TECHNICAL PAPER

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F400 with Flash Module Drives Using Hitachi Dynamic Provisioning Pools

Reference Architecture Guide

By Abhishek Dhanuka and Nobuo Yagi

June 2016

Hitachi Data Systems

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

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F400 with Flash Module Drives Using Hitachi Dynamic Provisioning Pools

Reference Architecture Guide

The purpose of this reference architecture guide is to help customers who prefer the SAP HANA Tailored Data Center Integration (TDI) approach to implement SAP HANA, rather than the SAP HANA appliance model. This reference architecture provides the storage requirements for the maximum number of supported active SAP HANA productive nodes on a Hitachi Virtual Storage Platform (VSP) F400 storage array with flash module drives (FMD) using Hitachi Dynamic Provisioning (HDP) pools.

On Hitachi storage systems, HDP provides wide striping and thin provisioning functionalities. Using Dynamic Provisioning is similar to 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 initially allocating any physical space.

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.

Unlike a SAP HANA appliance in which all hardware components are pre-configured by the hardware vendor, SAP HANA tailored data center integration deployments are customized solutions where the customer can choose any of the certified SAP HANA server vendors along with any certified SAP HANA enterprise storage to implement SAP HANA. This provides customers an opportunity to leverage their existing hardware and reduce TCO.

Using this reference architecture, SAP HANA solutions can be deployed using any certified SAP HANA server vendor and Hitachi Virtual Storage Platform F400. A list of SAP certified servers that are available for SAP HANA appliances can be found in the SAP Certified and Supported SAP HANA® Hardware Directory. SAP only allows using homogeneous compute server hardware from a single hardware partner in a SAP HANA tailored data center integration. Also, if a certificate provided by SAP is for a specific operating system, only that operating system can be used for SAP HANA implementation. Engineering validation for this solution has been performed using Hitachi Data Systems compute blade (CB) server blades.

Every SAP certified enterprise storage platform must meet TDI storage key performance indicators (KPIs) set by SAP. Testing showed that the storage design of VSP F400 with FMD for the SAP HANA Platform meets the TDI storage KPI requirements from SAP. It is not mandatory for customers to use the same storage design that was used for storage KPI testing, and this is demonstrated in this reference architecture guide. Refer to the SAP HANA Tailored Data Center Integration FAQ for more details about TDI.

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.

During validation, the scalability and storage KPI testing was performed using SAP Hardware Configuration Check Tool (HWCCT). Please refer to SAP Note 1943937 - Hardware Configuration Check Tool - Central Note.

- A maximum of sixteen SAP HANA scale-up systems passed the TDI KPIs on a single Hitachi Virtual Storage Platform F400 with FMD and RAID-10 (2D+2D) using HWCCT revision 112
- A maximum of five SAP HANA scale-up systems passed the TDI KPIs on a single Hitachi Virtual Storage Platform F400 with FMD and RAID-6 (6D+2P) using HWCCT revision 112

Note - Since the release of SAP HANA TDI in November 2013 several versions of the HWCCT have been published. To check whether or not the hardware configuration of your SAP HANA TDI infrastructure meets SAP KPIs, it is crucial that you use the same version of the 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 the HWCCT for your tests.

Solution Overview

This document provides an example configuration of the storage layout for SAP HANA platform nodes with variable sizes of main memory consolidated onto a Hitachi Virtual Storage Platform F400, tested within the Hitachi Data Systems lab environment.

This configuration uses the following storage components:

- **Hitachi Virtual Storage Platform F400**— Storage virtualization system designed to manage storage assets more efficiently

Figure 1 shows the server to storage configuration of this solution using VSP F400 with sixteen SAP HANA systems.

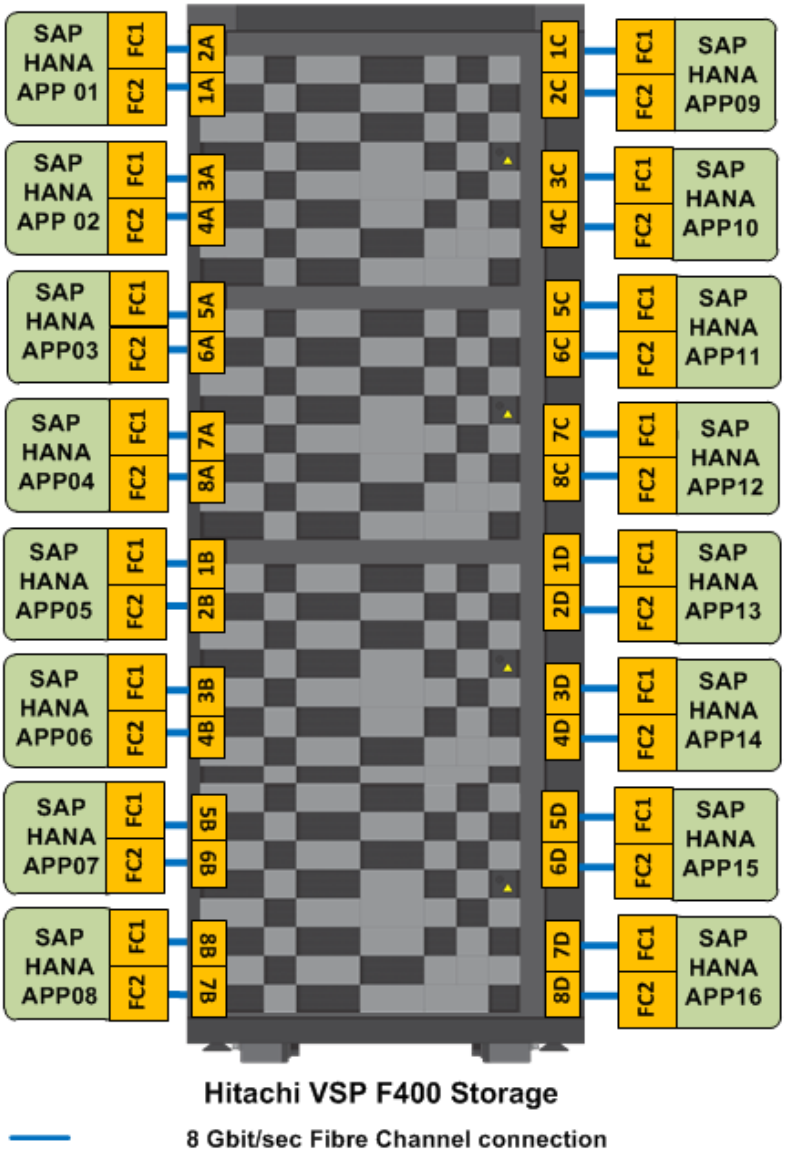

Figure 1

Key Solution Elements

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

Hardware Elements

Table 1 describes the hardware required to test the scalability of sixteen active nodes on VSP F400.

Table 1. Hardware Elements

Hardware	Quantity	Configuration	Role
Hitachi Virtual Storage Platform F400 with FMDs	1	▪ 8 Gb/sec channel blades	Block storage for SAP HANA platform nodes
SAP HANA Server	16	▪ Rack servers/blade chassis certified for SAP HANA with 256 GB (SAP HANA platform 1.0 SPS09, rev. 97 or later) nodes. A list of certified configurations can be found here	SAP HANA servers with 256 GB of main memory
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

SAP HANA Server

Only certified SAP HANA servers can be used in the SAP HANA TDI environment following the exact same bill of materials as the certified SAP HANA appliance server, but without the storage or local disks. A list of all certified servers and enterprise storage solutions can be found in the SAP HANA Hardware Directory. For more information regarding SAP HANA TDI, consult the SAP HANA TDI FAQ document on SCN.

Software Elements

Table 2 describes the software used to deploy the configuration for the testing.

Table 2. Software Elements

Software
SUSE Linux Enterprise Server for SAP Applications
SAP HANA platform
Hitachi Storage Navigator Modular 2
Hitachi Command Suite
Hitachi Virtual Storage Platform F400 Microcode

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

Solution Design

This is the detailed solution example for the SAP HANA tailored data center integration on Hitachi Virtual Storage Platform F400 using FMD with four active nodes. It includes components described in the following sections.

Fibre Channel Architecture

The SAP HANA platform nodes are directly attached to the designated Hitachi VSP target port.

Table 3 shows the storage port mapping.

Table 3. Storage Port Mapping

SAP HANA Platform Node	Fibre Channel Port	Hitachi Virtual Storage Platform F400 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:

- OS volume
- SAP HANA shared volume for the SAP HANA binaries and other configuration files.
- SAP HANA log volume
- SAP HANA data volume

This SAP HANA TDI setup utilizes an Hitachi Dynamic Provisioning (HDP) pool design for the storage layout that ensures maximum utilization and optimization at a lower cost than other solutions.

The layout can use one of the HDP pool layout options as listed in the Table 4, based on the required scalability.

Table 4. HDP Pools

Option	Hitachi Data Provisioning Pool Name	Purpose	Parity Group RAID Level and disks	Scalability
1	OS_SH_Log_Data_Pool	OS, SAP HANA Shared, Log and Data	RAID-10 (2D+2D) on 3.2 TB FMD; or	16 Nodes
2	OS_SH_Log_Data_Pool	OS, SAP HANA Shared, Log and Data	RAID-6 (6D+2P) on 3.2 TB FMD	5 Nodes

A minimum of one parity group is needed per HDP pool to fit in sixteen SAP HANA productive nodes on a Hitachi Virtual Storage Platform F400 using RAID-10 (2D+2D). Similarly, a minimum of one parity group is needed per HDP pool to fit in five SAP HANA productive nodes on a Hitachi Virtual Storage Platform F400 using RAID-6 (6D+2P). 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.

The following information shows an example layout of the HDP pool configuration on the Virtual Storage Platform F400 used in the SAP HANA tailored data center integration solution, with four active SAP HANA systems.

- System 1: 512 GB
- System 2: 1 TB
- System 3: 2 TB
- System 4: 4 TB

Storage for the systems should be provisioned as follows:

- The parity groups should be created first as shown in Table 5 if using the RAID-10 design, or Table 6 if using the RAID-6 design.
- A dynamic provisioning pool named **OS_SH_Log_Data_Pool** should be used to provision the operating system volume, SAP HANA shared volume, and data volume for the four SAP HANA systems on Virtual Storage Platform F400. All of the LDEVs should then be assigned to the pool.

Table 5. HDP Pool Provisioning with RAID-10

Hitachi Data Provisioning Pool	Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size	MPU Assign ment
OS_SH_Log_Data_ Pool	1	RAID-10 (2D+2D) on 3.2 TB FMD	00:00:01	OS_SH_LG_DA_Pool_1	2150.40 GB	MPU-10
			00:00:02	OS_SH_LG_DA_Pool_2	2150.40 GB	MPU-11
			00:00:03	OS_SH_LG_DA_Pool_3	2150.40 GB	MPU-20
OS_SH_Log_Data_ Pool	2	RAID-10 (2D+2D) on 3.2 TB FMD	00:00:04	OS_SH_LG_DA_Pool_4	2150.40 GB	MPU-21
			00:00:05	OS_SH_LG_DA_Pool_5	2150.40 GB	MPU-10
			00:00:06	OS_SH_LG_DA_Pool_6	2150.40 GB	MPU-11
OS_SH_Log_Data_ Pool	3	RAID-10 (2D+2D) on 3.2 TB FMD	00:00:07	OS_SH_LG_DA_Pool_7	2150.40 GB	MPU-20
			00:00:08	OS_SH_LG_DA_Pool_8	2150.40 GB	MPU-21
			00:00:09	OS_SH_LG_DA_Pool_9	2150.40 GB	MPU-10

Table 6. HDP Pool Provisioning with RAID-6

Hitachi Data Provisioning Pool	Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
OS_SH_Log_Data_P ool	1	RAID-6 (6D+2P) on 3.2 TB FMD	00:00:01	OS_SH_LG_DA_Pool_1	2457.50 GB	MPU-10
			00:00:02	OS_SH_LG_DA_Pool_2	2457.50 GB	MPU-11
			00:00:03	OS_SH_LG_DA_Pool_3	2457.50 GB	MPU-20
			00:00:04	OS_SH_LG_DA_Pool_4	2457.50 GB	MPU-21
			00:00:05	OS_SH_LG_DA_Pool_5	2457.50 GB	MPU-10
			00:00:06	OS_SH_LG_DA_Pool_6	2457.50 GB	MPU-11
			00:00:07	OS_SH_LG_DA_Pool_7	2457.50 GB	MPU-20
			00:00:08	OS_SH_LG_DA_Pool_8	2457.50 GB	MPU-21

- Virtual volumes for each of the nodes should then be provisioned as follows:
  - VVOLs for the OS, SAP HANA shared, Log and Data volumes should be created as shown in Table 7.

Table 7. VVOLs for the SAP HANA Nodes

Hitachi Data Provisioning Pool	VVOL ID	VVOL Name	VVOL Size	MPU Assignment
OS_SH_Log_Data_Pool	00:02:00	HANA_OS_N1	100 GB	MPU-10
	00:03:00	HANA_OS_N2	100 GB	MPU-11
	00:04:00	HANA_OS_N3	100 GB	MPU-20
	00:05:00	HANA_OS_N4	100 GB	MPU-21
	00:02:01	HANA_SH_N1	512 GB	MPU-10
	00:03:01	HANA_SH_N2	1024 GB	MPU-11
	00:04:01	HANA_SH_N3	1024 GB	MPU-20
	00:05:01	HANA_SH_N4	1024 GB	MPU-21
	00:02:02	HANA_LOG_N1_1	75 GB	MPU-10
	00:02:03	HANA_LOG_N1_2	75 GB	MPU-11
	00:02:04	HANA_LOG_N1_3	75GB	MPU-20
	00:02:05	HANA_LOG_N1_4	75 GB	MPU-21
	00:03:02	HANA_LOG_N2_1	150 GB	MPU-10
	00:03:03	HANA_LOG_N2_2	150 GB	MPU-11
	00:03:04	HANA_LOG_N2_3	150 GB	MPU-20
	00:03:05	HANA_LOG_N2_4	150 GB	MPU-21
	00:04:02	HANA_LOG_N3_1	150 GB	MPU-10
	00:04:03	HANA_LOG_N3_2	150 GB	MPU-11
	00:04:04	HANA_LOG_N3_3	150 GB	MPU-20
	00:04:05	HANA_LOG_N3_4	150 GB	MPU-21
	00:05:02	HANA_LOG_N4_1	150 GB	MPU-10
	00:05:03	HANA_LOG_N4_2	150 GB	MPU-11
	00:05:04	HANA_LOG_N4_3	150 GB	MPU-20
	00:05:05	HANA_LOG_N4_4	150 GB	MPU-21
	00:02:06	HANA_DATA_N1_1	155 GB	MPU-10
	00:02:07	HANA_DATA_N1_2	155 GB	MPU-11
	00:02:08	HANA_DATA_N1_3	155 GB	MPU-20
	00:02:09	HANA_DATA_N1_4	155 GB	MPU-21
	00:03:06	HANA_DATA_N2_1	310 GB	MPU-10
	00:03:07	HANA_DATA_N2_2	310 GB	MPU-11
	00:03:08	HANA_DATA_N2_3	310 GB	MPU-20
	00:03:09	HANA_DATA_N2_4	310 GB	MPU-21
	00:04:06	HANA_DATA_N3_1	620 GB	MPU-10
	00:04:07	HANA_DATA_N3_2	620 GB	MPU-11
	00:04:08	HANA_DATA_N3_3	620 GB	MPU-20
	00:04:09	HANA_DATA_N3_4	620 GB	MPU-21
	00:05:06	HANA_DATA_N4_1	1240 GB	MPU-10
	00:05:07	HANA_DATA_N4_2	1240 GB	MPU-11
	00:05:08	HANA_DATA_N4_3	1240 GB	MPU-20
	00:05:09	HANA_DATA_N4_4	1240 GB	MPU-21

- While mapping the LUN path assignment for each node, VVOLs should be added in the following order:
  - The OS volume for the specific SAP HANA platform node
  - The SAP HANA shared for the specific SAP HANA platform node
  - Finally the log volumes and data volumes for the specific SAP HANA platform node

Table 8 shows an example configuration of the LUN path assignment for Node001. The LUN assignment should be similar for all of the other nodes.

Table 8. Example LUN Path Assignment for Node001 for 512 GB SAP HANA Node Configuration

LUN ID	LDEV ID	LDEV Name
0000	00:02:00	HANA_OS_N1
0001	00:02:01	HANA_SH_N1
0002	00:02:02	HANA_LOG_N1_1
0003	00:02:03	HANA_LOG_N1_2
0004	00:02:04	HANA_LOG_N1_3
0005	00:02:05	HANA_LOG_N1_4
0006	00:02:06	HANA_DATA_N1_1
0007	00:02:07	HANA_DATA_N1_2
0008	00:02:08	HANA_DATA_N1_3
0009	00:02:09	HANA_DATA_N1_4

Best Practices of Storage Setup for SAP HANA TDI

- Create an HDP pool with a minimum of two parity groups whenever possible
- A parity group should be dedicated to one pool only, and should not be used for other purposes if one of its LDEVs is a Pool Volume
- Distribute the parity groups across at least two DBS trays if possible
- Create four VVOLs for log volumes per SAP HANA system and distribute them across the various MPUs
- Create four VVOLs for data volumes per SAP HANA system and distribute them across the various MPUs
- Use full allocation to provision DPVOLs whenever possible

Recommendations from SAP for I/O optimization

Starting with SAP HANA SPS10, SAP introduced parameters to adjust the I/O behavior and to optimize databases to work best with the file system and storage system used. Performance tests were conducted to define optimal values. The parameters listed in Table 9 must be set with the SAP hdbparam tool.

Table 9. hdbparam Parameters

Parameter	Value
async_read_submit	on
async_write_submit_blocks	all

For more information on how to use hdbparam, refer to documentation from SAP. To change the parameters during installation, refer to the SAP Note 2267798 - Configuration of the SAP HANA Database during Installation Using hdbparam.

Engineering Validation

Test Methodology

- HWCCT revision 112 was used for the scalability testing on Virtual Storage Platform F400.
- The initial testing started with one node and one parity group provisioned for the Hitachi Dynamic Provisioning pool. Additional storage was then provisioned to evaluate the requirements to fit that node.
- Additional nodes were added if the KPIs passed, or an additional parity group per HDP pool was added if they failed.

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

AS-513-00. June 2016