

Disaster Recovery Using Hitachi Universal Replicator on Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration With Hitachi Advanced Server DS7040

Reference Architecture Guide

By Yingping Niu

March 2019

Feedback

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

Revision History

Revision	Changes	Date
MK-SL-141-00	Initial release	March 5, 2018

Table of Contents

Solution Overview	3
Hardware Elements	6
Software Elements	9
Solution Design	10
Hitachi Advanced Server DS7040 Architecture	10
Storage Architecture	10
Management Server	11
Network Architecture	12
Journal for Hitachi Universal Replicator	12
Inter-site Configuration for Hitachi Universal Replicator	14
Disaster Recovery and Replication Components	15
Engineering Validation	22
Test Methodology	22
Test Results	22

Disaster Recovery Using Hitachi Universal Replicator on Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration With Hitachi Advanced Server DS7040

Reference Architecture Guide

Use this reference architecture guide to design a remote storage replication solution using Hitachi Universal Replicator to enable the disaster recovery for Hitachi Solution for the SAP HANA platform. This is a scale-out configuration with either of the following configurations:

- A 24 TB configuration using 1.5 TB SAP HANA nodes
- A 48 TB configuration using 3 TB SAP HANA nodes

For protection from the failure of SAP HANA when a data center is affected by an event that stops operations for a long period of time, use two sites in your disaster recovery strategy. With Hitachi Virtual Storage Platform technology, you maintain performance for the SAP HANA production instance and for site-to-site replication. This solution performs asynchronous replication of following on Hitachi Virtual Storage Platform (VSP) G700 or Virtual Storage Platform F700 to a secondary site:

- SAP HANA data volumes
- SAP HANA log volumes
- SAP HANA shared volumes

In addition to using it for failover production for SAP HANA from the primary site, the second site can handle your quality assurance environment for the SAP HANA landscape if you install additional disk drives.

This guide documents how to deploy this configuration using the following:

- Hitachi Advanced Server DS7040
- Hitachi Virtual Storage Platform G700 using solid state drives (SSD) or Virtual Storage Platform F700
- Hitachi Advanced Server DS120
- SUSE Linux Enterprise Server (SLES) or Red Hat Enterprise Linux (RHEL)
- Global File System 2 (GFS2)
- SAP HANA
- Hitachi Universal Replicator (HUR)

Data centers at each site must have almost identical hardware for this disaster recovery solution.

Implementing Hitachi Universal Replicator in this environment permits using your secondary site as a quality assurance environment, saving you the cost of creating a test environment. This additional use requires installing additional disk drives to the secondary site storage to run the quality assurance system.

The environment described in this architecture reference guide consists of Hitachi Solution for SAP HANA in a scale-out system using preconfigured hardware with SAP HANA software preinstalled. If you need changes to this architecture, contact the sales team at Hitachi Vantara or your Hitachi Vantara representative.

Engineering testing of this solution in the Hitachi Vantara lab was only on a 2+1 SAP HANA scale-out configuration. However, this reference architecture supports the scale-out configurations listed in [Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration using Hitachi Advanced Server DS7040 Reference Architecture](#). For information concerning your implementation of this environment, contact your Hitachi Vantara representative.

Use this document to understand an example disaster recovery architecture using Hitachi Universal Replicator on Unified Compute Platform for SAP HANA in a scale-out configuration. This technical paper assumes you have familiarity with the following:

- Storage area network-based storage systems
- General storage replication skills and concepts
- General network and virtual IP knowledge
- General WAN knowledge
- Advanced SAP HANA skills
- Disaster recovery scenarios

The following are important references used by this solution:

- [Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration using Hitachi Advanced Server DS7040 Reference Architecture](#)
- [Command Control Interface User and Reference Guide](#)
- [Hitachi Virtual Storage Platform F Series and G Series \(SVOS RF 8.3\) Universal Volume Manager User Guide](#)
- [Hitachi Virtual Storage Platform G130, G/F350, G/F370, G/F700, G/F900 Provision Guide](#)

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

The primary site, Site A, contains a production SAP HANA database instance. Implement this site as a scale-out configuration of Hitachi Solution for SAP HANA.

The secondary site, Site B, is an exact replica of primary site, with exception of optional additional storage disks if you need non-production instances. Site B houses the following:

- Failover production instance
- (Optional) Non-production instances that requires additional disks, such as the following:
 - Quality assurance
 - Development
 - Test, in-memory database of SAP HANA.

The design of the SAP HANA in-memory database enables this solution to use the same set of Hitachi Advanced Server DS7040 nodes for production and non-production instances.

In the Hitachi Engineering test environment for this solution, the Hitachi Virtual Storage Platform at the secondary site housed one set of disks for data, log, and /hana/shared LUNs for the following:

- The replicated production instance
- The quality assurance system (optional)

In this solution, install command control interface on the management servers on Site A and Site B, which performs the data replication operations within Hitachi Virtual Storage Platform at each site using Hitachi Open Remote Copy Manager (HORCM) instances. Each instance at the Site A and Site B management server has its own Open Remote Copy Manager configuration file that lists the following for production instance replication between two sites:

- SAP HANA data volumes
- SAP HANA log volumes
- SAP HANA shared volumes

Hitachi Virtual Storage Platform works with Hitachi Universal Replicator for disaster recovery of the SAP HANA environment. This solution uses a Fibre Channel over IP (FCIP) switch for the wide area network connections between each Virtual Storage Platform at the primary site and the secondary site. Take the options for long distance data replication in the landscape using the pre-existing enterprise network infrastructure and service provider into consideration for the SAP HANA disaster recovery setup.

With Hitachi Universal Replicator, updates on the SAP HANA nodes of the primary production volume in Virtual Storage Platform on Site A are copied to a local journal volume on primary site storage. Then, Virtual Storage Platform on Site B “pulls” data from the primary site journal volume across the inter-site wide area network connection to the secondary volume. The local system is free to perform its role as a transaction processing resource rather than as a replication engine.

Hitachi Universal Replicator offers a solution to avoid cases when a data center is affected by a disaster that stops operations for a long period of time. The Universal Replicator system locates a secondary storage system at Site B away from the primary storage system at the main data center on Site A. It copies data on the primary volumes (P-VOLs) at Site A to the secondary volumes (S-VOLs) at Site B asynchronously from the host write operations to the P-VOLs. Journal data is created synchronously with the updates to the P-VOL to provide a copy of the data written to the P-VOL. The journal data is managed at Site A and Site B to ensure the consistency of the primary and secondary volumes.

The redundancy provided by the RAID configuration (for example, RAID-1 or RAID-5) enables recovery from a P-VOL read failure. The primary storage system never reads the S-VOL used with Universal Replicator for data recovery.

In case of an outage or any component failure in Site A, an administrator initiates a manual failover to Site B. There are two different ways to enable client connection recovery using either virtual IP failover or DNS failover configuration. The actual implementation differs, based on the network and cluster management capabilities.

To replicate Hitachi Solution for SAP HANA, this reference solution uses the following on each site:

- **Hitachi Advanced Server DS7040**

Hitachi Advanced Server DS7040 provide high reliability, extreme performance and agile scalability.

This solution makes full use of the huge memory capacity of DS7040 to deliver real-world business benefits from data analytics. This solution uses a different number of DS7040 servers to provide the flexibility of multiple scale-out configuration options.

- **Hitachi Virtual Storage Platform G700 or VSP F700 using solid state drives (SSD)**

Hitachi Virtual Storage Platform scales for all data types, flexibly adapting for performance, capacity, and multi-vendor storage.

In this solution, the persistent storage of the SAP HANA server resides on Hitachi Virtual Storage Platform G700 or VSP F700 using solid state drives.

- **Hitachi Advanced Server DS120**

Optimized for performance, high density, and power efficiency in a dual-processor server, Hitachi Advanced Server DS120, a rack-mounted server, delivers a balance of compute and storage capacity.

- **Hitachi Universal Replicator**

With asynchronous replication over any distance, Hitachi Universal Replicator provides for business continuity, optimization of resource usage, and improves efficiency as well as resilience.

- **SAP HANA Platform**

SAP HANA is a multi-purpose, in-memory database appliance to analyze transactional and analytical data.

- **Cisco Nexus 93180YC-EX switch**

This 48-port switch provides 10 GbE connectivity to the appliance.

- **Cisco Nexus 3048 switch**

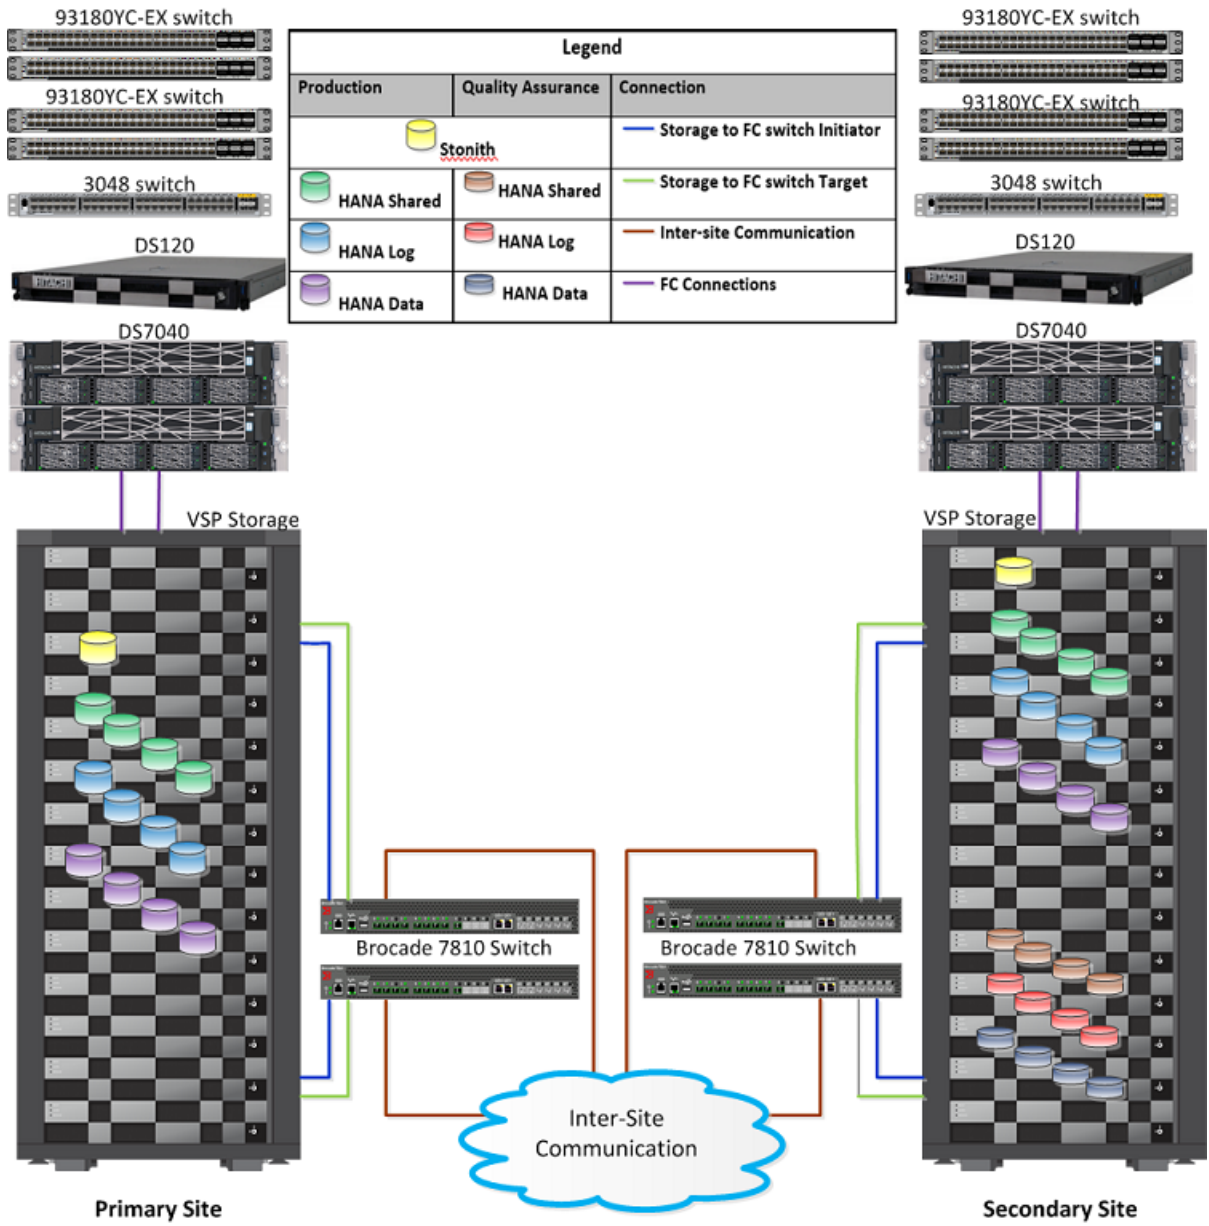
This 48-port 1 GbE switch provides management network to the appliance.

- **Brocade 7810 Extension Switch**

With twelve 32 Gb/s-capable Fibre Channel ports and six 1/10- GbE ports, this switch provides replication connectivity over long distance between data centers.

Figure 1 shows the configuration of this solution.

Figure 1



Hardware Elements

Table 1 describes the hardware used to deploy a scale-out configuration with two active nodes and one standby node at each site as an example (excluding any non-production system or systems).

TABLE 1. HARDWARE ELEMENTS EACH SITE

Hardware	Quantity	Configuration	Role
Hitachi Advanced Server DS7040	3	<ul style="list-style-type: none"> ■ Each server has the following: ■ 2 compute modules ■ CPU for each SAP HANA node: <ul style="list-style-type: none"> ■ 4 Intel Xeon Platinum 8176/8176M processors 28-core, 2.1 GHz, 165 W Or ■ 4 Intel Xeon Platinum 8180/8180M processors 28-core, 2.5 GHz, 205 W ■ 4 heat sinks for CPU 0/1 of both compute modules ■ Memory for each SAP HANA node: <ul style="list-style-type: none"> ■ 1.5 TB RAM with 48 × 32 GB DIMMs or 24 × 64 GB DIMMs per SAP HANA node Or ■ 3 TB RAM with 48 × 64 GB DIMMs or 24 × 128 GB DIMMs per SAP HANA node 	<p>Server blade chassis.</p> <ul style="list-style-type: none"> ■ One chassis is required for two SAP HANA nodes.
Hitachi Virtual Storage Platform G700 or VSP F700	1	<ul style="list-style-type: none"> ■ CTL — 1 pair ■ DKB — 2 pairs ■ CHB — 6 pairs ■ MPU — 1 pair ■ Cache — 512 GB ■ 1.9 TB solid state drives 	<p>Server blade for the SAP HANA nodes.</p>

TABLE 1. HARDWARE ELEMENTS EACH SITE (CONTINUED)

Hardware	Quantity	Configuration	Role
Mellanox Technologies MT27710 Family [ConnectX-4 Lx]	4	<ul style="list-style-type: none"> Mellanox Connectx4 dual port PCIe card for each HANA node 	10 GbE SUSE Enterprise Linux High Availability Extension (HAE) or Red Hat Enterprise Linux High Availability (HA) cluster communication network, SAP HANA inter-cluster network, client network
	6	<ul style="list-style-type: none"> Cisco SFP+ 3M twinaxial cables for each HANA node 	
Emulex LPe31002 M6 blade, 16 Gb/s Fibre Channel Dual HBA without cable	2	<ul style="list-style-type: none"> Only on Module 0 	Connectivity to the external storage sub-system
	2	<ul style="list-style-type: none"> SFP 	
Hitachi Advanced Server DS120	1	<ul style="list-style-type: none"> 2 Intel 4110 8-core processors; 85 W; 2.1 GHz 2 Samsung 32 GB DDR4 R-DIMM 2666 MHz memory 1 OCP mezzanine card — Quanta (Intel) X557-T4 10 GbE, 4 ports 1 PCIe — Intel XXV710 SFP28 10/25 GbE 2 SATADOM 128 GB MLC 1 Emulex LPE 31002 16 Gb/s 	Management server
Cisco Nexus 93180YC-EX switch	2	<ul style="list-style-type: none"> Two distinct VLANs, each dedicated to the Linux cluster network and SAP HANA inter-cluster network 	10 GbE SUSE HAE or RHEL HA cluster and inter-cluster network
	2	<ul style="list-style-type: none"> Two switches with one VLAN to provide uplink network to local area network infrastructure 	10 GbE client network
Cisco Nexus 3048 switch	2	<ul style="list-style-type: none"> 48 × 1 GbE ports 	1 GbE management network
Brocade 7810 switch	2	<ul style="list-style-type: none"> 12 × 32 Gb/s Fibre Channel ports 12 × 16 Gb/s Fibre Channel SFPs 6 × 1 GbE Ethernet ports 	FCoE switch for storage replication

Hitachi Advanced Server DS7000 Series

To take advantage of the latest developments in artificial intelligence (AI), data analytics and machine learning, you require an infrastructure with high reliability, extreme performance, and agile scalability. [Hitachi Advanced Server 7000 series servers](#) (PDF) deliver this with a unique modular architecture.

Your server can be configured and scaled to meet the needs of a wide variety of application workloads. This can be used from in-memory data analytics processing to virtualization and hybrid cloud.

The Advanced Server DS7000 series has several complementary models, each based on the Intel Xeon scalable processor. You can upgrade a model to the next model, preserving your hardware and software investment as you grow.

The DS7000 series is certified for SAP HANA. Use this server in either a tailored data center integration (TDI) solution or as part of a complete appliance.

- As part of a tailored data center integration (TDI) solution, each installation is customized by assembling hardware, operating system, and (optional) hypervisor from SAP-certified components.
- As part of an appliance deployment, the hardware vendor pre-assembles, pre-installs, and pre-configures hardware and software with a complete package certified by SAP for each installation.

These two methods of using a DS7000 series server protect your investment by providing a migration path for any architecture using SAP HANA, at any scale.

Hitachi Advanced Server DS120

Optimized for performance, high density, and power efficiency in a dual-processor server, [Hitachi Advanced Server DS120](#) delivers a balance of compute and storage capacity. This rack mounted server has the flexibility to power a wide range of solutions and applications.

The highly-scalable memory supports up to 3 TB using 24 slots of 2666 MHz DDR4 RDMM. DS120 is powered by the Intel Xeon scalable processor family for complex and demanding workloads. There are flexible OCP and PCIe I/O expansion card options available. This server supports up to 12 storage devices with up to 4 NVMe.

This solution each site uses one Hitachi Advanced Server DS120 as a management server for the scale-out appliance.

Hitachi Virtual Storage Platform

This solution uses one Hitachi Virtual Storage Platform G700 using solid state drives or Virtual Storage Platform F700 at each site.

Virtual Storage Platform G Series Family

The [Hitachi Virtual Storage Platform G series family](#) enables the seamless automation of the data center. It has a broad range of efficiency technologies that deliver maximum value while making ongoing costs more predictable. You can focus on strategic projects and to consolidate more workloads while using a wide range of media choices.

The benefits start with Hitachi Storage Virtualization Operating System RF. This includes an all new enhanced software stack that offers up to three times greater performance than our previous midrange models, even as data scales to petabytes.

Virtual Storage Platform G series offers support for containers to accelerate cloud-native application development. Provision storage in seconds, and provide persistent data availability, all the while being orchestrated by industry leading container platforms. Moved these workloads into an enterprise production environment seamlessly, saving money while reducing support and management costs.

Virtual Storage Platform F Series Family

Use [Hitachi Virtual Storage Platform F series family](#) storage for a flash-powered cloud platform for your mission critical applications. This storage meets demanding performance and uptime business needs. Extremely scalable, its 4.8 million random read IOPS allows you to consolidate more applications for more cost savings.

Hitachi Storage Virtualization Operating System RF is at the heart of the Virtual Storage Platform F series family. It provides storage virtualization, high availability, flash optimized performance, quality of service controls, and advanced data protection. This proven, mature software provides common features, management, and interoperability across the Hitachi portfolio. This means you can reduce migration efforts, consolidate assets, reclaim space, and extend life.

Reduce risks and solve problems faster. Integrated power analytics and automation features bring artificial intelligence to your data center. Cloud-accessible monitoring tools give your product support experts access wherever they have an internet connection for fast troubleshooting and remediation.

Data Center Switches

This solution uses the following switches to provide Ethernet connectivity and replication connectivity:

- Cisco Nexus 3048
- Cisco Nexus 93180YC-EX switch
- Brocade 7810 Extension Switch

[Cisco Nexus data center switches](#) are built for scale, industry-leading automation, programmability, and real-time visibility.

[Brocade 7810 extension switch](#) is purpose-built to handle modern workloads between data centers.

Software Elements

Table 2 describes the software products used to deploy two active node and one standby node configuration.

TABLE 2. SOFTWARE ELEMENTS

Software		Version
Hitachi Storage Navigator Modular 2		Microcode dependent
Hitachi Universal Replicator		Microcode dependent
Command control interface		01-49-03/00
Operating System ■ Used on Hitachi Advanced Server DS7040	SUSE Linux Enterprise Server for SAP Applications	SUSE Enterprise Linux 12 SP3
	Alternate: Red Hat Linux Enterprise	RHEL 7.4
SAP HANA		2.0 SPS03, Rev33 or later
Microsoft® Windows Server® 2016 ■ Used on Hitachi Advanced Server DS120		1607
Cisco Nexus 93180YC-EX switch		BIOS: 07.56 or later NXOS: 7.0 (3) I4 (2) or later

TABLE 2. SOFTWARE ELEMENTS (CONTINUED)

Software	Version
Cisco Nexus 3048 switch	6.0 (2) U3 (7) or later
Brocade 7810 switch	8.2.1 or later

Solution Design

The detailed design for using Hitachi Universal Replicator with this Hitachi Unified Compute Platform for the SAP HANA platform solution is based on specifications from SAP and is a 2+1 node configuration as an example that includes the following:

- “Hitachi Advanced Server DS7040 Architecture” on page 10
- “Storage Architecture” on page 10
- “Management Server” on page 11
- “Network Architecture” on page 12
- “Journal for Hitachi Universal Replicator” on page 12
- “Inter-site Configuration for Hitachi Universal Replicator” on page 14
- “Disaster Recovery and Replication Components” on page 15

Hitachi Advanced Server DS7040 Architecture

This solution uses three Hitachi Advance Server DS7040 each site. Each server has a 1.5 TB or 3 TB SAP HANA configuration with two active nodes plus one standby node.

Refer to *Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration using Hitachi Advanced Server DS7040 Architecture Reference Guide* for detailed information.

Storage Architecture

The central storage system for the SAP HANA scale-out cluster is a Hitachi Virtual Storage Platform G700 using solid state drives or Virtual Storage Platform F700.

The following reside on the storage for this solution:

- **Used for SAP HANA:**
 - Operating system LUNs
 - Linux cluster STONITH LUN
 - SAP HANA data LUNs
 - SAP HANA log LUNs
 - SAP HANA shared LUNs

- **Used for Hitachi Universal Replicator:**
 - Journal volume
- **Used for Hitachi Open Remote Copy Manager (HORCM):**
 - Command device

To properly size the storage for SAP HANA, refer to [Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration using Hitachi Advanced Server DS7040 Reference Architecture Guide](#).

At any time, the secondary site only has one live SAP HANA instance. The secondary site is for production or quality assurance. Normally the quality assurance instance is available for use at the secondary site. In case of a service outage on the primary site and failover occurs, the production instance will be available and the quality assurance instance need to be shutdown.

Use the server priority manager at the secondary site to do the following if you have quality assurance system:

- Designate the prioritized ports (replication) and non-prioritized ports (quality assurance).
- Set the upper limits and thresholds for the I/O activity of these ports to prevent low-priority activities from negatively affecting the high priority activities, such as replication.

Additional information is available in [Hitachi Virtual Storage Platform F Series and G Series Performance Guide](#) (PDF).

Management Server

This solution uses one Hitachi Advanced Server DS120 as the management server. This acts as a central device for managing SAP HANA on each site and running replication of the storage systems.

Figure 2 shows the management server network ports using Intel X557-T4 10GbE, 4-port mezzanine card.

- **Slot 1A Port 0** — Connect to the Cisco Nexus 3048 switch. It provides the 1 GbE management network to all components.
- **Slot 1A Port 1** — Connect to the uplink switch. It provides the 1 GbE network to the management server.

Figure 2



Connect the 10 GbE network ports to the two Cisco switches, Nexus-93180-A and Nexus-93180-B, to provide management access to the SAP HANA nodes from the management server.

The management server has the following additional components for running replication of the storage systems:

- One Emulex LightPluse 2-port 16 Gb/s Fibre Channel HBA on the PCIe slot

Table 3 shows the slots to place mezzanine and PCIe cards, including the Emulex LPE 31002 Fibre Channel HBA.

TABLE 3. MEZZANINE AND PCI-EX CARD SLOT POSITION ON HITACHI ADVANCED SERVER DS120

PCIe Slot	Description	Component
Default	BMC Port	Management Port for BMC-browser connection
Slot 3 (Blue)	PCI-Ex Slot	Emulex LPE 31002 M6 16 Gb/s
Slot 2 (Red)	PCI-Ex Slot	Intel XXV710 SFP28 10/25 GbE
Slot 1A (Green)	OCP Mezzanine Slot	Intel X557-T4 10 GbE, 4 ports

In addition to the software listed in *Hitachi Solution for the SAP HANA Platform in a Scale-Out Configuration Using Hitachi Advanced Server DS7040 Reference Architecture Guide*, install the following software on the management server if they do not exist:

- Command control interface (CCI)
- NTP server service

Network Architecture

Each site, the scale-out configuration of this SAP HANA solution requires the following four separate networks in order to run SAP HANA system:

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

Refer to *Hitachi Solution for the SAP HANA Platform in a Scale-out Configuration using Hitachi Advanced Server DS7040 Reference Architecture Guide* for detailed information.

Journal for Hitachi Universal Replicator

Create journal data synchronously with the updates to P-VOL to provide a copy of the data written to the P-VOL. It is managed at the primary and secondary sites to ensure the consistency of the primary and secondary volumes.

Use journals in Hitachi Universal Replicator to enable update sequence consistency maintained across a group of volumes. You need journal volumes on the primary and secondary storage systems.

The following information must be determined to properly size the journals on each site:

- The amount of changed data that the applications generate
- The maximum amount of time that the journals can accumulate updated data. This depends on your operation's recovery point objective (RPO)

Refer to *Hitachi Virtual Storage Platform F Series and G Series Universal Volume Manager User Guide* for sizing journal volumes.

This is the formula to calculate the **journal capacity**:

$$\text{Journal Capacity} = \text{Peak hour megabytes change} \times \text{Number of hours}$$

This solution uses the following for journal configuration:

- Registers two journal volumes in a journal in the primary storage system and in the secondary storage system. Use one journal volume in each system, but the second journal volume used as a reserve journal volume.
- As tested by Hitachi for a journal volume, with one RAID group configured as RAID-5 (3D+1P) using 1.92 TB SSD drives, an estimated throughput is 423 MB/s.
- For the base configuration of the solution, use at least two RAID groups configured as RAID-5 (3D+1P) with 1.92 TB SSD drives in a dynamic provisioning pool for use as a journal group.
- Journal sizing example:
 - Recovery point objective = 3.5 hours
 - Peak hour megabytes change = [423 MB/s] × [60 s/min] × [60 min/hour] = 1428 GB/hour
 - Basic journal volume size = 1428 GB/hour × 3.5 hours = 4998 GB

Table 4 shows the example of using the two parity groups configured as RAID-5 (3D+1P) with LDEVs to support the basic journal volume size of 4998 GB.

TABLE 4. LDEVs FOR USE WITH HUR_JOURNAL_GROUP DYNAMIC PROVISIONING POOL

Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size
1	RAID-5 (3D+1P) on 1.92 TB SSD drives	00:0B:01	Journal_1	2640 GB
		00:0B:02	Journal_2	2640 GB
2	RAID-5 (3D+1P) on 1.92 TB SSD drives	00:0B:03	Journal_3	2640 GB
		00:0B:04	Journal_4	2640 GB

Create a dynamic provisioning pool named HUR_Journal_Group and assign all the LDEVs created in Table 4 to this pool. Create two virtual volumes in Table 5, each with 5281 GB for use as two journal volumes.

TABLE 5. DYNAMIC PROVISIONING POOL IDS AND VIRTUAL VOLUME LDEV IDS

Dynamic Provisioning Pool ID	Dynamic Provisioning Pool Name	Virtual Volume Names	Virtual Volume LDEV ID	Virtual Volume Size	MPU Assignment
1	HUR_Journal_Group_1	Journal_VVOL_1	00:0C:01	5280.00 GB	MPU10
	HUR_Journal_Group_1	Journal_VVOL_2	00:0C:02	5280.00 GB	MPU20

During implementation of the Hitachi Universal Replicator service, Hitachi Vantara Global Services needs to right-size the journal volumes based on the actual workload, throughput, and RPO requirement. Then, the following needs to happen:

- Add additional disks to create a journal volume of RAID groups configured as RAID-5 (3D+1P) on 1.92 TB SSD drives.
- Create additional LDEVs and add them to the existing dynamic provisioning pool HUR_Journal_Group.
- Add RAID groups to the dynamic provisioning pool, create one virtual volume with the right size (for example, 5308 GB) to use as a journal volume, and second virtual volume to use as a reserve journal volume.

Inter-site Configuration for Hitachi Universal Replicator

Fibre Channel switches on either end of the replication network provides the additional credits necessary to overcome buffer shortages due to the network latency. This solution uses two Brocade 7810 Fibre Channel over IP (FCIP) SAN extension switches at each site for long distance replication.

If you have the SAN switches and long-distance amplifiers in the existing infrastructure, you can utilize the same infrastructure.

When new data exists in the P-VOL, the data is transferred through initiator ports and RCU (remote control unit) target ports at the primary and secondary systems. An initiator port in one system must be connected to an RCU target port in the other system.

Table 6 lists the target and initiator ports of the two storage systems at each site along with zoning alias configured on the SAN switches. Both sites need one zoning configuration.

TABLE 6. ZONING CONFIGURATION

Initiator Port	RCU Target Port	Zone Name
PrimA_Port3F	SecB_Port5F	PrimA_Port3F_SecB_Port5F
PrimA_Port4F	SecB_Port6F	PrimA_Port4F_SecB_Port6F
SecB_Port3F	PrimA_Port5F	SecB_Port3F_PrimA_Port5F
SecB_Port4F	PrimA_Port6F	SecB_Port4F_PrimA_Port6F

Table 7, “Hitachi Compute Rack Management Server Zoning,” on page 15 has the details on the zoning between the management server and VSP G700 or VSP F700.

- Create a single zone named MGMT_Pri_Sec.
- Add all four storage ports aliases and all the four management server Emulex port aliases in the MGMT_Pri_Sec zone.

With this configuration, command devices on the primary system and secondary system are accessible on the management servers at the primary site and secondary site.

TABLE 7. HITACHI COMPUTE RACK MANAGEMENT SERVER ZONING

Virtual Storage Platform Alias	Hitachi Compute Rack Management Server Alias	Zone Name
PrimA_Port1F	MGMT_PRI_Port0	MGMT_Pri_Sec
SecB_Port1F	MGMT_SEC_Port0	MGMT_Pri_Sec
PrimA_Port2F	MGMT_PRI_Port1	MGMT_Pri_Sec
SecB_Port2F	MGMT_SEC_Port1	MGMT_Pri_Sec

Disaster Recovery and Replication Components

In this reference architecture for SAP HANA disaster recovery, Hitachi Universal Replicator setup requires doing the following:

- “Install Command Control Interface (CCI)” on page 15
- “Configure Command Devices” on page 15
- “Configure Replication using Command Control Interface” on page 16
- “Setup Hitachi Universal Replicator” on page 19

Install Command Control Interface (CCI)

Command control interface enables you to perform storage system operations by issuing commands to Hitachi Virtual Storage Platform G700 using solid state drives or Virtual Storage Platform F700.

To implement this solution, install command control interface on the management server for the primary site and the secondary site. Command control interface has the following components residing in these locations:

- **Storage System** — Command devices and Hitachi Universal Replicator volumes (P-VOLs and S-VOLS)
- **DS120 Management Server** — Hitachi Open Remote Copy Manager (HORCM)

For detailed information about command control interface, see *Command Control Interface User and Reference Guide*.

Configure Command Devices

A command device is a dedicated logical volume on the storage system that functions as the interface to the storage system from the host. The command device accepts commands from the host that are executed on the storage system.

To implement this solution, create a 100 MB command device logical volume on the local and remote storage system. Each management server has one dual-port Emulex HBA card, connected through the Brocade 7810 switches to storage. Perform zoning and add the LUN path for the command devices in such a way that the primary site management server and the secondary site management server can access the command devices on both storage systems.

From each management server, enable and configure MultiPath IO (MPIO) using Server Manager in Microsoft Windows Server for the storage systems.

Configure Replication using Command Control Interface

A key aspect of this reference architecture on Hitachi Virtual Storage Platform G700 or Virtual Storage Platform F700 is defining the volume pair relationship for replication between storage systems. Define and manage storage replication relationships through the Hitachi Storage Navigator graphical user interface or on the management server running Hitachi Open Remote Copy Manager (HORCM).

Hitachi Open Remote Copy Manager operates as a daemon process on the host. When activated, Open Remote Copy Manager refers to its configuration files. The instance communicates with the storage sub-system and remote servers.

To implement this solution, install two instances Open Remote Copy Manager to operate Hitachi Universal Replicator.

- Install a Hitachi Open Remote Copy Manager instance on the primary management server that manages the P-VOLs.
- Install a Hitachi Open Remote Copy Manager instance on the secondary management server that manages the S-VOLs.

The Open Remote Copy Manager configuration file defines the communication path and the logical units to be controlled. Each instance has its own configuration file. The configuration file lists the following for replication:

- SAP HANA data volumes
- SAP HANA log volumes
- SAP HANA shared volumes

Figure 3 on page 17 shows an example of the configuration file (horcm04.conf) used by the Hitachi Open Remote Copy Manager instance on the primary management server.

Figure 3

```

horcm04 - Notepad
File Edit Format View Help
# data or be part of an LDM Volume Group.
# In order to prevent user access, do not assign a drive letter to
# a command device.
#
#/****** For HORCM_MON *****/
HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
172.17.171.127   horcm04         1000             3000

#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group      dev_name        Serial#      CU:LDEV(LDEV#)      MU#
FANAHR SKL_SO_HANA_Shared_VVOL_N1-4_1 415051 00:65 0
FANAHR SKL_SO_HANA_Shared_VVOL_N1-4_2 415051 00:66 0
FANAHR SKL_SO_HANA_Shared_VVOL_N1-4_3 415051 00:67 0
FANAHR SKL_SO_HANA_Shared_VVOL_N1-4_4 415051 00:68 0
FANAHR SKL_SO_Log_VVOL_N1 415051 01:01 0
FANAHR SKL_SO_Log_VVOL_N2 415051 02:01 0
FANAHR SKL_SO_Log_VVOL_N3 415051 03:01 0
FANAHR SKL_SO_Data_VVOL_N1_1 415051 01:02 0
FANAHR SKL_SO_Data_VVOL_N1_2 415051 01:03 0
FANAHR SKL_SO_Data_VVOL_N1_3 415051 01:04 0
FANAHR SKL_SO_Data_VVOL_N1_4 415051 01:05 0
FANAHR SKL_SO_Data_VVOL_N2_1 415051 02:02 0
FANAHR SKL_SO_Data_VVOL_N2_2 415051 02:03 0
FANAHR SKL_SO_Data_VVOL_N2_3 415051 02:04 0
FANAHR SKL_SO_Data_VVOL_N2_4 415051 02:05 0
FANAHR SKL_SO_Data_VVOL_N3_1 415051 03:02 0
FANAHR SKL_SO_Data_VVOL_N3_2 415051 03:03 0
FANAHR SKL_SO_Data_VVOL_N3_3 415051 03:04 0
FANAHR SKL_SO_Data_VVOL_N3_4 415051 03:05 0

#/****** For HORCM_INST *****/
HORCM_INST
#dev_group      ip address      service

```

Figure 4 shows an example of the configuration file (horcm06.conf) files used by the Hitachi Open Remote Copy Manager instance on the secondary management server.

Figure 4

```

horcm06 - Notepad
File Edit Format View Help
# data or be part of an LDM Volume Group.
# In order to prevent user access, do not assign a drive letter to
# a command device.
#
#/****** For HORCM_MON *****/
HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
172.17.171.147   horcm06          1000             3000
#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name        dev_name          dev_name
\\.\PHYSICALDRIVE1
#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group      dev_name          Serial#    CU:LDEV(LDEV#)  MU#
HANA_HUR SKL_SO_HANA_Shared_VVOL_N1-4_1 415012 01:06 0
HANA_HUR SKL_SO_HANA_Shared_VVOL_N1-4_2 415012 01:07 0
HANA_HUR SKL_SO_HANA_Shared_VVOL_N1-4_3 415012 01:08 0
HANA_HUR SKL_SO_HANA_Shared_VVOL_N1-4_4 415012 01:09 0
HANA_HUR SKL_SO_Log_VVOL_N1 415012 01:0A 0
HANA_HUR SKL_SO_Log_VVOL_N2 415012 02:0A 0
HANA_HUR SKL_SO_Log_VVOL_N3 415012 03:0A 0
HANA_HUR SKL_SO_Data_VVOL_N1_1 415012 01:0B 0
HANA_HUR SKL_SO_Data_VVOL_N1_2 415012 01:0C 0
HANA_HUR SKL_SO_Data_VVOL_N1_3 415012 01:0D 0
HANA_HUR SKL_SO_Data_VVOL_N1_4 415012 01:0E 0
HANA_HUR SKL_SO_Data_VVOL_N2_1 415012 02:0B 0
HANA_HUR SKL_SO_Data_VVOL_N2_2 415012 02:0C 0
HANA_HUR SKL_SO_Data_VVOL_N2_3 415012 02:0D 0
HANA_HUR SKL_SO_Data_VVOL_N2_4 415012 02:0E 0
HANA_HUR SKL_SO_Data_VVOL_N3_1 415012 03:0B 0
HANA_HUR SKL_SO_Data_VVOL_N3_2 415012 03:0C 0
HANA_HUR SKL_SO_Data_VVOL_N3_3 415012 03:0D 0
HANA_HUR SKL_SO_Data_VVOL_N3_4 415012 03:0E 0
#/****** For HORCM_INST *****/
HORCM_INST
#dev_group      ip_address      service
HANA_HUR 172.17.171.127 horcm04

```

Figure 5 lists the entries that need for adding to the services file of the management server on both sides for Open Remote Copy Manager to function.

Figure 5

```

horcm04          11004/udp          #Primary HORCM
horcm06          11006/udp          #Secondary HORCM

```

Setup Hitachi Universal Replicator

To implement this solution, set up Hitachi Universal Replicator as follows:

- For the failover to the secondary site, use two initiator ports on the storage at the primary site connected to two RCU target ports on the storage at the secondary site.
- For the failback to the primary site, use two initiator ports from the storage system at the secondary site connected to the two RCU target ports on the storage system at the primary site.
- On Site A, connect the initiator and RCU target ports on storage to the Brocade 7810 Fibre Channel switches.
- On Site B, connect the initiator and RCU target ports on storage to the Brocade 7810 Fibre channel switches.

Define the port attributes for initiator and target ports on the storage as listed in Table 8.

TABLE 8. PORT ATTRIBUTES FOR INITIATOR AND TARGET PORTS

Ports	Port Security	Port Speed	Fabric	Connection Type
Initiator Ports — 3F and 4F	Enabled	16 Gb/s	ON	P-to-P
Target Ports— 5F and 6F	Enabled	16 Gb/s	ON	P-to-P
Management server ports — 1F and 2F	Enabled	16 Gb/s	ON	P-to-P

Configure Hitachi Universal Replicator

Configure Hitachi Universal Replicator in Hitachi Unified Compute Platform for SAP HANA with these steps:

1. Define the Fibre Channel port attributes.

Define the port attributes for ports listed in Table 8 on Virtual Storage Platform G700 or Virtual Storage Platform F700.

2. Setup the remote connections between the primary system and the secondary system.

Associate the primary system and the secondary system in the Hitachi Universal Replicator relationships and add remote connections using the logical paths between them. This is required to perform the Universal Replicator pair operations.

3. Create parity groups and LDEVs for journal groups.

Create journal group RAID groups and LDEVs for use as journal volumes, as described in Journal for Hitachi Universal Replicator.

4. Create and register journal volumes in a journal.

Register every journal volume in the journal group.

5. Start the Hitachi Open Remote Copy Manager Instance

To start up the Hitachi Open Remote Copy Manager instance on the primary management server, issue the following example HORCM command ("04" refers to HORCM instance number):

```
# C:\horcm\etc\horcmstart 04
```

To start up the Open Remote Copy Manager instance on the secondary management server, issue the following example HORCM command ("06" refers to HORCM instance number):

```
# C:\horcm\etc\horcmstart 06
```

See the detailed implementation procedures in *Hitachi Virtual Storage Platform F Series and G Series Universal Volume Manager User Guide*, and *Hitachi Virtual Storage Platform G130, G/F350, G/F370, G/F700, G/F900 Provisioning Guide*. Contact your Hitachi Data Systems representative if you need these documents.

Initial Pair Copy Operation

Perform the initial data transfer, called the pair copy operation, between Site A and Site B by running the HORCM paircreate command on Site A. This is an example command:

```
# C:\horcm\etc\paircreate -IH04 -g HANADR -vl -f async -jp 0 -js 0
```

Executing the paircreate command starts the initial copy. The primary storage system copies all the data in sequence from the P-VOL directly to the S-VOL.

During the initial copy process, the status of the P-VOL and S-VOL is COPY. On completion of the initial copy process, the status of the P-VOL and S-VOL changes to PAIR.

Disaster Recovery Operations

To perform disaster recovery operations for the SAP HANA production site, use command control interface. This is for either a planned or unplanned failover and failback.

Note — A real disaster recovery event requires you to do additional activities than the following tested workflows, such as client connection recovery.

These are the tested workflows for a planned failover and failback.

■ To perform failover from Site A to Site B.

- (1) Ensure the pair is synchronized completely.
- (2) Shut down production on Site A.
- (3) On Site B, do the following:
 - i. If it exists, shut down the quality assurance or other non-production instance or instances.
 - ii. Run the command control interface pair split command with read and write.
 - iii. Check at the operating system level of the SAP HANA nodes if the /hana/shared file system is mounted.
 - iv. Start the SAP HANA system and validate if it successfully comes up.

■ **To perform failback from Site B to Site A.**

- (1) Shut down the production SAP HANA database at Site B.
- (2) Run the command control interface pair resync command, changing the S-VOL to P-VOL. Once the resync is 100% completed, start the failback.
- (3) At Site A, do the following:
 - i. Run the command control interface pair split command with read and write.
 - ii. Start the SAP HANA system and validate if it successfully comes up.
 - iii. Once validation completed, run the command control interface pair resync command, changing the S-VOL to P-VOL.

These are the tested workflows for an unplanned failover and failback.

■ **To perform failover from Site A to Site B.**

- (1) Ensure the pair is synchronized completely.
- (2) From Site B, do the following:
 - i. If it exists, shut down the quality assurance or other non-production instance or instances.
 - ii. Run the command control interface horctakeover command to switch operations between Site A and Site B.
 - iii. Check at the operating system level of the SAP HANA nodes if the /hana/shared file system is mounted.
 - iv. Start the SAP HANA system and validate if it successfully comes up.

■ **To perform failback from Site B to Site A.**

- (1) Shut down the production SAP HANA database at Site B.
- (2) Do a pair resync by changing the S-VOL to P-VOL. Once it is 100% completed, start the failback.
- (3) At Site A, do the following:
 - i. Run the command control interface pair split command with read and write.
 - ii. Start the SAP HANA system and validate if it successfully comes up.
 - iii. Once validation completed, run the command control interface pair resync command, with changing S-VOL to P-VOL.

For detailed information about command control interface commands and options, see *Command Control Interface User and Reference Guide*.

Engineering Validation

Validation of this reference architecture was conducted in the Hitachi Vantara laboratory. The procedures to failover to Site B and then to failback to Site A with soft and hard techniques using Hitachi Universal Replicator were tested.

Test Methodology

To test the setup, the following scenarios were executed in the lab:

- Planned failover to Site B
- Planned failback to Site A
- Planned failover to Site B during mid-flush and then failback to Site A
- Unplanned failover to Site B

Test Results

All the tests passed without issues.

For More Information

Hitachi Vantara Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the [Services](#) website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the [Resources](#) website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara [Training and Certification](#) website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the [Hitachi Vantara](#) website.

Hitachi Vantara



Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HitachiVantara.com | community.HitachiVantara.com

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

© Hitachi Vantara Corporation 2019. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd., VSP is a trademark or registered trademark of Hitachi Vantara Corporation. Microsoft and Windows Server are trademarks or registered trademarks of Microsoft Corporation. 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 Vantara.

MK-SL-141-00, March 2019.