SAP HANA Tailored Data Center Integration on Hitachi VSP E590 and VSP E790 with Hitachi Storage Virtualization Operating System RF

Reference Architecture Guide

MK-SL-211-01
January 2024

© 2024 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., 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.

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, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, 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.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at https://www.hitachivantara.com/en-us/company/legal.html or https://knowledge.hitachivantara.com/Documents/Open_Source_Software.

Feedback

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

Revision history

Changes	Date
VSP E590 and VSP 790 certification extension	January 2, 2024
Initial release	December 7, 2020

Reference Architecture Guide

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

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

- VSP E590, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)
- VSP E790, using Non-Volatile Memory express (NVMe) solid state drives (SSDs)

With a 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.

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

- Server blades from Hitachi were used. This solution is also valid for other Hitachi-supported servers.
- Testing showed that Hitachi VSP E590 and VSP E790 with NVMe SSDs meet the SAP enterprise storage certification key performance indicator (KPI) requirements for SAP HANA.
- Scalability and storage KPI testing was performed using SAP HANA Hardware and Cloud Measurement Tool (HCMT). See SAP Note 2493172 - Hardware and Cloud Measurement Tools (SAP user credentials required) for HCMT details.

The following table 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	Minimal Parity Group		HCMT Version
			Data	Log	
VSP E590	NVMe SSDs	19	1 RAID 5 (6D+1P)	1 RAID 5 (6D+1P)	HCMT-052
VSP E790	NVMe SSDs	34*	3 RAID 10 (2D+2D)	2 RAID 10 (2D+2D)	HCMT-052
Note — * Extrapolated scalability for maximum number of nodes using the engineering validation data and information.					

When implementing a SAP HANA TDI infrastructure on Hitachi 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: Since the release of SAP HANA TDI in November 2013, several versions of HCMT have been published. To verify that the hardware configuration of your SAP HANA TDI infrastructure meets the SAP KPIs, you must use the same version of HCMT used during the certification of the hardware, compute servers, and storage system for your tests. SAP Note 2493172 describes how to determine the right version of HCMT for your tests.

Note: Testing of this configuration was performed in a lab environment. Many 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:

- Hitachi Virtual Storage Platform E590 using NVMe SSDs
- Hitachi Virtual Storage Platform E790 using NVMe SSDs

The following table lists VSP E590 and VSP E790 features.

Item			Specifications
System	SVOS Version		Hitachi Storage Virtualization Operating System v9.5
	Max. Drives	NVMe SSD	24
	Drive Options	NVMe SSD	1.9 TB, 3.8 TB, 7.6 TB, 15.3 TB
	RAID Level		RAID 10, RAID 5, RAID 6
	RAID group configuration	RAID 10	2D+2D
		RAID 5	3D+1P, 4D+1P, 6D+1P, 7D +1P
		RAID 6	6D+2P, 12D+2P, 14D+2P
	Max. number of LDEVs		32768 (VSP E590), 49152 (VSP E790)
	Max. storage capacity		0.33 PB
	Max. external configuration		128 PB (VSP E590), 192 PB (VSP E790)
Memory	Cache memory capacity		768 GB
Device I/F	Supported channel type	Fibre Channel	24
		Data transfer rate	8 Gbps, 16 Gbps, 32 Gbps
	Max. number of CHB		6
Non-stop maintenance	Control PCB		Supported
	Cache memory		Supported
	Cache flash memory		Supported
	Power supply, fan		Supported
	Microcode		Supported
	Flash drive		Supported
	Flash module drive		Supported

The following figure shows a server-to-storage Fibre Channel switch indirect connection configuration for this solution. This example uses VSP E790 with 32 SAP HANA systems and NVMe SSDs.

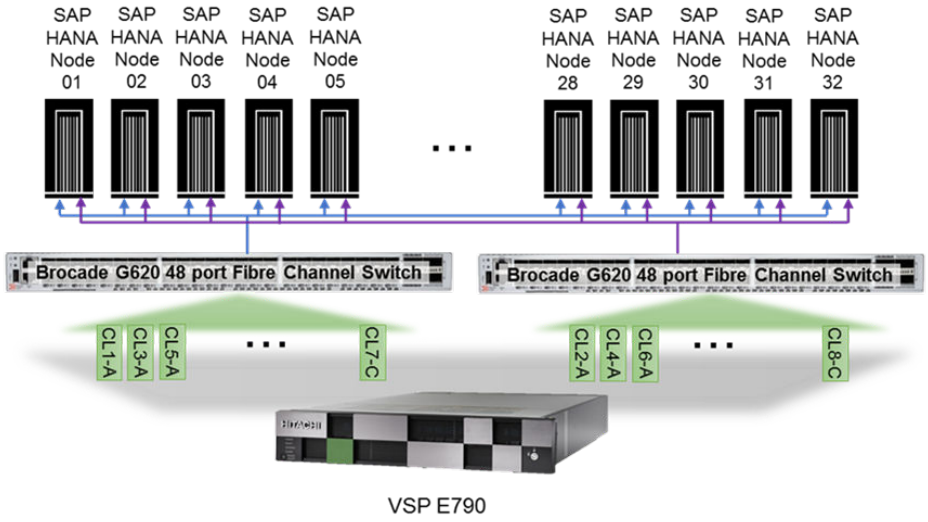

Key solution elements

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 E590 and VSP E790 using NVMe SSDs and passed SAP enterprise storage KPIs,

	Tested Number of SAP HANA Nodes	Storage Components	Fibre Channel Switch
VSP E590 (NVMe SSD)	19	▪ CTL: 1 pair ▪ 32 Gbps 4-port CHB: 3 pairs ▪ MPU: 1 pair ▪ Cache: 768 GB ▪ NVMe SSDs: 24 × 1.9 TB	48-port Brocade G620 Fibre Channel Switch: 1 pair
VSP E790 (NVMe SSD)	32	▪ CTL: 1 pair ▪ 32 Gbps 4-port CHB: 3 pairs ▪ MPU: 1 pair	48-port Brocade G620 Fibre Channel Switch: 1 pair

	Tested Number of SAP HANA Nodes	Storage Components	Fibre Channel Switch
		▪ Cache: 768 GB ▪ NVMe SSDs: 24 × 1.9 TB	

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

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

Optional Item	Components
Management Server, using this server	Hitachi Advanced Server DS120
1 GbE Management Network Switch, using either one of these switches	Brocade ICX 6430-48 port
	Cisco Nexus 3048
10 GbE Network Switch, using either one of these switches	Brocade VDX 6740-48 port
	Cisco 93180YC-EX

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 SPS05.

Solution design

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

- Hitachi Virtual Storage Platform E590 using NVMe SSDs
- Hitachi Virtual Storage Platform E790 using NVMe SSDs

The configuration uses RAID 5 (6D+1P).

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

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP E590 or VSP E790	Port Speed	Port Security
Node1	Port 0	16 Gbps	SW-1-P0	SW-1-P32	1A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P0	SW-2-P32	2A-Host Group 1	32 Gbps	Enabled
Node2	Port 0	16 Gbps	SW-1-P1	SW-1-P32	1A-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P1	SW-2-P32	2A-Host Group 2	32 Gbps	Enabled
Node3	Port 0	16 Gbps	SW-1-P2	SW-1-P33	3A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P2	SW-2-P33	4A-Host Group 1	32 Gbps	Enabled
Node4	Port 0	16 Gbps	SW-1-P3	SW-1-P33	3A-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P3	SW-2-P33	4A-Host Group 2	32 Gbps	Enabled
Node5	Port 0	16 Gbps	SW-1-P4	SW-1-P34	5A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P4	SW-2-P34	6A-Host Group 1	32 Gbps	Enabled
Node6	Port 0	16 Gbps	SW-1-P5	SW-1-P34	5A-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P5	SW-2-P34	6A-Host Group 2	32 Gbps	Enabled
Node7	Port 0	16 Gbps	SW-1-P6	SW-1-P35	7A-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P6	SW-2-P35	8A-Host Group 1	32 Gbps	Enabled
Node8	Port 0	16 Gbps	SW-1-P7	SW-1-P35	7A-Host Group 2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP E590 or VSP E790	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P7	SW-2-P35	8A-Host Group 2	32 Gbps	Enabled
Node9	Port 0	16 Gbps	SW-1-P8	SW-1-P36	1B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P8	SW-2-P36	2B-Host Group 1	32 Gbps	Enabled
Node10	Port 0	16 Gbps	SW-1-P9	SW-1-P36	1B-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P9	SW-2-P36	2B-Host Group 2	32 Gbps	Enabled
Node11	Port 0	16 Gbps	SW-1-P10	SW-1-P37	3B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P10	SW-2-P37	4B-Host Group 1	32 Gbps	Enabled
Node12	Port 0	16 Gbps	SW-1-P11	SW-1-P37	3B-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P11	SW-2-P37	4B-Host Group 2	32 Gbps	Enabled
Node13	Port 0	16 Gbps	SW-1-P12	SW-1-P38	5B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P12	SW-2-P38	6B-Host Group 1	32 Gbps	Enabled
Node14	Port 0	16 Gbps	SW-1-P13	SW-1-P38	5B-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P13	SW-2-P38	6B-Host Group 2	32 Gbps	Enabled
Node15	Port 0	16 Gbps	SW-1-P14	SW-1-P39	7B-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P14	SW-2-P39	8B-Host Group 1	32 Gbps	Enabled
Node16	Port 0	16 Gbps	SW-1-P15	SW-1-P39	7B-Host Group 2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP E590 or VSP E790	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P15	SW-2-P39	8B-Host Group 2	32 Gbps	Enabled
Node17	Port 0	16 Gbps	SW-1-P16	SW-1-P40	1C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P16	SW-2-P40	2C-Host Group 1	32 Gbps	Enabled
Node18	Port 0	16 Gbps	SW-1-P17	SW-1-P40	1C-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P17	SW-2-P40	2C-Host Group 2	32 Gbps	Enabled
Node19	Port 0	16 Gbps	SW-1-P18	SW-1-P41	3C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P18	SW-2-P41	4C-Host Group 1	32 Gbps	Enabled
Node20	Port 0	16 Gbps	SW-1-P19	SW-1-P41	3C-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P19	SW-2-P41	4C-Host Group 2	32 Gbps	Enabled
Node21	Port 0	16 Gbps	SW-1-P20	SW-1-P42	5C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P20	SW-2-P42	6C-Host Group 1	32 Gbps	Enabled
Node22	Port 0	16 Gbps	SW-1-P21	SW-1-P42	5C-Host Group 2	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P21	SW-2-P42	6C-Host Group 2	32 Gbps	Enabled
Node23	Port 0	16 Gbps	SW-1-P22	SW-1-P43	7C-Host Group 1	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P22	SW-2-P43	8C-Host Group 1	32 Gbps	Enabled
Node24	Port 0	16 Gbps	SW-1-P23	SW-1-P43	7C-Host Group 2	32 Gbps	Enabled

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group		
	Port Name	Port Speed	Host	Storage	VSP E590 or VSP E790	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P23	SW-2-P43	8C-Host Group 2	32 Gbps	Enabled
Node25	Port 0	16 Gbps	SW-1-P24	SW-1-P32	1A-Host Group 3	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P24	SW-2-P32	2A-Host Group 3	32 Gbps	Enabled
Node26	Port 0	16 Gbps	SW-1-P25	SW-1-P33	3A-Host Group 3	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P25	SW-2-P33	4A-Host Group 3	32 Gbps	Enabled
Node27	Port 0	16 Gbps	SW-1-P26	SW-1-P34	5A-Host Group 3	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P26	SW-2-P34	6A-Host Group 3	32 Gbps	Enabled
Node28	Port 0	16 Gbps	SW-1-P27	SW-1-P35	7A-Host Group 3	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P27	SW-2-P35	8A-Host Group 3	32 Gbps	Enabled
Node29	Port 0	16 Gbps	SW-1-P28	SW-1-P36	1B-Host Group 3	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P28	SW-2-P36	2B-Host Group 3	32 Gbps	Enabled
Node30	Port 0	16 Gbps	SW-1-P29	SW-1-P37	3B-Host Group 3	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P29	SW-2-P37	4B-Host Group 3	32 Gbps	Enabled
Node31	Port 0	16 Gbps	SW-1-P30	SW-1-P38	5B-Host Group 3	32 Gbps	Enabled
	Port 1	16 Gbps	SW-2-P30	SW-2-P38	6B-Host Group 3	32 Gbps	Enabled
Node32	Port 0	16 Gbps	SW-1-P31	SW-1-P39	7B-Host Group 3	32 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 E590 or VSP E790	Port Speed	Port Security
	Port 1	16 Gbps	SW-2-P31	SW-2-P39	8B-Host Group 3	32 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 utilizes the following two dynamic provisioning pools created with Hitachi Dynamic Provisioning for the storage layout. This ensures maximum utilization and optimization at a lower cost than other solutions.

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

The validated dynamic provisioning pool layout options with minimal disks and storage cache on Hitachi VSP E590 and VSP E790 storage are listed in the following table.

Storage	Cache	Validated Node Number	Number of Parity Groups in OS_SH_Data_Pool	Number of Parity Groups in LG_Pool
			RAID 5 (6D+1P)	RAID 5 (6D +1P)
VSP E590 (with NVMe SSD)	768 GB	19	1	1
VSP E790 (with NVMe SSD)	768 GB	15	1	1
		32	2	1

In the example, for Hitachi Virtual Storage Platform E590, you need the following:

- A minimum of one parity group is needed for OS_SH_DT_ Pool to fit 19 SAP HANA production nodes on one VSP E590 using RAID 5 (6D+1P).
- A minimum of one parity group is needed for LG_ Pool to fit 19 SAP HANA production nodes on one VSP E590 using RAID 5 (6D+1P).

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

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

- Determine the minimum sizes for operating system, data, log, and HANA shared using these formulas in the SAP white paper *SAP HANA Storage Requirements* from SAP Note 1900823:
  - Every SAP HANA node requires approximately 100 GB capacity for the operating system.
  - /hana/shared size uses the following formulas:

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

    Multi-node (scale-out) — Size = 1 × RAM of every 4 worker nodes
  - Data size requires at least 1 × RAM of each SAP HANA node
  - Log size uses the following formulas:
    - Systems with equal to or less than 512 GB memory — Size = 1/2 × RAM
    - Systems with greater than 512 GB memory — Size = 512 GB

    Note: Hitachi recommends these minimum sizes, as calculated using formulas. The SAP quicksizer and sizing reports might call for smaller sizes.
- Provision the storage:
  - Create two dynamic provisioning pools for the three SAP HANA systems on the storage:
    - Use OS_SH_DT_Pool to provision the operating system volume, SAP HANA shared volume, and data volume.
    - Use LG_Pool to provision the log volume.
  - For NVMe SSDs, create the parity groups first, using the example in the following table for Hitachi VSP E590 or VSP E790, using the RAID 5 storage design.

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
OS_SH_Data_Pool	1	RAID 5 (6D +1P) on 1.9 TB SSD	00:00:01	OS_SH_DT_DPVOL_1	2640 GB	MPU-10
			00:00:02	OS_SH_DT_DPVOL_2	2640 GB	MPU-20
			00:00:03	OS_SH_DT_DPVOL_3	2640 GB	MPU-10

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
			00:00:04	OS_SH_DT_DPVOL_4	2640 GB	MPU-20
Log_Pool	4	RAID 5 (6D +1P) on 1.9 TB SSD	00:00:13	LG_DPVOL_1	2640 GB	MPU-10
			00:00:14	LG_DPVOL_2	2640 GB	MPU-20
			00:00:15	LG_DPVOL_3	2640 GB	MPU-10
			00:00:16	LG_DPVOL_4	2640 GB	MPU-20

- Assign all LDEVs to the dedicated pool as shown in the examples in Table 7 for VSP E590 or VSP E790.
- Create virtual volumes (vVols) for the operating system, SAP HANA shared, log, and data volumes. The following table shows examples for SAP HANA systems with 384 GB, 768 GB, 1536 GB, and 3072 GB of memory.

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

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment	System Memory
	00:03:01	HANA_SH_N3	1536 GB	MPU-10	1536 GB
	00:04:01	HANA_SH_N4	3072 GB	MPU-20	3072 GB
	00:01:06	HANA_DATA_N1_1	96 GB	MPU-10	384 GB
	00:01:07	HANA_DATA_N1_2	96 GB	MPU-20	
	00:01:08	HANA_DATA_N1_3	96 GB	MPU-10	
	00:01:09	HANA_DATA_N1_4	96 GB	MPU-20	
	00:02:06	HANA_DATA_N2_1	192 GB	MPU-10	768 GB
	00:02:07	HANA_DATA_N2_2	192 GB	MPU-20	
	00:02:08	HANA_DATA_N2_3	192 GB	MPU-10	
	00:02:09	HANA_DATA_N2_4	192 GB	MPU-20	
	00:03:06	HANA_DATA_N3_1	384 GB	MPU-10	1536 GB
	00:03:07	HANA_DATA_N3_2	384 GB	MPU-20	
	00:03:08	HANA_DATA_N3_3	384 GB	MPU-10	
	00:03:09	HANA_DATA_N3_4	384 GB	MPU-20	
	00:04:06	HANA_DATA_N4_1	768 GB	MPU-10	3072 GB
	00:04:07	HANA_DATA_N4_2	768 GB	MPU-20	
	00:04:08	HANA_DATA_N4_3	768 GB	MPU-10	

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment	System Memory
	00:04:09	HANA_DATA_N4_4	768 GB	MPU-20	
Log_Pool	00:01:02	HANA_LOG_N1_1	48 GB	MPU-10	384 GB
	00:01:03	HANA_LOG_N1_2	48 GB	MPU-20	
	00:01:04	HANA_LOG_N1_3	48 GB	MPU-10	
	00:01:05	HANA_LOG_N1_4	48 GB	MPU-20	
	00:02:02	HANA_LOG_N2_1	96 GB	MPU-10	768 GB
	00:02:03	HANA_LOG_N2_2	96 GB	MPU-20	
	00:02:04	HANA_LOG_N2_3	96 GB	MPU-10	
	00:02:05	HANA_LOG_N2_4	96 GB	MPU-20	
	00:03:02	HANA_LOG_N3_1	128 GB	MPU-10	1536 GB
	00:03:03	HANA_LOG_N3_2	128 GB	MPU-20	
	00:03:04	HANA_LOG_N3_3	128 GB	MPU-10	
	00:03:05	HANA_LOG_N3_4	128 GB	MPU-20	
	00:04:02	HANA_LOG_N4_1	128 GB	MPU-10	3072 GB
	00:04:03	HANA_LOG_N4_2	128 GB	MPU-20	
	00:04:04	HANA_LOG_N4_3	128 GB	MPU-10	
	00:04:05	HANA_LOG_N4_4	128 GB	MPU-20	

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

The following table 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	00:01:00	HANA_OS_N1
0001	00:01:01	HANA_SH_N1
0002	00:01:02	HANA_LOG_N1_1
0003	00:01:03	HANA_LOG_N1_2
0004	00:01:04	HANA_LOG_N1_3
0005	00:01:05	HANA_LOG_N1_4
0006	00:01:06	HANA_DATA_N1_1
0007	00:01:07	HANA_DATA_N1_2
0008	00:01:08	HANA_DATA_N1_3
0009	00:01:09	HANA_DATA_N1_4

SAP HANA configuration

This 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 if using internal disks for the boot volume)
- SAP HANA data volume
- SAP HANA log volume
- SAP HANA shared volume

  For a scale-out solution, Hitachi uses multiple paths to the disks with Hitachi NAS Platform as NFS for `/hana/shared`. This solution is also valid for other file systems listed in SAP Note 405827 for SAP HANA.

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-length 0	
	uid_attribute	ID_SERIAL	
	failback	immediate	
	rr_weight	uniform	
	rr_min_io_rq	128	1
	features	0	
	no_path_retry	5	

SAP HANA persistent storage volume configuration

For both operating systems, SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux 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 1024
--name <lv name> <volume group>
```

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

	Number of LUNs	LV Name	VG Name
DATA	4 - following this reference architecture, or the number of assigned LUNs	lvdata	▪ Scale-up: vgdata ▪ Scale-out: vg<SID>data<node number>
LOG	4 - following this reference architecture, or the number of assigned LUNs	lvlog	▪ Scale-up: vglog ▪ Scale-out: vg<SID>log<node number>
SHARED (only applicable for scale-up SAP HANA deployments)	1 - following this reference architecture, or the number of assigned LUNs	lvshared	▪ Scale-up: vgshared

Note: If you only use 1 LUN to create the logical volumes for data, log, or shared, the options `--stripes` and `--stripesize` are not needed.

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

	System Type	Create Options	Mount Options	Mount Point
DATA	Scale-up	-F	inode64, nobarrier*	/hana/data
	Scale-out		N/A	/hana/data/<SID>
LOG	Scale-up	-F	inode64, nobarrier*	/hana/log
	Scale-Out		N/A	/hana/log/<SID>
SHARED	Scale-up	-F	inode64, nobarrier*	/hana/shared
	Scale-out	N/A	vers=3,proto=tcp,hard,intr,timeo=600,retrans=2,wsize=65536,rsize=65536 (These options were tested and verified on Hitachi NAS Platform.)	/hana/shared/<SID>
Note — * The nobarrier option is no longer valid from RHEL8.1.				

To create a file system, use the following command:

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

For example:

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

SAP HANA persistent storage volume configuration for scale-up deployments

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

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 once using the following command:

   ```
   mount -a
   ```

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

SAP HANA persistent storage volume configuration for scale-out deployments

This 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 HNAS or GFS2.
   - Hitachi NAS Platform (HNAS):
     a. Persist the SAP HANA shared file systems by adding the following entry to `/etc/fstab`:

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

     b. Mount the file system:

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

   - Global File System 2 (GFS2):

     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 25)
- SAP HANA software installation (on page 26)

SAP Storage Connector API Fibre Channel Client

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

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

- Database startup
- Failing-over nodes

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

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

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

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

SAP HANA software installation

After configuring the file system for the SAP HANA data volume, log volume, and SAP HANA shared volume, install SAP HANA on servers. Refer to SAP Note 2235581, the attachment SAP_HANA_OS_Release_Support_Matrix.pdf file outlines which combinations of OS RHEL Minor Releases or SLES Support Packages are supported with the various SAP HANA versions. By default, the SAP HANA database, as well as the SAP HANA client, need to be installed.

SAP HANA software installation in a scale-up environment

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

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

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

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

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

Create 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 at https://www.sap.com/documents/2016/06/84ea994f-767c-0010-82c7-eda71af511fa.html 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.

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.

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

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

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

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

```
/hana/shared/HIT/hdblcm/hdblcm –action=add_hosts
\ --sid=HIT
\ --install_hostagent
\ --sapmnt=/hana/shared
\ --addhosts=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 E590 and VSP E790 certified by SAP:

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

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-074.

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 ▪ 4 heat sinks for CPU 0/1 of both compute modules ▪ 3 TB RAM with 48 × 64 GB DIMMs or 24 × 128 GB DIMMs per SAP HANA node	SAP HANA server
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 VSP E790	1	▪ Controller board (CTL): 1 pair ▪ CPU: 4 (1 pair/CTL) ▪ Cache flash memory (CFM): 2 (1/CTL)	

Hardware	Quantity	Configuration	Role
		▪ Backup module (BKM): 2 (1/CTL) ▪ Battery: 2 (1/BKM) ▪ Fan: 12 (6/CTL) ▪ MPU: 1 pair ▪ Maximum cache: 768 GB ▪ Used 32 Gbps 4-port channel board (CHB): 1 pair ▪ Used Fibre Channel ports: 6 (six ports used by three SAP HANA nodes (2/node) ▪ NVMe SSD drives: 15 × 1.9 TB (one spare drive)	
Mellanox Technologies MT27710 Family [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 LPe31002 M6 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	1 GbE 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 DS7040 servers
Server Modules	Each 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
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 E790 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)

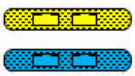

The following table shows the storage configuration and provisioning used for the 2+1 scale-out test.

Item	
Microcode Level	93-03-21-40/00
Cache Memory	768 GB
Number of Ports	6 × 32 Gbps ports used
Number of RAID Groups in Dynamic Provisioning Pool	1
Number of Dynamic Provisioning Pools	2
Number of LUs in Dynamic Provisioning Pools	4/PG
RAID Group Type	RAID 5 (6D+1P)
Number of Drives per PG	7
Drive Capacity	1.9 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	….	5	**6**	7	8	….	13	14	15	16	….	22	23
Parity Group, RAID 5 (6D +1P)	PG-1					PG-2					N/A				

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

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size*	MPU Assignment
OS_SH_DT_Pool	1	RAID 5 (6D+1P) on 1.9 TB SSDs	00:00:01	OS_SH_DT_1	2460 GB	MPU-10
			00:00:02	OS_SH_DT_2	2460 GB	MPU-20
			00:00:03	OS_SH_DT_3	2460 GB	MPU-10
			00:00:04	OS_SH_DT_4	2460 GB	MPU-20
LG_Pool	2	RAID 5 (6D+1P) on 1.9 TB SSDs	00:00:05	LG_1	2460 GB	MPU-10
			00:00:06	LG_2	2460 GB	MPU-20
			00:00:07	LG_3	2460 GB	MPU-10
			00:00:08	LG_4	2460 GB	MPU-20
* The LDEV size could be auto-sized when using the `raidcom` command to create an HDP Pool.						

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

Dynamic Provisioning Pool	vVol ID	vVol Name	vVol Size	MPU Assignment
OS_SH_DT_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-10
	00:02:00	HANA_OS_N2	100 GB	MPU-20
	00:03:00	HANA_OS_N3	100 GB	MPU-10
	00:01:01	HANA_SH_1	256 GB	MPU-10
	00:02:01	HANA_SH_2	256 GB	MPU-20
	00:03:01	HANA_SH_3	256 GB	MPU-10
	00:03:01	HANA_SH_4	256 GB	MPU-20
	00:01:06	HANA_DATA_N1_1	600 GB	MPU-10
	00:01:07	HANA_DATA_N1_2	600 GB	MPU-20
	00:01:08	HANA_DATA_N1_3	600 GB	MPU-10
	00:01:09	HANA_DATA_N1_4	600 GB	MPU-20
	00:02:06	HANA_DATA_N2_1	600 GB	MPU-10
	00:02:07	HANA_DATA_N2_2	600 GB	MPU-20
	00:02:08	HANA_DATA_N2_3	600 GB	MPU-10
	00:02:09	HANA_DATA_N2_4	600 GB	MPU-20
	00:03:06	HANA_DATA_N3_1	600 GB	MPU-10
	00:03:07	HANA_DATA_N3_2	600 GB	MPU-20
	00:03:08	HANA_DATA_N3_3	600 GB	MPU-10
	00:03:09	HANA_DATA_N3_4	600 GB	MPU-20
LG_Pool	00:01:02	HANA_LOG_N1	512 GB	MPU-10
	00:02:02	HANA_LOG_N2	512 GB	MPU-10
	00:03:02	HANA_LOG_N3	512 GB	MPU-10

Product descriptions

The following products were used in this reference architecture.

Hitachi Virtual Storage Platform E series family

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

- Improve IT agility: “Faster-to-market” for IT projects with proven high-performance infrastructure. Brings “enterprise-class” features and benefits to customers of all sizes whose business is outpacing their existing infrastructure and supports modern business processes like DevOPs.
- Financial elasticity that aligns costs to business goals, growth, and use: Customers can “have it their way” with purchase, lease, or cloud-like consumption models.
- Improved workforce efficiency: a better digital experience which boosts customer satisfaction (with both internal LOBs and end-users) and increases business productivity and profitability.

Servers for SAP HANA TDI environments

Only compute servers listed in the official SAP HANA Hardware Directory are supported. Find a description of all certified servers and enterprise storage solutions in Certified and Supported SAP HANA Hardware. For more information on SAP HANA TDI, see *SAP HANA Tailored Data Center Integration - Frequently Asked Questions* at https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact