

## **SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform One Block High End**

---

Reference Architecture Guide

© 2026 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., Hitachi Vantara, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at [https://support.hitachivantara.com/en\\_us/contact-us.html](https://support.hitachivantara.com/en_us/contact-us.html).

**Notice:** Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

**Notice on Export Controls.** The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, IntelliMagic, IntelliMagic Vision, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z17, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

The open source content used in Hitachi Vantara products may be found within the Product documentation or you may request a copy of such information (including source code and/or modifications to the extent the license for any open source requires Hitachi make it available) by sending an email to [OSS\\_licensing@hitachivantara.com](mailto:OSS_licensing@hitachivantara.com).

## Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to [Docs-Feedback@hitachivantara.com](mailto:Docs-Feedback@hitachivantara.com). To assist the routing of this message, use the paper number in the subject and the title of this guide in the text.

Thank you!

Changes	Date
▪ Initial release.	March 2026

---

# Contents

- Reference Architecture Guide.....4**
- Solution overview.....6
- Key solution elements.....8
- Solution design.....9
  - Fibre Channel architecture.....9
  - Storage architecture.....14
  - Virtual Storage Scale-out (VSSO) configuration.....19
  - SAP HANA configuration.....19
    - Multipath configuration.....19
    - SAP HANA persistent storage volume configuration.....21
  - SAP storage connector API Fibre Channel client.....24
  - SAP HANA software installation.....25
    - Install SAP HANA software in a scale-up environment.....25
    - Install SAP HANA software in a scale-out environment.....26
    - Configure the SAP HANA software.....28
- Engineering validation.....29
- Product descriptions.....35
  - Hitachi Virtual Storage Platform One Block High End.....35
  - Hitachi Storage Virtualization Operating System.....35

---

## Reference Architecture Guide

Use this reference architecture guide to implement an SAP® HANA Tailored Data Center Integration (TDI) solution for the SAP HANA platform using the Hitachi Virtual Storage Platform One Block High End storage system.

This solution provides the storage requirements and configurations for the maximum number of validated active production nodes in an SAP HANA deployment on the following VSP One Block High End configurations:

- VSP One Block High End 2 nodes using NVMe SSDs
- VSP One Block High End 4 nodes using NVMe SSDs
- VSP One Block High End 6 nodes using NVMe SSDs
- Hitachi Virtual Storage Scale-out (VSSO) to cluster multiple VSP One BHE 6-node storage systems

There are two types of SAP HANA deployments:

- SAP HANA appliance – a delivery model of preconfigured hardware with preinstalled software packages that can be implemented by an SAP HANA hardware partner or vendor.
- SAP HANA TDI – a delivery approach with more flexibility to choose hardware from any certified SAP HANA server vendor and any certified SAP HANA enterprise storage vendor, optimally integrate in your data center and reduce the total cost of ownership (TCO).

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

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

Engineering validation included and confirmed the following:

- Servers from Hitachi Vantara were used. This solution is also valid for other [Hitachi-supported servers](#).
- Testing showed that VSP One Block High End meets the SAP enterprise storage certification key performance indicator (KPI) requirements for SAP HANA.

- Scalability and storage KPI testing were performed using the SAP HANA Hardware and Cloud Measurement Tool (HCMT) to measure and analyze the hardware or cloud system before deploying SAP HANA or applying for SAP HANA certification. See SAP Note [2493172 - SAP HANA Hardware and Cloud Measurement Tool](#) (SAP user credentials required) for HCMT details.
- The storage platform passed the SAP enterprise storage KPIs. While running HCMT on SAP HANA hosts during testing, the storage system was fully dedicated to the test environment, and no other systems or applications were accessing the storage at that time.



**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 and Cloud Measurement Tool (HCMT) to meet SAP KPIs.

The following table gives an overview of the supported maximum number of SAP HANA systems with the VSP One Block High End storage system. Maximum SAP HANA node scalability values are projected based on engineering validation results.

Storage Model Name	Drive Type	Maximum SAP HANA Systems	Minimum Parity Group	
			Data	Log
VSP One BHE-2N (1 pair of nodes)	NVMe SSDs	84	9 × RAID 1 (2D+2D)	
VSP One BHE-4N (2 pairs of nodes)	NVMe SSDs	168	18 × RAID 1 (2D+2D)	
VSP One BHE-6N (3 pairs of nodes)	NVMe SSDs	252	27 × RAID 1 (2D+2D)	
VSSO for 4 × VSP One BHE-6N	NVMe SSDs	1008	108 × RAID 1 (2D+2D)	

When implementing an SAP HANA TDI infrastructure on VSP One Block High End, you are not required to replicate the exact storage design used in this reference architecture for enterprise storage KPI testing. However, the recommendation is to follow the guidelines outlined in this document for the setup and configuration of the storage system. See *SAP HANA Tailored Data Center Integration - Frequently Asked Questions* from the [SAP HANA Master Guide](#) for more details.



**Note:** This configuration was tested in a lab environment. Conduct proof-of-concept testing in a non-production environment that matches your production environment before implementing this solution.

## Solution overview

This reference architecture guide provides example configurations tested in the Hitachi Vantara lab, using NVMe SSD-based storage layouts for SAP HANA nodes on VSP One Block High End.

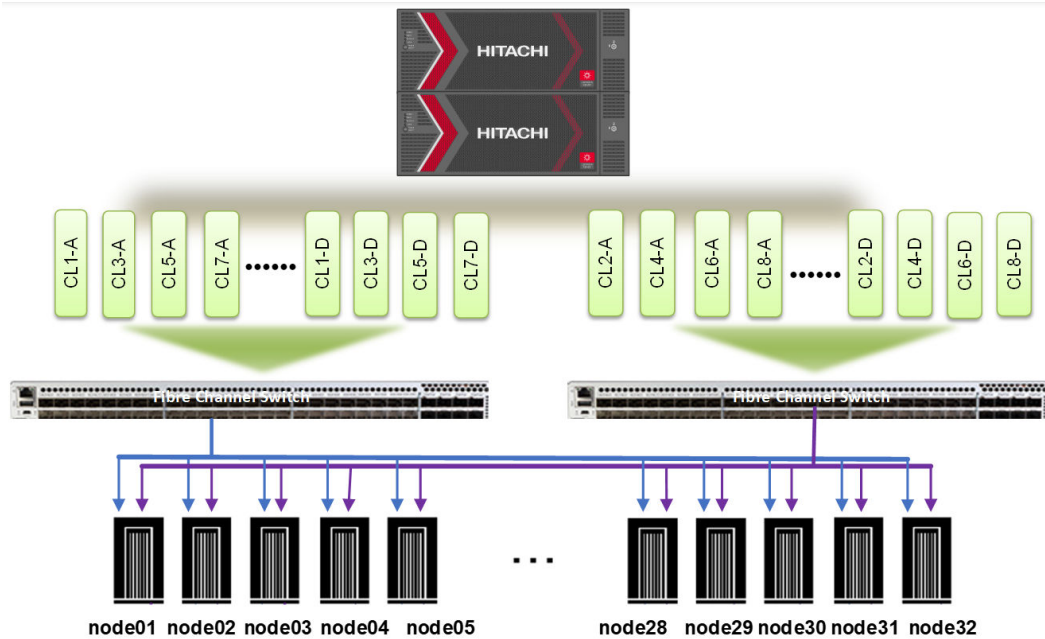
The following table lists the system specifications of VSP One Block High End.

Item		Specifications			
		VSP One Block High End (6CBX)	VSP One Block High End (4CBX)	VSP One Block High End (2CBX)	
System	Hitachi Storage Virtualization Operating System RF (SVOS RF) version	10.5.1 or higher			
	Number of nodes	6	4	2	
	Maximum NVMe SSD	288	192	96	
	NVMe SSD Drive Options	3.8TB, 7.6 TB, 15.3 TB, 30 TB, 60 TB			
	RAID Level	RAID 6, RAID 5, RAID 1*			
	RAID group configuration	RAID 1	2D+2D		
		RAID 5	3D+1P, 7D+1P		
		RAID 6	6D+2P, 14D+2P		
	Maximum number of LDEVs	65280			
	Maximum storage capacity (physical disk capacity)	15.1 PB	10.1 PB	5.0 PB	
Maximum external configuration	255 PB				
Maximum number of DBN2	12	8	4		
Memory	Cache memory capacity	6,144 GB		2,048 GB	
	Cache Flash Memory Module	BMB0/BMBE			
Storage Interface	DKC to DB Interface	NVMe (PCIe)/Dual Port			
	Data Transfer Rate	16 Gbps			
	Maximum Number of DKBs	24	16	8	
Device Interface	Fibre Channel	196	128	64	
	Transfer Speed	8/16/32/64 Gbps			

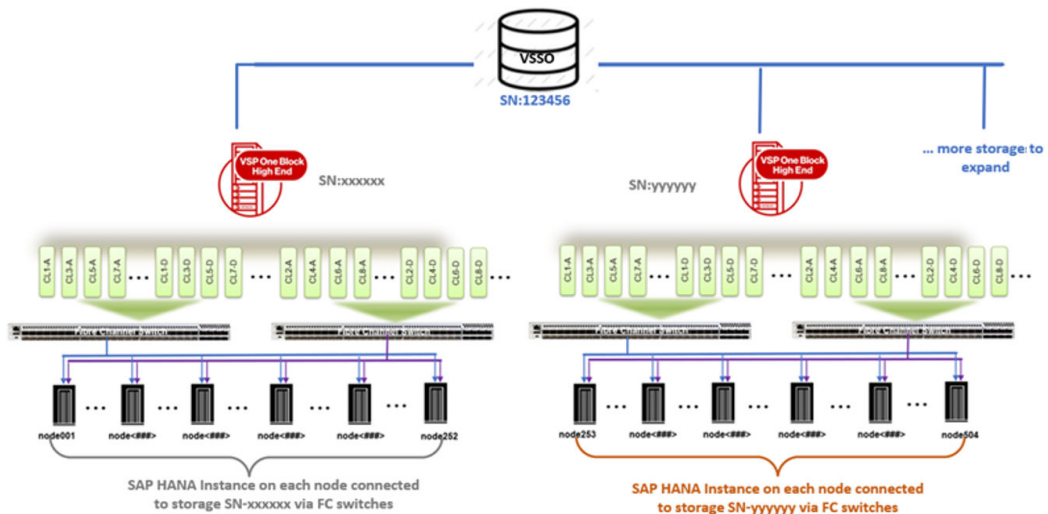
Item		Specifications		
		VSP One Block High End (6CBX)	VSP One Block High End (4CBX)	VSP One Block High End (2CBX)
	Maximum number of CHBs	48	32	16

\* RAID 1 supported by this storage system is commonly called RAID 1+0.

The following figure shows a server-to-storage Fibre Channel switch in a direct connection configuration. This example uses VSP One Block High End with 2 nodes and 32 SAP HANA systems.



The following figure shows an architecture using Hitachi Virtual Storage Scale-out (VSSO) for multi-clustered VSP BHE 6-node storage systems to run multiple SAP HANA instances.



The VSSO architecture provides a flexible and scalable storage solution for SAP HANA Enterprise Storage with the following key characteristics:

- Greater SAP HANA node scalability by adding up to the total number of SAP HANA nodes for multiple VSP One BHE storage systems with 6 nodes.
- No performance impact when adding more storage devices in a loosely clustered VSSO.
- Unified management of all storage systems from a single management tool – Hitachi Ops Center.
- A familiar connectivity model that uses the same physical connections between SAP HANA hosts and non-cluster storage for both SAN (using Fibre Channel switches) and direct-attached storage (DAS) without additional networking complexity.

## Key solution elements

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

### Hardware elements

The following table lists the tested active SAP HANA nodes that use NVMe SSDs for maximum scalability.

Hardware	Number of SAP HANA Nodes	Storage Components	Fibre Channel Switch
VSP One Block High End 2 nodes	32	CTL: 2 pairs	48-port Brocade G720 Fibre Channel Switch (1 pair)
		CHB: 4 pairs (32 Gbps)	
		MPU: 2 pairs	
		Cache: 2 TB	
		48 × 3.8 TB NVMe SSDs	

The following table shows the optional hardware that can be used.

Optional Items	Components
Management server	Hitachi Advanced Server HA810 G3
1 GbE management network switch, using either one of the Cisco switches	Cisco Nexus 92348 Cisco 93180YC-FX/FX3
10 GbE network switch, using either one of the Cisco switches	

### Server for an SAP HANA TDI environment

Only servers certified for use with SAP HANA can be used in an SAP HANA Tailored Data Center Integration (TDI) environment. This follows the same bill of materials as the certified SAP HANA appliance server, but without the storage or local disks.

Find a list of all certified servers and enterprise storage solutions in [Certified and Supported SAP HANA® Hardware Directory](#).

For more information about SAP HANA TDI, see *SAP HANA Tailored Data Center Integration - Frequently Asked Questions* from the [SAP HANA Master Guide](#).

### Software elements

The following software can be used in this environment:

- SAP HANA
- Operating system:
  - SUSE Linux Enterprise Server for SAP Applications (used in this solution)
  - Red Hat Enterprise Linux
- Storage management platforms:
  - Command Control Interface (CCI)
  - Hitachi VSP 360
  - Hitachi Ops Center (if using VSSO)

## Solution design

This is a detailed solution example for the SAP HANA Tailored Data Center Integration (TDI) on Hitachi Virtual Storage Platform One Block High End.

### Fibre Channel architecture

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

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

Distribute host connections across different storage controllers and channel boards to ensure high availability and optimal performance.

The following table shows example storage port mapping and Fibre Channel switches for 32 SAP HANA hosts on VSP One Block High End storage with 2 nodes.

SAP HANA Host	HBA Port		Fibre Channel Switch		Storage Target Port			
	Port	Speed (Gbps)	Host	Storage	Port	Host_ Grp Name	Speed (Gbps)	Port Security
node0 01	Port 0	16	SW-1 P16	SW-1- P0	CL1- A	node00 1	32/64	Enabled
	Port 1	16	SW-2-P16	SW-2- P0	CL2- A	node00 1	32/64	Enabled
node0 02	Port 0	16	SW-1-P17	SW-1- P0	CL1- A	node00 2	32/64	Enabled
	Port 1	16	SW-2-P17	SW-2- P0	CL2- A	node00 2	32/64	Enabled
node0 03	Port 0	16	SW-1-P18	SW-1- P1	CL3- A	node00 3	32/64	Enabled
	Port 1	16	SW-2-P18	SW-2- P1	CL4- A	node00 3	32/64	Enabled
node0 04	Port 0	16	SW-1-P19	SW-1- P1	CL3- A	node00 4	32/64	Enabled
	Port 1	16	SW-2-P19	SW-2- P1	CL4- A	node00 4	32/64	Enabled
node0 05	Port 0	16	SW-1-P20	SW-1- P2	CL5- A	node00 5	32/64	Enabled
	Port 1	16	SW-2-P20	SW-2- P2	CL6- A	node00 5	32/64	Enabled
node0 06	Port 0	16	SW-1-P21	SW-1- P2	CL5- A	node00 6	32/64	Enabled
	Port 1	16	SW-2-P21	SW-2- P2	CL6- A	node00 6	32/64	Enabled
node0 07	Port 0	16	SW-1-P22	SW-1- P3	CL7- A	node00 7	32/64	Enabled
	Port 1	16	SW-2-P22	SW-2- P3	CL8- A	node00 7	32/64	Enabled
node0 08	Port 0	16	SW-1-P23	SW-1- P3	CL7- A	node00 8	32/64	Enabled
	Port 1	16	SW-2-P23	SW-2- P3	CL8- A	node00 8	32/64	Enabled

SAP HANA Host	HBA Port		Fibre Channel Switch		Storage Target Port			
	Port	Speed (Gbps)	Host	Storage	Port	Host_ Grp Name	Speed (Gbps)	Port Security
node009	Port 0	16	SW-1-P24	SW-1-P4	CL1-B	node009	32/64	Enabled
	Port 1	16	SW-2-P24	SW-2-P4	CL2-B	node009	32/64	Enabled
node010	Port 0	16	SW-1-P25	SW-1-P4	CL1-B	node010	32/64	Enabled
	Port 1	16	SW-2-P25	SW-2-P4	CL2-B	node010	32/64	Enabled
node011	Port 0	16	SW-1-P26	SW-1-P5	CL3-B	node011	32/64	Enabled
	Port 1	16	SW-2-P26	SW-2-P5	CL4-B	node011	32/64	Enabled
node012	Port 0	16	SW-1-P27	SW-1-P5	CL3-B	node012	32/64	Enabled
	Port 1	16	SW-2-P27	SW-2-P5	CL4-B	node012	32/64	Enabled
node013	Port 0	16	SW-1-P28	SW-1-P6	CL5-B	node013	32/64	Enabled
	Port 1	16	SW-2-P28	SW-2-P6	CL6-B	node013	32/64	Enabled
node014	Port 0	16	SW-1-P29	SW-1-P6	CL5-B	node014	32/64	Enabled
	Port 1	16	SW-2-P29	SW-2-P6	CL6-B	node014	32/64	Enabled
node015	Port 0	16	SW-1-P30	SW-1-P7	CL7-B	node015	32/64	Enabled
	Port 1	16	SW-2-P30	SW-2-P7	CL8-B	node015	32/64	Enabled
node016	Port 0	16	SW-1-P31	SW-1-P7	CL7-B	node016	32/64	Enabled
	Port 1	16	SW-2-P31	SW-2-P7	CL8-B	node016	32/64	Enabled

SAP HANA Host	HBA Port		Fibre Channel Switch		Storage Target Port			
	Port	Speed (Gbps)	Host	Storage	Port	Host_ Grp Name	Speed (Gbps)	Port Security
node0 17	Port 0	16	SW-1-P32	SW-1- P8	CL1- C	node01 7	32/64	Enabled
	Port 1	16	SW-2-P32	SW-2- P8	CL2- C	node01 7	32/64	Enabled
node0 18	Port 0	16	SW-1-P33	SW-1- P8	CL1- C	node01 8	32/64	Enabled
	Port 1	16	SW-2-P33	SW-2- P8	CL2- C	node01 8	32/64	Enabled
node0 19	Port 0	16	SW-1-P34	SW-1- P9	CL3- C	node01 9	32/64	Enabled
	Port 1	16	SW-2-P34	SW-2- P9	CL4- C	node01 9	32/64	Enabled
node0 20	Port 0	16	SW-1-P35	SW-1- P9	CL3- C	node02 0	32/64	Enabled
	Port 1	16	SW-2-P35	SW-2- P9	CL4- C	node02 0	32/64	Enabled
node0 21	Port 0	16	SW-1-P36	SW-1- P10	CL5- C	node02 1	32/64	Enabled
	Port 1	16	SW-2-P36	SW-2- P10	CL6- C	node02 1	32/64	Enabled
node0 22	Port 0	16	SW-1-P37	SW-1- P10	CL5- C	node02 2	32/64	Enabled
	Port 1	16	SW-2-P37	SW-2- P10	CL6- C	node02 2	32/64	Enabled
node0 23	Port 0	16	SW-1-P38	SW-1- P11	CL7- C	node02 3	32/64	Enabled
	Port 1	16	SW-2-P38	SW-2- P11	CL8- C	node02 3	32/64	Enabled
node0 24	Port 0	16	SW-1-P39	SW-1- P11	CL7- C	node02 4	32/64	Enabled
	Port 1	16	SW-2-P39	SW-2- P11	CL8- C	node02 4	32/64	Enabled

SAP HANA Host	HBA Port		Fibre Channel Switch		Storage Target Port			
	Port	Speed (Gbps)	Host	Storage	Port	Host_ Grp Name	Speed (Gbps)	Port Security
node0 25	Port 0	16	SW-1-P40	SW-1- P12	CL1- D	node02 5	32/64	Enabled
	Port 1	16	SW-2-P40	SW-2- P12	CL2- D	node02 5	32/64	Enabled
node0 26	Port 0	16	SW-1-P41	SW-1- P12	CL1- D	node02 6	32/64	Enabled
	Port 1	16	SW-2-P41	SW-2- P12	CL2- D	node02 6	32/64	Enabled
node0 27	Port 0	16	SW-1-P42	SW-1- P13	CL3- D	node02 7	32/64	Enabled
	Port 1	16	SW-2-P42	SW-2- P13	CL4- D	node02 7	32/64	Enabled
node0 28	Port 0	16	SW-1-P43	SW-1- P13	CL3- D	node02 8	32/64	Enabled
	Port 1	16	SW-2-P43	SW-2- P13	CL4- D	node02 8	32/64	Enabled
node0 29	Port 0	16	SW-1-P44	SW-1- P14	CL5- D	node02 9	32/64	Enabled
	Port 1	16	SW-2-P44	SW-2- P14	CL6- D	node02 9	32/64	Enabled
node0 30	Port 0	16	SW-1-P46	SW-1- P14	CL5- D	node03 0	32/64	Enabled
	Port 1	16	SW-2-P45	SW-2- P14	CL6- D	node03 0	32/64	Enabled
node0 31	Port 0	16	SW-1-P46	SW-1- P14	CL7- D	node03 1	32/64	Enabled
	Port 1	16	SW-2-P46	SW-2- P14	CL8- D	node03 1	32/64	Enabled
node0 32	Port 0	16	SW-1-P47	SW-1- P14	CL7- D	node03 2	32/64	Enabled
	Port 1	16	SW-2-P47	SW-2- P14	CL8- D	node03 2	32/64	Enabled

## Storage architecture

Each SAP HANA node requires the following storage layout:

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

This SAP HANA TDI setup uses one dynamic provisioning pool for the storage layout, ensuring maximum resource use and optimization at a lower cost.

SAP\_HANA\_Pool is used for the following:

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

The validated dynamic provisioning pool layout with minimal disks and maximum storage cache is provided in the following table.

Storage	Cache	SAP HANA Node Number	Number of Parity Groups in SAP_HANA_Pool
Hitachi Virtual Storage Platform One Block High End 2 nodes	2 TB	Up to 32	6 × RAID 1 (2D+2D)

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

- The various combinations of SAP HANA node sizes
- The number of SAP HANA 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 and follow these guidelines:

- Determine the minimum sizes for operating system, data, log, and SAP HANA shared using these formulas in the SAP white paper [SAP HANA Master Guide for SAP HANA Storage Requirements](#) as follows:
  - Every SAP HANA node requires approximately 100 GB capacity for the operating system.
  - The /hana/shared size uses these formulas:  
 Single node (scale-up) — Size = MIN (1 × RAM; 1 TB)  
 Multi node (scale-out) — Size = 1 × RAM\_of\_worker per 4 worker nodes
  - Data size requires at least 1 × RAM on each SAP HANA node
  - Log size uses these formulas:  
 Systems with equal or less than 512 GB of memory — size = ½ × RAM  
 Systems with greater than 512 GB of memory — size = 512 GB
- Provision the storage:
  - Create one dynamic provisioning pool for the SAP HANA systems as follows:
    - Use SAP\_HANA\_Pool to provision the operating system volume, SAP HANA shared volume, data volume, and log volume.
    - Create the parity groups (PGs) and logical devices (LDEVs or volumes) first using the RAID 1 design as shown in the following table.

HDP	PG ID	PG RAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
SAP_HANA_Pool	1	RAID 1 (2D+2D) On 3.8 TB NVMe SSDs	00:00:01	1760 GB	MPU-010
			00:00:02	1760 GB	MPU-020
			00:00:03	1760 GB	MPU-110
			00:00:04	1760 GB	MPU-120
	2	RAID 1 (2D+2D) On 3.8 TB NVMe SSDs	00:00:05	1760 GB	MPU-010
			00:00:06	1760 GB	MPU-020
			00:00:07	1760 GB	MPU-110
			00:00:08	1760 GB	MPU-120
	3	RAID 1 (2D+2D) On 3.8 TB NVMe SSDs	00:00:09	1760 GB	MPU-010
			00:00:0A	1760 GB	MPU-020
			00:00:0B	1760 GB	MPU-110
			00:00:0C	1760 GB	MPU-120

HDP	PG ID	PG RAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
	4	RAID 1 (2D+2D) On 3.8 TB NVMe SSDs	00:00:0D	1760 GB	MPU-120
			00:00:0E	1760 GB	MPU-020
			00:00:0F	1760 GB	MPU-110
			00:00:10	1760 GB	MPU-120
	5	RAID 1 (2D+2D) On 3.8 TB NVMe SSDs	00:00:11	1760 GB	MPU-010
			00:00:12	1760 GB	MPU-010
			00:00:13	1760 GB	MPU-010
			00:00:14	1760 GB	MPU-010
	6	RAID 1 (2D+2D) On 3.8 TB NVMe SSDs	00:00:15	1760 GB	MPU-010
			00:00:16	1760 GB	MPU-020
			00:00:17	1760 GB	MPU-110
			00:00:18	1760 GB	MPU-120

- Assign all LDEVs to the dedicated pool from the previous table.
- Create virtual volumes for the operating system, SAP HANA shared, log, and data volumes. The following table lists examples for SAP HANA systems with 384 GB, 768 GB, 1536 GB, and 3072 GB of memory.

HDP	LDEV ID	LDEV Name	Size	MPU Assignment	System Memory
SAP_HANA_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-010	384 GB
	00:02:00	HANA_OS_N2	100 GB	MPU-020	768 GB
	00:03:00	HANA_OS_N3	100 GB	MPU-110	1536 GB
	00:04:00	HANA_OS_N4	100 GB	MPU-120	3072 GB
	00:01:01	HANA_SH_N1	384 GB	MPU-010	384 GB
	00:02:01	HANA_SH_N2	768 GB	MPU-020	768 GB
	00:03:01	HANA_SH_N3	1536 GB	MPU-110	1536 GB
	00:04:01	HANA_SH_N4	3072 GB	MPU-120	3072 GB

HDP	LDEV ID	LDEV Name	Size	MPU Assignment	System Memory
	00:01:02	HANA_LOG_N1_1	48 GB	MPU-010	384 GB
	00:01:03	HANA_LOG_N1_2	48 GB	MPU-020	
	00:01:04	HANA_LOG_N1_3	48 GB	MPU-110	
	00:01:05	HANA_LOG_N1_4	48 GB	MPU-120	
	00:01:06	HANA_DATA_N1_1	96 GB	MPU-010	
	00:01:07	HANA_DATA_N1_2	96 GB	MPU-020	
	00:01:08	HANA_DATA_N1_3	96 GB	MPU-110	
	00:01:09	HANA_DATA_N1_4	96 GB	MPU-120	
	00:02:02	HANA_LOG_N2_1	96 GB	MPU-010	768 GB
	00:02:03	HANA_LOG_N2_2	96 GB	MPU-020	
	00:02:04	HANA_LOG_N2_3	96 GB	MPU-110	
	00:02:05	HANA_LOG_N2_4	96 GB	MPU-120	
	00:02:06	HANA_DATA_N2_1	192 GB	MPU-010	
	00:02:07	HANA_DATA_N2_2	192 GB	MPU-020	
	00:02:08	HANA_DATA_N2_3	192 GB	MPU-110	
	00:02:09	HANA_DATA_N2_4	192 GB	MPU-120	
	00:03:02	HANA_LOG_N3_1	128 GB	MPU-010	1536 GB
	00:03:03	HANA_LOG_N3_2	128 GB	MPU-020	
	00:03:04	HANA_LOG_N3_3	128 GB	MPU-110	
	00:03:05	HANA_LOG_N3_4	128 GB	MPU-120	
	00:03:06	HANA_DATA_N3_1	384 GB	MPU-010	
	00:03:07	HANA_DATA_N3_2	384 GB	MPU-020	
	00:03:08	HANA_DATA_N3_3	384 GB	MPU-110	
	00:03:09	HANA_DATA_N3_4	384 GB	MPU-120	
	00:04:02	HANA_LOG_N4_1	128 GB	MPU-010	3072 GB
	00:04:03	HANA_LOG_N4_2	128 GB	MPU-020	
	00:04:04	HANA_LOG_N4_3	128 GB	MPU-110	

HDP	LDEV ID	LDEV Name	Size	MPU Assignment	System Memory
	00:04:05	HANA_LOG_N4_4	128 GB	MPU-120	
	00:04:06	HANA_DATA_N4_1	768 GB	MPU-010	
	00:04:07	HANA_DATA_N4_2	768 GB	MPU-020	
	00:04:08	HANA_DATA_N4_3	768 GB	MPU-110	
	00:04:09	HANA_DATA_N4_4	768 GB	MPU-120	

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

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

LUN ID	LDEV ID	Volume Name	File System Mount Point
0000	00:01:00	HANA_OS_N1	/(root)
0001	00:01:01	HANA_SH_N1	/hana/shared
0002	00:01:02	HANA_LOG_N1_1	/hana/log
0003	00:01:03	HANA_LOG_N1_2	
0004	00:01:04	HANA_LOG_N1_3	
0005	00:01:05	HANA_LOG_N1_4	
0006	00:01:06	HANA_DATA_N1_1	/hana/data
0007	00:01:07	HANA_DATA_N1_2	
0008	00:01:08	HANA_DATA_N1_3	
0009	00:01:09	HANA_DATA_N1_4	

## Virtual Storage Scale-out (VSSO) configuration

Virtual Storage Scale Out (VSSO) is a scale-out feature that allows multiple physical storage systems to be seen as a single mass storage system by defining the common virtual storage machine (VSM) for multiple physical storage systems. This does not change physical connections between the SAP HANA host and each storage system (using either DAS or SAN).

Use the following guidance to implement VSSO:

- Make physical connections between each SAP HANA host and a storage system, the same as for non-cluster storage.
- Storage provisioning:
  - Infrastructure level provisioning (dynamic provisioning, Virtual Volumes, server registration, and volume attachment) is configured for each storage system.
  - VSSO with multiple 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 single storage serial number. When a virtual storage machine (VSM) is created, a new storage resource group is generated for a VSSO to allow you to logically group multiple storage resources (such as pools, volumes and ports) within VSSO.
  - Logical provisioning (for example, expanding VSSO with additional 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](#).

## SAP HANA configuration

This section describes 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 and storage.

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

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

- SAP HANA server boot volume  
(not applicable if using internal disks for the boot volume).
- SAP HANA data volume

- SAP HANA log volume
- SAP HANA shared volume

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

- A VSP One File 30 series storage system as the NFS for /hana/shared.
- Global File System 2 (GFS2) for /hana/shared. This is for Hitachi-only TDI scale-out deployments where Hitachi provides the server and storage for the TDI landscape.

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

Section	Option	SUSE Linux Enterprise Server for SAP Applications	Red Hat Enterprise Linux
defaults		N/A	N/A
blacklist	devnode	$^(ram raw loop fd md dm- sr scd st)[0-9]^*$ $^hd[a-z]$ $^dcssblk[0-9]^*$	
devices/device	vendor	HITACHI	
	product	.*	
	user_friendly_names	No	
	path_checker	direction	
	path_grouping_policy	Multibus	
	path_selector	queue-length 0	
	uid_attribute	ID_SERIAL	
	failback	Immediate	
	rr_weight	Uniform	
	rr_min_io_rq	128	1
	features	0	
	no_path_retry	5	

## SAP HANA persistent storage volume configuration

For SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux operating systems, Hitachi Vantara uses an LVM-based storage layout. After the operating system is installed and multipathing is correctly configured, you can see the assigned LUNs in the following directories:

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

For example:

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

The last 6 digits indicate the LDEV ID used during the LUN assignment. In the example, 000101 maps to LDEV ID: 00:01:01.

### Procedure

1. Except for the LUN hosting the operating system, initialize the rest of the LUNs for use by LVM by running the `pvcreate` command, which is part of the `lvm2` rpm package.

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

2. After you have prepared the LUNs, configure the volume groups using the `vgcreate` command. The names for the volume groups 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 configuration using four physical disks/LUNs:

```
vgcreate vglog /dev/mapper/360060e801227fc00504027fc00000101[2,3,4,5]
```

3. To create additional volume groups, use the same syntax, revising the volume group name and the physical disks or LUNs.

After creating the volume groups, create a logical volume on top of the volume groups.

The following is the general syntax:

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

Use the following table when creating logical volumes.

	Number of LUNs	lv Name	vg Name
DATA	4- following this reference architecture or the number of assigned LUNs	lvdata	<ul style="list-style-type: none"> <li>Scale-up: vgdata</li> <li>Scale-out: vg&lt;SID&gt;data&lt;node number&gt;</li> </ul>
LOG	4- following this reference architecture or the number of assigned LUNs	lvlog	<ul style="list-style-type: none"> <li>Scale-up: vglog</li> <li>Scale-out: vg&lt;SID&gt;log&lt;node number&gt;</li> </ul>
SHARED <sup>1</sup>	1- following this reference architecture or the number of assigned LUNs	lvshared	<ul style="list-style-type: none"> <li>Scale-up: vgshared</li> </ul>
<sup>1</sup> Only applies to SAP HANA scale-up deployments.			



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

4. Create the file system on top of the logical volume.

The following table lists the XFS options to create and mount the file system.

	System Type	Create Options	Mount Options	Mount Point
DATA	Scale-up	-F	noatime,inode64,nodiratime	/hana/data
	Scale-out		N/A	/hana/data/<SID>
LOG	Scale-up	-F	noatime,inode64,nodiratime	/hana/log
	Scale-Out		N/A	/hana/log/<SID>
SHARED	Scale-up	-F	noatime,inode64,nodiratime	/hana/shared
	Scale-out*	N/A	vers=3,proto=tcp,hard,nolock,timeo=600,retrans=2,wsiz=65536,rsiz=65536	/hana/shared/<SID>

	System Type	Create Options	Mount Options	Mount Point
*The listed options apply to a VSP One File 30 series storage system as an NFS for /hana/shared. If you use GFS2 for /hana/shared, use <code>mkfs.gfs2</code> to create the file system. The GFS2 file system is managed by Linux Distributed Lock Manger (DLM).				

To create a file system, use the following command:

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

For example:

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

## SAP HANA persistent storage volume configuration for scale-up deployments

This procedure applies to scale-up systems only.

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

### Procedure

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

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

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

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

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

Example for scale-up:

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

3. Mount the file systems.

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

```
mount <mount point>
```

To mount the file systems all at the same time, use the following command:

```
mount -a
```

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

## SAP HANA persistent storage volume configuration for scale-out deployments

This procedure applies to scale-out systems only.

### 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 VSP One File series or GFS2.

- VSP One File 30 series

- a. Persist the SAP HANA shared file systems in `/etc/fstab` by adding the following entry to `/etc/fstab`:

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

- b. Mount the file system:

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

- Global File System 2 (GFS2)



**Note:** GFS2 is supported by SAP only when using Hitachi storage and servers in TDI scale-out deployments for SAP HANA shared file system.

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 the `df` or `mount` command.

SAP HANA mounts the file system used for log and data with the correct options during startup.

See the following for more details:

- [SAP storage connector API Fibre Channel client \(on page 24\)](#)
- [SAP HANA software installation \(on page 25\)](#)

## SAP storage connector API Fibre Channel client

For a scale-out configuration, SAP HANA offers a storage connector client for configurations 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 (Python calls) for the following:

- Database startup
- Node failover

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

The `fcClient/fcClientLVM` implementation uses `multipath` and `sg_persist` commands that are responsible for mounting the SAP HANA volumes and implementing an effective fencing mechanism during a host failover using SCSI-3 persistent reservations.

The SAP Storage Connector API configuration is included in the SAP `global.ini` file in the `/hana/shared/<SID>/global/hdb/custom/config` directory.

## SAP HANA software installation

After configuring the file system for the SAP HANA data volume, log volume, and shared volume, install the current version of SAP HANA supported by SAP on the SAP HANA server or servers. See SAP Note [2235581 - SAP HANA: Supported Operating Systems](#). The `SAP_HANA_OS_Release_Support_Matrix.pdf` file attachment outlines which combinations of RHEL or SLES operating systems are supported with the various SAP HANA versions.

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

### Install SAP HANA software in a scale-up environment

Follow this procedure to install SAP HANA in 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

#### 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> \
--saoadn_oasswird=<password> \
--datapath=/hana/data/HIT \
--logpath=/hana/log/HIT \
--hostname=saphanas.company.corp \
--certificates_hostmap= saphanas.company.corp=saphanas.company.corp
```

3. Provide passwords when prompted.

**Result**

After the installation is complete, continue with [Configure the SAP HANA software \(on page 28\)](#).

**Install SAP HANA software in a scale-out environment**

Follow this procedure to install SAP HANA in a scale-out environment. 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

**Before you begin**

Prepare a configuration file for use by the SAP Storage Connector API Fibre Channel client. This file describes the communication, persistence, and storage details for the SAP HANA 2+1 scale-out environment, similar to the following example.

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

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

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

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_* prtype = 5
partition_1_log lvmname = vgHITlog001-lvlog
partition_1_data lvmname = vgHITdata001-lvdata
partition_2_log lvmname = vgHITlog002-lvlog
partition_2_data lvmname = vgHITdata002-lvdata

[trace]
Ha_fcclient = info
```

See the [SAP HANA Fibre Channel Storage Connector Admin Guide](#) for more details. Use this configuration file during the installation (it can be removed later). Save the file as `/tmp/hana_install_HIT/global.ini`. Keep the file name (`global.ini`) the same, even if you choose a different directory.

## 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=/tmp/hana_install_<number> \
--hostname=hana001ic \
--certificates_hostmap=hana001ic=hana001
```

3. Provide passwords when prompted.
4. To add a worker node to the existing installation, use the following command:

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



**Note:** Make sure that your primary installation is aware of the node you are adding. That means your configuration needs to include the partition information. The primary node owns partition ID 1, the first additional node can be named node 2, and so forth.

5. To add a standby node, use the following command:



**Note:** You do not need to provide the partition information, because a standby node gets the needed information during a failover.

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

```
--sapmnt=/hana/shared \  
--sapadm_password=<password> -p <password>
```

## Result

After the installation is complete, continue with [Configure the SAP HANA software \(on page 28\)](#).

## Configure the SAP HANA software

The following is an example of the SAP HANA 2.0 `global.ini` file used during validation of the environment:

```
[communication]  
tcp_backlog = 2048  
listeninterface = .global  
  
[fileio]  
max_parallel_io_requests[DATA] = 64  
max_parallel_io_requests[LOG] = 64  
max_submit_batch_size[DATA] = 64  
max_submit_batch_size[LOG] = 64  
size_kernel_io_queue[DATA] = 512  
size_kernel_io_queue[LOG] = 512  
async_read_submit[DATA] = on  
async_read_submit[LOG] = on  
async_write_submit_blocks[DATA] = all  
async_write_submit_blocks[LOG] = all  
min_submit_batch_size[DATA] = 16  
min_submit_batch_size[LOG] = 16  
async_write_submit_active[DATA] = on  
async_write_submit_active[LOG] = on  
  
[multidb]  
mode = multidb  
database_isolation = low  
singletenant = yes  
  
[persistence]  
basepath_datavolumes = /hana/data/HIT  
basepath_logvolumes = /hana/log/HIT  
  
[storage]  
ha_provider = hdb_ha.fcClientLVM  
partition_*_*_prtype = 5  
partition_1_data lvmname = vgHITdata001-lvdata  
partition_1_log lvmname = vgHITlog001-lvlog  
partition_2_data lvmname = vgHITdata002-lvdata  
partition_2_log lvmname = vgHITlog002-lvlog  
partition_3_data lvmname = vgHITdata003-lvdata  
partition_3_log lvmname = vgHITlog003-lvlog
```

```
[trace]
ha_fcclientlvm = info
```

## Engineering validation

The SAP HANA hardware and cloud measurement tool (HCMT) revision hcmt-088 was used for testing and certification of the SAP HANA TDI 2+1 scale-out enterprise storage configuration. See SAP Note [2493172](#) for more information.

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


Hardware	Quantity	Configuration	Role
Hitachi Advanced Server DS7040	3	<ul style="list-style-type: none"> <li>▪ 2 compute modules</li> <li>▪ CPU for each SAP HANA node: <ul style="list-style-type: none"> <li>• 4 Intel Xeon Platinum 8176/8176M processors 28-core, 2.1 GHz, 165 W Or</li> <li>• 4 Intel Xeon Platinum 8180/8180M processors 28-core, 2.5 GHz, 205 W 4 heat sinks for CPU 0/1 of both compute modules</li> </ul> </li> <li>▪ 3 TB RAM with 48 × 64 GB DIMMs or 24 × 128 GB DIMMs per SAP HANA node</li> </ul>	SAP HANA server
Mellanox Technologies MT27710Family [ConnectX-4 Lx]	4	Mellanox Connectx4 dual port PCIe card for each SAP 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 SAP 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	

Hardware	Quantity	Configuration	Role
Management Server for Solution	1	<ul style="list-style-type: none"> <li>▪ 2 Intel(R) Xeon(R) Gold 6454S 32 cores; 270W; 2.2MHz</li> <li>▪ 2 HPE 32 GB 1Rx4 PC5-4800B-R DIMM</li> <li>▪ 1 Intel Eth Adptr I350T4 OCPv3</li> <li>▪ 1 PCIe — INT E810 10/25GbE 2p SFP28 Adptr</li> <li>▪ HPE NS204i-u Gen11 Ht Plg Boot Opt Dev</li> <li>▪ HPE 480GB NVMe RI M.2 MV SSD (pair)</li> </ul>	Used as the management server, and runs the following:  SAP HANA Studio
Hitachi Virtual Storage Platform One Block High End	1	<ul style="list-style-type: none"> <li>▪ Controller Chassis: 2 (1/ Node)</li> <li>▪ Controller board (CTL): 4 (1 pair/Node)</li> <li>▪ Disk board (DKB): 8 (2 pair/ Node)</li> <li>▪ CPU:8 (1 pair/CTL)</li> <li>▪ Cache flash memory (CFMM): 4 (2/Node)</li> <li>▪ Fan Module (FANM): 8 (2 pair/ Node)</li> <li>▪ Hitachi interconnect edge (HIE): 4 (1 pair/Node)</li> <li>▪ Used Fibre Channel ports: 4 × 32 Gbps</li> <li>▪ MPU: 2 pair</li> <li>▪ Maximum cache: 2048 GB (1024 GB/Node)</li> <li>▪ Used channel board (CHB): 1 pair</li> <li>▪ NVMe SSD drives: 13 × 3.8 TB (one spare drive)</li> <li>▪ Drive box: 4 (1/DKU)</li> </ul>	

Hardware	Quantity	Configuration	Role
Cisco Nexus 93180YC-FX switch	4	<ul style="list-style-type: none"> <li>▪ Two switches with distinct VLANs, each dedicated to cluster FS and SAP HANA inter-cluster network</li> <li>▪ Two switches with one VLAN to provide uplink network to customer network infrastructure</li> </ul>	10 GbE cluster FS and inter-cluster network 10 GbE client network
Cisco Nexus 92348 port switch	1	48 × 1 GbE ports	1 GbE management network

The following table and figure show the server components used for validation.

Feature	Three SAP HANA Nodes (2+1) Scale-out Configuration
Server	3 × DS7040 servers Each DS7040 server acts as one SAP HANA node with the following: <ul style="list-style-type: none"> <li>▪ Module 0 (primary)</li> <li>▪ Module 1 (non-primary)</li> </ul>
Network ports	2 × 2-port Mellanox Connect X-4 LX PCIe cards per compute module in the following locations: <ul style="list-style-type: none"> <li>▪ PCIE Slot 1</li> <li>▪ PCIE Slot 3</li> </ul>
Fibre Channel ports	2 × 2-port LPE31002-M6 16 Gbps PCIe cards on Module 0 for each SAP HANA node in the following locations: <ul style="list-style-type: none"> <li>▪ PCIE Slot 0</li> <li>▪ PCIE Slot 2</li> </ul>
Other interfaces	For all configuration sizes: <ul style="list-style-type: none"> <li>▪ 1 USB 3.0 port</li> <li>▪ KVM connector (VGA, COM, USB 2.0 port)</li> </ul>



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 Emulex LPe31002 M6 16 Gbps, 2-port Fibre Channel PCIe adapters. Use Fibre Channel cables to connect them with the designated VSP One Block High End Fibre Channel ports to eliminate single point of failure and achieve high performance.

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

Dynamic Provisioning Pool Name	Purpose
Microcode Level	A0-05-20-00/05
Cache Memory	2048 GB
Number of Ports	4 × 32 Gbps ports used

Dynamic Provisioning Pool Name	Purpose
Number of RAID Groups in Dynamic Provisioning Pools	3
Number of Dynamic Provisioning Pools	1
Number of LUs in Dynamic Provisioning Pools	4/PG
RAID Group Type	RAID 1 (2D+2D)
Number of Drives per PG	4
Drive Capacity	3.8 TB
Drive Type	NVMe SSD
Drive Box	4
Multi-Pathing Enabled	Yes

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

DB 00&01	0	1	2	3	....	9	10	11	12	13	14	....	19	20	21	22	23
DB 02&03	0	1	2	3	....	9	10	11	12	13	14	....	19	20	21	22	23
DB 04&05	0	1	2	3	....	9	10	11	12	13	14	....	19	20	21	22	23
DB 06&07	0	1	2	3	....	9	10	11	12	13	14	....	19	20	21	22	23
Parity Group, RAID 1 (2D+2D)									PG-1	PG-2	PG-3						

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

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
SAP_HANA_Pool	1	RAID 1 (2D+2D)	00:00:01	1760GB	MPU-010
			00:00:02	1760GB	MPU-020
			00:00:03	1760GB	MPU-110
			00:00:04	1760GB	MPU-120

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
	2	RAID 1 (2D +2D)	00:00:05	1760GB	MPU-010
			00:00:06	1760GB	MPU-020
			00:00:07	1760GB	MPU-110
			00:00:08	1760GB	MPU-120
	3	RAID 1 (2D +2D)	00:00:09	1760GB	MPU-010
			00:00:0A	1760GB	MPU-020
			00:00:0B	1760GB	MPU-110
			00:00:0C	1760GB	MPU-120

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

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment
SAP_HANA_Pool	00:01:00	HANA_OS_N1	100GB	MPU-010
	00:02:00	HANA_OS_N2	100GB	MPU-020
	00:03:00	HANA_OS_N3	100GB	MPU-110
	00:01:01	HANA_SH_1	256GB	MPU-010
	00:02:01	HANA_SH_2	256GB	MPU-020
	00:03:01	HANA_SH_3	256GB	MPU-110
	00:04:01	HANA_SH_4	256GB	MPU-120
	00:01:06	HANA_DATA_N1_1	768 GB	MPU-010
	00:01:07	HANA_DATA_N1_2	768 GB	MPU-020
	00:01:08	HANA_DATA_N1_3	768 GB	MPU-110
	00:01:09	HANA_DATA_N1_4	768 GB	MPU-120
	00:02:06	HANA_DATA_N2_1	768 GB	MPU-010
	00:02:07	HANA_DATA_N2_2	768 GB	MPU-020
	00:02:08	HANA_DATA_N2_3	768 GB	MPU-110

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment
	00:02:09	HANA_DATA_N2_4	768 GB	MPU-120
	00:03:06	HANA_DATA_N3_1	768 GB	MPU-010
	00:03:07	HANA_DATA_N3_2	768 GB	MPU-020
	00:03:08	HANA_DATA_N3_3	768 GB	MPU-110
	00:03:09	HANA_DATA_N3_4	768 GB	MPU-120
	00:01:02	HANA_LOG_N1	512GB	MPU-010
	00:01:03	HANA_LOG_N2	512GB	MPU-020
	00:01:04	HANA_LOG_N3	512GB	MPU-110

## Product descriptions

These products are used in this reference architecture.

### Hitachi Virtual Storage Platform One Block High End

Hitachi Virtual Storage Platform One Block High End (VSP One Block High End or VSP One B85) all-flash NVMe block storage systems deliver ultra-fast, highly reliable, and scalable data access for the most demanding enterprise applications.

Powered by Hitachi-engineered technology, they offer eight nines of availability, next-generation connectivity, comprehensive easy-to-use management, and seamless workload consolidation across open systems and mainframes. With end-to-end NVMe and best-in-class data reduction (4:1 ADR guaranteed), VSP One Block High End greatly enhances performance, scale and resilience in a simple, secure, sustainable way.

VSP One Block High End supports 2, 4, or 6-node configurations, and scales from 4 to 12 tightly coupled controllers with up to 288 NVMe SSDs.

See the *VSP One Block Specification Sheet* on the company website (<http://www.hitachivantara.com>) for more information.

### Hitachi Storage Virtualization Operating System

Hitachi Storage Virtualization Operating System 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 storage capacity, consolidating resources for a higher ROI and providing a high-speed front end to slower, less-predictable storage.

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

**Hitachi Vantara**



Corporate Headquarters  
2535 Augustine Drive  
Santa Clara, CA 95054 USA

[HitachiVantara.com/contact](https://HitachiVantara.com/contact)