HITACHI
Inspire the Next

WHITE PAPER

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform G200 with Solid State Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

By Yingping Niu

October 2017

Feedback

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

Revision History

Revision	Changes	Date
SL-002-00	Initial release	October 4, 2017

Table of Contents

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform G200 with Solid State Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

Use this reference architecture guide to implement a SAP HANA tailored data center integration (TDI) solution for the SAP HANA platform, rather than using the appliance model. This provides the storage requirements for the maximum number of validated active SAP HANA production nodes on Hitachi Virtual Storage Platform (VSP) G200 using solid state drive storage:

With a SAP HANA appliance deployment, each installation has the hardware vendor pre-configure all hardware components. With a SAP HANA tailored data center integration deployment, each installation is customized by choosing hardware from any certified SAP HANA server vendor and storage from any certified SAP HANA enterprise storage vendor. You may be able to leverage your 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 for this Hitachi system for SAP HANA TDI has the following points:

- Server blades from Hitachi were used.
- Testing showed that the storage design of Hitachi Virtual Storage Platform G200 with SSDs meets the SAP enterprise storage certification key performance indicator (KPI) requirements for SAP HANA.
- 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.
- There was up to eight SAP HANA systems connected in parallel to a single Virtual Storage Platform G200 using SSDs to pass the SAP enterprise storage KPIs. These were configured as RAID-10 (2D+2D). There are no other systems or applications using the storage while running HWCCT revision 112 on the eight SAP HANA systems during testing.

Note — If you need to have mixed applications running in parallel on the same storage array, SAP HANA and non-SAP HANA, you must perform testing for the expected workload using the SAP HANA Hardware Configuration Tool (HWCCT) to meet SAP KPIs.

When implementing a SAP HANA TDI infrastructure on Hitachi Virtual Storage Platform G200 using SSD, you do not have to use exactly same storage design described in this reference architecture guide that was used for enterprise storage KPI testing. However, the recommendation is to follow the directions and guidelines for the setup and configuration of the storage system. Refer to 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 the 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 reference architecture guide provides an example configuration of the storage layout for SAP HANA nodes with variable sizes of main memory consolidated on Hitachi Virtual Storage Platform G200, using solid state drives (SSD), tested in the Hitachi lab.

This configuration uses the following:

- **Hitachi Virtual Storage Platform G200 —** This highly reliable storage system offers high storage capacity with full redundancy to better protect data and manage storage operations.
- **Hitachi Storage Virtualization Operating System** — This integrates storage system software to provide system element management and advanced storage system functions.

Figure 1 on page 3 shows the server to storage configuration of this solution using Virtual Storage Platform G200 with eight SAP HANA systems.

Figure 1

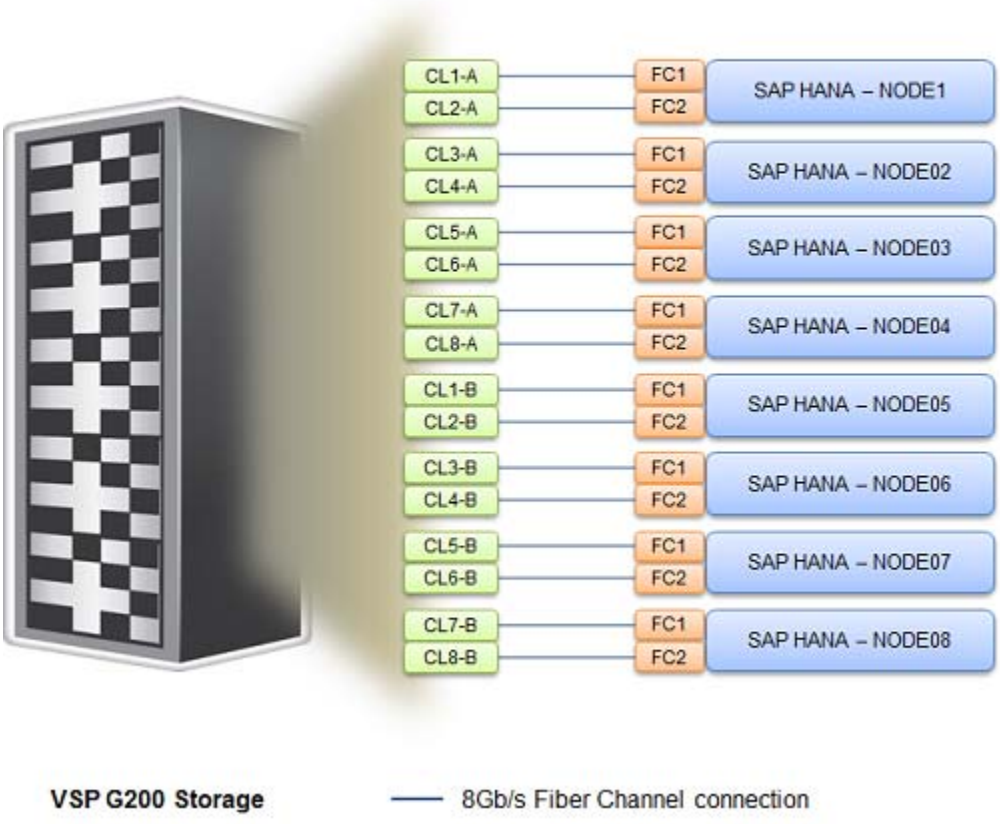
CL1-A
CL2-A
FC1
FC2
SAP HANA – NODE1
CL3-A
CL4-A
FC1
FC2
SAP HANA – NODE02
CL5-A
CL6-A
FC1
FC2
SAP HANA – NODE03
CL7-A
CL8-A
FC1
FC2
SAP HANA – NODE04
CL1-B
CL2-B
FC1
FC2
SAP HANA – NODE05
CL3-B
CL4-B
FC1
FC2
SAP HANA – NODE06
CL5-B
CL6-B
FC1
FC2
SAP HANA – NODE07
CL7-B
CL8-B
FC1
FC2
SAP HANA – NODE08
VSP G200 Storage
8Gb/s Fiber Channel connection

Key Solution Elements

These are the key hardware and software elements used for this environment.

Hardware Elements

Table 1, “Hardware Elements,” describes the hardware required to test the scalability of eight active nodes on Hitachi Virtual Storage Platform G200 using solid state drives (SSDs).

TABLE 1. HARDWARE ELEMENTS

Hardware	Quantity	Configuration	Role
Hitachi Virtual Storage Platform G200	1	▪ CTL: 1 pair ▪ CHB pairs: 2 × 8 Gb/s ▪ MPU: 2 pairs ▪ Cache: 64 GB ▪ 4 × 1.9 TB SSDs	Block storage for SAP HANA nodes
Certified server for SAP HANA (Note 1)	8	▪ 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

Note 1. 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.

This validation used Virtual Storage Platform G200.

Server for SAP HANA

Only servers certified for use with SAP HANA can be used in a 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 for this environment.

- 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 G200

Note — Scalability testing was carried out using SUSE Linux Enterprise Server for SAP Applications. However, this solution also supports 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 Hitachi Virtual Storage Platform G200 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 connecting 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 an example storage port mapping for eight nodes.

TABLE 2. STORAGE PORT MAPPING FOR 8 SAP HANA NODES

SAP HANA Node	Fibre Channel Port	Virtual Storage Platform G200 Port
Node001	Port 0	1A
	Port 1	2A
Node002	Port 0	3A
	Port 1	4A
Node003	Port 0	5A
	Port 1	6A
Node004	Port 0	7A
	Port 1	8A
Node005	Port 0	1B
	Port 1	2B
Node006	Port 0	3B
	Port 1	4B
Node007	Port 0	5B
	Port 1	6B
Node008	Port 0	7B
	Port 1	8B

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 8 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_DT_LG_Pool	■ OS Volume ■ SAP HANA Shared ■ SAP HANA Data Volume ■ SAP HANA Log Volume	1 × RAID-10 (2D+2D) on 1.9 TB SSD	1 to 8	64 GB

A minimum of one parity group is needed for OS_SH_DT_LG_Pool to fit eight SAP HANA production nodes on one Hitachi Virtual Storage Platform G200 using RAID-10 (2D+2D). Additional parity groups of the same type may need to be added. Drive boxes may be needed if the internal drives on Virtual Storage Platform G200 are not sufficient, depending on the following:

- The various combinations of node sizes
- The number of nodes to meet the capacity requirements

This is an example layout of the dynamic provisioning pool configuration on Virtual Storage Platform G200 using SSDs. While it is not limited to these systems, this SAP HANA tailored data center integration solution uses the following four active SAP HANA systems:

- **System 1** — 128 GB
- **System 2** — 512 TB
- **System 3** — 1 TB
- **System 4** — 2 TB

Provision the storage for the eight SAP HANA systems, as follows:

- Create the parity groups first, as shown in Table 4. Dynamic Provisioning Pool Provisioning with RAID-10 (2D+2D) for 8 of 128 GB Memory Nodes, using the RAID-10 design.
- Create one dynamic provisioning pool for the eigth SAP HANA systems on Virtual Storage Platform G200. Use **OS_SH_DT_LG_Pool** to provision the following volumes:
  - Operating system volume
  - SAP HANA shared volume
  - Data volume and log volume
- Assign all LDEVs to the dedicated pool.

- Determine the suggested minimum sizes for the data, log, and HANA shared volumes using these formulas provided by SAP:
  - Data = 1 × memory (RAM)
  - For systems less than or equal to 512 GB, Log = 0.5 × memory
  - For systems greater than 512 GB, Log = 512 GB
  - For a single node setup (scale-up), HANA Shared = minimum (1 × memory; 1 TB)
  - For 4 worker nodes (scale-out), HANA Shared = 1 × memory of workers

  Table 4 shows examples of the minimum sizes for the data, log and HANA shared volumes for four different memory sizes of HANA systems.

TABLE 4. DYNAMIC PROVISIONING POOL PROVISIONING WITH RAID-10 (2D+2D) FOR 8 OF 128 GB MEMORY NODES

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
OS_SH_DT_LG_Pool	1	RAID-10 (2D+2D) on 1.9 TB SSD	00:00:01	OS_SH_DT_LG_DPVOL_1	878 GB	MPU-10
			00:00:02	OS_SH_DT_LG_DPVOL_2	878 GB	MPU-11
			00:00:03	OS_SH_DT_LG_DPVOL_3	878 GB	MPU-20
			00:00:04	OS_SH_DT_LG_ DPVOL_4	878 GB	MPU-21

- For four different memory size HANA systems, on each of node, provision virtual volumes (VVOLs) for the following:
  - Operating system
  - SAP HANA shared
  - Log and Data volumes

  See Table 5 as an example.

TABLE 5. VVOLS FOR THE SAP HANA NODES FOR 4 MEMORY SIZE HANA SYSTEMS ON OS_SH_DT_LG_POOL DYNAMIC PROVISIONING POOL

VVOL ID	VVOL Name	VVOL Size	MPU Assignment
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	128 GB	MPU-10

TABLE 5. VVOLS FOR THE SAP HANA NODES FOR 4 MEMORY SIZE HANA SYSTEMS ON OS_SH_DT_LG_POOL DYNAMIC PROVISIONING POOL (CONTINUED)

VVOL ID	VVOL Name	VVOL Size	MPU Assignment
00:02:01	HANA_SH_N2	512 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	32 GB	MPU-10
00:01:07	HANA_DATA_N1_2	32 GB	MPU-11
00:01:08	HANA_DATA_N1_3	32 GB	MPU-20
00:01:09	HANA_DATA_N1_4	32 GB	MPU-21
00:02:06	HANA_DATA_N2_1	128 GB	MPU-10
00:02:07	HANA_DATA_N2_2	128 GB	MPU-11
00:02:08	HANA_DATA_N2_3	128 GB	MPU-20
00:02:09	HANA_DATA_N2_4	128 GB	MPU-21
00:03:06	HANA_DATA_N3_1	256 GB	MPU-10
00:03:07	HANA_DATA_N3_2	256 GB	MPU-11
00:03:08	HANA_DATA_N3_3	256 GB	MPU-20
00:03:09	HANA_DATA_N3_4	256 GB	MPU-21
00:04:06	HANA_DATA_N4_1	512 GB	MPU-10
00:04:07	HANA_DATA_N4_2	512 GB	MPU-11
00:04:08	HANA_DATA_N4_3	512 GB	MPU-20
00:04:09	HANA_DATA_N4_4	512 GB	MPU-21
00:01:02	HANA_LOG_N1_1	16 GB	MPU-10
00:01:03	HANA_LOG_N1_2	16 GB	MPU-11
00:01:04	HANA_LOG_N1_3	16 GB	MPU-20
00:01:05	HANA_LOG_N1_4	16 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

TABLE 5. VVOLS FOR THE SAP HANA NODES FOR 4 MEMORY SIZE HANA SYSTEMS ON OS_SH_DT_LG_POOL DYNAMIC PROVISIONING POOL (CONTINUED)

VVOL ID	VVOL Name	VVOL Size	MPU Assignment
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
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 Hitachi Virtual Storage Platform G200.

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 the 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.

- Normally, create a dynamic provisioning pool with a minimum of two parity groups whenever possible. However, with Hitachi Virtual Storage Platform G200, use one parity group for the SAP HANA TDI enterprise storage configuration to achieve a lower cost optimimzation when compared with other configurations.
- Dedicate a parity group to one pool only. Do not use the parity group 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) tested Hitachi Virtual Storage Virtual Storage Platform G200 for the enterprise storage certification, revision hwcct-112.
- The fsperf_mult_105.sh script tested these workloads:
  - Data volume
  - Log volume
  - MixedIO (A MixedIO test is mandary if the data volume and log volume are created in same parity group. Otherwise this test is optional.)
- For the optimal use of the system with a SAP HANA database, these suggested parameters were used:
  - "async_read_submit=on"
  - "async_write_submit_active=on"
  - "max_parallel_io_requests=256"
- Follow SAP Note 2267798 to setup these parameters using hdbparam tool for SAP HANA 1.0.
- Follow SAP Note 2399079 to setup these parameters defined in global.ini 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



[trace]
ha_fcclientlvm = info
```

For More Information

Hitachi Vantara 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.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the Resources website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantana Training and Certification website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the Hitachi Vantara website.

Hitachi Vantara

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

Regional Contact Information
Americas: +1 408 970 1000 or info@hitachivantara.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hitachivantara.com

SL-002-00, October 2017.