

# **Hitachi Unified Compute Platform 1000 for the SAP HANA Platform in a Scale-up Configuration Using a Rack Optimized Server for Solutions, 2U Single Node**

---

Reference Architecture Guide

© 2021 Hitachi Vantara LLC. All rights reserved.

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

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

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

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

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

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

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

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

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, 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, the Microsoft Corporate 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 at <https://www.hitachivantara.com/en-us/company/legal.html>.

## Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@HitachiVantara.com](mailto: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-168-04	Support for RHEL 7.9	September 21, 2021
MK-SL-168-03	Updated Engineering validation section with HANA HDB parameter used for RHEL 7.7	September 25, 2020
MK-SL-168-02	Updated for SLES 12 SP5	August 28, 2020

# Reference Architecture Guide


This scale-up converged solution is for Hitachi Unified Compute Platform 1000 for the SAP HANA Platform (UCP for SAP HANA). It uses either internal storage housed in the server or an external storage subsystem.

This SAP HANA infrastructure uses the following components:

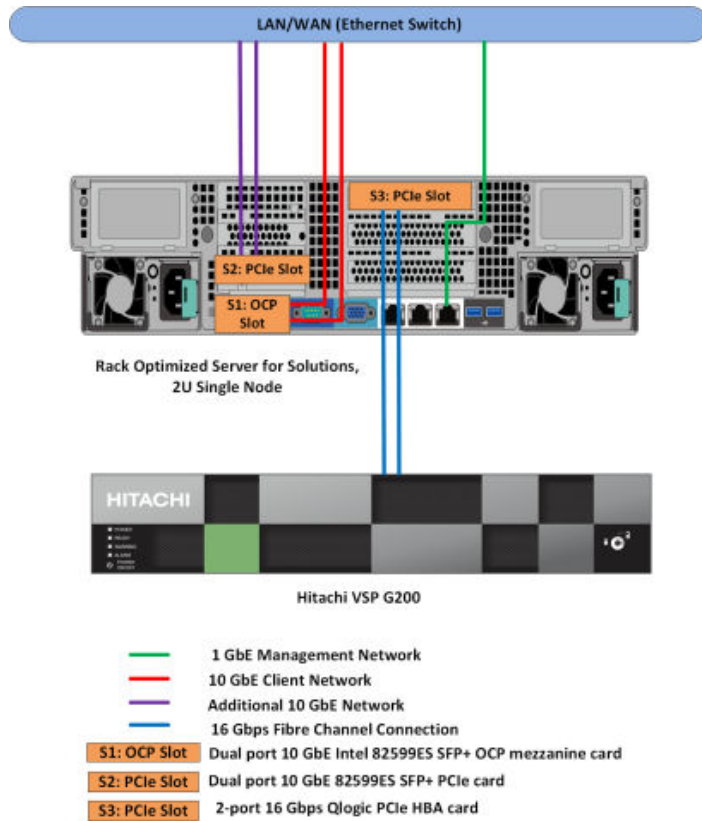
- Hardware
  - One rack optimized server for solutions, 2U single node, with or without drives providing internal storage
  - (Optional) External sub-system storage, such as Hitachi Virtual Storage Platform G200 (VSP G200) following the tailored data center integration (TDI) approach to implement SAP HANA
- Software
  - Preconfigured with SAP HANA to provide a converged solution for real time analytics

The validation of this environment with external storage was with Virtual Storage Platform G200. Your needs may require other 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 the rack optimized server for solutions, 2U single node.



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



Unified Compute Platform for the SAP HANA Platform is a preconfigured analytical system. 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 Scale-Up Configuration Sizes**

Number of Sockets	RAM Size	Storage
2 sockets	<ul style="list-style-type: none"> <li>▪ 64 GB (non-production only)</li> <li>▪ 128 GB</li> <li>▪ 256 GB</li> <li>▪ 384 GB</li> <li>▪ 512 GB</li> <li>▪ 768 GB</li> <li>▪ 1024 GB</li> <li>▪ 1536 GB</li> </ul>	Internal storage option:
		<ul style="list-style-type: none"> <li>▪ Internal drives on the rack optimized server for solutions, 2U single node</li> </ul>
		External storage option:
		<ul style="list-style-type: none"> <li>▪ Storage subsystem, such as Hitachi Virtual Storage Platform G200</li> </ul>



**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 Unified Compute Platform 1000 for SAP HANA for the different sized solutions.

Table 2 Hardware Elements

Hardware	Quantity	Configuration	Role
<p>Rack Optimized Server for Solutions, 2U single node (all solutions)</p> <p>This server has storage drives and network cards added, based on whether using an implementation with internal storage or an external storage subsystem</p>	1	<ul style="list-style-type: none"> <li>▪ CPU for the SAP HANA node, use only one of these processor types in an installation:               <ul style="list-style-type: none"> <li>• 2 Intel Xeon E5-2680 v3 processors (12 C, 2.5 GHz, 120 W)</li> <li>• 2 Intel Xeon E5-2680 v4 processors (14 C, 2.3 GHz, 120 W)</li> </ul> </li> <li>▪ 2 Heat sinks CPU 0/1</li> <li>▪ RAM per SAP HANA node:               <ul style="list-style-type: none"> <li>• 64 GB (4 × 16 GB DIMMS) for non-production use only</li> <li>• 128 GB (8 × 16 GB DIMMS)</li> <li>• 256 GB (16 × 16 GB DIMMS)</li> <li>• 384 GB (24 × 16 GB DIMMS)</li> <li>• 512 GB (16 × 32 GB DIMMS)</li> <li>• 768 GB (24 × 32 GB DIMMS)</li> <li>• 1024 GB (16 × 64 GB DIMMS)</li> <li>• 1536 GB (24 × 64 GB DIMMS)</li> </ul> </li> </ul>	SAP HANA server
Hitachi Virtual Storage Platform G200 (external storage implementations only)	1	<ul style="list-style-type: none"> <li>▪ Single frame</li> </ul>	Block storage when using an external storage sub-system
Storage drives (Internal storage implementations only)	9	<ul style="list-style-type: none"> <li>▪ 1.2 TB, 10k RPM, 6 Gbps SAS internal hard drives</li> </ul>	Block storage on the rack optimized server for solutions, 2U single node,

Hardware	Quantity	Configuration	Role
	2	<ul style="list-style-type: none"> <li>400 GB SATA solid state drives</li> </ul>	when using internal storage
RAID controller card (Internal storage implementations only)	1	<ul style="list-style-type: none"> <li>LSI SAS 3108 12 Gbps 1 GB RAID controller mezzanine card</li> </ul>	Required on the rack optimized server for solutions, 2U single node, when using internal storage
	1	<ul style="list-style-type: none"> <li>Activation key/card for LSI MegaRAID CacheCade Pro 2.0</li> </ul>	
OCP mezzanine network card (all implementations)	1	<ul style="list-style-type: none"> <li>Dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card</li> </ul>	For SAP HANA 10 GbE client network
	2	<ul style="list-style-type: none"> <li>10 Gbps 850nm Multimode Datacom SFP+ transceiver</li> </ul>	
Intel PCIe network cards (all implementations)	1	<ul style="list-style-type: none"> <li>Dual port 10 GbE 82599ES SFP+ PCIe card</li> </ul>	For additional 10 GbE network
	2	<ul style="list-style-type: none"> <li>Ethernet SFP+ SR optics dual rate 10GBASE-SR/1000BASE-SX</li> </ul>	
Qlogic PCIe HBA card (external storage solution implementations only)	1	<ul style="list-style-type: none"> <li>2-port 16 Gbps card</li> </ul>	Connectivity to the external storage sub-system
	2	<ul style="list-style-type: none"> <li>16 Gbps SFP</li> </ul>	
Brocade ICX 6430 24 port switch (optional)	1	<ul style="list-style-type: none"> <li>24 × 1 GbE ports</li> </ul>	Optional switch for management network
Brocade VDX 6740 48 port switches (optional)	2	<ul style="list-style-type: none"> <li>48 × 10 GbE ports</li> </ul>	Optional switches for the client network or additional backup network
Minkels Global Solutions Rack (optional)	1	<ul style="list-style-type: none"> <li>1 standard rack</li> </ul>	Optional rack for mounting server
PDUs (optional)	2	<ul style="list-style-type: none"> <li>12 outlet vertical PDUs</li> </ul>	Optional PDUs for solution

## Rack optimized server for solutions, 2U single node

The rack optimized server for solutions, 2U single node, is a rack-mounted server designed for optimal performance and power efficiency. The highly scalable memory capacity supports up to 1536 GB. It is powered by the Intel Xeon E5-2680 v3 processor or E5-2680 v4 processor for complex and demanding workloads. It supports flexible OCP and PCIe I/O expansion card options.

The configuration uses one rack optimized server for the different sized solutions listed in Table 1.

When implementing this solution using internal storage, add the drives listed in Table 2 to the rack optimized server for solutions, 2U single node.

## Hitachi Virtual Storage Platform Gx00 Models

Hitachi Virtual Storage Platform Gx00 models are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage systems. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

When validating this environment, Hitachi Virtual Storage Platform G200 was used as the external storage subsystem.

## Software elements

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

**Table 3 Software Elements**

Software	
Operating system	SUSE Linux Enterprise Server for SAP Applications
	Red Hat Enterprise Linux (alternate)
SAP HANA	
LSI MegaRAID CacheCade (internal storage solution implementation only)	



## Solution design

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


- [Rack optimized server for solutions, 2U single node, configuration \(on page 9\)](#)
- [Network architecture configuration \(on page 10\)](#)
- [LSI MegaRAID CacheCade configuration \(internal storage implementation only\) \(on page 10\)](#)
- [Storage architecture configuration \(on page 11\)](#)
- [SAP HANA configuration \(on page 14\)](#)

### Rack optimized server for solutions, 2U single node, configuration

The following figure shows the front and back view of the rack optimized server for solutions, 2U single node.



Rack Optimized Server for Solutions, 2U Single Node (Front)



Rack Optimized Server for Solutions, 2U Single Node (Back)

S1: OCP Slot	Dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card
S2: PCIe Slot	Dual port 10 GbE Intel 82599ES SFP+ PCIe card
S3: PCIe Slot	2-port 16 Gbps Qlogic PCIe HBA card (required for external storage only)

The solution uses one rack optimized server for solutions, 2U single node, with the following components:

- One dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card
- One Intel dual port 10 GbE Intel 82599ES SFP+ PCI card
- Required only with the internal storage option:
  - One LSI SAS 3108 12 Gbps 1 GB RAID controller mezzanine card
  - Nine 1.2 TB, 10k RPM, 6 Gbps, SAS hard drives
  - Two 400 GB SATA solid-state drives
- Required only with the external storage option:
  - One 2-port 16 Gbps Qlogic PCIe HBA card

## Network architecture configuration

Connect the 1 GbE management port on the rack optimized server for solutions, 2U single node, to a Brocade ICX 6430 24 port 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:

- Connect the following to Brocade VDX 6740 48 port switches or to any other external switches:
  - Port 0 of dual port 10GbE Intel 82599ES SFP+ OCP mezzanine card
  - Port 1 of dual port 10GbE Intel 82599ES SFP+ OCP mezzanine card
- 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=2 lacp_rate=fast
```

This acts as the client network for the SAP HANA node.

- Configure the external switch ports with a short LACP timeout value to speed up detecting corrupted connections.
- Make the connections in the following table on the additional dual port 10 GbE Intel 82599ES SFP+ PCI card, providing two additional ports for the 10 GbE network.

**Table 4 Network Setup**

Network Card	Port	Network Description
Slot 1	0	Client network for the SAP HANA node
	1	Client network for the SAP HANA node
Slot 2	0	Free for use as additional 10 GbE network for backup, SAP HANA system replication, or other purposes
	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 separate network switches, depending on the network environment.

## LSI MegaRAID CacheCade configuration (internal storage implementation only)



**Note:** Only use the LSI MegaRAID CacheCade when using the internal storage option.

Use two 400 GB SATA solid state drives to create a virtual drive using LSI MegaRAID CacheCade.

The use of MegaRAID CacheCade improves application performance. It expands the MegaRAID read caching capacity and introduces write caching.

CacheCade uses high-performing solid-state drives as a secondary tier of cache to provide faster reads and writes to maximize transactional I/O performance.

## Storage architecture configuration

This describes the two storage configurations for this environment:

- [Internal Storage Configuration \(on page 11\)](#)
- [External Storage Configuration \(on page 12\)](#)

### Internal storage option

These are the storage components needed to setup a scale-up configuration for SAP HANA using internal storage with the rack optimized server for solutions, 2U single node:

- One LSI SAS 3108 12 Gbps 1 GB RAID controller mezzanine card
- Nine 1.2 TB, 10k RPM, 6 Gbps SAS hard drives
- Two 400 GB SATA solid state drives

The sizing and configuring of the storage is the same for different memory sizes. This includes 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 5 Storage Configuration When Using Internal Storage**

Storage	Size and Type of Storage	Purpose
One RAID6 (6D+2P) parity group	8 × 1.2 TB 10k RPM SAS hard disk drives	<ul style="list-style-type: none"> <li>▪ Operating system volume: 100 GB</li> <li>▪ SAP HANA shared volume: 1.5 TB</li> <li>▪ Data volume: 4 TB</li> <li>▪ Log volume: 600 GB</li> </ul>
	1 × 1.2 TB 10k RPM SAS hard disk drive	Spare drive
One RAID1 parity group	2 × 400 GB SATA solid state drives	LSI MegaRAID CacheCade virtual drive

## External storage option



**Note:** Each implementation of this reference architecture can use a different storage architecture. Validation testing for this environment used for external storage used Hitachi Virtual Storage Platform G200. 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 a rack optimized server for solutions, 2U single node, using Hitachi Virtual Storage Platform G200 for external storage:

- One 2-port 16 Gbps Qlogic PCIe card
- One Hitachi Virtual Storage Platform G200
- Storage drive box trays (DBS)
- Spare drives

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 layout.

**Table 6 Dynamic Provisioning Pools**

Dynamic Provisioning Pool Name	Purpose	Parity Group RAID Level and Disks
OS_SH_Data_Pool	<ul style="list-style-type: none"> <li>▪ Operating system LUN</li> <li>▪ SAP HANA shared LUN</li> <li>▪ Data LUN</li> </ul>	RAID6 (14D+2P) on 600 GB, 10k RPM SAS drives
Log_Pool	<ul style="list-style-type: none"> <li>▪ Log LUN</li> </ul>	RAID6 (6D+2P) on 600 GB, 10k RPM SAS drives

This solution uses the dynamic provisioning pool layout as listed in the following table on Virtual Storage Platform G200.

Table 7 Dynamic Provisioning Pool Provisioning

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	RAID6 (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	RAID6 (6D +2P) on 600 GB, 10k RPM SAS drives	00:01:01	Log_Pool_1	750	MPU-10
			00:01:02	Log_Pool_2	750	MPU-11
			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, SAPA HANA shared, data, and log volumes as listed in the following table.

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	1600 GB	MPU-20
Log_Pool	00:02:02	HANA_LOG_1	150 GB	MPU-10
	00:02:03	HANA_LOG_2	150 GB	MPU-11
	00:02:04	HANA_LOG_3	150 GB	MPU-20
	00:02:05	HANA_LOG_4	150 GB	MPU-21

Dynamic Provisioning Pool	Virtual Volume ID	Virtual Volume Name	Virtual Volume Size	MPU Assignment
OS_SH_Data_Pool	00:02:06	HANA_DATA_1	400 GB	MPU-10
	00:02:07	HANA_DATA_2	400 GB	MPU-11
	00:02:08	HANA_DATA_3	400 GB	MPU-20
	00:02:09	HANA_DATA_4	400 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
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 section describes how to configure SAP HANA in this solution.

### File system

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

- Operating system volume
- SAP HANA shared volume
- Data volume
- Log volume

## Device-Mapper Multipath

If implementing this solution with internal storage on the rack optimized server for solutions, 2U single node, deactivate Device-Mapper Multipath.

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

## SAP HANA software installation

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

Install the following SAP HANA software components on the server:

- Database
- Client
- Agent

## Engineering validation

This scale-up converged solution is for Hitachi Unified Compute Platform 1000 for the SAP HANA platform with internal drives or external storage on SLES 12 SP4, SLES 12 SP5, RHEL 7.6, RHEL 7.7, and RHEL 7.9. This involved the following validation tasks:

- SAP HANA Hardware Configuration Check Tool (HWCCT) using FSPERF revision hwcct-212\_5 tested on these volumes for SLES 12 SP4, SLES 12 SP5 and RHEL 7.6:
  - Data volume
  - Log volume
- Appliance configuration with internal drive was tested using HWCCT 212\_5, using the fsperf\_single\_mix\_107.sh script for SLES 12 SP4, SLES 12 SP5, RHEL 7.6, RHEL 7.7, and RHEL 7.9:
  - The MixedIO on Data/Log Volume (A MixedIO test is mandatory if the data volume and log volume are created in the same parity group. Otherwise this test is optional.)

For optimal use of the system with a SAP HANA database, use the parameters listed in the corresponding appendix 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 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 SLES 12 SP4 and SLES 12 SP5 as an appliance solution using internal drives or as a TDI solution using external storage.

```
[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 7.6, Red Hat Enterprise Linux 7.7, and Red Hat Enterprise Linux 7.9 as an appliance solution using internal drives or as a TDI solution using external storage.

```
[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

```

## Product descriptions

These products are used in this reference architecture.

### 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](http://HitachiVantara.com) | [community.HitachiVantara.com](http://community.HitachiVantara.com)

Contact Information  
USA: 1-800-446-0744  
Global: 1-858-547-4526  
[HitachiVantara.com/contact](http://HitachiVantara.com/contact)