WHITE PAPER

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F400 or Virtual Storage Platform G400 with Solid State Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

By Ikki Mizumura and Yingping Niu

August 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.

Revision History

Revision	Changes	Date
AS-609-00	**Initial release**	**August, 2017**

Table of Contents

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F400 or Virtual Storage Platform G400 with Solid State Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

This reference architecture guide uses the SAP HANA tailored data center integration (TDI) approach to implement the SAP HANA platform, rather than using the appliance model. This reference architecture guide provides the storage requirements for the maximum number of validated active SAP HANA production nodes on the following Hitachi Virtual Storage Platform (VSP) storage:

- Virtual Storage Platform F400 using Solid State Drive (SSD)
- Virtual Storage Platform G400 using Solid State Drive (SSD)

With a SAP HANA appliance, all hardware components are pre-configured by the hardware vendor. With SAP HANA tailored data center integration deployment, each one is a customized solution and you can choose hardware from any certified SAP HANA server vendor, along with any certified SAP HANA enterprise storage vendor, to implement a SAP HANA platform. You can leverage existing hardware to reduce the total cost of ownership (TCO).

When deploying a SAP HANA TDI solution, SAP only allows using homogeneous compute server hardware from a single hardware partner in a single implementation.

If a certificate provided by SAP is for a specific operating system, you can only use that operating system for SAP HANA in that implementation.

Engineering validation at Hitachi Data Systems for this approach for SAP HANA has the following points:

- Server blades from Hitachi Data Systems were used during validation.
- Testing showed that the storage design of Virtual Storage Platform F400 or Virtual Storage Platform G400 with SSDs meets the SAP enterprise storage certification key performance indicator (KPI) requirements for SAP HANA.
- During validation, scalability and storage KPI testing was performed using SAP HANA Hardware Configuration Check Tool (HWCCT). Refer to SAP Note 1943937 - Hardware Configuration Check Tool - Central Note (SAP user credentials required) for details about HWCCT.
- A maximum of 16 SAP HANA systems connected to a single VSP F400 or VSP G400 using SSDs were validated to pass the SAP enterprise storage KPIs. These were configured as RAID-10 (2D+2D) while running HWCCT revision 112 on the 16 systems in parallel for testing.

When implementing a SAP HANA TDI infrastructure on VSP F600 or G600 using SSD, you do not have to use exactly the same storage design that Hitachi Data Systems used for enterprise storage KPI testing, as described in this reference architecture guide. But it is recommended to follow the directions and guidelines for the setup and configuration of the storage system. Refer to the SAP HANA Tailored Data Center Integration - Frequently Asked Questions for more details.

Note — Since the release of SAP HANA TDI in November 2013 several versions of HWCCT have been published. To check if the hardware configuration of your SAP HANA TDI infrastructure meets SAP KPIs, you must use the same version of HWCCT used during the certification of the hardware, compute servers and storage system, for your tests. SAP Note 1943937 describes how to determine the right version of HWCCT for your tests.

Note — Testing of this configuration was performed 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 matches your production environment before your production implementation of this solution.

Solution Overview

This document provides an example configuration of the storage layout for SAP HANA nodes with variable sizes of main memory consolidated on Hitachi Virtual Storage Platform F400 or Virtual Storage Platform G400 using solid state drives (SSD), tested in the Hitachi Data Systems lab.

This configuration uses the following:

- One of these for storage:
  - **Hitachi Virtual Storage Platform F400 —** This delivers superior all-flash performance for business-critical applications, with continuous data availability.
  - **Hitachi Virtual Storage Platform G400 —** This model scales for all data types, flexibly adapting for performance, capacity, and multi-vendor storage.
- **Hitachi Storage Virtualization Operating System** — Integrates storage system software to provide system element management and advanced storage system functions.

Figure 1 shows the server to storage configuration of this solution using Virtual Storage Platform G400 with sixteen SAP HANA systems. The configuration for Virtual Storage Platform F400 is similar.

Figure 1

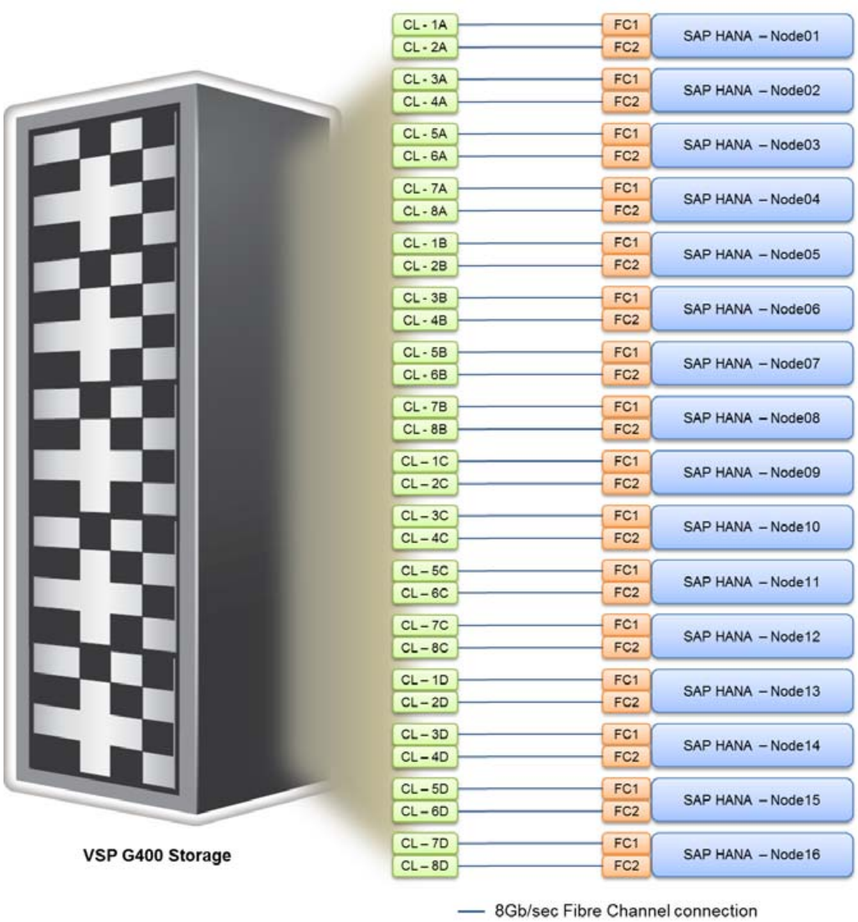

Key Solution Elements

These are the key hardware and software elements used for scalability testing.

Hardware Elements

Table 1 describes the hardware required to test the scalability of 16 active nodes on Hitachi Virtual Storage Platform F400 or Virtual Storage Platform G400 using solid state drives (SSDs).

TABLE 1. HARDWARE ELEMENTS

Hardware	Quantity	Configuration	Role
One of the following Hitachi Virtual Storage Platform models: ▪ VSP F400 using solid state drives ▪ VSP G400 using solid state drives	1	▪ CTL: 1 pair ▪ DKB: 1 pair ▪ CHB pairs: 4 × 8 Gb/sec ▪ MPU: 2 pairs ▪ Cache: 128 GB	Block storage for SAP HANA platform nodes
Certified server for SAP HANA	16	▪ Rack servers or server blade chassis ▪ 2-socket server	SAP HANA Nodes
Brocade ICX 6430-48 port switch (optional)	1	▪ 1 GbE ▪ 48 ports	1 GbE management network
Brocade VDX 6740-48 port switch (optional)	1	▪ 10 GbE ▪ 48 ports	10 GbE connectivity

See Certified and Supported SAP HANA Hardware.

Hitachi Virtual Storage Platform Gx00

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 Virtual Storage Platform Fx00 Models

Hitachi Virtual Storage Platform Fx00 models deliver superior all-flash performance for business-critical applications, with continuous data availability. High-performance network attached storage with nondisruptive deduplication reduces the required storage capacity by up to 90% with the power to handle large, mixed-workload environments.

Server for SAP HANA

Only servers certified for use with SAP HANA can be used in the SAP HANA TDI environment. This follows the exact same bill of materials as the certified SAP HANA appliance server, but without the storage or local disks.

Find a description of all certified servers and enterprise storage solutions in Certified and Supported SAP HANA Hardware. For more information on SAP HANA TDI, consult SAP HANA Tailored Data Center Integration - Frequently Asked Questions.

Software Elements

This is the software used to deploy the test configuration.

- SUSE Linux Enterprise Server for SAP Applications
- SAP HANA
- Hitachi Storage Navigator Modular 2
- Hitachi Command Suite
- Hitachi Storage Virtualization Operating System
  - Includes Hitachi Dynamic Provisioning
- Microcode for Hitachi Virtual Storage Platform F400 or Virtual Storage Platform G400

Note - Scalability testing was carried out using SUSE Linux Enterprise Server for SAP Applications. However, the solution also supports using Red Hat Enterprise Linux.

Hitachi Storage Virtualization Operating System

Hitachi Storage Virtualization Operating System (SVOS) spans and integrates multiple platforms. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Storage Virtualization Operating System includes standards-based management software on a Hitachi Command Suite (HCS) base. This provides storage configuration and control capabilities for you.

Storage Virtualization Operating System uses Hitachi Dynamic Provisioning (HDP) to provide wide striping and thin provisioning. Dynamic Provisioning provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) without initially allocating any physical space. Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs).

Solution Design

This is the detailed solution example for the SAP HANA tailored data center integration (TDI) on either of the following:

- Hitachi Virtual Storage Platform F400 using SSD
- Hitachi Virtual Storage Platform G400 using SSD

Fibre Channel Architecture

For each SAP HANA node, there are two dedicated Fibre Channel ports on the storage. The two Fibre Channel cables directly connect HBA ports on the node to the designated storage target ports to achieve the following:

- No single point of failure for high availability
- Expected throughput of data and log volume of SAP HANA

Table 2 shows the storage port mapping for 4 nodes as an example.

TABLE 2. STORAGE PORT MAPPING FOR 4 NODES

SAP HANA Node	Fibre Channel Port	VSP F400/VSP G400 Ports
Node001	Port 0	1A
Node001	Port 1	2A
Node002	Port 0	3A
Node002	Port 1	4A
Node003	Port 0	5A
Node003	Port 1	6A
Node004	Port 0	7A
Node004	Port 1	8A

Storage Architecture

Each SAP HANA node needs the following storage layout:

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

This SAP HANA TDI setup utilizes dynamic provisioning pools created with Hitachi Dynamic Provisioning for the storage layout. This ensures maximum utilization and optimization at a lower cost than other solutions.

The dynamic provisioning pool layout options with minimal disks and storage cache for 1 to 16 nodes are listed in Table 3.

TABLE 3. DYNAMIC PROVISIONING POOLS WITH DISKS AND STORAGE CACHE

Dynamic Provisioning Pool Name	Purpose	Parity Group RAID Level and Disks	Number of Nodes	Storage Cache
OS_SH_Data_Pool	▪ OS Volume ▪ SAP HANA Shared ▪ SAP HANA Data Volume	1 × RAID-10 (2D+2D) on 1.9 TB SSD	1 to 12	128 GB
		2 × RAID-10 (2D+2D) on 1.9 TB SSD	13 to 16	
Log_Pool	▪ SAP HANA Log Volume	1 × RAID-10 (2D+2D) on 1.9 TB SSD	1 to 16	

A minimum of two parity groups is needed for OS_SH_Data_Pool, and a minimum of one parity group is needed for Log_Pool to fit in 16 SAP HANA production nodes on a Hitachi Virtual Storage Platform F400 or Virtual Storage Platform G400 using RAID-10 (2D+2D). Additional parity groups of the same type may need to be added, depending on the various combination of node sizes as well as the number of nodes to meet the capacity requirements.

This is an example layout of the dynamic provisioning pool configuration on Virtual Storage Platform F400 or Virtual Storage Platform G400 using SSDs. It is used in this SAP HANA tailored data center integration solution with four active SAP HANA systems, but not limited to, the following:

- **System 1** — 512 GB
- **System 2** — 1 TB
- **System 3** — 2 TB
- **System 4** — 4 TB

Provision the storage for the 16 SAP HANA systems as follows:

- Create the parity groups first, as shown in Table 4 on page 7, using the RAID-10 design.
- Create two dynamic provisioning pools for the 16 SAP HANA systems on VSP F400 or VSP G400:
  - Use **OS_SH_Data_Pool** to provision the operating system volume, SAP HANA shared volume, and Data volume.
  - Use **Log_Pool** to provision the Log volume.
- Assign all LDEVs to the dedicated pool.
- Determine the suggested minimum sizes for the data, log, and HANA shared using these formulas, as provided by SAP:
  - Data = 1 × memory (RAM)
  - Log = 0.5 × memory, for systems less than or equal to 512 GB
  - Log = 512 GB, for systems greater than 512 GB
  - HANA Shared = minimum (1 × memory; 1 TB), for a single node setup (scale-up)
  - HANA Shared = 1 × memory of workers for 4 worker nodes (scale-out)

Table 5 on page 8 shows the examples of minimum sizes for data, log and HANA shared for four different memory sizes of HANA systems.

TABLE 4. DYNAMIC PROVISIONING POOL PROVISIONING WITH RAID-10 (2D+2D) FOR 16 NODES

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
OS_SH_Data_Pool	1	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:01	OS_SH_DATA_DPVOL_1	878 GB	MPU-10
			00:00:02	OS_SH_DATA_DPVOL_2	878 GB	MPU-11
			00:00:03	OS_SH_DATA_DPVOL_3	878 GB	MPU-20
			00:00:04	OS_SH_DATA_DPVOL_4	878 GB	MPU-21
	2	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:05	OS_SH_DATA_DPVOL_5	878 GB	MPU-10
			00:00:06	OS_SH_DATA_DPVOL_6	878 GB	MPU-11
			00:00:07	OS_SH_DATA_DPVOL_7	878 GB	MPU-20
			00:00:08	OS_SH_DATA_DPVOL_8	878 GB	MPU-21
Log_Pool	3	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:09	LOG_DPVOL_1	878 GB	MPU-10
			00:00:10	LOG_DPVOL_2	878 GB	MPU-11
			00:00:11	LOG_DPVOL_3	878 GB	MPU-20
			00:00:12	LOG_DPVOL_4	878 GB	MPU-21

- Provision virtual volumes for each of the nodes, as follows:
  - Create virtual volumes (VVOLs) for the operating system, SAP HANA shared, Log and Data volumes for four different memory size HANA systems. See Table 5 as an example.

TABLE 5. VVOLS FOR THE SAP HANA NODES FOR 4 MEMORY SIZE HANA SYSTEMS

Dynamic Provisioning Pool	VVOL ID	VVOL Name	VVOL Size	MPU Assignment
OS_SH_Data_Pool	00:01:00	HANA_OS_N1	100 GB	MPU-10
	00:02:00	HANA_OS_N2	100 GB	MPU-11
	00:03:00	HANA_OS_N3	100 GB	MPU-20
	00:04:00	HANA_OS_N4	100 GB	MPU-21
	00:01:01	HANA_SH_N1	512 GB	MPU-10
	00:02:01	HANA_SH_N2	1024 GB	MPU-11
	00:03:01	HANA_SH_N3	1024 GB	MPU-20
	00:04:01	HANA_SH_N4	1024 GB	MPU-21
	00:01:06	HANA_DATA_N1_1	128 GB	MPU-10
	00:01:07	HANA_DATA_N1_2	128 GB	MPU-11
	00:01:08	HANA_DATA_N1_3	128 GB	MPU-20
	00:01:09	HANA_DATA_N1_4	128 GB	MPU-21
	00:02:06	HANA_DATA_N2_1	256 GB	MPU-10
	00:02:07	HANA_DATA_N2_2	256 GB	MPU-11
	00:02:08	HANA_DATA_N2_3	256 GB	MPU-20
	00:02:09	HANA_DATA_N2_4	256 GB	MPU-21
	00:03:06	HANA_DATA_N3_1	512 GB	MPU-10
	00:03:07	HANA_DATA_N3_2	512 GB	MPU-11
	00:03:08	HANA_DATA_N3_3	512 GB	MPU-20
	00:03:09	HANA_DATA_N3_4	512 GB	MPU-21
	00:04:06	HANA_DATA_N4_1	1024 GB	MPU-10
	00:04:07	HANA_DATA_N4_2	1024 GB	MPU-11
	00:04:08	HANA_DATA_N4_3	1024 GB	MPU-20
	00:04:09	HANA_DATA_N4_4	1024 GB	MPU-21

TABLE 5. VVOLS FOR THE SAP HANA NODES FOR 4 MEMORY SIZE HANA SYSTEMS (CONTINUED)

Dynamic Provisioning Pool	VVOL ID	VVOL Name	VVOL Size	MPU Assignment
Log_Pool	00:01:02	HANA_LOG_N1_1	64 GB	MPU-10
	00:01:03	HANA_LOG_N1_2	64 GB	MPU-11
	00:01:04	HANA_LOG_N1_3	64 GB	MPU-20
	00:01:05	HANA_LOG_N1_4	64 GB	MPU-21
	00:02:02	HANA_LOG_N2_1	128 GB	MPU-10
	00:02:03	HANA_LOG_N2_2	128 GB	MPU-11
	00:02:04	HANA_LOG_N2_3	128 GB	MPU-20
	00:02:05	HANA_LOG_N2_4	128 GB	MPU-21
	00:03:02	HANA_LOG_N3_1	128 GB	MPU-10
	00:03:03	HANA_LOG_N3_2	128 GB	MPU-11
	00:03:04	HANA_LOG_N3_3	128 GB	MPU-20
	00:03:05	HANA_LOG_N3_4	128 GB	MPU-21
	00:04:02	HANA_LOG_N4_1	128 GB	MPU-10
	00:04:03	HANA_LOG_N4_2	128 GB	MPU-11
	00:04:04	HANA_LOG_N4_3	128 GB	MPU-20
	00:04:05	HANA_LOG_N4_4	128 GB	MPU-21

- While mapping the LUN path assignment for each node, add VVOLs in the following order:

1. **The operating volume for the specific SAP HANA node**
2. **The SAP HANA shared for the specific SAP HANA node**
3. **The log volumes and data volumes for the specific SAP HANA node**

Table 6 shows an example configuration of the LUN path assignment for Node001. Configure the LUN assignment similarly for all other nodes.

TABLE 6. EXAMPLE LUN PATH ASSIGNMENT FOR THE SAP HANA CONFIGURATION ON NODE001

LUN ID	LDEV ID	LDEV Name
0000	00:01:00	HANA_OS_N1
0001	00:01:01	HANA_SH_N1
0002	00:01:02	HANA_LOG_N1_1
0003	00:01:03	HANA_LOG_N1_2
0004	00:01:04	HANA_LOG_N1_3
0005	00:01:05	HANA_LOG_N1_4
0006	00:01:06	HANA_DATA_N1_1
0007	00:01:07	HANA_DATA_N1_2

TABLE 6. EXAMPLE LUN PATH ASSIGNMENT FOR THE SAP HANA CONFIGURATION ON NODE001 (CONTINUED)

LUN ID	LDEV ID	LDEV Name
0008	00:01:08	HANA_DATA_N1_3
0009	00:01:09	HANA_DATA_N1_4

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 VSP G400 or VSP F400.

Each node has two I/O paths connected with the storage. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available unless both paths fail.

SAP Storage Connector API Fibre Channel Client

The SAP HANA Storage Connector API Fibre Channel Client (fcClient) defines a set up interface functions call during normal SAP HANA cluster operation and failover handling. The scale-out configuration for SAP HANA uses the fcClientLVM implementation, which supports the use of Logical Volume Manager. SAP supports this solution to enable the use of high-performance Fibre Channel devices in a scale-out installation.

The fcClientLVM implementation uses standard Linux commands, such as multipath and sg_persist. Install and configure these commands.

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

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

Best Practice for Storage Setup

These are best practices for setting up storage in a SAP HANA TDI environment.

- Create a dynamic provisioning pool with a minimum of two parity groups whenever possible.
- Dedicate a parity group to one pool only. Do not use it for other purposes if one of its LDEVs is a Pool Volume.
- Distribute the parity groups across at least two drive trays if possible
- Create four DP-VOLs for log volumes for each SAP HANA system. Distribute the DP-VOLs across the various MPUs.
- Create four DP-VOLs for data volumes for each SAP HANA system. Distribute the DP-VOLs across the various MPUs.

Engineering Validation

The test methodology for validating this SAP HANA tailored datacenter integration (TDI) enterprise storage configuration used the following:

- SAP HANA Hardware Configuration Check Took (HWCCT) for testing Hitachi Virtual Storage Platform F400 and Virtual Storage Platform G400 for the enterprise storage certification, revision hwcct-112
- For the optimal use of the system with a SAP HANA database, these suggested parameters:
  - "async_read_submit=on"
  - "async_write_submit_blocks=all"

Follow SAP Note 2267798 to set up these parameters for SAP HANA 1.0

Follow SAP Note 2399079 to set up these parameters for SAP HANA 2.0

- The following is the *global.ini* file configured and used for validation of this solution:

```
[communication]
listeninterface = .global

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

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*__prtype = 5
partition_1_data__lvmname = vgHIQdata001-lvdata
partition_1_log__lvmname = vgHIQlog001-lvlog
partition_2_data__lvmname = vgHIQdata002-lvdata
partition_2_log__lvmname = vgHIQlog002-lvlog
partition_3_data__lvmname = vgHIQdata003-lvdata
partition_3_log__lvmname = vgHIQlog003-lvlog
partition_4_data__lvmname = vgHIQdata004-lvdata
partition_4_log__lvmname = vgHIQlog004-lvlog
partition_5_data__lvmname = vgHIQdata005-lvdata
partition_5_log__lvmname = vgHIQlog005-lvlog
partition_6_data__lvmname = vgHIQdata006-lvdata
partition_6_log__lvmname = vgHIQlog006-lvlog
```

```
partition_7_data__lvmname = vgHIQdata007-lvdata

partition_7_log__lvmname = vgHIQlog007-lvlog

partition_8_data__lvmname = vgHIQdata008-lvdata

partition_8_log__lvmname = vgHIQlog008-lvlog

partition_9_data__lvmname = vgHIQdata009-lvdata

partition_9_log__lvmname = vgHIQlog009-lvlog

partition_10_data__lvmname = vgHIQdata010-lvdata

partition_10_log__lvmname = vgHIQlog010-lvlog

partition_11_data__lvmname = vgHIQdata011-lvdata

partition_11_log__lvmname = vgHIQlog011-lvlog

partition_12_data__lvmname = vgHIQdata012-lvdata

partition_12_log__lvmname = vgHIQlog012-lvlog

partition_13_data__lvmname = vgHIQdata013-lvdata

partition_13_log__lvmname = vgHIQlog013-lvlog

partition_14_data__lvmname = vgHIQdata014-lvdata

partition_14_log__lvmname = vgHIQlog014-lvlog

partition_15_data__lvmname = vgHIQdata015-lvdata

partition_15_log__lvmname = vgHIQlog015-lvlog

partition_16_data__lvmname = vgHIQdata016-lvdata

partition_16_log__lvmname = vgHIQlog016-lvlog


[trace]

ha_fcclientlvm = info
```

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. 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 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.

Hitachi Data Systems

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

AS-609-00, August 2017.