

Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration using Advanced Server DS7040 and VSP E1090

Reference Architecture Guide

© 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., Hitachi Vantara, Ltd., or Hitachi Vantara Corporation (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, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, OS/390, PowerHA, PowerPC, 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>.

Feedback

Please send comments to doc.feedback@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Revision history

Changes	Date
<ul style="list-style-type: none">▪ RHEL 9.4 and SLES 15 SP6 recertification.	December 2024
<ul style="list-style-type: none">▪ RHEL 9.2 recertification.	September 2024
<ul style="list-style-type: none">▪ Support for SLES 15 SP5 and RHEL 8.8.▪ Added a note that the DS120 management server is end of service life (EOSL).▪ Updated the VSP E1090 storage illustration.▪ Updated network cable connection illustrations for HAE cluster network and client network.▪ Updated URL links.	February 19, 2024

Reference Architecture Guide

Use this reference architecture guide to design and deploy the Hitachi Solution for the SAP HANA platform in one of the following scale-out converged solutions:

- A 24 TB configuration using 1.5 TB SAP HANA nodes
- A 48 TB configuration using 3 TB SAP HANA nodes

Hitachi Solution for SAP HANA in a scale-out configuration is a pre-configured appliance ready to plug in to a network to provide real-time access to operational data for use in converged systems for real-time analytics.

This reference architecture uses the following components:

- Hitachi Advanced Server DS7040 using 2nd generation Intel Xeon scalable processors — These servers provide high reliability, extreme performance, and agile scalability.

This solution makes full use of the huge memory capacity of Advanced Server DS7040 to deliver real-world business benefits from data analytics. In this solution, you can use a number of Hitachi Advanced Server DS7040 nodes to provide the flexibility of various scale-out configuration options. Design your system with up to sixteen active HANA nodes with three standby nodes.

- Hitachi Virtual Storage Platform E1090 (VSP E1090) using NVMe solid state drives (SSD) — These scale for all data types, with storage management that flexibly adapts for performance, capacity, and multi-vendor storage.


In this solution, the persistent storage of the SAP HANA server resides on Hitachi Virtual Storage Platform E1090 using NVMe solid state drives.

- Hitachi Advanced Server DS120 — Optimized for performance, high density, and power efficiency in a dual-processor server, this rack-mount server delivers a balance of compute and storage capacity.

This solution uses one Advanced Server DS120 for use as a management server.

- SAP HANA — This is a multi-purpose, in-memory database to analyze transactional and analytical data.

The following figure shows the solution overview for the scale-out configuration using 1.5 TB SAP HANA nodes or 3 TB SAP HANA nodes.



Note: Testing of this configuration was 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 otherwise matches your production environment before your production implementation of this solution.

Key solution elements


These are the key hardware and software components used in this reference architecture. See [Product descriptions \(on page 35\)](#) for detailed component information.

Hardware elements

The following table describes the hardware used to deploy the sixteen active nodes and three standby nodes configuration.

Hardware	Quantity	Configuration	Role
Hitachi Advanced Server DS7040	19	<ul style="list-style-type: none"> ▪ 2 compute modules ▪ CPU for each SAP HANA node: <ul style="list-style-type: none"> • 2nd Gen Intel Xeon Scalable Processors • 4 Intel Xeon Platinum 8276 or 8276L processors, 28- core, 2.2 GHz, 165 W or • 4 Intel Xeon Platinum 8280 or 8280L processor, 28- core, 2.7 GHz, 205 W ▪ 4 heat sinks for CPU 0/1 in both compute modules ▪ 1.5 TB RAM with 48 × 32 GB DIMMs or 24 × 64 GB DIMMs for each SAP HANA node ▪ 3 TB RAM with 48 × 64 GB DIMMs or 24 × 128 GB DIMMs per SAP HANA node* 	SAP HANA server
Hitachi Virtual Storage Platform E1090	1	<ul style="list-style-type: none"> ▪ CTL — 1 pair ▪ DKBN — 4 pairs ▪ CHB — 4 pairs ▪ MPU — 1 pair ▪ Cache — 1024 GB ▪ 1.9 TB NVMe SSD for 1.5 TB HANA nodes ▪ 3.8 TB NVMe SSD for 3 TB HANA nodes 	Block storage for SAP HANA nodes
Mellanox Technologies MT27710 Family [ConnectX-4 Lx]	4	<ul style="list-style-type: none"> ▪ Mellanox ConnectX-4 dual port PCIe cards for each HANA node 	10 GbE SUSE

Hardware	Quantity	Configuration	Role
	6	<ul style="list-style-type: none"> ▪ Cisco SFP+ 3M twin axial cables for each HANA node 	Enterprise Linux Server High Availability Extension (HAE) or Red Hat Enterprise Linux High Availability (HA) cluster communication network, SAP HANA inter-cluster network, client network
Emulex LPe31002M6 blade, 16 Gbps Fibre Channel dual HBA without cable	2	<ul style="list-style-type: none"> ▪ Only on Module 0 	Connectivity to the external storage sub-system
	2	<ul style="list-style-type: none"> ▪ SFP for each HANA node 	
Hitachi Advanced Server DS120 G2**	1	<ul style="list-style-type: none"> ▪ 2 Intel 63424 24-core; 230 W; 2.8 GHz ▪ 2 Samsung 32 GB DDR4 R-DIMMs 2666 MHz memory ▪ 1 OCP Mezzanine card — Quanta (Intel) X710-T4L 10 GbE, 4 ports ▪ 1 PCIe — Intel E810SFP28 10/25 GbE ▪ 1 vROC with Mirrored M.2 256 GB SSD 	Optional management server
Cisco Nexus 92348 switch	1	<ul style="list-style-type: none"> ▪ 48 × 1 GbE ports 	1 GbE management network
Cisco Nexus 93180YC-FX3 or 93180YC-FX switches***	4	<ul style="list-style-type: none"> ▪ 48 × 10/25 GbE ports 	10 GbE SUSE Linux Enterprise Linux HAE or Red Hat Enterprise Linux HA cluster communication network, SAP HANA inter-cluster network, client network
Brocade G720 Fibre Channel switch	2	<ul style="list-style-type: none"> ▪ 48 × 32Gbps ports ▪ 32 Gbps SFPs 	SAN connections when there are more than 8 active HANA nodes in the configuration

Hardware	Quantity	Configuration	Role
<div style="background-color: #e0f2f1; padding: 10px; border: 1px solid black;">  Note: *L-type processors are required for 128 GB DIMM use. **DS120 is EOSL, replaced by DS120 G2. ***The solution supports both 93180YC-FX3 and 93180YC-FX switches. Information for the 93180YC-FX3 switch in this document also applies to the 91380YC-FX switch. </div>			

Software elements

This reference architecture uses the following software components:

- Hitachi Storage Virtualization Operating System RF
- SAP HANA
- SUSE Linux Enterprise Server (SLES)
- SUSE Linux Enterprise High Availability Extension (HAE)
- Red Hat Enterprise Linux (RHEL)
- Red Hat High-Availability Add On
- Global File System 2

Solution design

The detailed design for Hitachi Solution for SAP HANA in a scale-out configuration reference solution includes the following:

- Scale-out Supported Solutions
- Hitachi Advanced Server DS7000 series Configuration
- Storage Architecture
- Management Server
- SUSE Linux Enterprise Server or RHEL Server Configuration
- SAP HANA Configuration

Supported scale-out solutions

This solution supports the configurations listed in the following table.

Total RAM (1.5 TB)	Total RAM (3 TB)	Active Nodes + 0 Standby Node	Active Nodes + 1 Standby Node	Active Nodes + 2 Standby Nodes	Active Nodes + 3 Standby Nodes	1 Pair of FC Switches*
3	6	2+0	2+1	2+2	N/A	Optional
4.5	9	3+0	3+1	3+2	3+3	Optional
6	12	4+0	4+1	4+2	4+3	Optional
7.5	15	5+0	5+1	5+2	5+3	Optional
9	18	6+0	6+1	6+2	6+3	Optional
10.5	21	7+0	7+1	7+2	7+3	Optional
12	24	8+0	8+1	8+2	8+3	Optional
13.5	27	9+0	9+1	9+2	9+3	Yes
15	30	10+0	10+1	10+2	10+3	Yes
16.5	33	11+0	11+1	11+2	11+3	Yes
18	36	12+0	12+1	12+2	12+3	Yes
19.5	39	13+0	13+1	13+2	13+3	Yes
21	42	14+0	14+1	14+2	14+3	Yes
22.5	45	15+0	15+1	15+2	15+3	Yes
24	48	16+0	16+1	16+2	16+3	Yes



Note: “Optional” means the solution can use either 1 pair of Fibre Channel switches or use direct attached connection. “Yes” means the solution needs to use 1 pair FC switches.

Hitachi Advanced Server DS7000 series configuration




The Hitachi Advanced Server DS7000 series basic building block includes the following:



- 2 processors in one of the following configurations:
 - 28-core using 2nd Gen Intel Xeon Scalable Processors Platinum 8276 or 8276L
 - 28-core using 2nd Gen Intel Xeon Scalable Processors Platinum 8280 or 8280L processors
- 12 memory DIMMs for each CPU, with 24 DIMMs for each compute module
- Up to 7 PCIe Gen3 slots with 5 PCIe adapter modules x8 + 2 internal PCIe Gen3 x16
- 4 native 1/10 GbE copper Ethernet ports

This solution uses one Hitachi Advanced Server DS7040 as a SAP HANA node. It includes the following two compute modules:

- Module 0 acting as a master module
- Module 1, added as a subordinate module to add more resources

The following figure shows the front and back view of Advanced Server DS7040 with network and HBA cards.



-  Mellanox Connect-4x LX 2-port 10GbE Network Card
-  Emulex LPE 31002 16 GBps PCIe HBA Card

The following table shows the Hitachi Advanced Server DS7040 configuration of all SAP HANA nodes in an example 16+3 scale-out configuration.

Hitachi Advanced Server DS7040	SAP HANA Node Name	Role of SAP HANA Node
1	HANA Node 1	Master
2	HANA Node 2	Worker
3	HANA Node 3	Worker
4	HANA Node 4	Worker
5	HANA Node 5	Worker
6	HANA Node 6	Worker
7	HANA Node 7	Worker
8	HANA Node 8	Worker
9	HANA Node 9	Worker
10	HANA Node 10	Worker
11	HANA Node 11	Worker
12	HANA Node 12	Worker
13	HANA Node 13	Worker
14	HANA Node 14	Worker
15	HANA Node 15	Worker
16	HANA Node 16	Worker
17	HANA Node 17	Standby Node
18	HANA Node 18	Standby Node
19	HANA Node 19	Standby Node

This solution uses nineteen Hitachi Advanced Server DS7040 servers with the components listed in the following table.


Feature	Scale-out Configuration
Server	19 Advanced Server DS7040 servers
Server Modules	Each Advanced Server DS7040 serves as one SAP HANA node with the following: <ul style="list-style-type: none"> ▪ Module 0 (primary) ▪ Module 1 (non-primary)
Network ports	2 × 2-port Mellanox ConnectX-4 LX PCIe cards per compute module in the following locations: <ul style="list-style-type: none"> ▪ PCIe Slot 1 ▪ PCIe Slot 3
Fibre Channel Ports	2 × 2-port LPE31002-M6 16 Gbps PCIe cards on Module 0 for each HANA node in the following locations: <ul style="list-style-type: none"> ▪ PCIe Slot 0 ▪ PCIe Slot 2

Storage architecture

This section outlines the storage architecture for this solution.

Direct storage connection

For eight or fewer than eight active HANA nodes in a scale-out configuration, Fibre Channel switches are not necessary because the solution uses direct storage connection to connect the Fibre Channel HBA ports of the HANA nodes to the storage, as shown in the following figure.



Each Hitachi Advanced Server DS7040 server has two 2-port Emulex LPE31002-M6 16 Gbps PCIe HBA cards on the Master Module (0), PCIe slot 0, slot 2.

For each HANA node, port 0 of each of these Emulex HBA cards are directly attached to the 2 × 16 Gbps Fibre Channel ports on Hitachi Virtual Storage Platform.

The following are mandatory settings on the Hitachi Virtual Storage Platform E1090 storage system:


- Enable host mode options (HMO) 02, 94, and 109 for the corresponding storage port connected with the server. The following are the uses of the host mode options:
 - HMO 02 (Optional) — This setting is recommended when the system uses test-unit-ready (TUR) for path_checker in the `/etc/multipath.conf` file. Also, host mode option 02 is typically used when using fencing technology. With a scale-out system, this setting helps suppress unwanted warning messages related to a reservation conflict at the operating system level.
 - HMO 94 — This setting is mandatory for a direct connection between Hitachi Virtual Storage Platform E1090 and the Emulex Lpe31002 HBA to work. Without setting host mode option 94, SAN storage cannot be identified from the Emulex HBA.
 - HMO 109 — This setting fixes the SAN boot issue, where sometimes the system is not able to find the boot LUN after a reboot.
- The LUN ID must be 00 for the OS boot LUN.

Use the storage port properties listed in the following table.

For this setting	Use this value
Port Security	Disabled
Port Speed	16 Gbps
Fabric	OFF
Connection Type	P-to-P

Fibre Channel SAN storage connection

For more than eight active SAP HANA nodes in a scale-out configuration, the solution uses SAN connections to connect the Fibre Channel HBA ports of the SAP HANA node to the storage through Brocade G720 switches, as shown in the following figure.



Each Hitachi Advanced Server DS7040 server has two 2-port Emulex LPE31002-M6 16 Gbps PCIe HBA cards on the Master Module (0), PCIe slot 0, slot 2.

For each HANA node, port 0 of each of the two Emulex LPE31002-M6 16 Gbps PCIe HBA cards is connected to a port on each of two Brocade G720 Fibre Channel switches. The two switches are connected to the two 32 Gbps Fibre Channel ports on Hitachi Virtual Storage Platform E1090, and the switch zones are configured.

The following are mandatory settings on the Hitachi Virtual Storage Platform E1090 storage:

- In addition to Host Group (00), create second Host Group (01) to have two SAP HANA nodes share one storage port through the Fibre Channel switches.
- The LUN ID for the OS boot LUN must be 00.

Use the storage port properties listed in the following table.

For this setting	Use this value
Port Security	Enabled
Port Speed	Auto
Fabric	ON
Connection Type	P-to-P

RAID architecture

Hitachi Vantara uses a building block approach with four active nodes when designing the storage system for each SAP HANA node using Hitachi Virtual Storage Platform E1090.

Several usage aspects divide the space provided by Hitachi Virtual Storage Platform E1090 in the scale-out configuration, as follows:

- Operating system device provisioning for SAP HANA nodes
- STONITH device provisioning for storage-based fencing of the SUSE HAE cluster or the RHEL HA cluster
- HANA shared provisioning for storing HANA binaries and cluster-wide configuration files
- Log device provisioning for SAP HANA database
- Data device provisioning for SAP HANA database

Each SAP HANA node has its own data volume and log volume. Only active SAP HANA nodes need data volumes and log volumes. Standby nodes do not require these volumes.

This reference architecture uses a layout for the dynamic provisioning pool design for the storage that ensures maximum utilization and optimization at a lower cost.

1.5 TB SAP HANA node configuration

For the 1.5 TB HANA node configuration, build the storage system as follows:

- Shared storage-based STONITH volume and SAP HANA log volumes
 - Configure parity groups as RAID 6 (6D + 2P) to provision the shared storage based STONITH LUN and log LUNs for SAP HANA nodes on Hitachi Virtual Storage Platform E1090.
 - One parity group on 8 × 1.9 TB NVMe SSD drives created for up to 8 active HANA nodes.
 - Two parity groups on 16 × 1.9 TB NVMe SSD drives created for 9 up to 16 active HANA nodes.
 - Create four LDEVs with a capacity of 2640 GB from the single parity group.
 - Create one dynamic provisioning pool: Log_Stonith_HDP_Pool. Assign the four LDEVs from the parity group to the dynamic provisioning pool as pool volumes. Create the following virtual volumes from this pool:
 - One small virtual volume of 50 MB on Hitachi Virtual Storage Platform E1090 for STONITH of the SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network. This enables STONITH and fencing in clusters without using external power switches but using shared storage.
 - 600 GB virtual volumes for use as log volumes for each SAP HANA node.
- Operating System Volumes, SAP HANA Shared Volumes, and SAP HANA Data Volumes
 - For each four-node building block, configure two parity groups as RAID 6 (6D + 2P) on 16 × 1.9 TB NVMe SSD drives to provision the operating system LUN, HANA shared LUN, and data LUNs for SAP HANA nodes on Hitachi Virtual Storage Platform E1090.
- Create four LDEVs with a capacity of 2640 GB from each of these two parity groups.
- Create one dynamic provisioning pool: OS_HS_Data_HDP_Pool. Assign the four LDEVs from each of the two parity groups to the dynamic provisioning pool, as pool volumes. Create the following virtual volumes from this pool:
 - Create four virtual volumes (assuming you have four HANA nodes), each 100 GB, for use as operating system volumes. The installation of SUSE Linux Enterprise Server for SAP Applications or Red Hat Enterprise Linux resides on this boot LUN.
 - Create 4 virtual volumes, each 384 GB, for use as SAP HANA shared volumes. The capacity of each virtual volume will be increased by 384 GB for every four-node building block.
 - Create 16 volumes, each 1172 GB, for use as SAP HANA data volumes. Each SAP HANA node uses four volumes as its SAP HANA data volume.

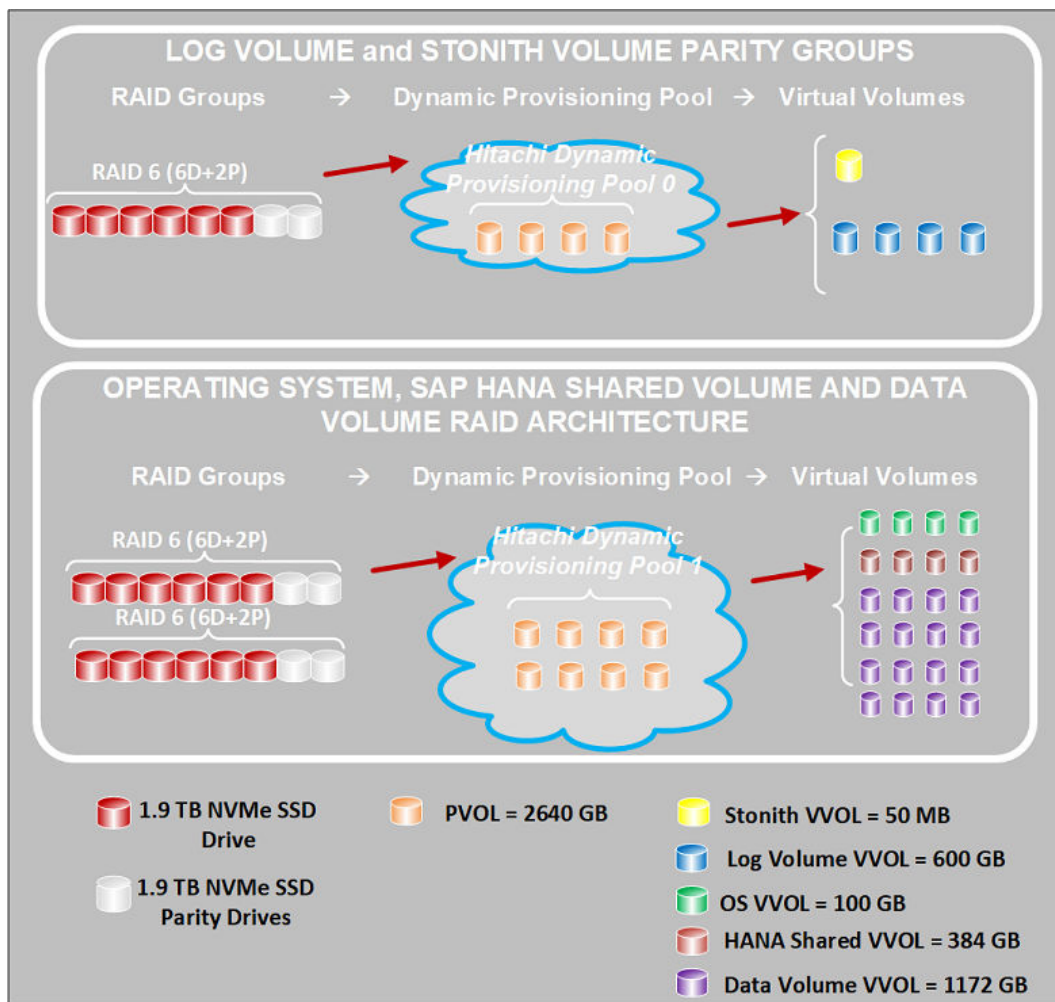
SAP HANA Node Building Blocks	VSP E1090	Log Volume and STONITH Drives, RAID Groups	Operating System, HANA Shared and Data Volume Drives, RAID Groups
4	1	8 × 1.9 TB NVMe SSD drives in 1 group configured as RAID 6 (6D+2P)	16 × 1.9 TB NVMe SSD drives in 2 groups configured as RAID 6 (6D+2P)
8	1	8 × 1.9 TB NVMe SSD drives in 1 group configured as RAID 6 (6D+2P)	32 × 1.9 TB NVMe SSD drives in 4 groups configured as RAID 6 (6D+2P)
12	1	16 × 1.9 TB NVMe SSD drives in 2 groups configured as RAID 6 (6D+2P)	48 × 1.9 TB NVMe SSD drives in 6 groups configured as RAID 6 (6D+2P)
16	1	16 × 1.9 TB NVMe SSD drives in 2 groups configured as RAID 6 (6D+2P)	64 × 1.9 TB NVMe SSD drives in 8 groups configured as RAID 6 (6D+2P)

Pool Name	Pool ID	Parity Group	Pool Volumes	Virtual Volumes
Log_Stonith_HDP_Pool	0	1	4 × 2640 GB	Log Volume: 4 × 600 GB STONITH: 1 × 50 MB
OS_HS_Data_HDP_Pool	1	2	4 × 2640 GB	Boot Volume: 8 × 100 GB HANA Shared: 4 × 384 GB Data Volume: 16 × 1172 GB
		3	4 × 2640 GB	

SAP HANA Node Building Block	Operating System LUNs	Log LUNs	Data LUNs	HANA Shared	STONITH
4	8 × 100 GB	4 × 600 GB	16 × 1172 GB	4 × 384 GB	1 × 50 MB
8	12 × 100 GB	8 × 600 GB	32 × 1172 GB	4 × 768 GB	1 × 50 MB

SAP HANA Node Building Block	Operating System LUNs	Log LUNs	Data LUNs	HANA Shared	STONITH
12	16 × 100 GB	12 × 600 GB	48 × 1172 GB	4 × 1152 GB	1 × 50 MB
16	19 × 100 GB	16 × 600 GB	64 × 1172 GB	4 × 1536 GB	1 × 50 MB

The following figure shows the RAID group configuration for Hitachi Virtual Storage Platform E1090 used in the 1.5 TB HANA node configuration with four active nodes and one standby node.



3 TB SAP HANA node configuration

For the 3 TB HANA node configuration, build the storage system as follows:

- Shared storage-based STONITH volume and SAP HANA log volumes
 - Configure one parity group as RAID 6 (6D + 2P) on 8 × 3.8 TB NVMe SSD drives that provisions the shared storage based STONITH LUN and log LUNs for the SAP HANA nodes on Hitachi Virtual Storage Platform E1090.
 - Create four LDEVs with a capacity of 2640 GB from the single parity group.
 - Create one dynamic provisioning pool: Log_Stonith_HDP_Pool. Assign the four LDEVs from the parity group to the dynamic provisioning pool as pool volumes. Create the following virtual volumes from this pool:
 - One small virtual volume of 50 MB on Hitachi Virtual Storage Platform E1090 for STONITH of the SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network. This enables STONITH and fencing in clusters without external power switches, but with shared storage.
 - Four 600 GB virtual volumes for use as log volumes for each SAP HANA node.
- Operating System Volumes, SAP HANA Shared Volumes, and SAP HANA Data Volumes
 - For each four-node building block, configure two parity groups as RAID 6 (6D + 2P) on 16 × 3.8 TB NVMe SSD drives provisions the operating system LUN, HANA shared LUN, and data LUNs for the SAP HANA nodes on Hitachi Virtual Storage Platform E1090.
 - Create four LDEVs with a capacity of 2640 GB from each of these two parity groups.
 - Create one dynamic provisioning pool: OS_HS_Data_HDP_Pool. Assign the four LDEVs from each of the two parity groups to the dynamic provisioning pool as pool volumes. Create the following virtual volumes from this pool:
 - Create four virtual volumes (assuming you have four HANA nodes), each 100 GB, for use as operating system volumes. The installation of SUSE Linux Enterprise Server for SAP Applications or Red Hat Enterprise Linux resides on this boot LUN.
 - Create 4 virtual volumes, each 768 GB, for use as SAP HANA shared volumes. The capacity of each virtual volume will be increased by 768 GB for every four-node building block.
 - Create 16 volumes each of size 2350 GB for use as SAP HANA data volumes. Each SAP HANA node uses four volumes as its SAP HANA data volume.


SAP HANA Node Building Blocks	VSP E1090	Log Volume and STONITH Drives, RAID Groups	Operating System, HANA Shared and Data Volume Drives, RAID Groups
4	1	8 × 1.9 TB NVMe SSD drives in 1 group configured as RAID 6 (6D+2P)	16 × 1.9 TB NVMe SSD drives in 2 groups configured as RAID 6 (6D+2P)

SAP HANA Node Building Blocks	VSP E1090	Log Volume and STONITH Drives, RAID Groups	Operating System, HANA Shared and Data Volume Drives, RAID Groups
8	1	8 × 3.8 TB NVMe SSD drives in 1 group configured as RAID 6 (6D+2P)	32 × 3.8 TB NVMe SSD drives in 4 groups configured as RAID 6 (6D+2P)
12	1	16 × 3.8 TB NVMe SSD drives in 2 groups configured as RAID 6 (6D+2P)	48 × 3.8 TB NVMe SSD drives in 6 groups configured as RAID 6 (6D+2P)
16	1	16 × 3.8 TB NVMe SSD drives in 2 groups configured as RAID 6 (6D+2P)	64 × 3.8 TB NVMe SSD drives in 8 groups configured as RAID 6 (6D+2P)

Pool Name	Pool ID	Parity Group	Pool Volumes	Virtual Volumes
Log_Stonith_HDP_Pool	0	4	4 × 2640 GB	Log Volume: 4 × 600 GB STONITH: 1 × 50 MB
OS_HS_Data_HDP_Pool	1	8	4 × 2640 GB	Boot Volume: 8 × 100 GB HANA Shared: 4 × 768 GB Data Volume: 16 × 2350 GB
		8	4 × 2640 GB	

SAP HANA Node Building Block	Operating System LUNs	Log LUNs	Data LUNs	HANA Shared	STONITH
4	8 × 100 GB	4 × 600 GB	16 × 2350 GB	4 × 768 GB	1 × 50 MB
8	12 × 100 GB	8 × 600 GB	32 × 2350 GB	4 × 1536 GB	1 × 50 MB
12	16 × 100 GB	12 × 600 GB	48 × 2350 GB	4 × 2304 GB	1 × 50 MB
16	19 × 100 GB	16 × 600 GB	64 × 2350 GB	4 × 3072 GB	1 × 50 MB

The following figure shows the RAID group configuration for Hitachi Virtual Storage Platform E1090 used in the 3 TB HANA node configuration with four active nodes and one standby node.



Build the Hitachi Virtual Storage Platform E1090 with the components as shown in the following table for 4 nodes, 8 nodes, 12 nodes, or 16 nodes.

VSP E1090 Component	4 Node Quantity	8 Node Quantity	12 Node Quantity	16 Node Quantity
Cache	1024 GB	1024 GB	1024 GB	1024 GB
MPU	1 pair	1 pair	1 pair	1 pair
DKBN	2 pairs	2 pairs	4 pairs	4 pairs
CHB	3 pairs	4 pairs	4 pairs	4 pairs

Map the LUNs in the following order for the LUN path assignment for each node:

1. The operating system volume for the specific SAP HANA node
2. The STONITH volume
3. The SAP HANA shared virtual volumes
4. The log volume of each SAP HANA node
5. The data volume of each SAP HANA node

The LUN assignment is the same for all nodes, except for the first LUN. The first LUN is the operating system volume for that specific node.

Network architecture

The scale-out configuration of this SAP HANA solution requires the following four separate networks:

- SAP HANA Inter-Cluster Network — This network provides the communication between the SAP HANA instances on the cluster.
- SAP HANA Client Network — This network is dedicated to the traffic between the SAP HANA database and its clients.
- SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability Cluster Network — This network is used for communication between cluster nodes.
- Management Network — This network is used for system management.

The SAP HANA inter-cluster network, client network, and SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network are required to provide the following:

- No single point of failure (NSPOF)
- At least 10 GbE equivalent throughput

To meet these requirements, this solution uses four Mellanox ConnectX4-Lx PCIe HBA cards installed on the PCIe Slot 1 and Slot 3 of Module 0 and Module 1 of the Hitachi Advanced Server DS7040 chassis.

To provide a redundancy bond, provide two ports from different PCIe network adapters at the operating system level using active-active network bonding mode, following the IEEE 802.3ad Link Aggregation standard for each of the following three networks:

- SAP HANA inter-cluster network
- SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network
- SAP HANA client network

The connections of each bond need to go to physically different Cisco Nexus 93180YC-FX3 switches. This way, if one switch fails, there is still another route to the corresponding host. Connect the two switches together using a virtual port-channel (vPC) peer link. This allows the links to be physically connected to two different switches but appear as a single device to a downstream device as part of a single port-channel.

Use an MTU size of 9216, following Cisco's best practices. Isolate it using the following:

- VLAN 100 for the HANA inter-cluster network
- VLAN 150 for the SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network

Switches

This solution requires the following two sets of switches:

- Cisco Nexus 93180YC-FX3 switch
 - For the SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network and the SAP HANA inter-cluster network within the scale-out configuration of the SAP HANA installation, use two vPC peers linked Cisco Nexus 93180YC-FX3 switches. In this solution, refer to the switches as Nexus-93180-A and Nexus-93180-B.
 - For the client network, use two vPC peers linked Cisco Nexus 93180YC-FX3 switches to provide connectivity. In this solution, refer to the switches as Nexus-93180-C and Nexus-93180-D.


Connect each of these two switches together using a vPC peer link. This lets both switches act together as one logical switch. If one switch fails, there still is a path to the hosts.

- Cisco Nexus 92348 management switch

Use this switch to connect the management ports of the hardware and the management server.

SAP HANA inter-cluster network

The following figure shows an example configuration for the SAP HANA inter-cluster network connection for a 1.5 TB or 3 TB SAP HANA node configuration.




Configure the SAP HANA inter-cluster network using operating system-level bonding on every node.

The following table lists an example of the inter-cluster network port mappings of the 10 GbE connection for a 1.5 TB or 3 TB SAP HANA node configuration.

DS7040 Server Chassis, Module, PCIe Slot, Port Number	Cisco 93180YC-FX3 Switch and Port Number	Bond
Chassis 1, Module 0, PCIe 1, Port 1	Nexus-93180-A, Port #1	Bond 0 of HANA Node
Chassis 1, Module 0, PCIe 3, Port 1	Nexus-93180-B, Port #1	Bond 0 of HANA Node

SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network

The following figure shows an example configuration for the SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network connection for a 1.5 TB or 3 TB SAP HANA node configuration.




Configure the SUSE Linux Enterprise High Availability Extension or Red Hat Enterprise Linux High Availability cluster network using operating system-level bonding on every node.

The following table lists an example of the SUSE HAE or RHEL HA network port mappings of the 10 GbE connection for a 1.5 TB or 3 TB SAP HANA node configuration.

DS7040 Server Chassis, Module, PCIe Slot, Port Number	Cisco 93180YC-FX3 Switch and Port Number	Bond
Chassis 1, Module 0, PCIe 1, Port 0	Nexus-93180-A, Port #21	Bond 1 of HANA Node
Chassis 1, Module 0, PCIe 3, Port 0	Nexus-93180-B, Port #21	Bond 1 of HANA Node

SAP HANA client network connection

The following figure shows an example of the SAP HANA client network connection for a 1.5 TB or 3 TB SAP HANA node configuration.



Configure the SAP HANA client network using operating system-level bonding on every node.

The following table lists an example of the SAP HANA client network port mappings of the 10 GbE connection for a 1.5 TB or 3 TB SAP HANA node configuration.

DS7040 Server Chassis, Module, PCIe Slot, Port Number	Cisco 93180YC-FX3 Switch and Port Number	Bond
Chassis 1, Module 1, PCIe 1, Port 0	Nexus-93180-C, Port #1	Bond 2 of HANA Node
Chassis 1, Module 1, PCIe 3, Port 0	Nexus-93180-D, Port #1	Bond 2 of HANA Node

Management network

The management network resides on a 1 GbE Cisco Nexus 92348 switch. The management network does not need to have a VLAN assigned to it. The Cisco Nexus 92348 switch uses the default switch configuration.

Management server

This solution uses one Hitachi Advanced Server DS120 G2 as the management server, which acts as a central device for managing the SAP HANA platform.

Manage the following from the management server:

- Hitachi Advanced Server DS7040
- Hitachi Virtual Storage Platform E1090
- NTP configuration
- Cisco network switches
- Brocade Fibre Channel switches (when needed)
- SAP HANA system using SAP HANA studio
- SSH access to the operating system running on the SAP HANA nodes

Place the mezzanine and PCI-Ex cards in the Advanced Server DS120 G2 management server slots, as shown in the following table and figure.

PCIe Slot	Description	Component
Slot 1 (Purple)	BMC Port	Management Port for BMC/Browser Connection
Slot 2 (Blue)	PCI-Ex Slot	Emulex LPE 35002 M2 32 Gbps
Slot 3 (Red)	PCI-Ex Slot	Intel E810 SFP28 10/25 GbE
Slot 4 (Green)	OCP Mezzanine Slot	Intel X710-T4L 10 GbE, 4 ports



The management server network ports use an Intel X557-T4 10GbE, 4-port mezzanine card.

- Slot 4 Port 0 — Connected to the Cisco Nexus 92348 switch. It provides the 1 GbE management network to all components.
- Slot 4 Port 1 — Connected to the uplink switch. It provides the 1 GbE network to the management server.

Connect the 10 GbE network ports to the two Cisco switches, Nexus-93180-A and Nexus-93180-B, to provide management access to the SAP HANA nodes from the management server.

Install the following software on the management server:

- PuTTY
- Teraterm
- JAVA 8.0.181 Build 13
- WinSCP
- SAP HANA Studio
- Foxit Reader
- iSMART
- Intel RSTe

SUSE Linux Enterprise Server or Red Hat Enterprise Linux Server configuration

Changing the configuration settings for SUSE Linux Enterprise Server or Red Hat Enterprise Linux Server is only supported using guidelines from SAP and the operating system distributor. Not following these guidelines can cause significant performance problems.

The following SAP Notes for SUSE Linux Enterprise Server or Red Hat Enterprise Linux are a good starting point for information on this topic (SAP S-user credentials required for SAP Notes access):

- [2235581 — SAP HANA: Supported Operating Systems](#)
- [2684254 — SAP HANA DB: Recommended OS settings for SLES 15 / SLES for Applications 15](#)
- [2578899 — SUSE Linux Enterprise Server 15: Installation Note](#)
- [1730999 — Configuration changes to SAP HANA system](#)
- [2009879 — SAP HANA Guidelines for Red Hat Enterprise Linux \(RHEL\) Operating System](#)
- [2777782 — SAP HANA DB: Recommended OS Settings for RHEL 8](#)
- [2772999 — Red Hat Enterprise Linux 8.x: Installation and Configuration](#)
- [3108302 — SAP HANA DB: Recommended OS Settings for RHEL 9](#)
- [3108316 — Red Hat Enterprise Linux 9.x: Installation and Configuration](#)

See *Updating and Patching the Operating System* in the SAP HANA Master Guide from the [SAP Help Portal](#) for more details.

SAP HANA configuration

Use this SAP HANA configuration for the Hitachi Solution for SAP HANA in a scale-out configuration.

Operating system configuration

A scale-out configuration for SAP HANA requires SAN boot.

The operating system LUN is the primary boot device for each scale-out node using the 16 Gbps 2-port Fibre Channel PCIe cards. The operating system LUN holds partitions for the following:

- Operating system
- `/usr/sap/` directory
- Linux swap space

Activate Device-mapper Multipath

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 Hitachi Advanced Server DS7040 and Hitachi Virtual Storage Platform E1090. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available unless all paths fail.

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

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

SAP HANA Shared Global File System 2 volume configuration

This solution uses Global File System 2 to store the cluster-wide SAP HANA binaries and configuration files of the in-memory database. The shared file system is `/hana/shared/<SID>` and accessible by all clustered nodes. `<SID>` is the system ID for the SAP HANA production database instance.

In this solution, the 1.5 TB and 3 TB SAP HANA node configurations use four virtual volumes. All virtual volumes are in one dynamic provisioning pool.

The logical volume manager (LVM) creates a single striped volume on which the GFS2 volume is created for the SAP HANA shared file system and managed by the LVM locking daemon (`lvmlockd`) for SLES 15 SPx, RHEL8.x, and RHEL 9.x.

The capacity of a 1.5 TB or 3 TB configuration is based on the number of SAP HANA nodes.

Configuration Size	Active HANA Nodes	Capacity
1.5 TB	Up to 4	1.5 TB
	Up to 8	3 TB
	Up to 12	4.5 TB
	Up to 16	6 TB
3 TB	Up to 4	3 TB
	Up to 8	6 TB
	Up to 12	9 TB
	Up to 16	12 TB



Note: Use other file systems (like NFS) for SAP HANA system backup, instead of the HANA shared GFS2 file system, to improve performance.

SAP HANA data volume configuration

The logical volume manager creates a single striped volume on which you create the XFS file system to store the SAP HANA data volume. The striped volume acts as the persistent layer for the SAP HANA server.

For each scale-out active HANA node, create a single striped volume across the data LUNs. Format the volume with the XFS file system to store the SAP HANA data volume.

SAP HANA log volume configuration

Logical Volume Manager creates a single striped volume on which you create the XFS file system to store the SAP HANA log volume.

For each scale-out active HANA node, create a single striped volume across the log LUNs. Format the volume with the XFS file system.

SAP Storage Connector API Fibre Channel client

The SAP HANA Storage Connector API Fibre Channel client defines a set of interface functions called during the following:

- Normal SAP HANA cluster operation
- Failover

The scale-out of Hitachi Unified Compute Platform for the SAP HANA Platform uses the fcClientLVM implementation, which supports the use of Logical Volume Manager. SAP supports this solution to enable the use of high-performance Fibre Channel devices for a scale-out installation.

The fcClientLVM implementation uses standard Linux commands, such as `multipath` and `sg_persist`.

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

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

See [SAP Note 1900823](#) (SAP HANA user credentials required) for detailed information about the SAP Fibre Channel Storage Connector API.

SAP HANA appliance software installation

After configuring the file system for the SAP HANA data and log volumes, install SAP HANA on the SAP HANA nodes.

Install the following SAP HANA software components:

- SAP HANA database
- SAP HANA client
- SAP host agent

Engineering validation

The test methodology for validating this scale-out appliance configuration using Hitachi Advanced Server DS7040 on Hitachi Virtual Storage Platform E1090 used the following:

- SAP HANA 2.0
- SAP HANA hardware cloud measurement tool (HCMT) used to test HANA data and log volumes for the following operating systems:
 - Red Hat Enterprise Linux 8.4, 8.6, 8.8, and 9.2
 - SUSE Linux Enterprise Server 15 SP3, SP4, and SP5

For optimal use of the system with a SAP HANA database, use the parameters listed in [Sample global.ini files \(on page 32\)](#) for your operating system release. Follow [SAP Note 2399079](#) to set up these parameters defined in `global.ini` for SAP HANA 2.0.

Sample global.ini files

These are the `global.ini` files used to test this environment on DS7040 server with 2nd Gen Intel Xeon Scalable Processors for engineering validation.

SUSE Enterprise Linux

The following is a sample `global.ini` file configured and used for validation of a system with default values (the default values are either specified in `[fileio]` or automatically taken if the `[fileio]` section is excluded in the `global.ini` file) of HDB parameters on SUSE Enterprise Linux 15 SP3, SP4, SP5, and SP6:

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

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

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

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

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_* prtype = 5
partition_1_data lvmname = vgHITdata001-lvdata
partition_1_log lvmname = vgHITlog001-lvlog
partition_2_data lvmname = vgHITdata002-lvdata
partition_2_log lvmname = vgHITlog002-lvlog
partition_3_data lvmname = vgHITdata003-lvdata
partition_3_log lvmname = vgHITlog003-lvlog
partition_4_data lvmname = vgHITdata004-lvdata
partition_4_log lvmname = vgHITlog004-lvlog
```

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

Red Hat Enterprise Linux

The following is a sample `global.ini` file configured and used for validation of a system with default values (the default values are either specified in `[fileio]` or automatically taken if the `[fileio]` section is excluded in the `global.ini` file) of HDB parameters on Red Hat Enterprise Linux 8.4, 8.6, 8.8, 9.2, and 9.4:

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

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

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

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

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_* prtype = 5
partition_1_data lvmname = vgHITdata001-lvdata
partition_1_log lvmname = vgHITlog001-lvlog
partition_2_data lvmname = vgHITdata002-lvdata
partition_2_log lvmname = vgHITlog002-lvlog
partition_3_data lvmname = vgHITdata003-lvdata
partition_3_log lvmname = vgHITlog003-lvlog
partition_4_data lvmname = vgHITdata004-lvdata
partition_4_log lvmname = vgHITlog004-lvlog
```

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

Product descriptions

The following products are used in this reference architecture.

Hitachi Advanced Server DS7000 series

To take advantage of the latest developments in artificial intelligence (AI), data analytics and machine learning, you require an infrastructure with high reliability, extreme performance, and agile scalability. [Hitachi Advanced Server DS7000 series](#) servers deliver this with a unique modular architecture.

Your server can be configured and scaled to meet the needs of a wide variety of application workloads. This can be used from in-memory data analytics processing to virtualization and hybrid cloud.

The Advanced Server DS7000 series has several complementary models, each based on the Intel Xeon scalable processor. You can upgrade a model to the next model, preserving your hardware and software investment as your business grows.

Hitachi Advanced Server DS120 G2

With support for two Intel Xeon Scalable processors in just 1U of rack space, the [Hitachi Advanced Server DS120 G2](#) delivers exceptional compute density. It provides flexible memory and storage options to meet the needs of converged and hyperconverged infrastructure solutions, as well as for dedicated application platforms such as internet of things (IoT) and data appliances.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 32 slots for high-speed DDR4 memory, allowing up to 4 TB memory capacity with RDIMM population (128 GB × 32) or 8 TB (512 GB × 16) of Intel Optane Persistent Memory. DS120 G2 supports up to 12 hot-pluggable, front-side-accessible 2.5-inch non-volatile memory express (NVMe), serial-attached SCSI (SAS), serial-ATA (SATA) hard disk drive (HDD), or solid-state drives (SSD). The system also offers 2 onboard M.2 slots.

With these options, DS120 G2 can be flexibly configured to address both I/O performance and capacity requirements for a wide range of applications and solutions.

Hitachi Virtual Storage Platform E1090

The Hitachi Virtual Storage Platform E1090 (VSP E1090) storage system is a high-performance, large-capacity data storage system. The VSP E1090 all-flash arrays (AFAs) support NVMe and SAS solid-state drives (SSDs). The VSP E1090H hybrid models can be configured with both SSDs and hard disk drives (HDDs).

- The NVMe flash architecture delivers consistent, low-microsecond latency, which reduces the transaction costs of latency-critical applications and delivers predictable performance to optimize storage resources.
- The hybrid architecture allows for greater scalability and provides data-in-place migration support.

Hitachi Storage Virtualization Operating System RF

Hitachi Storage Virtualization Operating System RF powers the Hitachi Virtual Storage Platform (VSP) family. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid arrays as storage capacity, consolidating resources for a higher ROI and providing a high-speed front end to slower, less-predictable arrays.

SAP HANA

SAP HANA converges database and application platform capabilities in-memory to transform transactions, analytics, text analysis, predictive and spatial processing so businesses can operate in real-time. This combines database, data processing, and application platform capabilities in a single in-memory platform. Also, the platform provides libraries for predictive, planning, text processing, spatial, and business analytics — all on the same architecture. This architecture comes from leading hardware partners of SAP, including Hitachi Vantara.

By eliminating the divide between transactions and analytics, SAP HANA allows you to answer any business question anywhere in real time.

As an SAP customer, you can [download more information](#), including the following:

- **SAP HANA Master Guide**

This is the central starting point for the technical implementation of SAP HANA. Use this guide for basic concepts and for planning.

- **SAP HANA Server Installation and Update Guide**

This guide provides an overview of how to install and update an SAP HANA system with the SAP HANA lifecycle management tools.

- **SAP HANA Administration Guide**

This guide explains how to configure, manage, maintain, and optimize your SAP HANA installation using SAP HANA administration tools.

[SAP HANA hardware directory](#) provides information about SAP HANA appliances certified by SAP hardware partners.

SUSE Linux Enterprise High Availability Extension

[SUSE Linux Enterprise High Availability Extension](#) is an industry-leading open source high availability clustering system designed to virtually eliminate unplanned downtime, be easy to use, and be deployed in physical and virtual environments.

Use high availability clustering to automate application and data recovery. You can use their flexible, policy-driven clustering solution to implement highly available Linux clusters and eliminate single points of failure. Your servers are continuously monitored, and when a fault or failure occurs, the workload is transferred from one server to another, or the application is automatically restarted on a known working system. This helps you maintain business continuity and minimize unplanned downtime.

Red Hat Enterprise Linux

Using the stability and flexibility of [Red Hat Enterprise Linux](#), reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability on military-grade security. Use Enterprise Linux to tailor your infrastructure as markets shift and technologies evolve.

Red Hat High-Availability Add-On

[RedHat High-Availability Add-On](#) allows a service to fail over from 1 node to another with no apparent interruption to cluster clients, evicting faulty nodes during transfer to prevent data corruption. This Add-On can be configured for most applications (both off-the-shelf and custom) and virtual guests, supporting up to 16 nodes. The High Availability Add-On features a cluster manager, lock management, fencing, command-line cluster configuration, and a Conga administration tool.

Global File System 2

Global File System 2 (GFS2) is a shared disk file system for Linux computer clusters. It allows all nodes to have direct concurrent access to the same shared block storage. GFS2 is included in [SUSE Linux Enterprise High Availability Extension](#) or [Red Hat High-Availability Add-On](#).

In order to use GFS2, a high-availability cluster framework, such as SUSE Linux Enterprise High Availability or Red Hat High Availability Add-On, needs to be implemented. However, the high availability of the SAP HANA scale-out instances is provided by SAP's host auto-failover feature, not by the operating system cluster framework itself.

This solution uses GFS2 to provide a shared binary file system in a scale-out configuration that is exposed to all the hosts of a system under `/hana/shared/<SID>`.

Cisco Nexus switches

The Cisco Nexus switch product line provides a series of solutions that make it easier to connect and manage disparate data center resources with software-defined networking (SDN). Leveraging the Cisco Unified Fabric, which unifies storage, data, and networking (Ethernet/IP) services, the Nexus switches create an open, programmable network foundation built to support a virtualized data center environment.

Brocade switches from Broadcom

Brocade and Hitachi Vantara have partnered to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

[Brocade Fibre Channel switches](#) deliver industry-leading performance with seventh generation 64Gb/sec Fibre Channel interfaces, simplifying scale-out network architectures. Get the high-performance, availability, ease of management, and support for the next generation of Hitachi Virtual Storage Platform storage systems on a solid storage network foundation that can grow as your need grows.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA



HitachiVantara.com/contact