One Block 20 storage systems (VSP One Block 24, VSP One Block 26, and VSP One Block 28). This doesn't change physical connections between SAP HANA host and each storage system (using either DAS or SAN).

The following provides guidance to implement VSSO:

- Make physical connections between each SAP HANA host and a storage array (via DAS or SAN), same as non-cluster storage.
- Storage provisioning:
  - Infrastructure level provisioning (creating DDP, Data Reduction Shared Volumes, server registration, and volume attachment) are configured from each storage system.
  - VSSO with multiple VSP One Block 24, or VSP One Block 26, or VSP One Block 28 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 (for example, expanding VSSO with more storage, creating volumes, adding host groups, and attaching volumes) is configured using Hitachi Ops Center.
  - See the VSSO setup video: Introduction to Virtual Storage Scale Out – YouTube for more information.

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 operating system volume (not applicable internal disks are used 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 Hitachi NAS Platform as NFS or GFS2 for `/hana/shared`. This solution is also valid for other file systems listed in SAP Note 405827 for SAP HANA.

Note: GFS2 for HANA shared filesystem is supported only if using Hitachi Virtual Storage Platform storage models and Hitachi servers in SAP HANA TDI Scale-out deployments.

Hitachi recommends the multipath settings listed in the following table.

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-length0	
	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 for SAP HANA, Hitachi uses an LVM-based storage layout. When the operating system is installed and multipathing is 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, which is part of the `lvm2 rpm` package, for example:

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

After you have prepared all the LUNs, 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, create a logical volume (LV) on top. The general syntax is the following:

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

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

	Number of LUNs	Stripe size*	LV Name	VG Name
DATA	4 - following this reference architecture, or the number of assigned LUNs	256 KB for VSP One Block 24 and Block 28 1024 KB for VSP One Block 26	lvdata	Scale-up: vgdata Scale-out: vg<SID>data<node number>
LOG	4 - following this reference architecture, or the number of assigned LUNs	512 KB for VSP One Block 28 and VSP One Block 26 256 KB for VSP One Block 24	lvlog	Scale-up: vglog Scale-out: vg<SID>log<node number>
SHARED (only applicable for scale-up SAP HANA deployments)	1 - following this reference architecture, or the number of assigned LUNs	N/A	lvshared	Scale-up: vgshared
Note: Recommended stripesize for maximum HANA node scalability using VSP One Block 20.				

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, and GFS2 as second option for `/hana/shared` scale-out solutions. The following table lists the options to create and mount the file system.

	System Type	Create Options	Mount Options	Mount Point
DATA	Scale-up	-t xfs -f	noatime, inode64	/hana/ data
	Scale-out		noatime, inode64	/hana/ data/ <SID>
LOG	Scale-up	-t xfs -f	noatime, inode64	/ hana/log
	Scale-out		noatime, inode64	/ hana/log/<SID>
SHARED	Scale-up	-t xfs -f	inode64	/hana/ shared

	System Type	Create Options	Mount Options	Mount Point
	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 VSP One File 30.)	/hana/ shared/ <SID>
Note:* For GFS2 as a second option, use `mkfs.gfs2` to create file system and `pacemaker` to mount it for `/HANA/shared/<SID>`.				

To create an xfs file system, use the following command:

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

For example:

```
mkfs -t 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 the operating system's startup file `/etc/fstab/` to mount the file systems automatically during boot operations.

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

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

2. Create the mount points using the following command:

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

   For example:

   ```
   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>
   ```

Mount them all at the same time using the following command:

```
mount -a
```

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

SAP HANA persistent storage volume configuration for scale-out deployments

This section is only for scale-out systems.

To configure the SAP HANA persistent storage volume, do the following.

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 VSP One File 30 or GFS2.
   - Hitachi VSP One File 30:
     a. Persist the SAP HANA shared file systems by adding the following entry to `/etc/fstab`:

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

     b. Mount the file system:

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

   - Global File System 2 (GFS2):

     Using GFS2 for HANA shared is for Hitachi-only TDI scale-out deployments where Hitachi provides servers and storage for the TDI landscape.

     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 cluster logical volume manager manages the shared file system 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 (on page 23)
- SAP HANA software installation (on page 24)

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 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 `/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 SAP HANA shared volume, install SAP HANA on the servers. See SAP Note 2235581 where the attachment file *SAP_HANA_OS_Release_Support_Matrix.pdf* 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 in a scale-up environment

Follow this configuration to complete your installation for a scale-up environment.

Before you begin

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

Procedure

1. Download a supported version of SAP HANA and follow the instructions to extract the archive.
2. Install SAP HANA using 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, provide various passwords.
4. After the installation is complete, continue with SAP HANA software configuration.

SAP HANA software installation in a scale-out environment

Before you begin

Before you can start with the installation in a scale-out environment, prepare a configuration file to be used by the SAP Storage Connector API for 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 a configuration file (see the following example as a reference):

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

[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
partition_3_log lvmname = vgHITlog003-lvlog
partition_3_data lvmname = vgHITdata003-lvdata

[trace] ha_fcclientlvm = 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 SAP 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.

Procedure

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, provide various passwords.
4. After the installation is complete, continue with SAP HANA software configuration.
   a. 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.

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

```
/hana/shared/HIT/hdblcm/hdblcm -action=add_hosts
\ --sid=HIT
```

```
\ --install_hostagent
\ --sapmnt=/hana/shared
\ --addhosts=hana002ic:role=standby:group=default
```

SAP HANA software configuration

The following recommended SAP HANA database parameter values were used to achieve maximum node scalability on VSP One Block 24, VSP One Block 26, and VSP One Block 28 that are certified by SAP:

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

The following table lists SAP HANA database parameter values as a reference for optimized performance.

	VSP One Block 28	VSP One Block 26	VSP One Block 24
max_parallel_io_requests [data]	64	64	64
max_submit_batch_size [data]	64	64	64
size_kernel_io_queue [data]	512	512	512
max_parallel_io_requests [log]	64	64	64
max_submit_batch_size [log]	64	64	64
size_kernel_io_queue [log]	512	512	512

To customize the values, follow SAP Note 2399079 to set up these parameters defined in *global.ini* or run the command for SAP HANA 2.0.

The following is an example of a *global.ini* file used with recommended SAP HANA database parameter values:

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

```
[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
partition_3_log lvmname = vgHITlog003-lvlog
partition_3_data lvmname = vgHITdata003-lvdata

[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-077.

The following table shows the hardware elements used during the 2+1 scale-out test including Network-HA, FC-HA, and SAP 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 processor 28-core,2.1 GHz, 165 W Or • 4 Intel Xeon Platinum 8180/8180M processor 28-core,2.5 GHz, 205 W	SAP HANA server

Hardware	Quantity	Configuration	Role
		▪ 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	
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 SAP HANA Studio
Hitachi Virtual Storage Platform One Block 28	1	▪ Controller board (CTL): 1 pair ▪ CPU: 4 (1 pair/CTL) ▪ Cache flash memory (CFM): 2 (1/CTL) ▪ Backup module (BKM): 2 (1/CTL) ▪ Battery: 2 (1/BKM) ▪ Fan:12 (6/CTL) ▪ MPU: 1 pair ▪ Maximum cache: 1024 GB ▪ Used 64 Gbps 4-port channel board (CHB): 1 pair ▪ Used Fibre Channel ports: 6 (six ports used by three SAP HANA nodes (2/node) ▪ NVMe SSD drives: 9 × 7.6 TB	

Hardware	Quantity	Configuration	Role
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	
Cisco Nexus 93180YC-FX Switch	4	48 × 10 GbE ports	10 GbE Linux HA cluster communication network, SAP HANA inter-cluster network, client network
Cisco Nexus 3048 switch	1	48 × 1 GbE ports	1GbE management network

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

Feature	SAP HANA Nodes (2+1) Scale-out Configuration
Server	3 Advanced Server DS7040
Server	Each Advanced Server DS7040 serves 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 each compute module in the following locations: ▪ PCIe Slot 1 ▪ PCIe Slot 3

Feature	SAP HANA Nodes (2+1) Scale-out Configuration
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

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
- SUSE HAE or RHEL HA cluster 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 to the designated Hitachi Virtual Storage Platform One B28 Fibre Channel ports to achieve no single point of failure and high performance.

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 following table shows the storage configuration and provisioning used for the 2+1 scale-out test.

Item	
Microcode Level	93-03-21-40/00
Cache Memory	1024 GB
Number of Ports	6× 64 Gbps ports used

Item	
Number of RAID Groups in Dynamic Provisioning Pool	1
Number of Dynamic Provisioning Pools	1
Number of LUs in Dynamic Provisioning Pools	18
RAID Group Type	DDP RAID 6 (6D+2P)
Number of Drives per PG	9
Drive Capacity	7.6 TB
Drive Type	NVMe SSD
Multi-Pathing Enabled	Yes

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

DB 00	0	1	….	7	8	9	10	….	16	17	18	19	….	22	23
DDP RAID Group, RAID 6 (6D+2P)	PG-1					N/A									

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

DDP Pool	Parity Group ID	DDP RAID Group RAID Level and disks	LDEV ID	LDEV Size*	MPU Assignment
Pool_for_HANA	1-1	DDP RAID6 (6D+2P) on 9 × 7.6TB NVMe SSDs	0	3071 GB	MPU-10
			1	2209 GB	MPU-20
			2	3071 GB	MPU-10
			3	2209 GB	MPU-20
			4	3071 GB	MPU-10
			5	2209 GB	MPU-20
			6	3071 GB	MPU-10

DDP Pool	Parity Group ID	DDP RAID Group RAID Level and disks	LDEV ID	LDEV Size*	MPU Assignment
			7	2209 GB	MPU-20
			8	3071 GB	MPU-10
			9	2209 GB	MPU-20
			10	3071 GB	MPU-10
			11	2209 GB	MPU-20
			12	3071 GB	MPU-10
			13	2209 GB	MPU-20
			14	3071 GB	MPU-10
			15	2209 GB	MPU-20
			16	3071 GB	MPU-10
			17	2209 GB	MPU-20
*The LDEV size can be auto-sized when creating LDEVs as pool volumes.					

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

DDP Pool	LDEV ID	vVol Name	vVol Size	MPU Assignment
Pool_for_HANA	110	HANA_OS_N1	100 GB	MPU-10
	120	HANA_OS_N2	100 GB	MPU-20
	130	HANA_OS_N3	100 GB	MPU-10
	100	STONITH	50 MB	MPU-20
	101	HANA_SH_1	256 GB	MPU-10
	102	HANA_SH_2	256 GB	MPU-20
	103	HANA_SH_3	256 GB	MPU-10
	104	HANA_SH_4	256 GB	MPU-20
	116	HANA_DATA_N1_1	600 GB	MPU-10
	117	HANA_DATA_N1_2	600 GB	MPU-20
	118	HANA_DATA_N1_3	600 GB	MPU-10

DDP Pool	LDEV ID	vVol Name	vVol Size	MPU Assignment
	119	HANA_DATA_N1_4	600 GB	MPU-20
	126	HANA_DATA_N2_1	600 GB	MPU-10
	127	HANA_DATA_N2_2	600 GB	MPU-20
	128	HANA_DATA_N2_3	600 GB	MPU-10
	129	HANA_DATA_N2_4	600 GB	MPU-20
	136	HANA_DATA_N3_1	600 GB	MPU-10
	137	HANA_DATA_N3_2	600 GB	MPU-20
	138	HANA_DATA_N3_3	600 GB	MPU-10
	139	HANA_DATA_N3_4	600 GB	MPU-20
	112	HANA_LOG_N1_1	128 GB	MPU-10
	113	HANA_LOG_N1_2	128 GB	MPU-20
	114	HANA_LOG_N1_3	128 GB	MPU-10
	115	HANA_LOG_N1_4	128 GB	MPU-20
	122	HANA_LOG_N2_1	128 GB	MPU-10
	123	HANA_LOG_N2_2	128 GB	MPU-20
	124	HANA_LOG_N2_3	128 GB	MPU-10
	125	HANA_LOG_N2_4	128 GB	MPU-20
	132	HANA_LOG_N3_1	128 GB	MPU-10
	133	HANA_LOG_N3_2	128 GB	MPU-20
	134	HANA_LOG_N3_3	128 GB	MPU-10
	135	HANA_LOG_N3_4	128 GB	MPU-20

Product descriptions

The following products were used in this reference architecture.

Hitachi Virtual Storage Platform One Block

The Hitachi Virtual Storage Platform One Block series simplifies system setup and management through Hitachi Clear Sight and VSP One Block Administrator. Dynamic Drive Protection reduces RAID complexity, and always-on compression and deduplication enhance simplicity.

Dynamic Carbon Reduction optimizes energy usage by switching CPUs to ECO mode during low activity. Adaptive Data Reduction (ADR) is always on, enhancing efficiency and reducing the overall CO2 footprint.

Thin Image Advanced (TIA) integrates with major snapshot ecosystems, prioritizing security by defending against threats and ensuring data confidentiality. CyberArk Privileged Access Manager plugins enhance block storage system security by prioritizing data confidentiality, ensuring compliance, and actively defending against security threats.

Hitachi Virtual Storage Platform One Block 20 includes 3 dedicated models. All have the same capacity (72 NVMe flash drives, the appliance and 2 × media trays) and they support Fibre Channel, iSCSI, and NVMe TCP connectivity. The new capabilities remove complexity: data reduction is always on, Dynamic Drive Protection removes complicated RAID setup, and Dynamic Carbon Reduction delivers real world reduction in power consumption. In addition, the models are FIPS compliant.

- VSP One Block 24 – 256 GB Cache + SW Advanced Data Reduction (ADR) + 24 cores
- VSP One Block 26 – 768GB Cache + 2 × Compression Accelerator Module (CAM) + 24 cores
- VSP One Block 28 – 1TB Cache + 4 × CAM + 64 cores

In short, the Hitachi Virtual Storage Platform One Block series combines simplicity, sustainability, and robust security features to optimize system management, energy efficiency, and data protection.

Hitachi Storage Virtualization Operating System RF

Hitachi Block Storage

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.

See https://www.hitachivantara.com/en-us/products/storage-platforms/storage-software for more information.

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.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact