

Disaster Recovery using Hitachi TrueCopy on Hitachi Unified Compute Platform for the SAP HANA Platform with Hitachi Compute Blade 2500, Hitachi Virtual Storage Platform G800, and Hitachi NAS Platform 4060

Reference Architecture Guide

By Abhishek Dhanuka

October 2016

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

Solution Overview	3
Key Solution Components.....	6
Hardware Elements.....	6
Software Elements.....	9
Solution Design.....	11
Storage Architecture	12
Inter-site Configuration	20
Data Replication and Disaster Recovery Component Configuration.....	20
Engineering Validation.....	26
Test Methodology	26
Test Results	26

Disaster Recovery using Hitachi TrueCopy on Hitachi Unified Compute Platform for the SAP HANA Platform with Hitachi Compute Blade 2500, Hitachi Virtual Storage Platform G800, and Hitachi NAS Platform 4060

Reference Architecture Guide

Use this reference architecture guide to determine the necessary components and overall layout for a disaster recovery solution. This environment supporting real time data processing uses Hitachi TrueCopy on Hitachi Unified Compute Platform for the SAP HANA platform in a scale-out configuration. This configuration for backup and recovery of SAP HANA environments uses the following:

- Hitachi Compute Blade 2500 (CB 2500) with 520X B2 or 520X B3 server blades
- A rack optimized server for solutions, 2U four nodes, using two of the four nodes
- Hitachi Virtual Storage Platform G800 (VSP G800)
- Hitachi NAS Platform 4060 (HNAS 4060)
- Hitachi TrueCopy
- SAP HANA
- Either of these operating systems:
 - SUSE Linux Enterprise Server
 - Red Hat Enterprise Linux

Solution testing was only on a 2+1 configuration with 1.5 TB SAP HANA nodes using 520X B2 server blades. However, this reference architecture supports the SAP HANA deployment of the scale-out configurations found in each of the following:

- [Hitachi Unified Compute Platform for the SAP HANA Platform using 1.5 TB SAP HANA Nodes in a Scale-Out 24 TB Configuration of 16 Active Nodes and 3 Standby Nodes with Hitachi Compute Blade 2500 Chassis, 520X B2 Server Blades, and Hitachi Virtual Storage Platform G800](#) (AS-437-02 or later, PDF)
- [Hitachi Unified Compute Platform 6000 for the SAP HANA Platform, Scale-out Configuration using 2 TB or 4TB SAP HANA Nodes with Intel Xeon E7-88xx v4 Processors](#) (AS-493-01 or later, PDF)

For help designing your implementation, contact your GSS representative.

The scale-out environment for Unified Compute Platform for SAP HANA is a preconfigured analytical appliance that provides real-time access to operational data for use in analytic models. Changes to this architecture require approval from the following:

- Sales
- Solution Engineering and Technical Operations
- Solution product management

Failure of SAP HANA may result in revenue loss. For protection from this loss, use two sites in the disaster recovery strategy. In addition to failover production, the second site can handle the quality assurance environment of the SAP HANA landscape.

The primary business problem this solution answers is disaster recovery for SAP HANA. This solution performs synchronous replication of SAP HANA data volumes and log volumes on Hitachi Virtual Storage Platform G800 to the secondary site. It also performs synchronous replication of the SAP HANA binaries and other configuration files stored in the /hana/shared file system on Hitachi NAS Platform 4060 to the secondary site.

Data centers at each site must have almost identical hardware for this disaster recovery solution. Implementing Hitachi Universal Replicator on this environment permits the additional use of the secondary site as a quality assurance environment. This additional use requires adding additional disk drives on the secondary site storage to run the quality assurance system.

Hitachi Virtual Storage Platform G800 technology permits maintaining sufficient performance for the SAP HANA production instance and site-to-site data replication.

This technical paper assumes you have familiarity with the following:

- Storage area network-based storage systems
- Network attached storage (NAS) systems
- General storage concepts
- General storage replication skills and concepts
- General network and virtual IP knowledge
- General WAN knowledge
- Advanced SAP HANA skills
- Disaster recovery scenarios
- Common IT storage practices

Note — Testing of this configuration occurred in the Hitachi 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 A contains a production SAP HANA database instance. Implement this site as a scale-out configuration of Hitachi Unified Compute Platform for the SAP HANA platform.

The secondary site B is an exact replica of primary site, with the exception of optional additional storage disks. Site B houses the following:

- Failover production instance
- (Optional) Non-production instances, such as the following that requires additional disks:
 - 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 Compute Blade 2500 nodes for production and non-production instances.

Hitachi Virtual Storage Platform G800 at the secondary site houses two sets of disks for data, log, and /hana/shared LUNs for the following:

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

In this solution, install command control interface on the management servers on the primary site and secondary site. This performs the data replication operations within Hitachi Virtual Storage Platform G800 at each site using Hitachi Open Remote Copy Manager instances. Each instance at the primary site and secondary site management server has its own Open Remote Copy Manager configuration file that lists the following for replication between two sites:

- SAP HANA data volumes
- SAP HANA log volumes
- Hitachi NAS Platform LUNs

Hitachi Virtual Storage Platform G800 works with Hitachi TrueCopy to enable SAP HANA disaster recovery. In this solution, use a Fibre Channel over IP (FCIP) switch for the wide area network connections between Hitachi Virtual Storage Platform G800 at primary site A and secondary site B.

For the synchronous replication of LUNs on Hitachi NAS Platform hosting the /hana/shared file system on Hitachi Virtual Storage Platform G800, perform a one-time initial configuration after completing the initial pair copy operations to configure the RAID mirror relationships between the LUNs on Hitachi NAS Platform at both sites. The Hitachi NAS Platform clusters at both the sites need to be made aware of the TrueCopy LUN relationship between them during this initial configuration by using the RAID mirroring (`sd-mirror-remotely`) command.

In case of an outage or any component failure in the primary site, an administrator initiates a manual failover to the secondary site using customized scripts. There are two different possibilities for enabling client connection recovery, either using virtual IP failover or DNS failover configuration. However, the actual implementation differs, based on the network and cluster management capabilities.

This solution supports four different disaster recovery options described below.

1. Add a Disaster Recovery Site

This option is for a primary site for production and a secondary site for disaster recovery.

2. Add a Quality Assurance or Development Instance

This option is for a primary site for production and a secondary site for disaster recovery site and a single quality assurance or development SAP HANA instance.

3. Disaster Recovery Connectivity Bundle

This option is for a Fibre Channel switch between two sites.

4. Only Disaster Recovery Site

This option is for adding a secondary site for a disaster recovery solution to an existing SAP HANA landscape. (The primary site for production already exists.)

To perform data replication with Hitachi Unified Compute Platform for SAP HANA, the reference solution uses the following:

- **Hitachi Compute Blade 2500 with 520X B2 or 520X B3 Server Blades**

- **Rack optimized server for solutions, 2U four nodes**

An ultra-dense design equipped with four independent nodes, the server creates the flexibility to set up different workloads independently in one 2U-shared infrastructure, providing optimal data center performance. This solution uses only two nodes out of the four available nodes.

- **Hitachi Virtual Storage Platform G800**

Virtual Storage Platform G800 is designed to manage storage assets more efficiently

- **Hitachi NAS Platform 4060**

NAS Platform 4060 is a network-attached storage solution used for file sharing, file server consolidation, data protection, and business-critical NAS workloads.

- **Hitachi TrueCopy**

For synchronous replication up to 43 miles (70 km), Hitachi TrueCopy provides no data-loss and rapid restart solution with real-time copies that are the same as the originals, reducing recovery time to minutes.

- **SAP HANA**

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

- **Brocade VDX 6740-48 switch**

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

- **Brocade ICX 6430-48 switch**

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

- **Brocade ICX 6430-24 switch**

This 24-port, 1 GbE switch provides the NAS Platform private network.

- **Brocade 6505 Extension switch**

A 24-port, entry-level switch that provides Fibre Channel connectivity between Hitachi Unified Storage on the primary site and the secondary site.

Figure 1 shows the configuration of this solution.




Figure 1

Key Solution Components

These are the key solution components in this reference architecture.

Hardware Elements

Table 1 describes the hardware needed to deploy the two active nodes and one standby node configuration at each site using the 520X B2 server blade in Hitachi Compute Blade 2500. All configurations for Hitachi Unified Compute Platform for the SAP HANA platform support using Hitachi TrueCopy.

Table 1. Hardware Elements

Hardware	Quantity	Configuration	Role
Hitachi Compute Blade 2500	2	<ul style="list-style-type: none"> ▪ 8 server blade chassis ▪ 2 management modules ▪ 10 cooling fan modules ▪ 5 power supply modules ▪ 2 FIVE-FX 16 Gb/sec 2-port Fibre Channel PCI-Ex card for each SAP HANA node ▪ 4 × 2-port 10GBASE-SR PCI-Ex card for each SAP HANA node 	Server blade chassis <ul style="list-style-type: none"> ▪ One chassis is required for three SAP HANA nodes.
520X B2 server blade	12	<ul style="list-style-type: none"> ▪ 2 × 18-core processors ▪ 768 GB RAM ▪ 2 pass-through mezzanine cards on Mezzanine Slot 2 and Slot 4 of each server blade 	Server blade for the SAP HANA nodes <ul style="list-style-type: none"> ▪ Two server blades are required for each SAP HANA node.
SMP connection board for 520X B2 server blade	6	<ul style="list-style-type: none"> ▪ 2 server blade SMP connection board ▪ SMP expansion module ▪ SMP connector cover 	SMP connector to turn two physical server blades into one SAP HANA node with 4 × 18-core processor and 1.5 TB of memory <ul style="list-style-type: none"> ▪ One SMP connector is required for each SAP HANA node.
Hitachi NAS Platform 4060	4	For every NAS Platform server <ul style="list-style-type: none"> ▪ 2 cluster ports ▪ 2 × 10 GbE ports ▪ 2 Fibre Channel ports ▪ 2 Ethernet ports 	Provide NFS shared file system for SAP HANA binaries, cluster-wide configuration files, and for storing one in-memory backup

Table 1. Hardware Elements (Continued)

Hardware	Quantity	Configuration	Role
Hitachi Virtual Storage Platform G800	2	<ul style="list-style-type: none"> ▪ CTL – 1 pair ▪ DKB (BED) – 2 pairs ▪ CHB (FED) – 3 pairs 	Block storage for SAP HANA nodes and NAS Platform
Rack optimized server for solutions, 2U four nodes	2	<ul style="list-style-type: none"> ▪ 2U, 2.5 inch bay chassis ▪ 2 server nodes, with each node having the following components: <ul style="list-style-type: none"> ▪ 2 Intel Xeon E5-2620 v3 processors, 2.4 GHz CPU, 32 GB RAM ▪ 2 × 500 GB 7200 RPM SATA drives ▪ 1 dual port 10 GbE Intel 82599ES SFP+OCP mezzanine card ▪ 1 dual port 1 GbE Base-T Intel i350 Mezzanine Card ▪ 1 Emulex Dual Port 8 Gb/sec Fibre Channel HBA ▪ 2 server filler 	<p>This server hosts the following:</p> <ul style="list-style-type: none"> ▪ Virtual SMU (vSMU) ▪ Management server running the following: <ul style="list-style-type: none"> ▪ NTP ▪ Hitachi Command Suite ▪ Hi-Track Remote Monitoring system ▪ SAP HANA Studio ▪ Command control interface
Brocade VDX 6740-48 port switch	8	<ul style="list-style-type: none"> ▪ Two distinct VLANs, each dedicated to NFS and SAP HANA intra-cluster network ▪ Two switches with one VLAN to provide uplink network to customer network infrastructure 	<p>10 GbE NFS and intra-cluster network</p> <p>10 GbE client network</p>
Brocade ICX 6430-48 port switch	2	<ul style="list-style-type: none"> ▪ 1 GbE ▪ 48 ports 	1 GbE Management Network
Brocade ICX 6430-24 port switch	2	<ul style="list-style-type: none"> ▪ 1 GbE ▪ 24 ports 	NAS Platform private network
Brocade 6505 switch	4	<ul style="list-style-type: none"> ▪ 24 port switch ▪ Fibre Channel switch 	Fibre Channel switch for storage connections

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 520X B2 or 520X B3 server blades in the Hitachi Compute Blade 2500 chassis. Table 2 shows the configuration for 520X B2 server blades.

Table 2. 520X B2 Server Blade Configuration

Feature	Configuration
Processors	<ul style="list-style-type: none"> ■ Intel Xeon E7-8880 v3 processor ■ 2 processors per server blade
Processor frequency	<ul style="list-style-type: none"> ■ 2.3 GHz
Processor cores	<ul style="list-style-type: none"> ■ 18 cores
Memory DIMM slots	<ul style="list-style-type: none"> ■ 48 slots populated
Memory	<ul style="list-style-type: none"> ■ 768 GB RAM ■ 16 GB DIMMs
Network ports	<ul style="list-style-type: none"> ■ 1 × USB 3.0 port ■ KVM connector (VGA, COM, USB 2.0 2-port)

Hitachi NAS Platform 4060

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

- Powerful hardware-accelerated file system with multi-protocol file services, dynamic provisioning, intelligent tiering, virtualization, and cloud infrastructure
- Seamless integration with Hitachi SAN storage, [Hitachi Command Suite](#), and [Hitachi Data Discovery Suite](#) for advanced search and index
- Integration with [Hitachi Content Platform](#) for active archiving, regulatory compliance, and large object storage for cloud infrastructure

This solution uses NAS Platform 4060 servers for file system sharing of the global binary and configuration SAP HANA files. There are four NAS Platform 4060 server nodes.

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

Hitachi Virtual Storage Platform 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. New management software improves ease of use to save time and reduce complexity. The infrastructure of Storage Virtualization Operating System creates a management framework for improved IT response to business demands.

This solution uses two Hitachi Virtual Storage Platform G800 storage platforms. The operating system LUNs, data LUNs, log LUNs, and LUNs for the Hitachi NAS Platform cluster reside on this storage array.

At any time, the secondary site only has one live SAP HANA instance. The secondary site is for disaster recovery or quality assurance. Normally, the quality assurance instance at the secondary site becomes the production instance in case of a service outage at the primary site.

Use the server priority manager at the secondary site to do the following:

- Designate the prioritized ports (replication) and non-prioritized ports (quality assurance).
- Set the upper limits and thresholds following best practices 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 the [Hitachi Virtual Storage Platform G800 Performance Guide](#).

Software Elements

This solution uses these software products.

Operating System Options

Use one of the following options for the operating system when implementing this solution:

- **SUSE Linux Enterprise Server (SLES)**

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

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

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

[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.

Changing the configuration settings is only supported when following guidelines from SAP and the operating system distributor. Otherwise, you could have significant performance problems. The following SAP Notes for are a good starting point for information on this topic:

- [1944799 – SAP HANA Guidelines for SLES Operating System Installation](#)
- [2009879 – SAP HANA Guidelines for Red Hat Enterprise Linux \(RHEL\) Operating System](#)

For more details, refer to “Updating and Patching the Operating System,” in the [SAP HANA Technical Operations Manual](#).

SAP HANA

[SAP HANA](#) converges database and application platform capabilities in-memory to transform transactions, analytics, text analysis, predictive and spatial processing so businesses can operate in real-time. This combines database, data processing, and application platform capabilities in a single in-memory platform. Also, the platform provides libraries for predictive, planning, text processing, spatial, and business analytics – all on the same architecture.

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

SAP customers can download more information on the SAP HANA platform at the [SAP Service Marketplace](#). See the installation and upgrade guides download section for SAP In-Memory Computing (SAP In-Memory Appliance – SAP HANA). The following are available:

- [SAP HANA Master Guide](#)

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

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

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

- [SAP HANA Technical Operations Manual](#)

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

- [SAP HANA Master Update Guide](#)

This explains how to update SAP HANA and its components.

- **SAP Integration and Certification Center (SAP ICC)**

This page provides information about SAP HANA appliances certified by SAP hardware partners.

Find all [SAP-related documentation](#) on the SAP website.

Solution Design

The detailed design for this solution using Hitachi TrueCopy with Hitachi Unified Compute Platform for SAP HANA in a scale-out configuration is based on specifications from SAP. It is a 2+1 node configuration that includes the following topics:

- Hitachi Compute Blade 2500 Chassis Configuration
- 520X B2 or 520X B3 Server Blade Architecture
- Fibre Channel SAN Architecture
- Storage Architecture
- Hitachi NAS Platform 4060 Architecture
- Network File System Design for Shared Binaries
- Management Server
- Network Architecture
- SAP Storage Connector API Fibre Channel Client
- Inter-site Configuration
- Disaster Recovery and Replication Components

This paper only discusses these topics:

- “Storage Architecture,” starting on page 12
- “Inter-site Configuration,” starting on page 20
- “Data Replication and Disaster Recovery Component Configuration,” starting on page 20

For detailed information on the topics not included in this reference architecture guide, see the following:

- **520X B2 server blade installation** — [Hitachi Unified Compute Platform for the SAP HANA Platform using 1.5 TB or 2 TB SAP HANA Nodes in a Scale-Out 24 TB or 32 TB Configuration of 16 Active Nodes and 3 Standby Nodes with Hitachi Compute Blade 2500 Chassis, 520X B2 Server Blades, and Hitachi Virtual Storage Platform G800 Reference Architecture Guide](#) (AS-437-02 or later, PDF)
- **520X B3 server blade installation** — [Hitachi Unified Compute Platform 6000 for the SAP HANA Platform, Scale-out Configuration using 2 TB or 4TB SAP HANA Nodes with Intel Xeon E7-88xx v4 Processors](#) (AS-493-01 or later, PDF)

Storage Architecture

The central storage system for the scale-out configuration cluster of this Hitachi Unified Compute Platform for SAP HANA is a Hitachi Virtual Storage Platform G800 storage platform. Divide the space on Virtual Storage Platform G800 among the following purposes:

- Boot device provisioning for SAP HANA nodes
- Log device provisioning for SAP HANA database
- Data device provisioning for SAP HANA database
- Block storage provisioning for the shared file system on Hitachi NAS Platform to store the SAP HANA binaries and cluster-wide configuration files

Figure 2 on page 14 shows a generic RAID group storage configuration of a building block of four active nodes used on similar appliances.

When active, every SAP HANA node has its own data volume and log volume. A standby node does not require a data volume and a log volume.

This design is based on a four node building block approach for the SAP HANA data volumes, log volumes, and shared binaries. Provision the parity groups shown in Table 3, starting on page 14, as follows.

- **Operating System LUN**
 - Each node has its own 100 GB LUN on Virtual Storage Platform G800 for the operating system volumes.
 - A single parity group configured as RAID-6 (6D + 2P) on 8 × 600 GB drives provisions the operating system LUN for SAP HANA nodes 1 to 19 on Virtual Storage Platform G800.
 - From this parity group, create 19 LDEVs, each with a capacity of 100 GB.
 - Map each LDEV exclusively to the corresponding SAP HANA node as follows: LUN number 00.
 - The installation of SUSE Linux Enterprise Server for SAP Applications or Red Hat Enterprise Linux resides on this LUN.
- **Hitachi NAS Platform 4060 Block Storage**
 - The block storage for Hitachi NAS Platform consists of two parity groups on Virtual Storage Platform G800, configured as RAID-6 (6D+2P) on 16 × 600 GB drives to store the shared binaries and configuration files of the SAP HANA database.
 - In each of the two parity groups, create two LDEVs of 1600 GB each.
 - Create a dynamic provisioning pool named HNAS_HDP_pool. Assign all the created Hitachi NAS Platform LDEVs to this pool. This allows the use of all the disks concurrently on the NAS Platform for better performance.
 - For a four node building block, create 4 virtual volumes, each with 1600 GB in HNAS_HDP_pool. Complete the LUN path assignment for these virtual volumes to the ports that are connected to Hitachi NAS Platform.

■ SAP HANA Log Volumes

- For the SAP HANA log volumes, first create two parity groups configured as RAID-6 (6D+2P) on 16 × 600 GB drives.
- In each of the two parity groups, create two LDEVs at 600 GB each.
- Map each SAP HANA log volume to all SAP HANA nodes at each port with the LUN ID of the specified host.

■ SAP HANA Data Volumes

- For the SAP HANA data volumes, create three parity groups configured as RAID-6 (14D+2P) on 48 × 600 GB drives.
- Create four LDEVs with a capacity of 1800 GB per each parity group. Table 3 on page 14 shows the parity groups and LDEVs created for data volumes.
- Assign three LDEVs for use as data volumes to each SAP HANA node, as shown in Table 5 on page 16.



Figure 2

Table 3 shows the parity groups and LDEV assignment for boot volumes, the Hitachi NAS Platform volumes, the SAP HANA log volumes, and the SAP HANA data volumes of the production system on both the sites.

Table 3. Groups and LDEV Assignment for Operating System Boot, Hitachi NAS Platform, SAP HANA Log Volumes, and SAP HANA Data Volumes on Both Sites

Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
1	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:01:00	HANA_Boot_LUN_N1	100 GB	MPU-10
		00:02:00	HANA_Boot_LUN_N2	100 GB	MPU-11
		00:03:00	HANA_Boot_LUN_N3	100 GB	MPU-20
		00:04:00	HANA_Boot_LUN_N4	100 GB	MPU-21

Table 3. Groups and LDEV Assignment for Operating System Boot, Hitachi NAS Platform, SAP HANA Log Volumes, and SAP HANA Data Volumes on Both Sites (Continued)

Parity Group ID	Parity Group RAID Level and Disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
2	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:00:01	HNAS_VOL_1	1600 GB	MPU-10
		00:00:02	HNAS_VOL_2	1600 GB	MPU-11
3	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:00:03	HNAS_VOL_3	1600 GB	MPU-20
		00:00:04	HNAS_VOL_4	1600 GB	MPU-21
4	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:01:01	HANA_LOG_N1	600 GB	MPU-21
		00:02:01	HANA_LOG_N2	600 GB	MPU-20
5	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:03:01	HANA_LOG_N3	600 GB	MPU-11
		00:04:01	HANA_LOG_N4	600 GB	MPU-10
6	RAID-6 (14D+2P) on 600 GB 10k RPM SAS HDD	00:01:02	HANA_DATA_N1_01	1800 GB	MPU-10
		00:02:02	HANA_DATA_N2_01	1800 GB	MPU-21
		00:03:02	HANA_DATA_N3_01	1800 GB	MPU-20
		00:04:02	HANA_DATA_N4_01	1800 GB	MPU-11
7	RAID-6 (14D+2P) on 600 GB 10k RPM SAS HDD	00:01:03	HANA_DATA_N1_02	1800 GB	MPU-11
		00:02:03	HANA_DATA_N2_02	1800 GB	MPU-10
		00:03:03	HANA_DATA_N3_02	1800 GB	MPU-21
		00:04:03	HANA_DATA_N4_02	1800 GB	MPU-20
8	RAID-6 (14D+2P) on 600 GB 10k RPM SAS HDD	00:01:04	HANA_DATA_N1_03	1800 GB	MPU-20
		00:02:04	HANA_DATA_N2_03	1800 GB	MPU-11
		00:03:04	HANA_DATA_N3_03	1800 GB	MPU-10
		00:04:04	HANA_DATA_N4_03	1800 GB	MPU-21

Table 4 shows the dynamic provisioning pool IDs and virtual volume LDEV IDs for Hitachi NAS Platform for the production system at both the sites.

Table 4. Dynamic Provisioning Pool IDs and Virtual Volume LDEV IDs for Hitachi NAS Platform

Dynamic Provisioning Pool ID	Dynamic Provisioning Pool Name	Virtual Volume Names	Virtual Volume LDEV ID for HNAS Shared Binaries	Virtual Volume Size for HNAS Shared Binaries	MPU Assignment
0	HNAS_HDP_Pool	HNAS_HANA_VVOL_1	00:0A:01	1600.00 GB	MPU-10
		HNAS_HANA_VVOL_2	00:0A:02	1600.00 GB	MPU-11
		HNAS_HANA_VVOL_3	00:0A:03	1600.00 GB	MPU-20
		HNAS_HANA_VVOL_4	00:0A:04	1600.00 GB	MPU-21

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

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

Table 5 shows an example configuration of the LUN path assignment for Node01. The LUN assignment would be the same for all nodes except for the first LUN, which would be the operating system LUN of that specific node.

Table 5. LUN Path Assignment

LUN ID	LDEV ID	LDEV Name
0000	00:01:00	hananode01
0001	00:01:01	LOG_1
0002	00:01:02	DATA_1_01
0003	00:01:03	DATA_1_02
0004	00:01:04	DATA_1_03
0005	00:02:01	LOG_2
0006	00:02:02	DATA_2_01
0007	00:02:03	DATA_2_02
0008	00:02:04	DATA_2_03
0009	00:03:01	LOG_3
0010	00:03:02	DATA_3_01
0011	00:03:03	DATA_3_02
0012	00:03:04	DATA_3_03
0013	00:04:01	LOG_4
0014	00:04:02	DATA_4_01
0015	00:04:03	DATA_4_02
0016	00:04:04	DATA_4_03

Figure 3 on page 17 shows the LUN and port assignment of the maximum SAP HANA server nodes that can be setup on Hitachi Virtual Storage Platform G800.

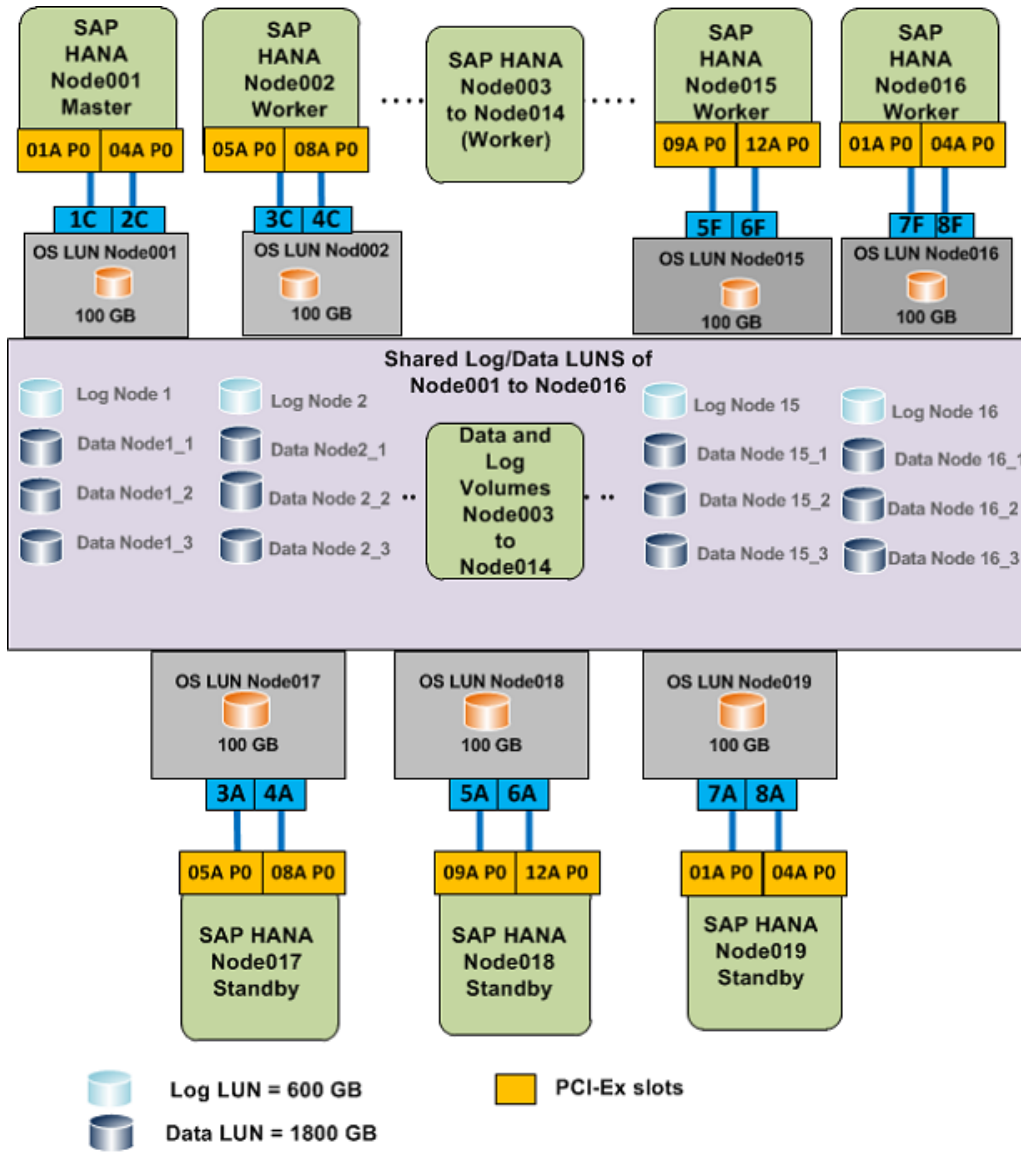


Figure 3

Table 6 shows the parity groups and LDEV assignment the Hitachi NAS Platform volumes, the SAP HANA log volumes, and the SAP HANA data volumes for quality assurance system at the secondary site.

Table 6. Groups and LDEV Assignment of Operating System Boot, Hitachi NAS Platform, SAP HANA Log Volumes, and SAP HANA Data Volumes for Quality Assurance System at Secondary Site

Parity Group ID	Parity Group RAID Level and disks	LDEV ID	LDEV Name	LDEV Size	MPU Assignment
09	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:00:17	HNAS_VOL_QA_1	1600 GB	MPU-10
		00:00:18	HNAS_VOL_QA_2	1600 GB	MPU-11
10	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:00:19	HNAS_VOL_QA_3	1600 GB	MPU-20
		00:00:20	HNAS_VOL_QA_4	1600 GB	MPU-21
11	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:01:05	HANA_LOG_QA_N1	600 GB	MPU-21
		00:02:05	HANA_LOG_QA_N2	600 GB	MPU-20
12	RAID-6 (6D+2P) on 600 GB 10k RPM SAS HDD	00:03:05	HANA_LOG_QA_N3	600 GB	MPU-11
		00:04:05	HANA_LOG_QA_N4	600 GB	MPU-10
13	RAID-6 (14D+2P) on 600 GB 10k RPM SAS HDD	00:01:06	HANA_DATA_QA_N1_01	1800 GB	MPU-10
		00:02:06	HANA_DATA_QA_N2_01	1800 GB	MPU-21
		00:03:06	HANA_DATA_QA_N3_01	1800 GB	MPU-20
		00:04:06	HANA_DATA_QA_N4_01	1800 GB	MPU-11
14	RAID-6 (14D+2P) on 600 GB 10k RPM SAS HDD	00:01:07	HANA_DATA_QA_N1_02	1800 GB	MPU-11
		00:02:07	HANA_DATA_QA_N2_02	1800 GB	MPU-10
		00:03:07	HANA_DATA_QA_N3_02	1800 GB	MPU-21
		00:04:07	HANA_DATA_QA_N4_02	1800 GB	MPU-20
15	RAID-6 (14D+2P) on 600 GB 10k RPM SAS HDD	00:01:08	HANA_DATA_QA_N1_03	1800 GB	MPU-20
		00:02:08	HANA_DATA_QA_N2_03	1800 GB	MPU-11
		00:03:08	HANA_DATA_QA_N3_03	1800 GB	MPU-10
		00:04:08	HANA_DATA_QA_N4_03	1800 GB	MPU-21

Table 7 shows the dynamic provisioning pool IDs and virtual volume LDEV IDs for Hitachi NAS Platform for quality system at the secondary site.

Table 7. Dynamic Provisioning Pool IDs and Virtual Volume LDEV IDs of Hitachi NAS Platform for Quality Assurance System at the Secondary Site

Dynamic Provisioning Pool ID	Dynamic Provisioning Pool Name	Virtual Volume Names	Virtual Volume LDEV ID for NAS Platform Shared Binaries	Virtual Volume Size for NAS Platform Shared Binaries	MPB Assignment
1	HNAS_QA_HDP_Pool	HNAS_QA_VVOL_1	00:0A:17	1600.00 GB	MPU-10
		HNAS_QA_VVOL_2	00:0A:18	1600.00 GB	MPU-11
		HNAS_QA_VVOL_3	00:0A:19	1600.00 GB	MPU-20
		HNAS_QA_VVOL_4	00:0A:20	1600.00 GB	MPU-21

Table 8 shows an example configuration of the LUN path assignment for Node01 on secondary site. The LUN assignment would be the same for all nodes except for the first LUN, which would be the operating system LUN of that specific node.

Table 8. LUN Path Assignment at Secondary Site

LUN ID	LDEV ID	LDEV Name
0000	00:01:00	hananode01
0001	00:01:01	LOG_1
0002	00:01:02	DATA_1_01
0003	00:01:03	DATA_1_02
0004	00:01:04	DATA_1_03
0005	00:02:01	LOG_2
0006	00:02:02	DATA_2_01
0007	00:02:03	DATA_2_02
0008	00:02:04	DATA_2_03
0009	00:03:01	LOG_3
0010	00:03:02	DATA_3_01
0011	00:03:03	DATA_3_02
0012	00:03:04	DATA_3_03
0013	00:04:01	LOG_4
0014	00:04:02	DATA_4_01
0015	00:04:03	DATA_4_02
0016	00:04:04	DATA_4_03
0017	00:01:05	LOG_QA_1
0018	00:01:06	DATA_QA_1_01
0019	00:01:07	DATA_QA_1_02
0020	00:01:08	DATA_QA_1_03
0021	00:02:05	LOG_QA_2
0022	00:02:06	DATA_QA_2_01
0023	00:02:07	DATA_QA_2_02
0024	00:02:08	DATA_QA_2_03
0025	00:03:05	LOG_QA_3
0026	00:03:06	DATA_QA_3_01
0027	00:03:07	DATA_QA_3_02
0028	00:03:08	DATA_QA_3_03
0029	00:04:05	LOG_QA_4
0030	00:04:06	DATA_QA_4_01
0031	00:04:07	DATA_QA_4_02
0032	00:04:08	DATA_QA_4_03

Inter-site Configuration

SAN switch and long distance amplifiers must be installed between the primary site and the secondary site, as and when applicable.

Customers who have the SAN switches and long distance amplifiers in the existing infrastructure can utilize the same infrastructure.

Table 9 lists the target and initiator ports of the two storage systems at each site along with zoning alias. Both sites need one zoning configuration.

Table 9. Zoning Configuration

Initiator Port	RCU Target Port	Zone Name
PrimA_Port3B	SecB_Port7B	PrimA_Port3B_SecB_Port7B
PrimA_Port4B	SecB_Port8B	PrimA_Port4B_SecB_Port8B
SecB_Port3B	PrimA_Port7B	SecB_Port3B_PrimA_Port7B
SecB_Port4B	PrimA_Port8B	SecB_Port4B_PrimA_Port8B

Table 10 has the details on the zoning between the management server and Hitachi Virtual Storage Platform G800.

- Create a single zone named MGMT_Pri_Sec.
- Add all four Virtual Storage Platform G800 ports aliases and all of the four management server Emulex port aliases in the single zone MGMT_Pri_Sec.

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 10. Management Server Zoning

Virtual Storage Platform G800 Alias	Hitachi Compute Rack Management Server Alias	Zone Name
PrimA_Port5B	MGMT_PRI_Port0	MGMT_Pri_Sec
SecB_Port5B	MGMT_SEC_Port0	MGMT_Pri_Sec
PrimA_Port6B	MGMT_PRI_Port1	MGMT_Pri_Sec
SecB_Port6B	MGMT_SEC_Port1	MGMT_Pri_Sec

Data Replication and Disaster Recovery Component Configuration

To provide data replication and disaster recovery for the backup and recovery of SAP HANA, Hitachi TrueCopy requires you to setup the following:

- “Install Command Control Interface” on page 21
- “Install Hitachi Dynamic Link Manager” on page 21
- “Configure Command Devices” on page 21
- “Configure Replication using Command Control Interface,” starting on page 21
- “Setup Hitachi TrueCopy” on page 24

Install Command Control Interface

Command control interface enables you to perform storage system operations by issuing commands to Hitachi Virtual Storage Platform G800.

In this solution, install command control interface on the management server for the primary site and the secondary site. Command control interface uses components residing on the following:

- **Storage System** — Command devices and Hitachi TrueCopy volumes (P-VOLs and S-VOLS)
- **Rack optimized server for solutions, 2U four nodes, node 1** — Hitachi Open Remote Copy Manager

Install Hitachi Dynamic Link Manager

Dynamic Link Manager provides the ability to distribute loads across multiple paths and switch to another path if there is a failure in the path currently being used, improving system availability and reliability.

Install Hitachi Dynamic Link Manager on the management server for the primary site and the secondary site. Dynamic Link Manager manages the access paths to the storage system.

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.

For this solution, create a 100 MB command device logical volume on the local storage system and the remote storage system.

Each management server has one dual-port Emulex HBA card, connected through the Brocade 6505 switch to Hitachi Virtual Storage Platform G800. 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.

Configure Replication using Command Control Interface

A key aspect of this reference architecture on Hitachi Virtual Storage Platform G800 is defining the volume pair relationship for data replication between storage systems. Define and manage storage replication relationships through either of the following:

- The graphical user interface in Hitachi Storage Navigator
- On the management server running Hitachi Open Remote Copy Manager

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

To be operational, Hitachi TrueCopy requires instances of Open Remote Copy Manager in the following places:

- On the primary management server to manage the P-VOLs
- On the secondary management server to manage 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
- Log volumes
- Hitachi NAS Platform LUNs

Figure 4 on page 22 shows the sample content of the configuration file (horcm04.conf) used by the Hitachi Open Remote Copy Manager instance on the primary management server.

Figure 5 on page 23 shows the sample configuration file (horcm06.conf) files used by the Hitachi Open Remote Copy Manager instance on the secondary management server.

Figure 6 on page 23 lists the sample entries that have to be added to the services file of the management server on both sides for Open Remote Copy Manager to function.



```

horcm04 - Notepad
File Edit Format View Help
#/****** For HORCM_MON *****/
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
172.17.171.104  horcm04      1000          3000

#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name        dev_name      dev_name
\\.\PhysicalDrive2

HORCM_CTQM
#groupinterval (10ms)
HANADR 100

#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group      dev_name      Serial#      CU:LDEV(LDEV#)  MU#
HANADR          LOG_1         310062      01:01           0
HANADR          LOG_2         310062      02:01           0
HANADR          DATA_1_01    310062      01:02           0
HANADR          DATA_1_02    310062      01:03           0
HANADR          DATA_1_03    310062      01:04           0
HANADR          DATA_2_01    310062      02:02           0
HANADR          DATA_2_02    310062      02:03           0
HANADR          DATA_2_03    310062      02:04           0
HANADR          HNAS_HANA_VVOL_1 310062      0A:01           0
HANADR          HNAS_HANA_VVOL_2 310062      0A:02           0
HANADR          HNAS_HANA_VVOL_3 310062      0A:03           0
HANADR          HNAS_HANA_VVOL_4 310062      0A:04           0

#/****** For HORCM_INST *****/
HORCM_INST
#dev_group      ip_address    service
HANADR          172.17.171.106 horcm06

```

Figure 4

```

horcm06 - Notepad
File Edit Format View Help
#/****** For HORCM_MON *****/
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
172.17.171.106  horcm06     1000          3000

#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name        dev_name      dev_name
\\.\PhysicalDrive2

HORCM_CTQM
#groupinterval (10ms)
HANADR 100

#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group      dev_name      Serial#      CU:LDEV(LDEV#)  MU#
HANADR          LOG_1         350063      01:01           0
HANADR          LOG_2         350063      02:01           0
HANADR          DATA_1_01    350063      01:02           0
HANADR          DATA_1_02    350063      01:03           0
HANADR          DATA_1_03    350063      01:04           0
HANADR          DATA_2_01    350063      02:02           0
HANADR          DATA_2_02    350063      02:03           0
HANADR          DATA_2_03    350063      02:04           0
HANADR          HNAS_HANA_VVOL_1 350063      0A:01           0
HANADR          HNAS_HANA_VVOL_2 350063      0A:02           0
HANADR          HNAS_HANA_VVOL_3 350063      0A:03           0
HANADR          HNAS_HANA_VVOL_4 350063      0A:04           0

#/****** For HORCM_INST *****/
HORCM_INST
#dev_group      ip_address    service
HANADR          172.17.171.104 horcm04
    
```

Figure 5

```

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

Figure 6

Setup Hitachi TrueCopy

The Hitachi TrueCopy setup for this SAP HANA disaster recovery solution is outlined in Table 1, starting on page 6. The reference solution setup is as follows:

- For the failover to the secondary site, this solution uses two initiator ports on Hitachi Virtual Storage Platform G800 at the primary site connected to two RCU target ports on Hitachi Virtual Storage Platform G800 at the secondary site.
- For the failback to the primary site, it uses 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.
- The initiator and RCU target ports on Hitachi Virtual Storage Platform G800 on the primary site connect to the Brocade 6505 Fibre Channel switch at the primary site.
- The initiator and RCU target ports on Hitachi Virtual Storage Platform G800 on the secondary site connect to the Brocade 6505 Fibre channel switch at the secondary site.
- Define the port attributes for initiator and target ports on Hitachi Virtual Storage Platform G800. Configure the storage system for Hitachi TrueCopy replication by defining the logical paths for data replication.
- Four ISL connections exist between the Brocade 6505 Fibre Channel switch on the primary site and the Brocade 6505 Fibre Channel switch on the secondary site.

Configure Hitachi TrueCopy

Configuring Hitachi TrueCopy for Hitachi Unified Compute Platform for SAP HANA requires the following steps.

- Configuration work flow
- Define port attributes
- Add remote connection
- Set number of volumes copied, path blockade, other options

The detailed implementation steps are provided in *Hitachi TrueCopy User Guide* (MK-94HM8019-XX). Contact your Hitachi Data Systems representative if you need this document.

Start the Hitachi Open Remote Control Manager Instance

After configuring Hitachi TrueCopy, start the Hitachi Open Remote Control Manager instances on both sides.

- Start up the Open Remote Control Manager instance on the primary management server. Ensure the correct instance number is used.
- Start up the Open Remote Control Manager instance on the secondary management server. Ensure the correct instance number is used.

Initial Paircreate Operation

Perform the initial data transfer, called the paircreate operation, between the primary site and secondary SAP HANA node volumes in data/log/HNAS volumes of HANA nodes. Execute the command on the primary site.

When executing the paircreate command, the initial copy happens. 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.

Initial Configuration for Replication of Hitachi NAS Platform 4060 LUNs

This is a one-time configuration to setup the mirror relationship between system drives on Hitachi NAS Platform.

The mirroring ensures that the registry information about the replicated system drives of the primary cluster is added on to secondary cluster. Mirroring also ensures that both Hitachi NAS Platform clusters are aware of the replication and mirror relationship.

If one site is permanently lost and the surviving LUs are promoted into the SSWS state, it is necessary to run 'sd-peg -up' on the surviving NAS Platform cluster to make it treat the S-VOLs as P-VOLs. Otherwise, sd-peg should never be used.

In addition, registry changes made on one cluster while it is in production always need recording when made and then copied to the other cluster after the next failover:

- If creating any new file systems, bind them to EVSs using the 'evsfs' command, and then export.
- If deleting any file systems on one cluster, delete them from the registry of the other cluster by the file system-forget-and-delete-nv-data command.
- If any exports have been created, deleted, or modified on one cluster, make the same changes on the other cluster.

Contact your Hitachi Data Systems representative for installation and configuration details.

Engineering Validation

Validation of this reference architecture was conducted in the Hitachi Data Systems laboratory. Tests of the steps of a failover to the secondary site and a failback to primary site for a planned and an unplanned shutdown using Hitachi TrueCopy were performed.

Test Methodology

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

- Planned failover to the secondary site
- Planned failback to the primary site
- Planned failover to the secondary site during mid-flush and then failback to the primary
- Unplanned failover to the secondary site
- Automated failover and failback using Hitachi Disaster Recovery Manager without quality assurance on the secondary site.

Test Results

All the tests passed without issues. The RPO was near zero as verified. The RTO was less than an hour in the lab environment.

Testing resulted in a time of less than 8 minutes for the server-storage failover operation and failback operation between the two sites.

- The overall time for the failover was 3 minutes 45 seconds.
- The overall time for the failback was 3 minutes, 45 seconds.

For More Information

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

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate [Resources](#) website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services [Education](#) website.

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

Hitachi Data Systems



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

© Hitachi Data Systems Corporation 2016. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. VSP and Hi-Track are trademarks or registered trademarks of Hitachi Data Systems Corporation. Microsoft, Windows Server, and Internet Explorer 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 Data Systems Corporation.

AS-467-02. October 2016.