Multi-Site Stretch Cluster and Disaster Recovery Protection Using VMware Cloud Foundation on the Hitachi Integrated Systems Platform with VSP One

A Reference Architecture and Configuration Guide

MK-SL-356-00
July 2025

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

The open source content used in Hitachi Vantara products may be found within the Product documentation or you may request a copy of such information (including source code and/or modifications to the extent the license for any open source requires Hitachi make it available) by sending an email to OSS_licensing@hitachivantara.com.

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to GPSE-Docs-Feedback@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.

Thank you!

Revision history

Changes	Date
Initial release	July 2025

Contents

A Reference Architecture and Configuration Guide

Introduction

In today's always-on digital landscape, enterprise IT requires infrastructure that ensures continuous availability, robust disaster recovery, and seamless scalability.

Without a robust business continuity and disaster recovery solution, businesses face significant operational and business risks:

- Extended Downtime: In the event of site failures or disasters, the inability to quickly failover workloads can lead to prolonged outages, impacting service availability and customer trust.
- Data Loss: Lack of real-time data replication and protection increases the risk of data corruption or loss, potentially causing irreparable damage to critical business information.
- Complex Recovery Processes: Manual or fragmented disaster recovery processes increase recovery time objectives (RTO) and recovery point objectives (RPO), making it difficult to meet stringent compliance and service-level agreements (SLAs).
- Increased Costs: Unplanned downtime and inefficient recovery translate into financial losses, potential regulatory penalties, and reputational damage, all of which could be mitigated with a resilient, integrated solution.

Hitachi Vantara has engineered the Hitachi Integrated Systems solution with VMware Cloud Foundation (VCF), the first value-added OEM (VAO) solution for VCF designed from the ground up to offer 100% data availability, full- stack hardware and software lifecycle management, and immutable snapshots for enhanced cyber resiliency.

VMware Cloud Foundation streamlines and optimizes infrastructure management, enabling faster innovation, and enhancing agility for both traditional virtual machines (VMs) and modern container workloads. It delivers a consolidated, software-defined approach to building and managing private and hybrid clouds, offering significant advantages in efficiency, operational control, and cost.

By elevating Hitachi Virtual Storage Platform One Block (VSP One Block) as principal storage for both the Management Workload Domain and Virtual Infrastructure (VI) Workload Domains, integrating into automation and orchestration tools such as Ansible, VCF Automation, VCF Operations, and VMware Live Site Recovery (VLSR, formerly Site Recovery Manager), true multi-site cloud resiliency is achieved across multiple availability zones and across multiple regions.

At its root, a 3 Datacenter (3DC) VMware Metro Storage Cluster with VMware Live Site Recovery architecture is built to ensure that operations in VCF Workload Domains, as well as in standalone VMware vSphere Foundation (VVF) clusters, remain online even in the event of a failure at one location. With this architecture, operations are carefully coordinated across three geographically distinct data centers so that they can collectively manage client requests, maintain data consistency, and support failover, all while optimizing performance. The 3 Data Center architecture is often implemented to enhance system availability and ensure robust disaster recovery capabilities.

Here's a summary of use cases offered by this 3DC architecture to help increase cloud resiliency and application availability within each respective VCF Workload Domain.

- Full Site outage protection (caused by power outages, planned maintenance, etc.) with automated site failover.
- Storage-related issues (controller failure, disk failure, path failure, etc.) with automated path/load re-balancing.
- ESXi host issues (storage connectivity, networking connectivity, power, etc.) with automated recovery of VMs on healthy ESXi hosts using placement policies such as affinity rules.
- Full region outage protection (caused by natural disasters, terrorist attacks, power grid overload, etc.) with orchestrated recovery plans, for both planned and unplanned events.

Real-world scenarios and their benefits

Here are several practical scenarios that highlight the importance and functionality of the 3DC architecture within a VCF Workload Domain.

Scenario	Explanation
Disaster Recovery (DR)	Imagine that the primary data center experiences a catastrophic event (such as a natural disaster or complete power loss). With workloads spread across three sites, the DR plan kicks in automatically, shifting active workloads to the secondary and tertiary sites with minimal downtime. This ensures business continuity even under severe disruptions.
Planned Maintenance/ Upgrades	When one data center needs scheduled maintenance or upgrades, the load is seamlessly migrated to the other two centers. The orchestration tools monitor the performance and maintain service levels, allowing maintenance without impacting overall availability.
Load Balancing and Performance Optimization	In applications with fluctuating demand, traffic can be distributed among three geographically dispersed locations. This not only enhances performance by reducing latency for end users but also ensures that a sudden spike in load on one site does not overwhelm a single point of the infrastructure.

Scenario	Explanation
Security Isolation and Compliance	Critical applications might have strict compliance requirements. By isolating workloads in separate physical locations, you can enforce stricter controls and security policies that limit exposure while still enabling seamless interconnectivity for application operations.

In each of these scenarios, the 3DC model ensures that if one component or site encounters issues, the overall service remains available and resilient. This leads to a significant reduction in Recovery Time Objective (RTO) and Recovery Point Objective (RPO) metrics.

Here's a breakdown of its features in relation to the availability of sites 1, 2, and 3:

Site 1 (Primary Site, Availability Zone 1 in Region 1):

- Acts as the main data center and manages most of the production workloads.
- Typically designed for high availability and real-time operations.
- In case of failure, operations can be switched to Site 2 (both sites are within the same region)

Site 2 (Secondary Site, Availability Zone 2 in Region 1):

- Synchronous mirroring with Site 1 ensures optimal data synchronization with no data loss during switchover.
- Ensures high availability during a primary site (site 1) outage.
- Note: Both site 1 and site 2 can represent “primary sites” for different workload domains (WLD), with the other site acts as the secondary site for each respective WLD. Both sites are within the same pseudo-cloud region (metro-distanced stretch cluster), and can be viewed as two distinct availability zones (AZs).

Site 3 (Tertiary Site in Region 2):

- Functions as a disaster recovery (DR) site or a backup site.
- Often used for asynchronous data backup or archiving.
- Adds an extra layer of redundancy.
- Ensuring that in extreme scenarios where both Site 1 and Site 2 become unavailable, Site 3 can still offer some operational continuity, although recovery time objectives and recovery point objectives might be lower compared to Site 2 in Region 1.

This architecture ensures seamless failover capabilities across the sites and is widely regarded for providing superior resilience in critical environments.

Use this design and configuration guide to bring up the solution, which integrates the following:

- Hitachi Integrated Systems
- VMware Cloud Foundation
- Hitachi Virtual Storage Platform

- Hitachi global-active device
  - Provides active-active replication for high availability across sites.
  - Ensures seamless failover and continuous application availability.
- Hitachi Universal Replicator (UR)
  - Offers asynchronous replication for disaster recovery.
  - Ensures data consistency and recovery across geographically dispersed sites.
- VMware Live Site Recovery Manger (VLRM) integrated with Hitachi Storage Replication Adapter (SRA)
  - Automates disaster recovery workflows.
  - Integrates with global-active device and HUR for seamless failover and failback operations.
- Hitachi Unified Compute Platform (UCP) Advisor
  - Simplifies infrastructure management by providing a unified interface for compute, storage, and network resources.
  - Integrates with VMware vSphere Lifecycle Manager (vLCM) for streamlined lifecycle management.
  - Enables policy-based provisioning and orchestration of infrastructure components.

The following diagrams show high-level overviews of this 3DC solution. The cities in this diagram are just for example only.

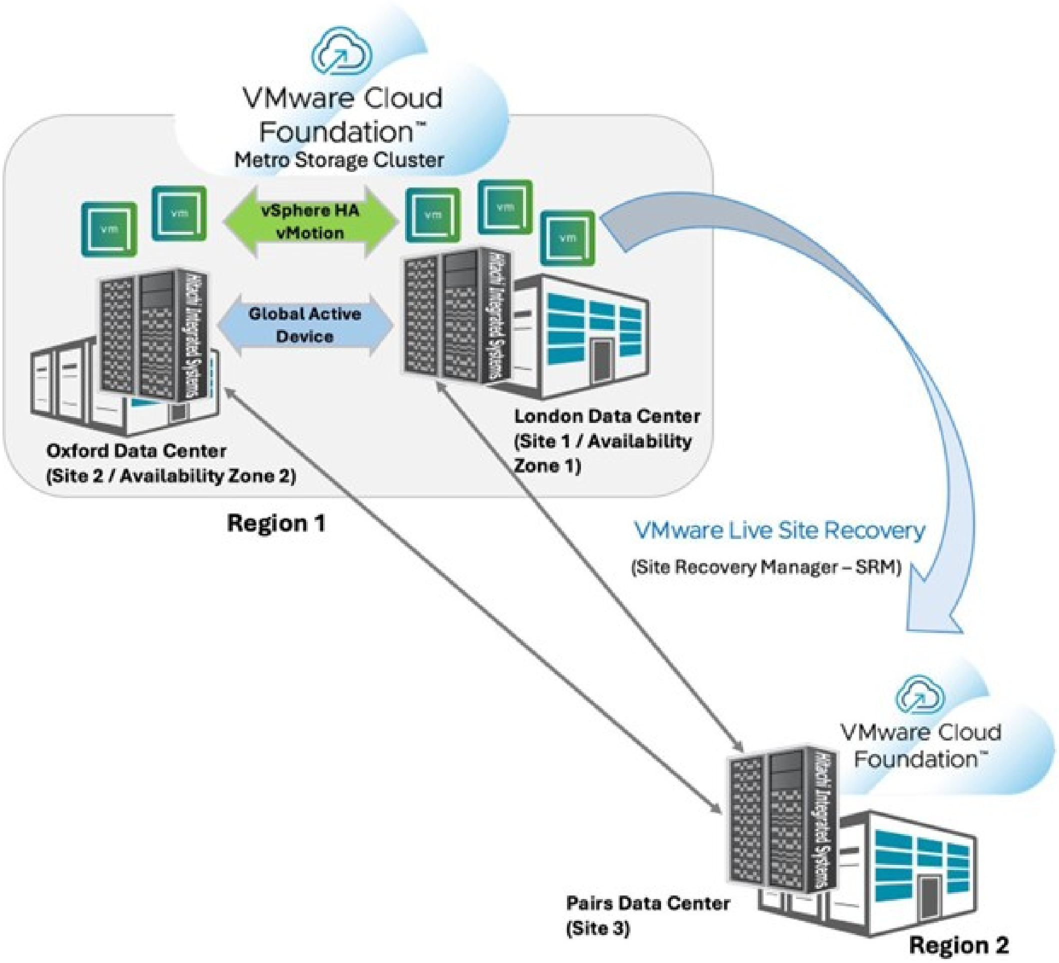
VMware Cloud Foundation™
Metro Storage Cluster
vm
vSphere HA
vMotion
Global Active Device
Hitachi Integrated Systems
Oxford Data Center
(Site 2 / Availability Zone 2)
London Data Center
(Site 1 / Availability Zone 1)
Region 1
VMware Live Site Recovery
(Site Recovery Manager – SRM)
VMware Cloud Foundation™
Pairs Data Center
(Site 3)
Region 2

Hitachi Integrated Systems solution with VMware Cloud Foundation is a highly configurable integrated infrastructure in which server, network, and storage can be scaled independently to optimize performance and eliminate overprovisioning costs. The following describes more details of the diagram:

- Hitachi Integrated Systems are deployed to 3 datacenters.
  - Each Hitachi Integrated Systems platform includes at least one Virtual Storage Platform
- vSphere Metro Storage Cluster is configured between site 1 and site 2, considered to be two availability zones within one region.
  - Global-active device is configured in between site1 VSP and site 2 VSP to provide a stretched active-active VMware datastore.
  - The distance between site 1 and site 2 is usually less than 100 km or 5 ms response time.
  - The vSphere HA become available in between site 1 and site 2. In case of a site failure (either site 1 or site 2) vSphere HA will bring up failed VMs on an available site automatically.
- vMotion is also available between site 1 and site 2.
  - VMware Live Site Recovery Manager (VLSR) is configured in between site 1 and site 3, located in a geographically-dispersed remote region.
- Hitachi Universal Replicator (HUR) is configured in between site 1 and site 3 to provide asynchronous data replication from site 1 to site 3
- HUR delta re-sync is also configured in between site 2 and site 3 to provide standby replication link in case of site 1 failure.
- When both site 1 and site 2 fail (large London metro area for example) VLSR is used to recover the VM workloads in site 3.

As you demand ever-faster delivery of new business services, there is the complexity and cost of deploying and managing the technology resources to support them. While moving some workloads to the public cloud is common practice, enterprise workloads, including business-critical and mission-critical applications, are untouchable and reign supreme on-premises. Many IT departments spend almost a quarter of their time and resources evaluating and installing increasingly disparate hardware components. Furthermore, the lack of a unified management framework and the need for highly specialized individuals who can design, configure, optimize, test, and manage each component, increases your cost, complexity and risk.

Hitachi Vantara and VMware by Broadcom have collaborated to address these on-premises challenges by introducing the Hitachi Integrated Systems solution with VCF. The purpose of a VCF solution is to bring value to our customers through reduced time-to-value (TTV), reduced total cost of ownership (TCO), increased agility, and enterprise-grade security and compliance. This is achieved by aligning the solution design with best practices according to VCF VMware Validated Solutions (VVS) specifications. It delivers Hitachi and VMware firmware and software updates within a committed release cycle while streamlining the lifecycle management experience for updating critical components, providing monitoring of these critical components across the full stack, and automating key operational tasks to maximize consistency and increase IT agility.

Our solution helps you run all virtualized or containerized workloads and business applications. Ordered as a single quote with the full BOM (hardware and software), and with its pre-validated, pre-built, and pre-assembled stack of compute, storage, and networking components, the solution delivers a hybrid cloud environment with predictable high performance, resiliency, and reliability. Coupled with a tightly integrated software architecture that enables full-stack deployment, operations, management, orchestration, monitoring, diagnostics, support, and lifecycle management, the Hitachi Integrated Systems solution with VCF is the simplest and most reliable path to private cloud, transforming a traditional data center into a modernized software-defined data center without the pain of dealing with long project delays and cost overruns.

In addition, while other VCF implementations only leverage VMware vSAN as the sole software-defined storage option, Hitachi Vantara has optimized its solution by natively integrating its enterprise-grade Virtual Storage Platform SAN storage into the management framework and across the data plane to achieve the highest level of performance and availability, and the most efficient capacity utilization.

This unique differentiator provides the flexibility of using the right storage platform for each individual workload being deployed in your hybrid cloud, to meet your business needs and service level objectives (SLOs), and without being constrained by the design limitations of plain-vanilla Hyperconverged Infrastructure (HCI) and Disaggregated HCI. The solution with the 3DC cluster addresses the challenges of deploying and managing on-premises technology resources, while maintaining true cloud resiliency normally only found with hyperscalers. By offering a comprehensive, integrated, and flexible solution, it ensures high performance, reliability, efficiency, and resiliency, transforming traditional data centers into modernized, software-defined private clouds. Finally, our Global Support and Services team provides comprehensive support for Full Stack software and hardware under the umbrella of One Support.

Note: From VCF 9.0, the Management Domain can be deployed entirely on VSP without requiring vSAN. Conversely, in VCF 5.2.1, achieving this same capability involves using the import tool to convert a VSP- connected vCenter and its clusters into a brownfield-ingested VCF management domain. In both scenarios, vSAN is unnecessary, meaning that our 3DC protection across all WLD Domains —including Management— relies solely on VMFS-FC datastores via our Hitachi VSP global-active device technology.

More details can be found at Importing the existing vSphere environment.

Solution overview

Gone are the days of just having DIY private clouds with multi-tier Lego blocks plagued with inconsistencies and project delays. Hitachi Integrated Systems solution with VCF is a true turnkey on-prem cloud solution, delivering an unprecedented out-of-box cloud experience, combining the scale and agility of public cloud with the security, resiliency, and performance of private cloud.

This solution integrates VMware Cloud Foundation (VCF) with Hitachi infrastructure and Hitachi software. The Three Data Center cluster is a robust high-availability and disaster recovery solution designed to ensure continuous data availability and protection across three geographically separated data centers (two availability zones within one metro region, and another geographically dispersed site in a remote region). This setup leverages advanced technologies such as VMware vSphere Metro Cluster (vMSC), global-active device, Hitachi Universal Replicator (HUR), and VMware Live Site Recovery to provide seamless data replication, failover, and disaster recovery capabilities.

The following illustration shows an overview of the 3DC architecture.

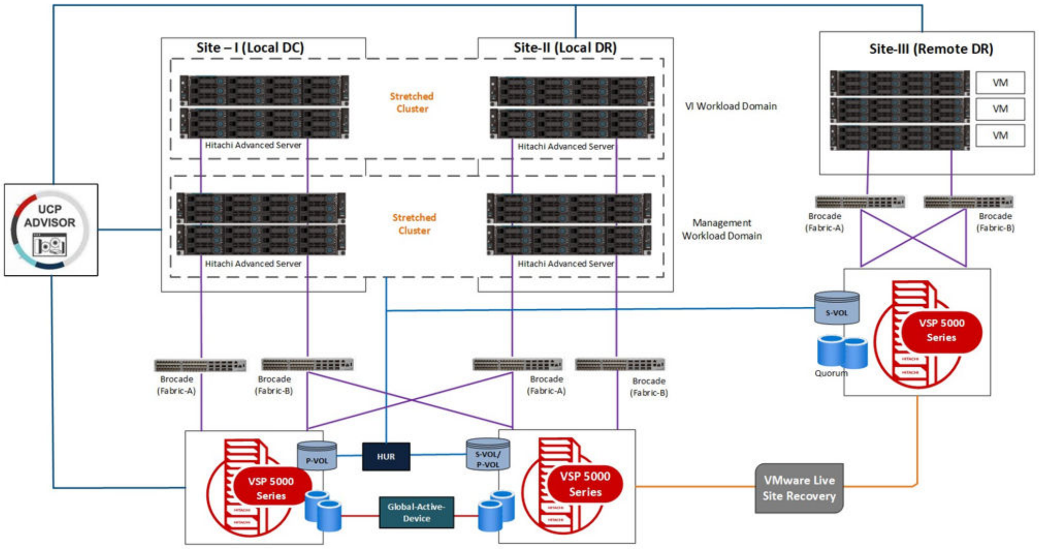

This illustration gives high level information about the 3DC architecture and its associated components which include global-active device/UR/VLSR/UCP Advisor.

The 3DC solution uses the Hitachi Integrated Systems solution with VCF and VSP storage to establish a production center, a local DR center, and a remote DR center. Data from the production center is synchronously replicated to the local DR center and asynchronously replicated to the remote DR center. The local DR center has equivalent service processing capabilities as the production center. Applications can be seamlessly switched to the local DR center without any data loss, achieving zero RPO and second-level RTO. In the event of rare natural disasters, which render both the production center and local DR center unavailable, applications can be transferred to the remote DR center. Routine DR drills ensure that applications can be recovered in the remote DR center within the acceptable time limit, maintaining service continuity and second-level RPO. Compared to solutions that deploy only a local DR center or a remote DR center, the multi-site redundant solution combines their benefits to handle larger-scale disaster scenarios. For smaller-scale regional disasters and larger-scale natural disasters, the DR system can respond more swiftly to prevent service data loss and achieve lower RPO and RTO. Consequently, the multi-site redundant solution is widely adopted.

The Hitachi Integrated Systems solution with VCF includes automation that enables the deployment of an entire cloud infrastructure in hours, not weeks or months, accelerating the time to value with a rapid and repeatable deployment process across all of your datacenter sites.

When deployed, the co-engineered and integrated management plane provides the following automated, policy-based IT operations:

- Move your workload across data centers to meet changing business needs.
- Manage your applications across private and public cloud from a common toolset to simplify operations.
- Scale your data center without increasing IT headcount.
- Automate your data center with policies to run IT at the speed of your business.
- Automate software-defined storage deployment with both VMware vSAN and Hitachi VSP.
- Automate software-defined networking deployment with VMware NSX.

The following illustration shows combinations offered with the Hitachi Integrated Systems solution with VCF.

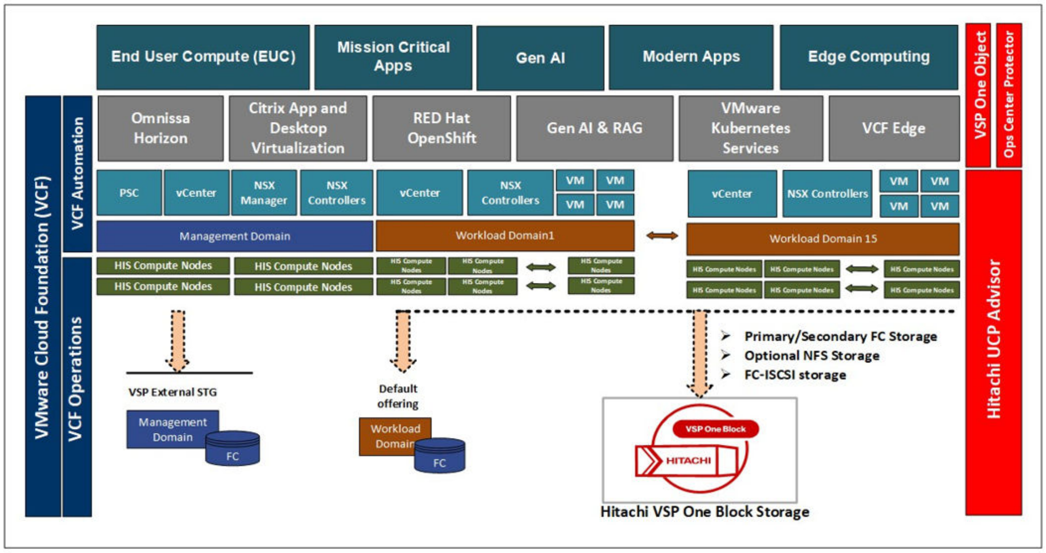

The following illustration shows the storage options for VMware Cloud Foundation Workload Domains.

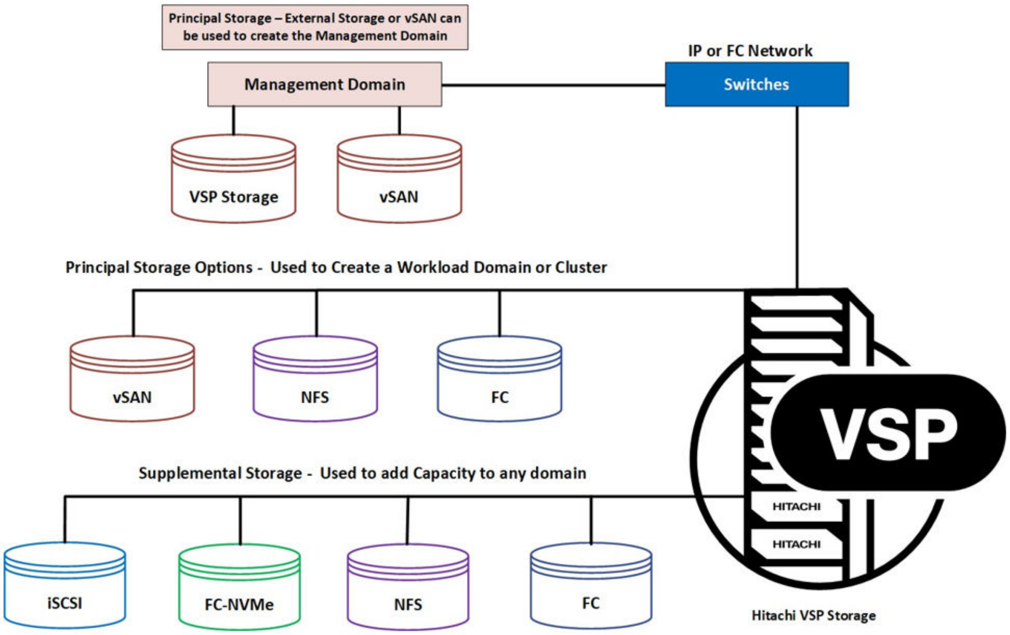
Principal Storage – External Storage or vSAN can be used to create the Management Domain
Management Domain
IP or FC Network
Switches
VSP Storage
vSAN
Principal Storage Options - Used to Create a Workload Domain or Cluster
vSAN
NFS
FC
VSP
Supplemental Storage - Used to add Capacity to any domain
HITACHI
HITACHI
iSCSI
FC-NVMe
NFS
FC
Hitachi VSP Storage

Note that support for vVols is deprecated as of vSphere 9.0, and will no longer be supported when vSphere 9.1 is released. Broadcom is no longer accepting new vVol storage certifications so not all VSP models can use vVols.

Overview of global-active device

Global-active device is a data mirroring technology developed by Hitachi that ensures high availability and disaster recovery for storage systems. Global-active device enables synchronous replication of data between storage systems located in different geographic locations, allowing for real-time data consistency and integrity. This technology is crucial for businesses that require continuous data availability and the ability to quickly recover from unexpected failures.

Global-active device provides a robust solution for maintaining high availability and disaster recovery in enterprise storage environments. By enabling active-active storage replication, global-active device ensures that data is always up-to-date and accessible across multiple data centers. This capability allows for seamless failover and failback operations, minimizing downtime and ensuring business continuity. Global-active device enables you to create and maintain synchronous, remote copies of data volumes.

A virtual storage machine is configured in the primary and secondary storage systems using the actual information of the primary storage system, and the global-active device primary and secondary volumes are assigned the same virtual LDEV number in the virtual storage machine. This enables the host to see the pair volumes as a single volume on a single storage system, and both volumes receive the same data from the host.

A quorum disk, which can be located in a third and external storage system or in an iSCSI server, including cloud-based virtual servers (for example, running in Azure), is used to monitor the global-active device pair volumes. The quorum disk acts as a heartbeat for the global-active device pair, with both storage systems accessing the quorum disk to check on each other. A communication failure between systems results in a series of checks with the quorum disk to identify the problem for the system able to receive host updates.

Alternate path software on the host runs in the Active/Active configuration. While this configuration works well at campus distances, at metro distances Hitachi Dynamic Link Manager is required to support preferred/nonpreferred paths and ensure that the shortest path is used.

If the host cannot access the primary volume (P-Vol) or secondary volume (S-Vol), host I/O is redirected by the alternate path software to the appropriate volume without any impact to the host applications.

Global-active device provides the following benefits:

- Continuous server I/O when a failure prevents access to a data volume
- Server failover and failback without storage impact
- Load balancing through migration of virtual storage machines without storage impact

Key features of global-active device include:

- Synchronous Replication: Ensures real-time data mirroring between primary and secondary storage systems, maintaining data consistency and integrity.
- High Availability: Provides continuous server I/O operations even during failures, ensuring that applications remain operational.
- Disaster Recovery: Facilitates rapid recovery from unexpected failures by enabling seamless failover and failback without impacting storage.
- Load Balancing: Allows for the migration of virtual storage machines without impacting storage, optimizing resource utilization.
- Active-Active Design: Enables production workloads on two systems simultaneously, ensuring full data consistency and protection.
- Zero Recovery Time Objectives: Offers zero downtime and no data loss for applications that require continuous operations.
- Global Storage Virtualization: Allows read and write copies of the same data across two systems or geographic locations.
- Simplified Operations: Automates high availability and simplifies distributed system design. Fault-Tolerant Infrastructure: Provides failover clustering and server load balancing without impacting storage.

Global-active device components

A typical global-active device system consists of storage systems, paired volumes, a consistency group, a quorum disk, a virtual storage machine, paths and ports, alternate path software, and cluster software.

Global-active device is an essential technology for organizations that require robust data protection and high availability across geographically dispersed data centers. It enhances the resilience and reliability of storage –, ensuring that critical data is always available and protected.

The following illustration shows the components of a typical global-active device system.

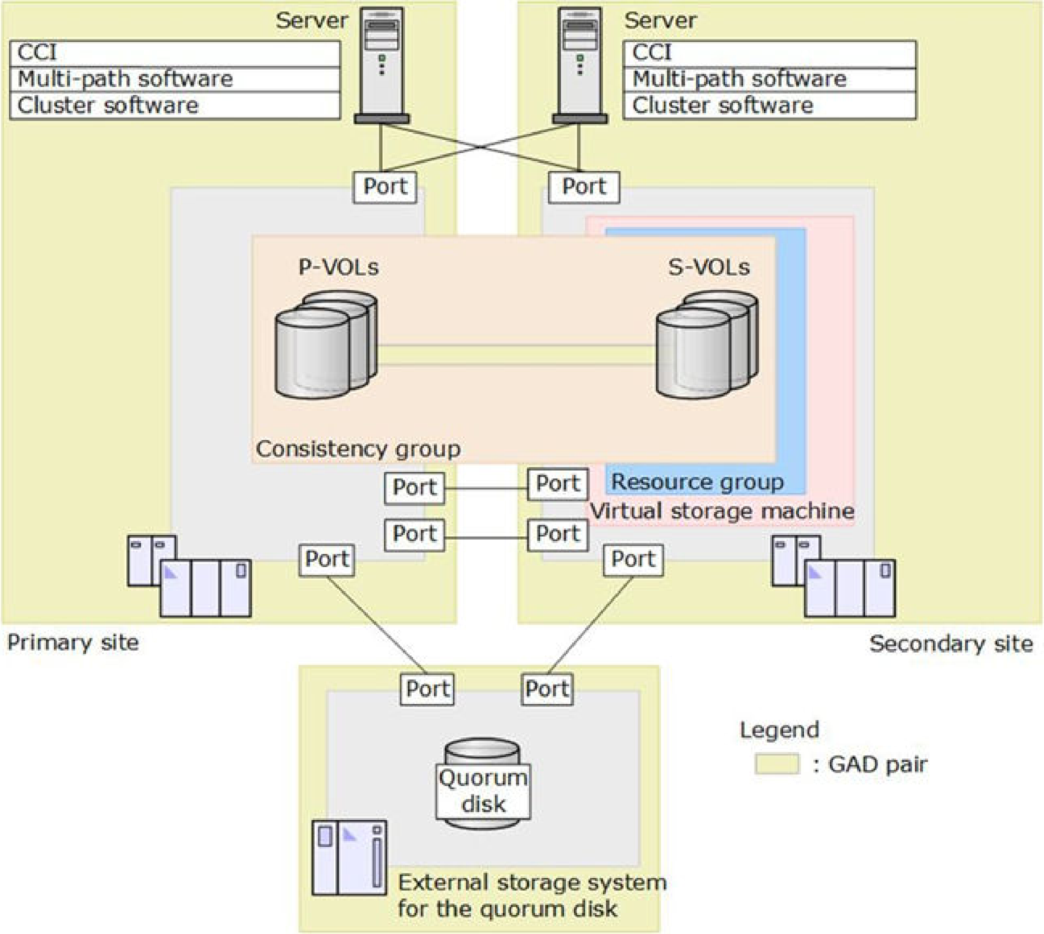

Storage system

The supported combination of storage systems at the primary and secondary sites differs, depending on the models and microcode versions of the storage systems. For details, see Requirements and restrictions.. An external storage system or iSCSI-attached or cloud-based server, which is connected to the primary and secondary storage systems using Universal Volume Manager, is required for the quorum disk.

Paired volumes

A global-active device pair consists of a P-Vol in the primary storage system and an S-Vol in the secondary storage system.

Consistency group

A consistency group consists of multiple global-active device pairs. By registering global-active device pairs to a consistency group, you can resynchronize or suspend the global-active device pairs by consistency group.

Quorum disk

The quorum disk is used to determine the storage system on which server I/O should continue when a storage system or path failure occurs. The quorum disk is virtualized from an external storage system that is connected to both the primary and secondary storage systems. Alternatively, a disk in an iSCSI server, including cloud- based virtual servers (for example, running in Azure), can be used as a quorum disk if the server is supported by Universal Volume Manager.

Virtual storage machine

A virtual storage machine (VSM) is configured in the secondary storage system with the same model and serial number as the (actual) primary storage system. The servers treat the virtual storage machine and the storage system at the primary site as one virtual storage machine.

You can create global-active device pairs using volumes in virtual storage machines. When you want to create a global-active device pair using volumes in VSMs, the VSM for the volume in the secondary site must have the same model and serial number as the VSM for the volume in the primary site.

Paths and ports

Global-active device operations are carried out between hosts and primary and secondary storage systems that are connected by data paths composed of one or more physical links.

The data path, also referred to as the remote connection, connects ports on the primary storage system to ports on the secondary storage system. Both Fibre Channel and iSCSI remote copy connections are supported. The ports have attributes that enable them to send and receive data. One data path connection is required, but you should use two or more independent connections for hardware redundancy.

Overview of Hitachi Universal Replicator (UR)

Hitachi Universal Replicator (UR) offers asynchronous remote copy for long-distance data replication, integrates microcode-based functionality for seamless operation, ensures disaster recovery with near-cloud setups, provides high resilience and performance for enterprise-scale applications, uses a disk-based journal for data consistency, supports multi-target and cascade configurations, delivers low-latency performance, is suitable for various use cases, enhances operational resilience, optimizes data management, drives IT agility, offers scalability for growing business needs, and supports sustainability efforts by reducing energy consumption and carbon footprint.

Hitachi Universal Replicator (UR) presents a solution to avoid cases when a data center is affected by a disaster that stops operations for a long period of time. In the Universal Replicator system, a secondary storage system is located at a remote site from the primary storage system at the main data center, and the data on the primary volumes (P-Vol s) at the primary site is copied to the secondary volumes (S-Vol s) at the remote site asynchronously from the host write operations to the P-Vol s. Journal data is created synchronously with the updates to the P-Vol to provide a copy of the data written to the P-Vol . The journal data is managed at the primary and secondary sites to ensure the consistency of the primary and secondary volumes.

The redundancy provided by the RAID configuration (for example, RAID1 or RAID5) enables recovery from a P-Vol read failure. The primary storage system never reads the Universal Replicator S-Vol for data recovery.

Replication operations

Remote replication for a UR pair is accomplished using the master journal volume on the primary storage system and the restore journal volume on the secondary storage system. As shown in the following figure, the P-Vol data and subsequent updates are transferred to the S-Vol by obtain journal, read journal, and restore journal operations involving the master and restore journal volumes.

The following illustration shows a UR logical diagram.

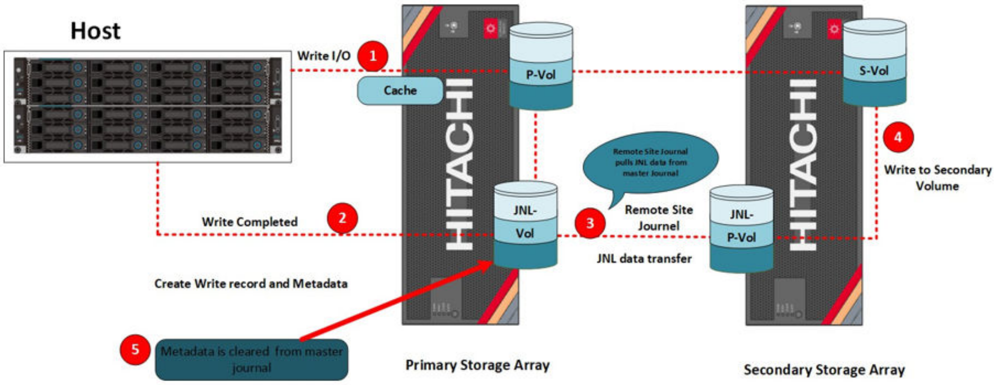

1. Data is being written to the P-Vol and new metadata written in the master journal volume.
2. The master journal data update is completed and host I/O completed.
3. The remote site journal pulls JNL data from the master journal.
4. Data is written to the secondary volume in the correct order and the sequence number is sent (notified) to the primary storage system.
5. The primary storage system clears the metadata associated with the notified sequence number from its master JNL volume.

Key Features and Benefits of UR

- Disk-Based Journal: HUR uses a disk-based journal to manage replication data, which helps in maintaining data consistency and integrity during the replication process.
- Multi-Target and Cascade Configurations: HUR supports multi-target and cascade configurations, allowing for flexible and scalable disaster recovery solutions. This means you can replicate data to multiple locations or in a cascading manner to ensure comprehensive data protection.
- Low Latency Performance: The software is designed to deliver low-latency performance, ensuring that the replication process does not significantly impact the performance of the primary storage system.
- Versatile Use Cases: HUR is suitable for various use cases, including data migration, business continuity, and disaster recovery. It provides flexibility to plan and implement different remote data center configurations.

- Operational Resilience: By ensuring data availability and integrity across multiple sites, HUR enhances operational resilience, making it a reliable solution for critical enterprise applications.
- Asynchronous Replication: Ensures data consistency between primary and secondary storage systems, even across long distances.
- Journal-Based Replication: Uses journal volumes to manage data updates and maintain consistency.
- Disaster Recovery: Provides robust protection against data center failures by replicating data to remote sites.
- Scalability: Supports large-scale environments with multiple storage systems and consistency groups.
- Flexibility: Compatible with various RAID configurations and storage systems.

Solution components

These are the key hardware and software components used in the solution.

Hardware components

The Hitachi Integrated Systems platform uses the following compute, networking and storage options. These are the key hardware components that this solution uses.

- Compute
  - Hitachi Advanced Server HA810 G3
  - Hitachi Advanced Server HA820 G3
  - Hitachi Advanced Server HA805 G3
  - Hitachi Advanced Server HA815 G3
  - Hitachi Advanced Server HA825 G3
- Fibre Channel SAN storage options
  - Hitachi Virtual Storage Platform 5000 series
  - Hitachi Virtual Storage Platform E1090
  - Hitachi Virtual Storage Platform One Block (VSP One Block 24, VSP One Block 26, VSP One Block 28)
- Cisco switches
  - Cisco Nexus C92348GC-X
  - Cisco Nexus C93180YC-FX3
  - Cisco Nexus C93600CD-GX
  - Cisco Nexus C9316D-GX
- Brocade switches from Broadcom
  - Brocade G720 Fibre Channel SAN Switches

Software components

These are the key software components used in this environment.

Use VMware Cloud Foundation to deploy and run a private cloud on top of the Hitachi Integrated Systems solution with VCF with Hitachi Virtual Storage Platform. It provides an integrated cloud infrastructure (compute, storage, networking and security) and cloud management service to run enterprise applications in private and public cloud environments.

See https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html to learn more about VMware Cloud Foundation.

The following table shows the VCF software stack and versions including VCF and UCP Advisor.

Software Component	Version	Build Number
VMware Cloud Builder	5.2.1	24307856
VMware SDDC Manager	5.2.1	24307856
VMware vCenter Appliance	8.0Update 3c	24305161
VMware ESXi	8.0Update 3b	24280767
VMware vSAN	8.0Update 3	24022510
VMware NSX	4.2.1	24304122
VMware Aria Suite	8.18	24029603
Hitachi Unified Compute Platform Advisor	4.7.0	-
VMware Live Site Recovery	9.0.2.8708	24401766
VCF Import Tool	5.2.1.0	24307788

See http://compatibility.hitachivantara.com for more information.

Solution design

The Hitachi Integrated Systems solution with VCF can start as a single rack and extend to multi-rack scenarios, supporting unlimited racks and nodes within VMware maximums. Each rack can accommodate up to 32 compute nodes, 2 leaf switches for network connectivity, spine switches for inter-rack connectivity, and a management network for out-of-band traffic.

Flexibility is a key benefit, allowing mixed and matched compute nodes for varied requirements. For example, you can add hybrid systems for more vSAN capacity or graphic resources.

SAN-based datastores are another advantage, supporting Hitachi Virtual Storage Platform for extended storage space and different storage tiers. This setup offers options for storage based on workload domains.

The following illustration shows an architecture view of the solution.

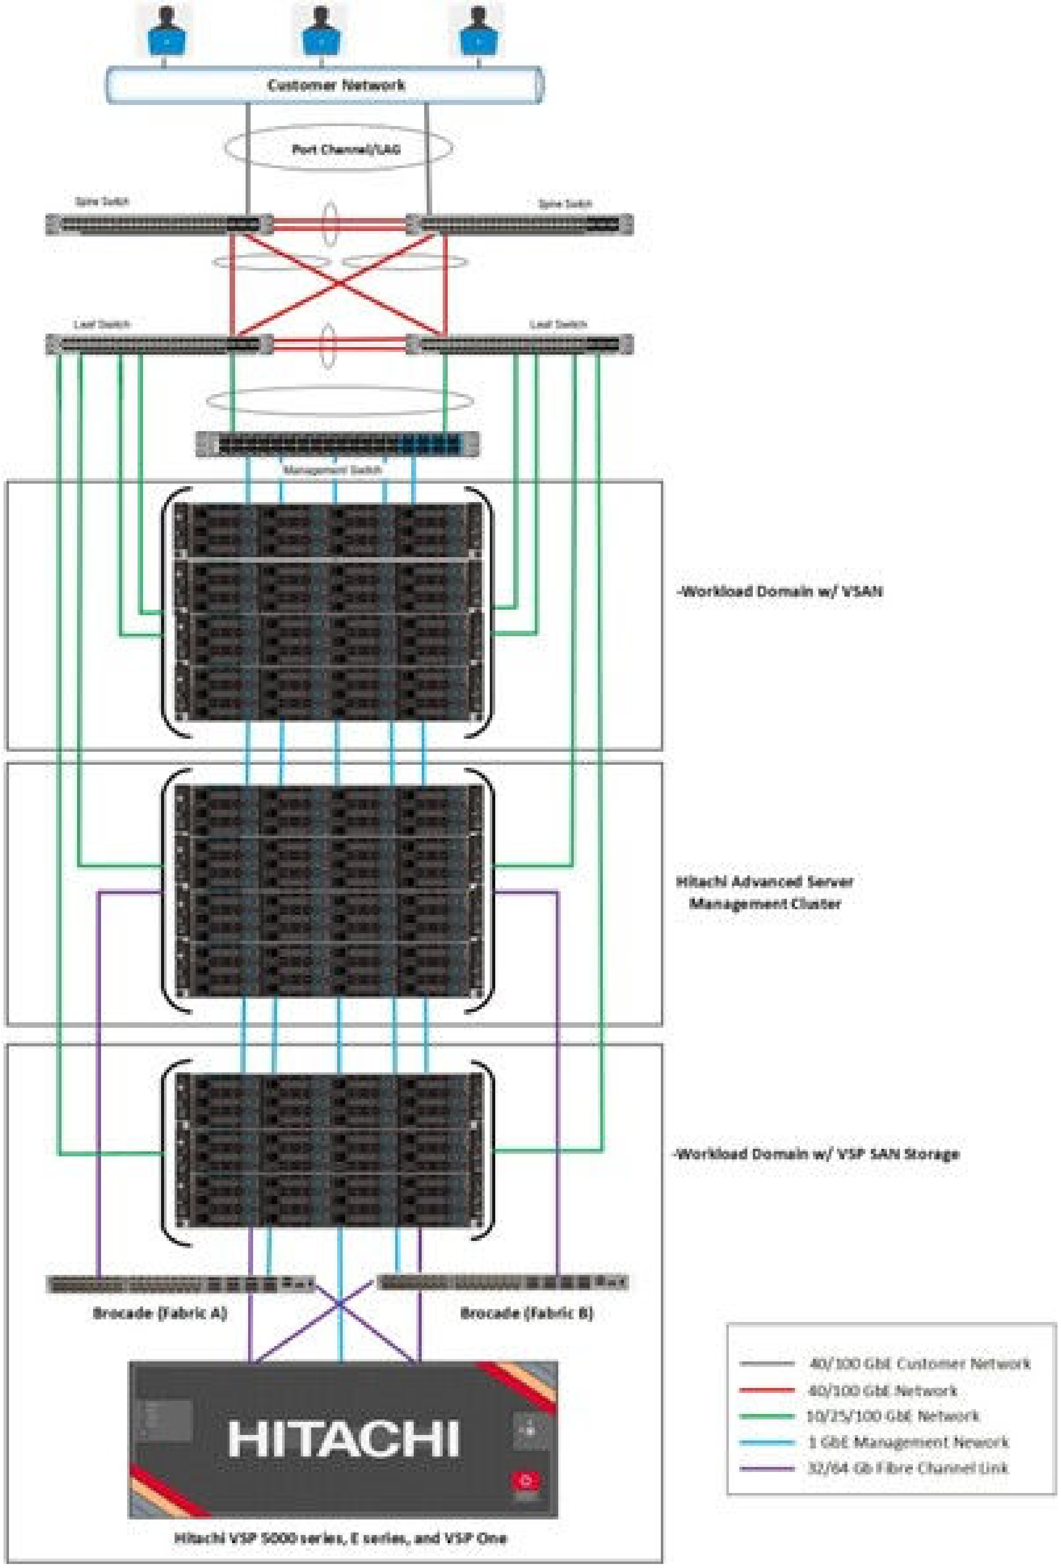

The illustration shows a high-level physical design of the solution with the following hardware:

- Cisco spine Ethernet switches
- Cisco leaf Ethernet switches
- Cisco management switch
- Hitachi Advanced Server compute nodes
- Hitachi Virtual Storage Platform

Solution validation

This section describes the test scenarios completed to validate the solution.

Note: Testing of this configuration was in a lab environment. Many factors 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.

Global-active device test scenarios

1. Host down (vSphere High Availability).

   This test is performed by immediate shutdown from BMC/ILO or pulling the power cable. In the following illustration, P-Vol represents the primary global-active device replicated volume in primary site and S-Vol represents the secondary volume in the secondary site.

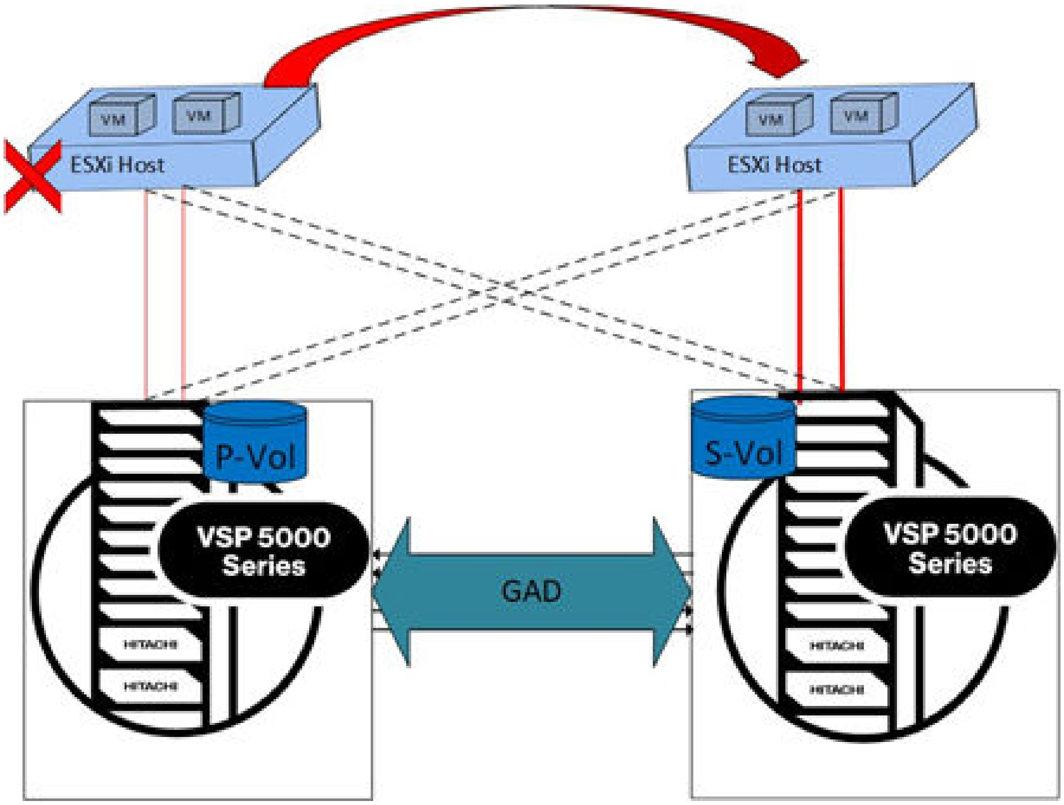

Observations:

- vSphere HA detects ESXi host down and triggers failover to another ESXi host.
- If there is an alive ESXi host in the Primary Site then VM is restarted there, otherwise on a ESXi host in the Secondary Site.
- VMware NMP/PSP uses ALUA information to select Active/Optimized path (which is Local Path).
- Because all ESXi hosts on both sites can access both the P-Vol and S-Vol on the storage system, VMware vSphere High Availability will restart virtual machines on available ESXi hosts at different sites during host failures.

2. Fibre Channel single path failure (local path partially down).

Fibre channel Link path down was simulated by removing the zone configuration from Fibre Channel switches or unplugging physical cables/ports.

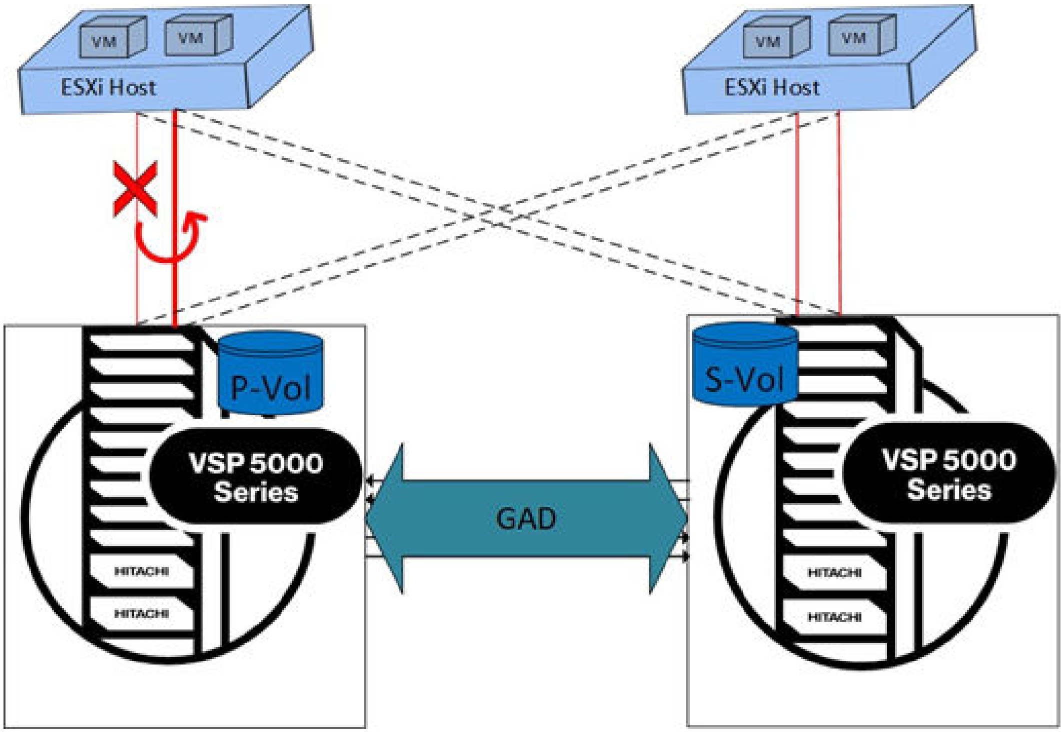

Observations:

- VMware NMP/PSP uses ALUA information to select the remaining Active/Optimized local path.
- When the path is restored, it automatically becomes available and active for use by the ESXi host.

3. Local path completely down (all active paths to local storage system fail for a single host on either site).

Fibre Channel Link path down was simulated by removing the zone configuration from Fibre Channel switches or unplugging physical cables/ports.

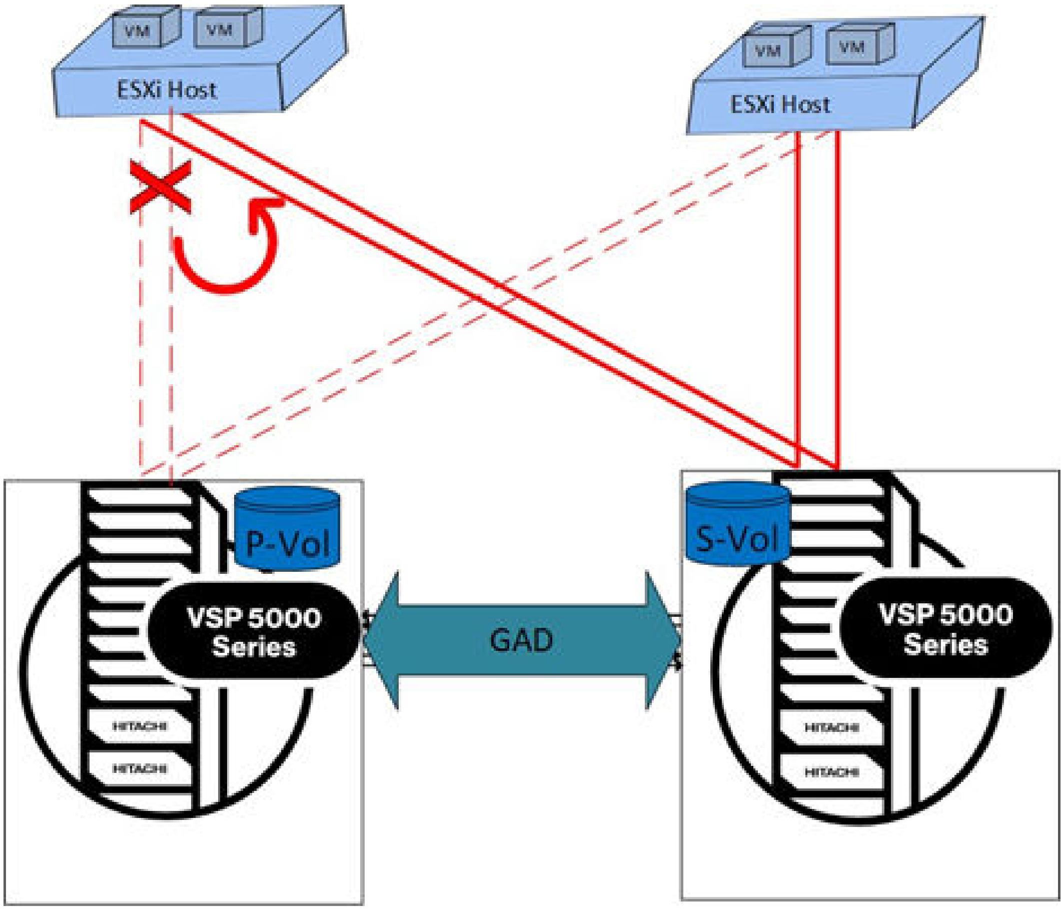

Observations:

- VMware NMP/PSP uses ALUA information to select the Active/Non-Optimized Cross Paths.
- Primary Site ESXi host IO goes to the S-Vol in the secondary site.
- When the path is restored, it automatically becomes available and active for use by the ESXi host.

4. Remote path and replication path down. (all paths down occur in any ESXi host in the cluster).

   Remote and replication path down was simulated by removing the zone configuration from Fibre Channel switches or unplugging physical cables/ports.

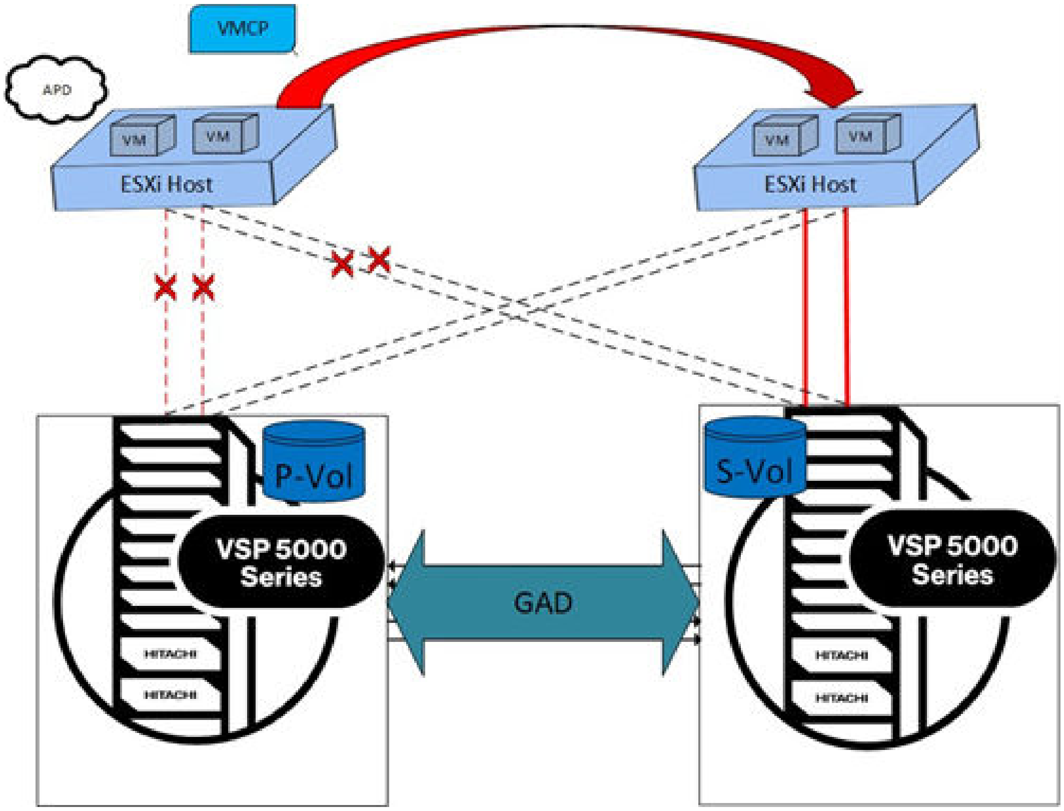

Observations:

- VMware HA detects all paths down and triggers failover to another ESXi host.
- If a functioning host is available at the Primary Site, the VM will be restarted there; otherwise, it will be restarted on a host at the Secondary Site.
- After failover, VMware NMP/PSP uses ALUA information to select local active paths.

5. Primary storage down (Site 1 storage failure).

Primary storage down was simulated by physically shutting down storage servers or hardware or disconnect power sources if necessary.

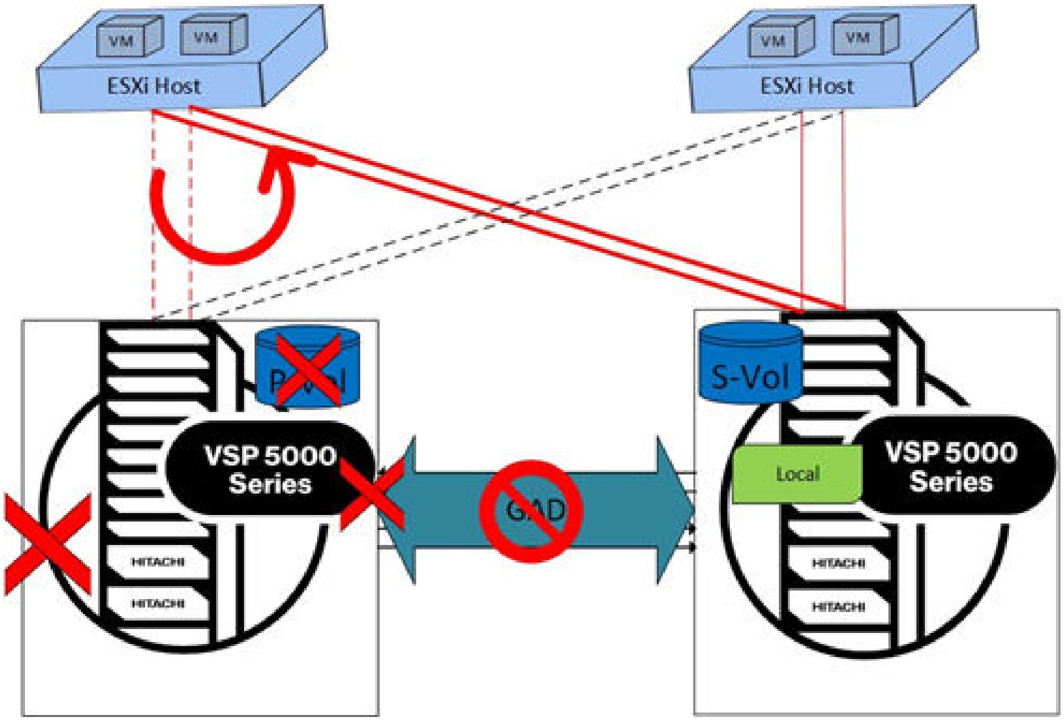

Observations:

- When ESXi hosts at Site 1 fail, vSphere High Availability automatically restarts the virtual machines on ESXi hosts at Site 2, No disruption is seen on the virtual machine.
- On the Site 1 ESXi hosts, active paths to the primary storage system are reported dead.

The following are the steps to recover from Storage Failover (Failback). The following assumes that global-active devices have triggered a storage failover to the S-Vol on Site 2. After Site 1 storage has been restored and is ready to serve production data, the following steps are required for restoring the metro storage cluster to its original state.

The recovery procedure is divided into two procedures depending on whether the global-active device is in a P- Local or S-Local state.

The following is the recovery procedure from S-Local after fixing the Site 1 failure:

a. Re-establish replication of S-VOL data to P-VOL.

b. Execute takeover-recovery (resync) of the global-active device pair.

c. Verify that the pair status is PAIR and swap the copy direction of the global-active device pair.

d. Restart the ESXi hosts in Site 1.

The combination of VMware NMP and ALUA automatically fails paths back to P-VOL on local storage in each site.

See the Global-active Device User Guide for more details.

The following illustration shows the recovery scenario with global-active device.

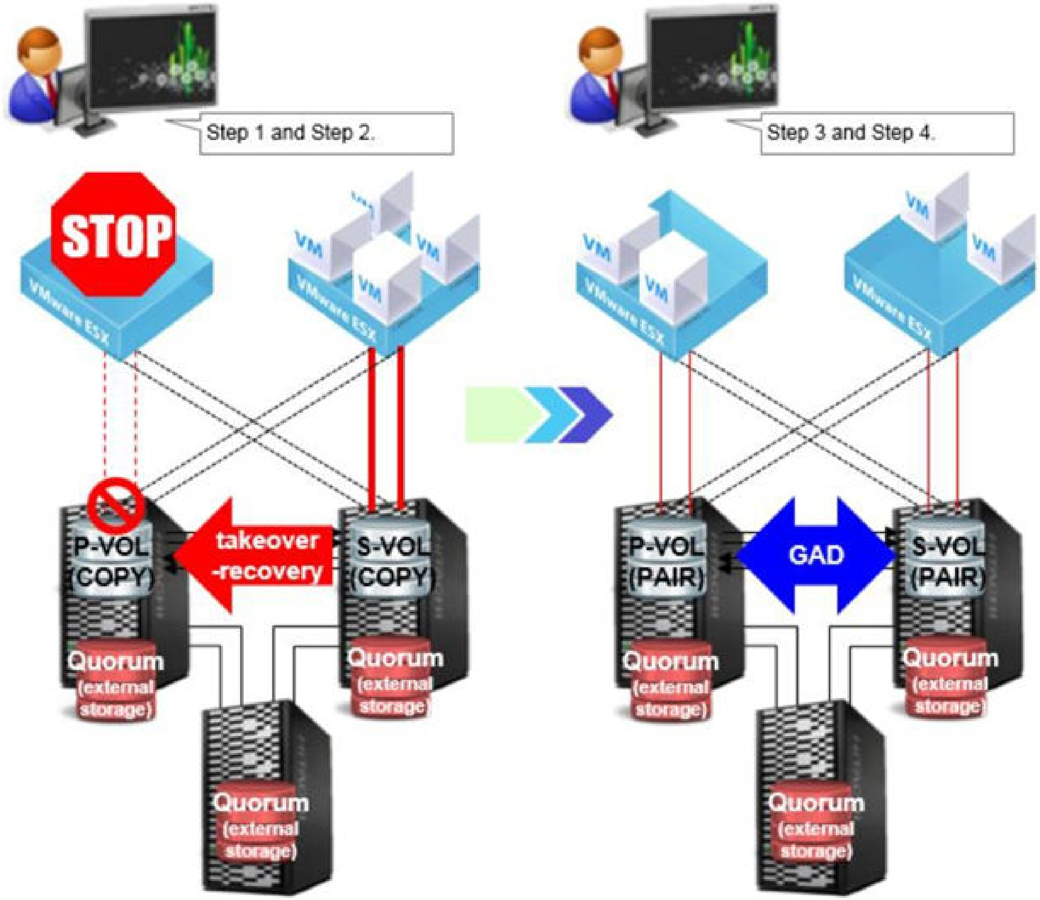

6. Global-active device Quorum down/Path to global-active device Quorum is down (Quorum disk failure).

   Global-active device quorum down was simulated by removing the zone configuration from Fibre Channel switches or unplugging physical cables/ports.

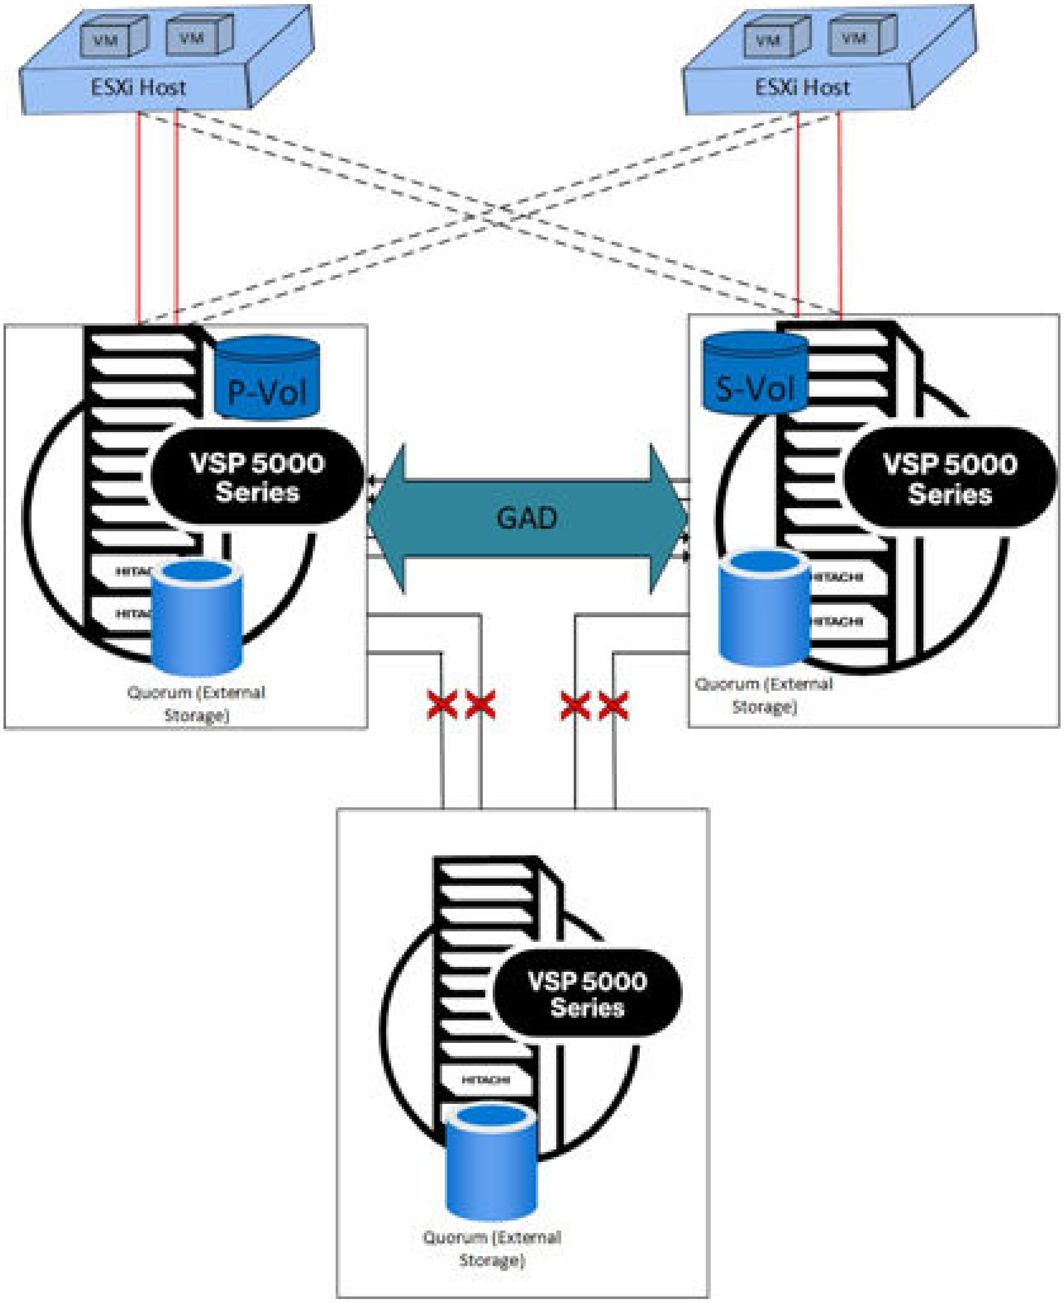

Observations:

- VMs will continue to operate normally.
- P-Vol and S-Vol will continue to replicate normally.
- Even if the quorum disk fails or all paths to the quorum disk are removed, replication between P-Vol and S-Vol will continue and the P-Vol and S-Vol will remain in pair state.

7. Replication link down (Storage Replication link failure).

   Storage Replication link failure is simulated by removing the zone configuration from Fibre Channel switches or unplugging physical cables/ports

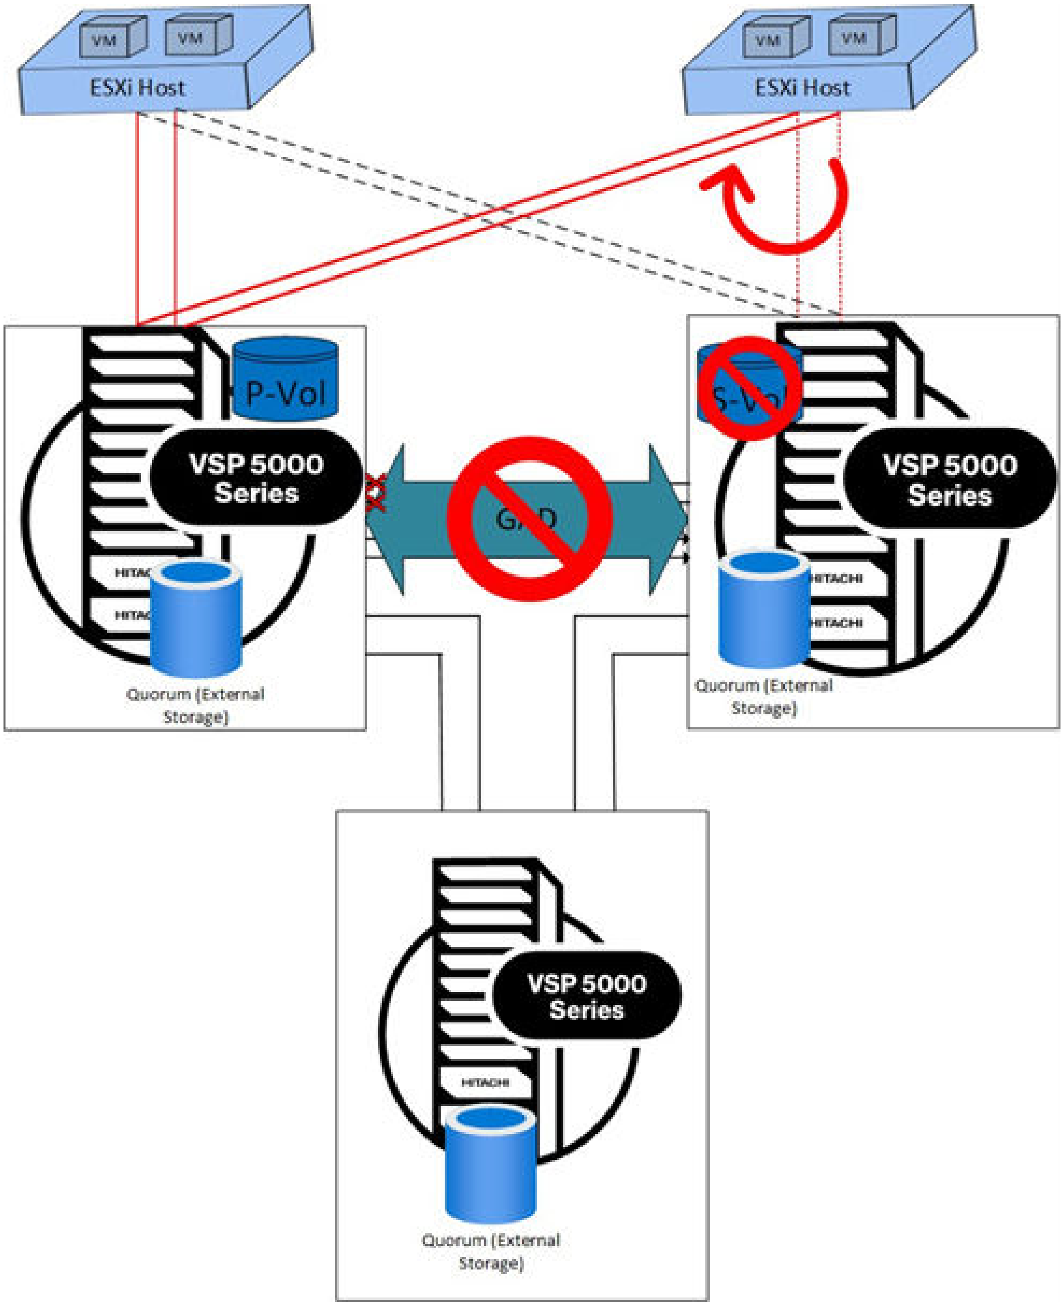

Observations:

- VMs continued to operate through either local preferred paths or remote nonpreferred paths.
- When replication links are interrupted, the default behavior is that whichever P-Vol or S-Vol receives the last write I/O becomes the active read/write volume, and the other volume becomes Block I/O Mode (no read/write access to the volume).
  - This might cause an unintended switchover between the preferred and nonpreferred paths.

- If I/O preference mode is set for a global-active device pair, the P-Vol will always be the active read/write volume and the S-Vol will be suspended with no read/write access when replication links become unavailable.
  - This prevents an unintended switchover between the preferred and nonpreferred paths.
  - I/O preference mode can be set using the CCI (Command Control Interface).
- See the Global-active Device User Guide – “I/O preference mode when remote path failed in a global- active device pair” section for more details.
- After replication link issues are resolved, resume replication using the Command Control Interface.

8. Primary site down (Site 1 failure).

Primary site failure is simulated by powering down ESXi hosts and the VSP storage at the same time in Site 1.

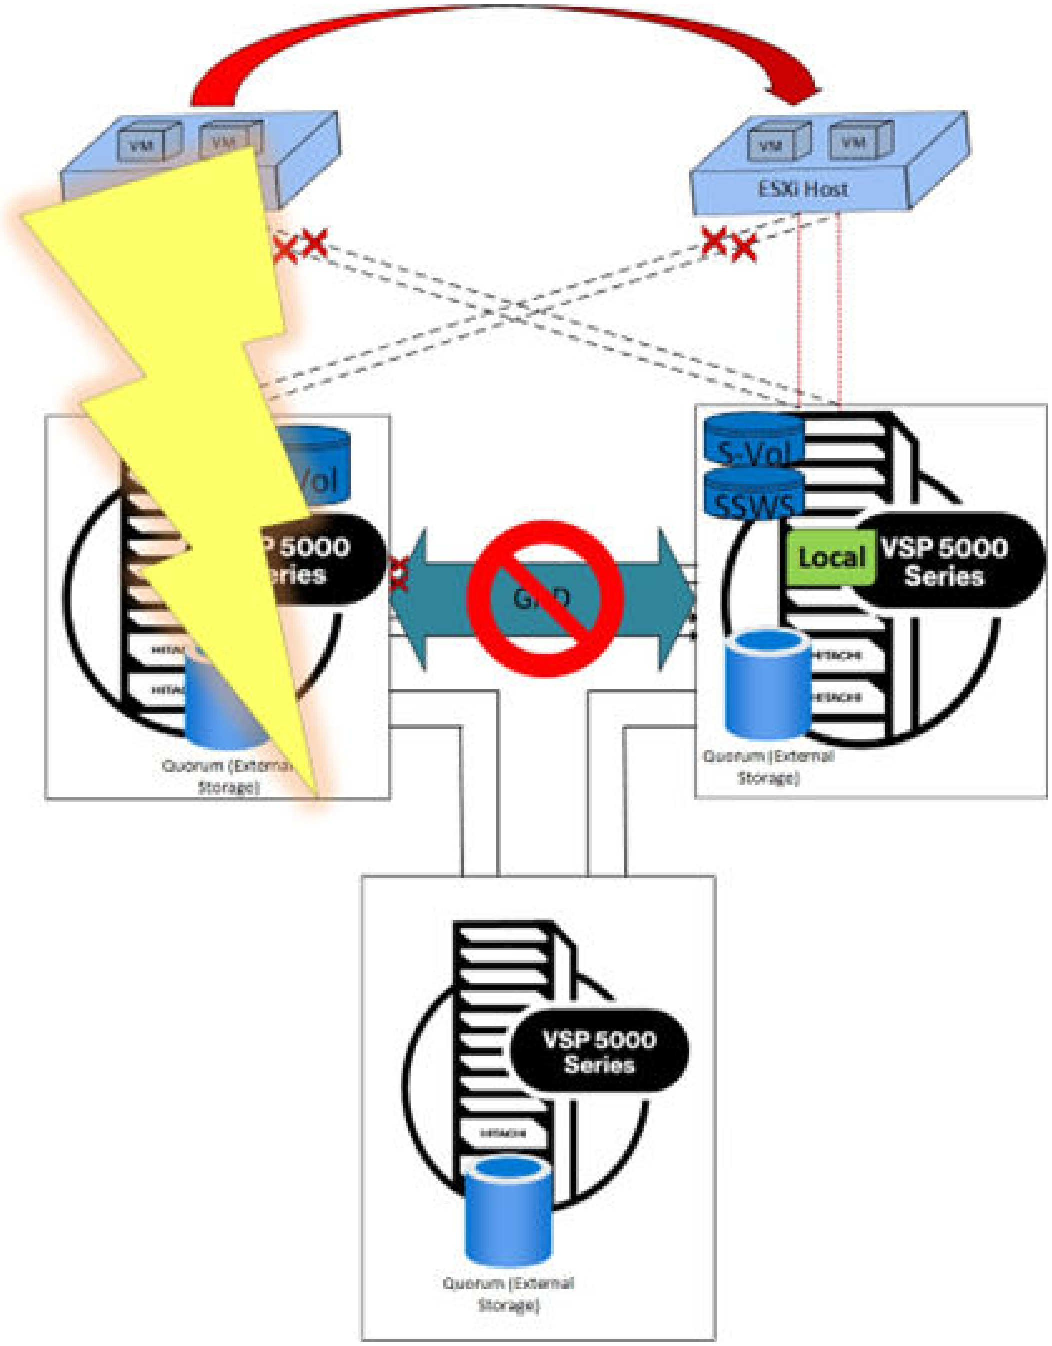

Observations:

- vSphere HA detects Hosts are down and triggers failover to another Hosts.
- VMs restart on Hosts in the Secondary Site.
- To recover see “The following are the steps to recover from Storage Failover (Failback)."

VMware live site recovery test case scenarios (between Region 1 and Region 2)

1. Recovery plan testing (Run recovery plan using the TEST option).

(The TEST option is necessary to verify whether the setup is functioning as expected without making any actual changes.)

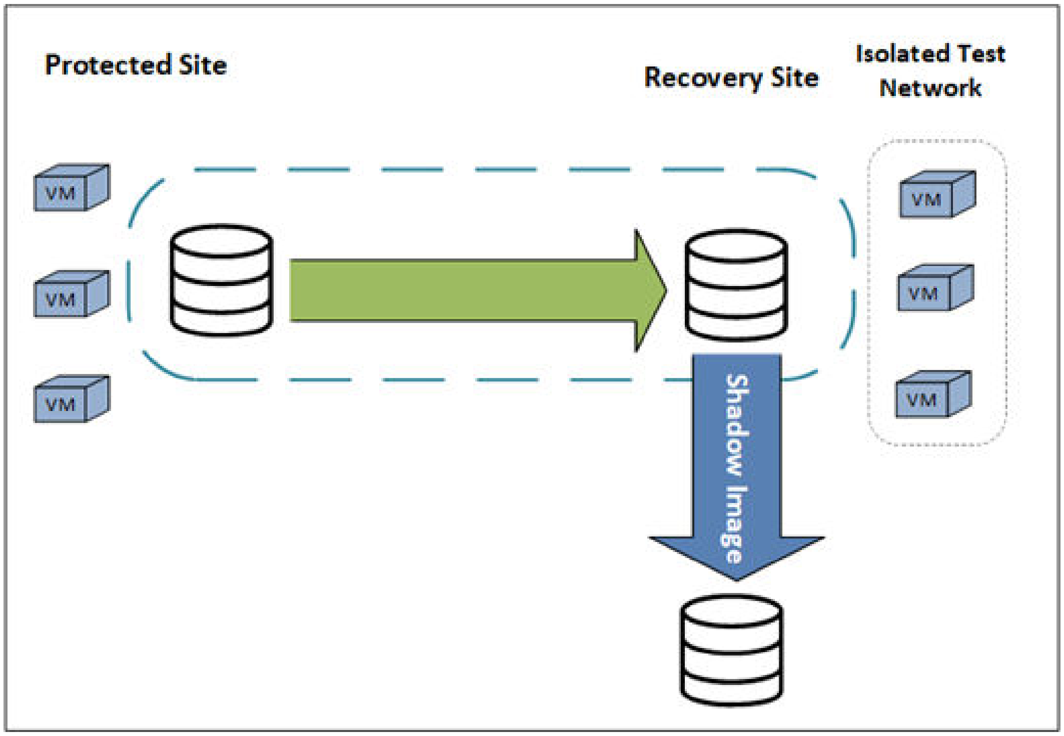

The Recovery plan test can be performed by starting the VM at Site 3 (Region 2) without stopping the VM workload that is running at the Protected site (Site 1). Shadow image is used to start VM at recovery site (Site 3 in Region 2).

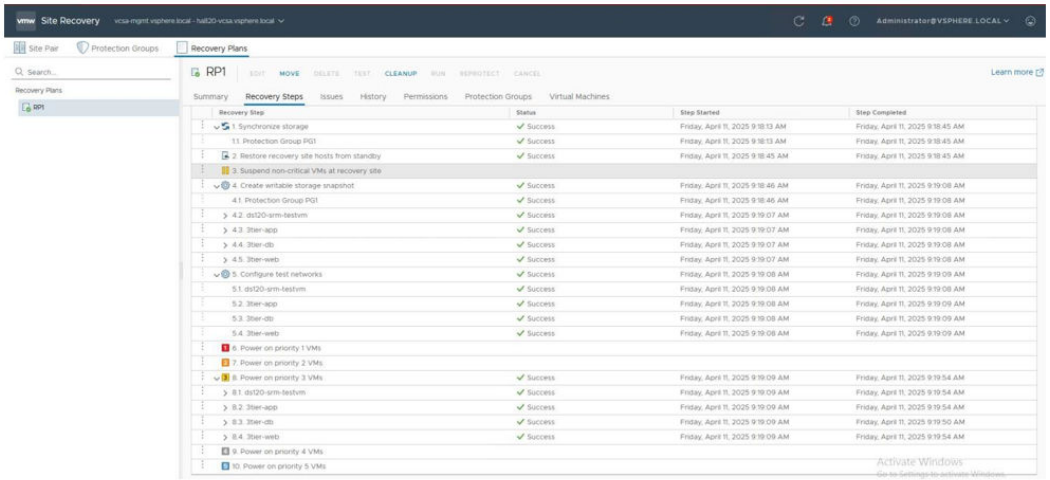

Recovery Step	Status	Step Started	Step Completed
1. Synchronize storage	✓ Success	Friday, April 11, 2025 9:18:13 AM	Friday, April 11, 2025 9:18:45 AM
1.1. Protection Group PG1	✓ Success	Friday, April 11, 2025 9:18:13 AM	Friday, April 11, 2025 9:18:45 AM
2. Restore recovery site hosts from standby	✓ Success	Friday, April 11, 2025 9:18:45 AM	Friday, April 11, 2025 9:18:45 AM
3. Suspend non-critical VMs at recovery site			
4. Create writable storage snapshot	✓ Success	Friday, April 11, 2025 9:18:46 AM	Friday, April 11, 2025 9:19:08 AM
4.1. Protection Group PG1	✓ Success	Friday, April 11, 2025 9:18:46 AM	Friday, April 11, 2025 9:19:08 AM
4.2. ds120-srm-testvm	✓ Success	Friday, April 11, 2025 9:19:07 AM	Friday, April 11, 2025 9:19:08 AM
4.3. 3tier-app	✓ Success	Friday, April 11, 2025 9:19:07 AM	Friday, April 11, 2025 9:19:08 AM
4.4. 3tier-db	✓ Success	Friday, April 11, 2025 9:19:07 AM	Friday, April 11, 2025 9:19:08 AM
4.5. 3tier-web	✓ Success	Friday, April 11, 2025 9:19:07 AM	Friday, April 11, 2025 9:19:08 AM
5. Configure test networks	✓ Success	Friday, April 11, 2025 9:19:08 AM	Friday, April 11, 2025 9:19:09 AM
5.1. ds120-srm-testvm	✓ Success	Friday, April 11, 2025 9:19:08 AM	Friday, April 11, 2025 9:19:08 AM
5.2. 3tier-app	✓ Success	Friday, April 11, 2025 9:19:08 AM	Friday, April 11, 2025 9:19:09 AM
5.3. 3tier-db	✓ Success	Friday, April 11, 2025 9:19:08 AM	Friday, April 11, 2025 9:19:09 AM
5.4. 3tier-web	✓ Success	Friday, April 11, 2025 9:19:08 AM	Friday, April 11, 2025 9:19:09 AM
6. Power on priority 1 VMs			
7. Power on priority 2 VMs			
8. Power on priority 3 VMs	✓ Success	Friday, April 11, 2025 9:19:09 AM	Friday, April 11, 2025 9:19:54 AM
8.1. ds120-srm-testvm	✓ Success	Friday, April 11, 2025 9:19:09 AM	Friday, April 11, 2025 9:19:54 AM
8.2. 3tier-app	✓ Success	Friday, April 11, 2025 9:19:09 AM	Friday, April 11, 2025 9:19:54 AM
8.3. 3tier-db	✓ Success	Friday, April 11, 2025 9:19:09 AM	Friday, April 11, 2025 9:19:50 AM
8.4. 3tier-web	✓ Success	Friday, April 11, 2025 9:19:09 AM	Friday, April 11, 2025 9:19:54 AM
9. Power on priority 4 VMs			
10. Power on priority 5 VMs			

- As shown, after running TEST in recovery plan, it did complete the operations.

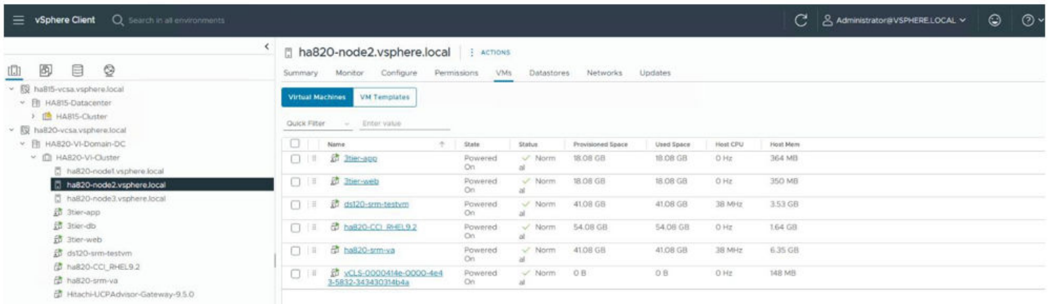

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
3tier-app	Powered On	✓ Normal	18.08 GB	18.08 GB	0 Hz	364 MB
3tier-web	Powered On	✓ Normal	18.08 GB	18.08 GB	0 Hz	350 MB
ds120-srm-testvm	Powered On	✓ Normal	41.08 GB	41.08 GB	38 MHz	3.53 GB
ha820-CCI_RHEL9.2	Powered On	✓ Normal	54.08 GB	54.08 GB	0 Hz	1.64 GB
ha820-srm-va	Powered On	✓ Normal	41.08 GB	41.08 GB	38 MHz	6.35 GB
vCLS-0000414e-0000-4e43-5832-343430314b4a	Powered On	✓ Normal	0 B	0 B	0 Hz	148 MB

- The previous image shows the details of the VMs running in the cluster.

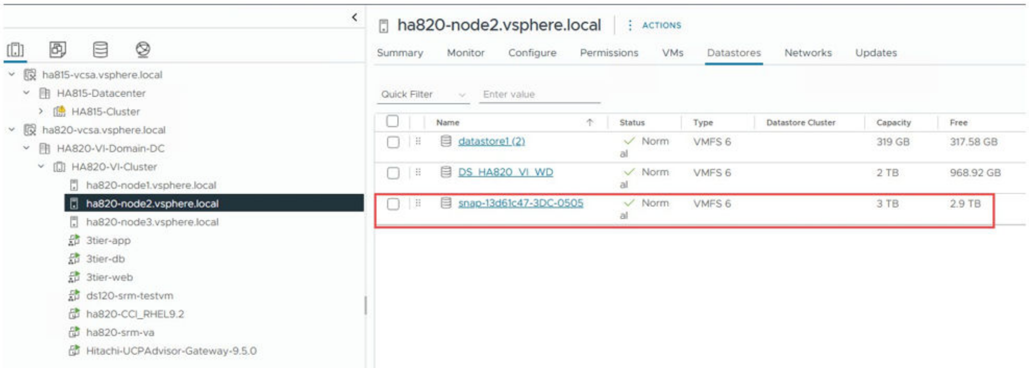

Name	Status	Type	Datastore Cluster	Capacity	Free
datastore1 (2)	✓ Normal	VMFS 6		319 GB	317.58 GB
DS_HA820_VI_WD	✓ Normal	VMFS 6		2 TB	968.92 GB
snap-13d61c47-3DC-0505	✓ Normal	VMFS 6		3 TB	2.9 TB

- As shown in the previous image, when we run the TEST plan it mounts the shadow image (snap datastore) on the recovery site and runs the VM workload on that datastore until we run the CLEANUP task.

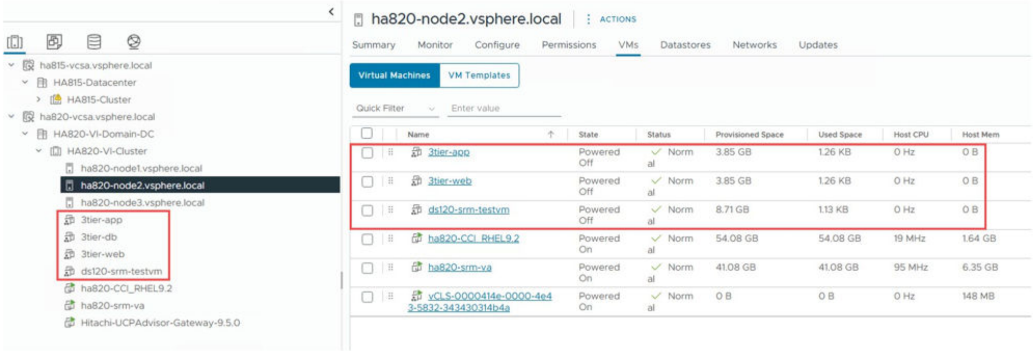

- ha815-vcsa.vsphere.local
  - HA815-Datacenter
    - HA815-Cluster
- ha820-vcsa.vsphere.local
  - HA820-VI-Domain-DC
    - HA820-VI-Cluster
      - ha820-node1.vsphere.local
      - ha820-node2.vsphere.local
      - ha820-node3.vsphere.local
      - 3tier-app
      - 3tier-db
      - 3tier-web
      - ds120-srm-testvm
      - ha820-CCI_RHEL9.2
      - ha820-srm-va
      - Hitachi-UCPAdvisor-Gateway-9.5.0

ha820-node2.vsphere.local | ACTIONS

Summary | Monitor | Configure | Permissions | VMs | Datastores | Networks | Updates

Virtual Machines | VM Templates

Quick Filter | Enter value

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
3tier-app	Powered Off	✓ Normal	3.85 GB	1.26 KB	0 Hz	0 B
3tier-web	Powered Off	✓ Normal	3.85 GB	1.26 KB	0 Hz	0 B
ds120-srm-testvm	Powered Off	✓ Normal	8.71 GB	1.13 KB	0 Hz	0 B
ha820-CCI_RHEL9.2	Powered On	✓ Normal	54.08 GB	54.08 GB	19 MHz	1.64 GB
ha820-srm-va	Powered On	✓ Normal	41.08 GB	41.08 GB	95 MHz	6.35 GB
vCLS-0000414e-0000-4e43-5832-343430314b4a	Powered On	✓ Normal	0 B	0 B	0 Hz	148 MB

- The previous image shows the status of the VMs on the protected site which were part of the TEST activity.

Observations:

- The Recovery plan test can be performed by starting VM at Site 3 without stopping the system that is running at the Protected site, and a shadow image is used to start the VM at the protected site.
- After running a test, VLSR removes temporary snapshots or logs and persists changes to the virtual machine disks.

2. Test Failover from Protected site to Recovery site.

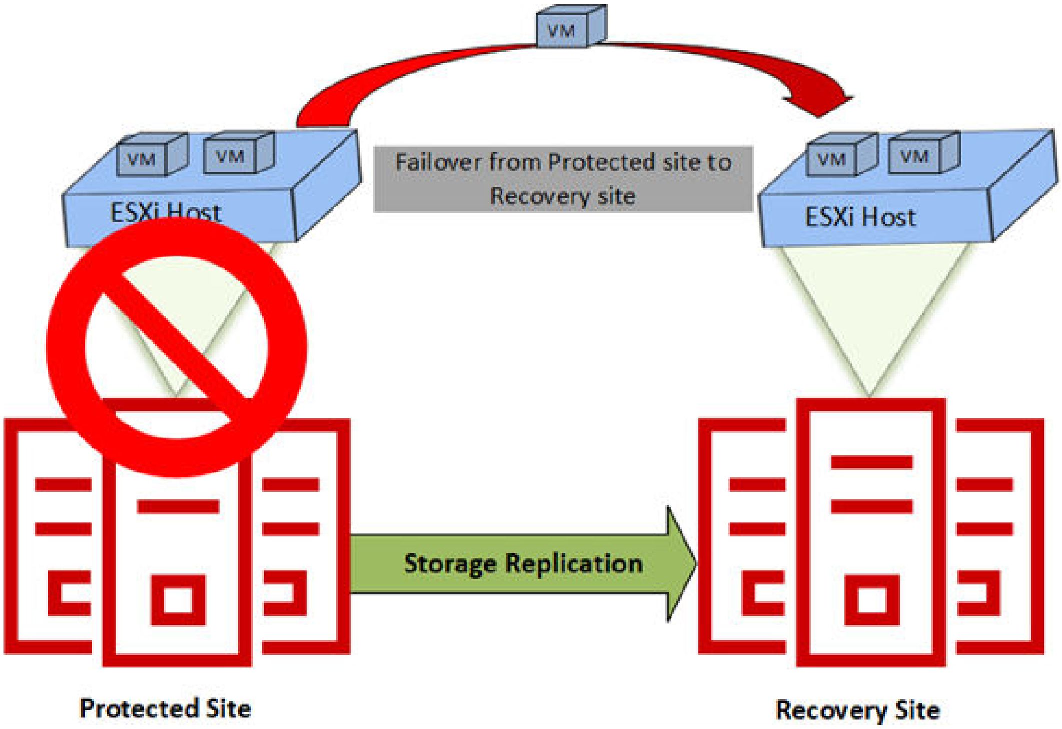

When failover was triggered during failover, VMs will be powered off on the protected site and migrated to the recovery site.

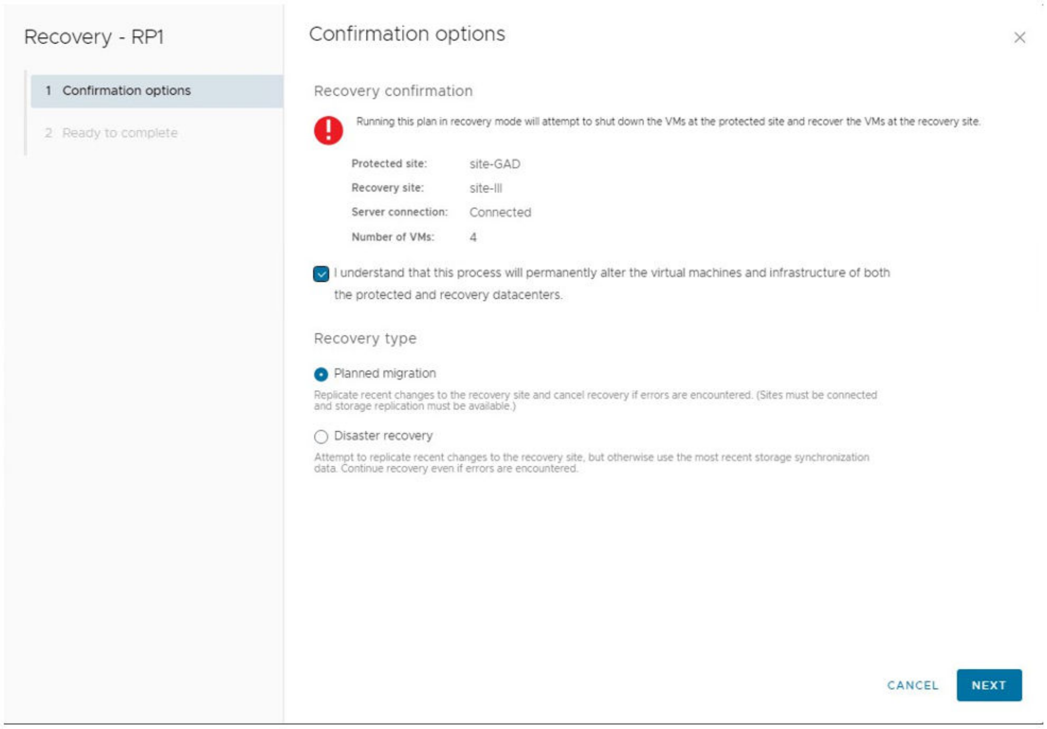

Recovery - RP1

1 Confirmation options

2 Ready to complete

Confirmation options

Recovery confirmation

Running this plan in recovery mode will attempt to shut down the VMs at the protected site and recover the VMs at the recovery site.

Protected site:	site-GAD
Recovery site:	site-III
Server connection:	Connected
Number of VMs:	4

☑ I understand that this process will permanently alter the virtual machines and infrastructure of both the protected and recovery datacenters.

Recovery type

● Planned migration

Replicate recent changes to the recovery site and cancel recovery if errors are encountered. (Sites must be connected and storage replication must be available.)

○ Disaster recovery

Attempt to replicate recent changes to the recovery site, but otherwise use the most recent storage synchronization data. Continue recovery even if errors are encountered.

- When you run an actual plan, you are prompted for the previous options to proceed.

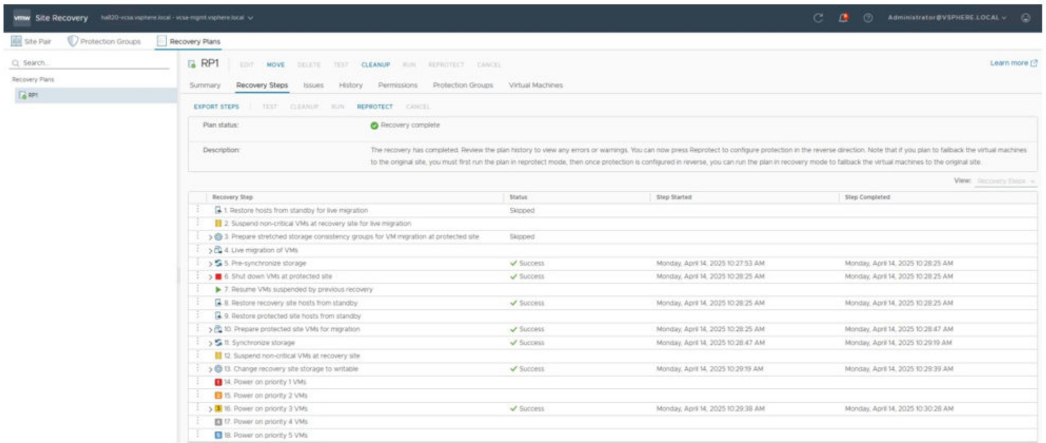

Plan status: Recovery complete

Description: The recovery has completed. Review the plan history to view any errors or warnings. You can now press Reprotect to configure protection in the reverse direction. Note that if you plan to fallback the virtual machines to the original site, you must first run the plan in reprotect mode, then once protection is configured in reverse, you can run the plan in recovery mode to fallback the virtual machines to the original site.

View: Recovery Steps

Recovery Step	Status	Step Started	Step Completed
1. Restore hosts from standby for live migration	Skipped		
2. Suspend non-critical VMs at recovery site for live migration			
3. Prepare stretched storage consistency groups for VM migration at protected site	Skipped		
4. Live migration of VMs			
5. Pre-synchronize storage	✓ Success	Monday, April 14, 2025 10:27:53 AM	Monday, April 14, 2025 10:28:25 AM
6. Shut down VMs at protected site	✓ Success	Monday, April 14, 2025 10:28:25 AM	Monday, April 14, 2025 10:28:25 AM
7. Resume VMs suspended by previous recovery			
8. Restore recovery site hosts from standby	✓ Success	Monday, April 14, 2025 10:28:25 AM	Monday, April 14, 2025 10:28:25 AM
9. Restore protected site hosts from standby			
10. Prepare protected site VMs for migration	✓ Success	Monday, April 14, 2025 10:28:25 AM	Monday, April 14, 2025 10:28:47 AM
11. Synchronize storage	✓ Success	Monday, April 14, 2025 10:28:47 AM	Monday, April 14, 2025 10:29:19 AM
12. Suspend non-critical VMs at recovery site			
13. Change recovery site storage to writable	✓ Success	Monday, April 14, 2025 10:29:19 AM	Monday, April 14, 2025 10:29:39 AM
14. Power on priority 1 VMs			
15. Power on priority 2 VMs			
16. Power on priority 3 VMs	✓ Success	Monday, April 14, 2025 10:29:38 AM	Monday, April 14, 2025 10:30:28 AM
17. Power on priority 4 VMs			
18. Power on priority 5 VMs			

- The previous image shows the detailed steps followed during the failover process when RUN tasks are performed.
- The following is the status from the CCI server for the pair associated with it.

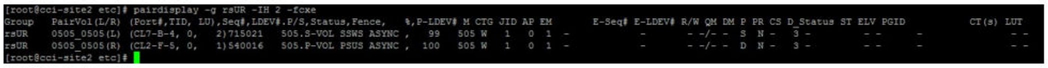

```
[root@cci-site2 etc]# pairdisplay -g rsUR -IH 2 -fcxe
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W QM DM P PR CS D_Status ST ELV PGID        CT(s) LUT
rsUR    0505_0505(L) (CL7-B-4, 0,    2)715021   505.S-VOL SSWS ASYNC ,   99    505 W   1   0  1 -          -       - -/- -  S  N -  3 -        - -      -             - -
rsUR    0505_0505(R) (CL2-F-5, 0,    1)540016   505.P-VOL PSUS ASYNC ,  100    505 W   1   0  1 -          -       - -/- -  D  N -  3 -        - -      -             - -
[root@cci-site2 etc]#
```

Observation:

- When failures occur at the protected site, the VM workload can be migrated to the recovery site by performing failover.

3. Failback from Recovery site to Protected site.

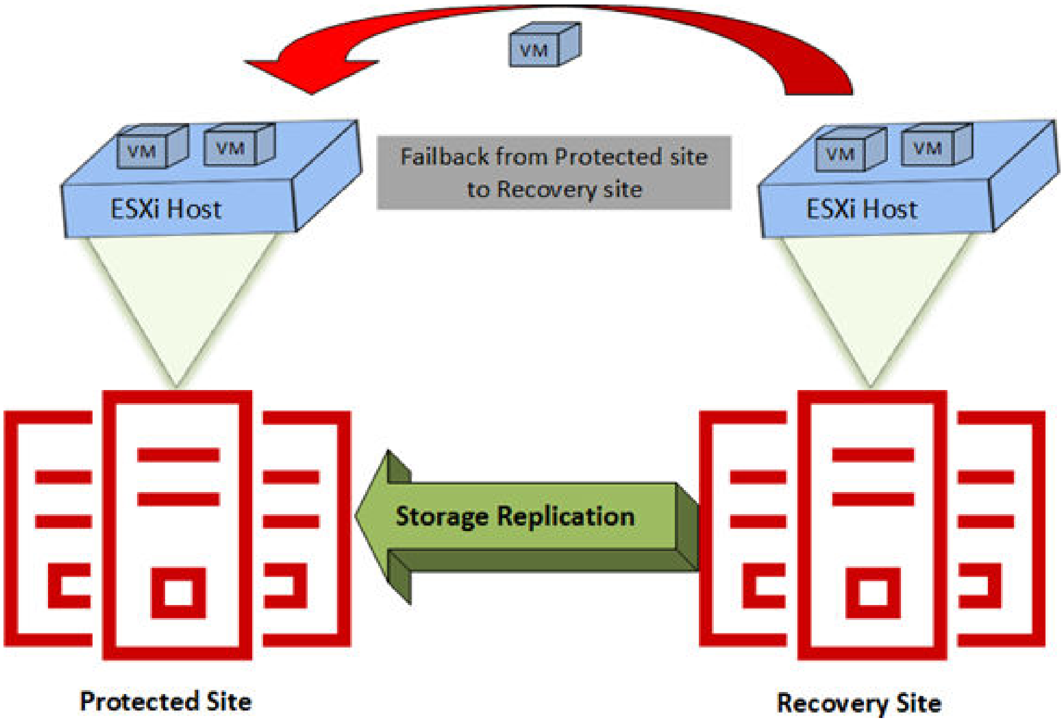

Observations:

- With global-active device/Universal Replicator/VMware Live Site Recovery, when failback (with disaster recovery) is performed, the system that was migrated to Site 3 can be restored to Site 1.
- Failback can be performed only when all sites are running without a failure and all the requirements for global-active device/Universal Replicator/VMware Live Site Recovery are met. See the VMware® Live Site Recovery Manager™ Deployment Guide for details.

Solution deployment

Solution demonstration involves implementing a 3DC solution with the Hitachi Integrated Systems platform and VSP to enhance disaster recovery across multiple sites; deploying VCF 5.2 using the VCF Import Tool to integrate existing vSphere environments into VMware Cloud Foundation for centralized management; setting up a vSphere Metro Storage Cluster with Hitachi VSP global-active device to ensure high availability and seamless VM migration across metro distances (two availability zones within the same region); and using VMware Live Site Recovery Manager with the 3DC solution to automate disaster recovery processes (across two regions), ensuring robust replication and failover capabilities across the three data centers.

The following sections cover the deployment of VMware Cloud Foundation Management Workload Domain using VCF Import tool and the other storage options with VI Workload Domain:

- VCF 5.2.1 Deployment using the VCF Import Tool (on page 37)
- vSphere Metro Storage Cluster Using VSP global-active device (on page 39)
- VMware Live Site Recovery with 3DC solution (on page 40)

VCF 5.2.1 Deployment using the VCF Import Tool

VMware Cloud Foundation (VCF) 5.2 introduces the VCF Import Tool, a command-line interface (CLI) designed to streamline the transition to a private cloud. This tool extends SDDC Manager's fleet management capabilities to existing vSphere environments, enabling seamless integration without impacting running workloads. This process outlines the steps to deploy VCF 5.2 using the VCF Import Tool, ensuring a smooth transition from an existing vSphere environment to a VCF-managed private cloud. The tool simplifies the deployment, configuration, and management of the cloud infrastructure, making it more efficient and automated.

Before you begin

1. Download and Extract VCF Import Tool:
   - Obtain the VCF Import Tool from the Broadcom software portal.
   - Extract the tool on the SDDC Manager appliance using the `tar -xvf vcf-brownfield-import-5.2.0.0-24108578.tar.gz` command.
2. Pre-checks:
   - Run pre-checks on your existing vSphere environment to ensure compatibility with VCF. This includes verifying vCenter inventories and configurations.

Procedure

1. Scan vCenter Inventory:
   Use the check parameter to scan the vCenter inventory for compatibility issues:

```
python vcf_brownfield.py check --vcenter <vcenter-ip> --sso-user <sso-user> --
sso- password <sso-password>
```

2. Register vCenter Server:
   Register the vCenter Server and its associated inventory with the SDDC Manager:

```
python vcf_brownfield.py import --vcenter <vcenter-ip> --sso-user <sso-user> --
sso- password <sso-password> --domain-name <domain-name>
```

3. Convert/Import:

Use the convert or import parameters to transition the existing vSphere environment into a VCF workload domain:

```
python vcf_brownfield.py convert --vcenter <vcenter-ip> --sso-user <sso-user> --
sso- password <sso-password> --domain-name <domain-name>
```

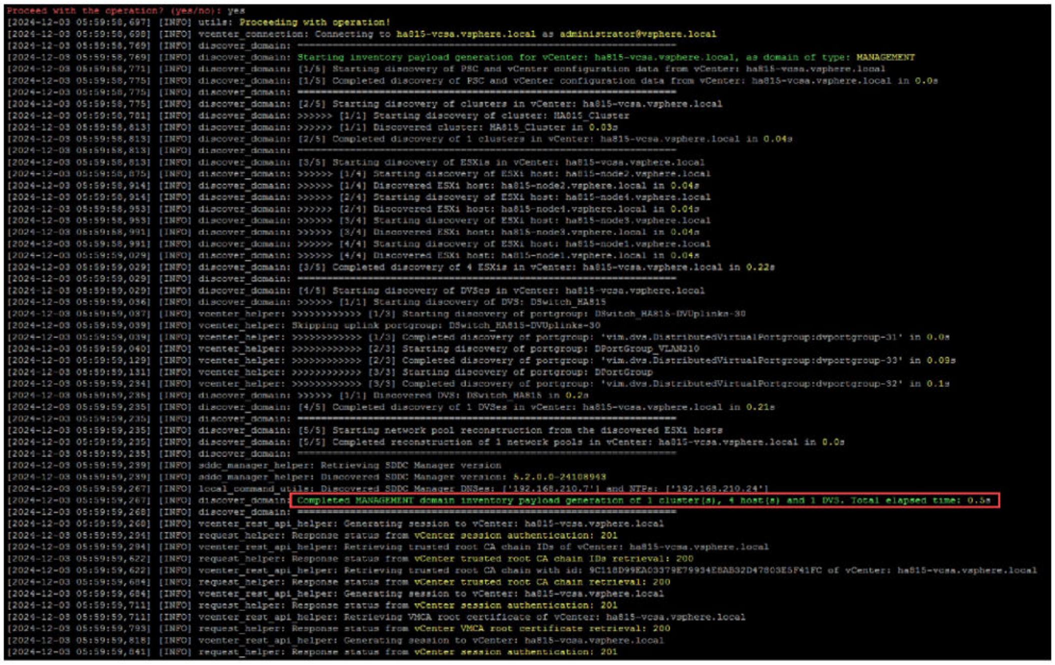

```
Proceed with the operation? (yes/no): yes
[2024-12-03 05:59:58,697] [INFO] utils: Proceeding with operation!
[2024-12-03 05:59:58,698] [INFO] vcenter_connection: Connecting to ha815-vcsa.vsphere.local as administrator@vsphere.local
[2024-12-03 05:59:58,769] [INFO] discover_domain: ------------------------------------------------------------------
[2024-12-03 05:59:58,769] [INFO] discover_domain: Starting inventory payload generation for vCenter: ha815-vcsa.vsphere.local, as domain of type: MANAGEMENT
[2024-12-03 05:59:58,771] [INFO] discover_domain: [1/5] Starting discovery of PSC and vCenter configuration data from vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:58,775] [INFO] discover_domain: [1/5] Completed discovery of PSC and vCenter configuration data from vCenter: ha815-vcsa.vsphere.local in 0.0s
[2024-12-03 05:59:58,775] [INFO] discover_domain: ==================================================================
[2024-12-03 05:59:58,775] [INFO] discover_domain: [2/5] Starting discovery of clusters in vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:58,781] [INFO] discover_domain: >>>>>> [1/1] Starting discovery of cluster: HA815_Cluster
[2024-12-03 05:59:58,813] [INFO] discover_domain: >>>>>> [1/1] Discovered cluster: HA815_Cluster in 0.03s
[2024-12-03 05:59:58,813] [INFO] discover_domain: [2/5] Completed discovery of 1 clusters in vCenter: ha815-vcsa.vsphere.local in 0.04s
[2024-12-03 05:59:58,813] [INFO] discover_domain: ------------------------------------------------------------------
[2024-12-03 05:59:58,813] [INFO] discover_domain: [3/5] Starting discovery of ESXis in vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:58,875] [INFO] discover_domain: >>>>>> [1/4] Starting discovery of ESXi host: ha815-node2.vsphere.local
[2024-12-03 05:59:58,914] [INFO] discover_domain: >>>>>> [1/4] Discovered ESXi host: ha815-node2.vsphere.local in 0.04s
[2024-12-03 05:59:58,914] [INFO] discover_domain: >>>>>> [2/4] Starting discovery of ESXi host: ha815-node4.vsphere.local
[2024-12-03 05:59:58,953] [INFO] discover_domain: >>>>>> [2/4] Discovered ESXi host: ha815-node4.vsphere.local in 0.04s
[2024-12-03 05:59:58,953] [INFO] discover_domain: >>>>>> [3/4] Starting discovery of ESXi host: ha815-node3.vsphere.local
[2024-12-03 05:59:58,991] [INFO] discover_domain: >>>>>> [3/4] Discovered ESXi host: ha815-node3.vsphere.local in 0.04s
[2024-12-03 05:59:58,991] [INFO] discover_domain: >>>>>> [4/4] Starting discovery of ESXi host: ha815-node1.vsphere.local
[2024-12-03 05:59:59,029] [INFO] discover_domain: >>>>>> [4/4] Discovered ESXi host: ha815-node1.vsphere.local in 0.04s
[2024-12-03 05:59:59,029] [INFO] discover_domain: [3/5] Completed discovery of 4 ESXis in vCenter: ha815-vcsa.vsphere.local in 0.22s
[2024-12-03 05:59:59,029] [INFO] discover_domain: ==================================================================
[2024-12-03 05:59:59,029] [INFO] discover_domain: [4/5] Starting discovery of DVSes in vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:59,036] [INFO] discover_domain: >>>>>> [1/1] Starting discovery of DVS: DSwitch_HA815
[2024-12-03 05:59:59,037] [INFO] vcenter_helper: >>>>>>>>>>>>> [1/3] Starting discovery of portgroup: DSwitch_HA815-DVUplinks-30
[2024-12-03 05:59:59,039] [INFO] vcenter_helper: Skipping uplink portgroup: DSwitch_HA815-DVUplinks-30
[2024-12-03 05:59:59,039] [INFO] vcenter_helper: >>>>>>>>>>>>> [1/3] Completed discovery of portgroup: 'vim.dvs.DistributedVirtualPortgroup:dvportgroup-31' in 0.0s
[2024-12-03 05:59:59,040] [INFO] vcenter_helper: >>>>>>>>>>>>> [2/3] Starting discovery of portgroup: DPortGroup_VLAN210
[2024-12-03 05:59:59,129] [INFO] vcenter_helper: >>>>>>>>>>>>> [2/3] Completed discovery of portgroup: 'vim.dvs.DistributedVirtualPortgroup:dvportgroup-33' in 0.09s
[2024-12-03 05:59:59,131] [INFO] vcenter_helper: >>>>>>>>>>>>> [3/3] Starting discovery of portgroup: DPortGroup
[2024-12-03 05:59:59,234] [INFO] vcenter_helper: >>>>>>>>>>>>> [3/3] Completed discovery of portgroup: 'vim.dvs.DistributedVirtualPortgroup:dvportgroup-32' in 0.1s
[2024-12-03 05:59:59,235] [INFO] discover_domain: >>>>>> [1/1] Discovered DVS: DSwitch_HA815 in 0.2s
[2024-12-03 05:59:59,235] [INFO] discover_domain: [4/5] Completed discovery of 1 DVSes in vCenter: ha815-vcsa.vsphere.local in 0.21s
[2024-12-03 05:59:59,235] [INFO] discover_domain: ==================================================================
[2024-12-03 05:59:59,235] [INFO] discover_domain: [5/5] Starting network pool reconstruction from the discovered ESXi hosts
[2024-12-03 05:59:59,235] [INFO] discover_domain: [5/5] Completed reconstruction of 1 network pools in vCenter: ha815-vcsa.vsphere.local in 0.0s
[2024-12-03 05:59:59,235] [INFO] discover_domain: ------------------------------------------------------------------
[2024-12-03 05:59:59,239] [INFO] sddc_manager_helper: Retrieving SDDC Manager version
[2024-12-03 05:59:59,239] [INFO] sddc_manager_helper: Discovered SDDC Manager version: 5.2.0.0-24108943
[2024-12-03 05:59:59,267] [INFO] local_command_utils: Discovered SDDC Manager DNSes: ['192.168.210.7'] and NTPs: ['192.168.210.24']
[2024-12-03 05:59:59,267] [INFO] discover_domain: Completed MANAGEMENT domain inventory payload generation of 1 cluster(s), 4 host(s) and 1 DVS. Total elapsed time: 0.5s
[2024-12-03 05:59:59,268] [INFO] discover_domain: ==================================================================
[2024-12-03 05:59:59,268] [INFO] vcenter_rest_api_helper: Generating session to vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:59,294] [INFO] request_helper: Response status from vCenter session authentication: 201
[2024-12-03 05:59:59,294] [INFO] vcenter_rest_api_helper: Retrieving trusted root CA chain IDs of vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:59,622] [INFO] request_helper: Response status from vCenter trusted root CA chain IDs retrieval: 200
[2024-12-03 05:59:59,622] [INFO] vcenter_rest_api_helper: Retrieving trusted root CA chain with id: 9C118D99EA03379E79934E8AB32D47803E5F41FC of vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:59,684] [INFO] request_helper: Response status from vCenter trusted root CA chain retrieval: 200
[2024-12-03 05:59:59,684] [INFO] vcenter_rest_api_helper: Generating session to vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:59,711] [INFO] request_helper: Response status from vCenter session authentication: 201
[2024-12-03 05:59:59,711] [INFO] vcenter_rest_api_helper: Retrieving VMCA root certificate of vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:59,793] [INFO] request_helper: Response status from vCenter VMCA root certificate retrieval: 200
[2024-12-03 05:59:59,818] [INFO] vcenter_rest_api_helper: Generating session to vCenter: ha815-vcsa.vsphere.local
[2024-12-03 05:59:59,841] [INFO] request_helper: Response status from vCenter session authentication: 201
```

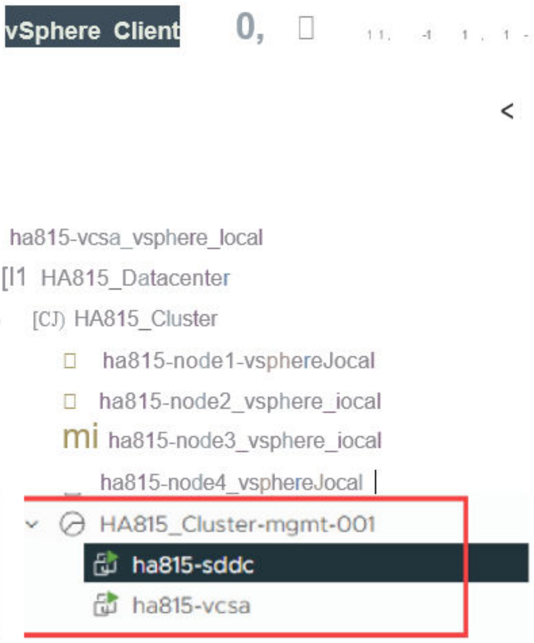

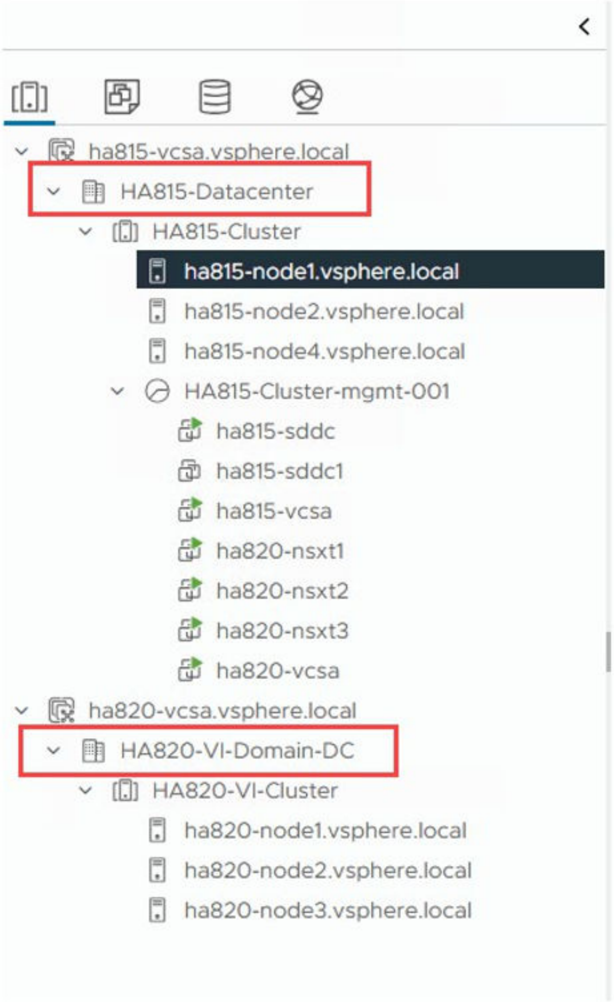

- ha815-vcsa.vsphere.local
  - HA815-Datacenter
    - HA815-Cluster
      - ha815-node1.vsphere.local
      - ha815-node2.vsphere.local
      - ha815-node4.vsphere.local
      - HA815-Cluster-mgmt-001
        - ha815-sddc
        - ha815-sddc1
        - ha815-vcsa
        - ha820-nsxt1
        - ha820-nsxt2
        - ha820-nsxt3
        - ha820-vcsa
- ha820-vcsa.vsphere.local
  - HA820-VI-Domain-DC
    - HA820-VI-Cluster
      - ha820-node1.vsphere.local
      - ha820-node2.vsphere.local
      - ha820-node3.vsphere.local

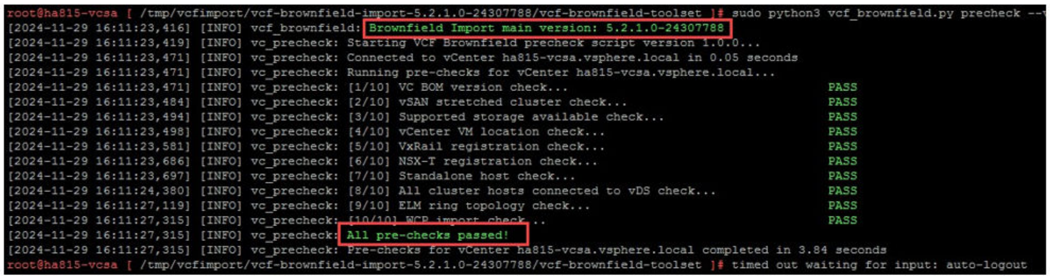

```
root@ha815-vcsa [ /tmp/vcfimport/vcf-brownfield-import-5.2.1.0-24307788/vcf-brownfield-toolset ]# sudo python3 vcf_brownfield.py precheck --v
[2024-11-29 16:11:23,416] [INFO] vcf_brownfield: Brownfield Import main version: 5.2.1.0-24307788
[2024-11-29 16:11:23,419] [INFO] vc_precheck: Starting VCF Brownfield precheck script version 1.0.0...
[2024-11-29 16:11:23,471] [INFO] vc_precheck: Connected to vCenter ha815-vcsa.vsphere.local in 0.05 seconds
[2024-11-29 16:11:23,471] [INFO] vc_precheck: Running pre-checks for vCenter ha815-vcsa.vsphere.local...
[2024-11-29 16:11:23,471] [INFO] vc_precheck: [1/10] VC BOM version check...                          PASS
[2024-11-29 16:11:23,484] [INFO] vc_precheck: [2/10] vSAN stretched cluster check...                  PASS
[2024-11-29 16:11:23,494] [INFO] vc_precheck: [3/10] Supported storage available check...            PASS
[2024-11-29 16:11:23,498] [INFO] vc_precheck: [4/10] vCenter VM location check...                     PASS
[2024-11-29 16:11:23,581] [INFO] vc_precheck: [5/10] VxRail registration check...                     PASS
[2024-11-29 16:11:23,686] [INFO] vc_precheck: [6/10] NSX-T registration check...                      PASS
[2024-11-29 16:11:23,697] [INFO] vc_precheck: [7/10] Standalone host check...                         PASS
[2024-11-29 16:11:24,380] [INFO] vc_precheck: [8/10] All cluster hosts connected to vDS check...      PASS
[2024-11-29 16:11:27,119] [INFO] vc_precheck: [9/10] ELM ring topology check...                       PASS
[2024-11-29 16:11:27,315] [INFO] vc_precheck: [10/10] WCP import check...                             PASS
[2024-11-29 16:11:27,315] [INFO] vc_precheck: All pre-checks passed!
[2024-11-29 16:11:27,315] [INFO] vc_precheck: Pre-checks for vCenter ha815-vcsa.vsphere.local completed in 3.84 seconds
root@ha815-vcsa [ /tmp/vcfimport/vcf-brownfield-import-5.2.1.0-24307788/vcf-brownfield-toolset ]# timed out waiting for input: auto-logout
```

See Importing the existing vSphere environment for more information

.

vSphere Metro Storage Cluster Using VSP global-active device

Integrating vSphere Metro Storage Cluster with VMware Cloud Foundation and Hitachi Virtual Storage Platform global-active device provides a robust solution for high availability, disaster recovery, and automated management across multiple data centers. This combination ensures your cloud infrastructure is resilient, scalable, and easy to manage.

vMSC enables active-active storage replication across geographically separated data centers, allowing virtual machines (VMs) to remain operational even if one data center fails. VCF leverages this capability to ensure continuous availability and disaster recovery for workloads. VSP global-active device enhances this by providing synchronous replication, ensuring data is mirrored in real-time across data centers, maintaining data consistency and integrity.

VCF integrates with vMSC and VSP global-active device to manage and orchestrate these replication processes, ensuring data is always up-to-date and available. VCF also automates lifecycle management for compute, storage, and network resources. When integrated with vMSC and VSP global-active device, VCF simplifies the deployment, configuration, and management of stretched storage clusters.

vMSC works with vSphere High Availability (HA) and Distributed Resource Scheduler (DRS) to provide automated failover and load balancing across data centers. VCF ensures workloads are automatically redistributed in case of failure, maintaining optimal performance and availability.

vMSC supports both uniform and non-uniform configurations. In a uniform setup, all hosts access the same storage devices across data centers. In a non-uniform setup, hosts access storage devices only within their local data center. VCF can be configured to support either setup based on deployment requirements.

VMware Live Site Recovery with 3DC solution

About Hitachi Storage Replication Adapter and Live Site Recovery Manager

The VMware vCenter VMware Live Site Recovery software application automates the disaster recovery process using storage-based replication. Hitachi Storage Replication Adapter (SRA) is the interface that integrates Hitachi storage systems and replication software with VMware vCenter Live Site Recovery processes.

Used together, VMware Live Site Recovery and Hitachi storage and software provide an automated and seamless disaster recovery solution within the VMware vCenter infrastructure.

VMware vCenter infrastructure

The VMware Live Site Recovery/Hitachi SRA solution on the VMware side consists of the following:

- VMware vSphere, the virtualization platform with data center infrastructure. vSphere includes VMware ESX/ESXi host, a virtualization platform that provides a data center infrastructure in which many virtual machines share hardware resources from a single physical machine. The ESX/ESXi host loads directly on a physical server. vCenter Server, which provides management of one or multiple vSphere environments. These vSphere elements are used at the protected and recovery sites.
- VMware Live Site Recovery which provides a disaster recovery solution that reduces planned and unplanned downtime of the vSphere infrastructure. Hitachi storage and replication software products. The Hitachi Storage Replication Adapter links VMware Live Site Recovery and Hitachi storage and replication software. The SRA/ VMware Live Site Recovery solution supports the Hitachi Virtual Storage Platform 5000 series (VSP 5000 series).

Note: For the latest information about Hitachi storage systems supported by SRA, see the VMware Compatibility Guide on the Broadcom website.

Hitachi remote and in-system replication are key features of the solution. Remote replication is used to backup protected site data at the recovery site in a remote location. In-system replication is used on the remote site to create a clone volume for testing the VMware Live Site Recovery-SRA solution.

The following remote replication products are supported:

- Hitachi Universal Replicator, which provides long-distance asynchronous replication across any distance without significant impact on host performance.
- Global-active device, which provides synchronous remote replication. The following in-system replication products are supported for creating a clone of the recovery site volume for testing.
- ShadowImage® (SI), which creates RAID-protected duplicate volumes within the storage system. With ShadowImage, you create a clone of the remote backup volume in the remote storage system.

The following illustration shows VMware vCenter Live Site Recovery SRA and Hitachi components.

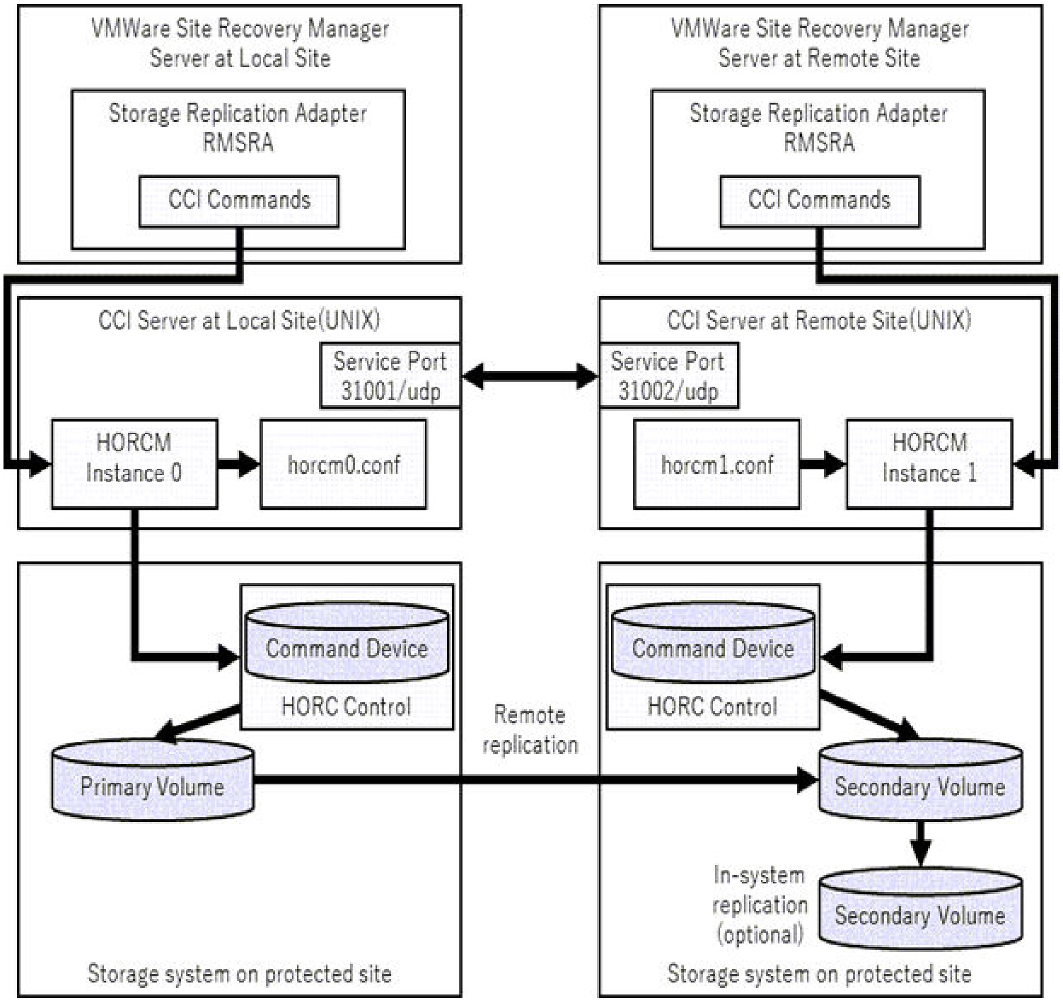
VMWare Site Recovery Manager Server at Local Site
Storage Replication Adapter RMSRA
CCI Commands
VMWare Site Recovery Manager Server at Remote Site
Storage Replication Adapter RMSRA
CCI Commands
CCI Server at Local Site(UNIX)
Service Port 31001/udp
HORCM Instance 0
horcm0.conf
CCI Server at Remote Site(UNIX)
Service Port 31002/udp
horcm1.conf
HORCM Instance 1
Command Device
HORC Control
Primary Volume
Remote replication
Command Device
HORC Control
Secondary Volume
In-system replication (optional)
Secondary Volume
Storage system on protected site
Storage system on protected site

Command Control Interface (CCI)

Hitachi remote and in-system replication software requires CCI to manage the pairs. The adapter plug-in links CCI with VMware Live Site Recovery.

There are two CCI components:

- CCI command devices, which reside on the storage systems. CCI uses the command device as the interface to the storage system from the host. The command device accepts CCI commands from the host and executes them on the storage system. The command device is a dedicated logical volume.

  > **Note:** The two methods for issuing CCI commands from a host are the inband method and the out-of-band method. In environments using VMware Live Site Recovery, the inband method is recommended because of performance considerations.

- Hitachi Open Remote Copy Manager (HORCM), which resides on the CCI server. HORCM operates as a daemon process. When activated, HORCM refers to the CCI configuration definition files, also located on the server. The HORCM instance communicates with the storage system and remote servers.

HORCM definition files describe the storage systems, pair volumes, and data paths. When a user issues a command, CCI uses the information in the HORCM files to identify which volumes are the targets of the command.

Two HORCM files are needed for each pair. One file describes the primary volumes (P-Vol s), which are also referred to as “protected volumes, and the other file describes the secondary volumes (S-Vol s), which are also referred to as “recovery volumes.

VMware Live Site Recovery and Hitachi components (as shown in the illustration) shows a two-server, two-HORCM-instance setup with optional in-system test copy.

How the VMware vCenter Live recovery /SRA solution works:

The VMware vCenter Live Site Recovery software coordinates processing with Hitachi storage and replication so that in a recovery condition, the virtual machines at the protected site are shut down and the replicated virtual machines are powered up.

Recovery is guided by a recovery plan that specifies the order in which the virtual machines are to be started up.

After a recovery is performed, the running virtual machines are no longer protected. The VMware vCenter Live recovery software provides a reprotect operation, which runs after the original protected site is back up. Reprotect activates CCI operations that reverse synchronize data in the storage systems from recovery site to protected site.

Finally, VMware vCenter Live recovery supports failback and reprotect operations allow you to restore protection back to the original configuration, with data flow from the protected site to the recovery site.

VMware vCenter Live recovery allows you to test recovery plans using an in-system copy of the replicated data without disrupting ongoing operations at either site.

ShadowImage

ShadowImage is a data protection and replication technology used in storage systems, particularly in Hitachi Virtual Storage Platforms. It enables local mirroring of data volumes within a storage system, creating full copies of data for purposes like backups, testing, or secondary host applications.

Key Features:

- Volume Pairs: ShadowImage works with pairs of volumes—a primary volume (P-Vol) and secondary volumes (S-Vol s). These pairs allow asynchronous updates, ensuring minimal disruption to production workloads.
- Cascaded Pairs: It supports advanced configurations, such as cascaded pairs, where multiple layers of secondary volumes are created.
- Copy Workflows: Includes initial copy and update copy workflows to replicate data efficiently.
- Non-Disruptive Operations: Host application input/output (I/O) on the production volume continues uninterrupted during replication.

This technology is widely used for data protection, disaster recovery, and ensuring high availability in enterprise environments.

Product descriptions

The products described in this section are part of the solution.

Hitachi Integrated Systems Platform

The Hitachi Integrated Systems platform is a high-performance, low-latency, integrated, converged solution using Hitachi Virtual Storage Platform One Block storage, Hitachi Advanced Server HA820 G3, as well as HA810 G3 with Sapphire Rapids Scalable Processors.

Hitachi Virtual Storage Platform One Block

The Hitachi Virtual Storage Platform One Block series simplifies system setup and management through the new VSP 360 management offering. Dynamic Drive Protection reduces RAID complexity, and always-on compression and deduplication enhance simplicity.

Dynamic Carbon Reduction optimizes energy usage by switching CPUs to ECO mode during low activity. Adaptive Data Reduction (ADR) is always on, enhancing efficiency and reducing the overall CO2 footprint.

Thin Image Advanced (TIA) integrates with major snapshot ecosystems, prioritizing security by defending against threats and ensuring data confidentiality. CyberArk Privileged Access Manager plugins enhance block storage system security by prioritizing data confidentiality, ensuring compliance, and actively defending against security threats.

Hitachi Virtual Storage Platform One Block includes the following 3 dedicated models:

- VSP One Block 24 – 256 GB Cache + SW Advanced Data Reduction (ADR) + 24 cores
- VSP One Block 26 – 768 GB Cache + 2x Compression Accelerator Module (CAM) + 24 cores
- VSP One Block 28 – 1 TB Cache + 4x CAM + 64 cores

All have the same drive count (72 NVMe flash drives, the appliance, and 2 × media trays) and they support Fibre Channel, iSCSI, and NVMe TCP connectivity. The new capabilities remove complexity such as data reduction always being on, Dynamic Drive Protection removes complicated RAID setup, and Dynamic Carbon Reduction delivers real-world reduction in power consumption. In addition, the models are FIPS compliant.

In short, the Hitachi Virtual Storage Platform One Block series combines simplicity, sustainability, and robust security features to optimize system management, energy efficiency, and data protection.

See https://www.hitachivantara.com/en-us/products/storage-platforms/block-storage/midrange/vsp-one-block for more information.

Hitachi Virtual Storage Platform 5000 series

This enterprise-class, flash array evolution storage, Hitachi Virtual Storage Platform 5000 series has an innovative, scale-out design optimized for NVMe and storage class memory. It achieves the following:

- **Agility using NVMe:** Speed, massive scaling with no performance slowdowns, intelligent tiering, and efficiency.
- **Resilience:** Superior application availability and flash resilience. Your data is always available, mitigating business risk.
- **Storage simplified:** Do more with less, integrate AI (artificial intelligence) and ML (machine learning), simplify management, and save money and time with consolidation.

See https://www.hitachivantara.com/en-us/products/storage-platforms/block-storage/enterprise for more information.

Hitachi Virtual Storage Platform E1090

The Hitachi Virtual Storage Platform E1090 (VSP E1090) storage system is a high-performance, large-capacity data storage system. The VSP E1090 all-flash arrays (AFAs) support NVMe and SAS solid-state drives (SSDs). The VSP E1090H hybrid models can be configured with both SSDs and hard disk drives (HDDs).

- The NVMe flash architecture delivers consistent, low-microsecond latency, which reduces the transaction costs of latency-critical applications and delivers predictable performance to optimize storage resources.
- The hybrid architecture allows for greater scalability and provides data-in-place migration support.

Hitachi Advanced Server HA820 G3

Hitachi Advanced Server HA820 is a high-performance two-socket rackmount server designed for optimal performance and power efficiency. This allows owners to upgrade computing performance without overextending power consumption and offers non-latency support to environments that require the maximum memory capacity. Hitachi Advanced Server HA820 G3 provides flexible I/O scalability for today's diverse data center application requirements.

Optimized for performance, high density, and power efficiency in a dual-processor server, HA820 G3 delivers a balance of compute and storage capacity. These rack mounted servers have the flexibility to power a wide range of solutions and applications.

The highly scalable memory supports up to 8 TB RAM using 32 slots of 2300 MHz DDR5 RDIMM. HA820 G3 is powered by the Intel Xeon Emerald Rapids scalable processor family for complex and demanding workloads. Flexible OCP and PCIe I/O expansion card options are available.

Network switches

All listed leaf, spine, and management switches in this solution can be replaced by any other VMware vSphere-supported network switches from vendors such as Cisco, Arista, and Extreme. This guide is not to maximum scale, but it is well-balanced for a typical enterprise-scale deployment.

The choice of switch models is based on a scale of five racks, each containing 32 dual-ported 10 Gbps NICs on Hitachi Integrated Systems. Network speed, inter-rack bandwidth, maximum node per rack, and maximum number of racks may differ with the leaf and spine switch models and the required port configuration.

You can select alternative switches to create a well-balanced building block as described in the following selection. Also, the building blocks can be expanded to more racks so there are more compute nodes by using larger spine switches with more ports and bandwidth. There is no limit to how many racks you can have and nodes that can be supported from a hardware perspective.

The Hitachi Integrated Systems platform is only limited by VMware maximum specifications. See https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html to learn more about VMware Cloud Foundation.

Brocade switches from Broadcom

Brocade and Hitachi Vantara have partnered to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

Brocade Fibre Channel switches deliver industry-leading performance with seventh generation 64Gb/sec Fibre Channel interfaces, simplifying scale-out network architectures. Get the high-performance, availability, ease of management, and support for the next generation of Hitachi Virtual Storage Platform storage systems on a solid storage network foundation that can grow as your need grows.

See https://www.broadcom.com/products/fibre-channel-networking/switches for more information.

Cisco Nexus switches

The Cisco Nexus switch product line offers a range of solutions that simplify the connection and management of disparate data center resources through software-defined networking (SDN). Leveraging the Cisco Unified Fabric, which unifies storage, data, and networking (Ethernet/IP) services, the Nexus switches create an open, programmable network foundation built to support a virtualized data center environment.

VMware Cloud Foundation

VMware Cloud Foundation is the hybrid cloud platform for managing VMs and orchestrating containers, built on full-stack hyperconverged infrastructure (HCI) technology. With a single architecture that is easy to deploy, VMware Cloud Foundation enables consistent, secure infrastructure and operations across private and public cloud. Increase enterprise agility and flexibility with the hybrid cloud that delivers it all.

Introducing the VCF Import Tool

The VMware Cloud Foundation (VCF) Import Tool is a command-line interface (CLI) utility introduced in VCF 5.2. It is designed to streamline the transition of existing vSphere environments into a VCF-managed private cloud. This tool extends the capabilities of the Software-Defined Data Center (SDDC) Manager, making it easier to integrate and manage your current infrastructure.

The VCF Import Tool is a powerful utility for organizations looking to modernize their data centers and transition to a private cloud environment with minimal disruption.

VMware Live Site Recovery

VMware Live Site Recovery (VLSR) is a disaster recovery solution that extends VMware vCenter to provide automated orchestration and management of disaster recovery plans. It integrates seamlessly with VMware vSphere and VMware Cloud on AWS, enabling businesses to protect and recover their applications without the need for a dedicated secondary site. By leveraging VMware's enterprise-class Software-Defined Data Center (SDDC) Manager technology, Site Recovery offers robust disaster recovery as a service (DRaaS) capability, ensuring business continuity and minimizing downtime.

Key features of VMware Site Recovery include automated orchestration, which simplifies the creation, testing, and execution of disaster recovery plans through automated workflows. Non-disruptive testing allows for regular testing of disaster recovery plans without impacting production environments, ensuring readiness and compliance. Seamless integration with VMware vSphere and VMware Cloud on AWS provides a unified management experience. Scalability supports a wide range of deployment topologies, including active-passive, active-active, and bi-directional configurations. Cost efficiency is achieved by eliminating the need for a dedicated secondary site, leveraging cloud resources to reduce capital and operational expenses. VMware Site Recovery is an essential tool for organizations seeking to enhance their disaster recovery capabilities, ensuring that critical applications and data are protected and can be quickly recovered in the event of a disaster.

Hitachi Unified Compute Platform Advisor

Hitachi Unified Compute Platform Advisor (UCP Advisor) is a comprehensive cloud infrastructure management and automation software that enables IT agility and simplifies day 0-N operations for edge, core, and cloud environments. The fourth-generation UCP Advisor accelerates application deployment and drastically simplifies converged and hyperconverged infrastructure deployment, configuration, life cycle management, and ongoing operations with advanced policy-based automation and orchestration for private and hybrid cloud environments.

The centralized management plane enables remote, federated management for the entire portfolio of converged, hyperconverged, and storage data center infrastructure solutions to improve operational efficiency and reduce management complexity. Its intelligent automation services accelerate infrastructure deployment and configuration, significantly minimizing deployment risk and reducing provisioning time and complexity, automating hundreds of mandatory tasks.

UCP Advisor improves predictability with guided lifecycle management capabilities for the complete data center infrastructure stack, including servers and switches from Arista, Brocade, and Cisco, and non-disruptively patches and upgrades infrastructure.

UCP Advisor provides deep integration with VMware management software, improving administrator productivity with intuitive and intelligent operations and automation. It complements VMware vRealize software to further streamline the administration and automation of Software-Defined Data Center (SDDC) Manager. Automated workflows deliver IT agility using UCP Advisor REST APIs when used with VMware Cloud Foundation Automation, enable self- services multi-cloud environments.

It provides comprehensive visibility and monitoring of the infrastructure for collective insight into health and operational efficiency. It automates network configuration operations and system monitoring including generating reports for compliance. UCP Advisor and the integrations with VMware Aria Options for Logs provide rich log analytics and auditability enabling comprehensive visibility of the infrastructure for better resource planning.

Conclusion

The solution is designed to ensure continuous operations by automatically addressing different failure scenarios. In the case of full site outages caused by power failures or planned maintenance, an automated site failover mechanism is activated. Storage issues such as controller, disk, or path failures are mitigated by automated path and load re-balancing, ensuring smooth operation without data loss. For ESXi host failures, virtual machines are promptly recovered on healthy hosts using strategic placement policies like affinity rules. Additionally, full region outages—triggered by events ranging from natural disasters to power grid overloads— are managed with orchestrated recovery plans, ensuring that the overall system remains resilient and operational even during severe disruptions.

The architecture is designed to maximize resiliency and availability through a multi-tiered site configuration. Site 1, also known as Availability Zone 1 in Region 1, serves as the primary data center managing most production workloads with an emphasis on high availability and real-time operations. In the event of a failure, operations are redirected to Site 2, or Availability Zone 2 in Region 1, which maintains synchronous mirroring with Site 1 to ensure perfect data synchronization and no loss during the switchover process.

Notably, both Site 1 and Site 2 can function as primary sites for different workload domains, operating as distinct availability zones within a pseudo-cloud region, with one acting as a backup for the other. Additionally, Site 3 in Region 2 functions as a disaster recovery or backup site, providing asynchronous data backup or archiving. This extra layer of redundancy ensures that even if both primary sites become unavailable, Site 3 can still sustain operational continuity, albeit with less aggressive recovery time and recovery point objectives. This seamless integration of multiple layers of protection leads to superior resilience, ensuring that critical environments maintain continuous operations despite disruptions.

Scenario	Test Results	Observed behavior
Host down (vSphere High Availability)	Pass	▪ vSphere HA detects that the ESXi host is down and triggers failover to another ESXi host. ▪ If there is an alive ESXi host in the Primary Site then the VM is restarted there, otherwise it is started on an ESXi host in the Secondary Site. ▪ VMware NMP/PSP uses ALUA information to select Active/Optimized path (which is a Local Path). ▪ Because all ESXi hosts on both sites can access both the P-Vol and S-Vol on the storage system, VMware vSphere High Availability will restart virtual machines on available ESXi hosts at different sites during host failures.
Fibre Channel Single path failure (Local path partially down)	Pass	▪ VMware NMP/PSP uses ALUA information to select the remaining Active/Optimized local path. ▪ When the path is restored, it automatically becomes available and active for use by the ESXi host.
Local path completely down (All active paths to local storage system fail for a single host on either site)	Pass	▪ VMware NMP/PSP uses ALUA information to select the Active/Non-Optimized Cross Paths. ▪ Primary Site ESXi host IO goes to the S-Vol in the secondary site. ▪ When the path is restored, it automatically becomes available and active for use by the ESXi host.
Remote path and replication path down. (all paths down occur in any ESXi host in cluster)	Pass	▪ VMware HA detects all paths down and triggers failover to another ESXi host. ▪ If a functioning host is available at the Primary Site, the VM will be restarted there; otherwise, it will be restarted on a host at the Secondary Site. ▪ After failover, VMware NMP/PSP uses ALUA information to select local active paths.

Scenario	Test Results	Observed behavior
Primary storage down (Site 1 storage failure)	Pass	▪ When ESXi hosts at Site 1 fail, vSphere High Availability automatically restarts the virtual machines on ESXi hosts at Site 2. No disruption is seen on the virtual machine. ▪ Active paths to the primary storage system are reported dead.
Global-active device Quorum down / Path to Quorum is down (Quorum disk failure)	Pass	▪ VMs will continue to operate normally. ▪ P-Vol and S-Vol will continue to replicate normally. ▪ Even if the quorum disk fails or all paths to the quorum disk are removed, replication between the P-Vol and S-Vol will continue and the P-Vol and S-Vol will remain in pair state.
Replication link down (Storage Replication link failure)	Pass	▪ Global-active device determines the volume with the most up-to-date data to operate in Local IO mode where the P-Vol becomes active and the other is in Block IO mode. ▪ If the P-Vol is in Local IO mode and S-Vol is in Block IO mode then, IO from the original ESXi Host in Primary Site will occur to P-Vol via Local Path. ▪ If the P-Vol is in Block IO mode and the S-Vol is in Local IO Mode, then IO from the original Host in Primary Site will occur to S-Vol via Cross Path. VMware NMP/PSP will switch from Local Path to Cross Path. ▪ After replication link issues are resolved, resume replication through the Command Control Interface.
Primary site down (Site 1 failure)	Pass	▪ vSphere HA detects Hosts are down and triggers failover to another Hosts. ▪ VMs restart on Hosts in the Secondary Site.

Reference documents

See the following documents for more information:

VMware Cloud Foundation Administration Guide

Global-Active Device User Guide

Storage Replication Adapter VMware® Live Site Recovery™ Deployment Guide

VMware vSphere Virtual Volumes (vVols) with Hitachi Virtual Storage Platform Quick Start vSAN OSA Compatibility Guide

VMware Compatibility Guide - VSAN OSA

VMware Compatibility Guide - VSAN ESA

VMware Cloud Foundation 5.2.1 Release Notes

Unified Compute Platform (UCP) Advisor Administration Guide

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact