

Hitachi Unified Compute Platform for the SAP HANA Platform with Logical Partitioning in a Scale-Up Configuration Using Hitachi Compute Blade 2500 and Hitachi Virtual Storage Platform

Reference Architecture Guide

By Stephan Kreitz, David Pascuzzi, and Francois Laforgia

July 2016

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.

Contents

Key Solution Elements	3
Hardware Elements.....	3
Software Elements.....	4
Solution Design	6
Logical Partitioning Supported Solutions	6
Hitachi Compute Blade 2500 Configuration	7
Server Blade Architecture	8
Storage Architecture	9
Network Architecture	11
Management Server.....	13
SAP HANA Node	14

Hitachi Unified Compute Platform for the SAP HANA Platform with Logical Partitioning in a Scale-Up Configuration Using Hitachi Compute Blade 2500 and Hitachi Virtual Storage Platform

Reference Architecture Guide

Hitachi Unified Compute Platform for the SAP HANA Platform with logical partitioning (LPAR) in a scale-up configuration is a pre-configured virtual appliance ready to plug into a network to provide real-time access to operational data for use in analytic models.

This reference architecture guide describes how to deploy Unified Compute Platform for SAP HANA Platform with LPAR in a scale-up configuration.

This reference architecture uses the following:

- **Hitachi Compute Blade 2500**

This has enterprise computing power and performance with flexible I/O architecture and logical partitioning. Multiple applications easily and securely co-exist in the same chassis.

This solution uses a varying number of 520X B2 server blades to provide the flexibility of multiple LPAR configuration options.

- **Hitachi Virtual Storage Platform Gx00 Models**

These scale for all data types, flexibly adapting for performance, capacity, and multi-vendor storage.

In this solution, the persistent storage of the SAP HANA server resides on Hitachi Virtual Storage Platform G600 (VSP G600)

- **Symmetric Multiprocessing Connector** - Connects multiple blades together into one unit

- **Brocade ICX 6430-24 switch** — 24-port 1 GbE switch that provides a management network to the appliance.

- **Brocade VDX 6740-48 switch** — 48-port switch that provides 10 GbE external connectivity to the appliance.

- **QuantaPlex T41S-2U server** — An ultra-dense design equipped with four independent nodes with the flexibility to set up different workloads independently in one 2U shared infrastructure. This solution uses one node of this server as a management server.

- **10 Gb 2-port LAN PCIe adapters**

- **Hitachi 16 Gb PCI-FX Fibre Channel adapters**

- **Logical Partitioning (LPAR)**

A feature of Hitachi Compute Blade family, this firmware-based hypervisor creates logical partitions on one physical server.

- SUSE Linux Enterprise Server
- Red Hat Enterprise Linux
- SAP HANA

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

Figure 1 depicts the hardware used in this solution.

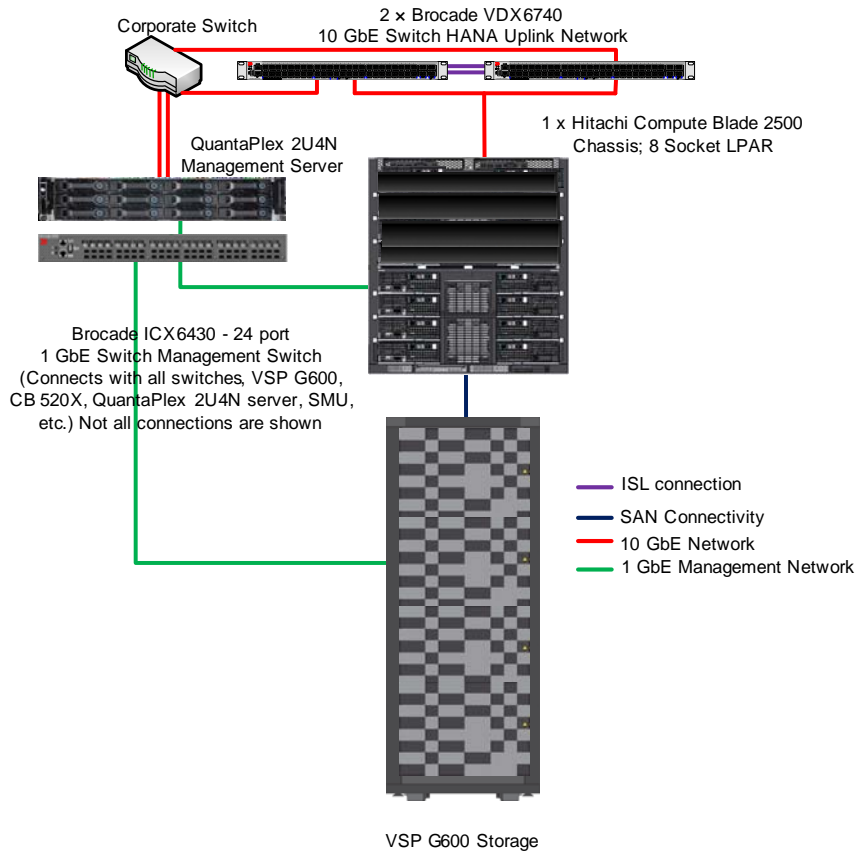


Figure 1

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

Hardware Elements

Based on Hitachi Unified Compute Platform for the SAP HANA Platform, it is possible to create scale-up configurations with logical partitioning.

Hitachi Compute Blade 2500

[Hitachi Compute Blade 2500](#) delivers enterprise computing power and performance with unprecedented scalability and configuration flexibility. Lower your costs and protect your investment.

Flexible I/O architecture and logical partitioning allow configurations to match application needs exactly with Hitachi Compute Blade 2500. Multiple applications easily and securely co-exist in the same chassis.

Add server management and system monitoring at no cost with Hitachi Compute Systems Manager. Seamlessly integrate with Hitachi Command Suite in Hitachi storage environments.

Different sizes of this solution use a different number of 520X B2 server blades (one, two, or four) in Hitachi Compute Blade 2500.

Symmetric Multiprocessing Connector

For multiple server blades, the solution uses symmetric multiprocessing (SMP) technology to combine multiple server blade resources into a single server.

The 520X B2 server blades use SMP in one of two ways, depending on the size of the solution:

- Combine two server blades with a 2-blade SMP connection board.
- Combine four server blades with a 4-blade SMP connection board.

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. New management software improves ease of use to save time and reduce complexity. The infrastructure of Storage Virtualization Operating System creates a management framework for improved IT response to business demands.

On Hitachi storage systems, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Dynamic Provisioning is like using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a logical size you specify of up to 60 TB created against it without allocating any physical space initially.

Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. Dynamic Provisioning distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Quanta Cloud Technology QuantaPlex T41S-2U Server

The QuantaPlex T41S-2U is an ultra-dense design equipped with four independent nodes. It creates the flexibility to set up different workloads independently in one 2U shared infrastructure, providing optimal data center performance per dollar. This is an optional component for the solution.

Brocade Switches

[Brocade and Hitachi Data Systems](#) partners 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.

This solution can include the following Brocade products:

- Brocade VDX 6740-48 port switch
- Brocade ICX 6430-24 port switch

Software Elements

Logical Partitioning

The logical partitioning feature from Hitachi partitions the physical resources of one single-blade or multi-blade server logically into independent server environments called LPARs (logical partitions).

CPUs, memory, and I/O devices can be assigned dedicated mode to LPARs to create fully isolated server environments for SAP HANA single node installations within each LPAR running on the same physical host without any noisy neighbor effects between the different server environments.

LPARs using one, two or four CPU sockets in dedicated mode, accessing the CPU sockets' local memory bank, can be created for running production SAP HANA instances. Combined with dedicated ports on the HBA cards for each LPAR, and separated controller ports in Hitachi Virtual Storage Platform, the LPAR manager enables the consolidation of multiple production SAP HANA instances on the same physical server using a single Virtual Storage Platform.

SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux

The following options are available as guest operating system for the LPARs:

■ 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](#)

For more details, see “Updating and Patching the Operating System,” in the [SAP HANA Technical Operations Manual](#).

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.

SAP customers can download more information on the SAP HANA Platform 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 Platform \(Core\) Documentation](#).

This page is the central starting point to find documentation for SAP HANA platform.

- [SAP Integration and Certification Center \(SAP ICC\)](#)

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

Find Information on production SAP HANA systems on Hitachi systems with LPAR in [SAP Note 2063057](#).

For SAP Community Network documents on Hitachi systems with LPAR, use the following documents as a starting point:

- [SAP on LPAR](#)
- [Best Practices Guide for Hitachi LPAR](#)

Solution Design

The detailed design for Hitachi Unified Compute Platform for the SAP HANA Platform with logical partitioning (LPAR) in a scale-up configuration reference solution includes the following:

Logical Partitioning Supported Solutions

SAP HANA on Hitachi solutions with logical partitioning are based on bare-metal appliance configurations with modifications that are required for running logical partitions (LPARs). This document describes the configurations for production LPAR SAP HANA instances. Table 1 lists the supported configurations for SAP HANA with logical partitioning.

This document uses the term “CPU socket” to differentiate between the LPAR size and the bare metal system size. CPU socket does not refer to an empty CPU socket, but includes the processor.

A partition on the CPU core level instead of the CPU socket level is not supported for LPARs running production SAP HANA instances.

Table 1. Supported Scale-up Configurations for SAP HANA with Logical Partitioning

Solution Size	Small (S)	Medium (M)	Large (L)
Number of CPU Sockets	2	4	8
Number of Server Blades	1	2 (using a 2-blade SMP)	4 (using a 4-blade SMP)
Total Memory (in GB) for the solution	256, 512, 768, 1024, 1536	512, 1024, 1536, 2048, 3072	1024, 2048, 3072, 4096, 6144
Number of LPARs for Production Usage using G600 storage	Any combination of LPARs that totals two sockets	Any combination of LPARs that totals four sockets	Any combination of LPARs that totals eight sockets.
Maximum Memory Size per LPAR	1-socket: 50% of the memory	25% of the memory	12.5% of the memory
	2-socket: 100% of the memory	50% of the memory	25% of the memory
	4-socket: N/A	100% of the memory	50% of the memory

Hitachi Compute Blade 2500 Configuration

This solution uses one Hitachi Compute Blade 2500 with the following components:

The LPAR SAP HANA scale-up configurations can have one, two, or four 520X B2 server blades. Refer to Table 2 for supported configurations.

There are two management modules on Hitachi Compute Blade 2500 connected to the management network.

Table 2. Chassis Configuration

Feature	Small (2-Socket) Configuration	Medium (4-Socket) Configuration	Large (8-Socket) Configuration
Blades	1 x 520X B2	2 x 520X B2	4 x 520X B2
Blade Location	<ul style="list-style-type: none"> ■ Blade 1 (primary) 	<ul style="list-style-type: none"> ■ Blade 3 (non-primary) ■ Blade 1 (primary) 	<ul style="list-style-type: none"> ■ Blade 7 (non-primary) ■ Blade 5 (non-primary) ■ Blade 3 (non-primary) ■ Blade 1 (primary)
SMP	none	1 x 2 Blade SMP	1 x 4 Blade SMP
Network ports	1 × 2-port 10GBASE-SR LAN PCI-E adapter on two I/O board modules for each server blade		
Fibre Channel Ports	1 Hitachi 16 Gb/sec, 2-port Fibre Channel PCI-E adapters on two I/O board modules for each server blade		
Other interfaces	For all sizes: <ul style="list-style-type: none"> ■ 1 USB 3.0 port ■ KVM connector (VGA, COM, USB 2.0 port) 		

A maximum of 28 I/O board modules (IOBD) can be installed on one Hitachi Compute Blade 2500. Depending on the size of the solution, between 4 and 16 I/O board modules are used.

Figure 2 shows the layout of the I/O board modules from the back of the Hitachi Compute Blade 2500 chassis.

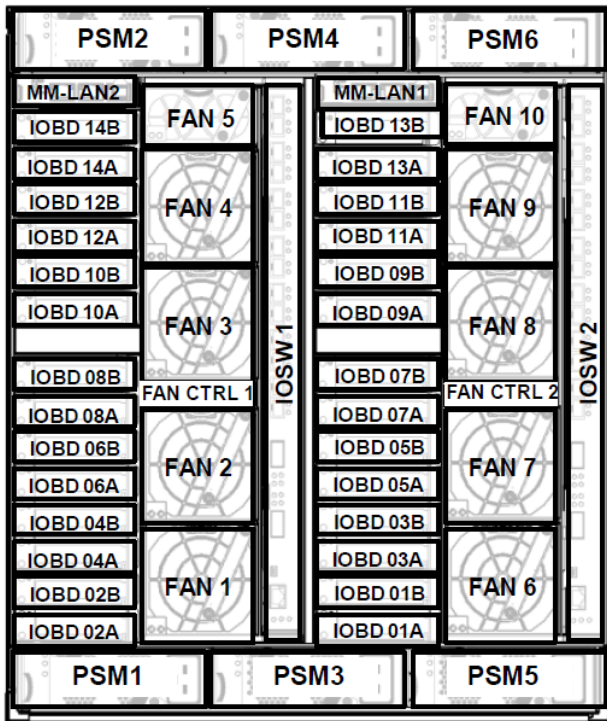


Figure 2

Server Blade Architecture

Each solution size has a different number of server blades, from one up to four full-width server blades. Multiple server blades use a two- or four-server blade SMP connection board to create a single four- or eight-socket SMP node with a maximum of 144 cores and different memory sizes. Table 3 shows the supported server blade configurations.

Table 3. Server Blade Configurations

	Small (2-Socket)	Medium (4-Socket)	Large (8-Socket)
Server Blades	1 server blade	2 server blades	4 server blades
Processor SKU	Intel Xeon E7-8880 v3 (all sizes)		
Processor Cores	18		
Processor Frequency	2.3 GHz		

Table 4 shows the different memory configurations.

Table 4. Memory Configurations

DIMMs	Memory Per Blade	Total Memory Small (2-Socket)	Total Memory Medium (4-Socket)	Total Memory Large (8-Socket)
16 × 16 GB DIMMs	256 GB	256 GB	512 GB	1024 GB
32 × 16 GB DIMMs	512 GB	512 GB	1024 GB	2048 GB
48 × 16 GB DIMMs	768 GB	768 GB	1536 GB	3072 GB
32 × 32 GB DIMMs	1024 GB	1024 GB	2048 GB	4096 GB
48 × 32 GB DIMMs	1536 GB	1536 GB	3072 GB	6144 GB

Storage Architecture

Hitachi Virtual Storage Platform G600 is used for this solution. Sizing and configuring of storage including storage drive box trays (DBS), spare drives, the operating system volume (OS), the existing SAP HANA shared volume (/hana/shared) varies for different size solutions by taking storage and I/O requirements into account.

Fibre Channel SAN Architecture

This solution uses between 4 and 16 × 16 Gb/sec Fibre Channel ports on Hitachi Virtual Storage Platform G600 directly attached to Hitachi Compute Blade 2500 via the Fibre Channel PCI-E adapters. Table 5 shows the Fibre Channel port mappings between the ports on Virtual Storage Platform G600 and the Fibre Channel adapters on the I/O board modules.

This configuration provides two dedicated paths from the LPAR within Hitachi Compute Blade 2500 to the ports on VSP G600 for each 1-socket LPAR, four dedicated paths for 2-socket LPARs and four dedicated paths for 4-socket LPARs.

Table 5. Fibre Channel Port Mapping

Server Blade Number	PCI-E Slot Number	Port Number	VSP G600 Port
Blade 1	IOBD 01A	0	1A
		1	2A
	IOBD 02A	0	3A
		1	4A
Blade 3	IOBD 03A	0	1B
		1	2B
	IOBD 04A	0	3B
		1	4B
Blade 5	IOBD 05A	0	1C
		1	2C
	IOBD 06A	0	3C
		1	4C

Table 5. Fibre Channel Port Mapping (Continued)

Server Blade Number	PCI-E Slot Number	Port Number	VSP G600 Port
Blade 7	IOBD 07A	0	1D
		1	2D
	IOBD 08A	0	3D
		1	4D

Shared Storage Architecture

Shared storage architecture can realize better return on investment and reduce total cost of ownership than having storage dedicated to one LPAR. This architecture increases the number of LPARs that can use the storage subsystem and reduces the number of disks required.

The following HDP pool configuration is used:

- One HDP Pool for OS, SAP HANA shared, and SAP HANA data
 - Depending on the LPAR system size, there are one or more parity groups with RAID-6 (14D + 2P) with 16 × 600 GB 10k RPM SAS drives
 - These raid groups are all assigned to one HDP pool
- One HDP Pool for SAP HANA log
 - Depending on the LPAR system size, there are one or more parity groups with RAID-6 (6D + 2P) with 8 × 600 GB 10k RPM SAS drives
 - These raid groups are all assigned to one HDP pool

Each LPAR in this solution has the following LUNs:

- One LUN to host the LPAR guest operating system
- One LUN to host SAP HANA's shared volume that matches the LPAR memory size
- Four LUNs to host SAP HANA's log volume
- Four or eight LUNs to host the HANA data volume

Network Architecture

This solution uses two logical networks. Each logical network maps to a separate physical network.

- **Client Network** — This provides the client access to the LPARs and the SAP HANA database running on the LPARs.
- **Management Network** — This provides a network used for the management of hardware and LPARs.

When providing additional features in a solution, you can configure more logical networks to isolate network traffic. Figure 3 shows the network connections for a 1 server blade with a 2 one socket LPAR configuration.




Figure 3

Switches

Two sets of switches are options for this solution.

■ Brocade VDX 6740-48 Top Of the Rack (TOR) Client Switches

These switches make client network connections to the SAP HANA node. The switches are uplinked to the client switch.

These two switches are connected together using ISL. This lets both switches act together as one single logical switch. If one switch fails, there still is a path to the hosts.

■ Brocade ICX 6430-24 Management Switch

This switch connects the management ports of the hardware and the management server. This provides an isolated network for managing the solution.

Client Network

This solution uses 10GBASE-SR 2-port LAN adapters installed on the PCI-E B-slots of each I/O board module. Between two and four LAN adapters are used for connecting to external 10 GbE switches. Each network controller is assigned to an LPAR providing the following:

- Two dedicated 10 GbE network ports for each 1-socket LPAR
- Four dedicated 10 GbE network ports for each 2-socket LPAR
- Four dedicated 10 GbE network ports for each 4-socket LPAR

To provide redundancy two network interfaces are configured at the operating system level to use active-active network bonding mode. This acts as the client network for the SAP HANA node. Connect these network interfaces to the TOR switches. To provide redundancy, one connection is made to each of the TOR client network switch.

Management Network

Logical partitioning manager requires two additional network ports for the LPAR management network path. Provide these two ports by two 1 GbE 4-port LAN mezzanine cards located in slot 1 and slot 3 of server blade 1. The LAN mezzanine cards to the management network through two 1/10Gb LAN switch modules located in the back of the Hitachi Compute Blade 2500 chassis.

Consolidate all 1 GbE management network connections using these 1G/10Gb internal switch modules, especially the LAN connection to the Hitachi Compute Blade 2500 chassis management modules. These connections are then up-linked to the management switch.

LPAR management network links are aggregated using active-backup mode. The IP address of LPAR manager and the chassis management modules (SVP) have to be in the same subnet.

The Hitachi Compute Blade 2500 chassis has two management modules to secure fault tolerance with the following:

- Manage the power supply of each module and monitor the status of the system unit.
- Support the management functionality of the network within the system unit for server blades and various modules.
- Connect between a management module and an external network through a management LAN module

Management Server

This solution uses a QuantaPlexx T41S 2U server for management. The management server acts as a central device for managing the SAP HANA platform. This server is optional.

Manage the following from the management server:

- Hitachi Compute Blade 2500 chassis
- 520X B2 server blades
- Brocade Switches
- LPAR
- SAP HANA node(s)
- Hitachi Virtual Storage Platform
- NTP configuration

Figure 4 shows the management server network ports using one dual port 1 GbE Base-T Intel i350 mezzanine card.

- **Slot 01 Port 2** — Connected to the Brocade ICX 6430. It provides the 1 GbE network to the management server.
- **Slot 01 Port 1** — Connected to the Brocade ICX 6430 port switch. It provides the 1 GbE management network to all components

The management server has the following additional components:

- One dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card
- One Emulex 2-port 8 Gb/sec Fibre Channel HBA on the PCIe slot




Figure 4

SAP HANA Node

This describes the SAP HANA volume and operating system configuration for the LPAR solution on Hitachi Unified Compute Platform for SAP HANA.

SAN Operating System Boot Configuration

A scale-up configuration for SAP HANA on LPAR requires SAN boot.

The operating system LUN is the primary boot device for each LPAR on the 16 Gb/sec 2-port Fibre Channel mezzanine cards. The operating system LUN holds partitions for the following:

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

The following sections that describe configurations and volume assignment hold true for each LPAR.

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 the server blades and Hitachi Virtual Storage Platform G600. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available unless all paths fail.

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

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

HANA Data Volume Configuration

Logical Volume Manager creates a single striped volume for each LPAR on which the XFS file system is created to store the SAP HANA data volume. The striped volume acts as the persistent layer for the SAP HANA server.

HANA Log Volumes Configuration

Logical Volume Manager creates a single striped volume for each LPAR on which the XFS file system is created to store the SAP HANA log volume.

SAP HANA Shared Volume

For the SAP HANA, Logical Volume Manager is used to create a single volume. This volume contains SAP HANA binaries, configuration, and traces.

SAP HANA Appliance Software Installation

After configuring the file system for the SAP HANA data and log volumes, SAP HANA 1.0 is installed on the SAP HANA server(s).

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 Hitachi Data Systems [Global 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 Hitachi Data Systems Corporate [Resources](#) website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems 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 Services [Education](#) website.

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

Hitachi Data Systems



Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 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. © Hitachi Data Systems Corporation 2016. All rights reserved. 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-423-03 July 2016.