End-to-End High Availability and Resiliency with Hitachi Virtual Storage Platform, Global-Active Device, and Microsoft Technologies

MK-SL-503-00
September 2026

Contents

Feedback....2
Executive Summary....3
Intended Audience....3
Solution Overview....3
Solution components....5
Physical infrastructure....6
Infrastructure at a glance....7
Network architecture....7
Storage foundation with VSP and GAD....9
Microsoft Hyper-V and Windows Server Failover Clustering....14
SQL Server Always On Availability Groups....19
Layered Protection Model....21
Architecture validation test cases....22
Business Value....30
Conclusion....30

Feedback

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

Thank you!

Revision history

Changes	Date
Initial release	September 2026

Executive Summary

Organizations running mission-critical applications require continuous access to data and services despite planned maintenance activities, infrastructure failures, and increasing operational complexity. Delivering this level of resilience requires a coordinated availability strategy that protects storage, compute, networking, clustering, and database services as a unified platform.

Key Takeaways

- Hitachi Global-Active Device (GAD) provides a resilient storage foundation through synchronous active-active replication across geographically separated Hitachi Virtual Storage Platform (VSP) systems.
- A layered architecture extends availability across the infrastructure stack through Microsoft Hyper-V, Windows Server Failover Clustering (WSFC), SQL Server Always on Availability Groups, redundant networking, and Fibre Channel SAN connectivity.
- The solution was validated across storage, network, compute, and database failure domains to demonstrate continuous application availability during planned and unplanned events.

This paper examines the architecture, design considerations, and validation results of a high-availability (HA) platform that helps organizations strengthen business continuity, reduce downtime risk, and improve operational confidence for mission-critical workloads

Intended Audience

This guide is intended for infrastructure architects, storage administrators, Windows Server administrators, virtualization architects, SQL Server administrators, and IT decision makers evaluating HA Microsoft workload architectures built on Hitachi Virtual Storage Platform and GAD.

Solution Overview

Modern enterprise applications depend on the continuous availability of data, infrastructure, and services. The architecture presented in this guide is built on Hitachi Virtual Storage Platform (VSP) systems with GAD, which provide the continuously available storage foundation for the solution. Building on this foundation, Microsoft technologies extend availability into the compute, cluster, and database layers to deliver end-to-end application resilience

Architecture Components

The solution is built from four key technology domains that collectively deliver end-to-end availability:

- Global-active device (GAD): Provides synchronous active-active storage replication and continuous access to enterprise data.
- Compute and Cluster Services: Microsoft Hyper-V and Windows Server Failover Clustering enable workload mobility, health monitoring, and automated recovery.
- Network and SAN Infrastructure: Redundant Ethernet and Fibre Channel connectivity eliminate communication and storage-path single points of failure.
- SQL Server Always On Availability Groups: Extends resilience to the database layer through synchronized replicas and transparent client failover.

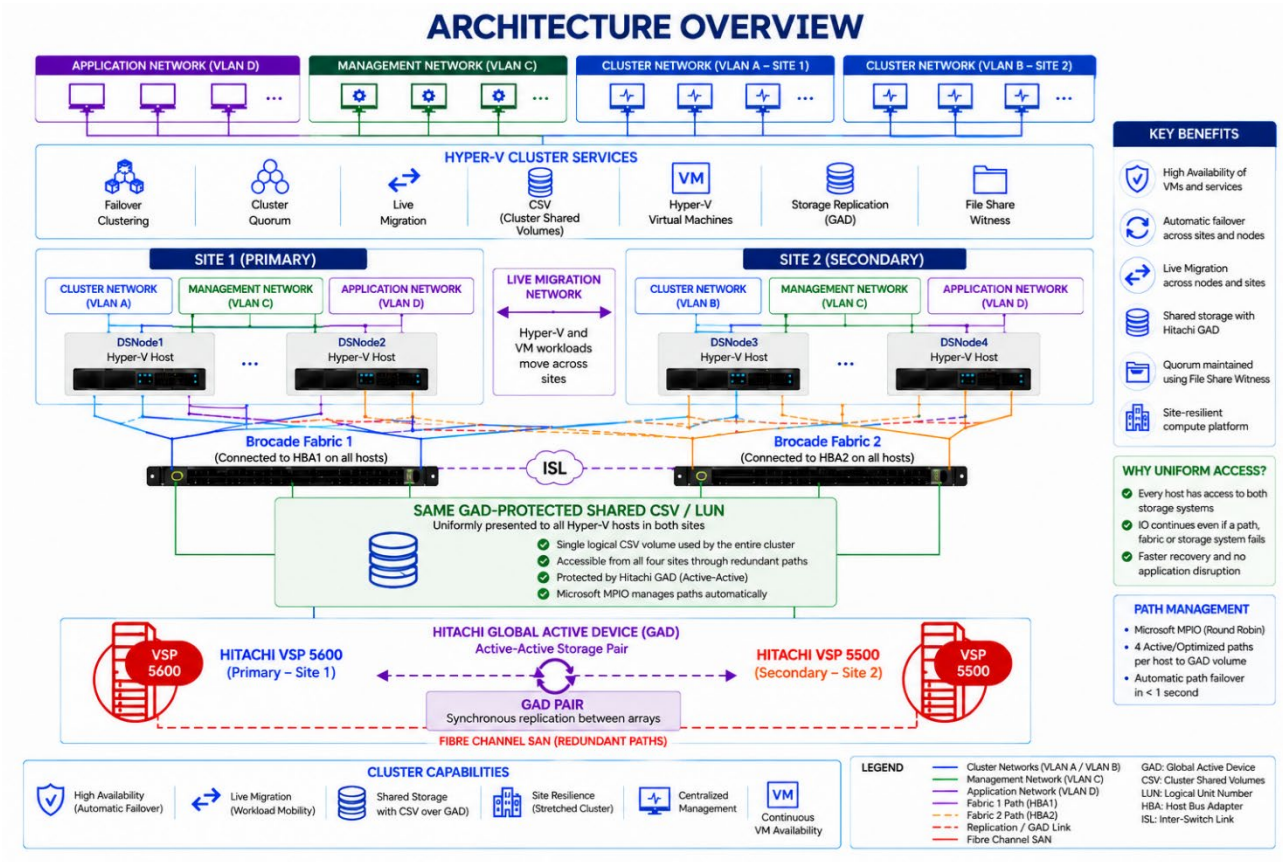

This high-level architecture illustrates how GAD provides the storage foundation for Microsoft Hyper-V, Windows Server Failover Clustering (WSFC), SQL Server Always On Availability Groups, and the supporting network and SAN infrastructure.

Key Solution Benefits

The architecture combines GAD with Microsoft high-availability technologies to deliver coordinated protection across storage, networking, compute, clustering, and database services. Rather than relying on a single availability mechanism, the solution adopts a layered protection model that helps organizations improve operational resilience and business continuity.

Capability	Business Benefit
Continuous Storage Availability	GAD provides synchronous active-active storage replication across sites, enabling continuous access to business-critical data.
End-to-End Resilience	Protection extends across storage, networking, virtualization, clustering, and database layers.
Operational Continuity	Planned maintenance and infrastructure upgrades can be performed with minimal service disruption.
Transparent Application Connectivity	Applications continue using a consistent SQL endpoint through the Availability Group Listener.
Validated Architecture	Demonstrated through operational and failure scenarios across storage, networking, compute, and database services.

Solution components

The enterprise high-availability architecture is built from a set of tightly integrated infrastructure components that collectively deliver resiliency across storage, networking, compute, clustering, and database services. Rather than relying on a single technology to maintain application availability, the solution adopts a layered approach where each component protects a specific failure domain while complementing the capabilities of the surrounding infrastructure.

At the core of the solution are Hitachi Virtual Storage Platform (VSP) systems and global-active device (GAD), which provide the continuously available storage foundation for business-critical workloads. Building upon this foundation, Microsoft Hyper-V, Windows Server Failover Clustering (WSFC), SQL Server Always On Availability Groups, redundant Ethernet networks, and Fibre Channel SAN infrastructure work together to extend availability across the compute, cluster, communication, and database layers.

This integrated design enables applications to remain available during storage disruptions, network failures, host maintenance activities, database failovers, and site-level events. By ensuring that each technology layer contributes to the overall resilience strategy, the architecture minimizes single points of failure and supports continuous business operations.

The following sections examine the key solution components and their role in delivering end-to-end application availability:

- Physical Infrastructure provides the compute, storage, and SAN foundation that supports the high-availability platform.
- Network Architecture delivers resilient communication pathways for storage synchronization, cluster operations, database replication, workload mobility, and application connectivity.
- Hitachi Virtual Storage Platform (VSP) and GAD provide the continuously available storage services that form the foundation of the architecture.
- Microsoft Hyper-V and Windows Server Failover Clustering (WSFC) provide workload mobility, infrastructure monitoring, cluster coordination, and automated recovery capabilities.
- SQL Server Always On Availability Groups extend resiliency to the database layer through synchronized replicas, automatic failover, and transparent client connectivity.

Together, these components establish a coordinated availability platform designed to maintain application access, preserve data integrity, and strengthen business continuity across multiple infrastructure failure domains.

Physical infrastructure

The Hitachi Integrated Solutions (HIS) compute platform serves as the foundation of the enterprise high-availability platform by distributing compute, storage, and SAN resources across two independent sites. This design helps eliminate single points of failure while providing the resiliency required to support mission-critical Microsoft workloads.

This architecture consists of Hitachi Integrated Solutions compute platforms, redundant Fibre Channel SAN fabrics, and Hitachi Virtual Storage Platform (VSP) systems configured with GAD. Together, these components provide the physical foundation upon which Microsoft Hyper-V, Windows Server Failover Clustering (WSFC), and SQL Server Always On Availability Groups deliver end-to-end application availability.

By separating resources across sites and incorporating redundancy at every layer, the infrastructure is designed to support planned maintenance activities, component failures, and site-level events while maintaining access to critical applications and data.

Physical architecture diagram

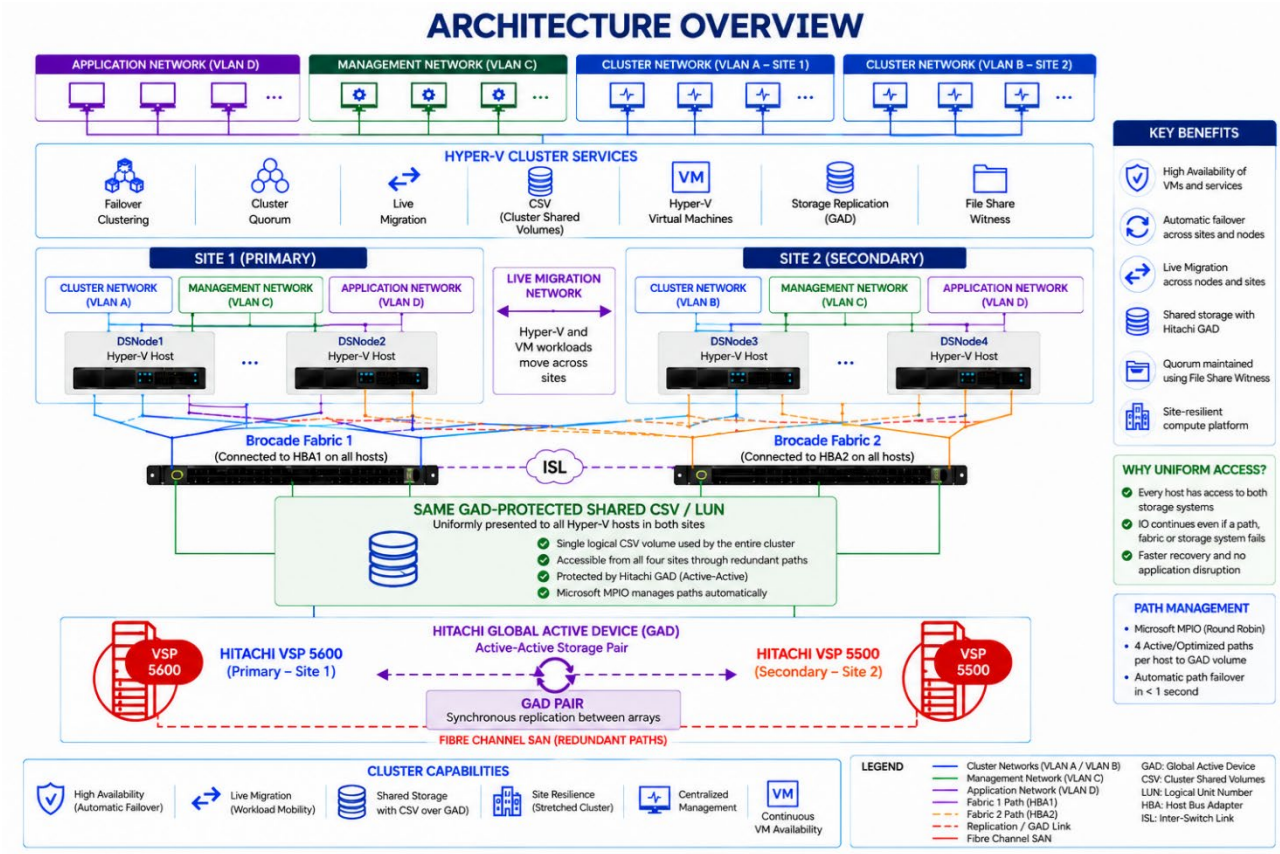

Figure 2. Physical deployment of the enterprise high-availability solution across two sites, illustrating redundant compute, SAN, network, and storage infrastructure.

Infrastructure at a glance

The physical infrastructure provides the foundation upon which the entire high-availability architecture operates. The following table summarizes the key infrastructure components deployed across the two-site environment and their role within the overall solution

Infrastructure Layer	Technology	Configuration
Storage	Hitachi Virtual Storage Platform	Global-active device (GAD)
Compute	Hitachi Integrated Solutions	Four Hyper-V hosts
Virtualization	Microsoft Hyper-V	Windows Server OS
SAN	Brocade Fibre Channel	Dual Fabric Architecture
Database	SQL Server Enterprise	Always On Availability Groups
Cluster	Windows Server Failover Clustering	High Availability Services
Identity	Active Directory	Authentication, DNS, and File Share Witness

Compute platform flexibility

While this architecture was validated using Hitachi Integrated Solutions compute infrastructure, the solution is not dependent on a specific server model. Organizations can deploy the architecture on supported Hitachi Integrated Solutions platforms that satisfy the performance, connectivity, and availability requirements of the target workload.

Similarly, GAD is supported across a broad range of Hitachi Virtual Storage Platform (VSP) systems, allowing organizations to align the architecture with existing infrastructure standards and business requirements.

Key design considerations

- Two-site architecture to improve infrastructure resiliency and provide business continuity.
- Redundant SAN connectivity to eliminate storage-path single points of failure.
- Hitachi VSP systems configured with GAD for continuous storage availability.
- Compute resources distributed across sites to support workload mobility and failover operations.
- Infrastructure designed to support planned maintenance and business continuity objectives.

Network architecture

A resilient network architecture is essential to maintaining continuous application availability across storage, compute, clustering, and database services. The network infrastructure supports multiple communication paths, including client access, cluster coordination, virtual machine mobility, storage connectivity, and application traffic. By separating these functions into dedicated network segments, the architecture improves resiliency, simplifies operations, and reduces the likelihood that congestion or failure within one network affects other critical services.

The network design complements the layered availability strategy implemented throughout the solution. While GAD provides the continuously available storage foundation, the network infrastructure ensures reliable communication between storage systems, Hyper-V hosts, Windows Server Failover Clustering (WSFC), SQL Server Always On Availability Groups, and client applications. Together, these technologies enable end-to-end application resilience across multiple failure domains.

To support operational flexibility and simplified application communication, the VLANs used throughout the environment are designed to be inter-routable where appropriate, allowing infrastructure services, clustered workloads, SQL Server Availability Group communications, and application traffic to communicate across designated network segments while maintaining logical separation and security boundaries.

Enterprise Network Architecture

The following figure shows the enterprise network architecture with dedicated communication paths for management, cluster operations, virtual machine traffic, live migration, storage connectivity, and application access.

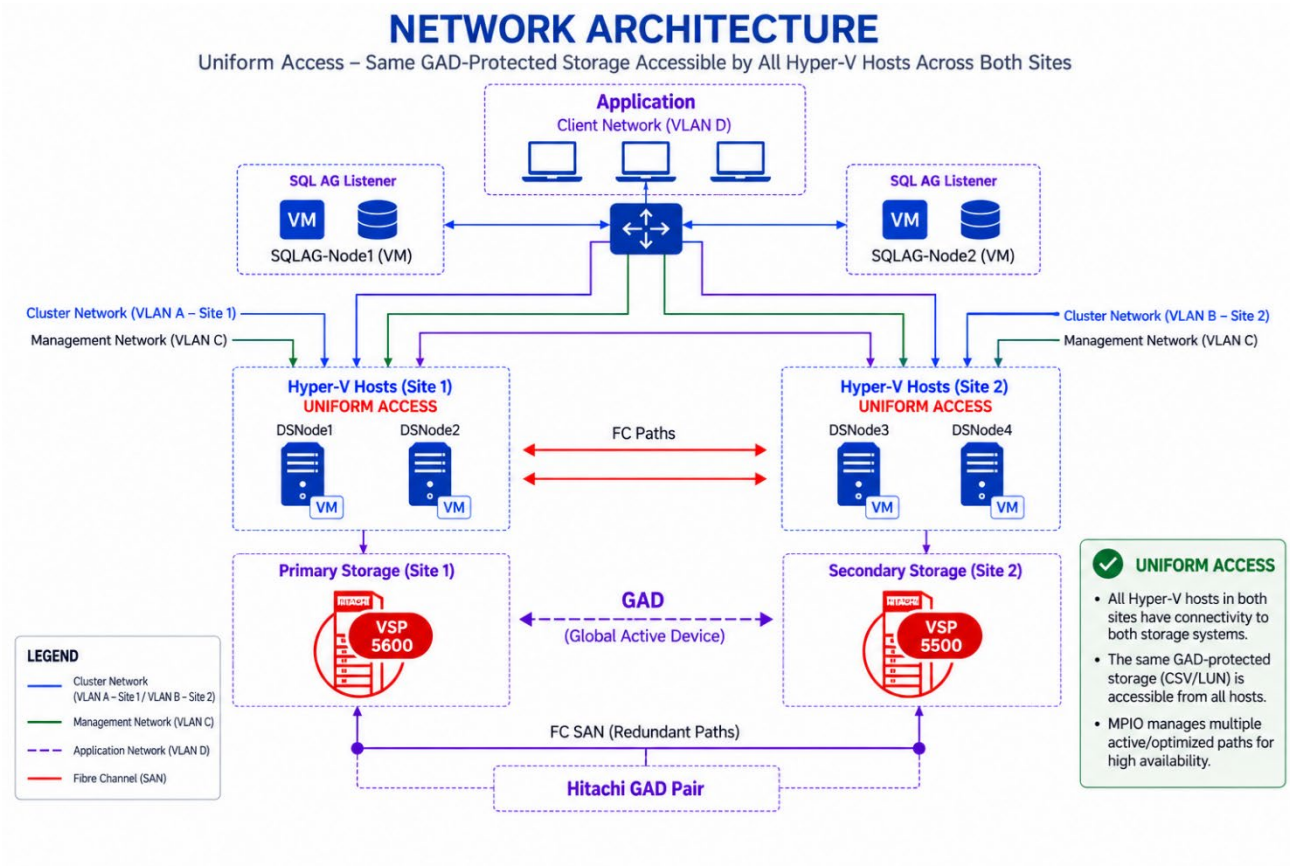

Network design at a glance

A resilient network architecture is essential to support storage synchronization, cluster communications, virtual machine mobility, database replication, and client connectivity. The following table outlines the purpose and availability requirements of each network segment within the solution

Network	Purpose
Management	Administration and monitoring
Client Access	Application and SQL client connectivity
Cluster Communication	WSFC heartbeat and cluster coordination
Live Migration	Hyper-V virtual machine mobility
Fibre Channel SAN	Connectivity to VSP storage systems
Virtual Machine Network	Production application and workload traffic

Network Design Best Practices

The following design principles were incorporated to support a resilient enterprise availability platform:

- Use dedicated network segments for management, application access, cluster communication, virtual machine traffic, and Live Migration.
- Implement redundant Ethernet and Fibre Channel connectivity to eliminate single points of failure.
- Maintain separate communication paths for storage access, cluster heartbeats, workload mobility, and client connectivity.
- Ensure VLANs are appropriately routed to support application communications while preserving traffic isolation and security requirements.
- Provide sufficient bandwidth and network resiliency to support GAD synchronization, SQL Server replication, Windows Server Failover Clustering (WSFC) communications, and Hyper-V operations.

Role in the High-Availability Architecture

Networking acts as the communication fabric that connects every component within the solution. It enables GAD storage synchronization, supports cluster health monitoring and failover operations, facilitates Hyper-V Live Migration, maintains SQL Server replica communications, and provides application access to business services.

By combining dedicated network segmentation, redundant communication paths, and resilient SAN connectivity, the architecture ensures that network failures remain isolated events rather than becoming application outages.

Storage foundation with VSP and GAD

Continuous application availability depends on multiple infrastructure layers working together. At the foundation of this architecture is a resilient storage platform, which provides the data availability required by virtualization, clustering, networking, and database services.

Every database transaction, virtual machine operation, and application request ultimately depends on reliable access to persistent data. While virtualization platforms, clustering technologies, and database replication provide
resiliency at higher layers of the infrastructure stack, their effectiveness depends on the availability of the underlying storage platform.

This solution is built on Hitachi Virtual Storage Platform (VSP) systems and GAD, which together provide a continuously available storage foundation for mission-critical workloads. VSP delivers enterprise-class performance, scalability, and reliability, while GAD extends those capabilities through synchronous active-active replication across geographically separated storage systems.

Global-active device

Global-Active Device (GAD) is the cornerstone of the architecture and the primary technology that enables continuous data availability across sites. Through synchronous active-active replication, GAD maintains identical copies of production volumes on two Hitachi VSP systems while presenting them to hosts as a single logical storage resource.

Unlike traditional replication architectures that rely on a primary storage system and a passive recovery copy, GAD enables both storage systems to actively participate in maintaining application data availability. This architecture helps organizations sustain planned maintenance activities, storage failures, communication interruptions, and site-level events while preserving data consistency and business continuity.

By providing continuously available storage services, GAD establishes the resilient foundation upon which Microsoft Hyper-V, Windows Server Failover Clustering (WSFC), SQL Server Always On Availability Groups, and the surrounding network infrastructure depend.

Why global-active device?

Global-active device extends beyond traditional storage replication by combining storage virtualization, synchronous active-active replication, quorum-based arbitration, and intelligent failure handling into a unified availability platform.

Key advantages include:

- Continuous storage availability across geographically separated sites.
- Synchronous active-active replication with real-time data consistency.
- Storage virtualization through a single logical storage view.
- Automated failure handling and pair-state management.
- Quorum-based protection against split-brain conditions.
- Integration with Microsoft Hyper-V, WSFC, SQL Server Always On Availability Groups, and Microsoft Multipath I/O (MPIO).

Together, these capabilities help organizations maintain business continuity while minimizing the operational impact of planned maintenance activities and unexpected infrastructure failures.

GAD Storage Architecture

Global-Active Device virtualizes storage across geographically separated Hitachi VSP systems, presenting them to hosts as a single logical storage resource. Through synchronous data replication, every write operation is committed to both storage systems, ensuring data consistency and enabling continuous availability across sites. From the host perspective, storage continues to be accessed through standard multipath connectivity, while GAD transparently manages storage synchronization behind the scenes.

The architecture leverages Virtual Storage Machines (VSMs) and Microsoft Multipath I/O (MPIO) to simplify host access and maintain continuous connectivity. This design enables applications and virtual machines to operate without awareness of the underlying storage replication mechanisms while benefiting from the resilience provided by GAD.

The following figure shows the global-active device architecture with synchronous active-active replication between geographically separated Hitachi Virtual Storage Platform systems, providing a resilient storage foundation for the enterprise high availability solution.

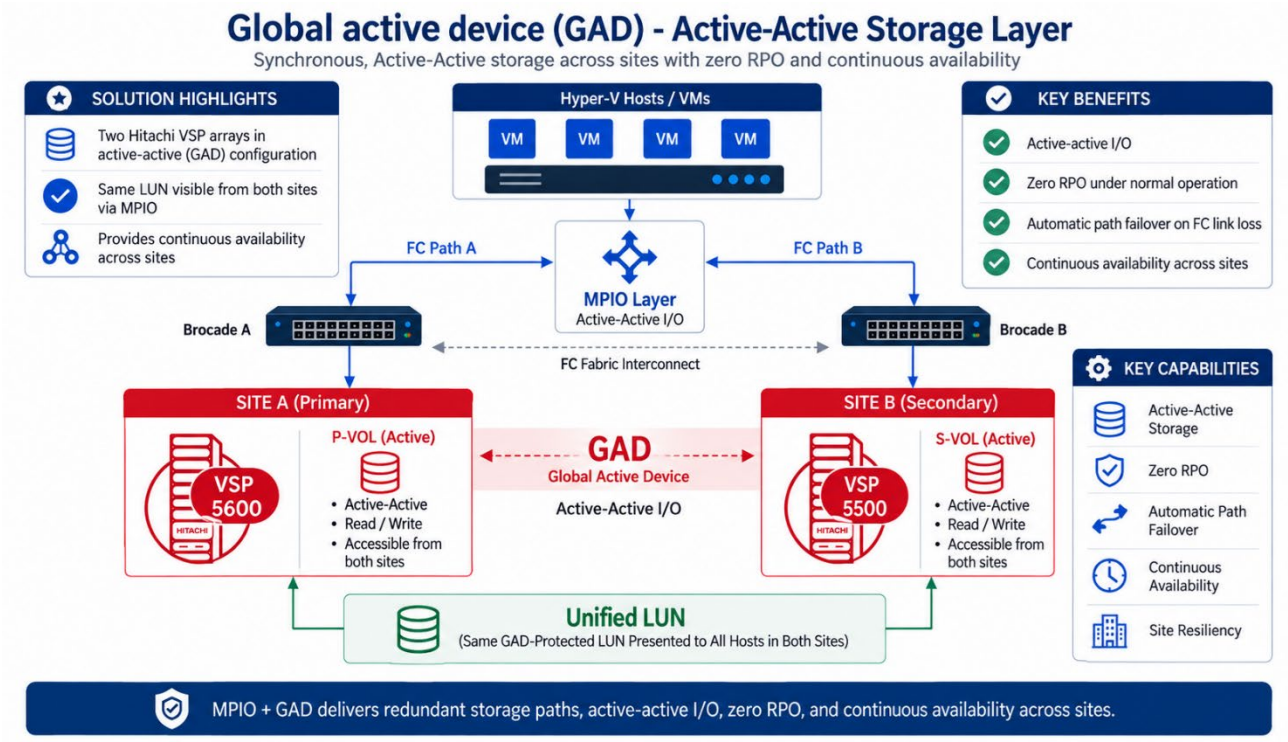

Storage access model

The validated architecture uses a **Uniform Access** storage model, in which Hyper-V hosts maintain storage connectivity to both Hitachi VSP systems participating in the GAD configuration. Through Microsoft Multipath I/O (MPIO) and redundant Fibre Channel SAN fabrics, hosts can access GAD-protected volumes through multiple available paths while GAD transparently maintains synchronized storage copies across both sites.

This approach improves storage resiliency by eliminating dependency on a single storage path or storage system. In the event of a path, fabric, or storage component disruption, workloads continue accessing the same storage resources through alternative available paths, supporting continuous application availability.

Operational view of global-active device

The following screenshots illustrate the operational state of the GAD environment used during validation testing. Together, they demonstrate synchronized active-active storage volumes and healthy pair relationships across both storage systems.

GAD pair synchronization status

The following figure from Hitachi Storage Navigator displays the Global-Active Device (GAD) pair status, showing synchronized storage volumes operating in the PAIR state. This status confirms active replication, data consistency, and reliable communication between the participating Hitachi VSP storage systems.

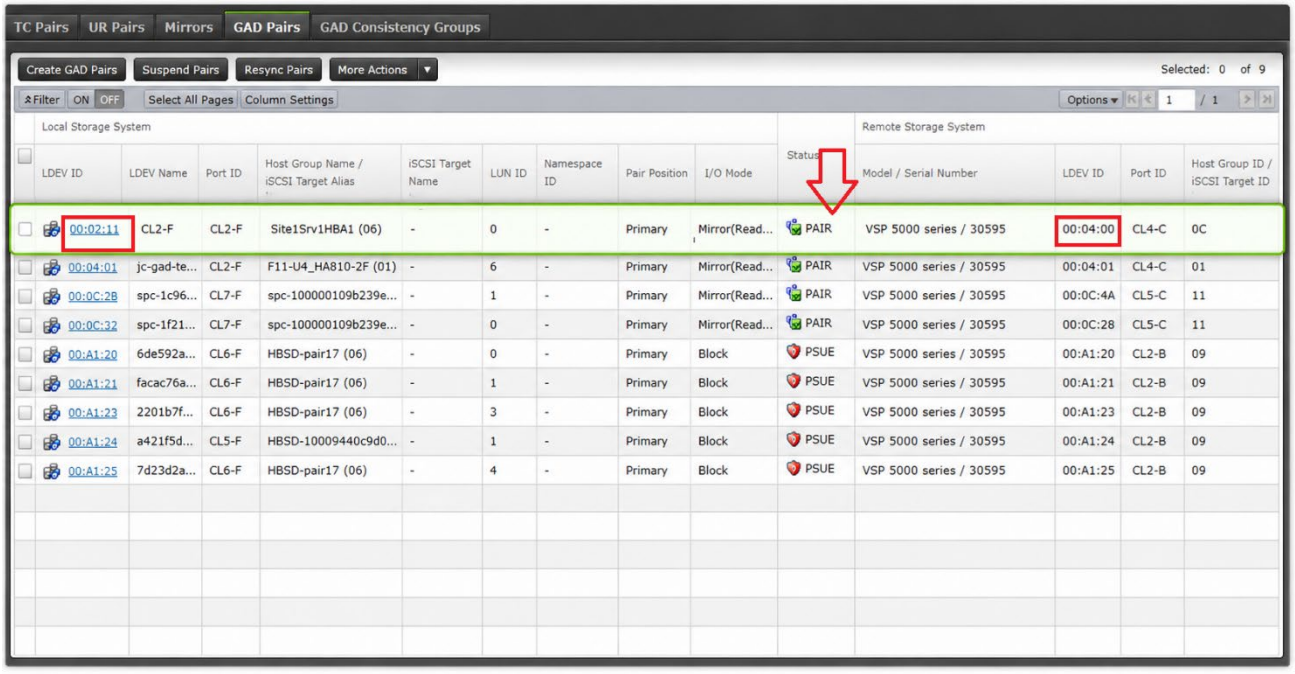

Host integration and storage presentation

Successful integration between GAD and Microsoft technologies depends on proper host presentation and interoperability settings. Host Mode and Host Mode Options ensure that storage resources are presented correctly to Windows Server, Hyper-V, WSFC, SQL Server, and MPIO environments.

The following figure shows the Host Group configuration in Hitachi Storage Navigator, with Host Mode Option 2C (Windows Extension) enabled for Microsoft Windows Server workloads.

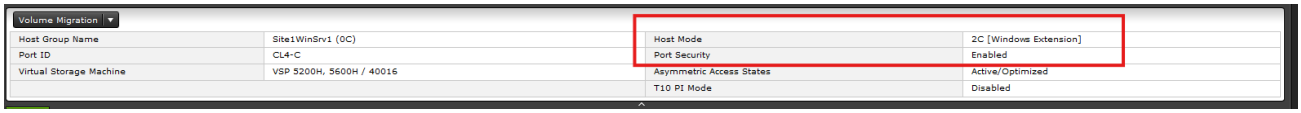

Local and remote volume relationship

The following figure shows the global-active device management view with local and remote volume relationships maintained through synchronous active-active replication across geographically separated storage systems.

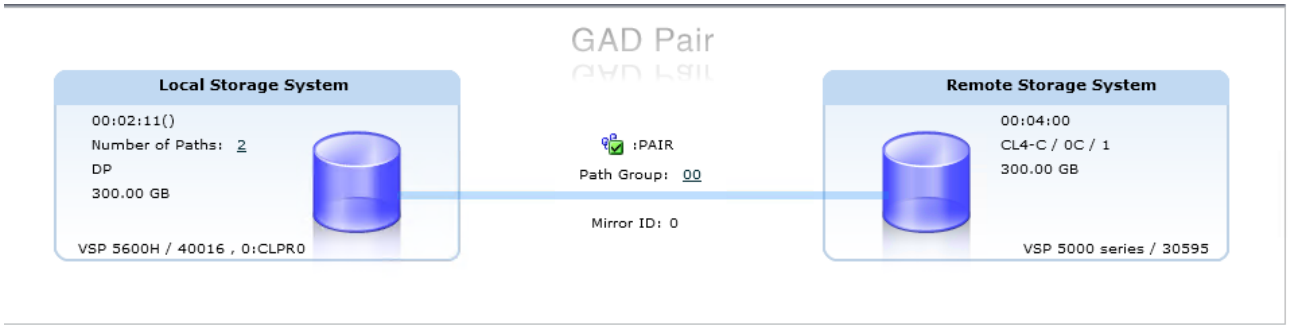

Pair Detail		
Status		PAIR
Processing Status		
Failure Factor		
Quorum Disk ID (LDEV ID)		03(-)
CTG ID		
Copy Pace		15 Tracks
Paired Time		2026/06/23 08:35:00
Last Update Time		2026/06/23 08:35:04
Pair Copy Time		000:00:02
Local Volume	I/O Mode	Mirror(Read Local)
	ALUA Mode	Disabled
Pair Operating Mode When Quorum Disk Blocked		-
I/O Preference Mode When Remote Path Failed		Disabled
Virtual Storage Machine	Model / Serial Number	VSP 5200H, 5600H / 40016
	LDEV ID	00:02:11
	Device Name	
	SSID	

GAD Failure Behavior

A key differentiator of global-active device is its deterministic failure-handling behavior.

To maintain storage availability while preserving data integrity, GAD uses quorum-based arbitration and intelligent pair-state management to determine which storage copy remains active during disruptive events.

Failure Response Summary

Failure Condition	GAD Response	Expected Outcome
Storage System Failure	Surviving copy continues Local I/O operations	Application access continues through the available storage system
Inter-Site Link Failure	Pair transitions to PSUE state	Local storage access remains available while replication is temporarily suspended
Quorum-Related Failure	Arbitration determines active copy	Split-brain conditions prevented and data consistency maintained
Communication Restored	Pair automatically resynchronized	Protected active-active operation restored

This behavior allows GAD to preserve both storage availability and data consistency while enabling higher-layer services such as Hyper-V, WSFC, and SQL Server Always On Availability Groups to continue operating during infrastructure disruptions.

Quorum and Arbitration

The GAD quorum provides an independent arbitration mechanism for the paired storage systems. It helps determine which storage system is permitted to continue host I/O when communication between the storage systems is disrupted or when multiple failure conditions occur. This arbitration is an important protection against split-brain conditions, where both storage systems could otherwise attempt to accept writes independently.

The quorum can be hosted on an external storage system located at a third site, providing an independent failure domain. Where a physical third site is not available, GAD Cloud Quorum provides an alternative architecture, including the option to deploy the quorum services through the Microsoft Azure Marketplace and AWS Marketplace.

For Azure Marketplace see *Global-Active Device Cloud Quorum* at https://marketplace.microsoft.com/en-us/product/hitachivantara.gad_cloud_quorum_v2-0-0?tab=Overview
For AWS Marketplace see *Global-Active Device Cloud Quorum* at
AWS Marketplace: Global-Active Device Cloud Quorum

The quorum does not replicate application data. Its role is to provide arbitration and I/O authority, helping GAD determine which storage copy should remain active during a failure.

This guide focuses on how GAD integrates with Microsoft technologies to provide end-to-end application availability. Readers interested in advanced architecture, supported configurations, quorum design, operational behavior, and implementation guidance.

See *Global Active Device for VSP One Block* on the Hitachi Vantara Product Documentation portal (docs.hitachivantara.com) for more information.

Microsoft Hyper-V and Windows Server Failover Clustering

While GAD provides the continuously available storage foundation for the architecture, application availability also depends on the ability of workloads to remain operational during server failures, maintenance activities, and infrastructure disruptions. To address these requirements, the solution integrates Microsoft Hyper-V and Windows Server Failover Clustering (WSFC) to deliver a resilient compute platform capable of sustaining failures across individual hosts and sites.

The following figure shows Microsoft Hyper-V and Windows Server Failover Clustering extending the availability capabilities of GAD through workload mobility, health monitoring, and automated recovery operations.

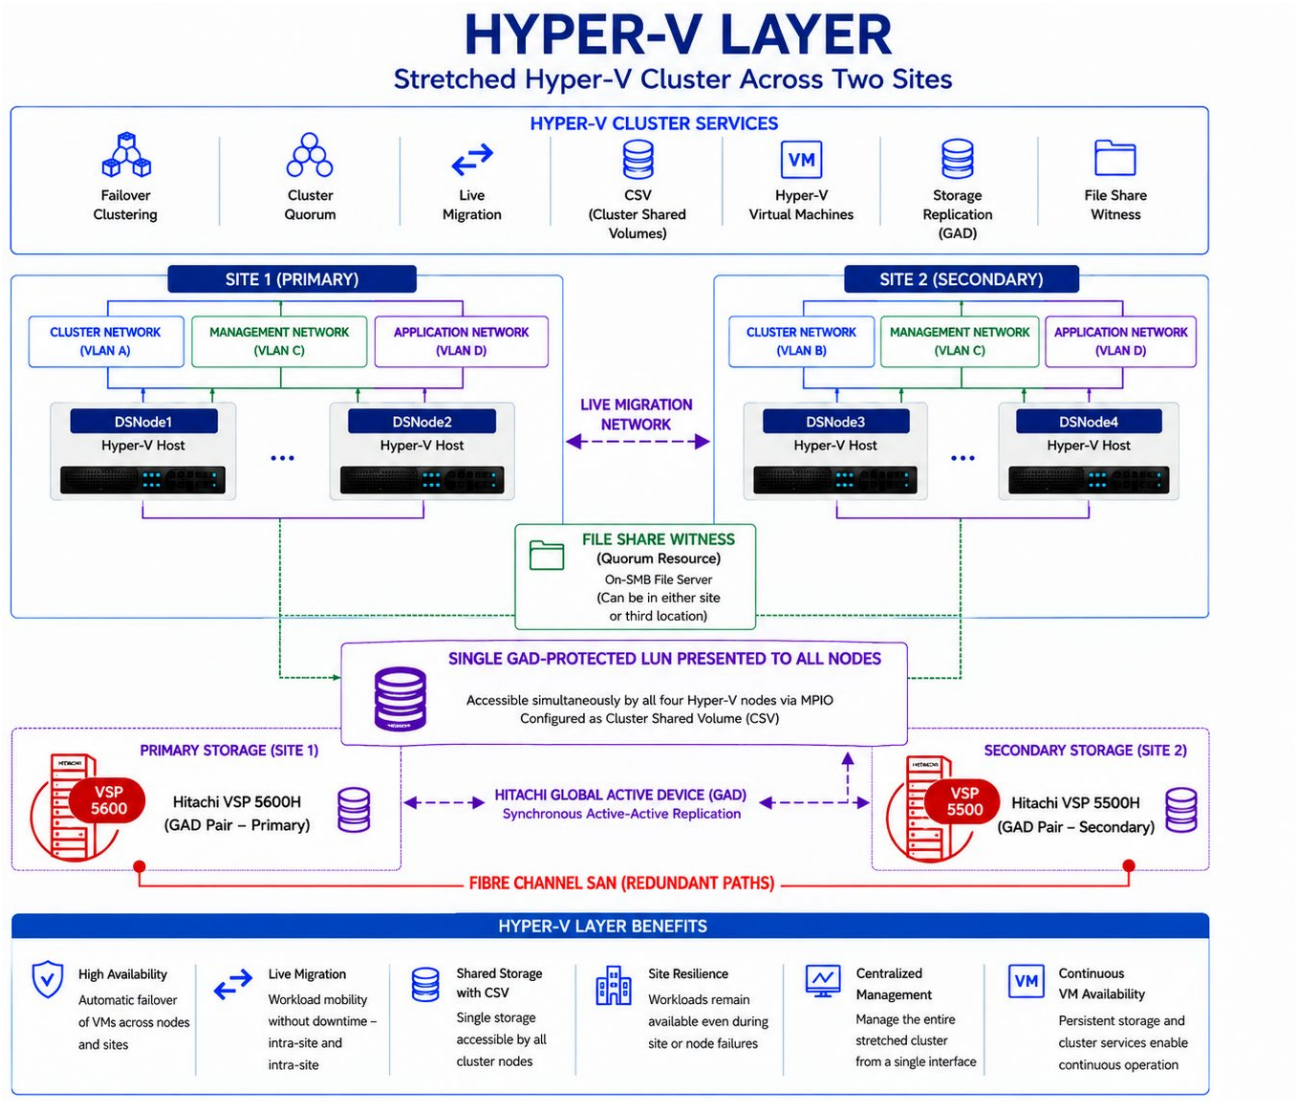

Best Practice

Combine Microsoft Hyper-V and Windows Server Failover Clustering with GAD protected storage to separate compute resiliency from storage resiliency. This design enables workloads to move, recover, and restart across cluster nodes while maintaining continuous access to application data through the underlying GAD storage platform.

Together, Hyper-V and WSFC extend the availability benefits of GAD into the compute layer, enabling virtual machines to move, restart, and recover without requiring changes to the underlying storage architecture. This separation of storage and compute resiliency helps minimize operational risk while supporting continuous business operations.

- Review supported **Hitachi VSP One Block** platforms and **Global-Active Device (GAD)** configurations in the Product Interoperability Matrix before deployment.
- Validate firmware, operating system, Hyper-V, WSFC, and SQL Server interoperability requirements before implementation.

Microsoft Hyper-V

Microsoft Hyper-V provides the virtualization platform for hosting SQL Server virtual machines and supporting application workloads. By abstracting workloads from the underlying hardware, Hyper-V enables workload mobility, simplified administration, and efficient utilization of compute resources. Because workloads continue accessing the same GAD-protected storage resources regardless of host location, virtual machines can move between cluster nodes without requiring storage migration or application reconfiguration.

Windows Server Failover Clustering

Windows Server Failover Clustering (WSFC) continuously monitors cluster health and coordinates workload recovery across the environment. When a host becomes unavailable or enters maintenance mode, WSFC automatically initiates recovery actions and redistributes workloads to available cluster nodes. Combined with the continuously available storage provided by GAD, this enables applications to remain available without requiring manual intervention.

Hyper-V Cluster validation environment

The following figure shows Windows Server Failover Cluster deployed across four Hyper-V hosts spanning two sites. The cluster coordinates workload placement, health monitoring, and recovery operations while leveraging GAD-protected storage.

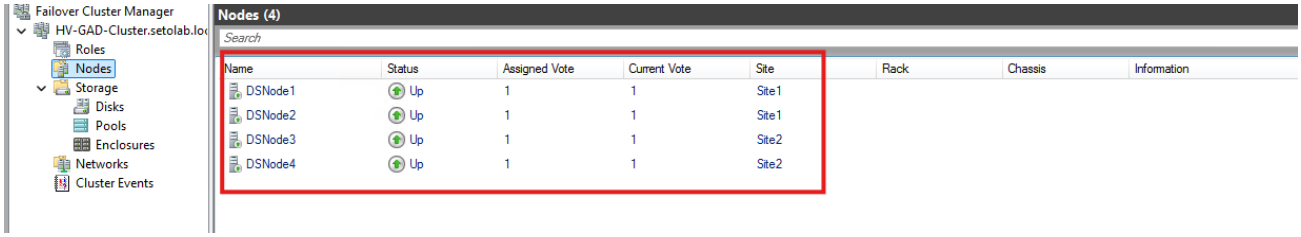

Resilient Storage Connectivity Through MPIO

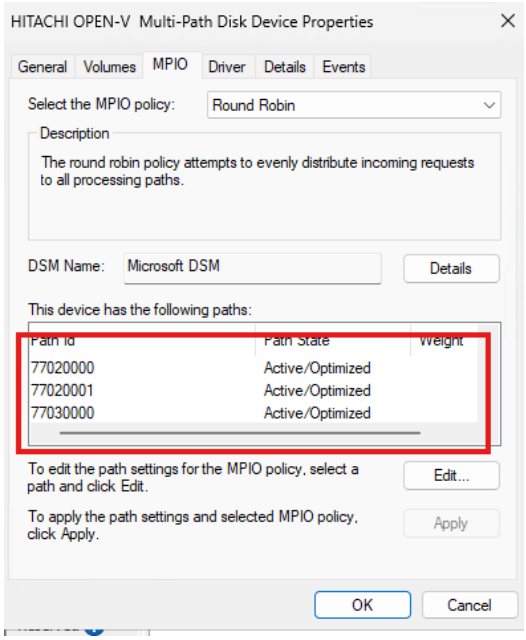

The following screen capture shows Microsoft Multipath I/O (MPIO) presenting four active storage paths to GAD-protected volumes. Multiple active paths help eliminate storage connectivity as a single point of failure and support continuous access to enterprise data.

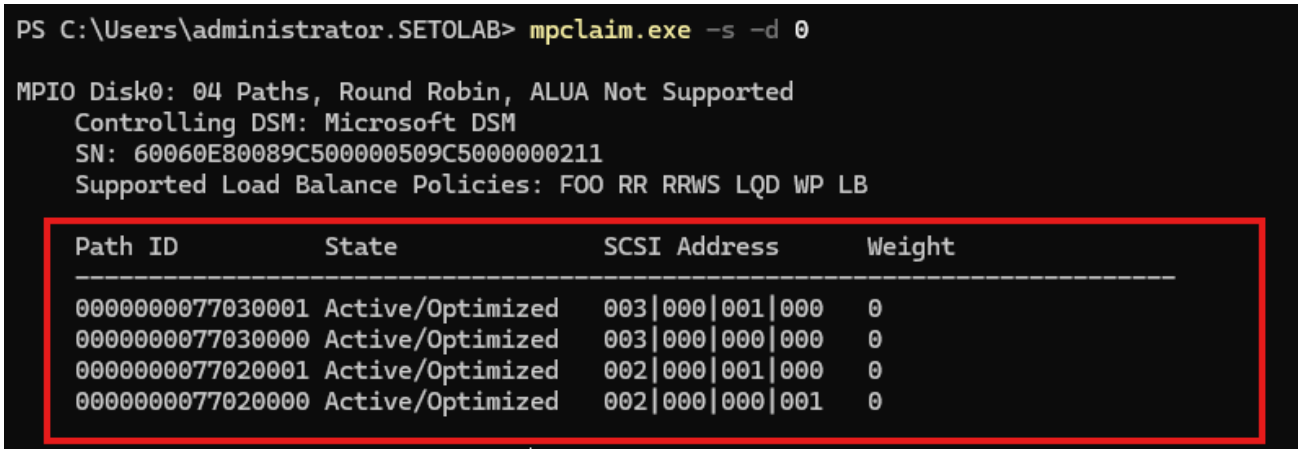

Cluster Shared Volumes and Virtual Hard Disks

Microsoft Hyper-V virtual machines are hosted on **Cluster Shared Volumes (CSV)** backed by **Global-active device (GAD)** protected storage. The virtual machine configuration files and **VHDX** disk files reside on shared storage accessible by all cluster nodes, enabling workloads to move between hosts while maintaining continuous access to application data.

This shared-storage architecture supports Hyper-V Live Migration, automated failover through Windows Server Failover Clustering (WSFC), and continuous availability of virtual machine workloads across both sites.

Cluster quorum

Cluster quorum is the mechanism used by Windows Server Failover Clustering to determine whether sufficient cluster resources remain available to continue providing services safely. Its purpose is to maintain cluster integrity and prevent split-brain conditions during node, network, or site-level failures.

In the validated architecture, four Hyper-V hosts are distributed evenly across two sites and participate in cluster voting. A File Share Witness provides an additional arbitration vote, resulting in a total of five available votes. As long as a majority of votes remain available, the cluster retains quorum and continues operating.

The following figure shows a quorum voting model for the stretched Hyper-V cluster. The File Share Witness provides an additional arbitration vote, enabling the surviving site to maintain cluster services during a site-level failure.

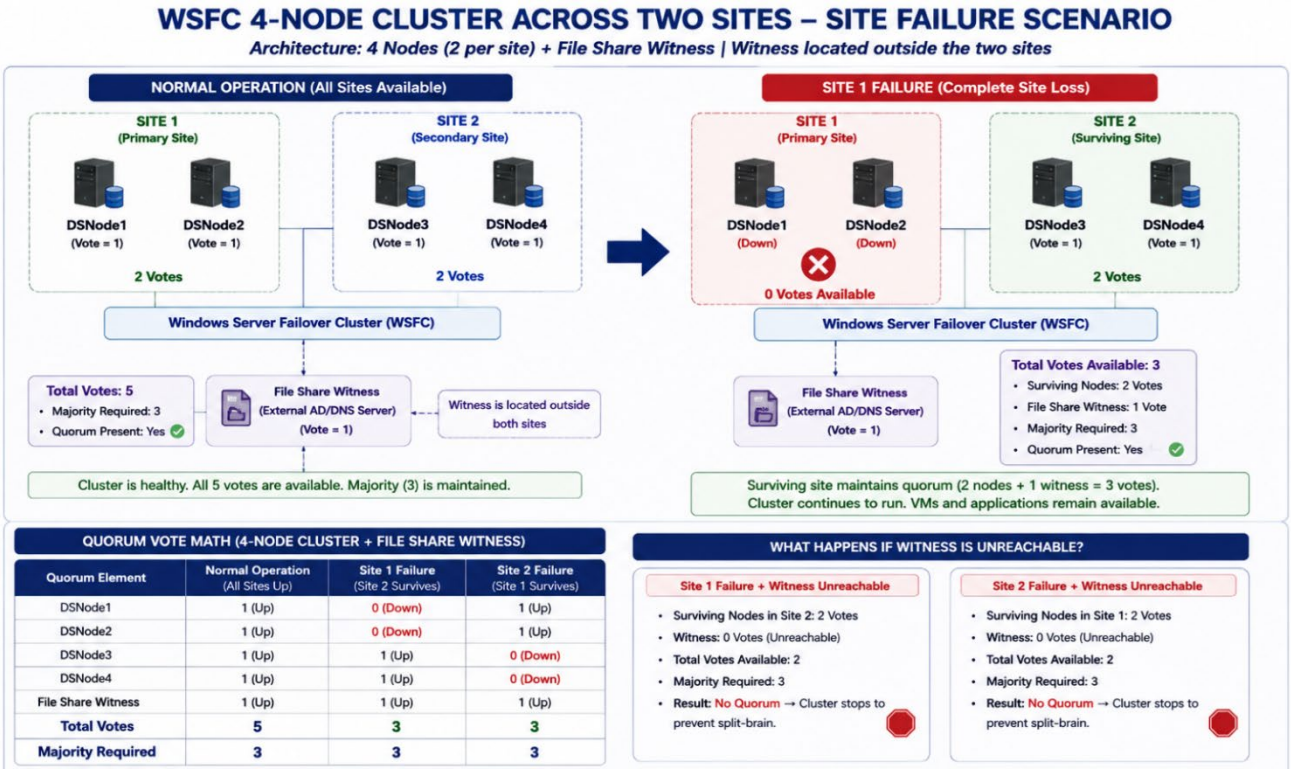

Hyper-V and WSFC Design Best Practices

The validated architecture incorporates several Microsoft and Hitachi design practices to improve workload availability, operational flexibility, and infrastructure resiliency when deploying Microsoft Hyper-V and Windows Server Failover Clustering together with GAD.

- **Distribute Hyper-V cluster nodes across multiple sites** to improve resiliency during host and site-level failures.
- **Leverage GAD-protected Cluster Shared Volumes (CSV)** to provide continuous access to virtual machine data during maintenance activities and infrastructure disruptions.
- **Use dedicated networks for cluster communications and Live Migration traffic** to improve performance, fault isolation, and operational stability.
- **Implement a quorum witness** to maintain cluster availability and prevent split-brain conditions during infrastructure failures.
- **Validate MPIO connectivity across redundant SAN fabrics** to ensure continuous storage access and automatic path failover.
- **Use Hyper-V Live Migration for planned maintenance operations** to minimize the impact of infrastructure updates and hardware servicing.

Microsoft Best Practice Reference

Microsoft recommends designing Hyper-V failover clusters with appropriate quorum configuration, dedicated cluster communications, resilient storage connectivity, and workload mobility capabilities to maximize service availability and support planned maintenance operations.
For additional guidance, see the following Microsoft resources:

- Windows Server Failover Clustering Documentation
  Failover Clustering topologies in Windows Server | Microsoft Learn
- Hyper-V Documentation
  https://learn.microsoft.com/windows-server/virtualization/hyper-v/
- Cluster Shared Volumes (CSV) Documentation
  https://learn.microsoft.com/windows-server/failover-clustering/failover-cluster-csvs
- Hyper-V Live Migration Documentation
  https://learn.microsoft.com/en-us/windows-server/virtualization/hyper-v/manage/live-migration-overview

SQL Server Always On Availability Groups

While global-active device ensures continuous storage availability and the Hyper-V cluster provides a resilient compute platform, business applications ultimately depend on uninterrupted access to the database. SQL Server Always On Availability Groups (AG) extend the high availability architecture to the database layer by maintaining synchronized database replicas and enabling automatic failover between SQL Server instances. Together with Windows Server Failover Clustering, Always On AG provides continuous database availability while presenting a single client connection endpoint through the Availability Group Listener, ensuring application transparency during failover events.

SQL Server Always On Availability Groups Architecture

The following figure shows SQL Server Always On Availability Groups extending availability to the database layer through synchronized replicas, automatic failover, and transparent client connectivity using the Availability Group Listener.

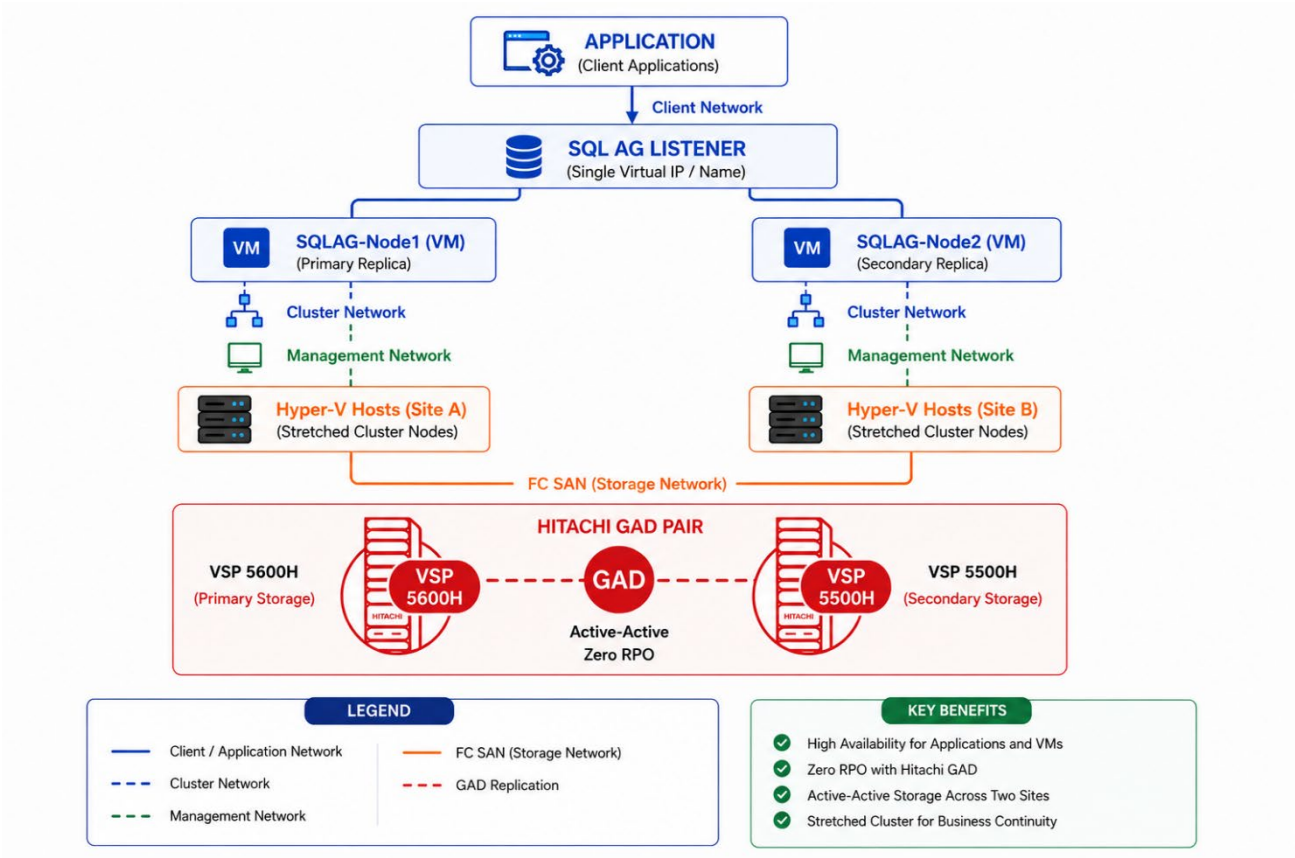

Note that **the application always connects to the Listener**, and never directly to an individual SQL node.

The following table lists architecture components.

Component	Purpose
Primary Replica	Processes all read/write transactions
Secondary Replica	Maintains a synchronized copy of the database
Availability Group Listener	Provides a single client connection endpoint
Windows Server Failover Clustering	Detects failures and coordinates replica failover
Hitachi GAD Storage	Provides continuously available storage for SQL virtual machines

How SQL AG and GAD work together

Although both technologies use synchronous replication mechanisms, they operate independently and provide protection for different failure scenarios.

GAD

- Protects storage accessibility and storage infrastructure.
- Maintains synchronized storage volumes across sites.
- Supports continuous storage availability during storage or site-level failures.

SQL Server Always On Availability Groups

- Protects the database service layer.
- Maintains synchronized database replicas.
- Supports automatic database failover and transparent client redirection through the AG Listener.

By combining GAD with SQL AG, the architecture protects both data accessibility and database service availability. If an SQL Server instance becomes unavailable, the synchronized database replica can assume the primary role while continuing to access storage protected by GAD. This layered approach helps maintain application continuity while reducing the impact of infrastructure and software failures

Unlike Failover Cluster Instances (FCIs), which rely on shared storage for a single SQL instance, Always On AG maintains independent SQL Server instances with synchronized databases, providing greater flexibility for maintenance and failover.

Validated Configuration and Design Best Practices

The validation environment incorporates the following design practices:

- Configure Availability Groups using **Synchronous Commit** mode for critical workloads.
- Use **Automatic Failover** between synchronized SQL replicas.
- Ensure all client applications connect through the **Availability Group Listener**.
- Integrate SQL Server Always On Availability Groups with Windows Server Failover Clustering.
- Deploy SQL Server virtual machines on Hitachi GAD-protected storage volumes.

Microsoft Reference

For SQL Server Always On Availability Groups design guidance, see the following:

- Microsoft SQL Server Always On Documentation
  https://learn.microsoft.com/sql/database-engine/availability-groups/windows/overview-of-always-on-availability-groups-sql-server
- Availability Group Listener Documentation
  https://learn.microsoft.com/sql/database-engine/availability-groups/windows/listeners-client-connectivity-application-failover

Layered Protection Model

Enterprise availability cannot be achieved through a single technology. While storage replication, virtualization, clustering, and database replication each provide protection within their respective domains, every technology addresses only a subset of potential failure scenarios. Continuous application availability is achieved when these technologies operate together as a coordinated platform, where each layer complements the capabilities of the layers above and below it.

The architecture presented in this paper adopts a layered protection model that begins with GAD and extends through SAN connectivity, networking, virtualization, clustering, and database services. By combining complementary technologies into a unified availability strategy, the solution minimizes single points of failure and helps ensure continuous access to business-critical applications and data during both planned maintenance activities and unexpected infrastructure events.

The following figure shows the layered protection architecture, illustrating how storage, networking, virtualization, clustering, and database services work together to deliver end-to-end application availability.

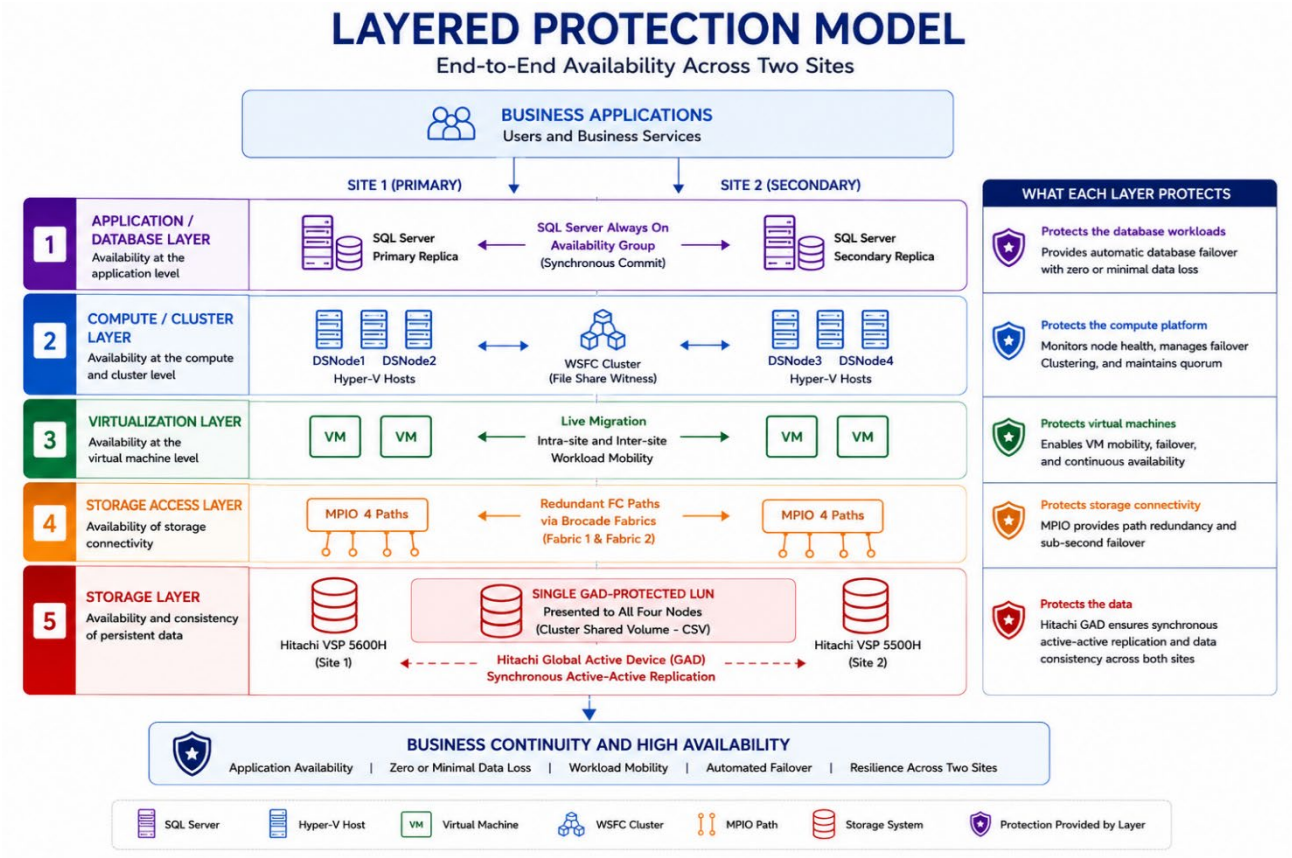

How the layers work together

At the foundation of the architecture, **GAD** maintains continuously available access to enterprise data through synchronous active-active replication across geographically separated Hitachi VSP systems. This storage foundation enables higher-layer technologies to focus on their respective protection domains without depending on a single storage system.

Building upon this foundation, redundant Fibre Channel SAN fabrics, Microsoft Multipath I/O (MPIO), and resilient network infrastructure maintain continuous communication between storage, compute, and database services. Microsoft Hyper-V provides workload mobility and virtualization capabilities, while Windows Server Failover Clustering (WSFC) continuously monitors infrastructure health and coordinates automated recovery actions when failures occur.

At the application layer, SQL Server Always On Availability Groups extend protection to the database service layer through synchronized replicas, automated failover, and transparent client connectivity using the Availability Group Listener. Together, these technologies transform isolated infrastructure failures into manageable operational events while maintaining access to critical applications and data.

Architecture validation test cases

A high-availability architecture should not only be designed according to established best practices but also be validated under realistic operating conditions. To demonstrate the effectiveness of the proposed solution, a series of test cases were run across the storage, network, compute, and database layers. The objective was to validate that Hitachi Global-Active Device (GAD), together with Microsoft Hyper-V, Windows Server Failover Clustering (WSFC), SQL Server Always On Availability Groups, and the supporting network infrastructure, maintained application availability and data accessibility during both planned maintenance activities and infrastructure disruptions.

The validation focused on common operational scenarios and failure conditions that organizations might encounter in production environments. These tests were designed to verify the effectiveness of the layered protection model and demonstrate how individual infrastructure layers work together to sustain business operations.

Validation summary

The following table summarizes the observed outcomes recorded during validation testing.

Test Case	Observed Result
Fibre Channel Storage Path Failure	Recovery completed in less than 1 second
Network Failure Recovery	Recovery completed in less than 1 second
SQL Server Always On AG Planned Failover (Intra-Site)	Application recovery completed in approximately 5-6 seconds
SQL Server Always On AG Planned Failover (Inter-Site)	Application recovery completed in approximately 7-9 seconds
SQL Primary Replica Forced Reboot	Application recovery completed in approximately 10-11 seconds
Hyper-V Live Migration (Intra-Site)	No transaction interruption observed
Hyper-V Live Migration (Inter-Site)	No transaction interruption observed

Note: Observed results represent measurements recorded within the validation environment and may vary depending on workload characteristics, infrastructure sizing, operating system settings, network performance, and application behavior.

Test Case 01 – Fibre Channel Storage Path Failure

Objective
Validate that application availability is maintained following the loss of a Fibre Channel storage path.

Validation activity
One Fibre Channel storage path was intentionally disabled to simulate a SAN connectivity failure while application workloads remained active.

The following figure shows the highlighted path that was intentionally disabled to simulate a SAN connectivity failure. Microsoft Multipath I/O (MPIO) automatically redirected storage traffic through the remaining active paths, maintaining uninterrupted access to GAD-protected storage volumes.

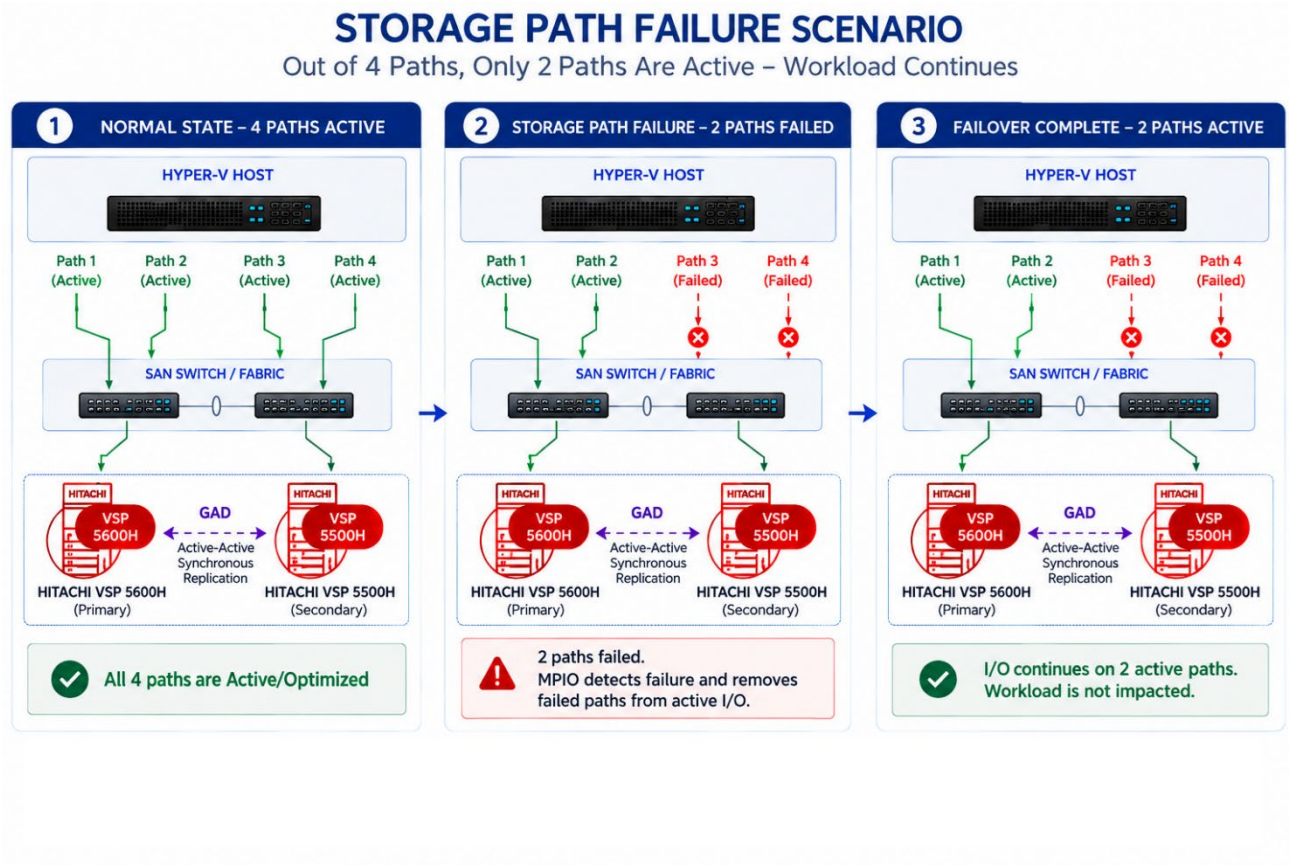

Results and Outcome
Microsoft Multipath I/O (MPIO) automatically redirected storage traffic through the remaining available paths. Hyper-V virtual machines retained uninterrupted access to GAD-protected storage volumes, SQL Server continued operating normally, and no sustained application interruption was observed.

Metric	Observed Result
Recovery Time	Less than 1 second
Storage Availability	Maintained
SQL Server Availability	Maintained
Application Availability	Maintained
Data Integrity	Preserved

Test Case 02 – Network Failure Recovery

Objective
Validate application continuity during a network communication failure.

Validation activity
A production network connection was intentionally disabled while application services remained operational.

The following figure shows one network connection that was intentionally disabled to simulate a communication failure. Redundant network paths maintained cluster communications, storage access, and application connectivity throughout the event.

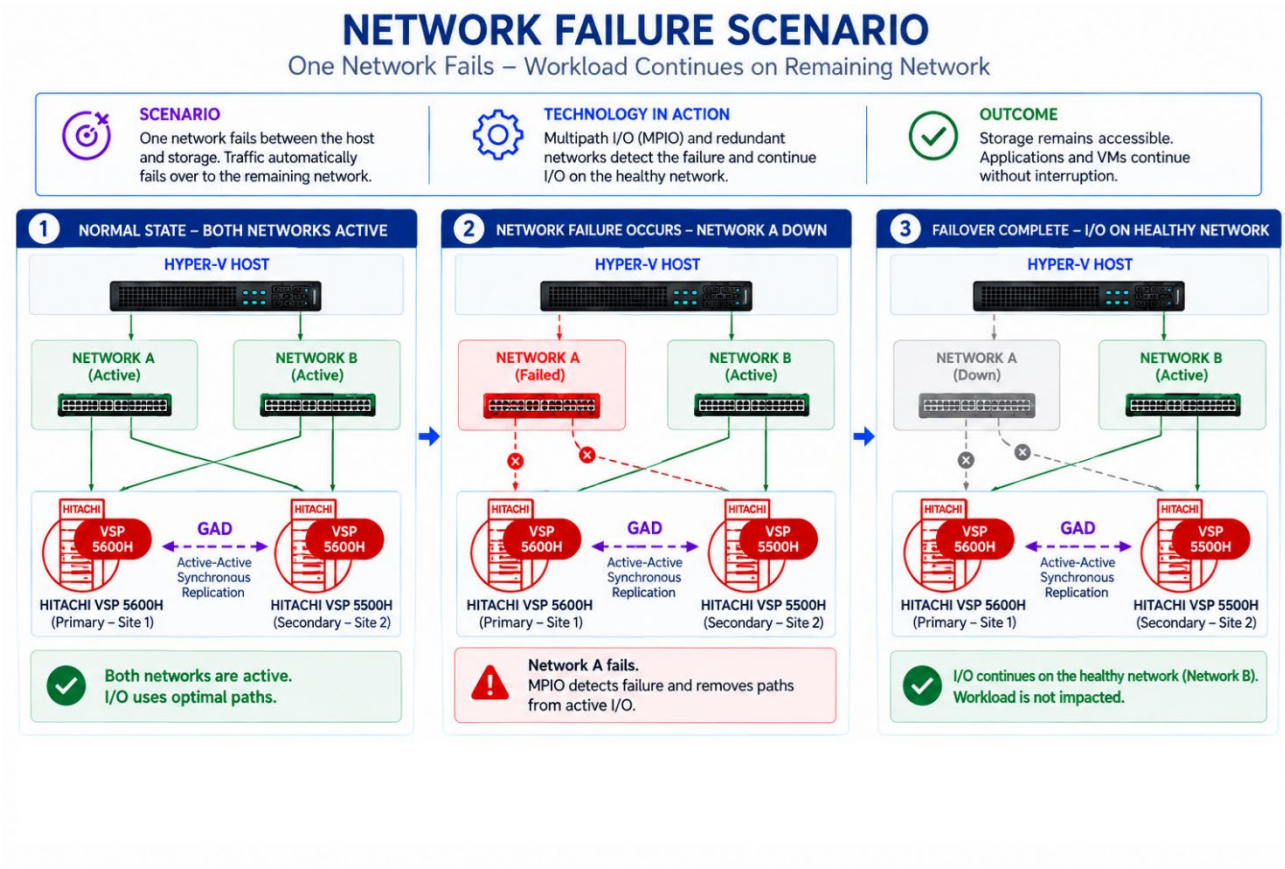

Results and outcome
Cluster communications remained healthy through available redundant paths, storage access remained unaffected, and SQL Server continued servicing application requests without interruption. The failure remained isolated to the affected network component.

Metric	Observed Result
Recovery Time	Less than 1 second
Cluster Health	Healthy
SQL Server Availability	Maintained
Application Availability	Maintained
Data Integrity	Preserved

Test Case 03 – SQL Server Always On AG Planned Failover (Intra-Site)

Objective
Validate a planned SQL Server Availability Group failover within the primary site.

Validation activity
A manual Availability Group failover was initiated between synchronized SQL replicas within the same site.

The following figure shows the primary SQL replica that was manually failed over to a synchronized secondary replica within the same site. The Availability Group Listener automatically redirected client connections without requiring application configuration changes.

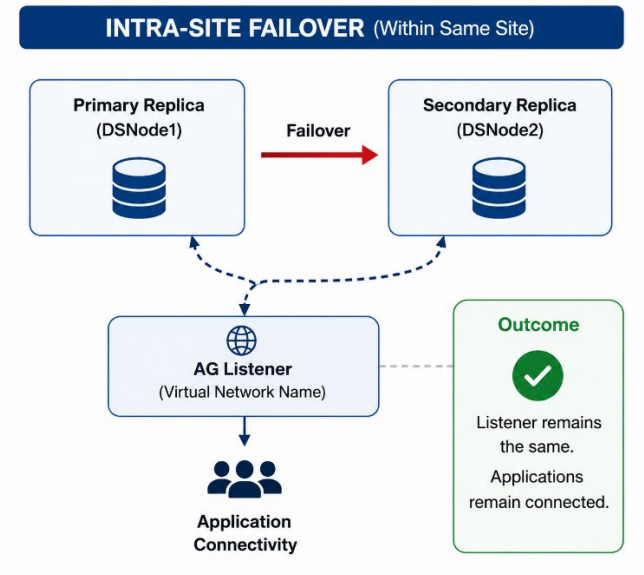

Results and outcome
The Availability Group Listener automatically redirected client connections to the new primary replica without requiring application configuration changes. Database synchronization remained healthy throughout the operation.

Metric	Observed Result
Test Runs	3
Recovery Time	Approximately 5-6 seconds
Listener Redirection	Automatic
Database Synchronization	Maintained
Application Availability	Maintained

Test Case 04 – SQL Server Always On AG Planned Failover (Inter-Site)

Objective
Validate database service continuity during a planned failover between sites.

Validation activity
A manual Availability Group failover was initiated between SQL replicas residing in separate sites.

The following figure shows that the SQL Server primary role was intentionally transferred between sites to validate business continuity. The Availability Group Listener redirected client connections while GAD continued providing access to synchronized storage resources across both locations.

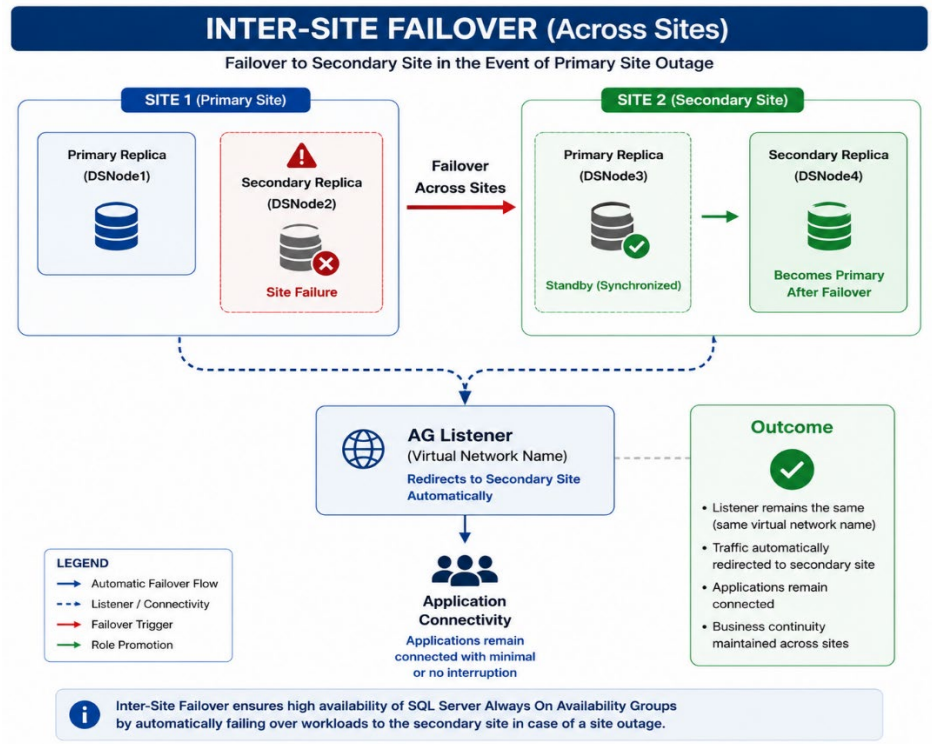

Results and outcome
SQL Server successfully transferred the primary role to the secondary site. The Availability Group Listener automatically redirected client connections while GAD continued providing synchronized storage services across both locations.

Metric	Observed Result
Test Runs	3
Recovery Time	Approximately 7-9 seconds
Listener Redirection	Automatic
Storage Availability	Maintained through Hitachi GAD
Application Availability	Maintained

Test Case 05 – SQL Primary Replica Forced Reboot
Objective
Validate automatic database recovery during an unexpected SQL Server failure.

Validation activity
The active SQL Server node hosting the primary replica was forcefully rebooted to simulate an unplanned infrastructure event.

The following figure shows that the active SQL Server node hosting the primary replica was forcefully rebooted to simulate an unexpected infrastructure failure. Windows Server Failover Clustering detected the outage and SQL Server Always On Availability Groups automatically promoted the synchronized secondary replica.

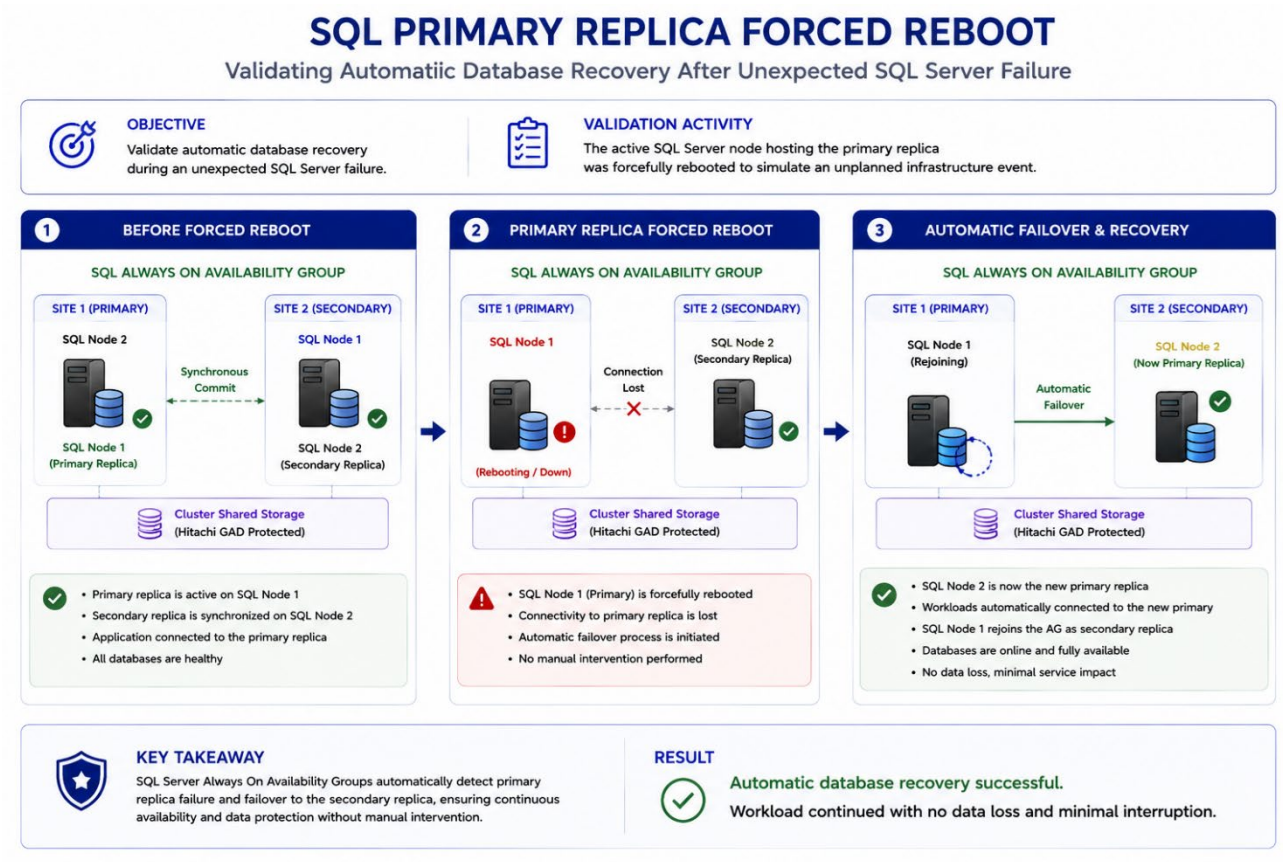

Results and outcome
Windows Server Failover Clustering detected the failure and SQL Server Always On Availability Groups automatically promoted the synchronized secondary replica. Application transaction processing resumed without manual intervention.

Metric	Observed Result
Test Runs	3
Recovery Time	Approximately 10-11 seconds
Failover Method	Automatic
Listener Redirection	Automatic
Data Integrity	Preserved
Application Availability	Restored Successfully

Test Case 06 – Hyper-V Live Migration

Objective
Validate workload mobility while maintaining continuous access to GAD-protected storage resources.

Validation activity
Hyper-V Live Migration was performed both within the primary site and across sites while application workloads remained active.

The following figure shows how the workload was migrated between Hyper-V cluster nodes while remaining online. Throughout the migration process, the virtual machine maintained uninterrupted access to GAD-protected storage and no transaction interruptions were observed.

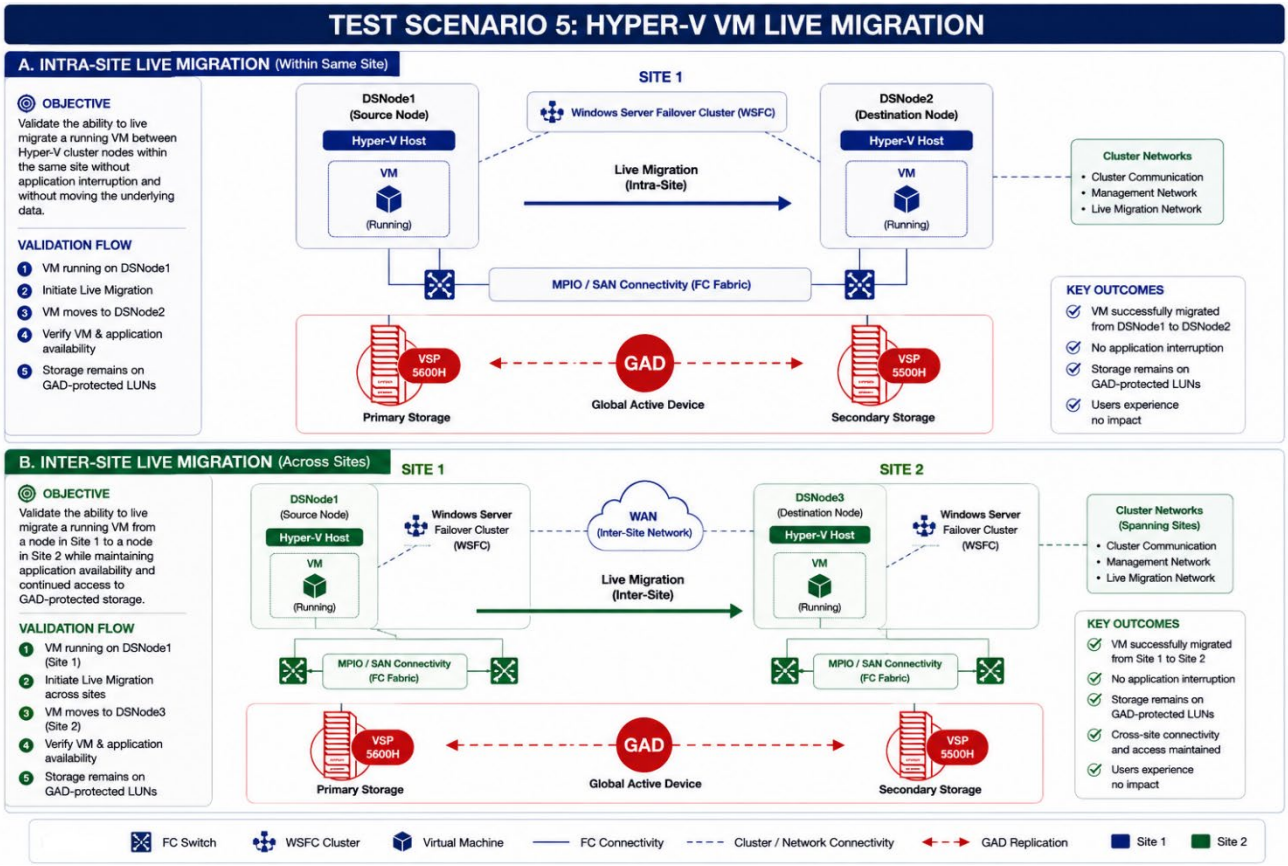

Results and outcome
Virtual machines successfully migrated between cluster nodes while maintaining access to GAD-protected storage volumes. No failed application transactions were observed during either migration scenario.

Metric	Observed Result
Intra-Site Migration	Successful
Inter-Site Migration	Successful
Storage Availability	Maintained
Application Availability	Maintained
Transaction Interruption	None Observed

Overall validation outcome

The validation results demonstrate the effectiveness of the layered protection strategy implemented throughout the solution. While individual technologies addressed specific failure domains, GAD consistently provided the continuously available storage foundation that enabled Hyper-V, Windows Server Failover Clustering, SQL Server Always On Availability Groups, and the supporting network infrastructure to sustain planned maintenance activities and infrastructure disruptions.

Across all test cases, the architecture maintained application availability, preserved data integrity, and demonstrated measurable recovery behavior across storage, network, compute, and database layers. These outcomes validate the value of combining Hitachi GAD with Microsoft high-availability technologies to create a coordinated enterprise availability platform capable of supporting mission-critical workloads.

Note : Observed recovery times reflect the results recorded within the validation environment and may vary depending on workload characteristics, infrastructure sizing, operating system settings, application behavior, and network conditions.

Business Value

The reference architecture demonstrates that high availability can be achieved through layered protection rather than dependence on a single technology. In the tested environment, storage-path and network failures were absorbed in less than one second with no sustained application interruption observed, while SQL Server Always On planned failover completed in 5–6 seconds within the site and 7–9 seconds across sites. A forced restart of the primary SQL replica resulted in application recovery within 10–11 seconds, while Hyper-V Live Migration produced no observed transaction interruption in the tested intra-site and inter-site scenarios.

At the foundation of these outcomes is global-active device (GAD), which provides the resilient storage layer across the Hitachi VSP Storage systems. By maintaining protected storage availability beneath the virtualization and database layers, GAD allows failures such as storage-path disruptions to be absorbed without unnecessarily triggering a database failover or application restart. Combined with MPIO, Hyper-V, WSFC, and SQL Server Always On Availability Groups, the architecture provides in depth, operational flexibility, and measurable recovery behavior, giving organizations greater confidence in maintaining mission-critical applications across infrastructure and site-level events.

Conclusion

Modern enterprises require more than isolated high-availability technologies to support mission-critical applications. Delivering continuous application availability requires a coordinated architecture capable of protecting storage, networking, compute, clustering, and database services across multiple failure domains. This paper demonstrates how GAD serves as the resilient storage foundation of that architecture, providing synchronous active-active replication and continuous access to enterprise data across geographically separated Hitachi Virtual Storage Platform (VSP) systems.

Building upon this foundation, Microsoft Hyper-V, Windows Server Failover Clustering (WSFC), SQL Server Always On Availability Groups, and resilient network infrastructure extend availability throughout the remainder of the application stack. Together, these technologies form a layered protection model that enables workloads to remain available during planned maintenance activities, infrastructure failures, and site-level events while preserving data integrity and application continuity. While the validation environment utilized specific Hitachi VSP systems and Hitachi Integrated Solutions infrastructure, the architectural principles described in this paper are applicable across supported Hitachi VSP One Block platforms and Hitachi Global-Active Device configurations, subject to applicable interoperability, firmware, and support requirements.

The validation results confirmed the effectiveness of this approach across storage, network, compute, and database failure domains. By combining GAD with proven Microsoft high-availability technologies, organizations can establish a validated platform that helps reduce downtime risk, simplify operations, strengthen business continuity, and improve confidence in supporting mission-critical applications. At the center of this strategy, global-active device provides the continuously available storage foundation that enables the broader architecture to deliver end-to-end resilience.

Additional Resources

Hitachi Vantara Resources

- See *Global Active Device for VSP One Block* on the Product Documentation portal (docs.hitachivantara.com) for more information.
- **Hitachi Product Interoperability Matrix**
  https://compatibility.hitachivantara.com/products/interop-matrix
- See *Virtual Storage Platform One Block* on the Hitachi Vantara Product Documentation portal (docs.hitachivantara.com) for more information.

Microsoft Resources

- **Microsoft Hyper-V Documentation**
  https://learn.microsoft.com/windows-server/virtualization/hyper-v/
- **Windows Server Failover Clustering Documentation**
  https://learn.microsoft.com/windows-server/failover-clustering/
- **Cluster Shared Volumes (CSV) Documentation**
  https://learn.microsoft.com/windows-server/failover-clustering/failover-cluster-csvs
- **SQL Server Always On Availability Groups Documentation**
  https://learn.microsoft.com/sql/database-engine/availability-groups/windows/
- **Availability Group Listener Documentation**
  https://learn.microsoft.com/sql/database-engine/availability-groups/windows/listeners-client-connectivity-application-failover

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact