HITACHI
Inspire the Next

WHITE PAPER

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F1500 All-Flash or Hitachi Virtual Storage Platform G1500 with Flash Module Drives using Dynamic Provisioning Pools

Reference Architecture Guide

By Abhishek Dhanuka

February 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.

Table of Contents

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform F1500 All-Flash or Hitachi Virtual Storage Platform G1500 with Flash Module Drives using Dynamic Provisioning Pools

Reference Architecture Guide

The purpose of this reference architecture guide is to help customers who prefer the Tailored Data Center Integration (TDI) deployment scenario to implement SAP HANA on a Hitachi Virtual Storage Platform F1500 (VSP F1500) all-flash or Hitachi Virtual Storage Platform G1500 (VSP G1500) with Flash Module Drives (FMDs) storage array.

This approach provides greater flexibility, lowers the initial investment, and lowers the operational costs when deploying SAP HANA in the data center compared to the traditional appliance model.

This technical paper assumes you have familiarity with the following:

- Storage area network (SAN) based storage systems
- Network attached storage (NAS) systems
- General storage concepts
- General network knowledge
- SAP HANA platform
- Common IT storage practices
- SAP HANA 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.

Solution Overview

Unlike a SAP HANA appliance, in which the hardware vendor preconfigures all hardware components, SAP HANA tailored data center integration deployments are customized solutions. You choose a server from any certified SAP HANA server vendors, along with storage from any certified SAP HANA enterprise storage, to implement SAP HANA. This provides you an opportunity to leverage your existing hardware to reduce the total cost of ownership (TCO).

Using this reference architecture, you can deploy SAP HANA solutions for real time data processing using hardware from any certified SAP HANA server vendor and VSP F1500 all-flash or VSP G1500 with FMDs. View a list of SAP-certified servers available for SAP HANA appliances in Certified 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. 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 server blades from Hitachi Data Systems.

Every SAP-certified enterprise storage platform must meet the TDI storage KPI requirements set by SAP. Testing showed that the storage design of VSP F1500 all-flash or VSP G1500 with FMDs meets the TDI storage KPI requirements as needed for the SAP HANA platform laid out by SAP.

It is not mandatory to use the same storage design that was used for storage KPI testing as demonstrated in this reference architecture guide. Refer to the SAP HANA Tailored Data Center Integration FAQ (PDF) for more details about TDI.

Note — SAP recommends that TDI customers run the HWCCT tool in their environment to ensure that their specific HANA TDI implementation meets SAP performance criteria.

During validation, extensive testing was carried out and scalability and storage KPI testing was performed using the HANA-HWC-ES-1.1 certification scenario. SAP HANA Hardware Configuration Check Tool (HWCCT) was used to perform the storage KPI testing. See SAP Note 1943937 — Hardware Configuration Check Tool — Central Note (user logon required):

- A maximum of 32 SAP HANA systems passed the TDI KPIs on a single Hitachi Virtual Storage Platform F1500 all-flash with 1 VSD (CPU) pair and 256 GB of cache using HWCCT revision 200.
- A maximum of 32 SAP HANA systems passed the TDI KPIs on a single Hitachi Virtual Storage Platform G1500 via FMDs with 1 VSD (CPU) pair and 256 GB of cache using HWCCT revision 200.

Note — Since the release of SAP HANA TDI in November 2013, several versions of HWCCT have been published. To check whether or not the hardware configuration of your SAP HANA TDI infrastructure meets the KPIs from SAP, 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 correct version of HWCCT for your tests.

VSP F1500 all-flash or VSP G1500 with FMDs can support up to 128 nodes. Table 1shows the node scalability based on the components added to the storage. This includes cache memory, microprocessor server blade, and controller.

TABLE 1. NODE SCALABILITY FOR VSP F1500 ALL-FLASH OR VSP G1500 WITH FMDS

Number of Nodes	Cache Memory (GB)	VSD (CPU) Pair	Number of Controllers
32	256	1	1
64	512	2	1
96	768	3	2
128	1024	4	2

Comparisons and more information about the storage and the flash module drives supported can be found in the links below:

- VSP F1500
- VSP G1500
- Hitachi Flash Module Drives

Solution Design

The following provides solution layouts for a SAP HANA systems implementation on VSP F1500 all-flash and VSP G1500 with FMDs.

Storage Architecture

Each SAP HANA node requires the following storage layout:

- Operating system (OS) volume
- SAP HANA shared volume for the SAP HANA binaries and other configuration files
- Log volume
- Data volume

The environment uses dynamic provisioning pools created with Hitachi Dynamic Provisioning (HDP) in a tailored data center integration (TDI) approach to implement SAP HANA.

The layout uses two dynamic provisioning pools with the specific type of parity groups listed in Table 2.

TABLE 2. DYNAMIC PROVISIONING POOLS FOR VSP F1500 ALL-FLASH AND VSP G1500 WITH FMDS

Dynamic Provisioning Pool Name	Purpose	Parity Group RAID Level and Disks on VSP F1500	Parity Group RAID Level and Disks on VSP G1500
OS_SH_Data_Pool	Operating system (OS), SAP HANA Shared, and Data	RAID-10 (2D+2D) on 7 TB FMD	RAID-10 (2D+2D) on 3.5 TB FMD
Log_Pool	Log	RAID-10 (2D+2D) on 7 TB FMD	RAID-10 (2D+2D) on 3.5 TB FMD

Note — VSP F1500 all-flash supports 7 TB FMD and 14 TB FMD only whereas VSP G1500 supports 1.7 TB FMD, 3.5 TB FMD, 7 TB FMD, and 14 TB FMD.

Table 3 shows the minimum number of number of parity groups needed per dynamic provisioning pool for the various combinations of nodes on VSP F1500 all-flash or VSP G1500 with FMDs using 1 VSD pair and 256 GB of cache. Additional parity groups might need to be added based on capacity and performance requirements.

TABLE 3. MINIMUM PARITY GROUPS NEEDED PER NODE UP TO 32 NODES ON VSP F1500 ALL-FLASH OR VSP G1500 WITH FMDS

Nodes	OS_SH_Data_Pool	Log_Pool	Cache (GB)	VSD Pair
1	1	1	256	1
2	1	1		
3	1	1		
4	1	1		
5	1	1		
6	1	1		
7	1	1		
8	1	1		
9	1	1		
10	1	2		
11	1	2		
12	2	2		
13	2	2		
14	2	2		
15	2	2		
16	2	2		
17	2	2		
18	3	2		
19	3	3		
20	3	3		
21	3	3		
22	3	3		
23	3	3		
24	3	3		
25	4	3		
26	4	3		
27	4	3		
28	4	4		
29	4	4		
30	4	4		
31	4	4		
32	5	4		

Table 4 shows the minimum number of number of parity groups needed per dynamic provisioning pool for the various combinations of nodes up to 128 nodes on VSP F1500 all-flash or VSP G1500 with FMDs.

TABLE 4. MINIMUM PARITY GROUPS NEEDED UP TO 128 NODES ON VSP F1500 ALL-FLASH OR VSP G1500 WITH FMDS

Nodes	OS_SH_Data_Pool	Log_Pool	Cache (GB)	VSD Pair	Controller
32	5	4	256	1	1
64	10	8	512	2	1
96	15	12	768	3	2
128	20	16	1024	4	2

The example layout below shows a dynamic provisioning pool configuration on VSP F1500 all-flash and VSP G1500 with FMDs used in a SAP HANA TDI solution, with a SAP HANA 2 TB 4-socket scale-out system in a 2+1 configuration.

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 per 4 worker nodes

Storage for the systems should be provisioned as follows:

- The parity groups should be created first as shown in Table 5 if using VSP F1500, or Table 6 if using VSP G1500 with FMDs.
- A dynamic provisioning pool named **OS_SH_DT_Pool** should be used to provision the operating system volume, SAP HANA shared volume, and Data volume.
- A dynamic provisioning pool named **Log_Pool should be used to provision the Log volume**

TABLE 5. HDP DYNAMIC POOL PROVISIONING FOR VSP F1500

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size	MPB Assignment
OS_SH_DT_Pool	1	RAID-10 (2D+2D) on 7 TB FMD	00:00:01	OS_SH_DT_1	1600 GB	MPB-0
			00:00:02	OS_SH_DT_2	1600 GB	MPB-4
			00:00:03	OS_SH_DT_3	1600 GB	MPB-0
			00:00:04	OS_SH_DT_4	1600 GB	MPB-4
			00:00:05	OS_SH_DT_5	1600 GB	MPB-0
			00:00:06	OS_SH_DT_6	1600 GB	MPB-4
			00:00:07	OS_SH_DT_7	1600 GB	MPB-0
			00:00:08	OS_SH_DT_8	1600 GB	MPB-4
Log_Pool	2	RAID-10 (2D+2D) on 7 TB FMD	00:01:01	LG_1	1600 GB	MPB-0
			00:01:02	LG_2	1600 GB	MPB-4
			00:01:03	LG_3	1600 GB	MPB-0
			00:01:04	LG_4	1600 GB	MPB-4
			00:01:05	LG_5	1600 GB	MPB-0
			00:01:06	LG_6	1600 GB	MPB-4
			00:01:07	LG_7	1600 GB	MPB-0
			00:01:08	LG_8	1600 GB	MPB-4

TABLE 6. HDP DYNAMIC POOL PROVISIONING FOR VSP G1500 WITH FMDS

Dynamic Provisioning Pool	Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size	MPB Assignment
OS_SH_DT_Pool	1	RAID-10 (2D+2D) on 3.5 TB FMD	00:00:01	OS_SH_DT_1	1600 GB	MPB-0
			00:00:02	OS_SH_DT_2	1600 GB	MPB-4
			00:00:03	OS_SH_DT_3	1600 GB	MPB-0
			00:00:04	OS_SH_DT_4	1600 GB	MPB-4
Log_Pool	2	RAID-10 (2D+2D) on 3.5 TB FMD	00:01:01	LG_1	1600 GB	MPB-0
			00:01:02	LG_2	1600 GB	MPB-4
			00:01:03	LG_3	1600 GB	MPB-0
			00:01:04	LG_4	1600 GB	MPB-4

- Virtual volumes for each of the nodes should then be provisioned as shown in Table 7.

TABLE 7. VMWARE VSPHERE VIRTUAL VOLUMES (VVOLS) FOR THE SAP HANA NODES IN A 2+1 SCALE-OUT CONFIGURATION

Dynamic Provisioning Pool	VVol ID	VVol Name	VVol Size	MPB Assignment
OS_SH_DT_Pool	00:02:00	HANA_OS_N1	100 GB	MPB-0
	00:03:00	HANA_OS_N2	100 GB	MPB-4
	00:04:00	HANA_OS_N3	100 GB	MPB-0
	00:02:01	HANA_SH_1	500 GB	MPB-0
	00:03:01	HANA_SH_2	500 GB	MPB-4
	00:04:01	HANA_SH_3	500 GB	MPB-1
	00:05:01	HANA_SH_4	500 GB	MPB-5
	00:02:06	HANA_DATA_N1_1	500 GB	MPB-0
	00:02:07	HANA_DATA_N1_2	500 GB	MPB-4
	00:02:08	HANA_DATA_N1_3	500 GB	MPB-0
	00:02:09	HANA_DATA_N1_4	500 GB	MPB-4
	00:03:06	HANA_DATA_N2_1	500 GB	MPB-0
	00:03:07	HANA_DATA_N2_2	500 GB	MPB-4
	00:03:08	HANA_DATA_N2_3	500 GB	MPB-0
	00:03:09	HANA_DATA_N2_4	500 GB	MPB-4

TABLE 7. VMWARE VSPHERE VIRTUAL VOLUMES (VVOLS) FOR THE SAP HANA NODES IN A 2+1 SCALE-OUT CONFIGURATION (CONTINUED)

Dynamic Provisioning Pool	VVol ID	VVol Name	VVol Size	MPB Assignment
LG_Pool	00:02:02	HANA_LOG_N1_1	128 GB	MPB-0
	00:02:03	HANA_LOG_N1_2	128 GB	MPB-4
	00:02:04	HANA_LOG_N1_3	128 GB	MPB-0
	00:02:05	HANA_LOG_N1_4	128 GB	MPB-4
	00:03:02	HANA_LOG_N2_1	128 GB	MPB-0
	00:03:03	HANA_LOG_N2_2	128 GB	MPB-4
	00:03:04	HANA_LOG_N2_3	128 GB	MPB-0
	00:03:05	HANA_LOG_N2_4	128 GB	MPB-4

- The VVols created for SAP HANA Share are mapped to the storage ports connected to the Hitachi NAS Platform. A shared filesystem is created which is then exported and mounted as /hana/shared/<SID> on all three of the nodes.
- For mapping the LUN path assignment for each node, add the VVols in the following order:

  (1) Map the operating system volume for the specific SAP HANA node

  (2) Map the log volume

  (3) Map the data volume.

Table 8 shows an example configuration of the LUN path assignment for Node 1. The LUN assignment is similar for all of the other nodes.

TABLE 8. EXAMPLE OF LUN PATH ASSIGNMENT FOR NODE01 IN A SAP HANA NODE CONFIGURATION

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

Best Practices of Storage Setup for SAP HANA TDI

Follow these best practices when setting up your storage in a SAP HANA TDI environment:

- Create a dynamic provisioning pool using Hitachi Dynamic Provisioning with a minimum of two parity groups if possible.
- Dedicate a parity group to one dynamic provisioning pool only. Do not use the parity group for other purposes if one of its LDEVs is a pool volume.
- Configure dynamic provisioning pools as RAID-10.
- Distribute the parity groups across at least two FMD trays.
- Create four VVols for Log volumes for each SAP HANA node. Distribute the VVols for the Log volumes across the various VSDs.
- Create four VVols for Data volumes for each SAP HANA system. Distribute the VVols for the Log volumes across the various VSDs.
- Use **full allocation** to provision VVols whenever possible.
- Use direct attached storage connections for up to 32 Nodes
- Each set of front end ports is shared with two hosts using a Fibre Channel switch for more than 32 Nodes up to a maximum of 64 Nodes per Controller Frame.

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 following parameters must be set with the SAP hdbparam tool.

TABLE 9. PARAMETERS FOR I/O OPTIMIZATION

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 the install refer to the SAP Note 2267798 - Configuration of the SAP HANA Database during Installation Using hdbparam.

Engineering Validation

Test Methodology

- HWCCT revision 200 was used for the scalability testing on VSP F1500 all-flash and VSP G1500 with FMDs.
- The initial testing started with one node and one parity group provisioned for the HDP 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 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-577-00, February 2017.