HITACHI
Inspire the Next

WHITE PAPER

Hitachi Unified Compute Platform 6000 for the SAP HANA Platform, Scale-out Configuration using 2 TB SAP HANA Nodes with Intel Xeon E7-88xx v4 Processors using GFS2

Reference Architecture Guide

By Yingping Niu

July 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-578-00	**Initial release**	**March, 2017**
AS-578-01	**Removed “for Tailored Datacenter Integration (TDI)” on page 1.** **Added a note that per SAP's recommendation, move the working directory for the XSA application instances to a local directory of a host on page 1.** **Removed “using unicast (UDPU)” on page 16.** **Added a note regarding using GFS2 for XSA BP restriction on page 26.**	**July, 2017**

Table of Contents

Hitachi Unified Compute Platform 6000 for the SAP HANA Platform, Scale-out Configuration using 2 TB SAP HANA Nodes with Intel Xeon E7-88xx v4 Processors using GFS2

Reference Architecture Guide

This reference architecture guide describes how to deploy Hitachi Unified Compute Platform 6000 for the SAP HANA Platform using Global File System 2 (GFS2) to provide a shared binaries file system (/hana/shared) in a scale-out configuration:

- A 32 TB configuration using 2 TB SAP HANA nodes

Note - When using SAP HANA XS Advanced (XSA), by default XSA is installed to /hana/shared/<SID>xs/ea_data. However, for this solution, do not install XSA to /hana/shared. Per SAP's recommendation, move the working directory for the XSA application instances to a local directory of a host.

This reference architecture uses the following:

- **Hitachi Compute Blade 2500**

  This has enterprise computing power and performance with flexible I/O architecture and logical partitioning.

  This solution uses a various number of 520X B3 server blades to provide the flexibility of multiple scale-out configuration options.

- **Hitachi Virtual Storage Platform G800**

  VSP G800 scales for all data types, flexibly adapting for performance, capacity, and multi-vendor storage.

  In this solution, the persistent storage of the SAP HANA server resides on Hitachi Virtual Storage Platform G800.

- **Symmetric Multiprocessing Connector** — Connects multiple blades together into one unit.
- **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.
- **Rack Optimized Server for Solutions, 2U Four Nodes** — An ultra-dense design equipped with four independent nodes with the flexibility to set up different workloads independently in one 2U shared infrastructure.
  - This solution uses one node for use as a management server.
- **10 GbE 2-port LAN PCIe adapters**
- **Hitachi 16 Gb/sec 2-port Fibre Channel PCIe adapters**
- **SUSE Linux Enterprise Server (SLES) with SUSE High Availability Extension (SUSE HAE)**
- **Red Hat Enterprise Linux (RHEL) with High Availability (HA) Add-On**
- **Global File System 2 (GFS2)**
- **SAP HANA** — This is a multi-purpose, in-memory database to analyze transactional and analytical data.

Figure 1 shows the 32 TB configuration using 2 TB HANA nodes.

Figure 1

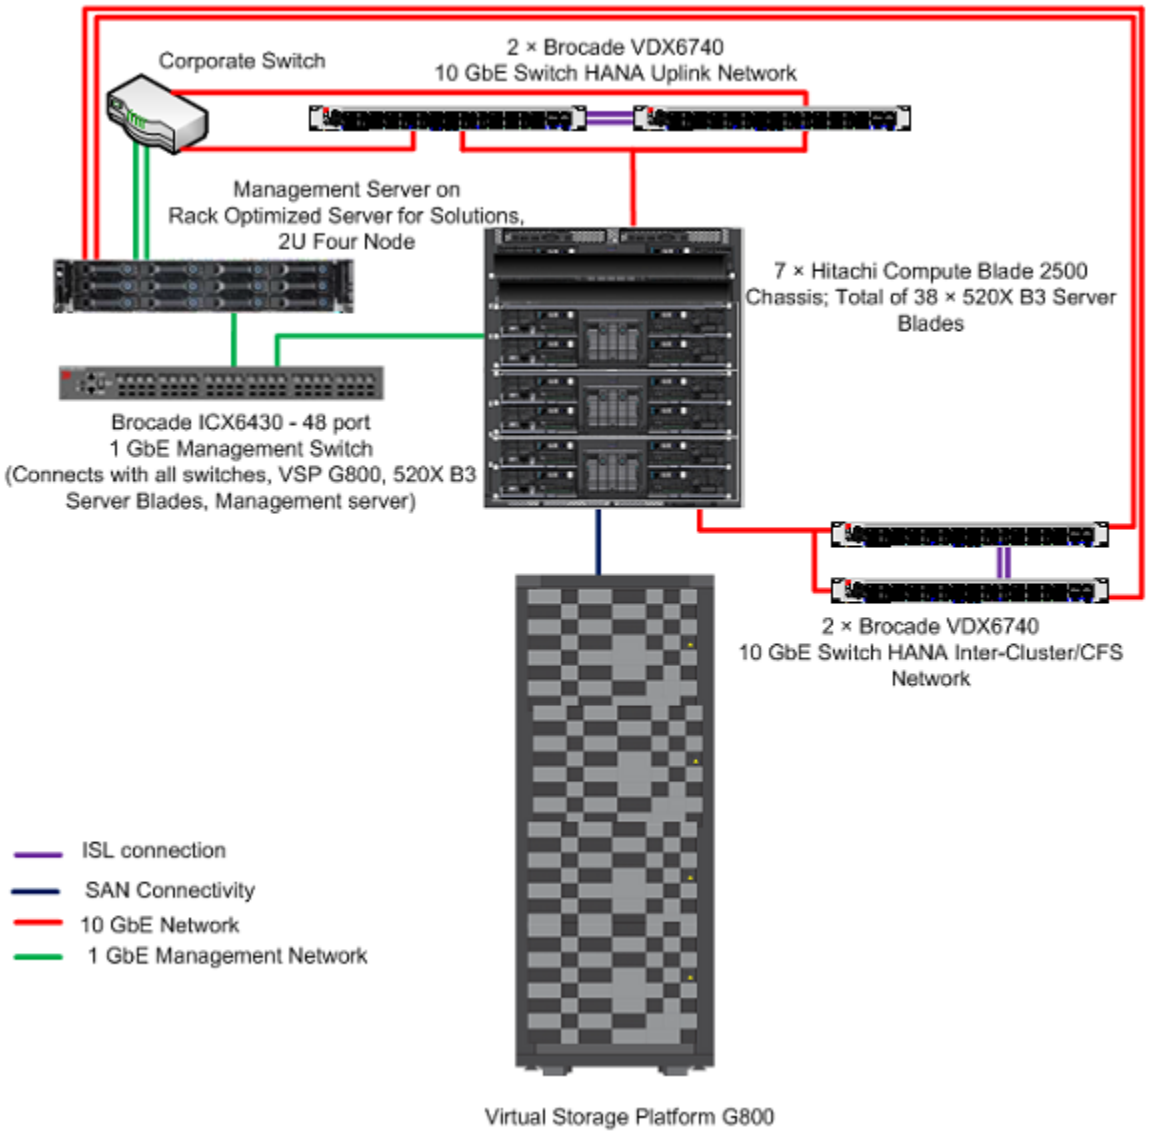

Note — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Key Solution Elements

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

Hardware Elements

Table 1 describes the hardware used to deploy the sixteen active node and three standby node configuration.

TABLE 1. HARDWARE ELEMENTS

Hardware	Quantity	Configuration	Role
Hitachi Compute Blade 2500 (CB 2500) chassis	7	■ 6 blades on chassis ■ 2 management modules ■ 10 cooling fan modules ■ 5 Power supply modules ■ 6 × Hitachi 16 Gb/sec 2-port FC PCIe cards (2 per node) ■ 12 × 2-port 10GBASE-SR PCIe cards (4 per node)	Server blade chassis
520X B3 server blade	38	■ 2 × 22-core or 24-core processors ■ 1024 GB RAM ■ 1 × Pass-through Mezzanine card on Mezzanine Slot 2 and 4	SAP HANA server
SMP connector module	19	■ 2-blade SMP connector board for 2 TB HANA nodes ■ SMP expansion module ■ SMP connector cover	SMP connector to turn two physical blades into one HANA node
Hitachi Virtual Storage Platform G800	1	■ CTL - 1 pair ■ DKB - 2 pairs ■ CHB - 3 pairs (up to 6 pairs for 16+3) ■ MPU - 2 pairs ■ Cache - 512 GB	Block storage for SAP HANA nodes Platform

TABLE 1. HARDWARE ELEMENTS (CONTINUED)

Hardware	Quantity	Configuration	Role
Rack Optimized Server for Solution, 2U Four Node	1	■ T41S 2U 2.5” Bay Chassis ■ 1 × T41S Server Node, with the following components: ■ 2 × Intel Xeon E5-2620 v4 (8C 2.1GHz 85W) ■ 1 × Heatsink CPU0 and CPU1 ■ 2 × 16GB DDR4 2,133MHz Memory Module ■ 2 × HDD SATA 500 GB, 7200rpm, 2.5 inch (6 Gb) ■ 1 × Dual port 10 GigE Intel 82599ES SFP+ OCP Mezzanine Card ■ 1 × Dual port 1 GigE Base-T Intel i350 Mezzanine Card ■ 1 × Emulex Dual Port 8 Gb/sec Fibre Channel HBA ■ 3 × Server Filler	Used as the management server that runs the following: ■ NTP ■ Hitachi Command Suite ■ Hi-Track Remote Monitoring system ■ SAP HANA Studio
Brocade VDX 6740-48 port switch	4	■ Two switches with distinct VLANs, each dedicated to SuSE HAE or RHEL HA cluster and SAP HANA inter-cluster network	10 GbE SuSE HAE or RHEL HA cluster communication and SAP HANA inter-cluster network
		■ Two switches with one VLAN to provide an uplink network to the customer network infrastructure	10 GbE client network
Brocade ICX 6430-48 port switch	1	■ 1 GbE ■ 48 ports	1 GbE Management Network

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.

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

- The 32 TB configuration uses 38 × 520X B3 server blades in 7 Hitachi Compute Blade 2500 chassis.

Symmetric Multiprocessing Connector

For multiple server blades, the solution uses symmetric multiprocessing (SMP) technology to combine multiple server blade resources into a single server.

- The 32 TB configuration uses a 2-blade SMP connection board to combine two 520X B3 server blades together in to a single server for use as a 2 TB HANA node.

Hitachi Virtual Storage Platform G800

Hitachi Virtual Storage Platform G800 (VSP G800) belongs to Hitachi Virtual Storage Platform family systems, which 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 G800, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System (SVOS) provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform G800. 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 uses a Virtual Storage Platform G800. The operating system LUNs, data LUNs, log LUNs, and LUNs for SAP HANA shared reside on this storage device.

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.

This solution uses the following Brocade products:

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

Software Elements

Host Operating System Options

The following options are available for use as an operating system with the solution:

- **SUSE Linux Enterprise Server (SLES) for SAP Applications**

  Compete more effectively though improved uptime, better efficiency, and accelerated innovation using SUSE Linux Enterprise Server. This is a versatile server operating system for efficiently, deploying highly available enterprise-class IT services in mixed IT environments with performance and reduced risk.

  SUSE Linux Enterprise Server was the first Linux operating system to be certified for use with SAP HANA. It remains the operating system of choice for the vast majority of SAP HANA customers.

- **Red Hat Enterprise Linux (RHEL)**

  Using the stability and flexibility of Red Hat Enterprise Linux, reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability on military-grade security. Use Enterprise Linux to tailor your infrastructure as markets shift and technologies evolve.

Changing the configuration settings is only supported along the guidelines of SAP and the operating system distributor and may otherwise cause significant performance problems. The following SAP Notes for SLES and RHEL are a good starting point for information on this topic:

- 1944799 - SAP HANA Guidelines for SLES Operating System Installation
- 2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL) Operating System

For more details, refer to “Updating and Patching the Operating System,” in the SAP HANA Technical Operations Manual from the SAP Help Portal.

SUSE Linux Enterprise High Availability Extension (HAE)

HAE is an integrated suite of open source clustering technologies that enables to implement highly available physical and virtual Linux cluster, and to eliminate single points of failure. High Availability Extension ships with a comprehensive set of tools to assist you in managing your cluster.

See SUSE Linux Enterprise HAE on the SUSE website for details.

Red Hat High-Availability Add-On

The High Availability Add-On for Red Hat Enterprise Linux Server provides continuous availability of services by using failover services between nodes within a cluster to eliminate single points of failure.

See High Availability Add-On Overview on the Red Hat website for the details.

Global File System 2 (GFS2)

Global File System 2 (GFS2) is a shared disk file system for Linux computer clusters. It allows all nodes to have direct concurrent access to the same shared block storage. GFS2 is included in High Availability Extension for SLES and in the Resilient Storage Add-On for RHEL.

The following links provide details of GFS2 on SLES and RHEL:

- https://www.suse.com/documentation/sle-ha-12/book_sleha/data/cha_ha_gfs2.html
- https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Global_File_System_2/ch-overview-GFS2.html

SAP HANA

SAP HANA converges database and application platform capabilities in-memory to transform transactions, analytics, text analysis, predictive and spatial processing so businesses can operate in real-time. This combines database, data processing, and application platform capabilities in a single in-memory platform. Also, the platform provides libraries for predictive, planning, text processing, spatial, and business analytics — all on the same architecture. This architecture comes from leading hardware partners of SAP, including Hitachi.

By eliminating the divide between transactions and analytics, SAP HANA allows you to answer any business question anywhere in real time.

SAP customers can download more information on the SAP HANA Platform at the SAP Service Marketplace and SAP Help Portal. The following resources are available:

- SAP HANA Master Guide

  This is the central starting point for the technical implementation of the SAP HANA platform. Use this for basic concepts and for planning the SAP HANA application system landscape.

- **SAP HANA Installation and Initial Configuration Guides**

  Use the various installation guides to install the required SAP In-Memory Database and the other software components for the different replication technologies. Refer to the SAP HANA Server Installation Guide for an overview of how to install SAP HANA.

- SAP HANA Technical Operations Manual

  This provides an end-to-end picture of the available SAP HANA appliance administration tools and the key tasks for a system administrator to perform.

- SAP HANA Master Update Guide

  This explains how to update SAP HANA and its components.

The following is a link to all SAP-related documentation:

http://help.sap.com/hana_platform/

The following is a link to information about SAP HANA appliances certified by SAP hardware partners:

http://scn.sap.com/community/icc

Solution Design

The detailed design for Hitachi Unified Compute Platform 6000 for the SAP HANA Platform in a scale-out configuration reference solution includes the following.

Scale-out Supported Solutions

This appliance supports the configurations listed in Table 2.

TABLE 2. SUPPORTED SCALE-OUT CONFIGURATIONS FOR SAP HANA

Total RAM	Active Nodes + 0 Standby Node	Active Nodes + 1 Standby Node	Active Nodes + 2 Standby Nodes	Active Nodes + 3 Standby Nodes
2	N/A	1+1	N/A	N/A
4	2+0	2+1	2+2	None
6	3+0	3+1	3+2	3+3
8	4+0	4+1	4+2	4+3
10	5+0	5+1	5+2	5+3
12	6+0	6+1	6+2	6+3

TABLE 2. SUPPORTED SCALE-OUT CONFIGURATIONS FOR SAP HANA (CONTINUED)

Total RAM	Active Nodes + 0 Standby Node	Active Nodes + 1 Standby Node	Active Nodes + 2 Standby Nodes	Active Nodes + 3 Standby Nodes
14	7+0	7+1	7+2	7+3
16	8+0	8+1	8+2	8+3
18	9 +0	9+1	9+2	9+3
20	10+0	10+1	10+2	10+3
22	11+0	11+1	11+2	11+3
24	12+0	12+1	12+2	12+3
26	13+0	13+1	13+2	13+3
28	14+0	14+1	14+2	14+3
30	15+0	15+1	15+2	15+3
32	16+0	16+1	16+2	16+3

This solution uses seven Hitachi Compute Blade 2500 chassis with the components as listed in Table 3.

TABLE 3. HITACHI COMPUTE BLADE 2500 CHASSIS COMPONENTS

Feature	Scale-out Configuration
Blades	6 × 520X B3
Blade Location	■ Blade 11 (non-primary) ■ Blade 9 (primary) ■ Blade 7 (non-primary) ■ Blade 5 (primary) ■ Blade 3 (non-primary) ■ Blade 1 (primary)
SMP	3 × 2 Blade SMP

TABLE 3. HITACHI COMPUTE BLADE 2500 CHASSIS COMPONENTS (CONTINUED)

Feature	Scale-out Configuration
Network ports	2 × 2-port 10GBASE-SR LAN PCIe adapter on two I/O board modules for each server blade in the following locations: ■ IOBD 01B ■ IOBD 03A ■ IOBD 05B ■ IOBD 07A ■ IOBD 09B ■ IOBD 11A ■ IOBD 02A ■ IOBD 04B ■ IOBD 06A ■ IOBD 08B ■ IOBD 10A ■ IOBD 12B
Fibre Channel Ports	2 × Hitachi 16 Gb/sec, 2-port Fibre Channel PCIe adapters on two I/O board modules for each HANA node in the following locations: ■ IOBD 01A ■ IOBD 05A ■ IOBD 09A ■ IOBD 04A ■ IOBD 08A ■ IOBD 12A
Other interfaces	For all sizes: ■ 1 USB 3.0 port ■ KVM connector (VGA, COM, USB 2.0 port)

Figure 2 shows the front and back view of a chassis in a scale-out configuration of 2 TB HANA nodes.

Figure 2

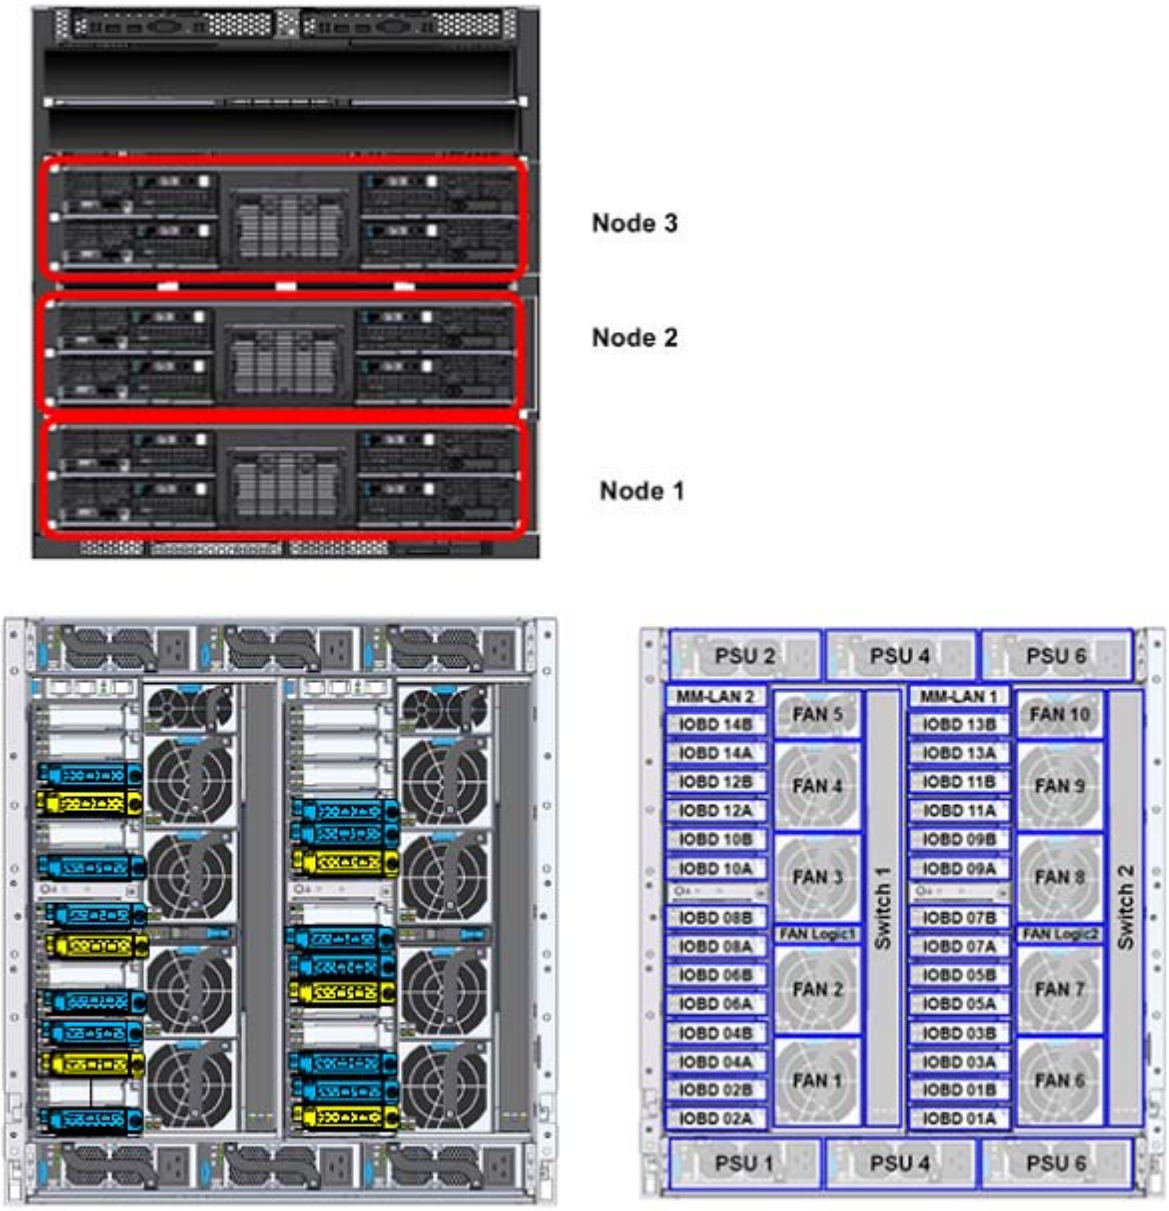

Table 4 shows the Hitachi Compute Blade 2500 chassis configuration of all of the SAP HANA nodes.

TABLE 4. HITACHI COMPUTE BLADE 2500 CHASSIS CONFIGURATION OF SAP HANA NODES

Chassis#	Blades#	SAP HANA Node Name	Role of SAP HANA Node
1	1 & 3	HANA Node 1	Master
1	5 & 7	HANA Node 2	Worker
1	9 & 11	HANA Node 3	Worker
2	1 & 3	HANA Node 4	Worker
2	5 & 7	HANA Node 5	Worker
2	9 & 11	HANA Node 6	Worker
3	1 & 3	HANA Node 7	Worker
3	5 & 7	HANA Node 8	Worker
3	9 & 11	HANA Node 9	Worker
4	1 & 3	HANA Node 10	Worker
4	5 & 7	HANA Node 11	Worker
4	9 & 11	HANA Node 12	Worker
5	1 & 3	HANA Node 13	Worker
5	5 & 7	HANA Node 14	Worker
5	9 & 11	HANA Node 15	Worker
6	1 & 3	HANA Node 16	Worker
6	5 & 7	HANA Node 17	Standby Node
6	9 & 11	HANA Node 18	Standby Node
7	1 & 3	HANA Node 19	Standby Node

Server Blade Architecture

For 2 TB SAP HANA nodes, the solution uses a total of 38 full-width server blades. Each SAP HANA node has two server blades that are connected using a two-blade SMP connection board to create a single four-socket SMP node with a total of either 88 or 96 cores and 2 TB of memory.

Refer to Table 5 for the server blade configuration.

TABLE 5. SERVER BLADE CONFIGURATION

	Server Blade Configuration
Total Server Blades	38
Processor SKU	Intel Xeon processor E7-8880 v4 or Intel Xeon processor E7-8890 v4
Processor Cores	22 or 24
Processor Frequency	2.2 GHz
Memory DIMM slots	■ 32 out of 48 are populated (per blade)
Memory	■ 1024 GB with 32 GB DIMMs

Storage Architecture

Hitachi Data Systems uses a building block approach with four active nodes when designing the storage system for each of the SAP HANA nodes using Virtual Storage Platform G800.

For the 2 TB HANA node configuration, build the storage system with drives and RAID groups as shown in Table 6, and with the disk storage as shown in Table 7.

TABLE 6. DRIVES AND RAID GROUPS FOR THE 2 TB HANA NODE CONFIGURATION

SAP HANA Node Building Blocks	Virtual Storage Platform G800	Operating System & Storage-based STONITH Drives, RAID Groups	Log Volume Drives, RAID Groups	Data Volume Drives, RAID Groups	HANA Shared Volume Drives, RAID Groups
4	1	8 × 600 GB 10k RPM SAS drives in 1 group configured as RAID-6 (6D+2P)	16 × 600 GB 10k SAS drives in 2 groups configured as RAID-6 (6D+2P)	48 × 1.2 TB 10k SAS drives in 3 groups configured as RAID-6 (14D+2P)	8 × 1.2 TB 10k RPM SAS drives in 1 group configured as RAID-6 (6D+2P)
8	1	8 × 600 GB 10k RPM SAS drives in 1 group configured as RAID-6 (6D+2P)	32 × 600 GB 10k SAS drives in 4 groups configured as RAID-6 (6D+2P)	96 × 1.2 TB 10k SAS drives in 6 groups configured as RAID-6 (14D+2P)	8 × 1.2 TB 10k RPM SAS drives in 1 group configured as RAID-6 (6D+2P)
12	1	8 × 600 GB 10k RPM SAS drives in 1 group configured as RAID-6 (6D+2P)	48 × 600 GB 10k SAS drives in 6 groups configured as RAID-6 (6D+2P)	144 × 1.2 TB 10k SAS drives in 9 groups configured as RAID-6 (14D+2P)	8 × 1.2 TB 10k RPM SAS drives in 1 group configured as RAID-6 (6D+2P)
16	1	8 × 600 GB 10k RPM SAS drives in 1 group configured as RAID-6 (6D+2P)	64 × 600 GB 10k SAS drives in 8 groups configured as RAID-6 (6D+2P)	192 × 1.2 TB 10k SAS drives in 12 groups configured as RAID-6 (14D+2P)	16 × 1.2 TB 10k RPM SAS drives in 2 groups configured as RAID-6 (6D+2P)

TABLE 7. DISK STORAGE FOR A 4-NODE STORAGE BUILDING BLOCK FOR THE 2 TB HANA NODE CONFIGURATION

SAP HANA Node Building Block	Operating System & Storage-based STONITH LUNs	Log LUNs	Data LUNs	HANA Shared LUNs
4	19 × 100 GB 1 × 50 MB	4 × 600 GB	12 × 2500 GB	4 × 1600 GB
8	19 × 100 GB 1 × 50 MB	8 × 600 GB	24 × 2500 GB	4 × 1600 GB
12	19 × 100 GB 1 × 50 MB	12 × 600 GB	36 × 2500 GB	4 × 1600 GB
16	19 × 100 GB 1 × 50 MB	16 × 600 GB	48 × 2500 GB	4 × 3200 GB

Build the Hitachi Virtual Storage Platform G800 with the components as shown in Table 8.

TABLE 8. HITACHI VIRTUAL STORAGE PLATFORM G800 COMPONENTS

Virtual Storage Platform G800 Component	4 Node Quantity	8 Node Quantity	12 Node Quantity	16 Node Quantity
Cache	128 GB	256 GB	512 GB	512 GB
MPU	2 pairs	2 pairs	2 pairs	2 pairs
DKB	2 pairs	2 pairs	2 pairs	2 pairs
CHB	3 pairs	4 pairs	5 pairs	6 pairs

Fibre Channel SAN Architecture

Connection Between Hitachi Compute Blade 2500 and VSP G800

The solution uses 38 × 8 Gb/sec Fibre Channel ports on Hitachi Virtual Storage Platform G800 directly attached to Hitachi Compute Blade 2500 via Fibre Channel PCIe adapters.

Each 520X B3 server blade has a pass-through mezzanine card on Mezzanine Slot 2 and Mezzanine Slot 4. These slots connect to the following:

- The 16 Gb/sec Hitachi FIVE-FX Fibre Channel PCI-Ex card
- The 10 GbE NIC card installed in the PCI-Ex slots through the backplane within the server chassis

Each SAP HANA node has two assigned dedicated 8 Gb/sec Fibre Channel ports on VSP G800. Using two Fibre Channel cables, connect the 16 Gb/sec Hitachi FIVE-FX Fibre Channel PCI-Ex cards on the chassis with the designated VSP G800 ports.

Set the port properties for the direct connection between Hitachi Compute Blade 2500 and Virtual Storage Platform G800 for all SAP HANA nodes as shown in Table 9.

TABLE 9. HITACHI VIRTUAL STORAGE PLATFORM G800 PORT PROPERTIES FOR HITACHI COMPUTE BLADE 2500 FIBRE CHANNEL CONNECTIONS

Property	Value
Port Security	**Disabled**
Port Speed	**Auto**
Fabric	**ON**
Connection Type	**P-to-P**

Hitachi FIVE-FX 16 Gb/sec HBA can emulate FC-SW virtually. Set the BIOS for Hitachi FIVE-FX HBA to **enable** for **Multiple Port ID**. Use the **Fabric** storage port setting to **ON** to set FC-SW virtual mode.

RAID Architecture

Several usage aspects divide the space provided by Virtual Storage Platform G800 in the Scale-out configuration, as follows:

- Operating System device provisioning for SAP HANA nodes
- STONITH device provisioning for storage-based fencing of the SuSE HAE cluster and the RHEL HA cluster
- HANA shared provisioning for storing SAP HANA binaries and cluster-wide configuration files
- Log device provisioning for SAP HANA database
- Data device provisioning for SAP HANA database

Each SAP HANA node has its own data volume and log volume. Only active SAP HANA nodes need data volumes and log volumes. Standby nodes do not require these volumes.

Figure 3 shows the RAID Group configuration for the Virtual Storage Platform G800 used in the 2 TB HANA node configuration with four active nodes and one standby node.

Provision the parity groups for SAP HANA node configuration as follows.

- **Operating System Volumes**
  - Each node has its own 100 GB LUN on Virtual Storage Platform G800 for the operating system volumes.
  - A single parity group configured as RAID-6 (6D + 2P) on 8 × 600 GB drives provisions the operating system LUN for SAP HANA nodes 1 to 19 on Virtual Storage Platform G800.
  - From this parity group, create 19 LDEVs, each with a capacity of 100 GB.
  - Map each LDEV exclusively to the corresponding SAP HANA node as follows: LUN number 00.
  - The installation of SUSE Linux Enterprise Server for SAP Applications or Red Hat Enterprise Linux resides on the boot LUN.
- **Shared Storage-based STONITH Volume**
  - A small LUN of 50 MB is created on Virtual Storage Platform 800 for STONITH of the SuSE HAE cluster and RHEL HA cluster. It provides a way to enable STONITH and fencing in clusters without external power switches, but with shared storage.
  - Map the volume to both ports of all SAP HANA nodes.
- **SAP HANA Shared Volumes**
  - For the 2 TB SAP HANA node configuration:
    - With 12 or fewer active nodes, create one parity group on Virtual Storage Platform G800, configured as RAID-6 (6D+2P) on 8 × 1.2 TB drives
    - From 13 up to 16 active nodes, create two parity groups on Virtual Storage Platform G800, configured as RAID-6 (6D+2P) on 16 × 1.2 TB drives
    - Each parity group has 4 × 1600 GB LDEVs.
    - Create a dynamic provisioning pool named HANA_SHARED_HDP. Assign all of the created LDEVs for the HANA Shared to this pool.
    - Create 4 virtual volumes in HANA_SHARED_HDP. Complete the LUN path assignment for these virtual volumes to each port for all SAP HANA nodes with the specified LUN ID.
- **SAP HANA Log Volumes**
  - For the SAP HANA log volumes, first create two parity groups configured as RAID-6 (6D+2P) on 16 × 600 GB drives.
  - In each of the two parity groups, create two 600 GB LDEVs.
  - Map each SAP HANA log volume to all SAP HANA nodes at each port with the LUN ID of the specified host.
- **SAP HANA Data Volumes**
  - For the 2 TB SAP HANA node configuration:
    - Create three parity groups configured as RAID-6 (14D+2P) on 48 × 1.2 TB drives.
    - Create four LDEVs with a capacity of 2500 GB per parity group.
    - Assign three LDEVs for use as data volumes to each SAP HANA node.

Figure 3

While mapping the LUN path assignment for each node, add the LUNs in the following order:

1. **Map the operating system volume for the specific SAP HANA node.**
2. **Map the STONITH volume.**
3. **Map the SAP HANA shared virtual volumes.**
4. **Map the log volume of each SAP HANA node.**
5. **Map the data volume of each SAP HANA node.**

The LUN assignment should be the same for all nodes except for the first LUN, which should be the operating system volume of that specific node.

Network Architecture

The scale-out configuration of the SAP HANA solution requires five separate networks as listed below:

- **SAP HANA Inter-Cluster Network** — This network provides the communication between the SAP HANA instances on the cluster.
- **SAP HANA Client Network** — This network is dedicated to handle the traffic between the SAP HANA database and its clients.
- **SuSE HAE/RHEL HA Cluster Network** — This network is used for communication between cluster nodes.
- **Management Network** — This network is used for management.

The SAP HANA inter-cluster network, client network, and SuSE HAE/RHEL HA cluster network are required to:

- Have no single point of failure (NSPOF)
- Provide at least 10 GbE equivalent throughput

To meet these requirements, the solution uses four 10GBASE-SR 2-port LAN adapters installed on the PCIe slots of each I/O board module per SAP HANA node. To provide a redundancy bond, two ports are provided from different PCIe network adapters at the operating system level using active-active network bonding mode and following the IEEE 802.3ad Link Aggregation standard for each of the three networks:

- SAP HANA inter-cluster network
- SuSE HAE or RHEL HA cluster network
- SAP HANA client network

Connections of each bond need to go to physically different VDX6740 switches. This way, if one switch fails there is still another route to the corresponding host. The two switches are connected together using ISL.

An MTU size of 9100 is used in accordance with Brocade best practices, and isolated using a VLAN of 100 for the HANA inter-cluster network and a VLAN of 150 is used for the SuSE HAE cluster network.

Switches

Two sets of switches are required for this solution.

- **Brocade VDX6740 Top Of the Rack (TOR) Switches**

  For the SuSE HAE cluster network and the SAP HANA inter-cluster network within the scale-out configuration of the SAP HANA installation, two ISL-paired Brocade VDX 6740 switches are used. In this solution, refer to the switches as VDX 6740-A and VDX 6740-B.

  For the client network, two ISL-paired Brocade VDX 6740 switches provide connectivity. In this solution, refer to the switches as VDX 6740-C and VDX 6740-D.

  Connect each of these two switches together using ISL. This lets both switches act together as one single logical switch. If one switch fails, there still is a path to the hosts.

- **Brocade ICX6430-48 port Management Switch**

  This 1 GbE switch is used to connect the management ports of the hardware and the management server.

SAP HANA Inter-Cluster Network

Figure 4 shows the SAP HANA inter-cluster network connection for nodes on a single chassis for the 2 TB SAP HANA node configuration.

Figure 4

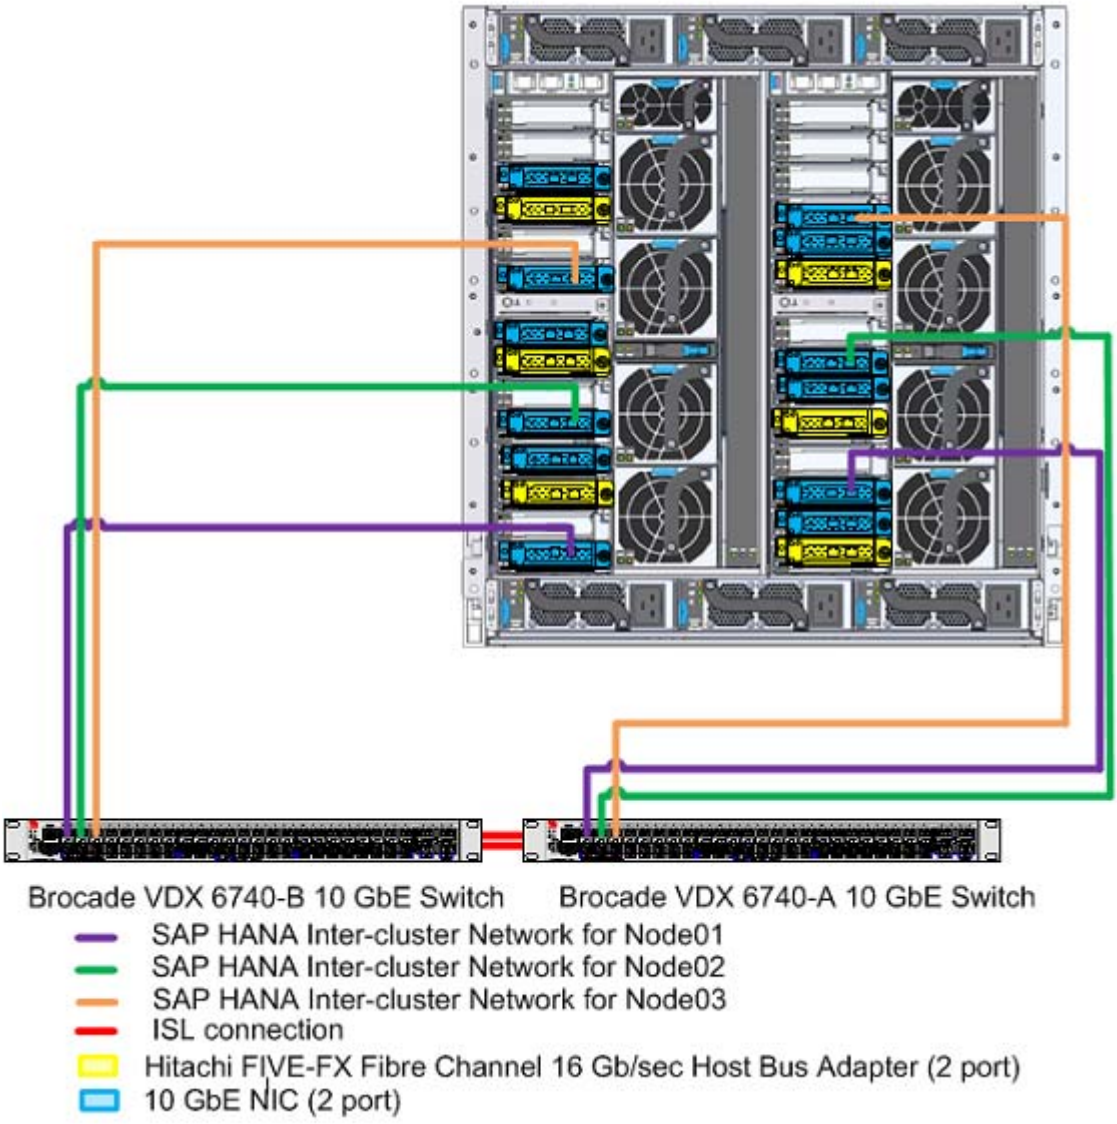

Configure the SAP HANA inter-cluster network using operating system level bonding on every node.

Table 10 lists the 10 GbE connection mappings for the three nodes on a single chassis for the 2 TB SAP HANA node configuration, using Chassis 1 as an example.

TABLE 10. SAP HANA INTER-CLUSTER PORT MAPPING FOR THE 2 TB HANA NODE CONFIGURATION

Chassis, PCIe Slot, Port Location	VDX 6740-48 Switch and Port Number	Bond
Chassis 1, IOBD 03A, Port 0	VDX6740-48A, Port #1	Bond 0 of Node 1
Chassis 1, IOBD 02A, Port 0	VDX6740-48B, Port #1	Bond 0 of Node 1
Chassis 1, IOBD 07A, Port 0	VDX6740-48A, Port #2	Bond 0 of Node 2
Chassis 1, IOBD 06A, Port 0	VDX6740-48B, Port #2	Bond 0 of Node 2
Chassis 1, IOBD 11A, Port 0	VDX6740-48A, Port #3	Bond 0 of Node 3
Chassis 1, IOBD 10A, Port 0	VDX6740-48B, Port #3	Bond 0 of Node 3

SuSE HAE or RHEL HA Cluster Network

Figure 5 shows the SuSE HAE or RHEL HA Cluster network connection for three nodes on a single chassis for the 2 TB SAP HANA node configuration.

Figure 5

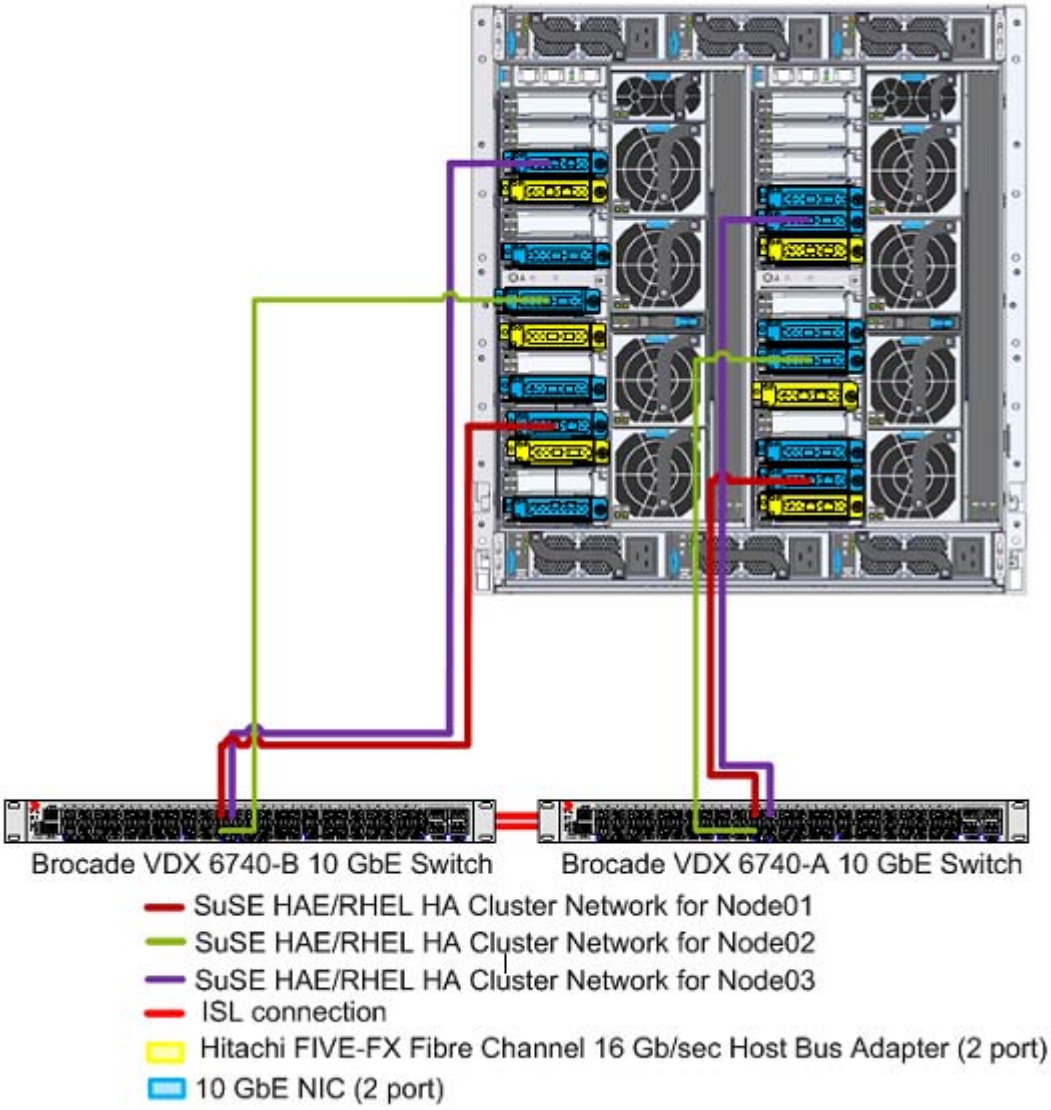

Configure the SuSE HAE or RHEL HA cluster network using operating system level bonding on every node.

Table 11 lists the port mappings of the 10 GbE connections for the three nodes on a single chassis for the 2 TB SAP HANA node configuration, using Chassis 1 as an example.

TABLE 11. SUSE HAE/RHEL HA CLUSTER NETWORK PORT MAPPINGS FOR THE 2 TB HANA NODE CONFIGURATION

Chassis, PCIe Slot, Port Location	VDX 6740-48 Switch and Port Number	Bond
Chassis 1, IOBD 01B, Port 1	VDX6740-48A, Port #21	Bond 1 of Node 1
Chassis 1, IOBD 04B, Port 1	VDX6740-48B, Port #21	Bond 1 of Node 1
Chassis 1, IOBD 05B, Port 1	VDX6740-48A, Port #22	Bond 1 of Node 2
Chassis 1, IOBD 08B, Port 1	VDX6740-48B, Port #22	Bond 1 of Node 2
Chassis 1, IOBD 09B, Port 1	VDX6740-48A, Port #23	Bond 1 of Node 3
Chassis 1, IOBD 12B, Port 1	VDX6740-48B, Port #23	Bond 1 of Node 3

SAP HANA Client Network Connection

Figure 6 shows the SAP HANA client network connection for three nodes on a single chassis for the 2 TB SAP HANA node configuration.

Figure 6

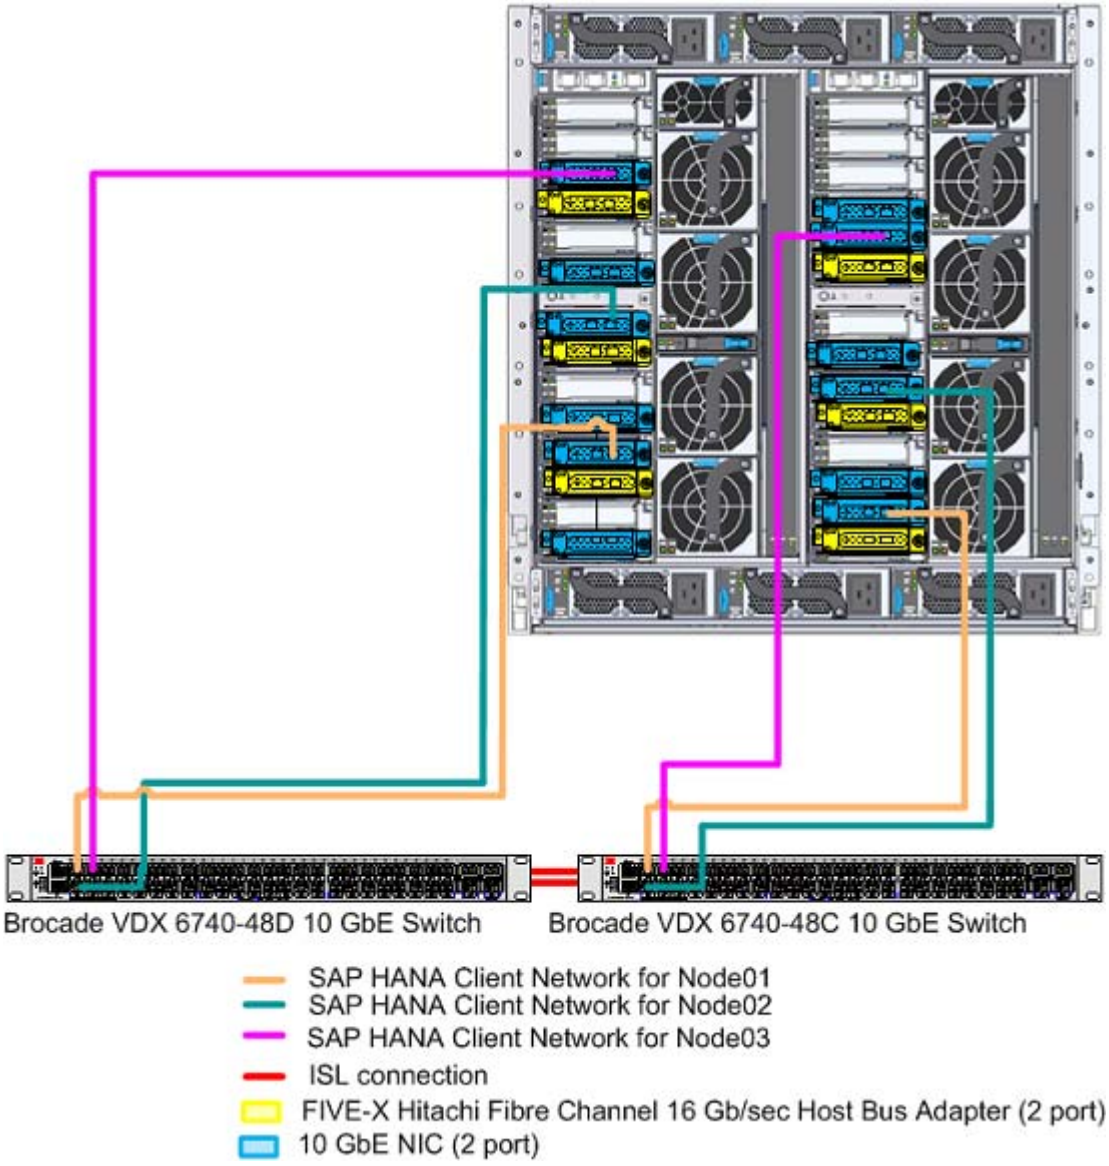

Configure the SAP HANA client network using operating system level bonding on every node.

Table 12 lists the port mappings of the 10 GbE connections for the three nodes on a single chassis for the 2 TB SAP HANA node configuration, using Chassis 1 as an example.

TABLE 12. SAP HANA CLIENT NETWORK PORT MAPPINGS FOR THE 2 TB HANA NODE CONFIGURATION

Chassis, PCIe Slot, Port Location	VDX 6740-48 Switch and Port Number	Bond
Chassis 1, IOBD 01B, Port 0	VDX6740-48C, port #1	Bond 2 of Node 1
Chassis 1, IOBD 04B, Port 0	VDX6740-48D, port #1	Bond 2 of Node 1
Chassis 1, IOBD 05B, Port 0	VDX6740-48C, port #2	Bond 2 of Node 2
Chassis 1, IOBD 08B, Port 0	VDX6740-48D, port #2	Bond 2 of Node 2
Chassis 1, IOBD 09B, Port 0	VDX6740-48C, port #3	Bond 2 of Node 3
Chassis 1, IOBD 12B, Port 0	VDX6740-48D, port #3	Bond 2 of Node 3

Management Network

Management network 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.

Management Server

This solution uses a node on the rack optimized server for solutions, 2U four node 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 B3 server blades
- Hitachi Virtual Storage Platform
- NTP configuration
- Brocade Switches
- SAP HANA node(s)

Figure 7 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 Brocade ICX 6430. It provides the 1 GbE network to the management server.
- **Slot 01 Port 1** — Connected to the Brocade ICX 6430 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

Connect the 10 GbE network ports to the two Brocade switches, VDX6740-A and VDX6740-B, to provide management access to the SAP HANA nodes from the management server on the rack optimized server for solutions, 2U four nodes.

Figure 7

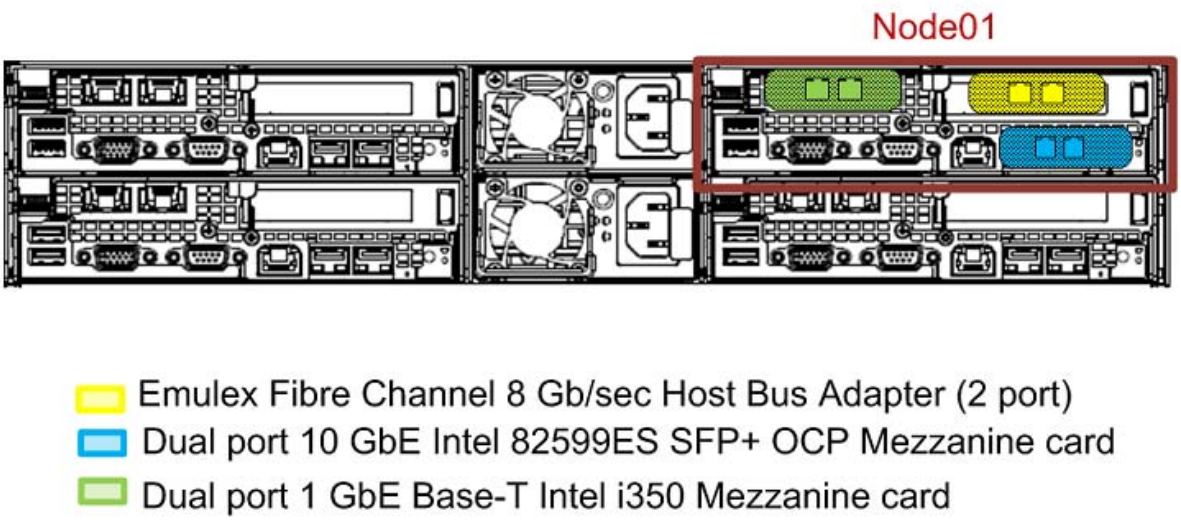

The following software is installed on the management server:

- Hitachi Command Suite
- PuTTY
- Teraterm
- JRE version jre-7u51-windows-i586 (no 64 bit)
- Adobe Flash Player
- WinSCP
- SAP HANA Studio

SAP HANA Configuration

The following sections describe configurations that are true for the scale-out solution on Hitachi Unified Compute Platform for the SAP HANA Platform.

SAN Operating System Configuration

A scale-out configuration for SAP HANA requires SAN boot.

The operating system LUN is the primary boot device for each scale-out node using the 16 Gb/sec 2-port Fibre Channel PCIe cards. The operating system LUN holds partitions for the following:

- Operating system
- /usr/sap/ directory
- Linux swap space

Activate Device-Mapper Multipath

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

Using Device-mapper Multipath allows the configuration of multiple I/O paths between the server blades and Hitachi Virtual Storage Platform G800. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available unless all paths fail.

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

- Operating system volume LUN
- SAP HANA shared volume LUN
- SAP HANA data volume LUN
- SAP HANA log volume LUN

HANA Shared GFS2 Volume Configuration

This solution uses global file system 2 (GFS2) to store the cluster-wide SAP HANA binaries and configuration files of the in-memory database. The shared file system is called **/hana/shared/<SID>** and accessible by all clustered nodes. "<SID>" is the *system ID* for the SAP HANA production database instance.

In this solution, the 2 TB SAP HANA node configuration uses four LDEVs/virtual volumes. All of the virtual volumes are under one HDP pool.

Logical Volume Manager (LVM) creates a single striped volume on which the GFS2 volume is created for the SAP HANA shared file system. The Clustered Logical Volume Manager (CLVM) is to manage shared storage using LVM.

Capacity based on the number of SAP HANA nodes:

- For up to 12 active HANA nodes: 6.25 TB
- For up to 16 active HANA nodes: 12.5 TB

Note — It is recommended to use other file systems (like NFS) for SAP HANA system backup instead of this HANA shared GFS2 file system in order to achieve improved performance.

HANA Data Volume Configuration

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.

For each scale-out active HANA node, a single striped volume is created across the data LUNS. The volume is formatted with the XFS file system to store the SAP HANA data volume.

HANA Log Volumes Configuration

Logical Volume Manager creates a single striped volume on which the XFS file system is created to store the SAP HANA log volume.

For each scale-out active HANA node, a single striped volume is created across the log LUNS. The volume is formatted with the XFS file system.

SAP HANA Appliance Software Installation

After configuring the file system for the SAP HANA data and log volumes, install SAP HANA on the SAP HANA nodes.

Install the following SAP HANA software components on the SAP HANA node:

- SAP HANA database
- SAP HANA client
- SAP Host Agent

Note - SAP HANA System Replication (HSR) automatic failover using Linux cluster has not been tested by Hitachi Data Systems SAP Solution Engineering, and is not supported by Hitachi Data Systems. Check with Linux vendors to implement this solution.

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-578-01, July 2017.