

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform G Series and VSP F Series with Hitachi Storage Virtualization Operating System

Reference Architecture Guide

By Hitachi Vantara

May 2020

Feedback

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

Revision History

Revision	Changes	Date
MK-SL-059-00	Initial release	May 2, 2018
MK-SL-059-01	Update CPU information and multipath information in document.	May 21, 2020

Table of Contents

Solution Overview	3
Key Solution Elements	6
Hardware Elements	7
Software Elements	9
Solution Design	10
Fibre Channel Architecture	10
Storage Architecture	14
SAP HANA Configuration	21
SAP Storage Connector API Fibre Channel Client	25
SAP HANA Software Installation	26
Engineering Validation	34

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform G Series and VSP F Series with Hitachi Storage Virtualization Operating System

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 models of Hitachi Virtual Storage Platform (VSP):

- VSP G350, using serial attached SCSI (SAS) hard disk drives (HDD)
- VSP G350 and VSP F350, using solid state drives (SSD)
- VSP G370, using serial attached SCSI (SAS) hard disk drives (HDD)
- VSP G370 and VSP F370, using solid state drives (SSD)
- VSP G700 using serial attached SCSI (SAS) hard disk drives (HDD)
- VSP G700 and VSP F700 using solid state drives (SSD)
- VSP G900 and VSP F900 using solid state drives (SSD)

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

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

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

- Server blades from Hitachi were used. This solution is also valid for other [Hitachi-supported servers](#).
- Testing showed that the following Hitachi storage meets the SAP enterprise storage certification key performance indicator (KPI) requirements for SAP HANA:
 - With serial attached SCSI (SAS) HDDs:
 - Virtual Storage Platform G350
 - Virtual Storage Platform G370
 - Virtual Storage Platform G700

- With SSDs:
 - Virtual Storage Platform G350
 - Virtual Storage Platform F350
 - Virtual Storage Platform G370
 - Virtual Storage Platform F370
 - Virtual Storage Platform G700
 - Virtual Storage Platform F700
 - Virtual Storage Platform G900
 - Virtual Storage Platform F900
- Scalability and storage KPI testing was performed using [SAP HANA Hardware Configuration Check Tool \(HWCCT\)](#). See [SAP Note 1943937 - Hardware Configuration Check Tool - Central Note](#) (SAP user credentials required) for details about HWCCT.
- Figure 1 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 HWCCT in the SAP HANA systems during testing, there are no other systems or applications using the storage, with the storage fully dedicated to this testing.

Note – If you need to have mixed applications running in parallel on the same storage, SAP HANA and non-SAP HANA, you must perform testing for the expected workload using the SAP HANA Hardware Configuration Tool (HWCCT) to meet SAP KPIs.

TABLE 1. MAXIMUM SUPPORTED NUMBER OF SAP HANA SYSTEMS

Storage	Drive Type	Maximum HANA Systems	Minimal Parity Group		HWCCT Version
			Data	Log	
VSP G350, G370	SAS HDDs	3	3 RAID-6 (14D+2P)	1 RAID-6 (6D+2P)	HWCCT-212
VSP G350, F350, G370, F370	SSDs	16	3 RAID-10 (2D+2D)	3 RAID-10 (2D+2D)	HWCCT-122
VSP G700	SAS HDDs	13	6 RAID-6 (14D+2P)	3 RAID-6 (6D+2P)	HWCCT-122
VSP G700, F700	SSDs	34 (Extrapolated scalability for maximum number of nodes using the engineering validation data and information.)	4 RAID-10 (2D+2D)	5 RAID-10 (2D+2D)	HWCCT-122
VSP G900, F900	SSDs	40 (Extrapolated scalability for maximum number of nodes using the engineering validation data and information.)	4 RAID-10 (2D+2D)	4 RAID-10 (2D+2D)	HWCCT-122

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

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

Note – Testing of this configuration was performed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that matches your production environment before your production implementation of this solution.

Solution Overview

This reference architecture guide provides example configurations tested in the Hitachi Vantara lab using SAS HDDs or 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 G350 using SAS HDDs or SSDs
- Virtual Storage Platform F350 using SSDs
- Virtual Storage Platform G370 using SAS HDDs or SSDs
- Virtual Storage Platform F370 using SSDs
- Virtual Storage Platform G700 using SAS HDDs or SSDs
- Virtual Storage Platform F700 using SSDs
- Virtual Storage Platform G900 using SSDs
- Virtual Storage Platform F900 using SSDs

Figure 1 shows Virtual Storage Platform models in the order of increasing performance and storage capacity.

Figure 1




Table 2 highlights the features of different Hitachi Virtual Storage Platform models.

TABLE 2. HIGHLIGHTED FEATURES COMPARISON OF HITACHI VIRTUAL STORAGE PLATFORM

Item		Specifications								
	Type	VSP G350	VSP F350	VSP G370	VSP F370	VSP G700	VSP F700	VSP G900	VSP F900	
System	SVOS Version	Hitachi Storage Virtualization Operating System v8.x								
	Maximum Drives	2U SFF	192		288		864		1152	
		4U LFF	252		372		1200		1440	
		FMD	84		132		432		576	
	Drive Options	SFF Hard Disk Drive	600 GB, 1.2 TB, 2.4 TB	n/a	600 GB, 1.2 TB, 2.4 TB	n/a	600 GB, 1.2 TB, 2.4 TB	n/a	600 GB, 1.2 TB, 2.4 TB	n/a
		LFF Hard Disk Drive	2.4 TB	n/a	2.4 TB	n/a	2.4 TB	n/a	2.4 TB	n/a
		SFF Solid State Drive	960 GB, 1.9 TB, 3.8 TB, 7.6 TB							
		Flash Module Drive	n/a	3.5 TB, 7 TB, 14 TB	n/a	3.5 TB, 7 TB, 14 TB	n/a	3.5 TB, 7 TB, 14 TB	n/a	3.5 TB, 7 TB, 14 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, 8D+2P, 10D+2P, 12D+2P, 14D+2P							
	Maximum number of LDEVs		16384		32767		49152		65280	
Maximum storage capacity		1,452 TB		2,178 TB		6,533 TB		8,711 TB		
Maximum external configuration		64 PB		128 PB		192 PB		255 PB		
Memory	Cache memory capacity	128 GB		256 GB		512 GB		1024 GB		
Storage I/F	DKC to drive interface	SAS/Dual Port								
	Data transfer rate	12 Gb/s								
	Maximum number of DKB	n/a		n/a		4		8		

TABLE 2. HIGHLIGHTED FEATURES COMPARISON OF HITACHI VIRTUAL STORAGE PLATFORM (CONTINUED)

Item		Specifications							
	Type	VSP G350	VSP F350	VSP G370	VSP F370	VSP G700	VSP F700	VSP G900	VSP F900
Device I/F	Supported channel type	Fibre Channel	16		16		48		48
		Data transfer rate	8 Gb/s, 16 Gb/s, 32 Gb/s						
	Maximum number of CHB	4		4		12		12	
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							

Figure 2 on page 5 shows a server-to-storage direct connection configuration of this solution. This example uses Hitachi Virtual Storage Platform G350 with three SAP HANA systems and SAS HDDs.

Figure 3 on page 6 shows a server-to-storage Fibre Channel switch indirect connection configuration of this solution. This example uses Virtual Storage Platform G350 or Virtual Storage Platform F350 with 16 SAP HANA systems and SSDs.

Figure 2





Figure 3



Key Solution Elements

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

Hardware Elements

Table 3, "Hardware Elements," on page 7 lists the hardware used to test the scalability of maximum active nodes on Hitachi Virtual Storage Platform G350, F350, G700, F700, and F900 using SAS HDDs or SSDs. Table 4, "Optional Hardware For Use With Any Storage System," on page 8 shows the optional hardware that can be used with any of these storage systems in your environment.

TABLE 3. HARDWARE ELEMENTS

	Certified Maximum Number of Nodes^a	Storage Components	Fibre Channel Switch
VSP G350, G370 (SAS HDD)	3	<ul style="list-style-type: none"> ▪ CTL: 1 pair ▪ 16 Gb/s 4-port CHB: 1 pair ▪ MPU: 1 pair ▪ Cache: 128 GB ▪ HDDs: 56 × 600 GB 	N/A
VSP G350, F350, G370, and F370 (SSD)	16	<ul style="list-style-type: none"> ▪ CTL: 1 pair ▪ 32 Gb/s 4-port CHB: 1 pair ▪ MPU: 1 pair ▪ Cache: 128 GB ▪ SSDs: 24 × 1.9 TB 	48-port Brocade G620 Fibre Channel Switch: 1 pair
VSP G700 (SAS HDD)	13	<ul style="list-style-type: none"> ▪ CTL: 1 pair ▪ 16 Gb/s 4-port CHB: 4 pair ▪ MPU: 1 pair ▪ Cache: 512 GB ▪ HDDs: 120 × 600 GB 	N/A
VSP G700 and F700 (SSD)	34	<ul style="list-style-type: none"> ▪ CTL: 1 pair ▪ 32 Gb/s 4-port CHB: 4 pair ▪ MPU: 1 pair ▪ Cache: 512 GB ▪ SSDs: 36 × 1.9 TB 	48-port Brocade G620 Fibre Channel Switch: 1 pair
VSP G900 and F900 (SSD)	40	<ul style="list-style-type: none"> ▪ CTL: 1 pair ▪ 32 Gb/s 4-port CHB: 4 pair ▪ MPU: 1 pair ▪ Cache: 1024 GB ▪ SSDs: 32 × 1.9 TB 	48-port Brocade G620 Fibre Channel Switch: 1 pair

a. View a list of certified configurations in the [Certified and Supported SAP HANA Hardware Directory](#).

TABLE 4. OPTIONAL HARDWARE FOR USE WITH ANY STORAGE SYSTEM

Optional Item	Components
Management Server, using either one of these servers	Rack mounted server for solutions, 2U
	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

Hitachi Virtual Storage Platform F Series Family

Use [Hitachi Virtual Storage Platform F series family](#) storage for a flash-powered cloud platform for your mission critical applications. This storage meets demanding performance and uptime business needs. Extremely scalable, its 4.8 million random read IOPS allows you to consolidate more applications for more cost savings.

Hitachi Storage Virtualization Operating System RF is at the heart of the Virtual Storage Platform F series family. It provides storage virtualization, high availability, flash optimized performance, quality of service controls, and advanced data protection. This proven, mature software provides common features, management, and interoperability across the Hitachi portfolio. This means you can reduce migration efforts, consolidate assets, reclaim space, and extend life.

Reduce risks and solve problems faster. Integrated power analytics and automation features bring artificial intelligence to your data center. Cloud-accessible monitoring tools give your product support experts access wherever they have an internet connection for fast troubleshooting and remediation.

Hitachi Virtual Storage Platform G Series Family

The [Hitachi Virtual Storage Platform G series family](#) enables the seamless automation of the data center. It has a broad range of efficiency technologies that deliver maximum value while making ongoing costs more predictable. You can focus on strategic projects and to consolidate more workloads while using a wide range of media choices.

The benefits start with Hitachi Storage Virtualization Operating System RF. This includes an all new enhanced software stack that offers up to three times greater performance than our previous midrange models, even as data scales to petabytes.

Virtual Storage Platform G series offers support for containers to accelerate cloud-native application development. Provision storage in seconds, and provide persistent data availability, all the while being orchestrated by industry leading container platforms. Moved these workloads into an enterprise production environment seamlessly, saving money while reducing support and management costs.

Engineering validation for this solution used Virtual Storage Platform G350, VSP F350, VSP G700, VSP F700 and VSP F900.

Server for a SAP HANA TDI Environment

Only servers certified for use with SAP HANA can be used in a SAP HANA TDI environment. This follows the exact same bill of materials as the certified SAP HANA appliance server, but without the storage or local disks.

Find a description of all certified servers and enterprise storage solutions in [Certified and Supported SAP HANA Hardware](#). For more information on SAP HANA TDI, consult [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#).

Software Elements

This environment uses the following software.

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

Hitachi Storage Virtualization Operating System RF

[Hitachi Storage Virtualization Operating System RF](#) (SVOS RF) spans and integrates multiple platforms. 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.

Storage Virtualization Operating System RF uses Hitachi Dynamic Provisioning (HDP) to provide wide striping and thin provisioning. Dynamic Provisioning provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) without initially allocating any physical space. Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs).

To enable essential management and optimization functions, this solution uses Hitachi Storage Navigator (SN), a part of Server Virtualization Operating System RF. Storage Navigator runs on most browsers. A command line interface is available.

Solution Design

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

- Hitachi Virtual Storage Platform G350, VSP G370, and VSP G700 using SAS HDD
- Hitachi Virtual Storage Platform G350, VSP F350, G370, VSP F370, VSP G700, VSP F700, VSP G900, and VSP F900 using SSD.

In the solution in this reference architecture guide, the following is true:

- The configuration for VSP G350 with HDD is applicable to VSP G370 with HDD.
- The configuration for VSP F350 with SSD is applicable to VSP G350, VSP G370, and VSP F370 with SSD.
- The configuration for VSP F700 with SSD is applicable to VSP G700 with SSD.
- The configuration for VSP F900 with SSD is applicable to VSP G900 with SSD.

Fibre Channel Architecture

For each SAP HANA node, there are two dedicated Fibre Channel ports on the storage. The two Fibre Channel cables connecting 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

Table 5 shows storage target port mapping with hosts for validated number of HANA systems using SAS HDDs.

TABLE 5. STORAGE PORT MAPPING FOR VALIDATED SAP HANA NODES USING SAS HDDS

SAP HANA Node	HBA Port	Virtual Storage Platform Target Port		
		VSP G350	VSP G700	Port Speed
Node1	Port 0	1A	1A	16 Gb/s
	Port 1	2A	2A	16 Gb/s
Node2	Port 0	3A	3A	16 Gb/s
	Port 1	4A	4A	16 Gb/s
Node3	Port 0	5A	5A	16 Gb/s
	Port 1	6A	6A	16 Gb/s
Node4	Port 0	N/A	7A	16 Gb/s
	Port 1	N/A	8A	16 Gb/s
Node5	Port 0	N/A	1B	16 Gb/s
	Port 1	N/A	2B	16 Gb/s
Node6	Port 0	N/A	3B	16 Gb/s
	Port 1	N/A	4B	16 Gb/s
Node7	Port 0	N/A	5B	16 Gb/s
	Port 1	N/A	6B	16 Gb/s
Node8	Port 0	N/A	7B	16 Gb/s
	Port 1	N/A	8B	16 Gb/s
Node9	Port 0	N/A	1C	16 Gb/s
	Port 1	N/A	2C	16 Gb/s
Node10	Port 0	N/A	3C	16 Gb/s
	Port 1	N/A	4C	16 Gb/s
Node11	Port 0	N/A	5C	16 Gb/s
	Port 1	N/A	6C	16 Gb/s

TABLE 5. STORAGE PORT MAPPING FOR VALIDATED SAP HANA NODES USING SAS HDDS (CONTINUED)

SAP HANA Node	HBA Port	Virtual Storage Platform Target Port		
		VSP G350	VSP G700	Port Speed
Node12	Port 0	N/A	7C	16 Gb/s
	Port 1	N/A	8C	16 Gb/s
Node13	Port 0	N/A	1D	16 Gb/s
	Port 1	N/A	2D	16 Gb/s

Table 6 shows storage port mapping, Fibre channel switches with hosts for validated number of HANA systems using SSDs.

TABLE 6. STORAGE PORT MAPPING FOR VALIDATED SAP HANA NODES USING 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 F350	VSP F700	VSP F900	Port Speed	Port Security
Node1	Port 0	16 Gb/s	SW-1-P0	SW-1-P32	1A-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P0	SW-2-P32	2A-Host Group 1			32 Gb/s	Enabled
Node2	Port 0	16 Gb/s	SW-1-P1	SW-1-P32	1A-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P1	SW-2-P32	2A-Host Group 2			32 Gb/s	Enabled
Node3	Port 0	16 Gb/s	SW-1-P2	SW-1-P33	3A-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P2	SW-2-P33	4A-Host Group 1			32 Gb/s	Enabled
Node4	Port 0	16 Gb/s	SW-1-P3	SW-1-P33	3A-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P3	SW-2-P33	4A-Host Group 2			32 Gb/s	Enabled
Node5	Port 0	16 Gb/s	SW-1-P4	SW-1-P34	5A-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P4	SW-2-P34	6A-Host Group 1			32 Gb/s	Enabled
Node6	Port 0	16 Gb/s	SW-1-P5	SW-1-P34	5A-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P5	SW-2-P34	6A-Host Group 2			32 Gb/s	Enabled
Node7	Port 0	16 Gb/s	SW-1-P6	SW-1-P35	7A-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P6	SW-2-P35	8A-Host Group 1			32 Gb/s	Enabled
Node8	Port 0	16 Gb/s	SW-1-P7	SW-1-P35	7A-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P7	SW-2-P35	8A-Host Group 2			32 Gb/s	Enabled
Node9	Port 0	16 Gb/s	SW-1-P8	SW-1-P36	1B-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P8	SW-2-P36	2B-Host Group 1			32 Gb/s	Enabled

TABLE 6. STORAGE PORT MAPPING FOR VALIDATED SAP HANA NODES USING SSDS (CONTINUED)

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group				
	Port Name	Port Speed	Host	Storage	VSP F350	VSP F700	VSP F900	Port Speed	Port Security
Node10	Port 0	16 Gb/s	SW-1-P9	SW-1-P36	1B-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P9	SW-2-P36	2B-Host Group 2			32 Gb/s	Enabled
Node11	Port 0	16 Gb/s	SW-1-P10	SW-1-P37	3B-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P10	SW-2-P37	4B-Host Group 1			32 Gb/s	Enabled
Node12	Port 0	16 Gb/s	SW-1-P11	SW-1-P37	3B-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P11	SW-2-P37	4B-Host Group 2			32 Gb/s	Enabled
Node13	Port 0	16 Gb/s	SW-1-P12	SW-1-P38	5B-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P12	SW-2-P38	6B-Host Group 1			32 Gb/s	Enabled
Node14	Port 0	16 Gb/s	SW-1-P13	SW-1-P38	5B-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P13	SW-2-P38	6B-Host Group 2			32 Gb/s	Enabled
Node15	Port 0	16 Gb/s	SW-1-P14	SW-1-P39	7B-Host Group 1			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P14	SW-2-P39	8B-Host Group 1			32 Gb/s	Enabled
Node16	Port 0	16 Gb/s	SW-1-P15	SW-1-P39	7B-Host Group 2			32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P15	SW-2-P39	8B-Host Group 2			32 Gb/s	Enabled
Node17	Port 0	16 Gb/s	SW-1-P16	SW-1-P40	N/A	1C-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P16	SW-2-P40	N/A	2C-Host Group 1		32 Gb/s	Enabled
Node18	Port 0	16 Gb/s	SW-1-P17	SW-1-P40	N/A	1C-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P17	SW-2-P40	N/A	2C-Host Group 2		32 Gb/s	Enabled
Node19	Port 0	16 Gb/s	SW-1-P18	SW-1-P41	N/A	3C-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P18	SW-2-P41	N/A	4C-Host Group 1		32 Gb/s	Enabled
Node20	Port 0	16 Gb/s	SW-1-P19	SW-1-P41	N/A	3C-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P19	SW-2-P41	N/A	4C-Host Group 2		32 Gb/s	Enabled
Node21	Port 0	16 Gb/s	SW-1-P20	SW-1-P42	N/A	5C-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P20	SW-2-P42	N/A	6C-Host Group 1		32 Gb/s	Enabled
Node22	Port 0	16 Gb/s	SW-1-P21	SW-1-P42	N/A	5C-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P21	SW-2-P42	N/A	6C-Host Group 2		32 Gb/s	Enabled
Node23	Port 0	16 Gb/s	SW-1-P22	SW-1-P43	N/A	7C-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P22	SW-2-P43	N/A	8C-Host Group 1		32 Gb/s	Enabled

TABLE 6. STORAGE PORT MAPPING FOR VALIDATED SAP HANA NODES USING SSDS (CONTINUED)

SAP HANA Node	HBA Port		Fibre Channel Switch Port Name		Virtual Storage Platform Target Port-Host Group				
	Port Name	Port Speed	Host	Storage	VSP F350	VSP F700	VSP F900	Port Speed	Port Security
Node24	Port 0	16 Gb/s	SW-1-P23	SW-1-P43	N/A	7C-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P23	SW-2-P43	N/A	8C-Host Group 2		32 Gb/s	Enabled
Node25	Port 0	16 Gb/s	SW-1-P24	SW-1-P44	N/A	1D-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P24	SW-2-P44	N/A	2D-Host Group 1		32 Gb/s	Enabled
Node26	Port 0	16 Gb/s	SW-1-P25	SW-1-P44	N/A	1D-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P25	SW-2-P44	N/A	2D-Host Group 2		32 Gb/s	Enabled
Node27	Port 0	16 Gb/s	SW-1-P26	SW-1-P45	N/A	3D-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P26	SW-2-P45	N/A	4D-Host Group 1		32 Gb/s	Enabled
Node28	Port 0	16 Gb/s	SW-1-P27	SW-1-P45	N/A	3D-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P27	SW-2-P45	N/A	4D-Host Group 2		32 Gb/s	Enabled
Node29	Port 0	16 Gb/s	SW-1-P28	SW-1-P46	N/A	5D-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P28	SW-2-P46	N/A	6D-Host Group 1		32 Gb/s	Enabled
Node30	Port 0	16 Gb/s	SW-1-P29	SW-1-P46	N/A	5D-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P29	SW-2-P46	N/A	6D-Host Group 2		32 Gb/s	Enabled
Node31	Port 0	16 Gb/s	SW-1-P30	SW-1-P47	N/A	7D-Host Group 1		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P30	SW-2-P47	N/A	8D-Host Group 1		32 Gb/s	Enabled
Node32	Port 0	16 Gb/s	SW-1-P31	SW-1-P47	N/A	7D-Host Group 2		32 Gb/s	Enabled
	Port 1	16 Gb/s	SW-2-P31	SW-2-P47	N/A	8D-Host Group 2		32 Gb/s	Enabled

Storage Architecture

Each SAP HANA node needs the following storage layout:

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

This SAP HANA TDI setup utilizes the following two dynamic provisioning pools created with Hitachi Dynamic Provisioning for the storage layout. This ensures maximum utilization and optimization at a lower cost than other solutions.

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

The validated dynamic provisioning pool layout options with minimal disks and storage cache on Hitachi Virtual Storage Platform G350, F350, G700, F700, and F900 storage are listed in Table 7.

TABLE 7. DYNAMIC PROVISIONING POOLS WITH DISKS AND STORAGE CACHE

Storage	Cache	Number of Nodes	Number of Parity Groups in OS_SH_Data_Pool		Number of Parity Groups in LG_Pool	
			RAID-6 (14D+2P)	RAID-10 (2D+2D)	RAID-6 (6D+2P)	RAID-10 (2D+2D)
VSP G350, G370 (with HDD)	128 GB	1	1	N/A	1	N/A
		2	2	N/A	1	N/A
		up to 3	3	N/A	1	N/A
VSP G350, F350, G370, F370 (with SSD)	VSP G350, VSP F350: 128 GB	up to 8	N/A	1	N/A	1
	VSP G370, VSP F370: 256GB	up to 15	N/A	2	N/A	2
		up to 16	N/A	3	N/A	3
VSP G700 (with HDD)	512 GB	up to 5	1	N/A	1	N/A
		up to 7	1	N/A	2	N/A
		up to 9	3	N/A	2	N/A
		up to 10	4	N/A	2	N/A
		up to 13	6	N/A	3	N/A
VSP G700, F700 (with SSD)	512 GB	up to 11	N/A	1	N/A	1
		up to 20	N/A	2	N/A	2
		up to 28	N/A	3	N/A	3
		up to 30	N/A	4	N/A	4
		up to 32	N/A	4	N/A	5

TABLE 7. DYNAMIC PROVISIONING POOLS WITH DISKS AND STORAGE CACHE (CONTINUED)

Storage	Cache	Number of Nodes	Number of Parity Groups in OS_SH_Data_Pool		Number of Parity Groups in LG_Pool	
			RAID-6 (14D+2P)	RAID-10 (2D+2D)	RAID-6 (6D+2P)	RAID-10 (2D+2D)
VSP G900, F900 (with SSD)	1024 GB	up to 17	N/A	1	N/A	1
		up to 23	N/A	2	N/A	2
		up to 31	N/A	3	N/A	3
		up to 32	N/A	4	N/A	3

As an example, Hitachi Virtual Storage Platform G350 listed in Table 7, you need the following:

- A minimum of three parity groups is needed for OS_SH_DT_Pool to fit three SAP HANA production nodes on one VSP G350 using RAID-6 (14D+2P).
- A minimum of one parity group is needed for LG_Pool to fit three SAP HANA production nodes on one VSP G350 using RAID-6 (6D+2P).

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

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

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

- **System 1** – 384 GB
- **System 2** – 768 GB
- **System 3** – 1536 GB
- **System 4** – 3072 GB

Provision the storage for the four SAP HANA systems listed above:

- Determine the minimum sizes for operating system, data, log, and HANA shared using these formulas in SAP white paper [SAP HANA Storage Requirements](#) as following:
 - Every HANA node requires approximately 100 GB capacity for the operating system.
 - /hana/shared size uses formulas:
Single node (scale-up) – Size = MIN (1 × RAM; 1 TB)
Multi-node (scale-out) – Size = 1 × RAM of every 4 worker nodes
 - Data size requires at least 1 × RAM of each HANA node
 - Log size uses formulas:
Systems with equal or less than 512 GB memory – Size = 1/2 × RAM
Systems with greater than 512 GB memory – Size = 512 GB

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

- Provision the storage:
 - Create two dynamic provisioning pools for the three SAP HANA systems on storage:
 - Use **OS_SH_DT_Pool** to provision the operating system volume, SAP HANA shared volume, and Data volume.
 - Use **LG_Pool** to provision the Log volume.
 - For HDDs, create the parity groups first, as an example shown in Table 8, “Dynamic Provisioning Pool with RAID-6 for 3 Nodes on Hitachi Virtual Storage Platform G350,” on page 17 for Hitachi Virtual Storage Platform G350, using the RAID-6 storage design.
 - For SSDs, create the parity groups first, as an example shown in Table 9, “Dynamic Provisioning Pool with RAID-10 (2D+2D) for 16 Nodes on Hitachi Virtual Storage Platform F350,” on page 18 for Hitachi Virtual Storage Platform F350, using the RAID-10 storage design.

TABLE 8. DYNAMIC PROVISIONING POOL WITH RAID-6 FOR 3 NODES ON HITACHI VIRTUAL STORAGE PLATFORM G350

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
OS_SH_DT_Pool	1	RAID-6 (14D+2P) on 600 GB SAS HDD	00:00:01	OS_SH_DT_DPVOL_1	1850 GB	MPU-10
			00:00:02	OS_SH_DT_DPVOL_2	1850 GB	MPU-20
			00:00:03	OS_SH_DT_DPVOL_3	1850 GB	MPU-10
			00:00:04	OS_SH_DT_DPVOL_4	1850 GB	MPU-20
	2	RAID-6 (14D+2P) on 600 GB SAS HDD	00:00:05	OS_SH_DT_DPVOL_5	1850 GB	MPU-10
			00:00:06	OS_SH_DT_DPVOL_6	1850 GB	MPU-20
			00:00:07	OS_SH_DT_DPVOL_7	1850 GB	MPU-10
			00:00:08	OS_SH_DT_DPVOL_8	1850 GB	MPU-20
	3	RAID-6 (14D+2P) on 600 GB SAS HDD	00:00:09	OS_SH_DT_DPVOL_9	1850 GB	MPU-10
			00:00:0A	OS_SH_DT_DPVOL_10	1850 GB	MPU-20
			00:00:0B	OS_SH_DT_DPVOL_11	1850 GB	MPU-10
			00:00:0C	OS_SH_DT_DPVOL_12	1850 GB	MPU-20
LG_Pool	4	RAID-6 (6D+2P) on 600 GB SAS HDD	00:01:01	LG_DPVOL_1	800 GB	MPU-10
			00:01:02	LG_DPVOL_2	800 GB	MPU-20
			00:01:03	LG_DPVOL_3	800 GB	MPU-10
			00:01:04	LG_DPVOL_4	800 GB	MPU-20

TABLE 9. DYNAMIC PROVISIONING POOL WITH RAID-10 (2D+2D) FOR 16 NODES ON HITACHI VIRTUAL STORAGE PLATFORM F350

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-10 (2D+2D) on 1.9 TB SSD	00:00:01	OS_SH_DT_DPVOL_1	878 GB	MPU-10
			00:00:02	OS_SH_DT_DPVOL_2	878 GB	MPU-20
			00:00:03	OS_SH_DT_DPVOL_3	878 GB	MPU-10
			00:00:04	OS_SH_DT_DPVOL_4	878 GB	MPU-20
	2	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:05	OS_SH_DT_DPVOL_5	878 GB	MPU-10
			00:00:06	OS_SH_DT_DPVOL_6	878 GB	MPU-20
			00:00:07	OS_SH_DT_DPVOL_7	878 GB	MPU-10
			00:00:08	OS_SH_DT_DPVOL_8	878 GB	MPU-20
	3	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:09	OS_SH_DT_DPVOL_9	878 GB	MPU-10
			00:00:10	OS_SH_DT_DPVOL_10	878 GB	MPU-20
			00:00:11	OS_SH_DT_DPVOL_11	878 GB	MPU-10
			00:00:12	OS_SH_DT_DPVOL_12	878 GB	MPU-20
Log_Pool	4	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:13	LG_DPVOL_1	878 GB	MPU-10
			00:00:14	LG_DPVOL_2	878 GB	MPU-20
			00:00:15	LG_DPVOL_3	878 GB	MPU-10
			00:00:16	LG_DPVOL_4	878 GB	MPU-20
	5	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:17	LG_DPVOL_5	878 GB	MPU-10
			00:00:18	LG_DPVOL_6	878 GB	MPU-20
			00:00:19	LG_DPVOL_7	878 GB	MPU-10
			00:00:20	LG_DPVOL_8	878 GB	MPU-20
	6	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:21	LG_DPVOL_9	878 GB	MPU-10
			00:00:22	LG_DPVOL_10	878 GB	MPU-20
			00:00:23	LG_DPVOL_11	878 GB	MPU-10
			00:00:24	LG_DPVOL_12	878 GB	MPU-20

- Assign all LDEVs to the dedicated pool as examples shown in Table 8, "Dynamic Provisioning Pool with RAID-6 for 3 Nodes on Hitachi Virtual Storage Platform G350," on page 17 for Virtual Storage Platform G350 and Table 9, "Dynamic Provisioning Pool with RAID-10 (2D+2D) for 16 Nodes on Hitachi Virtual Storage Platform F350," on page 18 for Virtual Storage Platform F350.

- Create virtual volumes (vVOLs) for the operating system, SAP HANA shared, log, and data volumes. Table 10, “vVOLs for SAP HANA Nodes for Four Memory Sizes of HANA Systems,” on page 19 shows examples for HANA systems with memory of 384 GB, 768 GB, 1536 GB, and 3072 GB.

TABLE 10. VVOLS FOR SAP HANA NODES FOR FOUR MEMORY SIZES OF HANA SYSTEMS

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
	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		
00:04:09	HANA_DATA_N4_4	768 GB	MPU-20		

TABLE 10. VVOLS FOR SAP HANA NODES FOR FOUR MEMORY SIZES OF HANA SYSTEMS (CONTINUED)

Dynamic Provisioning Pool	vVOL ID	vVOL Name	vVOL Size	MPU Assignment	System Memory
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

Table 11, "Example LUN Path Assignment for the SAP HANA Configuration on Node 1," on page 20 shows an example configuration of the LUN path assignment for Node 1. Configure the LUN assignment similarly for all other nodes.

TABLE 11. EXAMPLE LUN PATH ASSIGNMENT FOR THE SAP HANA CONFIGURATION ON NODE 1

LUN ID	LDEV ID	LDEV Name
0000	00:01:00	HANA_OS_N1
0001	00:01:01	HANA_SH_N1
0002	00:01:02	HANA_LOG_N1_1

TABLE 11. EXAMPLE LUN PATH ASSIGNMENT FOR THE SAP HANA CONFIGURATION ON NODE 1 (CONTINUED)

LUN ID	LDEV ID	LDEV Name
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 describes the SAP HANA volume and operating system configuration.

Multipath Configuration

This reference architecture uses Device-mapper Multipath, a native component of the Linux operating system.

Using Device-mapper Multipath allows the configuration of multiple I/O paths between the server blades and storage.

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

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

- SAP HANA server boot volume
Not applicable if using internal disks for the boot volume.
- SAP HANA data volume
- SAP HANA log volume
- SAP HANA shared volume
For a scale out solution, Hitachi uses multiple paths to the disks with Hitachi NAS Platform as NFS for /hana/shared. This solution is valid also for other the file systems listed in [SAP Note 405827](#) for SAP HANA.

Hitachi recommends the multipath settings in Table 12, "Recommended Multipath Settings," on page 22.

TABLE 12. RECOMMENDED MULTIPATH SETTINGS

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	directio	
	path_grouping_policy	multibus	
	path_selector	queue-length 0	
	uid_attribute	ID_SERIAL	
	failback	immediate	
	rr_weight	uniform	
	rr_min_io_rq	1 1284	128
	features	0	
	no_path_retry	5	
	prio (Only necessary for Hitachi LPAR using HBA in shared mode.)	weightedpath	
	prio_args4	hbtl 0:0:0:.* 1 1:0:0:.* 2	

HANA Persistent Storage Volume Configuration

For both operating systems, SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux, Hitachi uses an LVM-based storage layout. Once installing the operating system and correctly configuring multipathing, you can see the assigned LUNs in the 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 indicates the LDEV ID you have used during the LUN assignment. In the example above, 000101 maps to LDEV ID: 00:01:01.

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

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

After you have prepared all the LUNs, you need to configure the volume groups using the **vgcreate** command. The names for the volume group differs between scale-up and scale-out installations.

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

For creating other volume groups, use the same syntax, exchanging the volume group name as well as the physical disks or LUNs.

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

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

Use Table 13 to complete the creation of logical volumes.

TABLE 13. DETAILS FOR CREATING LOGICAL VOLUMES

	Number of LUNs	lv Name	vg Name
DATA	4 - following this reference architecture, or the number of assigned LUNs	lvdata	<ul style="list-style-type: none"> ▪ Scale-up: <code>vgdata</code> ▪ Scale-out: <code>vg<SID>data<node number></code>
LOG	4 - following this reference architecture, or the number of assigned LUNs	lvlog	<ul style="list-style-type: none"> ▪ Scale-up: <code>vglog</code> ▪ Scale-out: <code>vg<SID>log<node number></code>
SHARED (Only applicable for scale-up deployments of SAP HANA.)	1 - following this reference architecture, or the number of assigned LUNs	lvshared	<ul style="list-style-type: none"> ▪ Scale-up: <code>vgshared</code>

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 arrays use the XFS file system. In Table 14, find the options to create and mount the file system.

TABLE 14. FILE SYSTEM CREATE AND MOUNT OPTIONS

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

a. These options were tested and verified on Hitachi NAS Platform.

To create a file system, use the following command

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

For example:

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

SAP HANA Persistent Storage Volume Configuration for Scale-Up Deployments

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

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

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

See Table 13, "Details for Creating Logical Volumes," on page 23 and Table 14, "File System Create and Mount Options," on page 24 for volume group and logical volume names as well as the mount options.

2. Create the mount points using:

mkdir -p -m 755 <mount point>

Example for scale-up:

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

Or example for scale-out (assuming HIT is your SAP system ID)

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

3. Mount the file systems.

Mount them one at a time using this:

```
mount <mount point>
```

Or, mount them all at once using this:

```
mount -a
```

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

SAP HANA Persistent Storage Volume Configuration for Scale-Out Deployments Using Hitachi NAS Platform

This is for scale-out systems, only.

To configure the 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. Persist the HANA shared file systems in /etc/fstab by adding the following entry to /etc/fstab:

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

3. Mount the file system:

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

Check if the file system is mounted using either the **df** or **mount** command.

SAP HANA takes care about mounting the file system used for LOG and DATA with the correct options during HANA startup itself. See more details in 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 Fiber Channel-attached devices. This enables host auto-failover on block storages.

The Fiber Channel storage connector, fcClient/fcLientLVM, 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/fcLientLVM implementation uses standard Linux commands, such as **multipath** and **sg_persist**. Install and configure these commands.

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

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

SAP HANA Software Installation

After configuring the file system for the SAP HANA data volume, log volume and HANA shared volume, install the current version of SAP HANA supported by SAP on the SAP HANA server or servers.

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

SAP HANA Software Installation on a Scale-Up Environment

Follow this procedure to install SAP HANA on a scale-up environment. This procedure assumes HIT is your SAP System ID, 10 is your SAP System Number, and saphanas.company.corp is your fully qualified hostname).

To install SAP HANA software on 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 HANA by typing the following command

```
<path to HDB_LCM_LINUX_X86_64>/hdblcm --action install \  
--components=server,client \  
--install_hostagent \  
--number 10 \  
--sapmnt=/hana/shared \  
--sid=HIT \  
--datapath=/hana/data/HIT \  
--logpath=/hana/log/HIT \  
--hostname=saphanas.company.corp \  
--certificates_hostmap= saphanas.company.corp=saphanas.company.corp
```

3. During the installation you need to provide various passwords.
4. Once the installation is complete, continue with SAP HANA Software Configuration.

SAP HANA Software Installation on 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 HANA installation routine with the similar content to the following (assuming HIT is your SAP System ID).

Create this configuration file before following the procedure:

```
[communication]
```

```
listeninterface = .global
```

```
[persistence]
```

```
basepath_datavolumes = /hana/data/HIT
```

```
basepath_logvolumes = /hana/log/HIT
```

```
[storage]
```

```
ha_provider = hdb_ha.fcClientLVM
```

```
partition_*_*_prtype = 5
```

```
partition_1_log__lvmname = vgHITlog001-lvlog
```

```
partition_1_data__lvmname = vgHITdata001-lvdata
```

```
partition_2_log__lvmname = vgHITlog002-lvlog
```

```
partition_2_data__lvmname = vgHITdata002-lvdata
```

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

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

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

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

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

1. Download a supported version of SAP HANA and follow the instructions to extract the archive.
2. Install SAP HANA by running the following command

```
<path to HDB_LCM_LINUX_X86_64>/hdblcm --action install \  
--components=server,client \  
--install_hostagent \  
--number 10 \  
--sapmnt=/hana/shared \  
--sid=HIT \  
--storage_cfg=/tmp/hana_install_HIT \  
--hostname=hana001ic \  
--certificates_hostmap=hana001ic=hana001
```

3. During the installation you need to provide various passwords.
4. Once the installation is complete, continue with "SAP HANA Software Configuration" on page 29.

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

```
/hana/shared/HIT/global/hdb/install/bin/hdbaddhost \  
--sid=HIT \  
--install_hostagent \  
--sapmnt=/hana/shared \  
--role=worker \  
--hostname=hana002ic \  
--storage_partition=2
```

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

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

```
/hana/shared/HIT/global/hdb/install/bin/hdbaddhost \  
--sid=HIT \  
--install_hostagent \  
--sapmnt=/hana/shared \  
--role=standby \  
--hostname=hana003ic
```

SAP HANA Software Configuration

To achieve an optimized performance, the following suggested SAP HANA database parameters were used during engineering validation:

- For HDDs:

```
max_parallel_io_requests[data] = 512  
max_submit_batch_size[data] = 128  
size_kernel_io_queue[data] = 512  
async_read_submit[data] = on  
async_write_submit_blocks[data] = all  
min_submit_batch_size[data] = 16  
async_write_submit_active[data] = on  
max_parallel_io_requests[log] = 512  
max_submit_batch_size[log] = 128  
size_kernel_io_queue[log] = 512  
async_read_submit[log] = on  
async_write_submit_blocks[log] = all  
min_submit_batch_size[log] = 16  
async_write_submit_active[log] = on
```

- For SSDs:

```
max_parallel_io_requests[data] = 512
max_submit_batch_size[data] = 384
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] = 512
max_submit_batch_size[log] = 384
size_kernel_io_queue[log] = 1024
async_read_submit[log] = on
async_write_submit_blocks[log] = all
min_submit_batch_size[log] = 16
async_write_submit_active[log] = on
```

Follow SAP Note [2267798](#) to set up these parameters using hdbparam tool (examples as below) for SAP HANA 1.0.

```
hdbparam -paramset fileio[DATA].max_parallel_io_requests=256
```

```
hdbparam -paramset fileio[LOG].max_parallel_io_requests=256
```

Follow SAP Note [2399079](#) to set up these parameters defined in global.ini or run command (examples as below) for SAP HANA 2.0.

```
"ALTER SYSTEM ALTER CONFIGURATION('global.ini', 'SYSTEM') set ('fileio', 'max_parallel_io_requests[DATA]') = '256';"
```

```
"ALTER SYSTEM ALTER CONFIGURATION('global.ini', 'SYSTEM') set ('fileio', 'max_parallel_io_requests[LOG]') = '256';"
```

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

- Scale-up installation of SAP HANA 2.0 global.ini file:

```
[communication]
```

```
tcp_backlog = 2048
```

```
[fileio]
```

```
max_parallel_io_requests = 256
```

```
async_read_submit = on
```

```
async_write_submit_blocks = all
max_parallel_io_requests[DATA] = 256
max_parallel_io_requests[LOG] = 256
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
```

- Scale-out 2+1 installation of SAP HANA 2.0 global.ini file:

```
[communication]
```

```
tcp_backlog = 2048
```

```
listeninterface = .global
```

```
[fileio]
```

```
max_parallel_io_requests[DATA] = 256
```

```
max_parallel_io_requests[LOG] = 256
```

```
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
```

```
[trace]
```

```
ha_fcclientlvm = info
```

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

```
[persistence]
```

```
basepath_datavolumes=/hana/data/HIT
```

```
basepath_logvolumes=/hana/log/HIT
```

Scale-out 2+1 installation of SAP HANA 1.0 global.ini file:

```
[communication]
```

```
listeninterface = .global
```

```
[persistence]
```

```
basepath_datavolumes = /hana/data/HIT
```

```
basepath_logvolumes = /hana/log/HIT
```

```
[storage]
```

```
ha_provider = hdb_ha.fcClientLVM
```

```
partition_*_*__prtype = 5
```

```
partition_1_data__lvmname = vgHITdata001-lvdata
```

```
partition_1_log__lvmname = vgHITlog001-lvlog
```

```
partition_2_data__lvmname = vgHITdata002-lvdata
```

```
partition_2_log__lvmname = vgHITlog002-lvlog
```

```
[trace]
```

```
ha_fcclientlvm = info
```

Engineering Validation

The validation of this SAP HANA Tailored Datacenter Integration (TDI) enterprise storage configuration used the following:

- SAP HANA Hardware Configuration Check Tool (HWCCT) for testing Hitachi Virtual Storage Platform G350 for the enterprise storage certification, revision hwcct-212
- Table 15, "Hardware Elements in the Engineering Validation," on page 34 shows the hardware elements used on Hitachi Virtual Storage Platform G350 during engineering validation.

TABLE 15. HARDWARE ELEMENTS IN THE ENGINEERING VALIDATION

Hardware	Quantity	Configuration	Role
Hitachi Compute Blade 2500 (CB 2500)	1	<ul style="list-style-type: none"> ▪ 3 server blades on the chassis ▪ 2 management modules ▪ 10 cooling fan modules ▪ 6 power supply modules ▪ 3 Hitachi 16 Gb/sec 2-port Fibre Channel PCIe cards (1 per node) ▪ 3 × 2-port 10GBASE-SR PCIe cards (1 per node) 	Server blade chassis
520X B3 server blade	3	<ul style="list-style-type: none"> ▪ 2 × 22-core or 24-core processors ▪ 256 GB RAM per blade ▪ 1 pass-through mezzanine cards on Mezzanine Slot 2 	SAP HANA server (1 server blade per node)
Hitachi Compute Rack 220H	1	<ul style="list-style-type: none"> ▪ 1 server node with following components: <ul style="list-style-type: none"> ▪ 1 Intel Xeon E3-1281 v3 (4C 3.7GHz 82W) ▪ 1 Heatsink CPU0 and CPU1 ▪ 4 × 8 GB DDR3 1,600 MHz memory module ▪ 2 SATA HDD, 900 GB, 7200 RPM, 2.5 inch (6 Gb/s) ▪ 1 dual port 1 GbE Base-T PCIe card 	Management server runs for SAP HANA Studio


- Table 16 and Figure 3 on page 6 show the server components used on Hitachi Virtual Storage Platform G350 during engineering validation.

TABLE 16. SERVER COMPONENTS IN THE ENGINEERING VALIDATION OF HITACHI VIRTUAL STORAGE PLATFORM G350

Feature	Multiple Scale-up Configuration
Server Blades	3 × 520X B3
Server Blade Location	<ul style="list-style-type: none"> ▪ Blade 1 (primary for Node 1) ▪ Blade 3 (primary for Node 2) ▪ Blade 5 (primary for Node 3)
Network ports	2-port 10GBASE-SR LAN PCIe adapter on one I/O board modules for each server blade in the following locations: <ul style="list-style-type: none"> ▪ IOBD 01B ▪ IOBD 03B ▪ IOBD 05B
Fibre Channel Ports	Hitachi 16 Gb/sec, 2-port Fibre Channel PCIe adapters on one I/O board modules for each HANA node in the following locations: <ul style="list-style-type: none"> ▪ IOBD 01A ▪ IOBD 03A ▪ IOBD 05A
Other interfaces	For all sizes: <ul style="list-style-type: none"> ▪ 1 USB 3.0 port ▪ KVM connector (VGA, COM, USB 2.0 port)

- Figure 4 shows the server components used during engineering validation.

Figure 4



- Table 17 shows storage provision of dynamic provision pools on Hitachi Virtual Storage Platform G350 during engineering validation.

TABLE 17. DYNAMIC PROVISIONING POOL PROVISIONING DURING ENGINEERING VALIDATION ON HITACHI VIRTUAL STORAGE PLATFORM G350

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV (DPVOL) Name	DPVOL Size	MPU Assignment
OS_SH_DT_Pool	1	RAID-6 (14D+2P) on 600 GB SAS HDD	00:00:01	OS_SH_DT_DPVOL_1	3070 GB	MPU-10
			00:00:02	OS_SH_DT_DPVOL_2	3070 GB	MPU-20
	2	RAID-6 (14D+2P) on 600 GB SAS HDD	00:00:03	OS_SH_DT_DPVOL_3	3070 GB	MPU-10
			00:00:04	OS_SH_DT_DPVOL_4	3070 GB	MPU-20
	3	RAID-6 (14D+2P) on 600 GB SAS HDD	00:00:05	OS_SH_DT_DPVOL_5	3070 GB	MPU-10
			00:00:06	OS_SH_DT_DPVOL_6	3070 GB	MPU-20
LG_Pool	4	RAID-6 (6D+2P) on 600 GB SAS HDD	00:00:07	LG_DPVOL_1	1520 GB	MPU-10
			00:00:08	LG_DPVOL_2	1520 GB	MPU-20

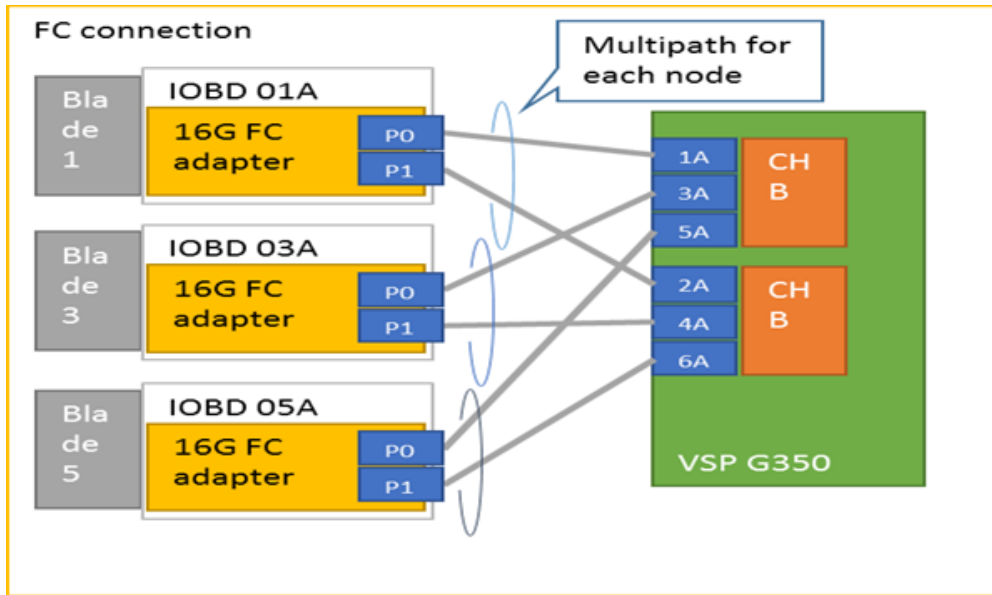
- Table 18 and Figure 5 show the Fibre Channel configurations used during engineering validation.

TABLE 18. FIBRE CHANNEL CONNECTION DURING ENGINEERING VALIDATION

Server Blade	PCIe Slot, Port Location	Virtual Storage Platform G350 Port
Blade 1 (Node 1)	IOBD 01A, Port 0	1A
	IOBD 01A, Port 1	2A
Blade 3 (Node 2)	IOBD 03A, Port 0	3A
	IOBD 03A, Port 1	4A
Blade 5 (Node 3)	IOBD 05A, Port 0	5A
	IOBD 05A, Port 1	6A

- Figure 5 shows the Fibre Channel connection during engineering validation.

Figure 5



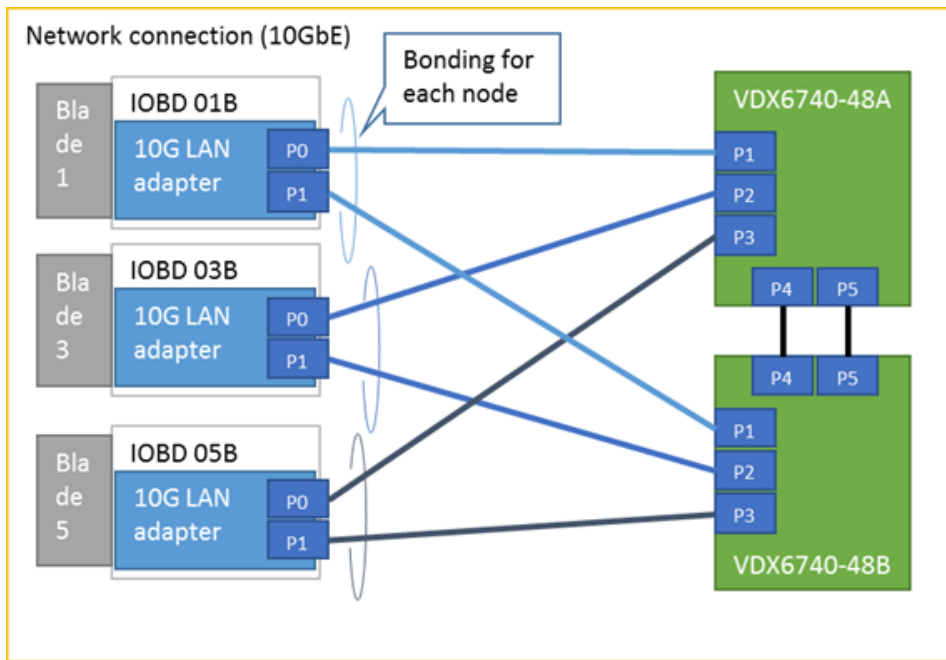
- Table 19 and Figure 6 on page 38 show the Fibre Channel configurations used during engineering validation.

TABLE 19. 10 GBE NETWORK PORT MAPPINGS DURING ENGINEERING VALIDATION

Server Blade	PCIe Slot, Port Location	VDX 6740-48 Switch and Port Number	Bond or ISL
Blade 1 (Node 1)	IOBD 01B, Port 0	VDX6740-48A, port #1	Bond 0 of Node 1
	IOBD 01B, Port 1	VDX6740-48B, port #1	
Blade 3 (Node 2)	IOBD 03B, Port 0	VDX6740-48A, port #2	Bond 0 of Node 2
	IOBD 03B, Port 1	VDX6740-48B, port #2	
Blade 5 (Node 3)	IOBD 05B, Port 0	VDX6740-48A, port #3	Bond 0 of Node 3
	IOBD 05B, Port 1	VDX6740-48B, port #3	
N/A	N/A	VDX6740-48A, port #4	ISL connection
		VDX6740-48B, port #4	
		VDX6740-48A, port #5	ISL connection
		VDX6740-48B, port #5	

- Figure 6 shows the network connection (10 GbE) used during engineering validation.

Figure 6



For More Information

Hitachi Vantara Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the [Services](#) website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the [Resources](#) website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara [Training and Certification](#) website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the [Hitachi Vantara](#) website.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact

© Hitachi Vantara Corporation 2018. All rights reserved. HITACHI and Hitachi Vantara are trademarks or registered trademarks of Hitachi, Ltd. VSP is a registered trademark or trademark of Hitachi Vantara Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

MK-SL-059-01. May 2020.