



# Disaster Recovery Using SAP HANA System Replication on Hitachi Unified Compute Platform for the SAP HANA Platform in a Scale-up Configuration Using Hitachi Compute Blade 2500 and Hitachi Virtual Storage Platform Family

## Reference Architecture Guide

By Milind Pathak

October 2015

## Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@hds.com](mailto:SolutionLab@hds.com). To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

---

# Contents

SAP HANA Disaster Recovery .....	1
SAP HANA System Replication .....	1
Hitachi Compute Blade 2500 .....	3
Hitachi Virtual Storage Platform Family .....	5
<b>Solution Overview .....</b>	<b>6</b>
<b>Key Solution Elements.....</b>	<b>8</b>
Hardware Elements.....	8
Software Elements.....	9
<b>Solution Design.....</b>	<b>10</b>
Hitachi Compute Blade 2500 Chassis Configuration .....	10
Network Design.....	11
Management Server.....	13
Planning and Prerequisites .....	15
Setting up SAP HANA System Replication .....	15
Takeover Process .....	19
Use Secondary Site for DEV/QA System.....	19
<b>Engineering Validation.....</b>	<b>20</b>

# Disaster Recovery Using SAP HANA System Replication on Hitachi Unified Compute Platform for the SAP HANA Platform in a Scale-up Configuration Using Hitachi Compute Blade 2500 and Hitachi Virtual Storage Platform Family

## Reference Architecture Guide

This reference architecture guide describes how to achieve disaster recovery for SAP HANA Haswell scale-up systems using SAP HANA System Replication.

### SAP HANA Disaster Recovery

SAP HANA database run mission-critical applications, and it is important that these systems remain available to users at all times. Thus, these systems must recover faster after a system component failure (High Availability) or after a disaster (Disaster Recovery). This should happen without any data loss (zero RPO) and in a very short or low recovery time objective (RTO). SAP and its partners offer the following disaster recovery mechanisms for SAP HANA. These solutions are based on completely redundant servers and/or storage.

- **Storage Replication:** Storage mirroring replicates data independently from the database software. Disks are mirrored without a control process from the SAP HANA system. SAP HANA hardware partners offer this solution. This solution needs additional servers and storage.
- **SAP HANA System Replication:** SAP HANA replicates all data to a secondary SAP HANA system constantly. Data can be constantly preloaded in the memory of the secondary system to minimize recovery (RTO). This solution needs additional servers and storage. Please refer to the [SAP HANA Disaster Recovery Support](#) document for more information.

This reference architecture guide explains how to achieve disaster recovery using SAP HANA system replication for SAP HANA scale-up systems installed on Hitachi Compute Blade 2500 with 520X B2 server blades, and Hitachi Virtual Storage Platform family storage array.

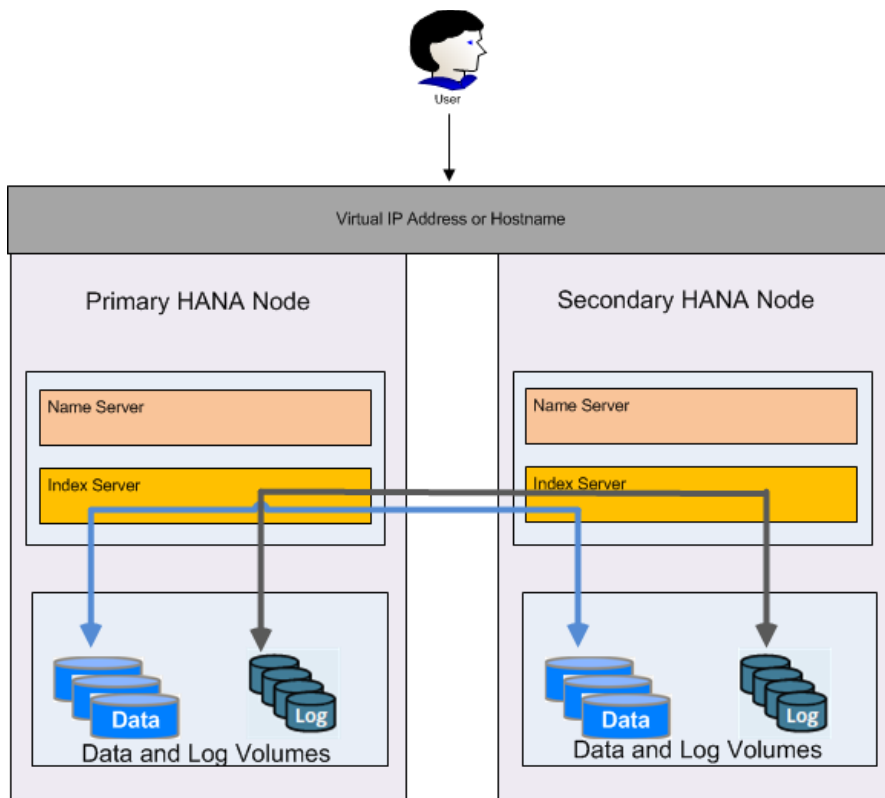
### SAP HANA System Replication

SAP HANA System Replication is implemented between two different SAP HANA systems with the same number of active nodes. After system replication is set up between the two SAP HANA systems, it replicates all of the data from the primary SAP HANA system to the secondary SAP HANA system (initial copy). After this, any logged changes in the primary system are also sent to the secondary system, but log entries are not replayed. Also, data snapshots are sent from primary to secondary at regular intervals. This means whenever the secondary has to take over, only log entries received after the last data snapshot need to be replayed. With the data snapshot, the primary system also sends information about the tables loaded in memory if the parameter 'preload\_column\_tables' is set to 'true'. If this parameter is also set to 'true' on the secondary system, these tables are preloaded in the memory of the secondary database. This reduces the RTO and makes SAP HANA system replication a faster disaster recovery solution in terms of recovery. The entire process of data replication occurs on the software level and is fully controlled by the SAP HANA database kernel.

The following replication modes are available for this procedure:

- **Synchronous on disk (mode=sync):** Transaction is committed after log entries are written on primary and secondary systems.
- **Synchronous in memory (mode=syncmem):** Transaction is committed after secondary system receives the logs but before they are written to disks.
- **Asynchronous (mode=async):** Transaction is committed after log entries are sent without any response from secondary.

Additionally, the synchronous replication mode (SYNC) can run with “full sync” enabled. In full sync operation, transaction processing on the primary site becomes blocked, when the secondary is not connected and newly created redo log buffers cannot be shipped to the secondary site. System replication offers synchronous and asynchronous modes to accommodate network latency. If the distance between your sites is less than 100 km you can use a synchronous replication mode: SYNC or SYNCMEM. For all data centers that are more than 100 km apart, the asynchronous replication mode ASYNC is recommended. The procedure described in this reference architecture is applicable to all three replication modes. System replication can be set up using SAP HANA Studio or command line. This reference architecture uses SAP HANA Studio for the setup and only explains the system replication process on SAP HANA scale-up systems. Refer to the [SAP HANA Administration Guide](#) to read more about system replication. Figure 1 shows an overview of SAP HANA system replication.



**Figure 1**

If the SAP HANA system fails at the primary site, the system administrator must perform manual takeover to the secondary site. Takeover can be performed using SAP HANA Studio or command line. Refer to the SAP [HANA system replication guide for details](#).

## 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.

One SAP HANA Scale-up configuration uses one, two, or four 520X B2 server blades in the Hitachi Compute Blade 2500 chassis for the different sized solutions. This Disaster Recovery reference architecture uses two SAP HANA scale-up nodes of the same size. Table 1 on page 4 lists the configurations for the 520X B2 server blades used in the various SAP HANA scale-up solutions.

**Table 1. 520X B2 Server Blade Configuration**

Feature		Configuration		
		Small (2-Socket)	Large (8-Socket) Medium (4-Socket)	
Processors		2 Intel Xeon E7-8880 processors	4 Intel Xeon E7-8880 processors	8 Intel Xeon E7-8880 processors
Processor SKU		Intel Xeon E7-8880 v3		
Processor frequency		2.3 GHz		
Processor cores		18 cores		
Number of blades		1	2	4
Number of DIMMs per blade	128 GB	8 × 16 GB DIMMs	8 × 8 GB DIMMs	N/A
	256 GB	16 × 16 GB DIMMs	8 × 8 GB DIMMs	8 × 8 GB DIMMs
	384 GB	24 × 16 GB DIMMs	24 × 8 GB DIMMs	N/A
	512 GB	32 × 16 GB DIMMs	16 × 16 GB DIMMs	16 × 8 GB DIMMs
	768 GB	48 × 16 GB DIMMs	24 × 16 GB DIMMs	24 × 8 GB DIMMs
	1024 GB	32 × 32 GB DIMMs	32 × 16 GB DIMMs	16 × 16 GB DIMMs
	1536 GB	48 × 32 GB DIMMs	48 × 16 GB DIMMs	24 × 16 GB DIMMs
	2048 GB	N/A	32 × 32 GB DIMMs	32 × 16 GB DIMMs
	3072 GB	N/A	48 × 32 GB DIMMs	48 × 16 GB DIMMs
	4096 GB	N/A	N/A	32 × 32 GB DIMMs
6144 GB	N/A	N/A	48 × 32 GB DIMMs	
Network Ports		2 × 2-port 10GBASE-SR LAN PCIe adapter on two I/O board modules: <ul style="list-style-type: none"> <li>■ IOBD 01B</li> <li>■ IOBD 02B</li> </ul>		
Fibre Channel Ports		2 × Hitachi 16 Gb/sec 2-port Fibre Channel adapters on two I/O board modules: <ul style="list-style-type: none"> <li>■ IOBD 01A</li> <li>■ IOBD 02A</li> </ul>		
Other Interfaces		1 USB 3.0 port KVM connector (VGA, COM, USB2.0 port)		

## Hitachi Virtual Storage Platform Family

Hitachi Virtual Storage Platform family delivers enterprise storage virtualization in a unified platform for midmarket to global enterprise organizations that need to manage information more efficiently. Existing storage from multiple vendors can now be centrally unified in a shared pool of data. A highly efficient architecture allows organizations to satisfy growth requirements and simplify operations to reduce the total cost of storage ownership. Including industry-leading global storage virtualization, Virtual Storage Platform family extends ease of migration, universal data replication and active mirroring to all storage assets. Provision and manage virtual storage machines across metro distances with active-active remote data center replication support. Combined with remote data center replication, this mirroring is an ideal solution to meet zero recovery point and time objectives for critical applications. With a central point of control, Virtual Storage Platform (VSP) family unifies block and file access, enabling organizations to consolidate workloads to further simplify management and defer the cost of additional storage acquisition. Using a single interface for managing both block and file storage streamlines administration.

VSP family systems are built on legendary Hitachi reliability, offering complete system redundancy, hot-swappable parts, outstanding data protection and non-disruptive updates to keep storage operations up and running at optimal performance. Additional data recovery and protection tools allow for application-aware recovery, simpler backup, restore, failover and consistency across copies, reducing business risk, downtime and migration concerns. VSP family complements virtualized server environments with its ability to consolidate multiple file and block workloads in a single system. Additional integration offloads storage-intensive processing from the server hosts to increase virtual machine density, improve performance and reduce workload contention. And it extends those benefits to legacy-attached storage via external storage virtualization. Five models in the VSP family, based on Hitachi Storage Virtualization Operating System (SVOS), provide a uniquely scalable, software-defined storage foundation. Powered with Hitachi global storage virtualization, new software capabilities unlock IT agility and enable the lowest storage total cost of ownership.

Validation of VSP G200, VSP G400 and VSP G600 has been done for SAP HANA scale-up systems. VSP G800 and VSP G1000 are validated for SAP HANA scale-out systems and are not discussed in this reference architecture guide.

For information about deploying SAP HANA scale-up systems using these servers and storage, refer to the following reference architecture guides:

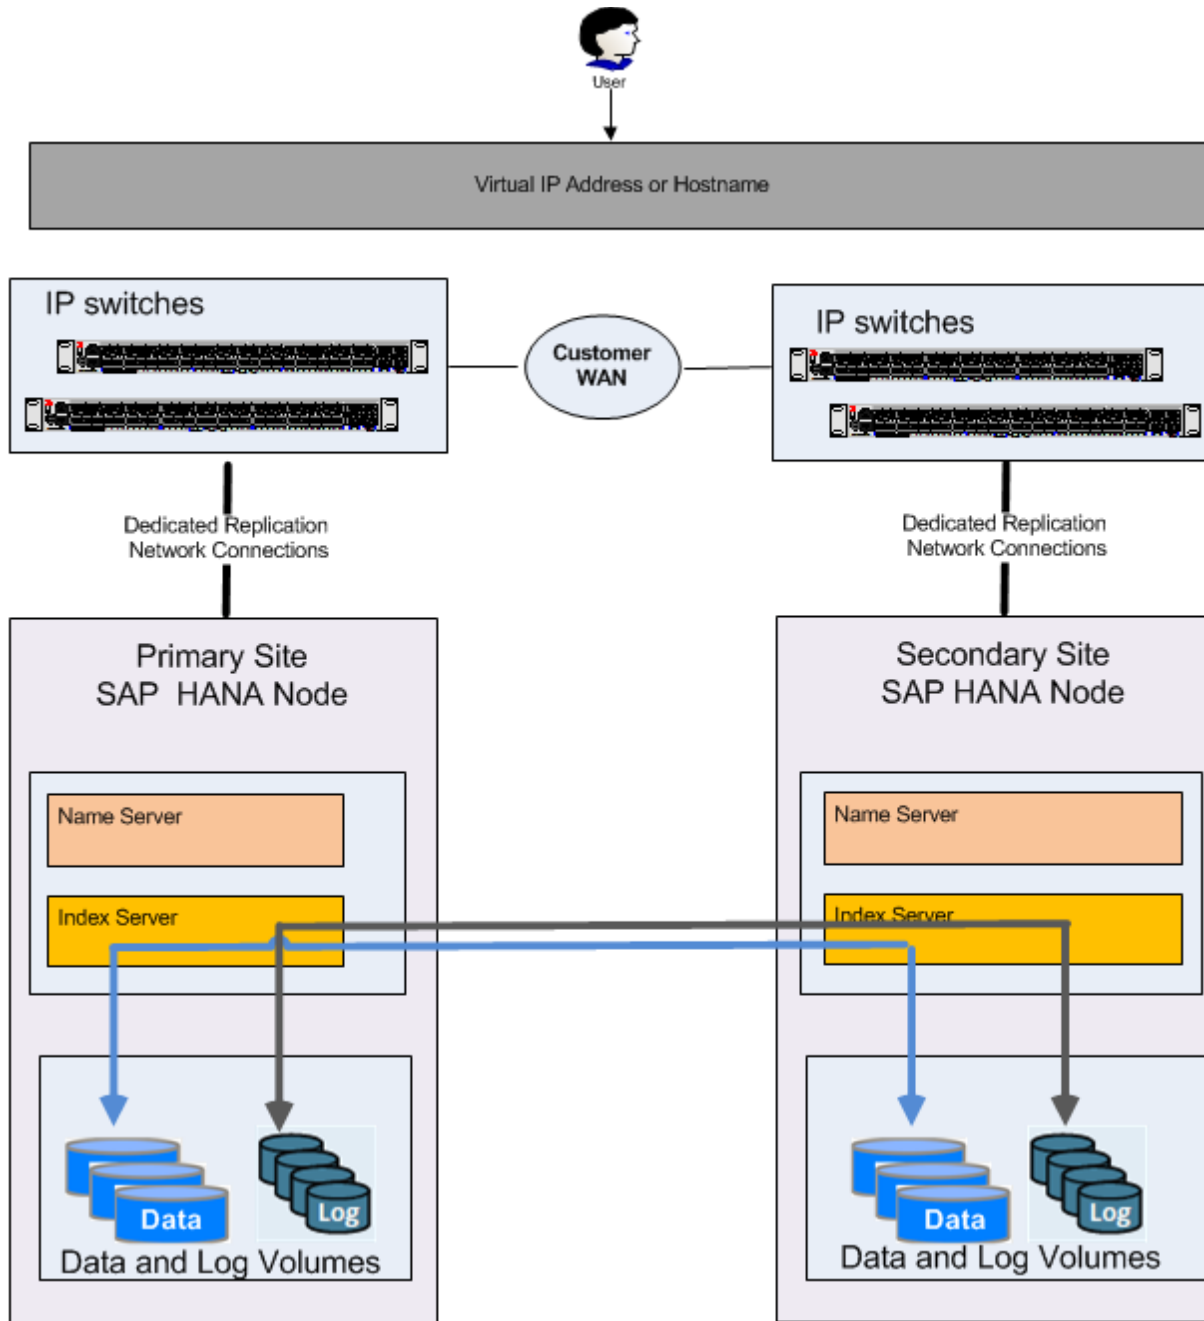
- **TDI With VSP G400 and VSP G600:** [Tailored Datacenter Integration implementation of Multiple SAP HANA Scale-Up HSW Appliances on Hitachi Virtual Storage Platform G400 and VSP G600](#). Multiple scale-up SAP HANA nodes can be deployed, VSP G400 supports a maximum of six Scale-Up nodes and VSP G600 supports a maximum of eight scale-up nodes.
- **Appliance With VSP G200:** [Hitachi Unified Compute Platform for the SAP HANA® Platform in a Scale-Up Configuration Using Hitachi Compute Blade 2500 and Hitachi Virtual Storage Platform G200](#).



## Solution Overview

This document provides an example configuration of SAP HANA disaster recovery using SAP HANA System Replication. SAP HANA scale-up systems in this guide follow the architecture defined by Hitachi Data Systems and use Hitachi Data Systems 520X B2 server blades and a VSP family storage array. The two SAP HANA scale-up systems are at two different datacenters and are referred to as "primary site" and "secondary site" in this reference architecture. There are dedicated replication network connections from the two SAP HANA nodes to Brocade VDX 6740 Ethernet switches on both primary and secondary sites. The Brocade VDX 6740 switches can be connected through the WAN that connects the two disaster recovery sites. This solution uses one QuantaPlex T41S-2U server with a single node at each site that acts as a central device for managing the SAP HANA appliance.

Two medium-sized (4-socket) SAP HANA nodes are used to validate the solution design in this reference architecture. Since SAP HANA System Replication is a solution offered by SAP, the procedure may be applied to all other Hitachi Data Systems and SAP-supported SAP HANA scale-up systems, but hardware and network configuration will change accordingly. Figure 2 provides an overview of this solution.



**Figure 2**

SAP HANA System Replication is set up from the SAP HANA node at the primary site to the SAP HANA node at the secondary site. The primary node can be a production SAP HANA node at the primary site or secondary high-availability node at the primary site (multi-tier replication if configured). The network design in this reference architecture is explained using two SAP HANA databases (not multi-tier).

## Key Solution Elements

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

### Hardware Elements

Table 2 describes the hardware used in this reference architecture for the two SAP HANA scale-up systems required for SAP HANA System Replication.

**Table 2. Hardware Elements**

Hardware		Quantity	Configuration (per unit)	Role
Hitachi Compute Blade 2500 (CB 2500) chassis		2	<ul style="list-style-type: none"> <li>■ 8-blade chassis</li> <li>■ 2 management modules</li> <li>■ 10 cooling fan modules</li> <li>■ Power Supply Modules (PSM)               <ul style="list-style-type: none"> <li>■ 2 PSM for 2-Socket</li> <li>■ 3 PSM for 4-Socket</li> <li>■ 4 PSM for 8-Socket</li> </ul> </li> <li>■ 4 I/O board modules</li> <li>■ 2 × 10 Gb 2-port LAN PCIe adapter</li> <li>■ 2 × Hitachi 16 Gb 2-port Fibre Channel adapters</li> </ul>	Server blade chassis
520X B2 server blade	2-Socket	4	<ul style="list-style-type: none"> <li>■ 2 × 18-core processors</li> <li>■ 48 × 16 GB DIMMs (1.5 TB)</li> <li>■ 1 × 2-port pass-through mezzanine card on mezzanine slot 2 and mezzanine slot 4</li> </ul>	SAP HANA server
SMP connection board for 520X B2 Blade Server	4-Socket	2	<ul style="list-style-type: none"> <li>■ SMP connection board               <ul style="list-style-type: none"> <li>■ 2-blade SMP connection board for 4-Socket</li> </ul> </li> <li>■ SMP expansion module</li> <li>■ SMP connector cover</li> </ul>	SMP connector
Hitachi Virtual Storage Platform G200/G400/G600		2	<ul style="list-style-type: none"> <li>■ Single frame</li> </ul>	Block storage for SAP HANA

**Table 2. Hardware Elements (Continued)**

Hardware	Quantity	Configuration (per unit)	Role
Brocade VDX 6740 port switch	4	<ul style="list-style-type: none"> <li>■ Two switches with distinct VLANs, each dedicated to SAP HANA Client network and replication network</li> </ul>	10 GbE Client network and replication network
Quanta Cloud Technology QuantaPlex T41S-2U server	2	<ul style="list-style-type: none"> <li>■ Intel Xeon E5-2620 v3 processor, 2.4 GHz CPU, 32 GB RAM</li> <li>■ 2 × 500 GB 7200 RPM SATA drives</li> <li>■ 1 dual-port 10GigE Intel 82599ES SFP+OCP Mezzanine card</li> <li>■ 1 dual-port 1 GigE Base-T Intel i350 Mezzanine Card</li> <li>■ Emulex Dual Port 8 Gb/sec Fibre Channel HBA</li> </ul>	One node of the server serves as the management server at each site running the following: <ul style="list-style-type: none"> <li>■ NTP</li> <li>■ Hitachi Command Suite</li> <li>■ Hi-Track® Remote Monitoring system</li> <li>■ SAP HANA Studio</li> </ul>

## Software Elements

Table 3 describes the software used in this reference architecture for two SAP HANA scale-up systems required for SAP HANA System Replication

**Table 3. Software Elements**

Software Component	Software Version
SAP HANA	SPS09 rev 97
SUSE Linux	SLES 11 SP3
Hitachi Virtual Storage Platform G400	80-02-01-00
Brocade VDX switches NOS version	4.1.3b

## Solution Design

This is the detailed solution design of this reference architecture. It includes the following sections:

- Hitachi Compute Blade 2500 Chassis Configuration
- Network Design
- Management Server
- Setting up SAP HANA System Replication

### Hitachi Compute Blade 2500 Chassis Configuration

The solution uses two Hitachi Compute Blade 2500 chassis and four 520X B2 server blades. There are two SAP HANA scale-up (medium-sized) systems with two 520X B2 server blades for each system as shown in Table 4.

**Table 4. 520X B2 Sever Blades**

Two Medium (4-Socket) Size SAP HANA Scale-Up	
Chassis 1 Server Blades	<ul style="list-style-type: none"> <li>■ Blade 3 (non-primary for SAP HANA Node 1)</li> <li>■ Blade 1 (primary for SAP HANA Node 1)</li> </ul>
Chassis 2 Server Blades	<ul style="list-style-type: none"> <li>■ Blade 3 (non-primary for SAP HANA Node 1)</li> <li>■ Blade 1 (primary for SAP HANA Node 1)</li> </ul>

- There are two management modules on the Hitachi Compute Blade 2500 chassis to connect to the management network.
- A maximum of 28 I/O board modules (IOBD) can be installed on one Hitachi Compute Blade 2500 chassis, but the solution only uses four I/O board modules.
  - Hitachi FIVE-FX 16 Gb/sec 2-port Fibre channel PCIe adapters are installed on IOBD 01A and IOBD 02A.
  - 10GBase-SR 2-port network PCIe adapters are installed on IOBD 01B and IOBD 02B.
  - Figure 3 shows the server blades for the two SAP HANA nodes in the Compute Blade 2500 chassis.




Figure 3

## Network Design

SAP recommends using a dedicated network for system replication. To understand network requirements for system replication, refer to [Network Recommendations for SAP HANA System Replication](#).

Hitachi 520X B2 server blades offer two 10 GbE ports dedicated to use for system replication. There are two 10GBASE-SR 2-port LAN adapters installed on the PCIe slots of the I/O board module of blade 1 of the Hitachi Compute Blade 2500 chassis. This solution uses two 10 GbE ports on the 10GBASE-SR 2-port LAN adapters for connectivity with the 10 GbE Brocade VDX 6740 switches.

The management module on the Hitachi Compute Blade 2500 chassis is connected to an external switch for management connectivity. Make the following network connections for replication network and client network from the Compute Blade 2500 chassis at the primary and secondary sites as shown in Figure 4.

- Port 0 of the I/O board module on PCIe slot IOBD 01B to port 1 of Brocade VDX 6740-48B.
- Port 0 of the I/O board module on PCIe slot IOBD 02B to port 1 of Brocade VDX 6740-48A.
- Port 1 of the I/O board module on PCIe slot IOBD 01B to port 0 of Brocade VDX 6740-48B.
- Port 1 of the I/O board module on PCIe slot IOBD 02B to port 0 of Brocade VDX 6740-48A.




Figure 4

This solution connects two **Brocade VDX 6740** switches using ISL. It enables both switches to act together as one single logical switch with the characteristics that, if one switch fails, there still is a path to the hosts. **Create separate VLANs for the ports used for client network and ports used for the replication network.** At the operating system level, the active-active network bond mode with options "mode= 802.3ad miimon=100 xmit\_hash\_policy=layer3+4 updelay=5000 lACP\_rate=fast" is used. The compute network setup uses the ports on the 10GBASE-SR 2-port LAN adapters. Create bonds at the operating system level using two network ports for the client network as well as for the replication network for each SAP HANA system as listed in Table 5.

**Table 5. Network Setup Using 10GBASE-SR 2-Port LAN Adapter**

SAP HANA Node	PCIe Slot Number	Switch Module Port	Network Description	Bond	IP Address
Primary	IOBD 01B	0	Client network for SAP HANA node 1	Bond0	192.168.150.111
	IOBD 02B	0	Client network for SAP HANA node 1		
	IOBD 01B	1	Replication network for SAP HANA node 1	Bond1	192.168.100.111
	IOBD 02B	1	Replication network for SAP HANA node 1		
Secondary	IOBD 01B	0	Client network for SAP HANA node 2	Bond0	192.168.150.112
	IOBD 02B	0	Client network for SAP HANA node 2		
	IOBD 01B	1	Replication network for SAP HANA node 2	Bond1	192.168.100.112
	IOBD 02B	1	Replication network for SAP HANA node 2		

## Management Server

This solution uses one node of a four-node Quanta Cloud Technology [QuantaPlex T41S-2U server](#) at each site for the management server. 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
- SAP HANA nodes
- Hitachi Virtual Storage Platform G200/G400/G600
- NTP configuration
- Hi-Track Remote Monitoring system from Hitachi Data Systems
- Hitachi Command Suite and management of the server blades
- SAP HANA Studio



Figure 5 shows the management server network ports using one dual-port 1 GbE Base-T Intel i350 mezzanine card. 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 two 10GbE and two 1GbE ports to customer switches.



- Emulex Fibre Channel 8 Gb/s Host Bus Adapter (2 port)
- Dual port 10GbE Intel 82599ES SFP+ OCP Mezanine card
- Dual port 1 GbE Base-T Intel i350 Mezanine card

**Figure 5**

Install the following software on the management server:

- Hitachi Command Suite
- Hitachi Compute Systems Manager software
- NTP server service
- PuTTY
- JRE version jre-7u51-windows-i586 (no 64 bit)
- Adobe Flash Player
- Remote client software
- SAP HANA Studio

## Planning and Prerequisites

This reference architecture guide assumes that the hardware is already set up using the network architecture as explained in the section Network Design, and the installation of the SAP HANA nodes on the primary and secondary servers has already been performed. No license key is required for the secondary SAP HANA node, as this information will be replicated from primary to secondary. Choose a replication mode (described in the section, "SAP SAP HANA System Replication") based on your requirements. If data preload is used, set the parameter 'preload\_column\_tables' on both SAP HANA nodes. If preload is used, the main parts of the column tables are already loaded into memory of the secondary node, as they were loaded in the primary. The delta only is loaded when takeover is initiated, which is typically much smaller than the main part and can be loaded within seconds in most cases. Other system parameters may be required for system replication to work efficiently and their value depends on the individual customer environment. Refer to [System Replication Configuration Parameters](#) for a list of parameters and their description.

Following are the prerequisites as described by SAP and SUSE for this solution:

- Both primary and secondary SAP HANA nodes must have the same SAP system ID and instance number but a different hostname.
- Both SAP HANA nodes have the same SAP HANA version.
- Changes to the .ini file configuration parameters made on one SAP HANA node are not replicated and should be manually made on the other system.

---

**Note** — The previous section does not list some of the prerequisites that are already addressed by using the described design. Please check the complete list of prerequisites described by SAP in the [SAP HANA Administration Guide](#).

---

Table 6 lists the information used to set up disaster recovery for SAP HANA.

**Table 6. Information Used for Disaster Recovery Setup**

	SAP HANA Node 1	SAP HANA Node 2
Hostname	saphanap	saphanas
SAP HANA SID	HIT	HIT
SAP HANA Instance Number	10	10
IP Address for client network	192.168.150.111	192.168.150.112
IP Address for replication network	192.168.100.111	192.168.100.112
Logical Name for system replication	HITSRP	HITSRS

## Setting up SAP HANA System Replication

This reference architecture assumes that installation of two SAP HANA nodes (primary and secondary nodes for system replication) has already been performed. Installation of primary and secondary nodes is performed using information provided in Table 6. Use the following steps to set up SAP HANA system replication between the primary and secondary nodes.

- Configure Name Resolution for Replication Network
- Set up SAP HANA System Replication From Primary SAP HANA Node to Secondary SAP HANA Node.

## Configure Name Resolution for Replication Network

The "Network Design" section describes the physical network connections required to provide a dedicated network for SAP HANA System Replication. Additionally, SAP HANA nodes must be configured to identify the replication network. This must be done before system replication is configured. Refer to [Network Configuration for SAP HANA System Replication](#) to configure name resolution. This is configured in the section **system\_replication\_hostname\_resolution** in the **global.ini** file on the Configuration tab of SAP HANA Studio, as shown in Figure 6. Configure the replication IP address of the secondary node on the primary node, and configure the replication IP address of the primary node on the secondary node. Figure 6 shows how the IP addresses are configured in SAP HANA Studio on the primary node.

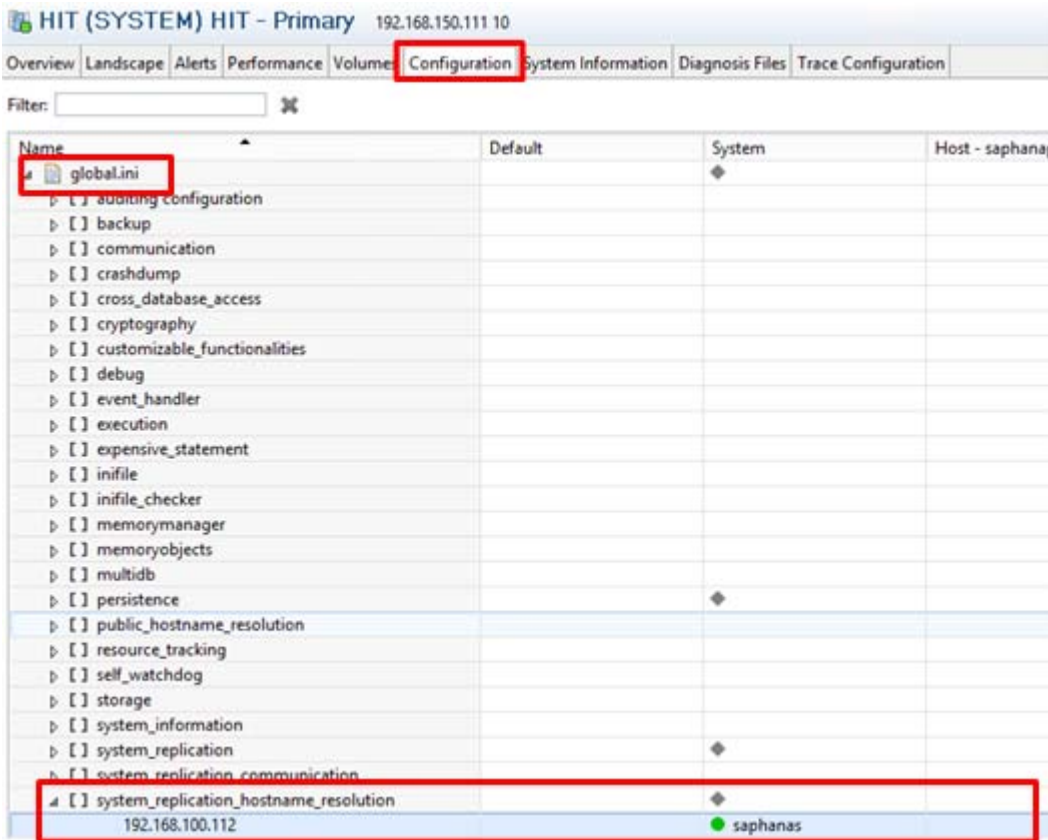


Figure 6

## Set up SAP HANA System Replication From Primary SAP HANA Node to Secondary SAP HANA Node

Perform system replication setup using the SAP [HANA system replication guide](#). The configuration can be further tuned by applying parameters described in [System Replication Configuration Parameters](#) and the SAP [HANA system replication guide](#) based on individual requirements. Validate that system replication is running by executing the command "hdbnsutil -sr\_state" on the primary SAP HANA server as user <sid>adm as shown in Figure 7.

```
saphanap:/usr/sap/HIT/HDB10> hdbnsutil -sr_state
checking for active or inactive nameserver ...

System Replication State
~~~~~

mode: primary
site id: 1
site name: SRPHIT

Host Mappings:
~~~~~

saphanap -> [SRPHIT] saphanap
saphanap -> [SRSHIT] saphanas

done.
```

Figure 7

Also, execute the following command to validate that the replication IP addresses (192.168.100.111 and 192.168.100.112) are used for replication.

```
hdbcons -e hdindexserver "replication info"
```

This command also gives the details of parameters such as whether or not "preload column tables" is used, replication status, and replication mode as highlighted in Figure 8.

```
saphanap:/usr/sap/HIT/HDB10> hdbcons -e hdbindexserver "replication info"
SAP HANA DB Management Client Console (type '\?' to get help for client com
mands)
Try to open connection to server process 'hdbindexserver' on system 'HIT',
instance '10'
SAP HANA DB Management Server Console (type 'help' to get help for server c
ommands)
Executable: hdbindexserver (PID: 2083)
[OK]
--
Dumping replication statistics ...
System Replication Primary Information
=====
System Replication Primary Configuration
[system_replication] logshipping timeout = 30
[system_replication] enable_full_sync = false
[system_replication] preload_column_tables = true
[system_replication] ensure_backup_history = true
[system_replication] enable_ssl = off
[system_replication] keep_old_style_alert = true
[system_replication] datashipping_snapshot_max_retention_time = 7200000000

- lastLogPos : 0x8dfe4a00
- lastLogPosTimestamp : 31.07.2015-15.56.17 (1438358177162445)
- lastSavepointVersion : 26272
- lastSavepointLogPos : 0x8dfe4702
- lastSavepointTimestamp : 31.07.2015-15.54.16 (1438358056887305)

1 session registered.

Session index 0
- SiteID : 2
- RemoteHost : 192.168.100.112

Log Connection
- ptr : 0x00007f38cd2e4c80
- channel : NetworkChannel FD 232 [0x00007f356768d5d8] {refCnt=
3, idx=1} 192.168.100.111/31103_tcp->192.168.100.112/13187_tcp Connected, [r
---]

- mode : ReplicationMode SynchronousMemory
- lastShippingStart : 31.07.2015-15.56.17 (1438358177162495)
- lastShippingEnd : 31.07.2015-15.56.17 (1438358177162804)
```

Figure 8

## Takeover Process

The takeover is the process of switching the current secondary system to the primary system. This is a manual process that can be performed using SAP HANA Studio or command line. Once the takeover command is executed, the former secondary system becomes the primary system, and it becomes ready to accept client connections. Some tasks must be performed before the system is fully available to users:

- Client connection recovery must be performed by the network administrator. The client or application servers must be able to reach the current primary SAP HANA database. There are many ways to implement this, such as IP redirection, DNS redirection, and use of switchover cluster software.
- Additional tasks are performed at the secondary system, for example row tables and column tables need to be loaded into memory (if preload table is used, the main parts of column tables are already loaded into memory), indices need to be rebuilt, and logs need to be replayed. For details, refer to SAP [HANA system replication guide](#).

## Use Secondary Site for DEV/QA System

It is possible to make use of the secondary site for running DEV/QA systems while the primary system is in production, but these details are not covered in this reference architecture guide. Please refer to SAP [HANA system replication guide](#).

## Engineering Validation

The failover tests listed in Table 7 were performed in the HDS lab to validate this solution.

**Table 7. Test Cases for Engineering Validation**

Test Case Number	Test Case Description	Result
1	Shut down SAP HANA on primary node by executing the command "HDB stop" and perform takeover by secondary node.	SAP HANA is started on secondary node.
2	After Step 1, start up the primary node again and re-register it for system replication as the secondary node.	Validated that system replication was active again, with the primary in Step 1 now acting as secondary.
3	Shut down the primary node in Step 2 and perform takeover to the secondary node in Step 2.	SAP HANA is started on secondary node.

 **Hitachi Data Systems**



Corporate Headquarters  
2845 Lafayette Street  
Santa Clara, CA 96050-2639 USA  
[www.HDS.com](http://www.HDS.com)    [community.HDS.com](http://community.HDS.com)

Regional Contact Information  
**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)

HITACHI is a trademark or registered trademark of Hitachi, Ltd., Other notices if required. All other trademarks, service marks and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

AS-429-00 October 2015.