HITACHI
Inspire the Next

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure for SAP HANA – Virtualization and Global-Active Device

Best Practices Guide

By Stephen Ma, Markus Berg

April 2020

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-196-00	Initial release	April 20, 2020

Table of Contents

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure for SAP HANA – Virtualization and Global-Active Device

Best Practices Guide

Use this to design and implement your Cisco and Hitachi Adaptive Solutions for Converged Infrastructure. It has a virtual server infrastructure (VSI) using Cisco ACI (application centric infrastructure) with global-active device on the Hitachi Virtual Storage Platform 5000 series (VSP). This creates your private cloud solution that is extended across multiple locations within metro distances (500 kilometers or about 300 miles).

Cisco Validated Designs are systems and solutions that are designed, tested, and documented to facilitate and accelerate your deployments. They incorporate a wide range of technologies, products and best-practices into a portfolio of solutions that address your business needs.

Cisco and Hitachi work together to deliver a converged infrastructure solution that helps enterprise businesses meet their current challenges to position themselves for the future. These best practices expand the previously-released Cisco and Hitachi Adaptive Solutions with Cisco ACI, which is a virtual server infrastructure incorporating Cisco ACI.

This solution does the following:

- Create a stretched data center using the Cisco ACI MultiPod design for a seamless network between locations within metro distance
- Use global-active device on the Hitachi Virtual Storage Platform 5000 series for active-active storage clustering across these locations
- Incorporate other features of the Virtual Storage Platform 5000 series

The recommended architecture is built on the following:

- Cisco Unified Computing System (Cisco UCS) using the unified software release to support the Cisco UCS hardware platforms for the Cisco UCS B-Series Blade Server
- Cisco UCS 6400 or UCS 6300 Fabric Interconnects
- Cisco Nexus 9000 Series switches
- Cisco MDS 9000 Multilayer switches
- Hitachi Virtual Storage Platform 5000 series along with other enterprise Virtual Storage Platform systems.

This provides guidance to you on how to implement a stretched datacenter across metro distances with Cisco and Hitachi Adaptive Solutions for Converged Infrastructure. It is a best practice datacenter architecture extended between locations as a seamless environment for the underlying network, storage hypervisor-based compute infrastructure.

It expands the Cisco and Hitachi Adaptive Solutions with Cisco ACI using the Cisco ACI MultiPod design for a uniform network that incorporates global-active device on Virtual Storage Platform for storage resiliency between locations. Also, it complies with the VMware vSphere Metro Storage Cluster (vMSC) specifications.

This design features the introduction of the Hitachi Virtual Storage Platform 5000 series for NVMe drives and SAS flash storage. Virtual Storage Platform connects through the Cisco MDS Multilayer Switch to Cisco UCS. It is enabled within the ACI network using the same Cisco Nexus family of switches.

These best practices cover specifics of products utilized within the Cisco validation lab. They are considered relevant for equivalent supported components listed in the Cisco and Hitachi Vantara published compatibility matrixes. Supported adjustments from the example validated build must be evaluated with care, as their implementation instructions may differ.

Note – These procedures were developed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

Solution Components

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a stretched datacenter is a validated architecture incorporating Cisco and Hitachi products. It produces a resilient uniform datacenter infrastructure spanning multiple locations, keeping within the VMware vSphere Metro Storage Cluster (vMSC) recommended practices.

When validating the best practices in this guide, the IP network used was a series of 75 kilometer Fibre spools with the SAN network connection using shorter multimode Fibre. It relied on the independent validation of long distance Fibre Channel connections supported by Cisco and Hitachi.

Cisco and Hitachi Adaptive Solutions for the SAP HANA platform with Cisco Unified Computing System server blades, the Hitachi Virtual Storage Platform family, and the VMware virtualization technology, builds an infrastructure for SAP HANA in a tailored data center integration (TDI) environment. It is validated for SUSE Linux Enterprise Server (SLES) or Red Hat Enterprise Linux (RHEL) as the guest operating systems.

This is an architecture that covers specifics of products utilized within the Cisco lab (Figure 1 on page 3). However, equivalent supported products can be replaced, which are listed within Cisco's and Hitachi's published compatibility matrixes.

Figure 1

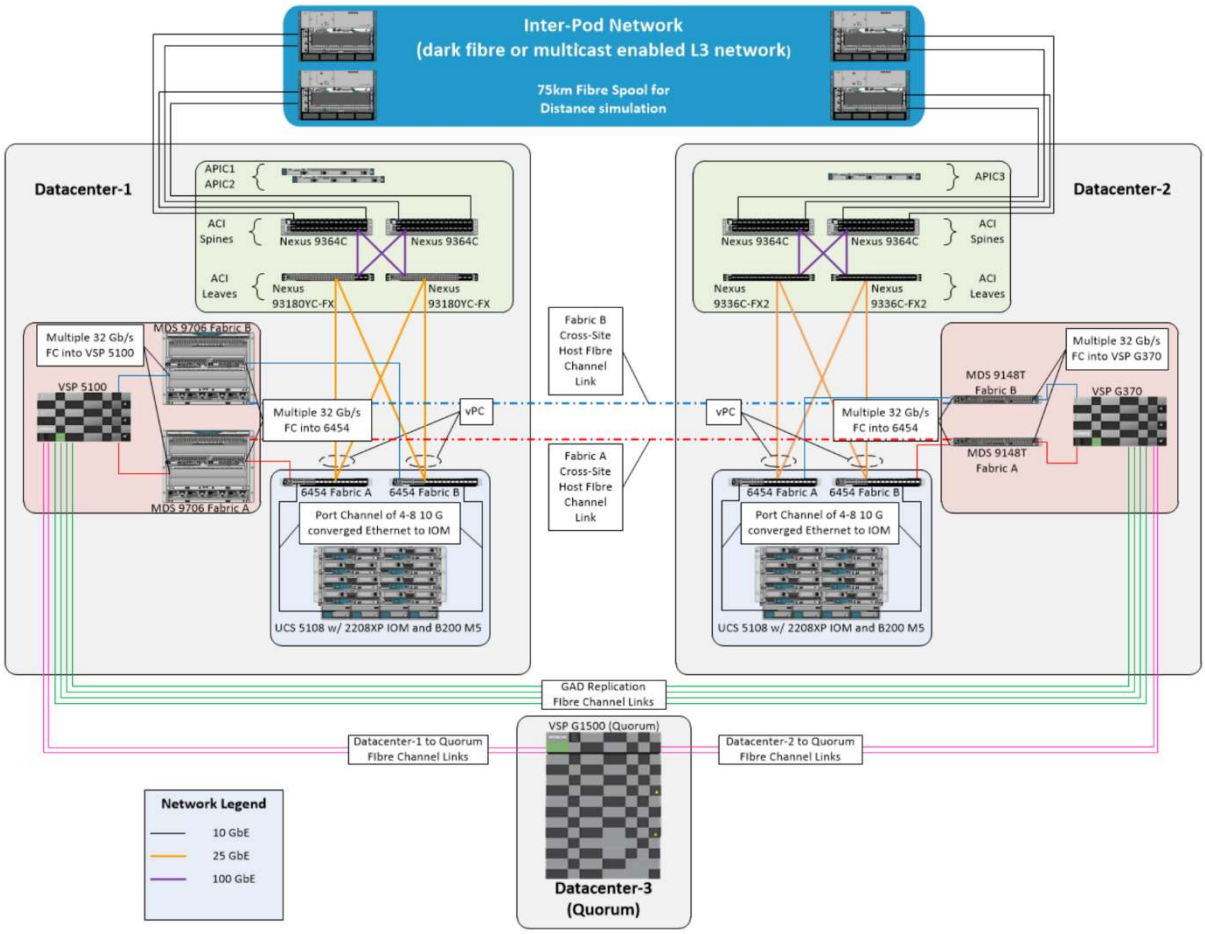

For more information, read the design guide and a more detailed deployment guide on the following:

- **Cisco and Hitachi Adaptive Solutions for Converged Infrastructure**
  - Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide
  - Cisco and Hitachi Adaptive Solutions for Converged Infrastructure
- **Cisco and Hitachi Adaptive Solutions for SAP HANA TDI**
  - Cisco and Hitachi Adaptive Solutions for SAP HANA Tailored Data Center Integration Design Guide
  - Cisco and Hitachi Adaptive Solutions for SAP HANA TDI on Cisco UCS M5 Servers with SLES 12 SP4 and RHEL 7.5

These design and deployment guides have details on the following:

- **Compute**

  This incorporates servers the current processor generation of Intel Xeon Scalable Platform processors. The servers are available in Cisco's Advanced blade factor, managed by Cisco UCS Manager.

- **Network**

  The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gb/s Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric.

  The unified fabric uses the innovative single connect technology to lower costs by reducing the number of network adapters, switches, and cables. This lowers the power and cooling needs of the system.

- **Storage access**

  Cisco UCS system provides consolidated access to SAN storage and network attached storage over the unified fabric. This provides storage choices and investment protection.

  Also, server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.

- **Management**

  The system uniquely integrates compute, network and storage access subsystems, enabling it to be managed as a single entity through Cisco UCS Manager. This increases IT staff productivity by enabling storage, network, and server administrators to collaborate on service profiles that define the desired physical configurations and infrastructure policies for applications.

  Service profiles increase business agility by enabling IT to automate and provision resources in minutes instead of days.

Software Components

These are some major software components used in a Cisco and Hitachi Adaptive Solutions for Converged Infrastructure environment.

VMware vSpehere

VMware vSphere is a virtualization platform that provides a datacenter infrastructure. It helps you get the best performance, availability, and efficiency from your infrastructure and applications. Virtualize applications with confidence using consistent management.

VMware vSphere has the following components:

- VMware vSphere ESXi

  This hypervisor loads directly on a physical server. ESXI provides a robust, high-performance virtualization layer that abstracts server hardware resources and makes them shareable by multiple virtual machines.

- VMware vCenter Server

  This provides a centralized platform for managing your VMware vSphere environments so you can automate and deliver a virtual infrastructure with confidence:

  - VMware vSphere vMotion
  - VMware vSphere Storage vMotion
  - VMware vSphere Distributed Resource Scheduler
  - VMware vSphere High Availability
  - VMware vSphere Fault Tolerance

A VMware vSphere Metro Storage Cluster (vMSC) architecture on Hitachi Virtual Storage Platform uses global-active device (active-active stretched clustering technology) to provide an ideal solution for maximizing availability and uptime by clustering physical data centers within metro distances. This usually means distances of less than 500 kilometers (about 300 miles).

A Metro Storage Cluster solution from Hitachi has storage systems presenting replicated storage as a single LUN from geographic sites. The design enables high availability of services by allowing virtual machine migration between sites with no downtime.

At a minimum, you need to use one of the following options:

- First option:
  - VMware vSphere ESXi 6.5, or vSphere 6.7
  - VMware vSphere Metro Storage Cluster
  - Global-active device with VMware Native Multi-Pathing (NMP) with ALUA or
- Second option:
  - VMware vSphere ESXi vSphere 6.5, or vSphere 6.7
  - VMware vSphere Metro Storage Cluster
  - Global-active device with Hitachi Dynamic Link Manager

VMware vSphere ESXi hosts using either iSCSI or Fibre Channel protocol are fully supported for Metro Storage Cluster with global-active device configuration on Hitachi Virtual Storage Platform systems.

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.

As a SAP customer, you can log on and download more information about SAP HANA, including the following:

- **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 Server Installation and Update Guide**

  This provides an overview of how to install and update a SAP HANA system with the SAP HANA lifecycle management tools.

- **SAP HANA Administration Guide**

  This explains how to configure, manage, maintain and optimize your SAP HANA installation using SAP HANA administration tools.

SAP HANA hardware directory provides information about SAP HANA appliances certified by SAP hardware partners.

Guest Operating System Options

The following options are available as the operating system:

- **SUSE Linux Enterprise Server**

  Compete more effectively though 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 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**

  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 along the guidelines of SAP and the operating system distributor and may otherwise cause significant performance problems. The following SAP Notes for SUSE Linux Enterprise Server and Red Hat Enterprise Linux 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

Hardware Components

You can use the Hitachi Virtual Storage Platform models in Table 1, "Hitachi Virtual Storage Platform F Series," on page 7, and Table 2, "Hitachi Virtual Storage Platform G Series," on page 8 for Cisco and Hitachi Adaptive Solutions for Converged Infrastructure environment.

TABLE 1. HITACHI VIRTUAL STORAGE PLATFORM F SERIES

	VSP F350	VSP F370	VSP F700	VSP F900	VSP F1500	VSP F5500
Performance (IOPS)	600K	1.2M	1.4M	2.4M	4.8M	21M
Maximum Raw Internal Capacity	2.8 PB (15 TB SSD)	4.3 PB (15 TB SSD)	6.0 PB (14 TB FMD) 13.0 PB (15 TB SSD)	8.1 PB (14 TB FMD) 17.3 PB (15 TB SSD)	8.1 PB (14 TB FMD) 34.6 PB (15 TB SSD)	69PB
Total Efficiency Guarantee Ratio *	Up to 7:1	Up to 7:1	Up to 7:1	Up to 7:1	2:1 Data Reduction	Up to 7:1
Host Port Counts (Maximum without drives)	16 Fibre Channel 8 iSCSI	16 Fibre Channel 8 iSCSI	48 Fibre Channel 24 iSCSI	64 Fibre Channel 32 iSCSI	176 Fibre Channel 176 IBM® FICON® 88 iSCSI	192 Fibre Channel 192 IBM® FICON® 96 iSCSI
Value Added Bundled Features	100% Data Availability Guarantee Adaptive Data Reduction Storage Virtualization In System Replication Copy Management Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Adaptive Data Reduction Storage Virtualization In System Replication Copy Management Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Adaptive Data Reduction Storage Virtualization In System Replication Copy Management Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Adaptive Data Reduction Storage Virtualization In System Replication Copy Management Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Adaptive Data Reduction Storage Virtualization In System Replication Copy Management Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Adaptive Data Reduction Storage Virtualization In System Replication Copy Management Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee

TABLE 2. HITACHI VIRTUAL STORAGE PLATFORM G SERIES

	VSP G350	VSP G370	VSP G700	VSP G900	VSP G1500	VSP G5500
Performance (IOPS)	600K	1.2M	1.4M	2.4M	4.8M	21M
Maximum Raw Internal Capacity	2,467 TB (10 TB LFF HDD) 442 TB (2.4 TB SFF HDD) 2,889 TB (15 TB SSD)	3,642 TB (10 TB LFF HDD) 664 TB (2.4 TB SFF HDD) 4,333 TB (15 TB SSD)	11.7 PB (10 TB LFF HDD) 1.9 PB (2.4 TB SFF HDD) 6.0 PB (14 TB FMD) 13.0 PB (15 TB SSD)	14.0 PB (10 TB LFF HDD) 2.6 PB (2.4 TB SFF HDD) 8.1 PB (14 TB FMD) 17.3 PB (15 TB SSD)	6.7 PB (6 TB LFF HDD) 5.3 PB (2.4 TB SFF HDD) 8.1 PB (14 TB FMD) 34.6 PB (15 TB SSD)	69PB
Total Efficiency Guarantee Ratio *	Up to 7:1	Up to 7:1	Up to 7:1	Up to 7:1	2:1 Data Reduction	Up to 7:1
Host Port Counts (Max. without drives)	16 Fibre Channel 8 iSCSI	16 Fibre Channel 8 iSCSI	64 Fibre Channel 32 iSCSI	80 Fibre Channel 40 iSCSI	192 Fibre Channel 176 IBM FICON 192 Fibre Channel over Ethernet 96 iSCSI	192 Fibre Channel 192 IBM® FICON® 96 iSCSI
Value Added Bundled Features	100% Data Availability Guarantee Storage Virtualization In System Replication Copy Management Dynamic Tiering Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Storage Virtualization In System Replication Copy Management Dynamic Tiering Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Storage Virtualization In System Replication Copy Management Dynamic Tiering Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Storage Virtualization In System Replication Copy Management Dynamic Tiering Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Storage Virtualization In System Replication Copy Management Dynamic Tiering Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee	100% Data Availability Guarantee Storage Virtualization In System Replication Copy Management Dynamic Tiering Infrastructure Analytics Non-disruptive Migration Total Efficiency Guarantee

Best Practices for Virtualized SAP HANA Systems with the Highest Availibility

SAP HANA tailored datacenter integration (TDI) offers great flexibility for SAP business solutions in production and non-production environments. SAP HANA on VMware vSphere 6.5 or 6.7 follows the SAP HANA TDI guidelines. It offers the power of SAP HANA while benefitting from the flexibility of a virtualized environment.

VMware has certified VMware vSphere 6.5 and vSphere 6.7 on Intel Skylake processors and Cascade Lake processors for virtualized SAP HANA deployments. See SAP Note 2393917 for more information.

SAP HANA TDI increases flexibility, providing alternatives to deploy SAP HANA appliances in 5 phases, including many kinds of virtualization technology. Your need to understand the possibilities and requirements of an SAP HANA TDI environment is crucial. One of these requirements defined by SAP is to consider all virtualized environments for SAP HANA as SAP HANA TDI.

SAP provides the following documentation about SAP HANA TDI environments that explain the 5 phases of SAP HANA TDI, as well as the hardware and software requirements for the whole stack:

- SAP HANA Tailored Data Center Integration 2017 - Overview – This defines the states and steps to implement a SAP HANA TDI solution.
- SAP HANA Tailored Data Center Integration - Frequently Asked Questions
- SAP HANA TDI-Storage Requirements
- SAP HANA Network Requirements

To add higher availability and a 100% data guaranty, Hitachi expands the virtualized SAP HANA solution with global active device. With global-active device set on Hitachi Virtual Storage Platform, you can reach an "always on" SAP HANA Infrastructure.

Global-Active Device

As part of Hitachi Storage Virtualization Operating System RF, global-active device ensures two physical systems are logically presented as one system. Implementing cross-mirrored storage volumes between two Hitachi Virtual Storage Platform systems, global-active device accepts read/write I/Os on both sides that are continuously updated.

If a disk controller failure occurs at one site, the controller at the other site automatically takes over and accepts read/write I/Os. It enables production workloads on both systems while maintaining full data consistency and protection. Global-active device assures that an active and up to date storage volume is available to a production application despite the loss of a virtualized controller, system, or site.

During an outage, global-active device, by design, does the following:

- Blocks I/O to the site where the failure occurred by pushing status PSUE to the volume
- Redirects all I/O to the alternate site automatically by honoring ALUA settings with the use of NPM

For more information on the states and statues of global-active device, see Global-Active Device.

During the outage where pairs are suspended to either PSUE or PSUS state, you must correct the cause of the outage and then resynchronize the active-active volume set, so the data is consistent at the primary site and the secondary site again. For information on disaster recovery of global-active device, refer to Disaster recovery of global-active device.

Optionally, you have the choice between using datacenter-specific storage or highly available global-active device datasets. Within this deployment you can use other Hitachi management products to allocate a VMFS datastore to your virtual environment.

Hitachi Storage Virtualization Operating System RF provides global-active device, which simplifies distributed system design and operations for SAN and NAS workloads. Global-active device allows you to achieve zero recovery time and recovery point objectives (RTO and RPO) by enabling synchronous replication up to 500 kilometers (about 300 miles).

Global-active device supports read/write copies of the same data in two places at the same time. Its active-active design implements cross-mirrored storage volumes between two Hitachi Virtual Storage Platform systems that accept read/write I/Os on both sides which are continuously updated.

Figure 2

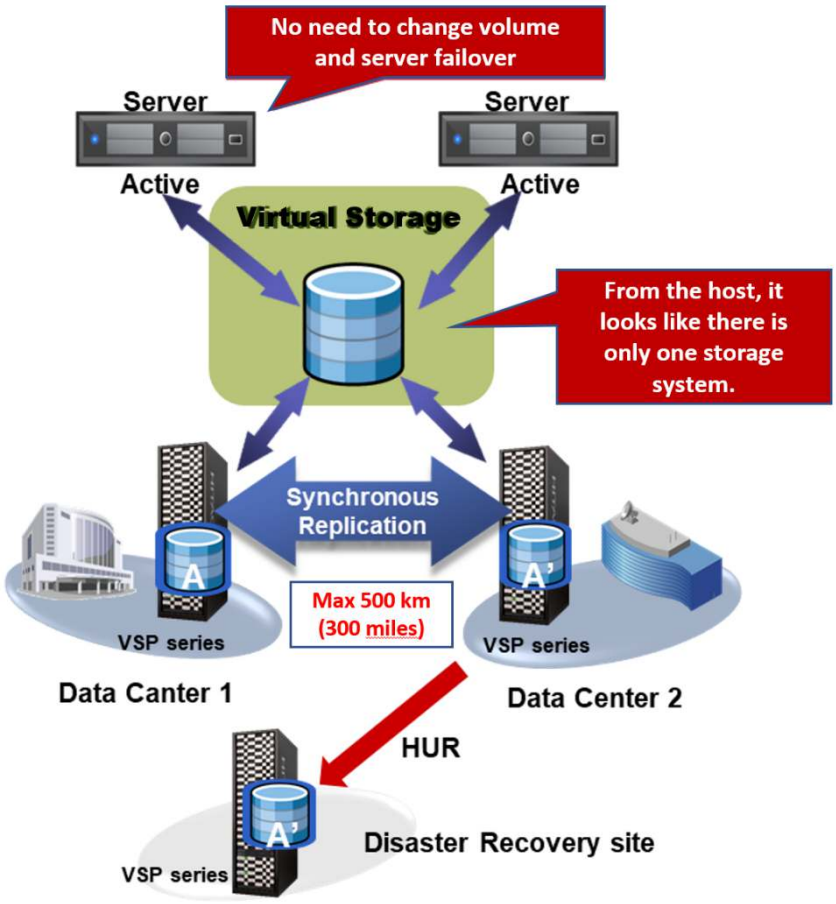

Global-active device supports Virtual Storage Platform model intermix. For example, a Virtual Storage Platform 5000 series system can be replicated to a Virtual Storage Platform F900 system. Check the matrix for model intermix compatibility.

If a disk controller failure occurs at one site, the controller at the other site automatically takes over and accepts read/write I/Os. Global-active device assures that an up-to-date storage volume is always available and enables production workloads on both systems while maintaining full data consistency and protection.

Global-active device simplifies and automates high availability to ensure continuous operations for your most mission-critical data and applications for both SAN and NAS deployments, whether they are block or file. By leveraging some of the unique capabilities of the Virtual Storage Platform family, global-active device assures that an active and up-to-date storage volume is available to a production application despite the loss of a virtualized controller.

In addition, the embedded NAS modules can be configured to mirror NVRAM transactions across the system chassis to enable a stretched NAS cluster that can also be used for nondisruptive workload migration and automated failover across sites. If a controller fails at either site, the NAS services running on that controller are automatically failed over to the surviving site.

VMware vSphere Metro Storage Cluster

A VMware vSphere Metro Storage Cluster environment on Hitachi Virtual Storage Platform F series and VSP G series systems provide an ideal solution for maximizing availability and uptime. It clusters physical data centers within metro distances (up to 500 kilometers or about 300 miles).

This Metro Storage Cluster solution from Hitachi Vantara consists of storage systems presenting replicated storage as a single LUN from different geographically distributed sites. This design enables high availability of services by allowing virtual machine migration between sites with no downtime. A combination of software and hardware products from Hitachi Vantara provides the following key functions to a VMware vSphere infrastructure:

- Host multipathing
- Internal and externalized storage provisioning
- Synchronous storage replication across metro cluster distances
- Transparent storage failover
- Host access using a uniform (recommended) or non-uniform topology

These functions work together with VMware vSphere vMotion, VMware vSphere High Availability, and VMware vSphere Distributed Resource Scheduler to build this reference architecture for vSphere Metro Storage Cluster. The advanced functions found in Hitachi Virtual Storage Platform F series and G series systems do the following:

- Fulfill the requirements of a virtual infrastructure
- Lessen the need for additional hardware that may be required in traditional Metro Storage Cluster solutions

vSphere Metro Storage Cluster supports stretched storage, leveraging global-active device on Virtual Storage Platform which provides a single stretched datastore across the datacenter. Within this design, VM Component Protection (VMCP) from VMware is used to protect virtual machines from storage related events, such as permanent device loss (PDL) and all paths down (APD).

This guide shows the pre-designed, pre-configured, and pre-validated solution is comprised of a VMware vSphere stack and stretched storage stack, leveraging vSphere Metro Storage Cluster and global-active device on Cisco Unified Computing System virtual server instances infrastructure.

VMware Support

Hitachi Virtual Storage Platform F series and VSP G series are aligned with the VMware software-defined storage vision, providing the following support:

- **VMware vSphere Metro Storage Cluster**

  Using global-active device on Hitachi Virtual Storage Platform, you can create and maintain synchronous, remote copies of data volumes.

  Virtual storage instances are configured in the primary and secondary storage systems using the actual information of the primary storage system. Global-active device primary and secondary volumes are assigned the same virtual LDEV number in the virtual storage instances. This enables the host to see the pair volumes as a single volume on a single storage system with both volumes receiving the same data from the host.

  Configuring global-active device as the backend storage for a vSphere Metro Storage Cluster provides an ideal solution for maximizing availability and uptime by clustering physical datacenters that reside within metro distances (up to 500 kilometers, about 300 miles) of each other.

- **Hitachi Storage Provider for VMware vCenter**

  Hitachi Storage Provider for VMware vCenter works with VMware vSphere API for Storage Awareness (VASA) to provide access to Hitachi Virtual Storage Platform F series and VSP G series storage. Storage Provider for VMware vCenter enables policy-based storage management using VMware Storage Policy-based Management (SPBM) and VMware Virtual Volumes (vVols).

  Storage Provider for VMware vCenter provides a simplified method for VMware administrators and storage administrators to deliver effective storage that meets advanced virtual machine requirements.

- **VMware vSphere Storage APIs - Array Integration**

  VMware vSphere Storage APIs – Array Integration (VAAI) enables multiple storage functions (primitives) within vSphere to be offloaded to certified storage hardware. This reduces VMware ESXi processor utilization by allowing certified storage hardware to perform these functions on the storage systems themselves.

  In many cases, VAAI accelerates these functions allowing them to complete in less time as compared to performing the functions within software on the ESXi host.

- **Hitachi Storage Replication Adapter**

  Hitachi Storage Replication Adapter with VMware Site Recovery Manager provides a disaster recovery solution that works with your storage environment and your VMware environment. Storage at both sites are "paired" during Site Recovery Manager configuration. VMware administrators use Site Recovery Manager to manage the configuration and definition of the disaster recovery plan.

- **VMware vSphere API for Multipathing**

  Hitachi Virtual Storage Platform F series and VSP G series support VMware vStorage vSphere API for Multipathing to provide enhanced control of I/O path selection and failover.

- **VMware vStorage API for Data Protection**

  Hitachi Virtual Storage Platform F series and VSP G series support VMware vStorage API for Data Protection to enable backup applications to perform file-level or virtual machine-level backup of running virtual machines.

Multipathing

VMware Native Multipathing Plug-in is multipathing software that integrates with global-active device on Hitachi Virtual Storage Platform to provide load balancing, path optimization, path failover, and path failback capabilities for VMware vSphere hosts.

Native Multipathing Plug-in load balances I/O between all available preferred paths (active) from P-VOL and keep all paths to S-VOL as active non-optimized paths. Native Multipathing Plug-in rules for Virtual Storage Platform are now enabled by default within vSphere. This allows for a reduction in time to production when deploying new VMware vSphere ESXi clusters connected to Hitachi storage. The rules handle devices configured with ALUA to be used for global-active device.

Virtual Storage Machine

Use a virtual storage machine to manage virtualized storage system resources. You create a virtual storage machine within a physical storage system, so the host server recognizes the two storage systems as one virtual storage system.

When using global-active device on Hitachi Virtual Storage Platform to provide nondisruptive host access to volumes that reside on different storage systems, create a virtual storage machine for the storage system that contains the primary volumes (P-VOLs) of the global-active device pairs in the secondary storage system. For global-active device, the primary storage system is the virtualized storage system.

Figure 3 shows the relationship between physical storage systems and virtual storage machines. In this example, the physical storage in Datacenter 1, a Hitachi Virtual Storage Platform 5000 series system, has a virtual VSP G370 created on it. The physical storage in Datacenter 2, a Hitachi Virtual Storage Platform G370, has a virtual VSP G370 created on it. The serial numbers on the virtual storage machines in Datacenter 1 and Datacenter 2 are given the same serial number, which is different from the serial number of either physical storage system.

Figure 3

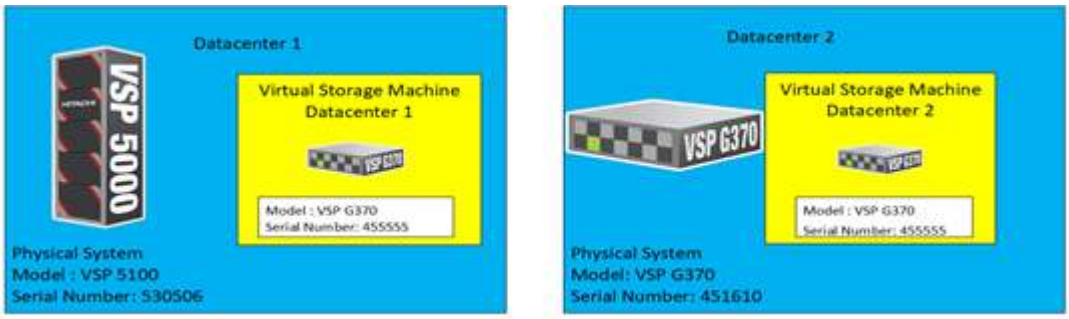

This example emulated a common virtual storage machine of a Hitachi Virtual Storage Platform G370. Click to learn more about virtual storage machine emulation types.

Quorum

To avoid split-brain scenarios, normally a quorum entity external to either system in a separate location determines operational control when specific failures occur.

In a VMware vSphere Metro Storage Cluster on a Hitachi Virtual Storage Platform, there are various options for providing quorum services. This includes a separate storage system, which can be any supported third party storage that can be attached to a Virtual Storage Platform F series or VSP G series system. Or, present an iSCSI disk from a physical or virtual machine from a third site or the cloud.

The recommendation is to keep the quorum system at a third site. This site is used in disaster recovery situations where the primary or secondary datacenter is offline with the quorum used to maintain data consistency between each Virtual Storage Platform in global-active device replication and to verify fail over.

For example, to create an external storage system, create a 20 GB LUN on external storage, such as Virtual Storage Platform G1500 or other supported third-party storage for use as a quorum disk. Present this LUN to Virtual Storage Platform on Site 1 and on Site 2 storage as externalized storage using Hitachi Storage Virtualization Operating System RF with its virtualization device capability. The quorum disk stores continually updated information about data consistency in global-active device P-VOLs and S-VOLs for use during site failover operations. Global-active device uses the information in the event of a failure to direct host operations to the other volume within the pair.

Replication Links Main Control Unit and Remote Control Unit

A storage replication link consists of bidirectional ports on the datacenter 1 storage system connected to a remote-control unit bidirectional port that is defined on the datacenter 2 storage system.

It represents a bidirectional remote copy connection from the primary data volume (P-VOL) on the datacenter 1 storage system to the secondary data volume (S-VOL) on the datacenter 2 storage system. This design defines two main control unit and remote control unit paths for storage replication.

Quorum Links

An example of quorum links is when you configure something as the following:

- Two external paths between a Hitachi Virtual Storage Platform 5100 system in Datacenter 1 storage with the Virtual Storage Platform G1500 system used for quorum storage in a third location
- Two external paths between a Hitachi Virtual Storage Platform G370 in Datacenter 2 with the Virtual Storage Platform G1500 system used for quorum storage in a third location

Use Cases and Your Benefit

These are sample use cases with the benefit you receive when implementing such an environment.

Continuous Operations

For many mission-critical applications, any failure that prevents access to data results in application interruption and possibly manual failover to another copy of data, potentially at a second site.

Global-active device can now provide active-active stretched clusters over local distances and metro distances (up to 500 kilometers or about 300 miles). Multipath software allows application access to replicated data from the shortest path for highest performance.

Applications can avoid manual failovers and access the second copy of data either locally or between metro-distance sites. This enables high availability and continuous operations with access to data, assuring critical business application availability and integrity.

Figure 4 on page 15 shows a failover for the virtual machines running SAP HANA on the left from using the P-VOL to using the S-VOL.

Figure 4

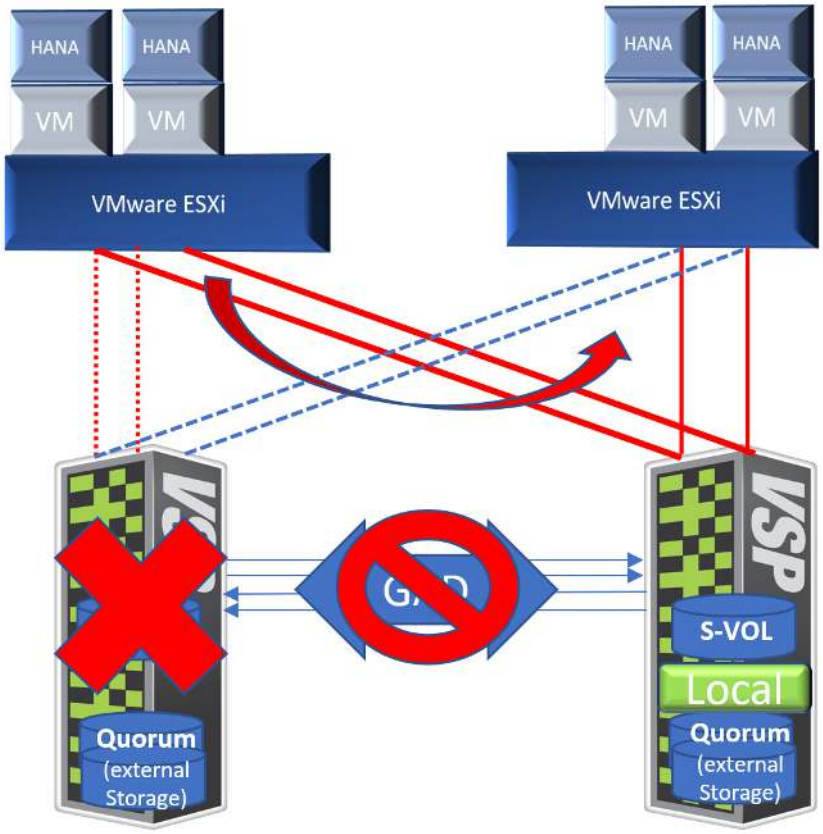

Workload Mobility

There does not need to be any downtime of application workloads when there is a need to move workloads. You can migrate application workloads non-disruptively to another physical location within metro distance (up to 500 kilometers or about 300 miles).

To balance a processing workload during peak usage periods, you can move an application to a second server. The concurrent data mirroring capability of global-active device on Hitachi Virtual Storage Platform makes data immediately available to servers moved to a second site within metro distances without application interruption.

Mirroring of the NVRAM across the NAS modules allows the NAS cluster logic to be extended across sites, allowing enterprise virtual server (EVS) workloads to be non-disruptively balanced across sites to provide an active/active workload distribution.

This process could also be used to take servers offline temporarily for routine maintenance and updates, and to perform non-disruptive upgrades.

Figure 5 on page 16 shows moving a virtual machine with SAP HANA from one VMware ESXi instance to another.

Figure 5

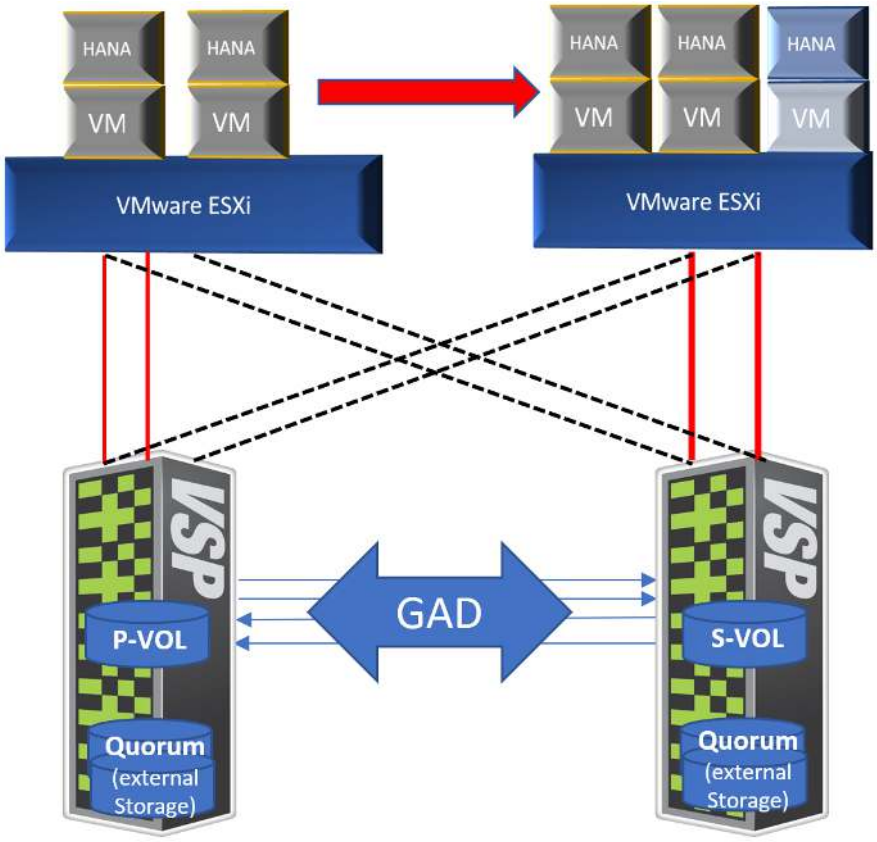

Nondisruptive Data Migration

Global-active device on Hitachi Virtual Platform is a powerful tool when installing and brining online a new Virtual Storage Platform family system when the primary system has critical application volumes in use.

Migrate data volumes from the primary to the new Virtual Storage Platform without disruption to normal operations to provide simpler and automated business continuity. Global-active device makes nonstop operations and workload mobility between systems and data centers a reality.

Global-active device provides greater flexibility, reliability, and scalability to delivers unmatched resiliency by eliminating downtime for mission-critical data and applications.

Figure 6 on page 17 shows how global-active device simplifies and automates data migration.

Figure 6

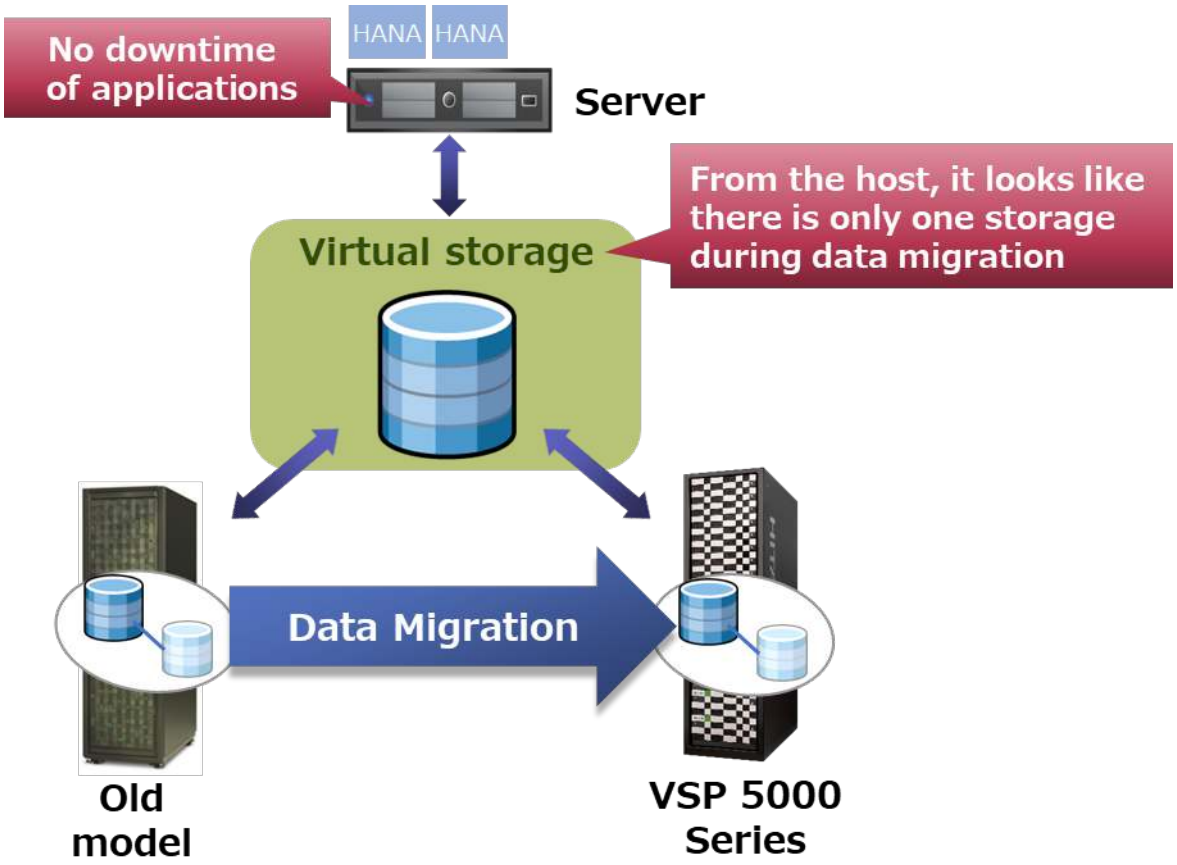

Planned Site Maintenance

As a component of Hitachi Storage Virtualization Operating System RF, use Hitachi Ops Center Administrator to manage resources. The wizard-driven graphical user interface allows anyone to manage, monitor, and maintain Hitachi Virtual Storage Platform systems using a status dashboard. REST-based APIs extend operations. The benefits of global active device in Virtual Storage Platform lie in fault tolerance, failover, and load balancing without storage access impact.

From a host perspective, the volumes from two physically separate storage systems are involved in a global active device pair. If one volume or system had a failure that prevents the host from accessing that volume, host I/O continues by accessing the other volume at the other system. Fault-tolerance is achieved at the storage level.

At the server level, failover can be enhanced through clustering software in which one or more servers have access to the global active device volume through the virtualized storage machine such that, if failure occurred on one server, failover to another server happens without any volume pair configuration actions needed.

Since the primary and secondary storage systems hold the same data in each global active device pair, load balancing becomes simpler. There is no need to execute storage operations to prepare the data in a paired volume for use at the secondary site. The volume can already be seen by a second, remote host. This could host, for example, a virtual machine which eases the I/O burden from the primary storage system and primary host.

Figure 7 shows how resources can be moved for planned site maintenance and then moved back when finished (Figure 8 on page 19, Figure 9 on page 20).

Figure 7

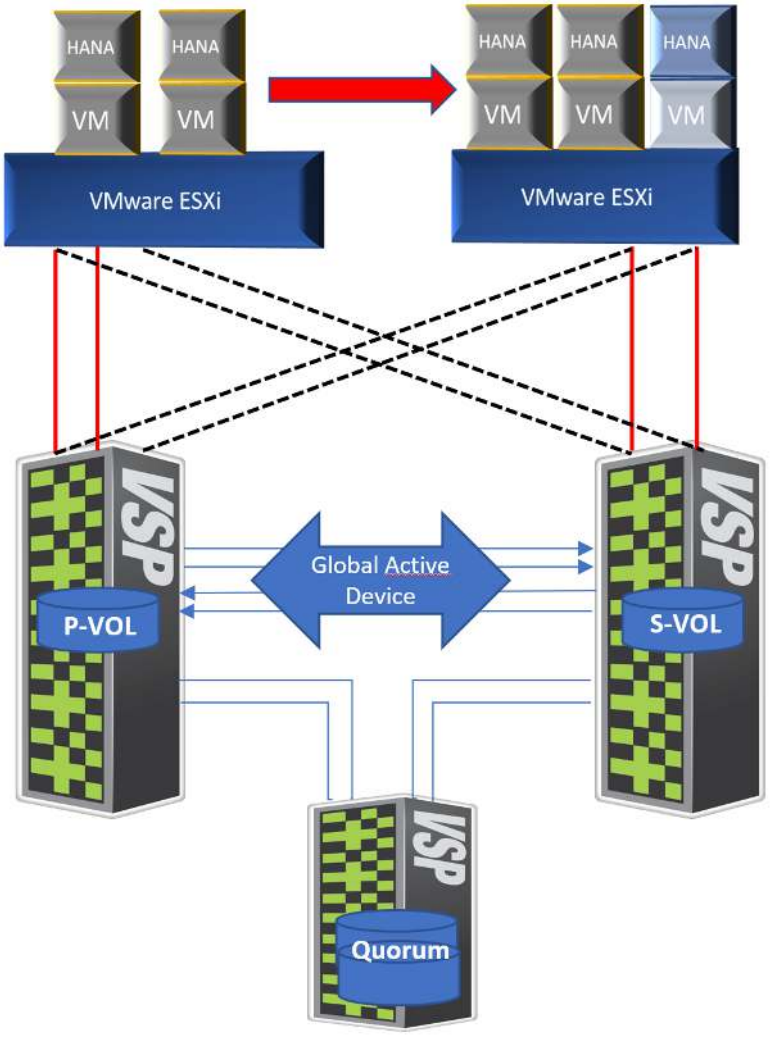

Figure 8

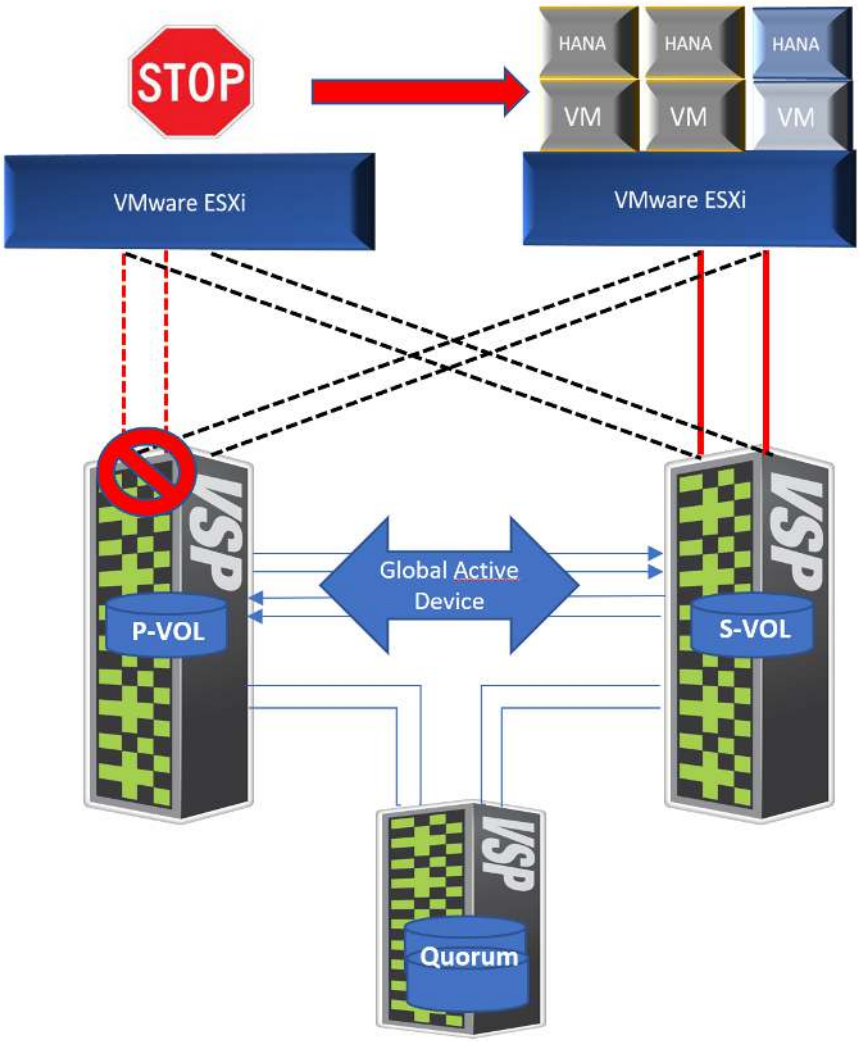
STOP
HANA
HANA
HANA
VM
VM
VM
VMware ESXi
VMware ESXi
VSP
P-VOL
Global Active Device
VSP
S-VOL
VSP
Quorum

Figure 9

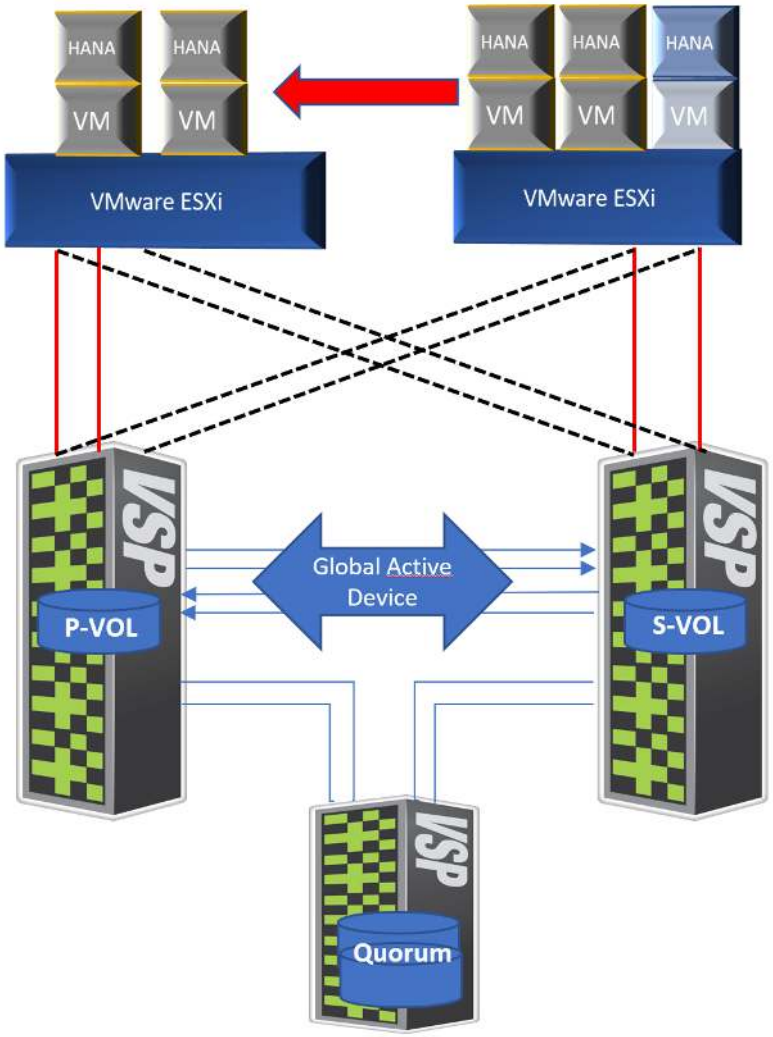
HANA
HANA
VM
VM
VMware ESXi
HANA
HANA
HANA
VM
VM
VM
VMware ESXi
VSP
P-VOL
Global Active Device
VSP
S-VOL
VSP
Quorum

System Requirements

This describes the system hardware and software requirements for the various components for a sample environment.

Hardware Options

The hardware in Table 3 is sample architecture for a two datacenter solution.

TABLE 3. EXAMPLE SYSTEM HARDWARE

Datacenter 1	Datacenter 2
Hitachi Virtual Storage Platform 5100	Hitachi Virtual Storage Platform G370
Cisco UCS B200 M5 Blade Servers	Cisco UCS B200 M5 Blade Servers
Cisco UCS 6454 Fabric Interconnects	Cisco UCS 6454 Fabric Interconnects
Cisco ACI Application Policy Infrastructure Controllers	Cisco ACI Application Policy Infrastructure Controller
Cisco Nexus 93180YC-FX Switches (Leaf)	Cisco Nexus 93180YC-FX Switches (Leaf)
Cisco Nexus 9364C Switches (Spine)	Cisco Nexus 9364C Switches (Spine)
Cisco Nexus 93180YC-EX Switches (NX-OS Standalone Mode Data Center Interconnect/Existing Network connectivity as implemented for lab simulation)	Cisco Nexus 93180YC-EX Switches (NX-OS Standalone Mode Data Center Interconnect/Existing Network connectivity as implemented for lab simulation)

SAP HANA Database Operating System Options

There are the three supported operating systems:

- SUSE Linux Enterprise Server 12SP3-HANA 2.0 SPS03 Scale-up MDC
- Red Hat Enterprise Linux 7.4-HANA 2.0 SPS03 Scale-up MDC
- SUSE Linux Enterprise Server 12SP3-HANA 2.0 SPS02 Scale-Out MDC

SAP HANA Volumes

You need these volumes for SAP HANA.

Figure 10

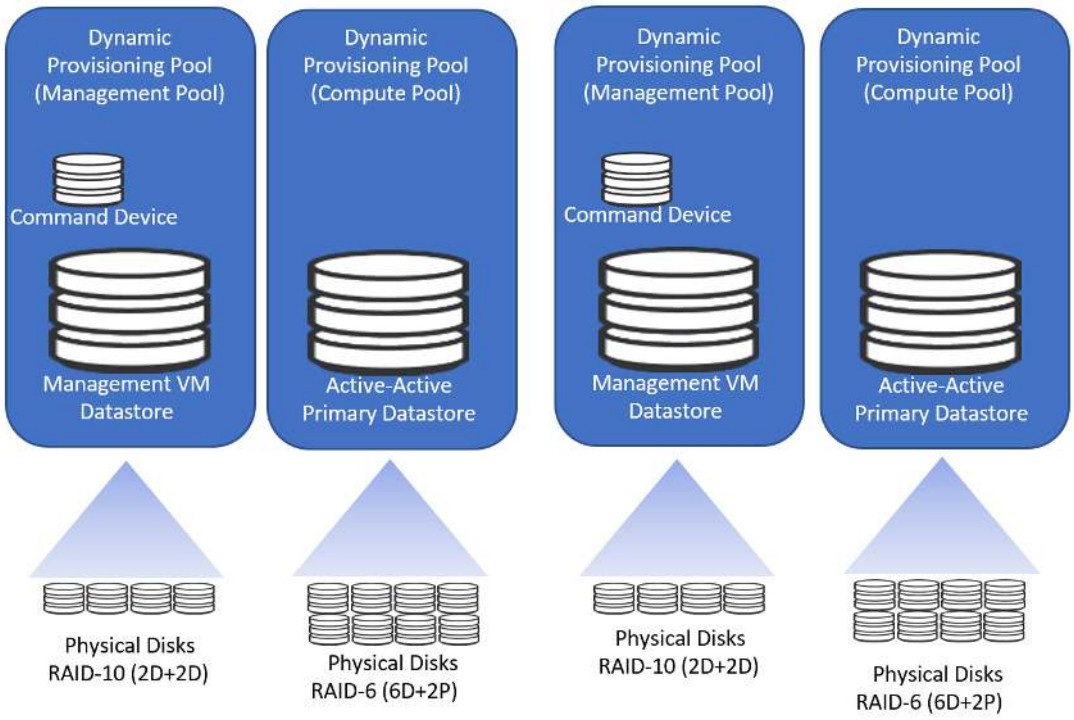

References

Read these design guides to see various options for your environment:

- Cisco and Hitachi Adaptive Solutions for SAP HANA Tailored Data Center Integration Design Guide
- Cisco and Hitachi Adaptive Solutions for SAP HANA TDI on Cisco UCS M5 Servers with SLES 12 SP4 and RHEL 7.5

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 Vantana 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 95054 USA
www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact

MK-SL-196-00. April 2020.