

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F600 or Virtual Storage Platform G600 with Flash Module Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

By 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-611-00	Initial release	August 1, 2017

Table of Contents

- Solution Overview 2**
- Key Solution Elements 3**
 - Hardware Elements..... 3
 - Software Elements..... 4
- Solution Design 4**
 - Fibre Channel Architecture 5
 - Storage Architecture..... 5
 - Best Practice for Storage Setup 10
- Engineering Validation 10**

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F600 or Virtual Storage Platform G600 with Flash Module Drives using Hitachi Dynamic Provisioning

Reference Architecture Guide

Use this reference architecture guide with the [SAP HANA tailored data center integration](#) (TDI) approach to implement the SAP HANA platform, rather than using the appliance approach. 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 F600
- Virtual Storage Platform G600 using flash module drives (FMD)

With a SAP HANA appliance, all hardware components are pre-configured by the hardware vendor. With a 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 Solutions 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 F600, or Virtual Storage Platform G600 with FMDs, 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 28 SAP HANA systems passed the SAP enterprise storage KPIs on a single Virtual Storage Platform F600, or Virtual Storage Platform G600 with FMDs. These were configured as RAID-10 (2D+2D) while using HWCCT revision 112 for testing.

When implementing a SAP HANA TDI infrastructure, you do not need to use the same storage design that Hitachi Data Systems used for enterprise storage KPI testing, as described in this reference architecture guide. 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 F600, or Virtual Storage Platform G600 with flash module drives (FMD), tested in the Hitachi Data Systems lab.

This configuration uses the following:

- One of these for storage:
 - **Hitachi Virtual Storage Platform F600** — This delivers superior all-flash performance for business-critical applications, with continuous data availability.
 - **Hitachi Virtual Storage Platform G600** — 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 F600 with twenty-eight SAP HANA systems. The configuration for Virtual Storage Platform G600 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 28 active nodes on Hitachi Virtual Storage Platform F600 or Virtual Storage Platform G600 with flash module drives (FMDs).

TABLE 1. HARDWARE ELEMENTS

Hardware	Quantity	Configuration	Role
One of the following Hitachi Virtual Storage Platform models: <ul style="list-style-type: none">VSP F600VSP G600 using flash module drives	1	<ul style="list-style-type: none">CTL: 1 pairDKB: 1 pairCHB pairs: 7 × 8 Gb/sMPU: 2 pairsCache: 256 GB	Block storage for SAP HANA nodes
Certified server for SAP HANA (Note 1)	28	<ul style="list-style-type: none">Rack servers or server blade chassisServer memory: 256 GB	SAP HANA Nodes
Brocade ICX 6430-48 port switch (optional)	1	<ul style="list-style-type: none">1 GbE48 ports	1 GbE management network
Brocade VDX 6740-48 port switch (optional)	1	<ul style="list-style-type: none">10 GbE48 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.

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 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 F600 or Virtual Storage Platform G600

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 F600
- Hitachi Virtual Storage Platform G600 using FMD

Fibre Channel Architecture

Directly attach the SAP HANA nodes to the designated Hitachi Virtual Storage Platform target port.

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	Hitachi Virtual Storage Platform 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 pool design 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 in the Table 3 are for 28 nodes.

TABLE 3. DYNAMIC PROVISIONING POOLS

Dynamic Provisioning Pool Name	Purpose	Parity Group RAID Level and Disks	Number of Nodes
OS_SH_Data_Pool	<ul style="list-style-type: none"> ■ OS Volume ■ SAP HANA Shared ■ SAP HANA Data Volume 	2 × RAID-10 (2D+2D) on 1.6 TB FMD	28
Log_Pool	<ul style="list-style-type: none"> ■ SAP HANA Log Volume 	2 × RAID-10 (2D+2D) on 1.6 TB FMD	

A minimum of two parity groups is needed for each dynamic provisioning pool to fit in 28 SAP HANA production nodes on Hitachi Virtual Storage Platform F600 or Virtual Storage Platform G600 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 F600, or Virtual Storage Platform G600 with FMDs. 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 28 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 28 Nodes,” on page 7, using the RAID-10 design.
- Create two dynamic provisioning pools for the four SAP HANA systems on Virtual Storage Platform:
 - 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 4. DYNAMIC PROVISIONING POOL PROVISIONING WITH RAID-10 (2D+2D) FOR 28 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 1.6 TB FMD	00:00:01	OS_SH_DA_Pool_1	800 GB	MPU-10
			00:00:02	OS_SH_DA_Pool_2	800 GB	MPU-11
			00:00:03	OS_SH_DA_Pool_3	800 GB	MPU-20
			00:00:04	OS_SH_DA_Pool_4	800 GB	MPU-21
	2	RAID-10 (2D+2D) on 1.6 TB FMD	00:00:05	OS_SH_DA_Pool_5	800 GB	MPU-10
			00:00:06	OS_SH_DA_Pool_6	800 GB	MPU-11
			00:00:07	OS_SH_DA_Pool_7	800 GB	MPU-20
			00:00:08	OS_SH_DA_Pool_8	800 GB	MPU-21
Log_Pool	3	RAID-10 (2D+2D) on 1.6 TB FMD	00:00:09	LG_Pool_1	800 GB	MPU-10
			00:00:10	LG_Pool_2	800 GB	MPU-11
			00:00:11	LG_Pool_3	800 GB	MPU-20
			00:00:12	LG_Pool_4	800 GB	MPU-21
	4	RAID-10 (2D+2D) on 1.6 TB FMD	00:00:13	LG_Pool_5	800 GB	MPU-10
			00:00:14	LG_Pool_6	800 GB	MPU-11
			00:00:15	LG_Pool_7	800 GB	MPU-20
			00:00:16	LG_Pool_8	800 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 as shown in Table 5. VVOLs for the SAP HANA Nodes for Four HANA Systems for four different memory size HANA systems, as an example.

TABLE 5. VVOLS FOR THE SAP HANA NODES FOR FOUR 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 FOUR 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 the 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

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.

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 Tool (HWCCT) for testing Hitachi Virtual Storage Platform F600 and Virtual Storage Platform G600 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"
 - "max_parallel_io_requests=256"

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.

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 Data Systems Corporation 2017. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. HDS and VSP are registered trademarks or trademarks of Hitachi Data Systems Corporation. 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-611-00, August 2017.