High Availability in a Scale-up Environment for SAP HANA with Auto Host-Failover using Red Hat GFS2

TECHNICAL PAPER

Reference Architecture Guide

By Milind Pathak and Abhishek Dhanuka

January 2017

Hitachi Data Systems

Feedback

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

Contents

High Availability in a Scale-up Environment for SAP HANA with Auto Host-Failover using Red Hat GFS2

Reference Architecture Guide

Use this reference architecture guide to achieve high availability on Hitachi storage solutions and servers for scale-up deployments of SAP HANA for real time data processing. It uses host auto-failover in a Red Hat Global File System 2 environment.

You want your mission-critical applications in a SAP HANA database to make fast recovery after system component failure (high availability) or after a disaster (disaster recovery). This needs to happen without any data loss (zero RPO) and in very short recovery time (low RTO).

Achieve high availability or disaster recovery with zero RPO and low RTO using high availability features of the following:

- Hitachi Compute Blade 2500 (CB 2500) with 520X B2 or 520X B3 server blades
- Hitachi Virtual Storage Platform family (VSP)
- SAP HANA
- Red Hat Linux High Availability Add-on in a Global File System 2 environment

The hardware in this solution provides the following redundant hardware components to provide fault tolerance:

- **Hitachi Compute Blade 2500 with 520X B2 or 520X B3 server blades**

  This includes items such as redundant power supplies and fans, two hot-swappable management modules, and multiple Ethernet and Fibre Channel HBA interfaces.

- **Hitachi Virtual Storage Platform family storage**

  This includes items such as dual controllers, redundant front-end and back-end I/O modules, and power supply units. The storage design uses striping and parity to provide redundancy for automatic recovery from disk failures.

The software in this solution provides fault recovery with SAP HANA. This includes a watchdog function in SAP HANA that automatically restarts configured services in case of their failure, such as the index server, name server, and so forth.

In addition to these features, SAP and its partners offer the following high availability mechanisms for SAP HANA. These solutions are based on completely redundant servers and/or storage.

- **Host Auto-Failover**

  In a host auto-failover solution, you add one or more standby nodes to the SAP HANA platform. Configure these nodes to work in standby mode.

  In case of failure, data and log volumes of a failed worker node are taken over by a standby node, becoming a worker node by taking over the user load.

  The host auto-failover solution does not need additional storage, only additional servers. The database shared binaries file system must be available on all servers.

  Implementing this solution is discussed in this document.

- **Storage Replication**

  In a storage replication solution, you achieve data replication by means of storage mirroring that is independent from the database software. Mirror the disks without a control process from the SAP HANA platform. SAP HANA hardware partners offer this solution.

  The storage replication solution needs additional servers and storage.

  Implementing this solution is *not* discussed in this document.

- **SAP HANA System Replication**

  The SAP HANA system replication solution replicates all data to a secondary SAP HANA system constantly. Data can be constantly pre-loaded in the memory of the secondary system to minimize the recovery time objective (RTO).

  The SAP HANA system replication solution needs additional servers and storage.

  Implementing this solution is *not* discussed in this document.

Refer to FAQ: High Availability for SAP HANA to read more about high availability for SAP HANA.

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 reference architecture guide provides an example high availability configuration of SAP HANA using host auto-failover. This scale-up system for SAP HANA follows the architecture defined in Hitachi Unified Compute Platform 6000 for the SAP HANA Platform in a Scale-up Configuration with Intel Xeon E7-88xx v4 Processors Reference Architecture Guide (AS-503-01 or later, PDF). This solution uses one additional scale-up SAP HANA system to act as the standby node.

Host auto-failover is an "N+M" host fault recovery solution. This recovery solution adds one or more hosts to a single host or distributed SAP HANA system. The added hosts or hosts are configured to work in standby mode.

As long as the added hosts are in standby mode, the services on these hosts do not contain any data and do not accept requests or queries.

When an active worker host fails, a standby host automatically takes its place. Since the standby host may take over operation from any of the worker hosts, it needs access to all the database volumes. Accomplish this access in one of two ways:

- Use a shared networked storage server with a distributed file system
- Use vendor-specific solutions that use a SAP HANA programmatic interface (the Storage Connector API) to dynamically detach and attach (mount) networked storage upon failover. This networked storage can be attached using block storage by Fiber Channel.

Refer to SAP HANA Host Auto-Failover for further details about this topic.

Figure 1 on page 4 shows an overview of host auto-failover in a SAP HANA environment.

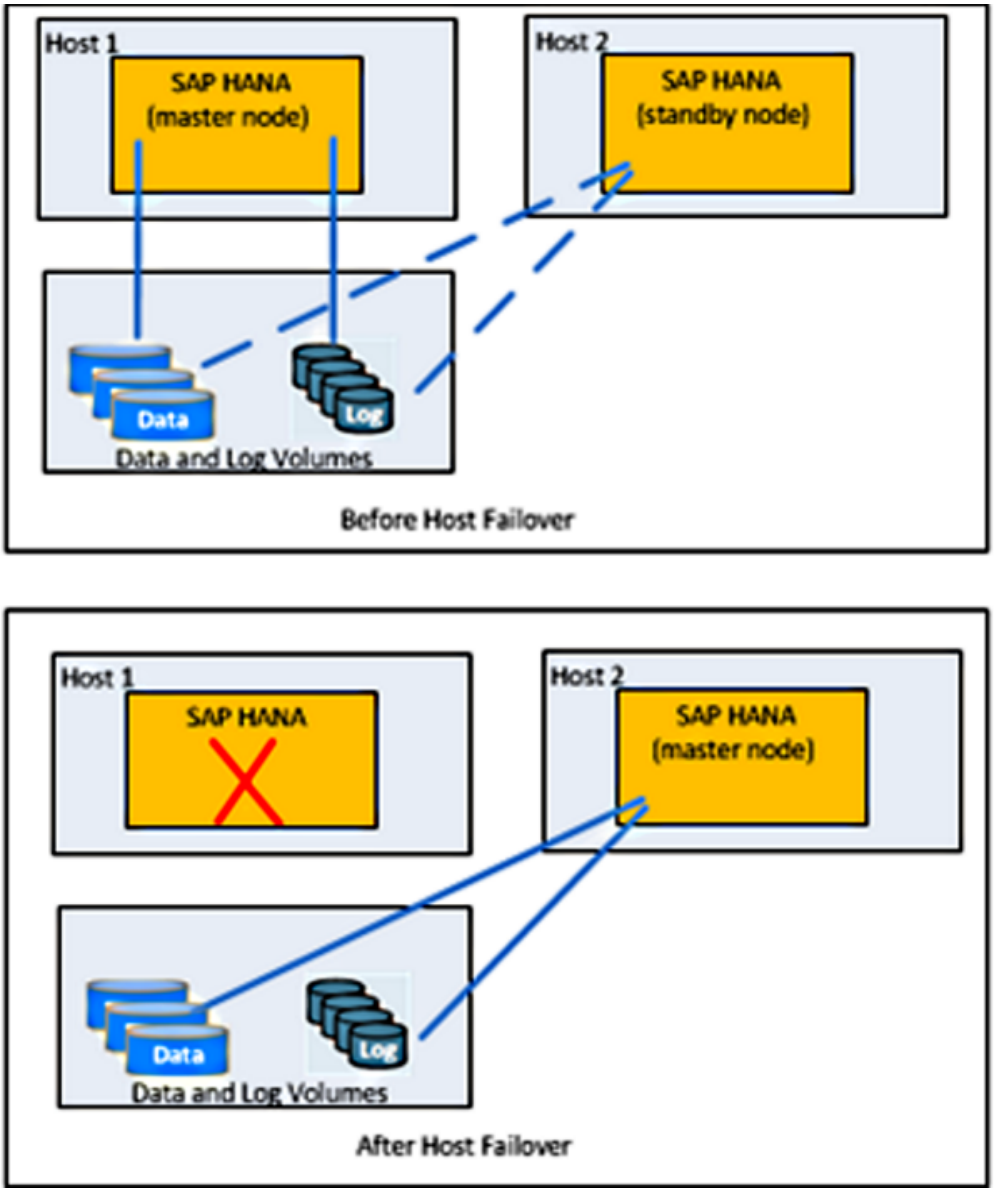

Figure 1

This reference architecture guide uses the Storage Connector API. If the primary SAP HANA database on Host 1 fails (the worker node), the data and log file system are unmounted from this host and are mounted on Host 2 (the standby node). Host auto-failover needs a shared file system to store the SAP HANA database shared binaries.

As shown in Figure 2, the Red Hat Global File System 2 cluster file system is used in this reference architecture to achieve the shared file system.

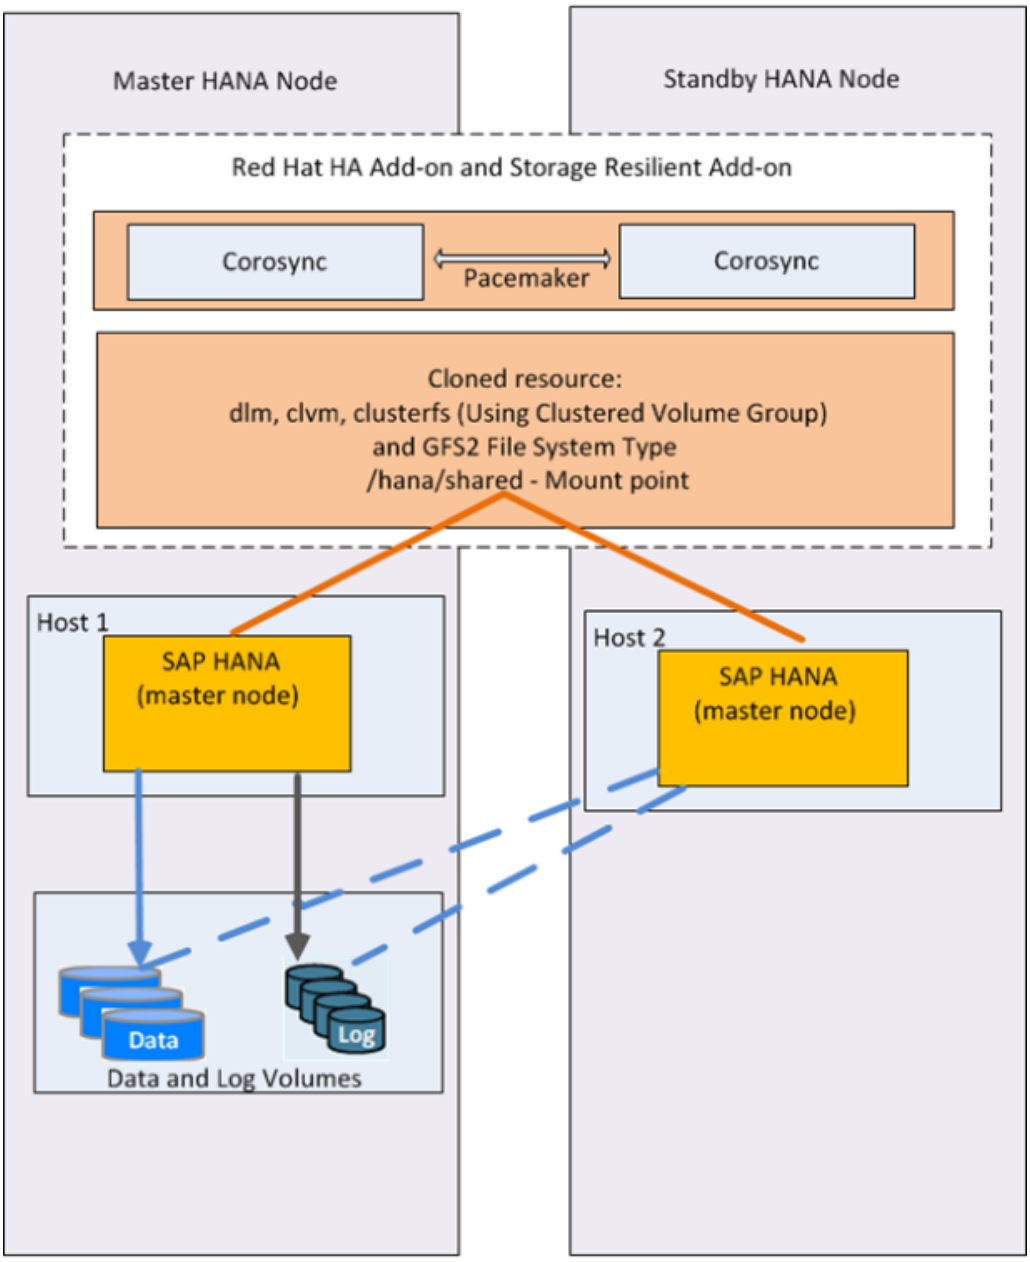

Figure 2

Key Solution Elements

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

Hardware Elements

Table 1 describes the hardware used in this reference architecture for two scale-up SAP HANA systems. Refer to 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 Reference Architecture Guide (AS-386-05 or later, PDF) for details about the architecture and design of a scale-up SAP HANA system.

Table 1. Hardware Elements

Hardware	Quantity	Configuration (per unit)	Role

Hitachi Compute Blade 2500 chassis	1	■ 8 server blade chassis
		■ 2 management modules
		■ 10 cooling fan modules
		■ Power supply modules
		  ■ 2 modules for 2-socket
		  ■ 3 modules for 4-socket
		  ■ 4 modules for 8-socket
		■ 4 I/O board modules
		■ 2 × 10 GbE 2-port LAN PCIe adapter
		■ 2 Hitachi 16 Gb/sec 2-port Fibre Channel adapters
			Server blade chassis

520X B2 server blade (4-Socket)	4	■ 2 × 18-core processors ■ RAM for all listed memory sizes ■ 1 × 2-port pass through mezzanine card on mezzanine slot 2 and mezzanine slot 4 of server blade 1 and server blade 2 (on single server blade for 2-socket)	SAP HANA server master and standby node

SMP connection board for 520X server blade (4-Socket)	2	■ SMP connection board
		  ■ 2-server blade SMP connection board for 4-socket
		  ■ SMP expansion module
		  ■ SMP connector cover
			SMP connector

Hitachi Virtual Storage Platform G200	1	■ Single frame	Block storage for SAP HANA nodes

Hitachi Compute Blade 2500

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

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

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

This solution uses two 520X B2 server blades in Hitachi Compute Blade 2500.

Hitachi Virtual Storage Platform Gx00 Models

Hitachi Virtual Storage Platform Gx00 models are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

This solution uses Hitachi Virtual Storage Platform G200. The operating system LUNs, SAP HANA data, and log LUNs reside on this storage array, as well as the shared LUNs for SAP HANA configuration files, binaries, and traces.

Software Elements

Table 2 describes the software products used to deploy the two High Availability nodes.

Table 2. Software Elements

Software Component	Software Version
SAP HANA	SPS12 or later
Red Hat Linux	■ Red Hat Enterprise Linux for SAP HANA with High Availability and Smart Management 7.2 ■ Red Hat Enterprise Linux Resilient Storage Add-On

Note — Use latest version of Red Hat High Availability Add-on and Red Hat Resilient Storage Add-On.

Red Hat Enterprise Linux

Red Hat Enterprise Linux delivers military-grade security, 99.999% uptime, support for business-critical workloads, and so much more. Ultimately, the platform helps you reallocate resources from maintaining the status quo to tackling new challenges.

Red Hat Enterprise Linux Server for SAP HANA provides an open, reliable, and scalable foundation for your most demanding data solutions. This ready-to-use environment is preconfigured for performance and optimized for SAP HANA.

The only support for changing the configuration settings is along the guidelines from SAP and Red Hat. Other changes may cause significant performance problems. The following SAP Note for Red Hat Enterprise Linux is a good starting point for information on this topic:

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

For more details, see section 2.1.4.1, “Updating and Patching the Operating System” in the *SAP HANA Technical Operations Manual*.

This solution requires Red Hat Enterprise Linux for SAP HANA with High Availability and Smart Management. This is a variant of Red Hat Enterprise Linux for SAP HANA, which enhances native SAP HANA replication and fail-over technology to automate the takeover process. It consists of these components:

- Red Hat Enterprise Linux for SAP HANA
- Red Hat Enterprise Linux High Availability Add-on
- Red Hat Enterprise Linux Smart Management Add-on

Red Hat Enterprise Linux for SAP HANA with High Availability and Smart Management

The Red Hat High Availability Add-On is an integrated set of software components to deploy in a variety of configurations to suit your needs for performance, high availability, load balancing, scalability, file sharing, and economy.

This add-on is based on Pacemaker. A cluster configured with Pacemaker comprises separate component daemons that do the following:

- Monitor cluster membership
- Scripts that manage the services
- Resource management subsystems that monitor the disparate resources

Refer to High Availability Add-On Overview for the components of the Pacemaker architecture.

Red Hat Enterprise Linux Resilient Storage Add-on

Red Hat Enterprise Linux for Resilient Storage Add-On is required for Global File System 2 (GFS2).

The Red Hat GFS2 file system is included in the Red Hat Resilient Storage Add-On. It is a native file system that interfaces directly with the Linux kernel file system interface (VFS layer). When implemented as a cluster file system, GFS2 employs distributed metadata and multiple journals.

Red Hat supports the use of GFS2 file systems only as implemented in High Availability Add-On. GFS2 is based on a 64-bit architecture, which can theoretically accommodate an 8 EB file system. However, the current supported maximum size of a GFS2 file system for 64-bit hardware is 100 TB. The current supported maximum size of a GFS2 file system for 32-bit hardware is 16 TB.

In this solution, use the GFS2 file system for SAP HANA database shared binaries.

Solution Design

The following is the detailed solution design of this reference architecture guide:

Planning and Prerequisites

This reference architecture guide assumes that the hardware is already setup using the network architecture as explained in "Network Design" on page 10. Refer to Red Hat Enterprise Linux 7 Global File System 2 (PDF) for pre-requisites and best practices for setting up the GFS2 file system.

The following items are required for the setup of this high availability solution:

- Red Hat Cluster Communication IP addresses and host names
- SAP HANA client network IP addresses and host names
- A shared storage device (LUN) with Fiber Channel path added to both servers
- SAP HANA shared, data and log volumes for the SAP HANA database with the Fiber Channel path added to both servers
- SAP HANA SID and Instance numbers

Table 3 lists the information used to setup high availability in this reference architecture.

Table 3. Information used for High Availability Setup

	HANA Node 1	HANA Node 2
Public Hostname	saphanan1	saphanan2
IP Address for client network	192.168.150.201	192.168.150.202
Cluster Hostname	node001	node002
IP Address for Cluster Communication network	192.168.100.201	192.168.100.202
SAP HANA SID	HIT	
SAP HANA Instance Number	10	

STONITH Device (Fencing)

In a Red Hat High Availability Cluster terminology, **STONITH** is an acronym for *shoot the other node in the head*.

STONITH protects your data from being corrupted by rogue nodes or concurrent access. Just because a node is unresponsive, this does not mean it is not accessing your data. The only way to be 100% sure that your data is safe before allowing the data to be accessed from another node is to fence the node using STONITH. This makes certain that the node is truly offline.

STONITH also has a role to play in the event that a clustered service cannot be stopped. In this case, the cluster uses STONITH to force the whole node offline. This makes it safe to start the service elsewhere.

A cluster without the STONITH mechanism is not supported.

Red Hat Cluster Suite is able to use SCSI persistent reservations as a fencing method. SCSI registration occurs when a node registers a unique key with a device (in this case, a LUN). A device can have many registrations. For this solution, each node creates a registration on each device.

When a node failure occurs, the **fence_scsi** agent installed with Red Hat High Availability Add-on removes the failed node's key from all devices. This prevents the node from being able to write to those devices. Refer to Using SCSI Persistent Reservation Fencing (fence_scsi) with pacemaker in a RHEL 6 or 7 High Availability cluster for pre-requisites and further details.

Create a small LUN (50 MB) on the storage array that is shared between the cluster members. Map this LUN to the master and the standby SAP HANA servers through storage ports. Note the SCSI identifier of the block device (/dev/disk/by-id/scsi*) of this LUN. The SCSI identifier should be the same on both primary and secondary HANA servers. It is possible to add more than one device in a cluster for redundancy.

Network Design

The network architecture used in this reference architecture can be classified in two categories:

Management Network

This network provides management network for the solution using Brocade VDX and ICX switches. Using this network, perform network administrative tasks such as configuration of switches of accessing the BMC console of the server blades.

Connect the 1 GbE management port of the management module on Hitachi Compute Blade 2500 to a Brocade ICX 6430 24 port switch or to any other external 1 GbE switch for management network connectivity. The connectivity is shown in Figure 3 on page 11.

Compute Network:

This solution requires two separate compute networks:

- **SAP HANA Client Network**

  This network provides communication between the SAP HANA production system and the SAP production application servers.

- **Cluster Communication Network**

  This network is used for cluster communication network for Red Hat Enterprise Linux High Availability Add-on.

Configure the cluster communication network and SAP HANA client network as dedicated network bonds. There are two 10GBASE-SR 2-port LAN adapters installed on the PCIe slots of the I/O board module of server 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 external switches. Make the following network connections for client and replication networks as shown:

- 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 3 of Brocade VDX 6740-48B
- Port 1 of the I/O board module on PCIe slot IOBD 02B to port 3 of Brocade VDX 6740-48A
- Port 0 of the I/O board module on PCIe slot IOBD 05B to port 2 of Brocade VDX 6740-48B
- Port 0 of the I/O board module on PCIe slot IOBD 06B to port 2 of Brocade VDX 6740-48A
- Port 1 of the I/O board module on PCIe slot IOBD 05B to port 4 of Brocade VDX 6740-48B
- Port 1 of the I/O board module on PCIe slot IOBD 06B to port 4 of Brocade VDX 6740-48A

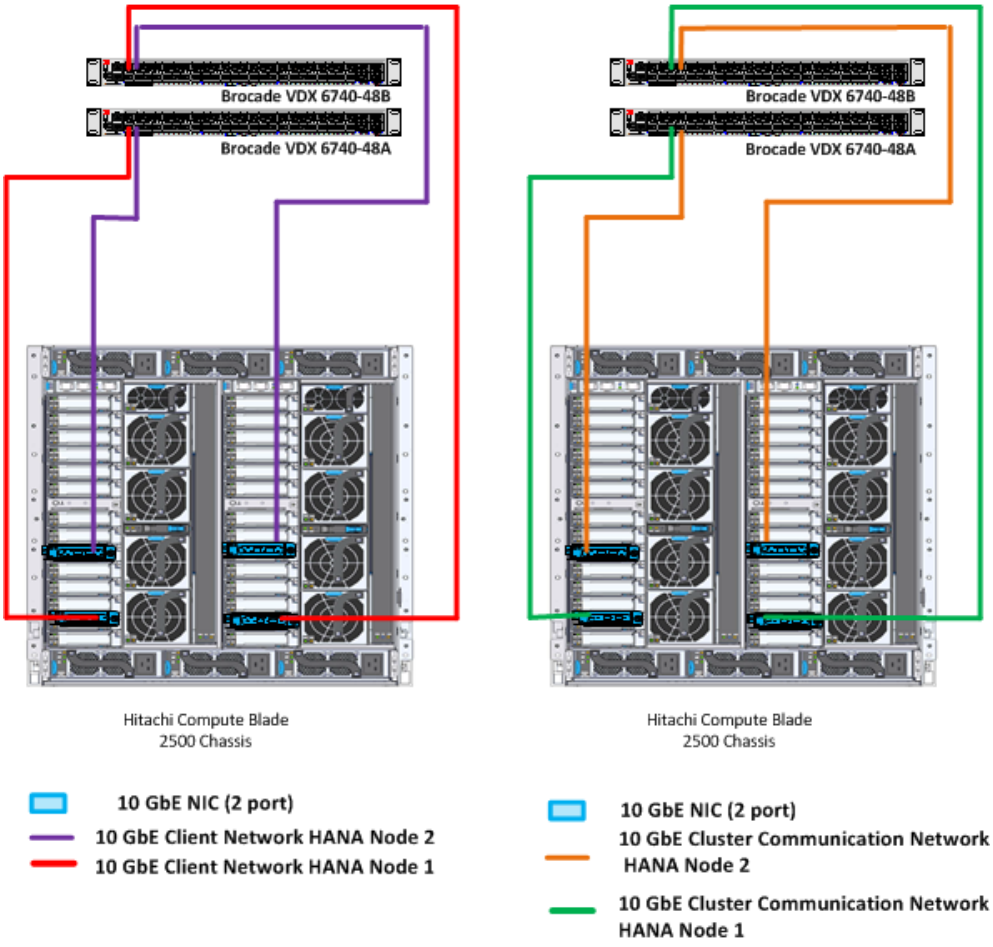

Figure 3

This solution connects two **Brocade VDX 6740** switches together 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 the ports used for the cluster communication network.
- At the operating system level, create an **active-active** network bond mode with these options:

  mode= 802.3ad miimon=100 xmit_hash_policy=layer3+4 updelay=5000 lacp_rate=fast

  The compute network setup uses the ports on the 10GBASE-SR 2-port LAN adapters.
- Create bonds at operating system level using two network ports for client network as well as for cluster communication network for each SAP HANA system, as listed in Table 4

Table 4. Network Setup Using 10GBASE-SR 2-Port LAN Adapter

HANA Node	PCIe Slot Number	PCIe Slot Port	Network Description	Bond	IP Address
Primary SAP HANA server	IOBD 01B	0	Client network for SAP HANA node 1	Bond0	192.168.150.201
	IOBD 02B	0	Client network for SAP HANA node 1		
	IOBD 01B	1	Cluster communication network for SAP HANA node 1	Bond1	192.168.100.201
	IOBD 02B	1	Cluster communication network for SAP HANA node 1		
Secondary SAP HANA server	IOBD 05B	0	Client network for SAP HANA node 2	Bond0	192.168.150.202
	IOBD 06B	0	Client network for SAP HANA node 2		
	IOBD 05B	1	Cluster communication network for SAP HANA node 2	Bond1	192.168.100.202
	IOBD 06B	1	Cluster communication network for SAP HANA node 2		

Red Hat High Availability Add-on Configuration

This describes how to install and configure Red Hat High Availability Add-on to automate the failover process in SAP HANA system replication.

Installation

Install the Red Hat High Availability Add-on and Resilient Storage Add-on packages. Use the latest versions for Red Hat Enterprise Linux 7.2 of the following:

- Corosync
- Pacemaker
- PCS
- fence-agents
- Global File System 2
- LVM2 (and their dependencies)

Perform this installation on the primary and the secondary SAP HANA servers.

Configuration

Note — The validation of this reference architecture used SCSI-based fencing. However, it is possible to implement IPMI-based fencing, as well. For such an implementation, both SAP HANA nodes must be able to access each other's BMC IP address.

Refer to Creating a Red Hat High-Availability Cluster with Pacemaker for detailed steps to configure two-node clusters.

The following are the high level steps to configure Red Hat Availability Add-on:

1. Create a two node cluster.
2. Set up the required cluster and resource parameters.
3. Create the SCSI fencing resources.

The device ID used is created as described in “STONITH Device (Fencing)” on page 10.

This completes the basic cluster configuration on the master and the standby SAP HANA servers.

Red Hat Global File System 2 Configuration

To configure the Red Hat Global File System 2 (GFS2), do the following.

1. Create this directory structure on both cluster nodes: /hana/shared
2. Create a physical volume and volume group for the SAP HANA shared volumes. The volume group must be created with type **cluster**.
3. Create cluster resources for distributed lock manager (dlm) and cluster logical volume manager (clvm).

4. Create GFS2 using the volume group created earlier in this procedure.
5. Create a cluster file system resource in the Red Hat cluster using GFS2 and this mount point: /hana/shared

   The cluster now mounts GFS2 on mount point /hana/shared.

Commands and best practices for creating the file system are provided in Red Hat documentation Configuring a GFS2 File System in a Cluster.

Provision the necessary space and create the filesystems for the data and log volumes to be used for SAP HANA following the architecture defined in Hitachi Unified Compute Platform 6000 for the SAP HANA Platform in a Scale-up Configuration with Intel Xeon E7-88xx v4 Processors Reference Architecture Guide (AS-503-01 or later, PDF).

SAP HANA Database Installation on Master and Standby Nodes

Perform installation of the SAP HANA database master node and the standby node on the Red Hat cluster nodes *saphanan1* and *saphanan2.*

Use Red Hat Global File System 2, as created in “Red Hat Global File System 2 Configuration,” starting on page 13, for shared binaries of the SAP HANA database. Refer to SAP HANA Server Installation and Update Guide (PDF) for installation instructions.

Engineering Validation

The failover tests listed in Table 5 were performed in the Hitachi Data Systems lab to validate this solution.

Table 5. Test Cases for Engineering Validation

Test Case	Description	Results
Master node reboot (soft shutdown)	Reboot the master server by running the **reboot** command at the operating system level.	All resources on the server stopped and the server rebooted. When the server came back up, all resources started and Red Hat GFS2 mounted again. SAP HANA failed over to the standby node.
Master node reboot (abrupt shutdown at the operating system level)	Reboot master server by running the **echo b > /proc/sysrq-trigger** command at the operating system level.	All resources on the server are stopped and the server rebooted. When the server came back up, all resources started and Red Hat GFS2 file system mounted again. SAP HANA failed over to the standby node.
Network single link network failure at the server	Remove one of the Red Hat cluster network cables from the NIC card of the server.	No impact on the cluster or SAP HANA.
Network single link network failure at the switch	Remove one of the Red Hat cluster network cables from the Brocade VDX 6740 switch.	No impact on the cluster or SAP HANA.
Switch failure (reboot a Brocade VDX 6740 switch that provides the Red Hat cluster)	Reboot one of the switches by running the **reload** command at the switch level.	No impact on the cluster or SAP HANA.
Manual fencing	Manually fence the master node using the **pcs stonith fence <node hostname>** command from the standby node.	All resources on the fenced server stopped and the server rebooted. Other cluster nodes are not affected. When the server came back up, all resources started and Red Hat GFS2 mounted again. SAP HANA failed over to the standby node.
Trigger reboot of master server while /hana/shared is in use	Trigger a backup of the SAP HANA database on Red Hat GFS2 and then reboot the HANA master server using the **echo b > /proc/sysrq-trigger** command	The backup is cancelled and the master server rebooted with the standby node taking over the master role. When the original master came back up, all resources started and Red Hat GFS2 is mounted again.
Failing of cluster manager on master node	Run the **kill -9 pacemaker** command on all servers.	The cluster node is fenced and the server rebooted. Other cluster nodes are not affected. When the server came back up, all resources started and Red Hat GFS2 mounted again. SAP HANA failed over to the standby node

Cluster Configuration

This is the complete cluster configuration used to validate this solution in the Hitachi Data Systems lab. Further tuning of certain parameters may be required to make this work in your environment. Also, use your IP addresses and hostnames, based on your environment.

Cluster Name: hanagfs2

Corosync Nodes:

node001 node002

Pacemaker Nodes:

node001 node002

Resources:

Stonith Devices:

Resource: hana_scsi (class=stonith type=fence_scsi)

Attributes: devices=/dev/disk/by-id/dm-name-360060e80124e560050404e56000004ff pcmk_host_list="node001 node002" pcmk_reboot_action=off

Meta Attrs: provides=unfencing

Operations: monitor interval=60s (hana_scsi-monitor-interval-60s)

Fencing Levels:

Location Constraints:

Ordering Constraints:

Colocation Constraints:

Ticket Constraints:

Alerts:

No alerts defined

Resources Defaults:

default-resource-stickness: 1000

default-migration-threshold: 5000

Operations Defaults:

timeout: 600s

Cluster Properties:

cluster-infrastructure: corosync

cluster-name: hanagfs2

dc-version: 1.1.15-9.el7-e174ec8

have-watchdog: false

last-lrm-refresh: 1476924731

no-quorum-policy: ignore

stonith-watchdog-timeout: 0

Quorum:

Options:

IPMI-based STONITH Device

This solution was validated using a shared storage based fencing mechanism. However, IPMI based fencing can be used, also. To implement IPMI-based fencing, both cluster nodes must be able to reach each other's BMC IP address.

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 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 96050-2639 USA
www.HDS.com community.HDS.com

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

AS-571-00. January 2017.