

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

By Archana Kuppuswamy, Prasad Patkar, Abhishek Dhanuka

April 2017

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.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.

Table of Contents

- Key Solution Elements 3**
 - Hardware Elements..... 3
 - Software Elements..... 6
- Solution Design..... 8**
 - Rack Optimized Server for Solutions, 2U Single Node, Configuration..... 8
 - Network Architecture Configuration 9
 - LSI MegaRAID CacheCade Configuration (Internal Storage Implementation Only) 10
 - Storage Architecture Configuration 10
 - SAP HANA Configuration..... 13

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

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 sub-system.

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)
- **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.

Figure 1 shows the topology of this reference solution using internal drives on the rack optimized server for solutions, 2U single node. Figure 2 shows the topology of this reference solution using external drives on a storage subsystem, such as Virtual Storage Platform G200.

Figure 1

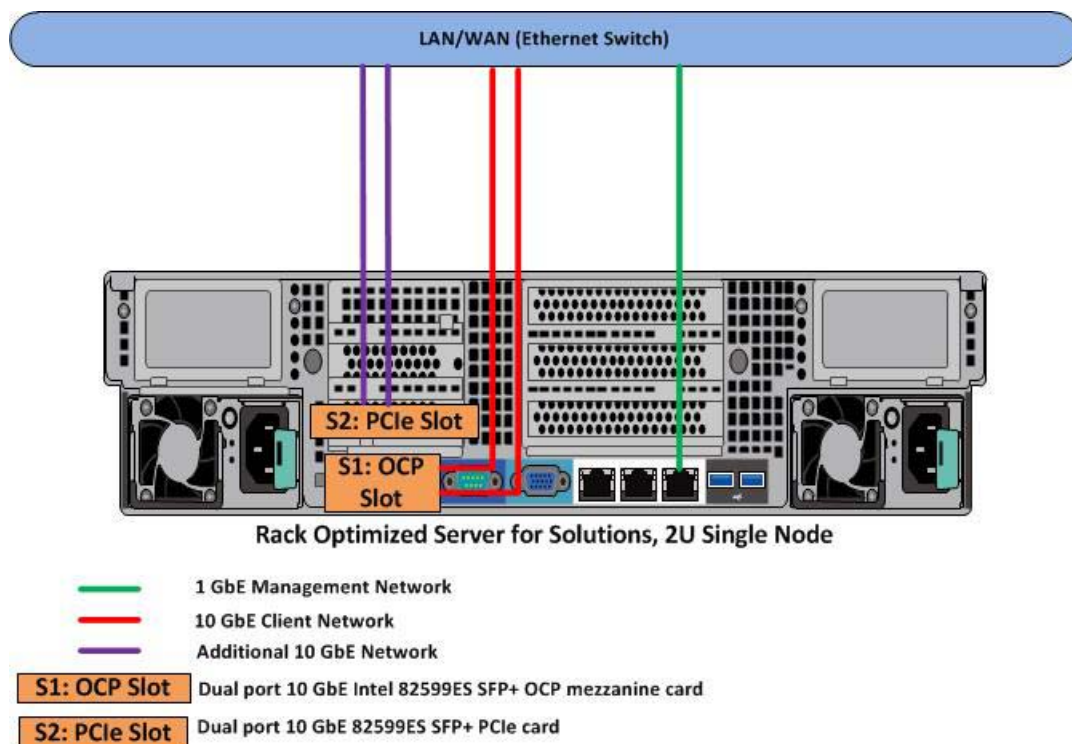
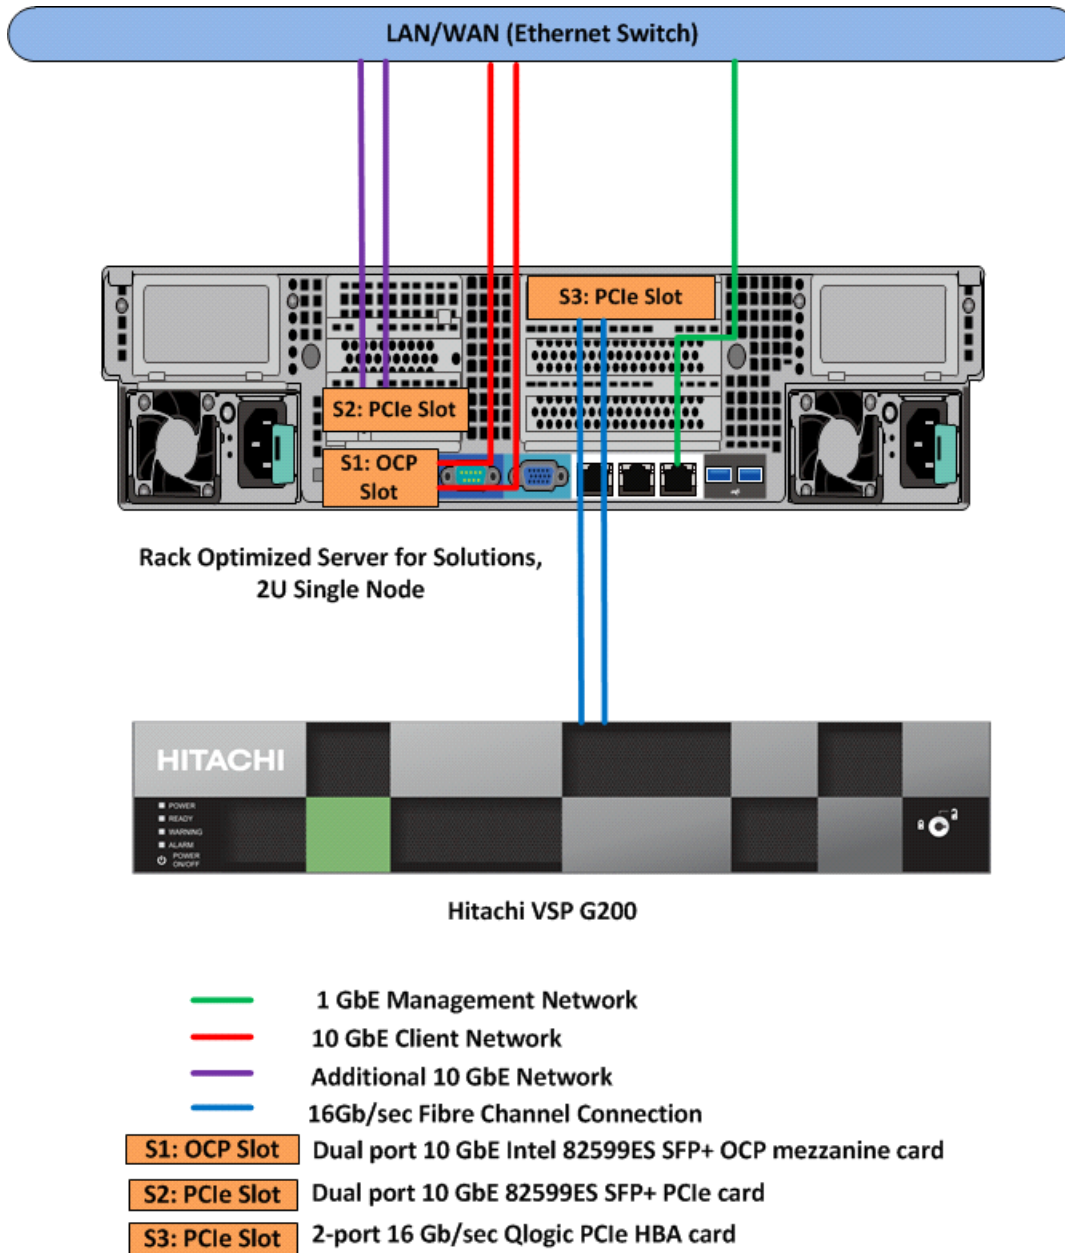


Figure 2



Unified Compute Platform for SAP HANA 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 Table 1.

TABLE 1. SUPPORTED SCALE-UP CONFIGURATION SIZES

Number of Sockets	RAM Size	Storage
2 sockets	<ul style="list-style-type: none">64 GB (non-production only)128 GB256 GB384 GB512 GB768 GB1024 GB1536 GB	Internal storage option: <ul style="list-style-type: none">Internal drives on the rack optimized server for solutions, 2U single node
		External storage option: <ul style="list-style-type: none">Storage subsystem, such as Hitachi Virtual Storage Platform G200

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 Data Systems account representative. Changing this layout can require manual configuration of the network and/or using different components.

Hardware Elements

Table 2 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 sub-system</p>	1	<ul style="list-style-type: none"> ■ CPU for the SAP HANA node, use only one of these processor types in an installation: <ul style="list-style-type: none"> ■ 2 Intel Xeon E5-2680 v3 processors (12 C, 2.5 GHz, 120 W) or ■ 2 Intel Xeon E5-2680 v4 processors (14 C, 2.3 GHz, 120 W) ■ 2 Heat sinks CPU 0/1 ■ RAM per SAP HANA node: <ul style="list-style-type: none"> ■ 64 GB (4 × 16 GB DIMMS) for non-production use only ■ 128 GB (8 × 16 GB DIMMS) ■ 256 GB (16 × 16 GB DIMMS) ■ 384 GB (24 × 16 GB DIMMS) ■ 512 GB (16 × 32 GB DIMMS) ■ 768 GB (24 × 32 GB DIMMS) ■ 1536 GB (24 × 64 GB DIMMS) ■ 1024 GB (16 × 64 GB DIMMS) ■ 1536 GB (24 × 64 GB DIMMS) 	SAP HANA server
Hitachi Virtual Storage Platform G200 (external storage implementations only)	1	<ul style="list-style-type: none"> ■ Single frame 	Block storage when using an external storage sub-system

TABLE 2. HARDWARE ELEMENTS (CONTINUED)

Hardware	Quantity	Configuration	Role
Storage drives (Internal storage implementations only)	9	■ 1.2 TB, 10k RPM, 6 Gb/sec SAS internal hard drives	Block storage on the rack optimized server for solutions, 2U single node, when using internal storage
	2	■ 400 GB SATA solid state drives	
RAID controller card (Internal storage implementations only)	1	■ LSI SAS 3108 12 Gb/sec 1 GB RAID controller mezzanine card	Required on the rack optimized server for solutions, 2U single node, when using internal storage
	1	■ Activation key/card for LSI MegaRAID CacheCade Pro 2.0	
OCP mezzanine network card (all implementations)	1	■ Dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card	For SAP HANA 10 GbE client network
	2	■ 10 Gb/sec 850nm Multimode Datacom SFP+ transceiver	
Intel PCIe network cards (all implementations)	1	■ Dual port 10 GbE 82599ES SFP+ PCIe card	For additional 10 GbE network
	2	■ Ethernet SFP+ SR optics dual rate 10GBASE-SR/1000BASE-SX	
Qlogic PCIe HBA card (external storage solution implementations only)	1	■ 2-port 16 Gb/sec card	Connectivity to the external storage sub-system
	2	■ 16 Gb/sec SFP	
Brocade ICX 6430 24 port switch (optional)	1	■ 24 × 1 GbE ports	Optional switch for management network
Brocade VDX 6740 48 port switches (optional)	2	■ 48 × 10 GbE ports	Optional switches for things such as the client network or additional backup network
Minkel Global Solutions Rack (optional)	1	■ 1 standard rack	Optional rack for mounting server
PDUs (optional)	2	■ 12 outlet vertical PDUs	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 arrays. 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 sub-system.

Software Elements

Table 3 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)	

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.

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 on SAP HANA at the [SAP Service Marketplace](#). See the installation and upgrade guides download section for SAP In-Memory Computing (SAP In-Memory Appliance — SAP HANA). The following are available:

- **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 Installation and Initial Configuration Guides**

Use the various installation guides to install the required SAP In-Memory Database and the other software components for the different replication technologies. Refer to the [SAP HANA Server Installation Guide](#) for an overview on how to install SAP HANA.

- **SAP HANA Technical Operations Manual**

This provides an end-to-end picture of the available SAP HANA administration tools and the key tasks for a system administrator to perform.

- **SAP HANA Master Update Guide**

This explains how to update SAP HANA and its components.

- **SAP Integration and Certification Center (SAP ICC)**

This page provides information about SAP HANA appliances certified by SAP hardware partners.

This is a link to all SAP-related documentation: http://help.sap.com/hana_platform/

Operating System Choices

The scale-up configuration of Hitachi Unified Compute Platform for the SAP HANA Platform can run on the following Linux operating systems:

- **SUSE Linux Enterprise Server (SLES) for SAP Applications**

Compete more effectively through improved uptime, better efficiency, and accelerated innovation using [SUSE Linux Enterprise Server](#). 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 the vast majority of SAP HANA customers.

- **Red Hat Enterprise Linux (RHEL)**

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.

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 SLES and RHEL 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](#)

The initially delivered configuration of the operating system should persist. Do not make any modifications to the operating system, except as noted or approved by SAP.

Solution Design

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

- Rack Optimized Server for Solutions, 2U Single Node, Configuration
- Network Architecture Configuration
- LSI MegaRAID CacheCade Configuration (Internal Storage Implementation Only)
- Storage Architecture Configuration
- SAP HANA Configuration


Rack Optimized Server for Solutions, 2U Single Node, Configuration

Figure 3 shows the front and back view of the rack optimized server for solutions, 2U single node.

Figure 3



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 82599ES SFP+ PCIe card
- S3: PCIe Slot** 2-port 16 Gb/sec 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 Gb/sec 1GB RAID controller mezzanine card
 - Nine 1.2 TB, 10k RPM, 6 Gb/sec, SAS hard drives
 - Two 400 GB SATA solid-state drives
- Required only with the external storage option:
 - One 2-port 16 Gb/sec 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=layer3+4 updelay=5000 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 Table 4 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

Figure 1 shows the standard network configuration used for this solution when using internal storage. Figure 2 shows the standard network configuration for this solution when using external storage.

Note — The management network and client network can be on the same network switch or a 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 Option
- 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 the rack optimized server for solutions, 2U single node:

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

The sizing and configuring of the storage is 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).

Table 5 lists the storage configuration for different memory sizes.

TABLE 5. STORAGE CONFIGURATION WHEN USING INTERNAL STORAGE

Storage	Size and Type of Storage	Purpose
One RAID-6 (6D+2P) parity group	8 × 1.2 TB 10k RPM SAS hard disk drives	<ul style="list-style-type: none"> ■ Operating system volume: 100 GB ■ SAP HANA shared volume: 1.5 TB ■ Data volume: 4 TB ■ Log volume: 600 GB
	1 × 1.2 TB 10k RPM SAS hard disk drive	Spare drive
One RAID-1 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 Gb/sec 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 Table 6 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"> ■ Operating system LUN ■ SAP HANA shared LUN ■ Data LUN 	RAID-6 (14D+2P) on 600 GB, 10k RPM SAS drives
Log_Pool	<ul style="list-style-type: none"> ■ Log LUN 	RAID-6 (6D+2P) on 600 GB, 10k RPM SAS drives

This solution uses the dynamic provisioning pool layout in Table 7 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	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 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 following Table 8.

TABLE 8. VIRTUAL VOLUMES FOR THE SAP HANA NODES

Dynamic Provisioning Pool	Virtual Volume ID	Virtual Volume Name	Virtual Volume Size	MPU Assianment
OS_SH_Data_Pool	00:02:00	HANA_OS	100 GB	MPU-10
	00:02:01	HANA_SH	1600 GB	MPU-20

TABLE 8. VIRTUAL VOLUMES FOR THE SAP HANA NODES

Dynamic Provisioning Pool	Virtual Volume ID	Virtual Volume Name	Virtual Volume Size	MPU Assianment
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
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

Table 9 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 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 sub-system:

- 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 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 latest SAP HANA SPS stack on the server.

Install the following SAP HANA software components on the server:

- Database
- Client
- Agent

For More Information

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

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the [Resources](#) website.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. HDS Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Academy [Training and Certification](#) website.

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

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

HITACHI is a trademark or registered trademark of Hitachi, Ltd., Other notices if required. All other trademarks, service marks and company names are properties of their respective owners.

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

AS-486-05, April 2017.