TECHNICAL PAPER

Hitachi Unified Compute Platform 6000 for SAP HANA Dynamic Tiering on Hitachi Compute Blade 2500

Reference Architecture Guide

By David Pascuzzi

April 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

Key Solution Components 4

Hardware Components 4
Software Elements 10

Solution Design 13

Hitachi Compute Blade 2500 Chassis Configuration 13
520X B2 Server Blade Architecture 13
Fibre Channel Architecture 14
Network Architecture 15
Management Server 16
Storage Architecture 17
Parity Group, RAID, and LUN Configuration 18
Impact on SAP HANA Node Sizing 23
Operating System Storage Configuration 23

Hitachi Unified Compute Platform 6000 for SAP HANA Dynamic Tiering on Hitachi Compute Blade 2500

Reference Architecture Guide

This reference architecture guide describes the configuration used for a SAP HANA Tailored Datacenter Integration (TDI) solution to extend Hitachi Unified Compute Platform for the SAP HANA platform (UCP for SAP HANA) to provide SAP HANA dynamic tiering. This architecture works in scale-up configurations using either SUSE Linux Enterprise Server for SAP Applications (SLES) or Red Hat Enterprise Linux for SAP HANA (RHEL) with the Resilient Storage Add-On. This solution can be adapted to work in a scale-out environment, but it isn't covered here.

This reference architecture does not apply to virtualized implementations and has not been tested in scale-out configurations.

Unified Compute Platform for SAP HANA is a preconfigured analytical appliance, ready to plug in to a network to provide real-time access to operational data for use in analytic models.

SAP HANA dynamic tiering allows SAP HANA to offload data to an Extended Storage Server referred to as the SAP HANA dynamic tiering node. This allows SAP HANA to process data that is normally accessed (hot) without the overhead of less active data (warm). This solution can greatly extend the amount of data that SAP HANA can process by moving warm data to the SAP HANA dynamic tiering node.

This reference architecture describes a bolt-on extension to a scale-up solution for Unified Compute Platform for SAP HANA that shares common components.

The solution without high availability adds one server blade for the SAP HANA dynamic tiering node. For non-production instances or in cases when high availability is provided differently, only the one blade needs to be added

SAP HANA database runs mission critical applications and it is important that these systems remain available to users at all times. This requires that these systems are able to make fast recovery after system component failure. This reference architecture guide has an option to include high availability for SAP HANA systems when using SAP HANA dynamic tiering. To provide high availability, add a second server blade with SAP HANA dynamic tiering and a second server blade with SAP HANA are added to the solution.

Information regarding the scale-up solution for Hitachi Unified Compute Platform for SAP HANA is found in the appropriate reference architecture guide. This solution can apply to any scale-up solution for Hitachi Unified Compute Platform for SAP HANA. For this reference architecture a 1 TB scale-up solution is used.

This reference architecture includes the following:

- **Hitachi Compute Blade 2500** — This has enterprise computing power and performance with flexible I/O architecture and logical partitioning. Multiple applications easily and securely co-exist in the same chassis.
  - This solution uses 520X B2 server blades.
- **Hitachi Virtual Storage Platform Gx00 models —** These scale for all data types, flexibly adapting for performance, capacity, and multi-vendor storage The base SAP HANA solution needs its storage replaced with Virtual Storage Platform G400.
- **SAP HANA with SAP Dynamic Tiering** — This is a multi-purpose, in-memory database to analyze transactional and analytical data.
- **QuantaPlex T41S-2U server** — Provides a central point for management of all components used in this solution.
- **Brocade ICX 6430-48 switch** — 48-port 1 GbE switch that provides a management network to the appliance.
- **Brocade VDX 6740-48 switch** — 48-port switch that provides 10 GbE external connectivity to the appliance.
- One of the following operating systems:
  - SUSE Linux Enterprise Server for SAP Applications
  - Red Hat Enterprise Linux for SAP HANA

Figure 1 shows the topology for this reference architecture for a base solution using a single blade scale-up solution for Hitachi Unified Compute Platform for SAP HANA dynamic tiering.

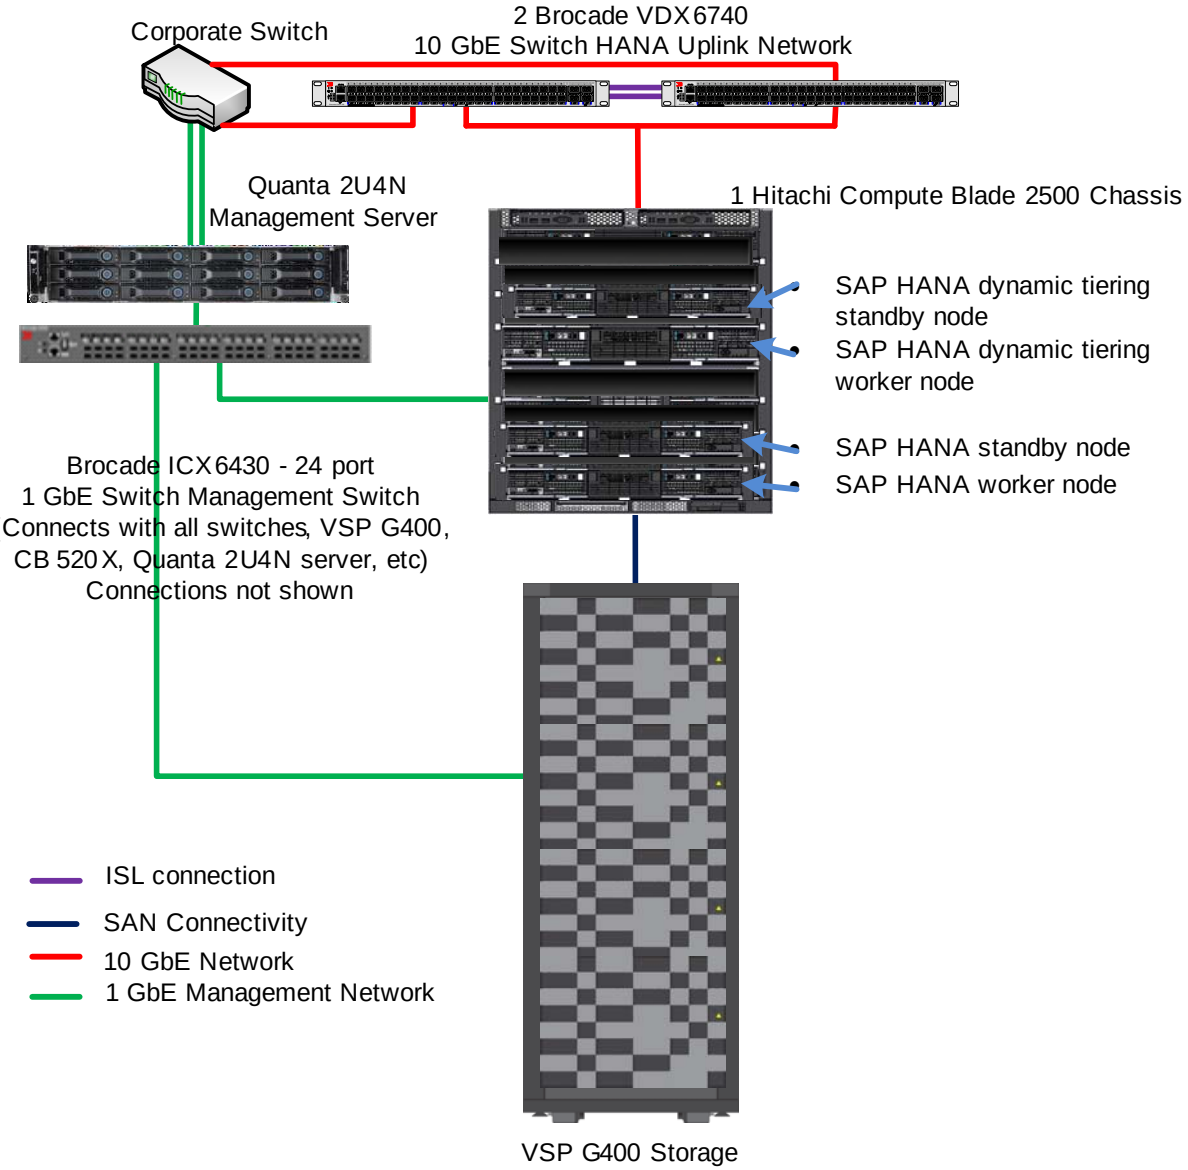

Figure 1

Key Solution Components

The following are the key hardware and software components used in this reference architecture.

Hardware Components

This solution consists of one scale-up appliance for SAP HANA extended by the hardware needed for SAP HANA dynamic tiering. Optionally, high availability can be provided with extra server blades for SAP HANA and SAP HANA dynamic tiering. This reference architecture uses a 2-socket 1 TB appliance, as shown in Table 1.

Table 1. SAP HANA Scale-up Appliance Key Hardware Elements

Hardware	Quantity	Configuration	Role
Hitachi Compute Blade 2500	1	■ 8 Server blade chassis ■ 2 Management modules ■ 10 Cooling fan modules ■ 3 Power supply modules ■ 4 I/O board modules ■ 2 × 10 GbE 2-port LAN PCIe adapter ■ 2 Hitachi 16 Gb/sec 2-port Fibre Channel adapters	Server blade chassis
520X B2 server blade (2 socket)	1	■ 2 × 18-core Intel 8880v3 Haswell EX processors ■ 1 TB ■ 1 × 2-port pass through mezzanine card on mezzanine slot 2 and mezzanine slot 4 of server blade 1 and server blade 2 (on single server blade for 2-Socket)	SAP HANA server
Hitachi Virtual Storage Platform G400	1	■ Single frame ■ This solution replaces the standard VSP G200 with a VSP G400	Block storage for SAP HANA and dynamic tiering nodes

To provide SAP HANA dynamic tiering, extend the solution by adding a server blade, switches, and management server. Table 2 lists the hardware used to deploy SAP HANA dynamic tiering without high availability. Table 3 lists the hardware used to deploy SAP HANA dynamic tiering with high availability.

Table 2. Additional Hardware Elements for SAP HANA Dynamic Tiering Base Solution Without High Availability

Hardware	Quantity	Configuration (per Unit)	Role
520X B2 server blade (2-socket)	1	■ 2 × 18-core Intel 8880v3 Haswell EX processors ■ 512 GB RAM	SAP HANA dynamic tiering worker node
Items added to the CB 2500 chassis		■ 2 × 10 GbE 2-port LAN PCIe adapter ■ 2 Hitachi 16 Gb/sec 2-port Fibre Channel adapters	
QuantaPlex T41S-2U server 2U4N server	1	■ T41S 2U4N 2.5 inch bay chassis ■ 1 T41S server node ■ 3 Server fillers ■ 2 Intel Xeon E5-2620 v3 (6C 2.4GHz 85W) processors ■ 1 Heat sink, CPU0 and CPU1 ■ 2 × 16 GB DDR4 2,133 MHz memory module ■ 2 × 500 GB, 7200 RPM, 2.5 inch SATA HDD ■ 1 Dual-port 10 GigE Intel 82599ES SFP+ OCP mezzanine card ■ 1 Dual-port 1 GigE Base-T Intel i350 mezzanine card ■ 1 Dual-port 8 Gb/sec Fibre Channel Emulex HBA	Management server
Brocade ICX 6430-48 port switch	1	■ 1 GbE	1 GbE management network
Brocade VDX 6740-48 port switch	2	■ 10 GbE	Client network

Table 3. Additional Hardware Elements for SAP HANA Dynamic Tiering Base Solution With High Availability

Hardware	Quantity	Configuration (per Unit)	Role
520X B2 server blade (2-socket)	1	▪ 2 × 18-core Intel 8880v3 Haswell EX processors ▪ 512 GB RAM	SAP HANA dynamic tiering standby node
520X B2 server blade (2-socket)	1	▪ 2 × 18-core Intel 8880v3 Haswell EX processors ▪ 1024 GB RAM	SAP HANA standby node
Items added to chassis		▪ 2 × 10 GbE 2-port LAN PCIe adapter per node ▪ 2 Hitachi 16 Gb/sec 2-port Fibre Channel adapters per node	

Hitachi Compute Blade 2500

Hitachi Compute Blade 2500 delivers enterprise computing power and performance with unprecedented scalability and configuration flexibility. Lower your costs and protect your investment.

Flexible I/O architecture and logical partitioning allow configurations to match application needs exactly with Hitachi Compute Blade 2500. Multiple applications easily and securely co-exist in the same chassis.

Add server management and system monitoring at no cost with Hitachi Compute Systems Manager. Seamlessly integrate with Hitachi Command Suite in Hitachi storage environments.

This configuration adds one 520X B2 server blade in the existing Hitachi Compute Blade 2500 chassis. In the event that there are no free slots, provide another Hitachi Compute Blade 2500 chassis for this solution. When providing high availability, add a second server blade for SAP HANA dynamic tiering and add a second server blade for SAP HANA.

Table 4 lists the configuration for the 520X B2 server blades used for the SAP HANA worker node and SAP HANA standby node.

Table 4. SAP HANA node 520X B2 Server Blade Configuration

Feature	Configuration
Processors	2 Intel Xeon E7-8880 processors
Processor SKU	Intel Xeon E7-8880 v3
Processor frequency	2.3 GHz
Processor cores	18 cores
Number of server blades	1
Number of DIMMS	48 × 16 GB DIMMs, 1 TB

Table 4. SAP HANA node 520X B2 Server Blade Configuration (Continued)

Feature	Configuration

Network Ports	2 × 2-port 10GBASE-SR LAN PCIe adapter on two I/O board modules:
	■ Worker node
	  ■ IOBD 09B
	  ■ IOBD 10B
	■ Standby node (HA solutions)
	  ■ IOBD 11B
	  ■ IOBD 12B

Fibre Channel Ports	2 Hitachi 16 GB/sec 2-port Fibre Channel adapters on two I/O board modules:
	■ Worker node
	  ■ IOBD 09A
	  ■ IOBD 10A
	■ Standby node (HA solutions)
	  ■ IOBD 11A
	  ■ IOBD 12A

Table 5 lists the configuration for the 520X B2 server blades used for the SAP HANA dynamic tiering worker node and standby node.

Table 5. SAP HANA Dynamic Tiering Node 520X B2 Server Blade Configuration

Feature	Configuration
Processors	2 Intel Xeon E7-8880 processors
Processor SKU	Intel Xeon E7-8880 v3
Processor frequency	2.3 GHz
Processor cores	18 cores
Number of server blades	1
Number of DIMMS	32 × 16 GB DIMMs, 512 GB

Table 5. SAP HANA Dynamic Tiering Node 520X B2 Server Blade Configuration (Continued)

Feature	Configuration

Network Ports	2 × 2-port 10GBASE-SR LAN PCIe adapter on two I/O board modules
	■ Worker node
	  ■ IOBD 09B
	  ■ IOBD 10B
	■ Standby node (HA solutions)
	  ■ IOBD 11B
	  ■ IOBD 12B

Fibre Channel Ports	2 Hitachi 16 GB/sec 2-port Fibre Channel adapters on two I/O board modules:
	■ Worker node
	  ■ IOBD 09A
	  ■ IOBD 10A
	■ Standby node (high availability solutions)
	  ■ IOBD 11A
	  ■ IOBD 12A

Figure 2 shows the layout of the I/O board modules from the back of Hitachi Compute Blade 2500 chassis.

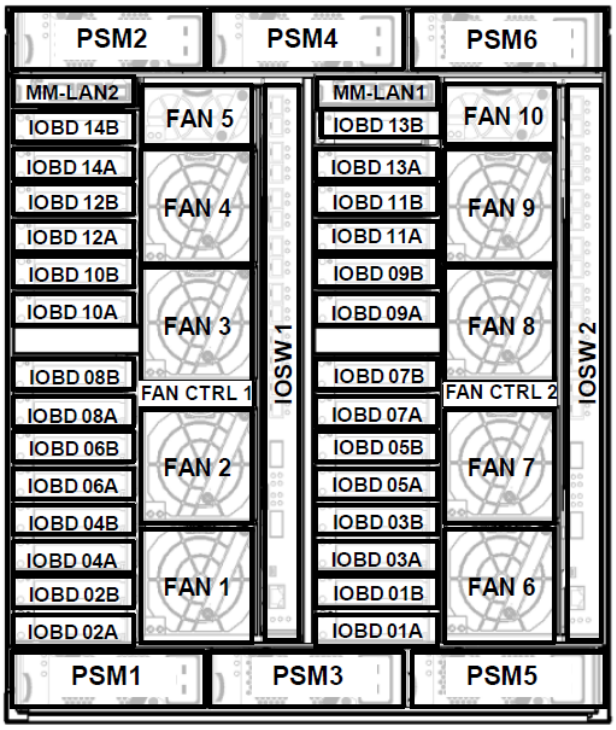

Figure 2

Hitachi Virtual Storage Platform Gx00 Models

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 Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort. New management software improves ease of use to save time and reduce complexity. The infrastructure of Storage Virtualization Operating System creates a management framework for improved IT response to business demands.

This solution requires Hitachi Virtual Storage Platform G400 or higher.

The SAP HANA operating system LUNs, SAP HANA data, and SAP HANA log LUNs reside on this storage array, as well as the shared LUNs for SAP HANA configuration files, binaries, and traces.

This solution creates an extra LDEV on the parity group used for the SAP HANA operating system for the SAP HANA Dynamic Tiering work node. Also, this solution uses new parity groups for SAP HANA dynamic tiering storage data and log files.

When providing high availability, create two more LDEVs on the parity group used for the SAP HANA operating system: one for each of the standby nodes for the operating system.

Brocade Switches

Brocade and Hitachi Data Systems partner to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

The solution using the following Brocade products:

- Brocade VDX 6740-48 port switch
- Brocade ICX 6430-48 port switch

QuantaPlex T41S-2U Server

QuantaPlex T41S-2U is an ultra-dense design equipped with four independent nodes. It creates the flexibility to set up different workloads independently in one 2U shared infrastructure, providing optimal data center performance per dollar.

Software Elements

SAP HANA

The SAP HANA platform is flexible, multipurpose in-memory software. It combines SAP software components optimized to specific hardware. These components come from leading hardware partners of SAP, including Hitachi Data Systems.

The SAP HANA platform allows customers to analyze large volumes of data in real time. It is also a development platform, providing an infrastructure and tools for building high-performance applications based on SAP HANA Extended Application Services (SAP HANA XS). It is the foundation of various SAP HANA editions, like the SAP HANA Platform Edition, providing core database technology, and the SAP HANA Enterprise Edition, bundling additional components for data provisioning. The SAP HANA Platform Edition integrates a number of SAP components, including the SAP HANA database, SAP HANA studio and SAP HANA clients.

As a SAP customer, you can get more information on the SAP HANA platform at the SAP Service Marketplace and help.sap.com.

See **Installation and Upgrade Information** on SAP HANA Platform for the following guides and all other SAP-related documentation:

- **The SAP HANA Master Guide** — This guide is the entry point for planning the installation of your SAP HANA system landscape
- **The SAP HANA Master Update Guide** — This guide describes how to update the SAP HANA platform and related components
- **SAP HANA Server Installation and Update Guide** — This guide describes how to install and update an SAP HANA system with the SAP HANA life-cycle management tools.
  - There are various related installation guides to install the required SAP in-memory database and the other software components for the different replication technologies.
- **SAP HANA Technical Operations Manual** — This guide is the entry point for administering and operating your SAP HANA system landscape.

SAP Integration and Certification Center (SAP ICC) provides information about SAP Certified Appliance Hardware for SAP HANA.

SAP HANA Dynamic Tiering

SAP HANA dynamic tiering provides the ability to keep data in either memory or on the disk in a columnar format. This helps you choose memory for hot data and disk for warm data, helping to strike the right price versus performance balance. In normal operations SAP HANA dynamic tiering is expected to have less than six concurrent queries. Because of this, the SAP HANA dynamic tiering node doesn't have the same demands as the SAP HANA node; it can be a smaller size node.

See **Installation and Upgrade Information** on SAP HANA dynamic tiering for the following guides:

- **The SAP HANA Dynamic Tiering Master Guide** — This guide is the entry point for planning the installation of your SAP HANA dynamic tiering system landscape.
- **SAP HANA Dynamic Tiering Installation Guide** — This guide describes how to install SAP HANA dynamic tiering option with the SAP HANA life-cycle management tools.
- **SAP HANA Dynamic Tiering Administration Guide** — This guide is the entry point for administering and operating your SAP HANA Dynamic Tiering node.

The SAP HANA dynamic tiering node and its storage do not require certification.

SUSE Linux Enterprise Server and Red Hat Enterprise Linux

Hitachi Unified Compute Platform for the SAP HANA Platform in a scale-up configuration can run on either of the following:

- 64-bit SUSE Linux Enterprise Server for SAP Applications
- 64-bit Red Hat Enterprise Linux for SAP HANA with Resilient Storage Add-on

The SAP HANA dynamic tiering node must run the same operating system version as the SAP HANA node.

The initially delivered configuration of the operating system should persist. Changing the configuration settings can cause significant performance problems to occur.

Do not make any modifications to the operating system, except as noted or approved by SAP. To modify your operating system configuration or the installed software packages, follow the guidelines given by SAP and the operating system distributor. See the following SAP Notes for SLES and RHEL:

- 1944799 — SAP HANA Guidelines for SLES Operating System Installation
- 2136965 — SAP HANA DB: Recommended OS settings for RHEL 6.6

For more details, see the section “Updating and Patching the Operating System,” in the SAP HANA Technical Operations Manual.

High Availability

This solution can optionally provide high availability. To provide high availability, Use the Fiber Channel storage connector from SAP HANA. Find more information at http://scn.sap.com/docs/DOC-64726.

To provide fault recovery, SAP HANA software includes a watchdog function. This automatically restarts configured services (index server, name server, and so on) in case of their failure. In addition to these features, SAP and its partners offer another high availability mechanism for SAP HANA. These solutions are based on completely redundant servers and/or storage.

Add one standby node for SAP HANA and another standby node for SAP HANA dynamic tiering to a SAP HANA system. In case of failure, the standby node takes over the data and log volumes of a failed worker node to become a worker node. This solution does not need additional storage for high availability, only the servers.

Data Backup

This reference architecture does not cover backup and recovery. The solution can be extended to provide data backup by using one of the following:

- Any SAP HANA dynamic tiering certified backint program
- Add shared storage to provide an area to store the backups

See the SAP HANA dynamic tiering guides for information on supported SAP HANA dynamic tiering backups.

Solution Design

The detailed design for this reference solution extension to Hitachi Unified Compute Platform for SAP HANA Dynamic Tiering in a scale-up configuration includes the following:

- Hitachi Compute Blade 2500 Chassis Configuration
- 520X B2 Server Blade Architecture
- Fibre Channel Architecture
- Network Architecture
- Storage Architecture
- Parity Group, RAID, and LUN Configuration
- Impact on SAP HANA Node Sizing
- Operating System Storage Configuration

Hitachi Compute Blade 2500 Chassis Configuration

This solution uses one Hitachi Compute Blade 2500 chassis with the following components added to the chassis:

- One or two 520X B2 server blades for SAP HANA. See Table 4, “SAP HANA node 520X B2 Server Blade Configuration,” on page 6
- One or two 520X B2 server blades for SAP HANA dynamic tiering. See Table 5, “SAP HANA Dynamic Tiering Node 520X B2 Server Blade Configuration,” on page 7
- Extend the configuration by mounting four to eight I/O board modules (IOBD), as follows:
  - Install Hitachi FIVE-FX 16 Gb/sec 2-port Fibre channel PCI-E adapters on the A-slot of each I/O board module.
  - Install 10GBase-SR 2-port network PCI-E adapters on the B-slot of each I/O board module.

520X B2 Server Blade Architecture

The base solution uses one server blade for the SAP HANA worker node and one server blade for the SAP HANA dynamic tiering worker node. When providing high availability, add a SAP HANA standby node and SAP HANA dynamic tiering standby node.

In normal operations SAP HANA dynamic tiering is expected to have less than six concurrent queries. Because this node doesn't have the same demands as the SAP HANA node, it can be a smaller size node.

Table 4 on page 6 lists the information for the SAP HANA server blade or server blades configuration.

Table 5 on page 7 lists the information for the SAP HANA dynamic tiering server blade or server blades configuration.

Fibre Channel Architecture

As shown in Table 6, the solution uses four to eight Hitachi 16 GB/sec 2-port Fibre Channel adapters installed on the PCIe slot of the I/O board module of each server blade. This solution uses eight to sixteen 16 GB/sec Fibre Channel ports on Hitachi Virtual Storage Platform G400 directly attached to Hitachi Compute Blade 2500 using the Fibre Channel PCIe adapters.

Table 6. Fibre Channel Port Mapping

Usage	Blade Number	PCIe Slot Number	Port Number	Hitachi Virtual Storage Platform Port
SAP HANA Dynamic Tiering worker node	Blade 9	IOBD 09A	0	1B
			1	2B
		IOBD 10A	0	3B
			1	4B
SAP HANA worker node	Blade 1	IOBD 01A	0	1A
			1	2A
		IOBD 02A	0	3A
			1	4A
SAP HANA Dynamic Tiering standby node	Blade 11	IOBD 11A	0	1D
			1	2D
		IOBD 12A	0	3D
			1	4D
SAP HANA standby node	Blade 3	IOBD 03A	0	1C
			1	2C
		IOBD 04A	0	3C
			1	4C

This configuration supports high availability by providing multiple paths from the host within Hitachi Compute Blade 2500 to multiple ports on Virtual Storage Platform G400.

Set the port properties for the point-to-point connection between Hitachi Compute Blade 2500 and Virtual Storage Platform G400 storage, as shown in Table 7.

Table 7. Port Properties

Property	Value
Port Security	**Disabled**
Port Speed	**Auto (16 Gb/sec)**
Fabric	**ON**
Connection Type	**P-to-P**

Network Architecture

The SAP HANA dynamic tiering configuration of the SAP HANA solution requires two separate networks.

- **Client Network** — Dedicated to traffic between the SAP HANA database, SAP HANA dynamic tiering, and client software
  - The SAP HANA dynamic tiering nodes and SAP HANA nodes are connected to the client network using a Brocade VDX 6740 switch.
    - Set an MTU size of 9100 in accordance with Brocade best practices.
  - Connect each of the ports on both VDX switches to two ports of each server blade, utilizing the active-active bonding mode.
  - Connect the uplink ports of the Brocade VDX 6740-48 switches to your network.
- **Management Network —** Used for hardware management in the solution
  - This resides on a 1 GbE Brocade ICX 6430-48 port switch.
  - The management network does not need to have a VLAN assigned to it.
  - The Brocade ICX 6430-48 port switch uses the default switch configuration.

There are two 10GBASE-SR 2-port LAN adapters installed on the PCIe slots of the I/O board module of each server blade on the Hitachi Compute Blade 2500 chassis. The solution has two 10 GbE ports on the 2 × 10GBASE-SR 2-port LAN adapters on each server blade, connected to the external switch. The compute network setup uses the ports on the 10GBASE-SR 2-port LAN adapters, as listed in Table 8 on page 16.

To connect to the client network, the operating system is configured to provide active-active network bond, as follows:

- Bond the two ports connected to the client netowork,eth9901 and eth9902, as bond0
- Use active-active network bond mode with options "mode= 802.3ad miimon=100 xmit_hash_policy=layer3+4 updelay=5000 lacp_rate=fast"

Table 8. Network Setup Using 10GBASE-SR 2-Port LAN Adapters

Usage	Server Blade	PCIe Slot Number	Switch Module Port	Network Description
SAP HANA dynamic tiering worker node	Blade 9	IOBD 09B	0	Client network
			1	Unused
		IOBD 10B	0	Client network
			1	Unused
SAP HANA worker node	Blade 1	IOBD 01B	0	Client network
			1	Unused
		IOBD 02B	0	Client network
			1	Unused
SAP HANA dynamic tiering standby node (HA)	Blade 11	IOBD 11B	0	Client network
			1	Unused
		IOBD 12B	0	Client network
			1	Unused
SAP HANA standby node (HA)	Blade 3	IOBD 03B	0	Client network
			1	Unused
		IOBD 04B	0	Client network
			1	Unused

Management Server

This solution uses a Quanta T41S 2U server for management. The management server acts as a central device for managing the SAP HANA platform.

Manage the following from the management server:

- Hitachi Compute Blade 2500 chassis
  - 520X B2 server blades
- Brocade ICX 6430 — 48 port switch
- Brocade VDX 6740 — 48 port switches
- SAP HANA node or nodes

- SAP HANA dynamic tiering node or nodes
- Hitachi Virtual Storage Platform G400
- NTP configuration
- Hitachi Command Suite and management of the server blades
- SAP HANA Studio

Figure 3 shows the management server network ports using one dual port 1 GbE Base-T Intel i350 mezzanine card.

- **Slot 01 Port 2** — Connected to the client provided external switch. It provides the 1 GbE network to the management server.
- **Slot 01 Port 1** — Connected to the Brocade ICX 6430-48 port switch. It provides the 1 GbE management network to all components

The management server has the following additional components:

- One dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card
- One Emulex 2-port 8 Gb/sec Fibre Channel HBA on the PCIe slot

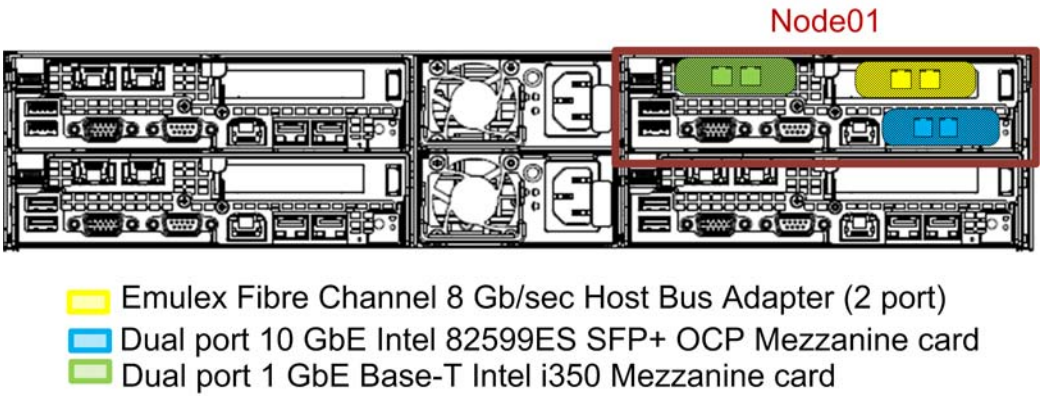

Figure 3

Storage Architecture

Hitachi Virtual Storage Platform G400 is used for this SAP HANA dynamic tiering solution. Sizing and configuring of storage varies for different size solutions, taking requirements into account. This includes the following:

- Storage drive box trays (DBS)
- Spare drives
- Operating system volume (OS)
- Existing SAP HANA shared volume (/hana/shared)
- SAP HANA dynamic tiering log volume (/hana/log_es)
- SAP HANA Dynamic Tiering data volume (/hana/data_es)

For this solution, the changes to the standard SAP HANA parity groups are the addition of one to three 100 GB LUN on the operating system parity group. If Red Hat Enterprise Linux with the Resilient Storage Add-On is used, add one 50 MB LUN on the operating system parity group to be used as a STONITH device.

To provide for storage for SAP HANA dynamic tiering, add additional disk drives to the storage array.

Parity Group, RAID, and LUN Configuration

This reference architecture uses various configurations of parity group, RAID level, and LUN on Hitachi Virtual Storage Platform G400 for different size solutions.

For this solution using SAP HANA dynamic tiering, Hitachi Data Systems adds three types of storage:

- SAP HANA dynamic tiering Data
- SAP HANA dynamic tiering Log
- An optional STONITH LUN

SAP HANA Worker Node

For the SAP HANA worker node, use the standard parity group and LDEV configuration. Table 9 lists the configuration.

Table 9. SAP HANA Parity Group and LDEV

Parity Group	RAID Level and Disks	LDEV ID	LDEV Size	HOST ID	MPU ID	Description
1	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	0:01:00	100 GB	000	MPU10	Operating system volume
		0:00:01	3 TB	001	MPU11	SAP HANA shared volume
2	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	0:01:01	150 GB	002	MPU20	SAP HANA log volume
	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	0:01:02	150 GB	003	MPU10	
	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	0:01:03	150 GB	004	MPU21	
	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	0:01:04	150 GB	005	MPU11	

Table 9. SAP HANA Parity Group and LDEV (Continued)

Parity Group	RAID Level and Disks	LDEV ID	LDEV Size	HOST ID	MPU ID	Description
3	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:02:01	750 GB	006	MPU21	SAP HANA data volume
		00:02:02	750 GB	007	MPU11	
		00:02:03	750 GB	008	MPU20	
		00:02:04	750 GB	009	MPU10	

SAP HANA Dynamic Tiering Worker Node

For data storage, 3 parity groups are used for each 18 TB of SAP HANA dynamic tiering compressed data. The recommended configuration is based upon the standard six concurrent query workload and storage system throughput requirements. This configuration will provide ample performance.

The sizing of the log partition needs to be large enough to store the largest SAP HANA dynamic tiering data modification (load, update, delete), and in most cases one parity group will be adequate. To load more than 3 TB of data at a time, parity groups will need to be added to the solution.

This reference architecture uses 600 GB hard disk drives; they can be replaced with solid state drives.

Table 10 shows the raw storage sizes and how much compressed data can be stored on them. SAP HANA dynamic tiering can see up to a 2.5 data compression factor. The sizing presented is uncompressed and includes the following:

- Data storage
- SAP HANA dynamic tiering overhead
- 20% temporary storage, this size is the SAP recommended size. Depending on the queries it could be significantly larger.
- 5% metadata storage
- 16 GB per core storage requirement.

Table 10. LDEVs for Various Data Storage Sizes

Raw Physical Storage Size	Compressed data size	Recommended number of Parity Groups (6D+2P) by Hitachi Data Systems	
		Data	Log
9 TB	18 TB	3	1
18 TB	36 TB	6	1
27 TB	54 TB	9	1
36 TB	72 TB	12	1
45 TB	90 TB	15	1

Table 11 shows the parity group configuration when using Red Hat Enterprise Linux. Figure 4 shows a sample LUN layout for a 9 TB raw storage solution. The STONITH LUN is not part of the solution when using SUSE Linux Enterprise Server with one LUN for /hana/shared.

Table 11. Storage Configuration

Parity Group	RAID Level and Disks	LDEV ID	LDEV Size	LUN ID	MPU ID	Description	Raw physical storage size in TB
Existing Parity Group	Existing operating system disks	0:05:01	100 GB	000	MPU10	One 100 GB operating system LDEV provided for each node added to the solution	9,18,27,36,45
Existing Parity Group	Existing HANA Shared	Existing from HANA installation	Existing from HANA installation	001	Existing from HANA installation	Existing from HANA installation	9,18,27,36,45
Existing Parity Group	Existing operating system disk	0:05:03	50 MB	002	MPU10	STONITH LUN: Only for Red Hat Enterprise Linux with Resilient Storage Add-on. This LUN is shared by the base and high availability nodes	9,18,27,36,45 RHEL Solutions
3	RAID-6 (6D+2P)	0:05:04	.5 TB	003	MPU20	.75 TB Log	9,18,27,36,45
		0:05:05	.5 TB	004	MPU21	.75 TB Log	9,18,27,36,45
		0:05:06	.5 TB	005	MPU10	.75 TB Log	9,18,27,36,45
		0:05:07	.5 TB	006	MPU11	.75 TB Log	9,18,27,36,45
4	RAID-6 (6D+2P)	0:05:09	3 TB	007	MPU20	3 TB Data LDEV	9,18,27,36,45
5	RAID-6 (6D+2P)	0:05:09	3 TB	008	MPU21	3 TB Data LDEV	9,18,27,36,45

Table 11. Storage Configuration

Parity Group	RAID Level and Disks	LDEV ID	LDEV Size	LUN ID	MPU ID	Description	Raw physical storage size in TB
6	RAID-6 (6D+2P)	0:05:09	3 TB	009	MPU10	3 TB Data LDEV	9,18,27,36,45
7	RAID-6 (6D+2P)	0:05:0A	3 TB	010	MPU11	3 TB Data	18,27,36,45
8	RAID-6 (6D+2P)	0:05:0B	3 TB	011	MPU20	3 TB Data	18,27,36,45
9	RAID-6 (6D+2P)	0:05:0C	3TB	012	MPU21	3 TB Data	18,27,36,45
10	RAID-6 (6D+2P)	0:05:0D	3TB	013	MPU10	3 TB Data	27,36,45
11	RAID-6 (6D+2P)	0:05:0E	3TB	014	MPU11	3 TB Data	27,36,45
12	RAID-6 (6D+2P)	0:05:0F	3TB	015	MPU20	3 TB Data	27,36,45
13	RAID-6 (6D+2P)	0:05:10	3TB	016	MPU21	3 TB Data	36,45
14	RAID-6 (6D+2P)	0:05:11	3TB	017	MPU10	3 TB Data	36,45
15	RAID-6 (6D+2P)	0:05:12	3TB	018	MPU11	3 TB Data	36,45
16	RAID-6 (6D+2P)	0:05:13	3TB	019	MPU20	3 TB Data	45
17	RAID-6 (6D+2P)	0:05:14	3TB	020	MPU21	3 TB Data	45
18	RAID-6 (6D+2P)	0:05:15	3TB	021	MPU10	3 TB Data	45

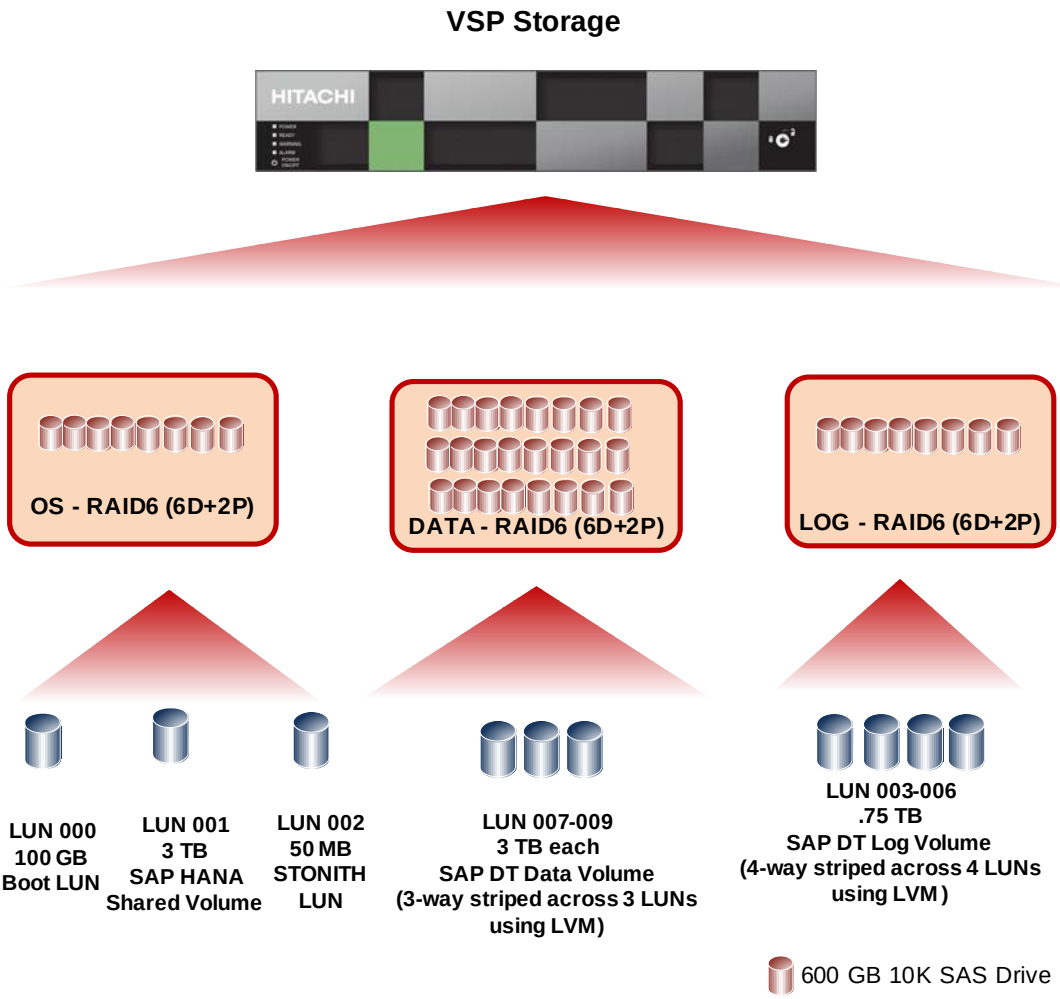
VSP Storage
HITACHI
OS - RAID6 (6D+2P)
DATA - RAID6 (6D+2P)
LOG - RAID6 (6D+2P)
LUN 000
100 GB
Boot LUN
LUN 001
3 TB
SAP HANA
Shared Volume
LUN 002
50 MB
STONITH
LUN
LUN 007-009
3 TB each
SAP DT Data Volume
(3-way striped across 3 LUNs
using LVM)
LUN 003-006
.75 TB
SAP DT Log Volume
(4-way striped across 4 LUNs
using LVM)
600 GB 10K SAS Drive

Figure 4

Standby Nodes

When providing high availability, use a SAP HANA standby node and a SAP HANA dynamic tiering standby node. These nodes need their own operating system LDEV. Table 12 shows the extra LDEVs created for the standby nodes.

Except for the operating system LDEV, connect the standby nodes to the same LDEVS as the corresponding worker node.

Table 12. High Availability Nodes LDEVs

Parity Group	RAID Level and Disks	LDEV ID	LDEV Size	LUN ID	MPU ID	Description
Existing	Existing OS Disks	0:06:01	100 GB	000	MPU10	100 GB OS LDEV for the SAP HANA standby node
Existing	Existing OS Disks	0:07:01	100 GB	000	MPU10	100 GB OS LDEV for the SAP HANA dynamic tiering standby node

Impact on SAP HANA Node Sizing

When accessing SAP HANA dynamic tiering data, all requests and responses come through the SAP HANA node. These results are then processed by SAP HANA. This requires free memory on the SAP HANA node and temporary storage in the SAP HANA dynamic tiering node. If SAP HANA runs out of free memory, active queries will be canceled.

The recommendation is to size the SAP HANA node with at least 256 GB more memory than is needed for SAP HANA. Monitor the memory usage by the SAP HANA node. If needed, the size of the SAP HANA node can be increased.

When designing the hot and warm data layout, understand data access. Frequently accessed data should be in the SAP HANA node. For best performance, minimize joining data between the tiers.

Operating System Storage Configuration

This explains the SAP HANA operating system configuration.

Operating System Volume Configuration

Each node in this solution requires an operating system volume. Each volume is a 100 GB LUN on Hitachi Virtual Storage Platform.

The four ports of the 16 GB Hitachi FIVE-FX Fibre Channel 2-port adapters for each node have the 100 GB operating system LUN configured as the primary boot device. The operating system LUN holds partitions for SUSE Linux for SAP Applications, /usr/sap, and the Linux swap space.

Activate Device-mapper Multipath

This reference architecture uses Device-mapper Multipath, a component of the native Linux operating system.

Using Device-mapper Multipath allows the configuration of multiple I/O paths between the server blades. Hitachi Virtual Storage Platform multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available, unless all four paths fail.

Use Device-mapper Multipath for the following I/O paths:

- Operating system LUN of the SAP HANA node
- SAP HANA data volume LUNs
- SAP HANA log volume LUNs
- SAP HANA shared volume LUNs
- Operating system LUN of the SAP HANA dynamic tiering node
- SAP HANA dynamic tiering data volume LUNs
- SAP HANA dynamic tiering log volume LUNs
- SAP HANA dynamic tiering STONITH LUN

SAP HANA Data Volume Configuration

The Logical Volume Manager configures the SAP HANA persistent storage volumes. Table 9, “SAP HANA Parity Group and LDEV,” on page 18 shows the size of the SAP HANA data volume LUNs defined by the solution.

Logical Volume Manager creates a single striped volume on which the XFS file system is created to store the SAP HANA data volume. The striped volume acts as the persistent layer for the SAP HANA server.

Table 14, “File System Options for Data and Log Volumes,” on page 27 contains options used by striped volume creation, XFS file system creation, and mount for data volume to optimize performance. When using high availability, connect these LUNs to both SAP HANA nodes. SAP HANA mounts these LUNs to the active node.

SAP HANA Log Volumes Configuration

This reference architecture uses Logical Volume Manager to configure the SAP HANA log volumes.

With the four 150 GB SAP HANA log volume LUNs, Logical Volume Manager creates a single four-way striped volume on which the XFS file system is created to store the SAP HANA log volume.

Table 14, “File System Options for Data and Log Volumes,” on page 27 contains options used by striped volume creation, XFS file system creation, and mount for log volume and data volume to optimize performance. When using high availability, connect these LUNs to both SAP HANA nodes. SAP HANA mounts them to the active node.

SAP HANA Shared Volume Configuration

This solution uses shared storage for the SAP HANA shared volume. This volume is accessed simultaneously by all nodes. There are three options that can be used.

Network File System (NFS)

This reference architecture can use network file system (NFS) on Hitachi NAS Platform. NFS is a distributed file system. In a scale-out configuration, Hitachi Unified Compute Platform for SAP HANA already has an NAS Platform server providing NFS to the solution for the SAP HANA shared volume. This option is not covered in this reference architecture.

Oracle Clustered File System (OCFS2)

This reference architecture can use Oracle Clustered File System (OCFS2), and an optional clustered logical volume manager (cLVM). These are components of the SUSE Linux Enterprise Server operating system. They provide a shared clustered file system across nodes. OCFS2 needs to use cLVM when you need more than one LUN. Table 13, “Clustered Logical Volume Manager File System Configuration for OCFS2 and GFS2,” on page 26 lists the cLVM configuration.

Using OCFS2 allows the configuration to have a shared clustered file system for the SAP HANA shared volume without using an NFS server.

OCFS2 has the following limitations:

- It can only be used with SUSE Linux Enterprise Server.
- Maximum number of nodes is 32.
- Maximum file system size is 4 petabytes.

Global File System (GFS2)

This reference architecture can use Global File System (GFS2) and clustered logical volume manager (cLVM). These are components of the Red Hat Enterprise Linux and SUSE Linux Enterprise Server 12 operating systems. They provide a shared clustered file system across nodes.

To use GFS2 with Red Hat Enterprise Linux, you must purchase a license for Resilient Storage Add-On for all nodes.

Using GFS2 allows the configuration to have a shared clustered file system for the SAP HANA shared volume without using an NFS server. Table 13 on page 26 lists the cLVM configuration.

GFS2 has the following limitations:

- A STONITH LUN is required.
- Maximum file system size is 100 TB.
- Maximum number of nodes is 16.

Table 13. Clustered Logical Volume Manager File System Configuration for OCFS2 and GFS2

Option	Shared Volume
Logical volume stripe type	Striped across all Data LUNs
Logical volume extents	Using 100% of the free space
LVM stripe size	512
OCFS2 block size	4k
OCFS2 cluster size	4k
GFS2 lock table name	Hanadt:shared
GFS2 number of journals	4

SAP HANA Dynamic Tiering Data Volume Configuration

Logical Volume Manager configures the SAP HANA dynamic tiering persistent storage volumes. Refer to Table 10, “LDEVs for Various Data Storage Sizes,” on page 19 for the size of SAP HANA dynamic tiering data volume LUNs.

Logical Volume Manager creates a single striped volume on which the XFS file system is created to store the SAP HANA dynamic tiering data volume. The striped volume acts as the persistent layer for the SAP HANA dynamic tiering node.

Table 8, “Network Setup Using 10GBASE-SR 2-Port LAN Adapters ,” on page 16 contains options used by striped volume creation:

- XFS file system creation
- Mount options for data and log volumes.

When using high availability, connect these LUNs to both SAP HANA dynamic tiering nodes. SAP HANA mounts these nodes to the active node.

SAP HANA Dynamic Tiering Log Volume Configuration

Logical Volume Manager configures the persistent log volumes for SAP HANA dynamic tiering. Refer to Table 10, “LDEVs for Various Data Storage Sizes,” on page 19 for the size of the data volume LUNs for SAP HANA dynamic tiering. The recommended number of LUNs for log volumes is one. It creates a single striped volume on which the XFS file system is created to store the log volume for SAP HANA Dynamic Tiering. The striped volume acts as the persistent layer for the SAP HANA dynamic tiering node.

Table 14 on page 27 contains options used by striped volume creation:

- XFS file system creation
- Mount options for data and log volumes

When using high availability, connect these LUNs to both SAP HANA dynamic tiering nodes. SAP HANA mounts these nodes to the active node.

Table 14. File System Options for Data and Log Volumes

Option	Data Volume	Log Volume
Logical volume stripe type	Striped across all data LUNs	Striped across all log LUNs
Logical volume extents	Use 100% of the free space	Use 100% of the free space
LVM stripe size	512	512
XFS stripe unit	1024	1024
XFS stripe width	XFS stripe unit times number of non-parity disks	XFS stripe unit times number of non-parity disks
Write barrier support	Disabled	Disabled
64-bit inode support	Enabled	Enabled

SAP HANA STONITH LUN

Using Global File System or Oracle Clustered File System with clustered logical volume manager requires a STONITH LUN to provide fencing. This is a 50 MB LUN that all nodes can access. This is a raw device that is not formatted or mounted.

SAP HANA Software Installation

This solution uses the SAP HANA software components that were installed in the shared directory during the SAP HANA node installation.

This solution uses the SAP landscape management tools on the SAP HANA node to install the SAP HANA dynamic tiering node from the SAP HANA node. When providing high availability, this solution uses the SAP landscape management tools on the SAP HANA node to install the standby nodes from the SAP HANA node.

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-450-02. April 2016