TECHNICAL PAPER

HITACHI
Inspire the Next

SAP HANA Tailored Datacenter Integration with Hitachi Virtual Storage Platform Using 1 TB SAP HANA Nodes

Reference Architecture Guide

By Archana Kuppuswamy

February 11, 2014

SAP Global Partner

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.

Table of Contents

SAP HANA Tailored Datacenter Integration with Hitachi Virtual Storage Platform Using 1 TB SAP HANA Nodes

Reference Architecture Guide

This reference architecture guide describes the storage design of Hitachi Virtual Storage Platform for use in SAP HANA Tailored Datacenter Integration with 1 TB SAP HANA nodes.

SAP HANA Tailored Datacenter Integration deployments are customized solutions in which customers can choose any of the certified SAP HANA server vendors along with any certified SAP enterprise storage.

- SAP certifies the minimum enterprise storage layout only, as described in this document.
- Hitachi Data Systems and the server vendor define the final complete configuration for the customer solution.

Using the family of enterprise storage products from Hitachi Data System, including Hitachi Virtual Storage Platform, with SAP HANA, has the following benefits:

- Increased performance when loading data into SAP HANA
- Scalable deployments of SAP HANA
- Disaster recovery with minimal performance impact to the production instance

The SAP HANA Tailored Datacenter Integration solution released by SAP allows the following for you:

- Reduced hardware and operational costs
- Lowers risk
- Optimizes time-to-value of existing hardware
- Shortens implementation cycles

The SAP HANA Storage Connector API for Fiber Channel is the only supported API for use with this solution.

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

Build the storage system with RAID groups, disk storage, and Virtual Storage Platform components using a four node building approach, as listed in Table 1 on page 2, Table 2 on page 2, and Table 3 on page 3. Each column shows you the final number of components, but not the change between each building block.

Table 1. Drives and RAID Groups

SAP HANA Node Building Blocks	*Operating System Drives, RAID Groups*	*Log Volume Drives, RAID Groups*	*Data Volume Drives, RAID Groups*	*Hitachi NAS Platform Volume Drives, RAID Groups*
4	8 × 900 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)	32 × 600 GB 10k SAS drives in 2 groups configured as RAID-6 (14D+2P)	16 × 900 GB 10k RPM SAS drives in 2 groups configured as RAID-6 (6D+2P)
8	8 × 900 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)	64 × 600 GB 10k SAS drives in 4 groups configured as RAID-6 (14D+2P)	16 × 900 GB 10k RPM SAS drives in 2 groups configured as RAID-6 (6D+2P)
12	8 × 900 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)	96 × 600 GB 10k SAS drives in 6 groups configured as RAID-6 (14D+2P)	32 × 900 GB 10k RPM SAS drives in 4 groups configured as RAID-6 (6D+2P)
16	8 × 900 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)	128 × 600 GB 10k SAS drives in 8 groups configured as RAID-6 (14D+2P)	32 × 900 GB 10k RPM SAS drives in 4 groups configured as RAID-6 (6D+2P)

Table 2. Disk Storage for a 4 Node Storage Building Block

SAP HANA Node Building Block	*VSP Frame*	*HNAS Servers*	*Operating System LUNs*	*Log LUNs*	*Data LUNs*	*HNAS LUNs*
4	Single	2	16 × 100 GB	4 × 1500 GB	4 × 3600 GB	4 × 2400 GB
8	Single	2	16 × 100 GB	8 × 1500 GB	8 × 3600 GB	4 × 2400 GB
12	Dual	2	16 × 100 GB	12 × 1500 GB	12 × 3600 GB	8 × 2400 GB
16	Dual	2	16 × 100 GB	16 × 1500 GB	16 × 3600 GB	8 × 2400 GB

Table 3. Hitachi Virtual Storage Platform Family Components

Virtual Storage Platform Component	*4 Node Quantity*	*8 Node Quantity*	*12 Node Quantity*	*16 Node Quantity*
Cache	96 GB	128 GB	192 GB	256 GB
VSD (MP-PK)	1 Pair	1 Pair	2 Pairs	2 Pairs
Cache (CP-PK)	2 Pairs	2 Pairs	3 Pairs	4 Pairs
DKA (BED)	1 Pair	2 Pairs	3 Pairs	4 Pairs
CHA (FED)	3 Pair	4 Pairs	5 Pairs	6 Pairs

This reference architecture is for solution architects, SAP HANA administrators, storage administrators, and SAP HANA technical architects. It assumes you have familiarity with the following areas:

- Storage network based storage systems
- Network attached storage (NAS) systems
- General storage concepts
- Common IT storage concepts
- Linux file system
- Multipath configuration of Linux systems
- SAP storage connector API
- SAP HANA
- Management console using a Hitachi Compute Rack 210H server or similar hardware

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.

Solution Overview

This document provides an example configuration of the storage layout of an eight-node cluster using seven active nodes and one standby node, each containing 1 TB of main memory. Validation was within the Hitachi Data System lab environment.

. This configuration uses the following Hitachi storage components:

- **Hitachi Virtual Storage Platform** — The first 3-D scaling storage platform designed for all data types that flexibly adapts for performance, capacity, and multi-vendor storage
- **Hitachi NAS Platform 3080** — A network-attached storage solution used for file sharing, file server consolidation, data protection, and business-critical NAS workloads

Figure 1 on page 5 shows the server to storage configuration of this solution.

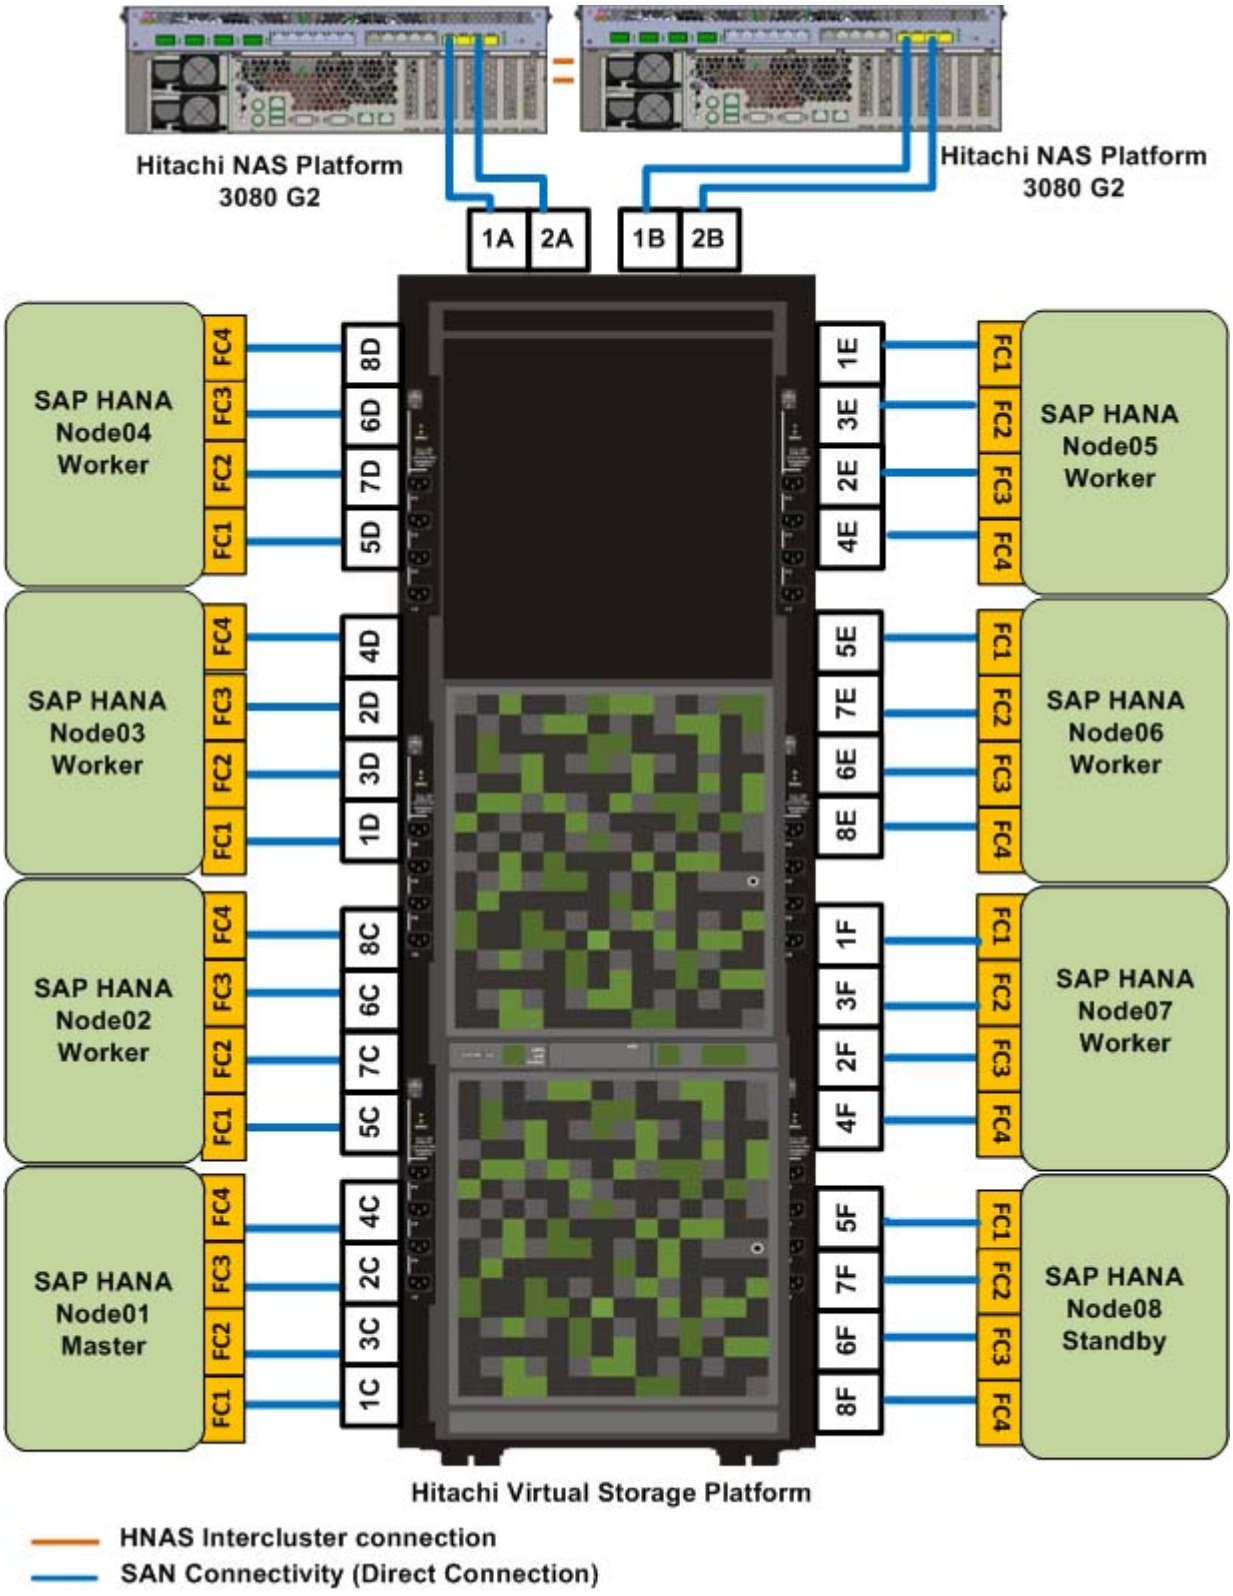
Hitachi NAS Platform
3080 G2
Hitachi NAS Platform
3080 G2
1A
2A
1B
2B
SAP HANA
Node04
Worker
SAP HANA
Node05
Worker
SAP HANA
Node03
Worker
SAP HANA
Node06
Worker
SAP HANA
Node02
Worker
SAP HANA
Node07
Worker
SAP HANA
Node01
Master
SAP HANA
Node08
Standby
FC1
FC2
FC3
FC4
8D
6D
7D
5D
1E
3E
2E
4E
4D
2D
3D
1D
5E
7E
6E
8E
8C
6C
7C
5C
1F
3F
2F
4F
4C
2C
3C
1C
5F
7F
6F
8F
Hitachi Virtual Storage Platform
HNAS Intercluster connection
SAN Connectivity (Direct Connection)

Figure 1

Key Solution Elements

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

Hardware Elements

Table 4 describes the hardware used to deploy the seven active node and one standby node configuration.

The hardware listed below is the recommended configuration for SAP HANA Tailored Datacenter Integration deployments with Hitachi storage. Additional hardware, such as storage area network and 10 GbE Network switches, may be required depending on which server vendor is picked for SAP HANA Tailored Datacenter Integration.

Table 4. Hardware Elements

Hardware	*Quantity*	*Configuration*	*Role*
Hitachi NAS Platform 3080	2	For every NAS Platform server ■ 2 cluster ports ■ 2 × 10 Gb ports ■ 2 Fibre Channel ports ■ 2 Ethernet ports	Provide NFS shared file system for SAP HANA binaries, cluster-wide configuration files.
SMU	1	■ 2 Intel Core 2 Duo E7500 model processor, 2.93 GHz CPU, 4 GB RAM	NAS Platform cluster management
Hitachi Virtual Storage Platform	1	■ Single frame	Block storage for SAP HANA nodes and NAS Platform
Hitachi Compute Rack 210H server ■ Optional configuration for management and disaster recovery ■ Can be replaced with other hardware vendor's servers	1	■ Intel Xeon E5-2620 model processor, 2.0 GHz CPU, 32 GB RAM ■ 2 × 300 GB 10k RPM SAS drives ■ 8 Gb/sec dual port Emulex Fibre Channel HBA	Management server
Rack servers or server blade chassis running SAP HANA	8	■ Rack servers or server blade chassis certified for SAP HANA with 1 TB SAP HANA nodes. A list of certified configurations can be found on https://service.sap.com/pam	Servers for SAP HANA with 1 TB of main memory

Hitachi NAS Platform 3080

Hitachi NAS Platform is an advanced and integrated network attached storage (NAS) solution. It provides a powerful tool for file sharing, file server consolidation, data protection, and business-critical NAS workloads.

This solution uses Hitachi NAS Platform 3080 for file system sharing of the global binary and configuration files of HANA. Optionally additional storage can be added and presented to the Hitachi NAS Platform for use of SAP HANA backups. There are two Hitachi NAS Platform 3080 server nodes that are clustered together.

The system management unit (SMU) provides front-end server administration and monitoring tools for NAS Platform. It supports clustering and acts as a quorum device in a cluster.

Hitachi Virtual Storage Platform

Hitachi Virtual Storage Platform is the first 3-D scaling storage platform designed for all data types. Its storage architecture flexibly adapts for performance, capacity, and multi-vendor storage. Combined with the unique Hitachi Command Suite management platform, it transforms the data center.

- **Scale Up** — Meet increasing demands by dynamically adding processors, connectivity, and capacity in a single unit. Provide the highest performance for both open and mainframe environments.
- **Scale Out** — Meet multiple demands by dynamically combining multiple units into a single logical system with shared resources. Support increased demand in virtualized server environments. Ensure safe multi-tenancy and quality of service through partitioning of cache and ports.
- **Scale Deep** — Extend storage value by virtualizing new and existing external storage systems dynamically. Extend the advanced functions of Hitachi Virtual Storage Platform to multivendor storage. Offload less demanding data to external tiers to save costs and to optimize the availability of tier-one resources.

The operating system LUNs, data LUNs, log LUNs, and LUNs for the Hitachi NAS Platform cluster reside on this storage device.

This solution uses a one frame or a dual frame Virtual Storage Platform, as follows:

- One single-frame Virtual Storage Platform for up to eight SAP HANA nodes
- One dual-frame Virtual Storage Platform for up to 16 SAP HANA nodes

Server for SAP HANA

The server for SAP HANA refers to the same bill of material as the certified SAP HANA appliance from any certified SAP HANA hardware vendor but without the storage.

Management Server

The Hitachi Compute Rack 210H server acts as the management server in this solution. This is an optional configuration for management and for a disaster recovery solution.

This hardware can be replaced with a server from another hardware vendor. Use this server to manage all the other components of this solution.

This server runs the following:

- Hitachi Command Suite
- Hi-Track remote monitoring system
- SAP HANA Studio
  - To manage the SAP HANA scale-out instance.
- Adobe Flash
  - To log on to Hitachi Virtual Storage Platform using Hitachi Command Suite

Software Elements

Table 5 describes the software products used to deploy the seven active node and one standby node configuration.

Table 5. Software Elements

Software	*Version*
Hitachi Virtual Storage Platform	70-06-02-00
Hitachi Command Suite	7.6.0-03
Hitachi Storage Navigator Modular 2	Microcode dependent
Hitachi NAS Platform firmware	11.2.3319.16
SMU software	11.2.3319.04
SMU operating system	CentOS-6.2
SAP HANA platform	1.0
SUSE Linux Enterprise Server for SAP Applications	11 SP2
Microsoft® Windows Server® 2008 R2	Standard Edition ▪ For the Hitachi Compute Rack 210H server

Solution Design

This is the detailed design for the scale out of a SAP HANA Tailored Datacenter Integration solution where Hitachi Virtual Storage is the preferred storage. The solution contains the following design elements:

Fibre Channel SAN Architecture

Each SAP HANA node needs to support four 8 Gb/sec Fibre Channel ports.

For the eight node configuration, with seven active nodes and one standby node, the Fibre Channel SAN architecture has 32 Fibre Channel cables directly connecting each Fibre Channel port to the designated target port on Hitachi Virtual Storage Platform.

This direct-attached storage configuration is the preferred SAN architecture. It has been validated for use with SAP HANA Fiber Channel Storage Connector (fcClient) with this solution. Refer to SAP Storage Connector API Fibre Channel Client for more details.

The use of a Fiber Channel switch is allowed, but there is no necessity to use it with this configuration. Consider the best practices of the SAN switch provider when designing or implementing your Fibre Channel zones.

Table 6 shows the storage port mapping.

Table 6. Storage Port Mapping

SAP HANA Node, Slot, Port	*Virtual Storage Platform Ports*
Node 1, FC1	1C
Node 1, FC2	3C
Node 1, FC3	2C
Node 1, FC4	4C
Node 2, FC1	5C
Node 2, FC2	7C
Node 2, FC3	6C
Node 2, FC4	8C
Node 3, FC1	1D

Table 6. Storage Port Mapping (Continued)

SAP HANA Node, Slot, Port	*Virtual Storage Platform Ports*
Node 3, FC2	3D
Node 3, FC3	2D
Node 3, FC4	4D
Node 4, FC1	5D
Node 4, FC2	7D
Node 4, FC3	6D
Node 4, FC4	8D
Node 5, FC1	1E
Node 5, FC2	3E
Node 5, FC3	2E
Node 5, FC4	4E
Node 6, FC1	5E
Node 6, FC2	7E
Node 6, FC3	6E
Node 6, FC4	8E
Node 7, FC1	1F
Node 7, FC2	3F
Node 7, FC3	2F
Node 7, FC4	4F
Node 8, FC1	5F
Node 8, FC2	7F
Node 8, FC3	6F
Node 8, FC4	8F
NAS Platform Server 1, FC Port 1	1A
NAS Platform Server 1, FC Port 3	2A
NAS Platform Server 2, FC Port 1	1B
NAS Platform Server 2, FC Port 3	2B

The port properties for the direct connection between the SAP HANA servers and Hitachi Virtual Storage Platform are as shown in Table 7

Table 7. Port Properties for Direct Attached Storage

Property	*Value*
Port Attribute	Target
Port Security	Disable
Port Speed	Auto (8 Gbps)
Fabric	Off
Connection Type	FC-AL

The port properties for the SAN with Fibre Channel switches between the SAP HANA servers and Hitachi Virtual Storage Platform are as shown in Table 8.

Table 8. Hitachi Virtual Storage Platform Port Properties for SAN with Fibre Channel Switches

Property	*Value*
Port Attribute	Target
Port Security	Disable
Port Speed	Auto (8 Gbps)
Fabric	On
Connection Type	P-to-P

Storage Architecture

The central storage system for the whole SAP HANA scale-out cluster is a Hitachi Virtual Storage Platform storage system. Several usage aspects divide the space provided by Virtual Storage Platform, namely:

- Boot LUN provisioning for SAP HANA nodes
- Log device provisioning for SAP HANA database
- Data device provisioning for SAP HANA database
- Block storage provisioning for NAS Platform shared file system

Figure 2 on page 12 shows the RAID Group configuration and components need for the Hitachi Virtual Storage Platform architecture used in the SAP HANA Tailor Datacenter Integration with seven active nodes and one standby node.

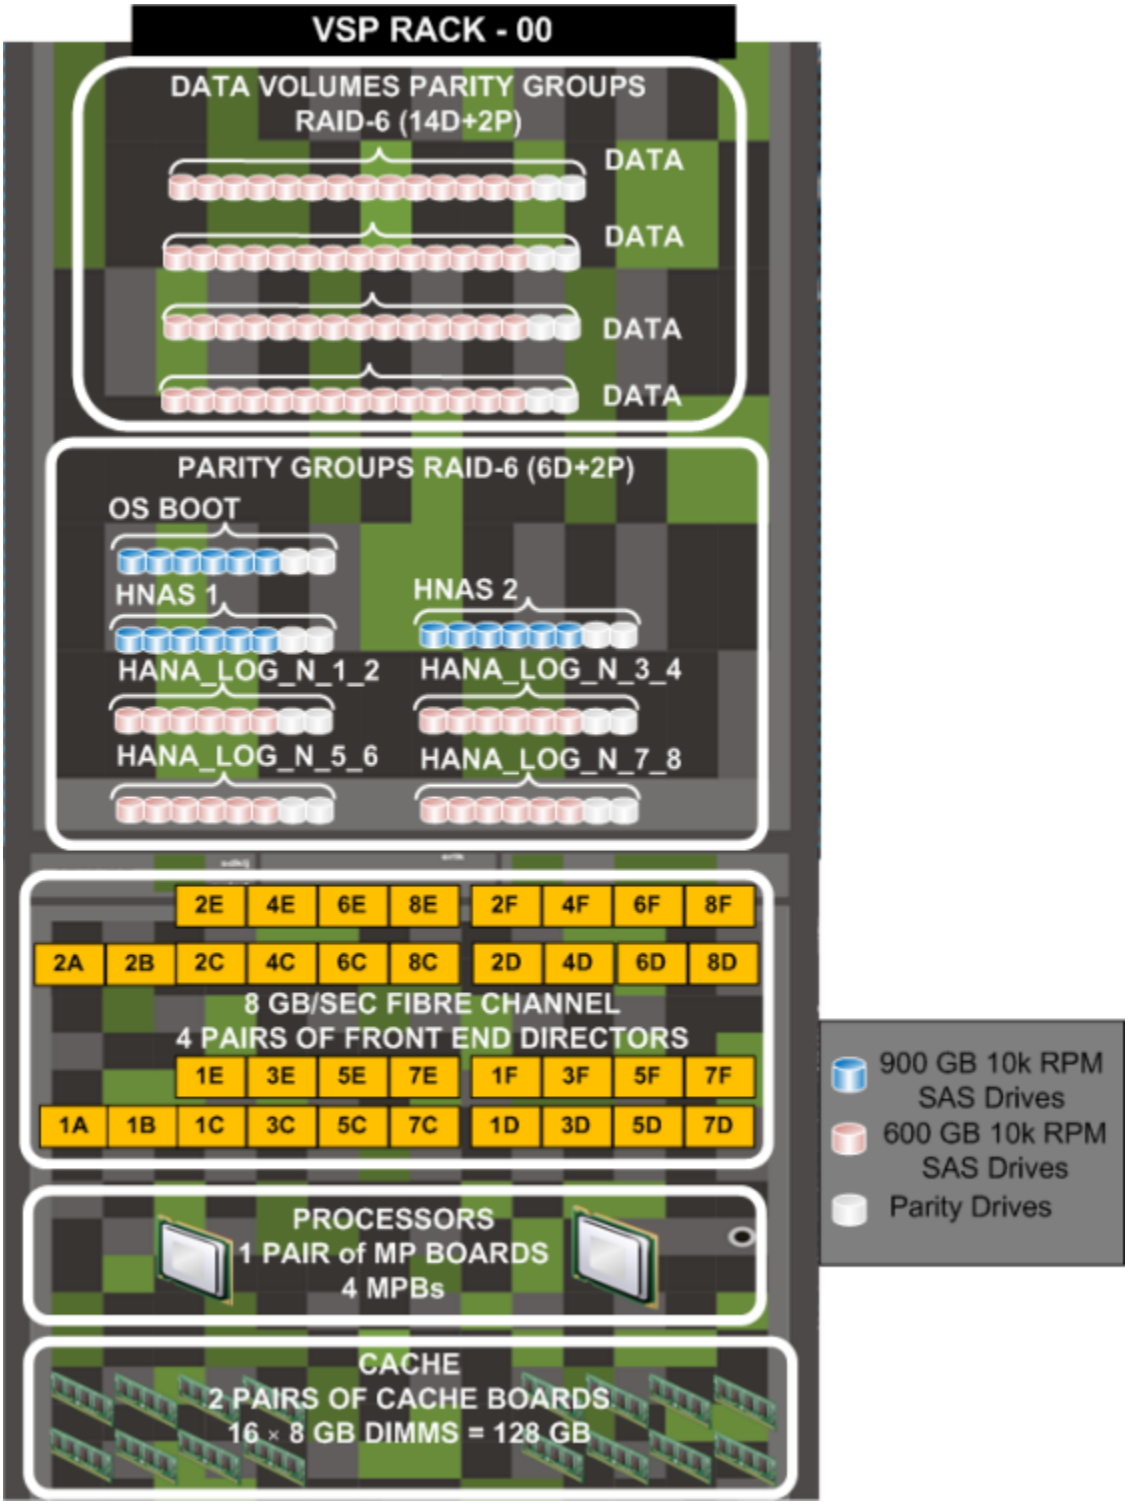
VSP RACK - 00
DATA VOLUMES PARITY GROUPS
RAID-6 (14D+2P)
DATA
DATA
DATA
DATA
PARITY GROUPS RAID-6 (6D+2P)
OS BOOT
HNAS 1
HNAS 2
HANA_LOG_N_1_2
HANA_LOG_N_3_4
HANA_LOG_N_5_6
HANA_LOG_N_7_8
2E 4E 6E 8E 2F 4F 6F 8F
2A 2B 2C 4C 6C 8C 2D 4D 6D 8D
8 GB/SEC FIBRE CHANNEL
4 PAIRS OF FRONT END DIRECTORS
1E 3E 5E 7E 1F 3F 5F 7F
1A 1B 1C 3C 5C 7C 1D 3D 5D 7D
PROCESSORS
1 PAIR of MP BOARDS
4 MPBs
CACHE
2 PAIRS OF CACHE BOARDS
16 x 8 GB DIMMS = 128 GB
900 GB 10k RPM
SAS Drives
600 GB 10k RPM
SAS Drives
Parity Drives

Figure 2

Provision the parity groups in Figure 2 on page 12 as follows:

- **Operating System Boot LUN (OS BOOT)**

  A single parity group configured as RAID-6 (6D + 2P) on eight 900 GB drives provisions the operating system boot LUN.

  - From this parity group, create eight LDEVs, each with a capacity of 100 GB.
  - Map each LDEV exclusively to the corresponding SAP HANA node as follows: **LUN number 00**

- **SAP HANA Log Volumes (HANA_LOG)**

  For each of the SAP HANA log volumes, create four parity groups, each configured as RAID-6 (6D+2P) on 32 × 600 GB drives.

  - In each of the four parity groups, create two LDEVs at 1500 GB each.
  - Map each SAP HANA log volume to all SAP HANA nodes at each port with the specified host LUN ID.

- **Hitachi NAS Platform 3080 Block Storage (HNAS)**

  The block storage for the Hitachi NAS Platform consists of two parity groups, each configured as RAID-6 (6D+2P) on 16 × 900 GB drives.

  - In each of the two parity groups, create two LDEVs at 2400 GB each.

- **SAP HANA Data Volumes**

  For the SAP HANA data volumes, create four parity groups configured as RAID-6 (14D+2P) on 64 × 600 GB drives.

  - Create four LDEVs with a capacity of 1869.98 GB per each parity group.
  - Create four dynamic provisioning pools and assign the four LDEVs from a single parity group to each dynamic provisioning pool.
  - Create two virtual volumes, each with 3600 GB per pool. Map each SAP HANA data volume to all SAP HANA nodes.

Figure 3 shows the configuration of the dynamic provisioning pools for the data volumes.

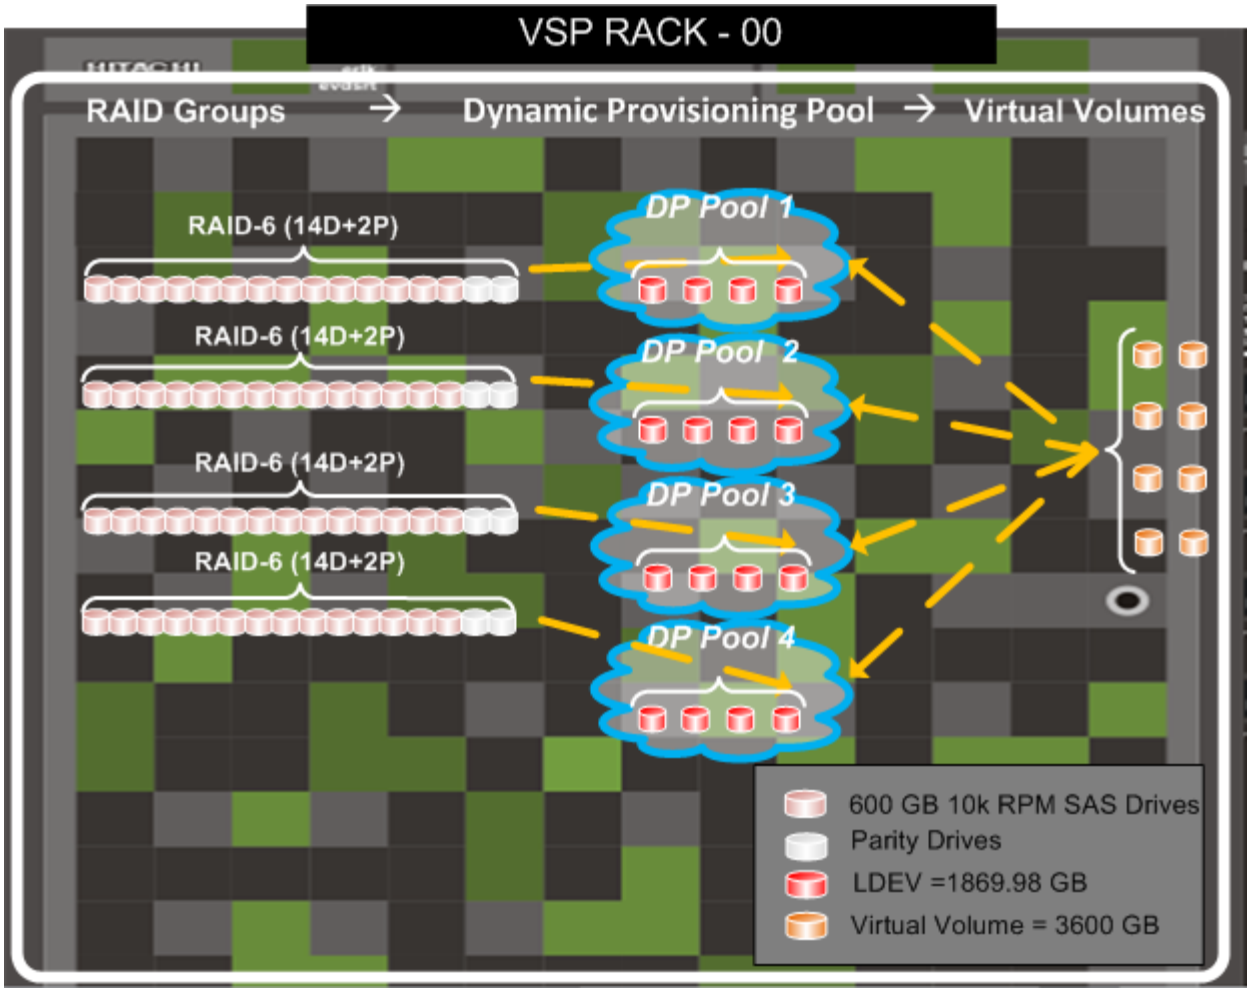

Figure 3

Table 9 shows the parity groups and LDEV assignment for boot volumes, Hitachi NAS Platform volumes, and SAP HANA log volumes.

Table 9. Parity Groups and LDEV Assignment for Operating System Boot, Hitachi NAS Platform, and SAP HANA Log Volumes

Parity Group	*Parity Group RAID Level and Disks*	*LDEV ID*	*LDEV Size*
1	RAID-6 (6D+2P) on 900 GB 10k RPM SAS drives	00:01:00	100.00 GB
		00:02:00	100.00 GB
		00:03:00	100.00 GB
		00:04:00	100.00 GB
		00:05:00	100.00 GB
		00:06:00	100.00 GB
		00:07:00	100.00 GB
		00:08:00	100.00 GB
2	RAID-6 (6D+2P) on 900 GB 10k RPM SAS drives	00:00:01	2400.00 GB
		00:00:02	2400.00 GB
3	RAID-6 (6D+2P) on 900 GB 10k RPM SAS drives	00:00:03	2400.00 GB
		00:00:04	2400.00 GB
4	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:01:01	1500.00 GB
		00:02:01	1500.00 GB
5	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:03:01	1500.00 GB
		00:04:01	1500.00 GB
6	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:05:01	1500.00 GB
		00:06:01	1500.00 GB
7	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:07:01	1500.00 GB
		00:08:01	1500.00 GB

Table 10 shows the parity groups and LDEV assignments for dynamic provisioning data volumes.

Table 10. Parity Groups and LDEV Assignments for Dynamically Provisioned Data Volumes

Parity Group	*Parity Group RAID Level and Disks*	*LDEV ID*	*LDEV Size*	*Dynamic Provisioning Pool ID*	*Dynamic Provisioning Pool Capacity*
8	RAID-6 (14D+2P) on 600 GB 10k RPM SAS drives	00:0A:01	1869.98 GB	1	7.30 TB
		00:0A:02	1869.98 GB		
		00:0A:03	1869.98 GB		
		00:0A:04	1869.98 GB		

Table 10. Parity Groups and LDEV Assignments for Dynamically Provisioned Data Volumes (Continued)

Parity Group	Parity Group RAID Level and Disks	LDEV ID	LDEV Size	Dynamic Provisioning Pool ID	Dynamic Provisioning Pool Capacity
9	RAID-6 (14D+2P) on 600 GB 10k RPM SAS drives	00:0A:05	1869.98 GB	2	7.30 TB
		00:0A:06	1869.98 GB		
		00:0A:07	1869.98 GB		
		00:0A:08	1869.98 GB		
10	RAID-6 (14D+2P) on 600 GB 10k RPM SAS drives	00:0A:09	1869.98 GB	3	7.30 TB
		00:0A:0A	1869.98 GB		
		00:0A:0C	1869.98 GB		
		00:0A:0D	1869.98 GB		
11	RAID-6 (14D+2P) on 600 GB 10k RPM SAS drives	00:0A:0E	1869.98 GB	4	7.30 TB
		00:0A:0F	1869.98 GB		
		00:0A:10	1869.98 GB		
		00:0A:11	1869.98 GB		

Table 11 shows the dynamic provisioning pool IDs and virtual volume LDEV IDs for SAP HANA data volumes.

Table 11. Dynamic Provisioning Pool IDs and Virtual Volume LDEV IDs

HDP ID	LDEV ID for SAP HANA Data Volumes	LDEV Size for SAP HANA Data Volumes
1	00:01:02	3600 GB
	00:02:02	3600 GB
2	00:03:02	3600 GB
	00:04:02	3600 GB
3	00:05:02	3600 GB
	00:06:02	3600 GB
4	00:07:02	3600 GB
	00:08:02	3600 GB

While eight SAP HANA data and log volumes are available to the SAP HANA appliance for scale-out, it only uses seven of those pairs for this reference architecture. The eighth node is a standby node.

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

1. Map the boot LUN for the specific SAP HANA node.
2. Map the log volume and data volume of each SAP HANA node.

Table 12 shows an example configuration of the LUN path assignment for Node01.

Table 12. LUN Path Assignment

LUN ID	*LDEV ID*	*LDEV Name*
0000	00:01:00	hananode01
0001	00:01:01	LOG_1
0002	00:01:02	DATA_1
0003	00:02:01	LOG_2
0004	00:02:02	DATA_2
0005	00:03:01	LOG_3
0006	00:03:02	DATA_3
0007	00:04:01	LOG_4
0008	00:04:02	DATA_4
0009	00:05:01	LOG_5
0010	00:05:02	DATA_5
0011	00:06:01	LOG_6
0012	00:06:02	DATA_6
0013	00:07:01	LOG_7
0014	00:07:02	DATA_7
0015	00:08:01	LOG_8
0016	00:08:02	DATA_8

Figure 4 shows the LUN assignment for each SAP HANA server node.

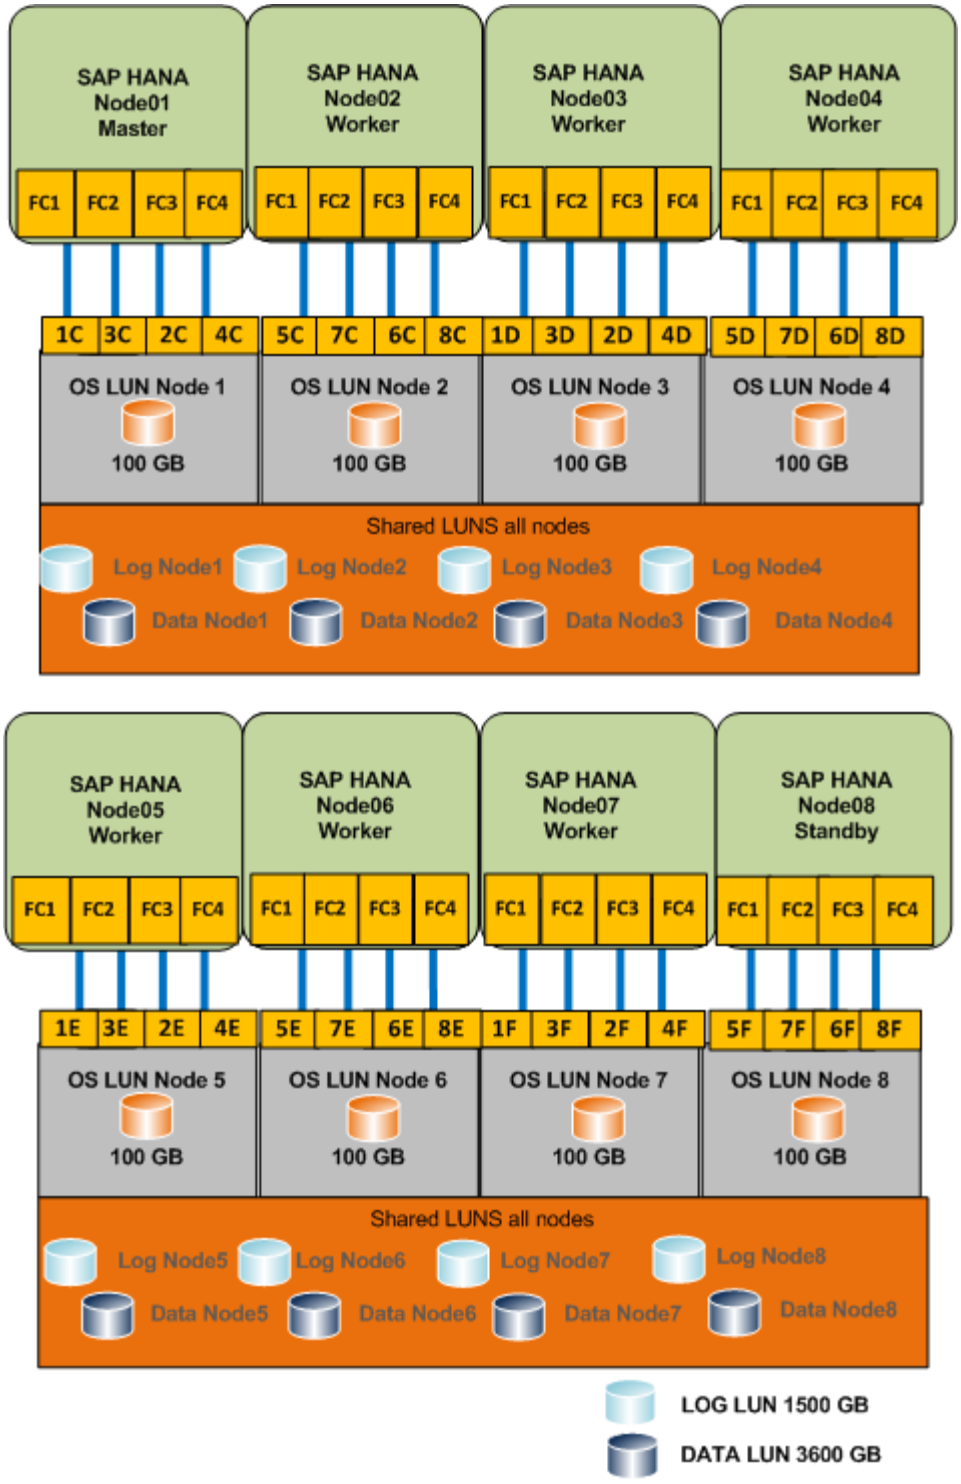

Figure 4

This configuration uses a minimum of the following as hot spare disks:

- Four 600 GB 10k RPM SAS drives
- Two 900 GB 10k RPM SAS drives

Hitachi NAS Platform 3080 Architecture

This architecture uses Hitachi NAS Platform 3080.

System Management Unit

Web Manager, the graphical user interface of the system management unit (SMU), provides front-end server administration and monitoring tools. It supports clustering and acts as a quorum device in a cluster. This solution uses an external SMU that manages two Hitachi NAS Platform servers.

Use one of the following browsers to run Web Manager:

- Microsoft Internet Explorer®: version 9.0, or later
- Mozilla Firefox: version 6.0, or later

G2 Servers

This solution uses two G2 servers in the cluster configuration. The two Hitachi NAS Platform servers are cluster interconnected with two 10 GbE links.

Private Management Network

Connect the private management interfaces of the G2 servers and the SMU to a dedicated one GbE management switch for private heart beat network. Devices connected to this private management switch are only accessible through the SMU.

Public Data Network

The public data network consists of the public Ethernet port of the SMU connected to a 1 GbE management switch.

Storage Subsystem

This solution uses Hitachi Virtual Storage Platform as the storage subsystem. Hitachi NAS Platform has direct attached Fibre Channel connections with the Hitachi Virtual Storage Platform target ports using two Fibre Channel cables.

Server Connections

Figure 5 shows the back of the G2 server.

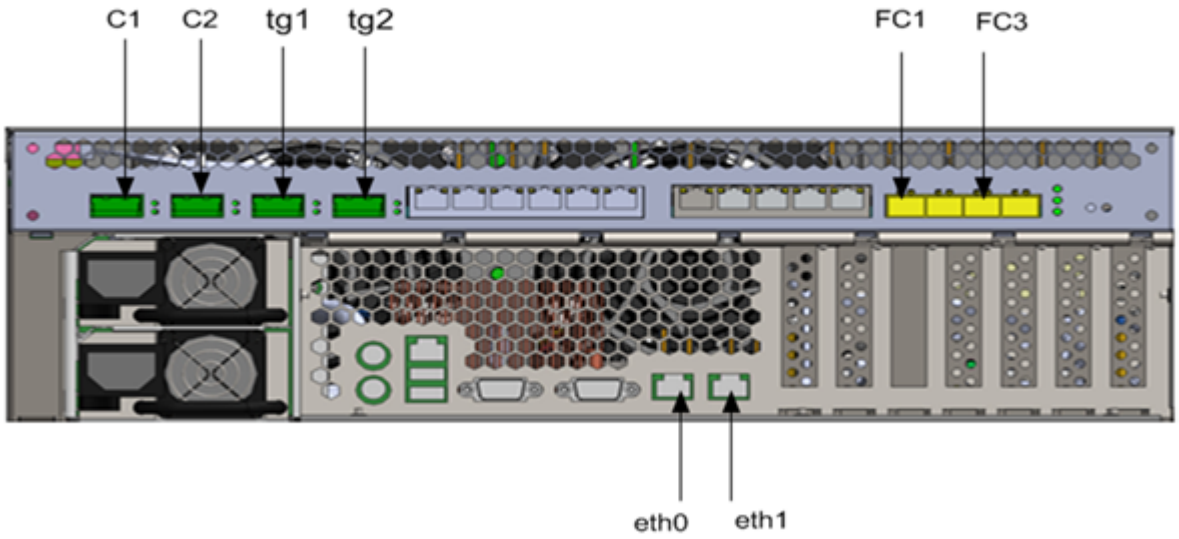

Figure 5

Port C1 and Port C2 are the NAS Platform 3080 cluster ports. To enable clustering, do the following:

- Connect Port C1 of first NAS Platform server to Port C1 of the second NAS Platform server.
- Connect Port C2 of first NAS Platform server to Port C2 of the second NAS Platform server.

Port tg1 and Port tg2 are 10 GbE ports. Link aggregate and connect these ports to the 10 GbE switch for the HANA NFS network connection between the SAP HANA nodes and Hitachi NAS platform. Hitachi recommends that the 10 GbE switch supports the following:

- Jumbo frames
- LACP
- The ability to segregate this network traffic from any other VLAN

Connect Fibre Channel Port FC1 and Port FC3 directly to the Hitachi Virtual Server Platform ports, as follows:

- 1A and 2A for the first NAS Platform server
- 1B and 2B for the second NAS Platform server

Connect Port eth1 of the NAS platform server to the dedicated 1 GbE management switch for a private heart beat network.

For the direct connection between NAS Platform and Virtual Storage Platform, set the port properties as shown in Table 13.

Table 13. Hitachi Virtual Storage Platform Port Properties

Property	*Value*
Port Attribute	Target
Port Security	Disable
Port Speed	Auto (4 Gbps)
Fabric	Off
Connection Type	FC-AL

The Hitachi Virtual Storage Platform port properties for the SAN with Fibre Channel switches between NAS Platform and Hitachi Virtual Storage Platform are as shown in Table 14.

Table 14. Port Properties for SAN with Fibre Channel Switches

Property	*Value*
Port Attribute	Target
Port Security	Disable
Port Speed	Auto (8 Gbps)
Fabric	On
Connection Type	P-to-P

Network File System Design for Shared Binaries

This solution requires a network file system to store cluster-wide HANA binaries and configuration files of the in-memory database. Host this shared file system called **/hana/shared/<SID>**, on Hitachi NAS Platform. Mount this file system on all SAP HANA nodes.

Table 15 shows the parity groups setup and the four LDEVs used for NAS Platform in this configuration.

Table 15. Parity Groups Setup

Parity Group	*Parity Group RAID Level and Disks*	*LDEV ID*	*LDEV Size*
2	RAID-6 (6D+2P) on 900 GB 10k RPM SAS drives	00:00:01	2400.00 GB
		00:00:02	2400.00 GB
3	RAID-6 (6D+2P) on 900 GB 10k RPM SAS drives	00:00:03	2400.00 GB
		00:00:04	2400.00 GB

This solution uses four LDEVs for SAP HANA shared binaries, as listed in Table 15 on page 21.

- Refer to each LDEV as a system drive.
- Create two system drive groups and assign the two LDEVs (system drives) from a single parity group to each system drive group.
- With these system drive groups, create a single storage pool called **HANABIN_PROD** with a capacity of 9.37 TB.

Configure two EVSs on the NAS Platform nodes as follows:

- EVS on NAS Platform node 1 as **HNASEVS1**
- EVS on NAS Platform node 2 as **HNASEVS2**

Create the shared file system hana_shared_<SID> using the storage pool **HANABIN_PROD** with the following:

- Capacity of 9.3TB
- Block size of 32 KB
- Auto expansion disabled

Mount and then export the file system. Mount the NFS export / **hana_shared_<SID>** on the file system path "**/hana/shared/<SID>**" on all eight SAP HANA nodes, where SID is the system ID for the SAP HANA production database instance.

Set the MTU size to 9000 on both NAS Platform nodes.

SAP Storage Connector API Fibre Channel Client

The SAP HANA Storage Connector API Fibre Channel client defines a set of interface functions called during the following:

- Normal SAP HANA cluster operation
- Failover handling

Storage Connector clients implement the functions defined in the Storage Connector API.

SAP HANA can make the needed storage partition for an SAP HANA node and ensure proper fencing in a failover case using these storage connector clients. SAP currently ships two implementations.

- fcClient
- iSCSIclient

This solution with scale-out uses the **fcClient** implementation. SAP supports this solution to enable the use of high-performance Fibre Channel devices for a scale-out installation. This solution does not support iSCSIclient.

The fcClient implementation uses standard Linux packages, such as multipath-tools and sg3_utils. Install and configure these packages.

The following is true for each data and log volume:

- It resides on a LUN of its own.
- It is identified by the name seen in /dev/mapper on the operating system.

The fcClient implementation is responsible for mounting the SAP HANA volumes. It also implements a proper fencing mechanism during a failover by means of SCSI-3 persistent reservations.

Configuration of the SAP Storage Connector API is contained within the SAP global.ini file within the **/hana/shared/<SID>** mount point. Refer to the sample global.ini file in “Appendix 1: global.ini” on page 26.

To find the wwid of the log and data volumes, do the following:

- Access the /dev/disk/by-id folder at the SUSE Linux operating system level.
- Verify the dm device number.
  - The last three digits of the dm device correspond to the actual LDEV ID on Hitachi Virtual Storage Platform.

SAP HANA Node Configuration

The scale-out of the SAP HANA solution consists of three types of nodes:

- **Master Node** — Initial node where the first partition of the SAP HANA database is installed
- **Worker Node** — Secondary nodes with their own database partitions.
- **Standby Node** — A node without a database partition
  - The standby nodes are a pool of computing resources that will be used in case of a failure of an active node (master or worker).

This solution has the following:

- One master node
- Six worker nodes
- One standby node

SAN Operating System Boot Configuration and SUSE Installation

This eight-node SAP HANA configuration uses SAN boot for each node.

Each node has its own 100 GB LUN on Hitachi Virtual Storage Platform for the operating system boot LUNs. Map the boot LUNs to each node according to Table 9 on page 15.

The installation of SUSE Linux Enterprise Server for SAP Applications resides on the boot LUN. Configure the SUSE OS to use multipathing for accessing the direct-attached storage devices. The multipath.conf file needs to be set with appropriate options as specified in Appendix 2: multipath.conf.

HANA Node Network Configuration

Each SAP HANA node has four different networks required for the following:

- **Management of the operating system** — one 1 GbE network

  This network is a 1 GbE operating system management network. This non-redundant network is for management only.

  This is not vital to the SAP HANA services.

- **SAP HANA inter-cluster network** — One fully redundant 10 GbE network

  This network is for node-to-node communication within the SAP HANA appliance. This is not meant for public access.

- **SAP HANA NFS network** —- One fully redundant 10 GbE network

  This network is for /hana/shared. Every node in the cluster must be able to access this shared resource for SAP HANA binaries. This is not meant for public access.

- **SAP HANA client network** —- One fully redundant 10 GbE network

  This is for connection between the SAP HANA database and its clients.

  The SAP HANA inter-cluster, NFS, and client networks are required as a must for the HANA appliance.

SAP HANA Data and Log Configuration

SAP HANA uses a plain file system to store the contents of the database. The database state is held in data files and log files that are stored in data and log volumes. The fcClient is used through the SAP Storage Connector API for this solution.

Each database partition has its own data and log volume that consist of a single LUN with a standard XFS file system. Each node has the LUN path for all of the other nodes' data and log volumes for high availability.

For achieving optimal performance, Hitachi Data Systems recommends that you use the following options to create the data and log file system with the mkfs command and the corresponding mount options.

- To create file system for data volume:

  ```
  mkfs.xfs -f -d sunit=2048,swidth=28672 /dev/mapper/…
  ```

- To create file system for log volume:

  ```
  mkfs.xfs -f -d sunit=2048,swidth=12288 /dev/mapper/…
  ```

Sudoers File Configuration

On the master and worker nodes, Hitachi Data Systems recommends that the /etc/suoders file have the following options specified:

> <SID>adm ALL=NOPASSWD: /sbin/multipath, /sbin/multipathd, /etc/init.d/multipathd, /usr/bin/sg_persist, /bin/mount, /bin/umount, /bin/kill, /usr/bin/lsof

SID stands for the HANA system ID.

Appendix 1: global.ini

The following is a sample global.ini file:

```
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT
[storage]
ha_provider = hdb_ha.fcClient
ha_provider_path = /usr/sap/HIT/HDB10/exe/python_support
partition_*_*__prtype = 5
partition_*_log__mountoptions = -o sunit=2048,swidth=12288,inode64,nobarrier,largeio,swalloc
partition_*_data__mountoptions = -o sunit=2048,swidth=28672,inode64,nobarrier,largeio,swalloc
partition_1_log__wwid = 360060e8006da98000000da9800000101
partition_1_data__wwid = 360060e8006da98000000da9800000102
partition_2_log__wwid = 360060e8006da98000000da9800000201
partition_2_data__wwid = 360060e8006da98000000da9800000202
partition_3_log__wwid = 360060e8006da98000000da9800000301
partition_3_data__wwid = 360060e8006da98000000da9800000302
partition_4_log__wwid = 360060e8006da98000000da9800000401
partition_4_data__wwid = 360060e8006da98000000da9800000402
partition_5_log__wwid = 360060e8006da98000000da9800000501
partition_5_data__wwid = 360060e8006da98000000da9800000502
partition_6_log__wwid = 360060e8006da98000000da9800000601
partition_6_data__wwid = 360060e8006da98000000da9800000602
partition_7_log__wwid = 360060e8006da98000000da9800000701
partition_7_data__wwid = 360060e8006da98000000da9800000702
```

Appendix 2: multipath.conf

The following is a sample multipath.conf file:

```
devices {
    device {
        vendor          HITACHI
        product         DF600F
    }
    device {
        vendor          HITACHI
        product         OPEN*
    }
}

defaults {
    user_friendly_names     no
    path_checker            directio
    path_grouping_policy    multibus
    path_selector           "queue-length 0"
    getuid_callout          "/lib/udev/scsi_id --whitelisted --device=/dev/%n"
    failback                immediate
    rr_weight               uniform
    rr_min_io               1000
    features                "1 queue_if_no_path"
    max_fds                 max
}

blacklist {
    devnode "^fio[a-z]"
}
multipaths {
```

```
    multipath {
        wwid                    360060e8006da98000000da9800000100
        path_selector           "round-robin 0"
    }
}
```

The *multipaths* section defines a different multipath setup for the boot LUN. Instead of path selector "queue-length 0," the path selector "round-robin 0" must be used.

You can identify the correct boot LUN with the last eight digits, which is the hexadecimal representation of the LDEV ID. This example shows a LUN ID 00000100 which is LDEV ID 00:01:00 on the Hitachi Virtual Storage Platform, and follows the definition pinned to Node 1.

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 Corporation

Corporate Headquarters
2845 Lafayette Street, Santa Clara, California 95050-2627 USA
www.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-274-00, February 2014