HITACHI
Inspire the Next

Hitachi Solution for the SAP HANA Platform in a Scale-up Configuration Using Hitachi Advanced Server DS220 with Intel Xeon Platinum 8176 or 8176M Processors

Reference Architecture Guide

MK-SL-015-12
December 2023

Feedback

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

Revision history

Revision	Changes	Date
MK-SL-015-12	Support for SLES 15 SP5 and RHEL 8.8	December 8, 2023
MK-SL-015-11	Support for SLES 15 SP4 and RHEL 8.6	February 22, 2023
MK-SL-015-10	Support for SLES 15 SP3 and RHEL 8.4	March 16, 2022

Reference Architecture Guide

Read about the scale-up Hitachi Solution for the SAP HANA platform using Hitachi Advanced Server DS220. This solution uses either of the following:

- Internal storage housed in the server serving as an appliance configuration
- Connected to an external storage subsystem following the tailored data center integration (TDI) approach to implement SAP HANA

This SAP HANA infrastructure uses the following components:

- Hardware
  - One Hitachi Advanced Server DS220 with internal drives providing internal storage
  - (Optional) External sub-system storage, such as Hitachi Virtual Storage Platform G200 (VSP) or VSP G350
- Software
  - Preconfigured with SAP HANA to provide a converged solution for converged solutions for real time analytics

The validation of this environment with external storage was with Hitachi Virtual Storage Platform G200 or VSP G350. Your needs may require other storage options. Contact your account representative for details and implementation services whenever you require using external storage.

The following figure shows the topology of this reference solution using internal drives on Hitachi Advanced Server DS220.

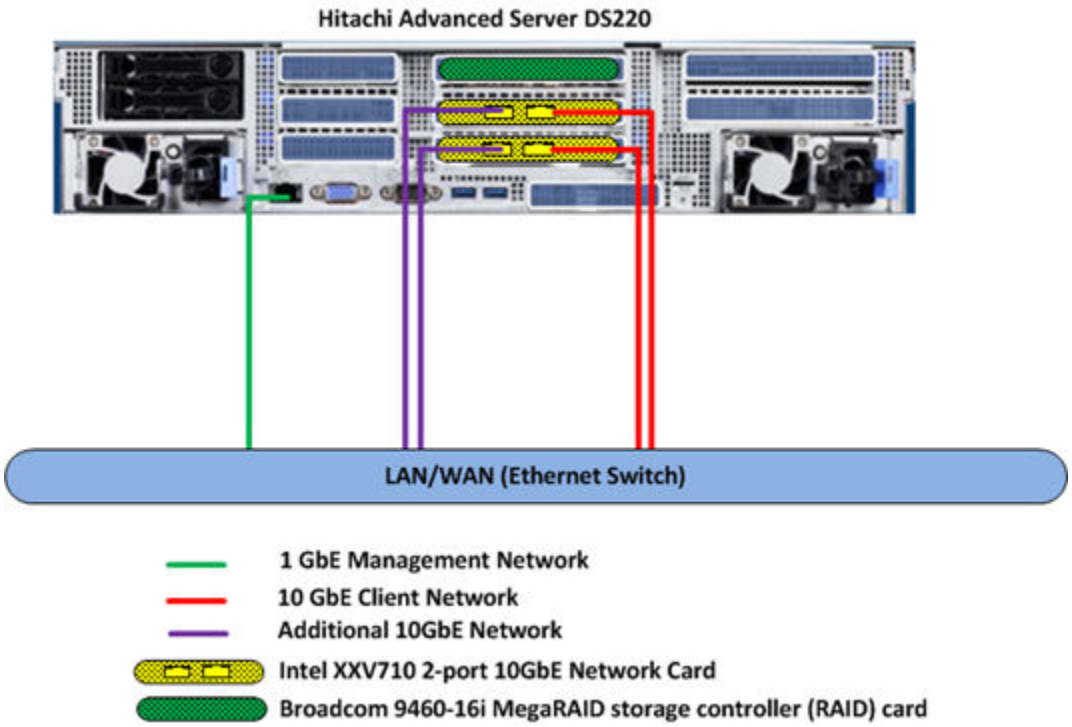

The following figure shows the topology of this reference solution using external drives on a storage subsystem, such as Virtual Storage Platform G200.

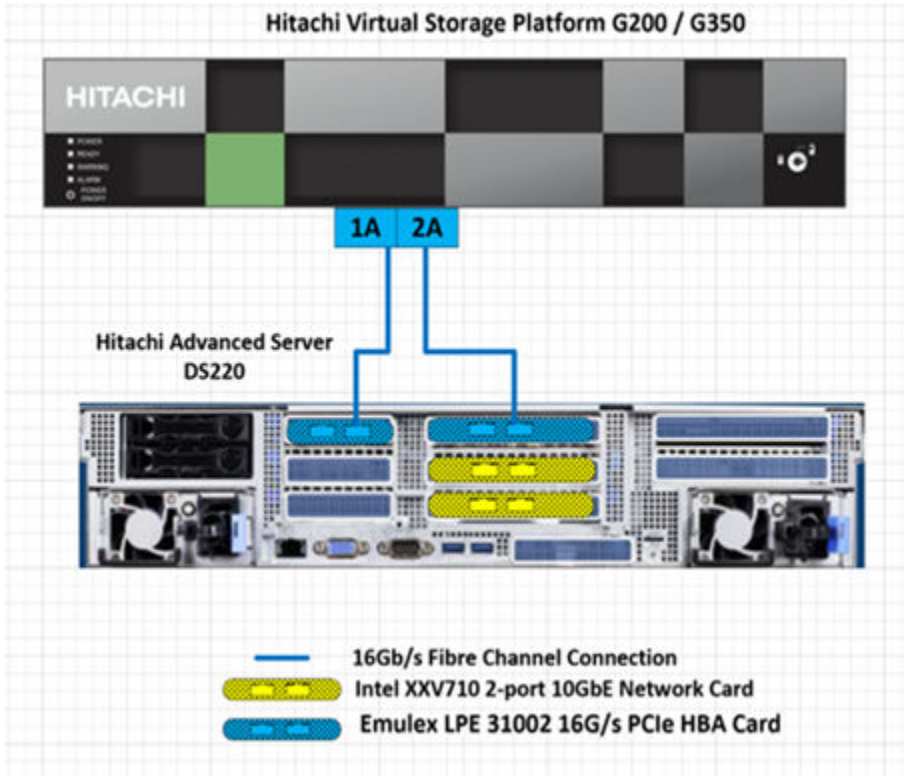

Solution for SAP HANA is a preconfigured converged system for real-time analytics. It is ready to plug into your network to provide real-time access to operational data for use in analytic models.

This system supports the configurations listed in the following table.

Table 1 Supported Configuration

Number of Sockets	RAM Size	Storage
2 sockets	▪ 384 GB ▪ 768 GB ▪ 1536 GB ▪ 3072 GB	Internal storage option: ▪ Internal drives on the Hitachi Advanced Server DS220
		External storage option: ▪ Storage subsystem, such as Hitachi Virtual Storage Platform G200 or VSP G350

This technical paper assumes that you have familiarity with the following:

- Storage area network (SAN)-based storage systems
- General storage concepts
- Common IT storage practices
- SAP HANA

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

Key solution elements

The following are the key hardware and software components used in this reference architecture.

Note: Do not change the layout of any of the components in this environment without consulting your Hitachi Vantara account representative. Changing this layout can require manual configuration of the network and/or using different components.

Hardware elements

The following table lists the hardware used to deploy the specific scale-up configuration of Hitachi Solution for the SAP HANA platform for the different sized solutions.

Hitachi Advanced Server DS220 has storage drives, RAID card, and HBA card added, based on whether using an implementation with internal storage or an external storage sub-system, as indicated in the following table.

Table 2 Hardware Elements

Hardware	Quantity	Configuration	Role	Implementation Type
Hitachi Advanced Server DS220	1	▪ CPU per SAP HANA node: • 2 Intel Xeon Platinum 8176 or 8176M Processor 28-core, 2.1GHz, 165W ▪ 2 heat sinks CPU 0/1 ▪ RAM per SAP HANA node: • 384 GB (12 × 32 GB DIMMS) • 768 GB (24 × 32 GB DIMMS or 12 × 64 GB DIMMS) • 1536 GB (24 × 64 GB DIMMS or 12 × 128 GB DIMMS) • 3072 GB (24 × 128 GB DIMMS)	SAP HANA server	All implementations
Intel PCIe network cards	2	▪ Intel XXV710 dual port SFP28 (LP-MD2) PCIe card	For SAP HANA 10 GbE client network and additional 10 GbE network	All implementations
	2	▪ Cisco SFP+ 3M Twinaxial cables		
Hitachi Virtual Storage Platform G200 or VSP G350	1	▪ Single frame	Block storage when using an external storage sub-system	External storage only
Broadcom PCIe HBA card	1	▪ 2-port LPE31002-M6 16 Gb/s card	Connectivity to the external storage sub-system	External storage only
	2	▪ 16 Gb/s SFP		

Hardware	Quantity	Configuration	Role	Implementation Type
RAID controller card	1	▪ Broadcom 9460-16i MegaRAID storage controller (RAID) card	Required on the Hitachi Advanced Server DS220, when using internal storage	Internal storage only
Storage drives	5	▪ 1.92 TB Intel S4500 SATA SSD	Supports less than or equal to 1536 GB RAM per HANA node.	Internal storage only
	9		Supports greater than or equal to 1536 GB RAM per HANA node	
	1		Spare drive	
Cisco Nexus 3048 switch	1	▪ 48 × 1 GbE ports	Optional switch for management network	All implementations (optional)
Cisco Nexus 93180YC-EX switch	2	▪ 48 × 10 GbE ports	Optional switches for things such as the client network or additional backup network	All implementations (optional)
Minkels Global Solutions Rack	1	▪ 1 standard rack	Optional rack for mounting server	All implementations
PDUs	6	▪ Vertical PDUs	Optional PDUs for solution	All implementations

Software elements

The following table describes the software products used to deploy this solution.

Software	
Operating system	SUSE Linux Enterprise Server for SAP Applications

Software	
	Red Hat Enterprise Linux for SAP (alternate)
SAP HANA	

Solution design

The detailed design for this scale-up configuration of Hitachi Solution for the SAP HANA Platform for this reference solution includes the following:

- Hitachi Advanced Server DS220 Configuration
- Network Architecture
- Storage Architecture Configuration
- SAP HANA Configuration

Hitachi Advanced Server DS220 configuration

The following figure shows the front and back view of Hitachi Advanced Server DS220 when using internal drives.

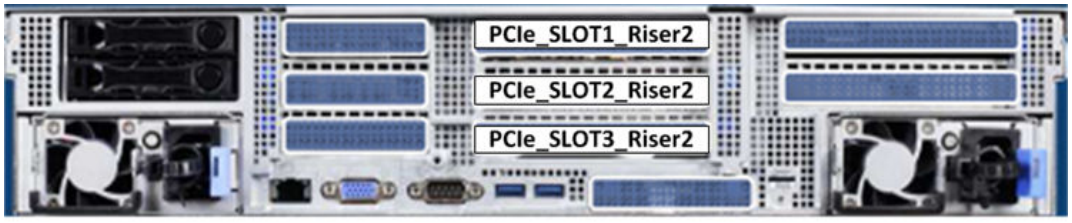

Hitachi Advanced Server DS220 with internal drives - PCIe Slots used

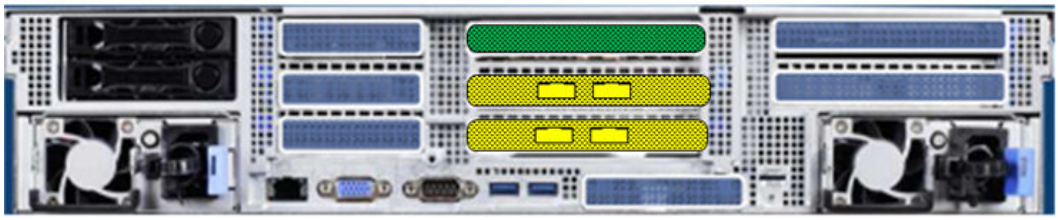

Hitachi Advanced Server DS220 with internal drives (Back)

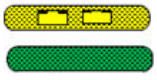

Intel XXV710 2-port 10GbE Network Card

Broadcom 9460-16i MegaRAID storage controller (RAID) card

The following figure shows the front and back view of the Hitachi Advanced Server DS220 when using external drives on a storage subsystem, such as Hitachi Virtual Storage Platform G200.

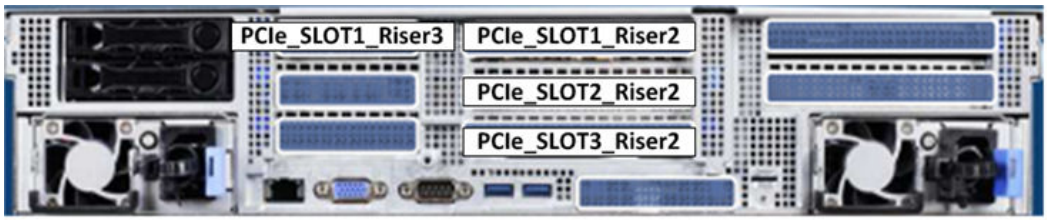

Hitachi Advanced Server DS220 with external storage – PCIe Slots used (Back)

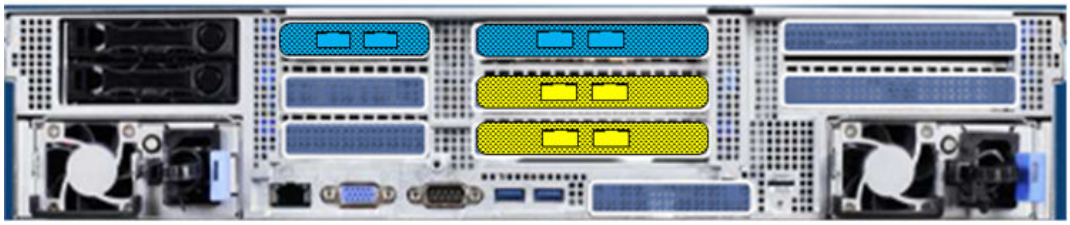

Hitachi Advanced Server DS220 with external Storage (Back)

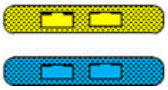

Intel XXV710 2-port 10GbE Network Card
Emulex LPE 31002 16G/s PCIe HBA Card

This solution uses one Hitachi Advanced Server DS220, with the following components:

- 2 dual port 10 GbE Intel XXV710 SFP28 PCIe card on the PCIe_SLOT2_Riser2 and PCIe_SLOT3_Riser2
- Required only with the internal storage option:
  - 1 Broadcom 9460-16i MegaRAID storage controller (RAID) PCIe card on the PCIe_SLOT1_Riser2
  - 5 × 1.92 TB Intel S4500 SATA SSD drives on the front bay, supporting less than or equal to 1536 GB RAM per SAP HANA node
  - 9 × 1.92 TB Intel S4500 SATA SSD on the front bay, supporting greater than 1536 GB RAM per SAP HANA node
  - 1 × 1.92 TB Intel S4500 SATA SSD drives used as spare drive
- Required only with the external storage option:
  - 2 × 2-port Broadcom LPE31002-M6 16 Gb/s PCIe HBA card on PCIe_SLOT1_Riser2 and PCIe_SLOT1_Riser3

Network architecture configuration

Connect the 1 GbE management port on Hitachi Advanced Server DS220 to a Cisco Nexus 3048 switch or to any other external 1 GbE switch for management connectivity.

Make the following 10 GbE network connections for the client network setup of the SAP HANA node as an uplink network following table:

- Connect the following to Cisco Nexus 93180YC-EX switches or to any other external switches:
  - Port 0 of 10 GbE Intel XXV710 SFP28 PCIe card to the two different Cisco Nexus 93180YC-EX switches
  - Bond the corresponding two ports eth9901 and eth9902 as bond0 at the operating system level using active-active network bond mode with the following options:

    mode= 802.3ad miimon=100 xmit_hash_policy=layer3+4 lacp_rate=fast

  This acts as the client network for the SAP HANA node.
- If additional 10 GbE network connections are required, connect the Port 1 of both the 10 GbE Intel XXV710 SFP28 PCIe cards to the Cisco Nexus 93180YC-EX switches or to any other external switches.

Table 3 Network Setup

Network Card	Port	Network Description
PCIe_SLOT3_Riser2	0	Client network for the SAP HANA node (eth9902)
	1	Free for use as additional 10 GbE network for backup, SAP HANA system replication, or other purposes
PCIe_SLOT2_Riser2	0	Client network for the SAP HANA node (eth9901)
	1	Free for use as additional 10 GbE network for backup, SAP HANA system replication, or other purposes

Note: The management network and client network can be on the same network switch or a separate network switch, depending on the network environment.

Storage architecture configuration

This describes the two storage configurations for this environment:

- Internal Storage
- External Storage Option

Internal storage option

These are the storage components needed to setup a scale-up configuration for SAP HANA using internal storage with Hitachi Advanced Server DS220:

- 1 Broadcom 9460-16i MegaRAID storage controller (RAID) PCIe card
- 5 × 1.92 TB Intel S4500 SATA SSD drives configured as RAID-5(4D+1P), supporting less than or equal to 1536 GB RAM per SAP HANA node including the spare drive
- 9 × 1.92 TB Intel S4500 SATA SSD drives configured as RAID-5(8D+1P), supporting greater than 1536 GB RAM per SAP HANA node including the spare drive
- 1 × 1.92 TB Intel S4500 SATA SSD drives used as a spare drive

Storage sizing and configuring varies for the different sizes by considering requirements for I/O and capacity, including the following:

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

The following table lists the storage configuration for different memory sizes.

Table 4 Storage Configuration When Using Internal Storage

Configuration	Memory Less Than or Equal to 1.5 TB	Memory Greater Than 1.5 TB
Parity Group	▪ 1 × RAID-5 (4D+1P) using 5 × 1.92 TB Intel S4500 SATA SSD	▪ 1 × RAID-5 (8D+1P) using 9 × 1.92 TB Intel S4500 SATA SSD
OS, Data, Log and HANA Shared	▪ Operating system volume: 100 GB ▪ SAP HANA shared volume: 1536 GB ▪ Data volume: 4915 GB ▪ Log volume: 600 GB	▪ Operating system volume: 100 GB ▪ SAP HANA shared volume: 3072 GB ▪ Data volume: 9830 GB ▪ Log volume: 600 GB
Spare drive	1 × 1.92 TB Intel S4500 SATA SSD	

External storage option

Note: Each implementation of this reference architecture can use a different storage architecture. Validation for this environment used Hitachi Virtual Storage Platform G200 and VSP G350 as external storage.

The following storage configuration and LUN layout is for Virtual Storage Platform G200. However, VSP G350 uses the same configuration as VSP G200, except for the RAID level and disk type. On VSP G350, the RAID level was RAID-10 (2D+2D) with 1.9 TB SSD drives.

Contact your account representative for details and implementation services when you want an environment using external storage.

These are the storage components you need to implement a scale-up SAP HANA system with Hitachi Advanced Server DS220 using Hitachi Virtual Storage Platform G200 for external storage:

- 2 × 2-port Broadcom LPE31002-M6 16 Gb/s PCIe HBA card
- 1 Hitachi Virtual Storage Platform G200
- Storage drive box trays (DBS)
- Spare drives

In case of the external storage option with direct connection between the Virtual Storage Platform G200 storage and the Emulex HBA on the DS220 server, the following is mandatory:

- Use the System Mode 847 on the storage array.
- Enable the Host Mode Option 02, Host Mode Option 94, and Host Mode Option 109 for the corresponding storage port connected with the server.
  - Host Mode Option 109 — Fixes the SAN boot issue, where sometimes the system cannot find the Boot LUN after a reboot.
  - Host Mode Option 94 —Mandatory for direct connection between the Hitachi Virtual Storage Platform and the Emulex Lpe31002 HBA to work. Without setting up Host Mode Option 94, SAN storage cannot be identified from the Emulex HBA.
  - Host Mode Option 02 — This is required because the system uses Test-Unit-Ready (TUR) for path_checker in the /etc/multipath.conf file.
- The Host Group ID must be 00 for SAN boot.

Make sure to use the port properties listed in the following table.

Table 5 Port Properties on Virtual Storage Platform G200

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

The SAP HANA node needs the following storage layout:

- Operating system volume
- SAP HANA shared volume for the SAP HANA binaries and other configuration files
- Log volume
- Data volume

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

Use two dynamic provisioning pools with the specific parity groups listed in the following table for the storage layout.

Table 6 Dynamic Provisioning Pools

Dynamic Provisioning Pool Name	Purpose	Parity Group RAID Level and Disks
OS_SH_Data_Pool	Operating system LUN SAP HANA shared LUN Data LUN	RAID-6 (14D+2P) on 600 GB, 10k RPM SAS drives
Log_Pool	Log LUN	RAID-6 (6D+2P) on 600 GB, 10k RPM SAS drives

The example layout in the following table uses the dynamic provisioning pool layout on Virtual Storage Platform G200 for a SAP HANA TDI solution with 1.5TB scale-up system.

Table 7 Dynamic Provisioning Pool for a SAP HANA TDI Solution

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size (GB)	MPU Assignment
OS_SH_Data_Pool	1	RAID-6 (14D +2P) on 600 GB, 10k RPM SAS drives	00:00:01	OS_SH_DA_Pool_1	1800	MPU-10
			00:00:02	OS_SH_DA_Pool_2	1800	MPU-11
			00:00:03	OS_SH_DA_Pool_3	1800	MPU-20
			00:00:04	OS_SH_DA_Pool_4	1800	MPU-21
Log_Pool	2	RAID-6 (6D +2P) on 600 GB, 10k RPM	00:01:01	Log_Pool_1	750	MPU-10
			00:01:02	Log_Pool_2	750	MPU-11

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size (GB)	MPU Assignment
		SAS drives	00:01:03	Log_Pool_3	750	MPU-20
			00:01:04	Log_Pool_4	750	MPU-21

Provision the virtual volumes for the operating system, SAP HANA shared, data, and log volumes following table for a SAP HANA TDI solution with 1.5 TB scale-up system.

Table 8 Virtual Volumes for the SAP HANA Nodes

Dynamic Provisioning Pool	Virtual Volume ID	Virtual Volume Name	Virtual Volume Size	MPU Assignment
OS_SH_Data_Pool	00:02:00	HANA_OS	100 GB	MPU-10
	00:02:01	HANA_SH	1024 GB	MPU-20
Log_Pool	00:02:02	HANA_LOG_1	128 GB	MPU-10
	00:02:03	HANA_LOG_2	128 GB	MPU-11
	00:02:04	HANA_LOG_3	128 GB	MPU-20
	00:02:05	HANA_LOG_4	128 GB	MPU-21
OS_SH_Data_Pool	00:02:06	HANA_DATA_1	384 GB	MPU-10
	00:02:07	HANA_DATA_2	384 GB	MPU-11
	00:02:08	HANA_DATA_3	384 GB	MPU-20
	00:02:09	HANA_DATA_4	384 GB	MPU-21

The following table has the LUN path assignment used when validating this environment.

Table 9 Example LUN Path Assignment

LUN ID	LDEV ID	LDEV Name
0000	00:02:00	HANA_OS
0001	00:02:01	HANA_SH

LUN ID	LDEV ID	LDEV Name
0002	00:02:02	HANA_LOG_1
0003	00:02:03	HANA_LOG_2
0004	00:02:04	HANA_LOG_3
0005	00:02:05	HANA_LOG_4
0006	00:02:06	HANA_DATA_1
0007	00:02:07	HANA_DATA_2
0008	00:02:08	HANA_DATA_3
0009	00:02:09	HANA_DATA_4

SAP HANA configuration

This describes how to configure SAP HANA in this solution.

File system

These volumes utilize the BTRFS file system for an implementation of this solution either with internal drives or with an external storage sub-system:

- BTRFS Filesystem
  - Operating system volume (SLES Only)
- XFS Filesystem
  - Operating system volume (RHEL Only)
  - SAP HANA shared volume
  - Data volume
  - Log volume

Device-Mapper Multipath

If implementing this solution with internal storage on Hitachi Advanced Server DS220, deactivate Device-Mapper Multipath.

If implementing this solution with an external storage sub-system, activate Device-Mapper Multipath.

SAP HANA software installation

After configuring the file system for the SAP HANA data volume and log volume, install the SAP HANA 2.0 SP06 or latest stack on the server.

Install the following SAP HANA software components on the server:

- Database
- Client
- Agent

Engineering validation

The test methodology for validating the appliance configuration using Hitachi Advanced Server DS220 with Intel S4500 SATA SSD drives and SAP HANA tailored data center integration (TDI) enterprise storage configuration with Hitachi Virtual Storage Platform G200 and VSP G350 used the following:

- SAP HANA Hardware Configuration Check Tool (HWCCT) with FSPERF revision `hwcct-212_4` was tested on these volumes for SUSE Linux Enterprise Server 12 SP2:
  - Data volume
  - Log volume
- SAP HANA Hardware Configuration Check Tool (HWCCT) with FSPERF revision `hwcct-212_5` was tested on these volumes for SUSE Linux Enterprise Server 12 SP3, SUSE Linux Enterprise Server 12 SP4, SUSE Linux Enterprise Server 12 SP4, Red Hat Enterprise Linux 7.3, Red Hat Enterprise Linux 7.4, and Red Hat Enterprise Linux 7. 6:
  - Data volume
  - Log volume
- For additional testing, a configuration with Intel S4500 SATA SSD internal drives with `fsperf_single_mix_106.sh` script except for `fsperf_single_mix_107.sh` script with SUSE Linux Enterprise Server 12 SP4:
  - MixedIO on Data/Log Volume (A MixedIO test is mandatory if the data volume and log volume are created in same parity group. Otherwise this test is optional.)
- SAP HANA Hardware and Cloud Measurement Tools HCMT-046_0 was tested on these volumes for SUSE Linux Enterprise Server 15 SP1:
  - Data volume
  - Log volume
- For either an appliance or TDI solution, SAP HANA Hardware and Cloud Measurement Tool (HCMT) revision hcmt-047_0 was tested on these volumes for Red Hat Enterprise Linux 8.1:
  - Data volume
  - Log volume
- For either an appliance or TDI solution, SAP HANA Hardware and Cloud Measurement Tool (HCMT) revision `hcmt-055_0` was tested on these volumes for Red Hat Enterprise Linux 8.1 and SUSE Linux Enterprise Server 15 SP2:
  - Data volume
  - Log volume

- For either an appliance or TDI solution, SAP HANA Hardware and Cloud Measurement Tool (HCMT) revision hcmt-056_0 was tested on these volumes for Red Hat Enterprise Linux 8.2:
  - Data volume
  - Log volume
- For either an appliance or TDI solution, SAP HANA Hardware and Cloud Measurement Tool (HCMT) revision hcmt-060_0 was tested on these volumes for SLES 15 SP3 and Red Hat Enterprise Linux 8.4:
  - Data volume
  - Log volume
- For either an appliance or TDI solution, SAP HANA Hardware and Cloud Measurement Tool (HCMT) revision hcmt-065_0 was tested on these volumes for SLES 15 SP4 and Red Hat Enterprise Linux 8.6:
  - Data volume
  - Log volume
- For either an appliance or TDI solution, SAP HANA Hardware and Cloud Measurement Tool (HCMT) revision hcmt-073_0 was tested on these volumes for SLES 15 SP5 and Red Hat Enterprise Linux 8.8:
  - Data volume
  - Log volume
  - Shared volume

For optimal use of the system with a SAP HANA database, use the parameters listed in in the corresponding appendix for your operating system release. Follow SAP Note 2399079 to setup these parameters defined in `global.ini` for SAP HANA 2.0.

Sample global.ini files

These are sample `global.ini` files for the two different operating systems.

SUSE Enterprise Linux Server for SAP Applications

This is the sample `global.ini` file used when validating this solution with SUSE Enterprise Linux Server for SAP Applications 12 SP2, 12 SP3, 12 SP4, 15 SP1, 15 SP2, 15 SP3, 15 SP4, and 15 SP5 as an appliance or a TDI solution.

[communication]

```
tcp_backlog = 2048
```

[fileio]

```
max_parallel_io_requests[data] = 64
max_submit_batch_size[data] = 64
```

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

[multidb]

```
mode = multidb
database_isolation = low
singletenant = yes
```

[persistence]

```
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
```

Red Hat Enterprise Linux

This is the sample `global.ini` file used when validating this solution with Red Hat Enterprise Linux 7.3, 7.4, 7.6, 8.2, 8.4, 8.6, and 8.8 as an appliance solution.

[communication]

```
tcp_backlog = 2048
```

[fileio]

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

[multidb]

```
mode = multidb
database_isolation = low
singletenant = yes
```

[persistence]

```
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
```

This is the sample `global.ini` file used when validating this solution with Red Hat Enterprise Linux 7.6, Red Hat Enterprise Linux 8.1, Red Hat Enterprise Linux 8.2, Red Hat Enterprise Linux 7.6, and Red Hat Enterprise Linux 8.1, as a TDI solution.

[communication]

```
tcp_backlog = 2048
```

[fileio]

```
max_parallel_io_requests[data] = 64
max_submit_batch_size[data] = 64
size_kernel_io_queue[data] = 512
async_read_submit[data] = on
async_write_submit_blocks[data] = all
min_submit_batch_size[data] = 16
async_write_submit_active[data] = auto
max_parallel_io_requests[log] = 256
max_submit_batch_size[log] = 64
size_kernel_io_queue[log] = 512
async_read_submit[log] = on
async_write_submit_blocks[log] = all
min_submit_batch_size[log] = 16
async_write_submit_active[log] = auto
```

[multidb]

```
mode = multidb
database_isolation = low
singletenant = yes
```

[persistence]

```
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
```

This is the sample `global.ini` file used when validating this solution with Red Hat Enterprise Linux 8.2, Red Hat Enterprise Linux 8.4, Red Hat Enterprise Linux 8.6, and Red Hat Enterprise Linux 8.8 as a TDI solution.

[communication]

```
tcp_backlog = 2048
```

[fileio]

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

[multidb]

```
mode = multidb
database_isolation = low
singletenant = yes
```

[persistence]

```
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
```

This is the sample `global.ini` file used when validating this solution with Red Hat Enterprise Linux 8.1 as an appliance solution.

[communication]

```
tcp_backlog = 2048
```

[fileio]

```
max_parallel_io_requests[data] = 128
max_submit_batch_size[data] = 64
size_kernel_io_queue[data] = 512
async_read_submit[data] = on
async_write_submit_blocks[data] = all
min_submit_batch_size[data] = 16
async_write_submit_active[data] = auto
max_parallel_io_requests[log] = 128
max_submit_batch_size[log] = 64
```

```
size_kernel_io_queue[log] = 512
async_read_submit[log] = on
async_write_submit_blocks[log] = all
min_submit_batch_size[log] = 16
async_write_submit_active[log] = auto
```

[multidb]

```
mode = multidb
database_isolation = low
singletenant = yes
```

[persistence]

```
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
```

Product descriptions

The following information describes the hardware and software components used in this reference architecture.

Hitachi Advanced Server DS220

With a combination of two Intel Xeon Scalable processors and high storage capacity in a 2U rack-space package, Hitachi Advanced Server DS220 delivers the storage and I/O to meet the needs of converged solutions and high-performance applications in the data center.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 24 slots for high-speed DDR4 memory, allowing up to 3 TB of memory per node when 128 GB DIMMs are used. This server supports up to 12 large form factor storage devices and an additional 2 small form factor storage devices.

This server has three storage configuration options:

- 12 large form factor storage devices and an additional 2 small form factor storage devices in the back of the chassis
- 16 SAS or SATA drives, 8 NVMe drives, and an additional 2 small form factor storage devices in the back of the chassis
- 24 SFF devices and an additional 2 SFF storage devices in the back of the chassis

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 consolidating 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

Hitachi 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. Move these workloads into an enterprise production environment seamlessly, saving money while reducing support and management costs.

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. For more information, see https://www.sap.com/products/hana.html.

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

As a 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 for basic concepts and for planning.

- SAP HANA Server Installation and Update Guide

  Use the various installation guides to install the required SAP In-Memory Database and the other software components for the different replication technologies.

- SAP HANA Administration Guide

  This provides the central operations documentation for the on-premises deployment of the SAP HANA Platform.

SAP HANA hardware directory provides information about SAP HANA appliances certified by SAP hardware partners.

Operating system options for SAP HANA

SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux for SAP HANA are available operating systems when running SAP HANA.

- SUSE Linux Enterprise Server (SLES) for SAP Applications

  Compete more effectively through improved uptime, better efficiency, and accelerated innovation using SUSE Linux Enterprise Server for SAP Applications. This is a versatile server operating system for efficiently deploying highly available enterprise-class IT services in mixed IT environments with performance and reduced risk.

  SUSE Linux Enterprise Server was the first Linux operating system to be certified for use with SAP HANA. It remains the operating system of choice for most SAP HANA customers.

- Red Hat Enterprise Linux (RHEL) for SAP HANA

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

Changing the configuration settings is only supported along the guidelines of SAP and the operating system distributor and may otherwise cause significant performance problems. The following SAP Notes for SUSE Linux Enterprise Server and Red Hat Enterprise Linux are a good starting point for information on this topic:

- 1944799 - SAP HANA Guidelines for SLES Operating System Installation
- 2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL) Operating System

For more details, see "Updating and Patching the Operating System" by searching in the "View SAP HANA document" from Technical Information and Best Practices.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact