HITACHI

Hitachi Vantara

Hitachi iQ MGX Solution with Hammerspace and Hitachi Virtual Storage Platform One Block

Reference Architecture Guide

MK-SL-455-00
April 2026

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 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!

Changes	Date
Initial release.	April 2026

Contents

Reference Architecture Guide

Purpose

In today's fast-evolving digital landscape, businesses are racing to harness the power of artificial intelligence (AI) to drive innovation, streamline operations, and stay ahead of the competition. However, building an AI-ready infrastructure that delivers optimal performance at the right cost can be a complex puzzle. By seamlessly integrating cutting-edge hardware, scalable software, and tailored services, Hitachi Vantara empowers customers to tackle their AI projects with confidence, efficiency, and precision.

Hitachi iQ bridges the gap between ambition and execution, delivering AI-ready solutions that fit both the job and the budget. Hitachi iQ allows organizations to automate, expedite, and streamline their business through intelligent, performant, scalable and flexible AI infrastructure and solutions.

Whether you are looking for industry-specific AI solutions or just starting to identify general purpose capabilities, Hitachi iQ has the power to automate your business processes and improve your customer experience.

Unlike other approaches on the market, Hitachi iQ goes beyond the basics of storage and infrastructure by layering industry-specific AI outcomes for industries such as finance, energy, transportation, manufacturing, and more. These solutions provide relevance and simplification to customers by accelerating the development and adoption of AI solutions and services into their ecosystem.

From entry level to enterprise capacity, Hitachi iQ is a versatile solution for AI needs. This reference architecture focuses on the Hitachi iQ MGX solution with Hammerspace and VSP One Block storage.

The intended audience of this document is:

- Data scientists and data engineers
- AI developers and Architects
- Data analysts
- Storage administrators
- System administrators
- IT professionals

This technical paper assumes that you are already familiar with the following:

- GPU-based server products
- NFS storage concepts
- Common IT storage practices

- Kubernetes Orchestration Platform
- General networking concepts

Solution overview

Hitachi Vantara has integrated Hammerspace software with the VSP One storage platform as part of the Hitachi iQ program. This integration expands Hitachi iQ's ability to handle diverse data management requirements for AI dataset creation, processing, governance, and protection. Hammerspace's data orchestration capabilities are crucial for optimizing data for AI workflows and pipelines, especially in large enterprises dealing with unstructured data across siloed environments.

The Hitachi iQ MGX with Hammerspace solution targets entry level to midrange AI/ML and GenAI workloads, utilizing a compact 2U PCIe GPU platform with the flexibility of the MGX modular architecture. Powered by NVIDIA RTX 6000 GPUs and Hammerspace’s global data orchestration, the solution delivers scalable compute and seamless access to distributed data across sites and clouds through a unified namespace, backed by a strategic Hitachi Vantara and Hammerspace partnership.

The Hitachi iQ solution architecture includes all the components necessary to build a fully functioning and optimized AI infrastructure for both entry, midrange, and high-end AI and Predictive Analytics use cases. This includes the GPU Servers, high-speed reliable storage, and high-speed networking components and well as a full management stack. Hitachi iQ with Hammerspace provides flexible options to tailor price and performance to each customer's unique workload or use case needs.

The solution described in this reference architecture offers significant benefits to businesses. The following are the primary benefits:

- Scalability
  - Compute Scalability
    - Early-stage AI projects often rely on a few high performance PCIe GPUs (NVIDIA RTX PRO 6000 Blackwell) to balance high VRAM capacity with cost-efficiency.
    - As models grow or user demand increases, the MGX Modular Architecture allows for seamless, incremental GPU expansion or high availability deployments within smaller, dense failover domains.
    - Customers gain long-term infrastructure protection through the MGX platform's forward compatibility with next-generation PCIe accelerators and high bandwidth interconnects.
  - Storage Scalability
    - AI projects typically start with lower-capacity storage and scale up based on performance and data volume needs.
    - Customers expect a modular upgrade path to higher-performance, higher-capacity storage systems.
    - While Virtual Storage Platform One File (VSP One File) may be sufficient for lower end workloads, the Hammerspace portfolio is the focal point for scalable and intelligent data orchestration in this reference architecture.
- Affordability
  - Affordable Compute
    - Customers look for cost-effective GPU solutions that support mixed workloads and scale with demand.
    - PCIe-based GPUs offer a balance of performance and affordability.
  - Affordable Storage
    - Customers aim to generate value from existing datasets before scaling to premium solutions.
    - Storage investments grow in tandem with proven AI outcomes, driving demand for affordable, high-performance storage.
  - Affordable Networking
    - Initial AI phases leverage existing Ethernet infrastructure to reduce costs.
    - As performance needs grow, InfiniBand is introduced for high-speed, low-latency networking.

- Consolidated Workloads and Enterprise Support
  - Consolidated Infrastructure
    - Customers prefer unified infrastructure built on standardized servers to simplify operations and maximize resource utilization across projects.
    - Mixed workload environments—language, graphics, inferencing, model tuning—can all operate on a shared infrastructure.
  - Unified Enterprise Support
    - A key barrier to adoption is the complexity of supporting diverse technologies and open-source software.
    - The Hitachi iQ platform delivers unified support across compute, storage, and networking components—accelerating issue resolution and minimizing operational overhead.
    - Certified by leading partners like NVIDIA, and supporting platforms such as RHEL and container orchestration tools, Hitachi provides SLA-bound global support, simplifying the GenAI/AI journey.

Business overview

With the Hitachi iQ portfolio, customers will benefit from flexible infrastructure that has been specifically engineered to provide the highest performing AI platform, while simplifying the enterprise AI journey, leading to improved time to value, better economics at scale, and the flexibility required to scale to meet the demands of the enterprise AI workload, regardless of the size. All while being backed by Hitachi Vantara's award-winning support and services organizations.

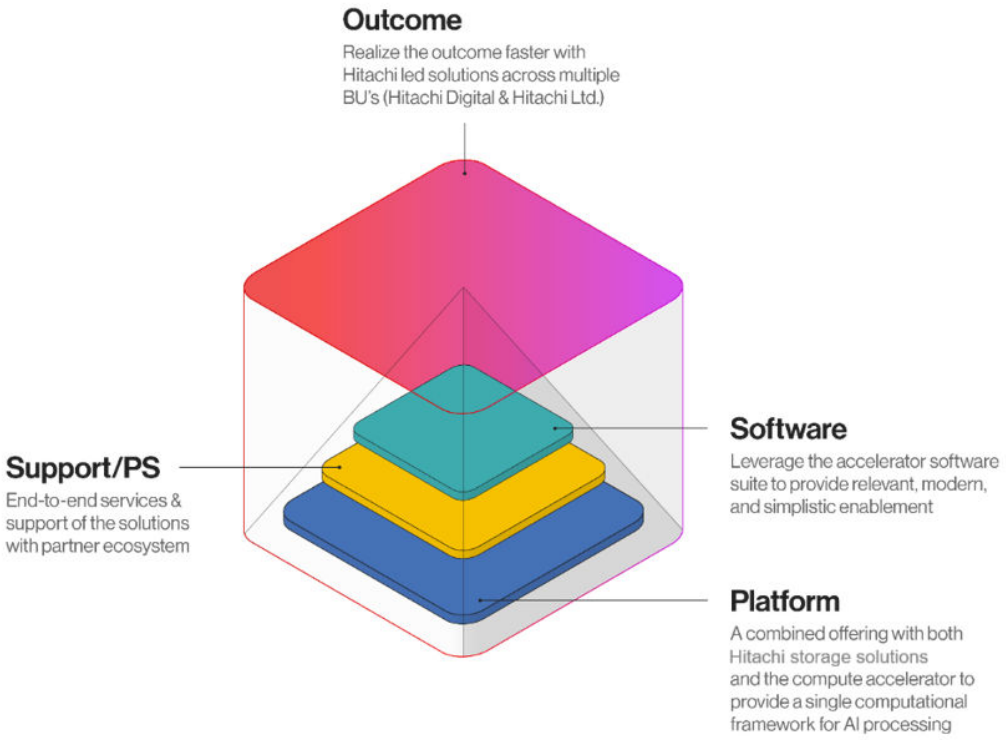

This reference architecture demonstrates a reference configuration of the Hitachi iQ MGX solution with Hammerspace and VSP One Block Storage.

Key components and technologies

Hammerspace and Hyperscale NAS

Hammerspace is a software-defined data orchestration and storage solution that provides unified file access via a high-performance Parallel Global File System that can scale out, meaning it can grow in capacity and performance by adding more storage nodes. This is essential for the large datasets and high-performance requirements of AI/GenAI workloads. It is designed to work seamlessly with Hitachi Vantara storage hardware, specifically VSP One Block.

Hammerspace software creates a Global Data Environment that may span multiple sites and multiple cloud environments.

The Hammerspace solution consists of four primary elements:

- Anvil Nodes: They are the brains of the operation, controlling the organization of and access to the data being stored, and tracking all file system metadata. No file I/O is managed by Anvil.
- DSX Nodes: Manages I/O operations, including data orchestration services, and are designed to scale out to accommodate any performance and capacity requirements.
- Storage: Hammerspace consumes storage capacity for the filesystem through external NFSv3 exports, external S3 Object Buckets, or Internal Block storage connected to the DSX systems. They are the devices that are responsible for providing storage for all data under Hyperscale NAS Management.
- Clients: They are the consumers of the storage services made available by the Hyperscale NAS system. Typical NAS clients are individual workstations or application servers or clusters.

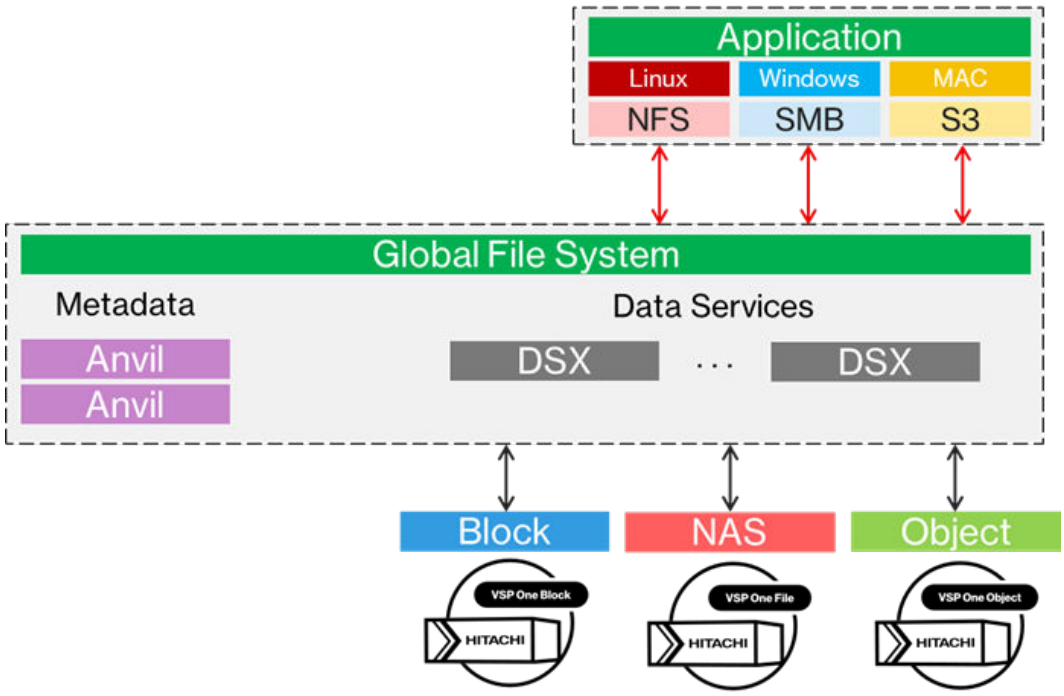

The diagram shows a layered architecture where applications on the client (in Linux, Windows, and macOS) use standard protocols (NFS, SMB, S3) to access a Hammerspace Global File System. This Global File System uses separate metadata (managed by Anvil servers) from data services (managed by DSX servers), and it integrates with Hitachi storage solutions: VSP One Block (block storage), VSP One File (NAS), and VSP One Object (object storage). It shows how applications interact with a unified file system that spans different types of Hitachi Vantara VSP storage systems.

Hyperscale NAS is a shared storage architecture. It has two defining characteristics:

- High performance typical of specialized parallel file systems, delivered via standard pNFS with the Flexible Files layout type (Flex Files). Separation of the metadata and data paths is a key attribute of this approach.
- Ease-of-use and features typical of enterprise scale-out NAS appliances, such as snapshots, replication, auditing, etc.

Hitachi VSP One Block

Virtual Storage Platform One (VSP One) Block offers hassle-free, enterprise-class block storage with certified sustainability for mission-critical applications.

With VSP One Block, you have the flexibility to choose the optimal solution to meet the specific demands of your workloads.

These are some of the key features:

- COMPACT: All NVMe storage platforms with the ability to support up to 3.6 PBe in just a 2U minimizing rack space, with reduced power and cooling costs.
- SIMPLE: At its core, embedded management with ClearSight, Dynamic Drive Protection always on compression and deduplication (ADR).
- SUSTAINABLE: Sustainability dashboard reports energy usage and CO2 impact, ‘always on compression’ with Advanced Data Reduction (ADR), bezels now contain 40% recycled materials, and intelligent power consumption with automatic ECO mode during low utilization periods.

For specifications see virtual-storage-platform-one-block-matrix-specifications.pdf.

Hitachi MGX compute node

The Hitachi iQ MGX compute node includes SYS-221GE-FNR2 systems that are DP Intel 2U PCIe GPU systems with up to 4 double-width PCIe GPUs cards.

Key Features:

- Dual 5th Gen Intel® Xeon® Scalable Gold 6500 series processors, up to 250W TDP
- Support for up to 4 double-width PCIe GPU accelerator cards. The system can support up to 4 NVIDIA RTX PRO™ 6000 or 4500 GPUs.
- Up to 32 DIMMs supporting up to 4TB DDR5-5600 in 1DPC or 8TB DDR5-4400 in 2DPC.
- Up to 4 PCIe 5.0 x16 FHFL double width + 4 AOC slots (1xPCIe 5.0 x16 HHFL and 3 PCIe 5.0 x16 FHFL slots).
- Up to 8 front hot-swap E1.S NVMe drive bays.
- 4 Redundant 2000W Titanium Level power supplies

For additional details about the server see the Supermicro Documentation and Datasheet.

Cisco Nexus 9300-GX Series Switches

The Cisco Nexus® 9300-GX switches are the next generation of fixed Cisco Nexus 9000 Series Switches capable of supporting 400 Gigabit Ethernet (GE). With the increase in use cases for applications requiring Artificial Intelligence (AI) and Machine Learning (ML), the platform addresses the need for high-performance, power-efficient, compact switches in the networking infrastructure. These switches are designed to support 100G and 400G fabrics.

Cisco Nexus 93600CD-GX Switch used in this solution for storage data network is a 1RU switch that supports 12 Tbps of bandwidth and 4.0 bpps across 28 fixed 40/100G QSFP-28 ports and 8 fixed 10/25/40/50/100/200/400G QSFP-DD ports.

Kubernetes

Kubernetes is an open-source container orchestration platform for deployment automation, scaling, and management of containerized applications.

NVIDIA GPU Operator

The NVIDIA GPU Operator uses the operator framework within Kubernetes to automate the management of all NVIDIA software components needed to provision NVIDIA GPUs. These components include the NVIDIA drivers (to enable CUDA), Kubernetes device plugin for GPUs, the NVIDIA Container Runtime, automatic node labelling, NVIDIA Data Center GPU Manager-based monitoring, and more.

NVIDIA Network Operator

An analog to the NVIDIA GPU Operator, the NVIDIA Network Operator simplifies scale-out network design for Kubernetes by automating aspects of network deployment and configuration that would otherwise require manual work. It loads the required drivers, libraries, device plugins, and CNIs on any cluster node with an NVIDIA network interface. Paired with the NVIDIA GPU Operator, the Network Operator enables GPUDirect RDMA, a key technology that accelerates cloud-native AI workloads by orders of magnitude. The NVIDIA Network Operator uses Kubernetes CRD and the Operator Framework to provision the host software needed for enabling accelerated networking.

NVIDIA AI Enterprise

NVIDIA AI Enterprise is a comprehensive, cloud-native software platform designed to accelerate data science pipelines and streamline the development and deployment of production-grade co-pilots and generative AI applications. It offers easy-to-use microservices that deliver optimized model performance, coupled with enterprise-grade security, support, and stability.

This ensures a seamless transition from prototype to production for businesses that rely on AI to operate.

NVIDIA Base Command Manager

NVIDIA Base Command™ Manager, a part of NVIDIA AI Enterprise and NVIDIA Mission Control, provides rapid deployment and comprehensive management for diverse AI and high-performance computing clusters, whether at the edge, in data centers, or in multi and hybrid-cloud environments.

It automates provisioning and administration of clusters ranging in size from a couple of nodes to hundreds to thousands, supports NVIDIA GPU-accelerated and other systems.

NVIDIA® Magnum IO™ GPUDirect® Storage (GDS)

NVIDIA GPUDirect Storage enables a direct data path between local or remote storage, such as NVMe or NVMe over Fabric (NVMe-oF), and GPU memory.

It avoids extra copies through a bounce buffer in the CPU's memory, enabling a direct memory access (DMA) engine near the NIC or storage to move data on a direct path into or out of GPU memory — all without burdening the CPU.

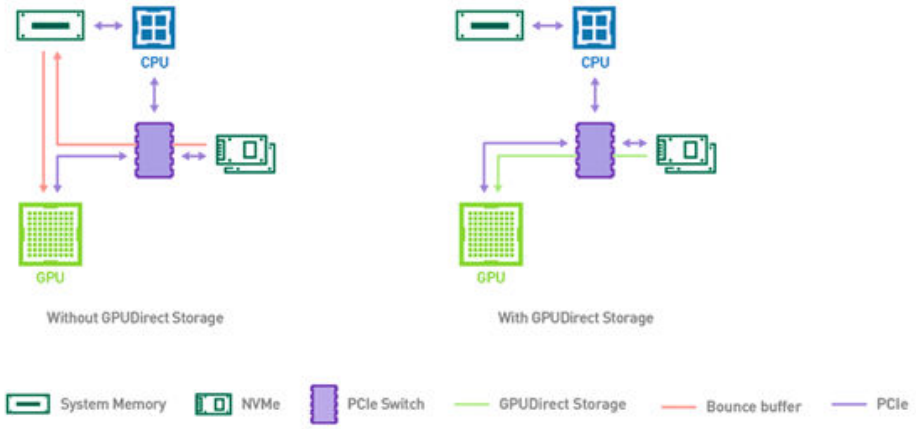

GPUDirect Storage offers the following capabilities:

- Direct Memory Access (DMA) Engine: DMA capabilities allow for direct communication between GPU memory and storage devices, bypassing the need to copy data through system memory. This reduces latency and enhances overall system performance.
- RDMA Capabilities: GPUDirect Storage utilizes remote direct memory access (RDMA) technology to efficiently access data stored in remote memory locations without involving the CPU, enabling data transfer between GPUs and storage devices across the network.
- NVIDIA Kernel Extensions and Drivers: These extensions and drivers facilitate the integration of GPUDirect Storage, enabling efficient data transfer paths between storage and GPU memory.
- Coherent Memory Access: GPUDirect Storage ensures consistent memory access and data integrity between GPUs and storage devices during data transfers.

Solution components

These are the key hardware and software components recommended and used during our testing in a lab environment. For the latest configuration and version details, consult your Hitachi technical representative to verify the most current information.

Hardware components

The following table lists the hardware components tested in this reference architecture.

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	Hitachi Advanced HA810 G3 SFF Server	Management servers CPU: 2 × Intel® Xeon® Gold 5418Y(185W, 24C, 2.0GHz), Total cores - 48 RAM: 16 × 32GB DIMMs (512 GB) Drive ▪ Internal SSD - 1 × 6.4 TB NVMe (optional) ▪ M.2 boot device - 2 × 480 GB NVMe SSD (NS204i-u Boot controller) NIC: ▪ 1 × 100 GEb 2p QSFP28 CX-6 card (Optional, use when NVIDIA Base Command Manager is used) ▪ 1 × 10 GbE 2-port BASE-T OCP3 Adapter (BCM57416)	SPV: 7.42 OS: RHEL 9.4	5
Hitachi MGX server	SYS-221GE-FNR2	GPU Compute Server CPU: ▪ 2 × Intel® Xeon® Gold 6548Y 32 cores RAM: ▪ 16 × 64 GB DDR5 DIMMs (1024 GB) Drives: ▪ 4 × 3.84 TB NVMe ▪ 2 × 960 GB M.2 NVMe	BMC FW: 01.05.08 BIOS: 2.7 OS: RHEL9.4	2

Vendor	Hardware	Detail Description	Version	Quantity
		NIC: ▪ 1 × AOC-STGS-I2T-P (Dual port 10 GbE) ▪ 1 × NVIDIA® ConnectX-7 NDR200 (Dual-port 200GbE, QSFP112, VPI) GPU (either one): ▪ 2 to 4 NVIDIA RTX PRO 6000 Server Edition ▪ 2 to 6 NVIDIA RTX PRO 4500 Blackwell Server Edition		
Hitachi Vantara	Hitachi Advanced HA815 G3 Server	Hammerspace Anvil Node CPU: ▪ CPU: 2 × AMD EPYC 9224 24-Core Processor ▪ 2.5 GHz ▪ Total Cores 48, Threads 96 (Intel GPU option also available) RAM: ▪ 1 TB with 16 × 64 GB DDR5-4800 Drives: ▪ NS204i-u – 480 GB NVMe SSD for boot drive ▪ 2 × 7.68 TB SSD NIC: ▪ INT E810 100 GbE 2p QSFP28 PCIe Adapter ▪ BCM 57416 10 GbE 2p BASE-T OCP3 Adapter ▪ BCM 57414 10/25 GbE 2p SFP28 OCP3 Adapter	SPV: 7.42 OS: v5.1.24-318	2
Hitachi Vantara	Hitachi Advanced Server HA815 G3 Server	Hammerspace DSX Node	SPV: 7.42 OS: v5.1.24-318	2

Vendor	Hardware	Detail Description	Version	Quantity
		CPU: ▪ CPU: 2 × AMD EPYC 9354 32-Core Processor ▪ 3.25 GHz ▪ Total Cores 64, Threads 128 (Intel GPU option also available) RAM: ▪ 512 GB with 16 × 32 GB DDR5-4800 Drives: ▪ NS204i-u – 480 GB NVMe SSD for boot drive ▪ 2 × 7.68 TB SSD NIC: ▪ INT E810 100 GbE 2p QSFP28 PCIe Adapter ▪ MLX MCX623106AS 100 GbE 2p QSFP56 PCIe Adapter ▪ BCM 57416 10 GbE 2p BASE-T OCP3 Adapter		
Hitachi Vantara	VSP One Block 28	▪ Drives: 24 * 3.8 TB ▪ CHB: 4 ▪ 8 × NVMe TCP 100G ▪ Cache: 1 TiB	SVOS 10.3.1	1
Cisco	Cisco 93600CD switch	▪ 28 fixed 40/100G QSFP-28 ports ▪ 8 fixed 10/25/40/50/100/200/400G QSFP-DD ports	BIOS: 05.43 NXOS: 10.4(3)	2
Cisco	Cisco 93108TC-FX3 switch	▪ 48 × 100M/1/10 Gbps BASE-T ports ▪ 6 × 40/100 Gbps QSFP28 ports	BIOS: 01.03 NXOS: 10.4(3)	2

Software components

The following table lists the key software components tested in this reference architecture.

NVIDIA Base Command Manager version	10 (RHEL 9.4)
OS	Management Nodes: RHEL 9.4 MGX Nodes: RHEL 9.4 Hammerspace Nodes: CentOS 8.5.2111
Kubernetes and Operator	Upstream Kubernetes (K8s) 1.30 NVIDIA GPU Operator 24.9.2 NVIDIA Network Operator 24.7.0
NVIDIA Drivers and Tools	Linux Kernel: 5.14.0-427.57.1 NVIDIA Open Driver Version: 575.51.03-1 CUDA Version: 12.9 NVIDIA GDS toolkit: 12.9 GDS release version: 1.14.0.30 NVIDIA_fs version: 2.25.6 libcufile version: 2.12
OFED	MLNX_OFED_LINUX-24.07-0.6.1.0
NVIDIA Driver	28.39.2048
Hammerspace Software	5.1.24-318
VSP One B28	SVOS 10.3.1
Cisco Nexus Switches	NXOS 10.4

Solution design

This is a detailed solution example of how the Hitachi iQ MGX solution with Hammerspace and VSP One Block Storage is configured.

The following is a reference rack elevation design for 16 × MGX compute nodes. This reference design offers a scalable and modular architecture to start as small as 2 nodes and grow up to 16 nodes. The rack layout can be adjusted to meet local data center requirements, such as maximum power per rack and rack layout between system, storage and network components to meet local needs for power and cooling distribution.

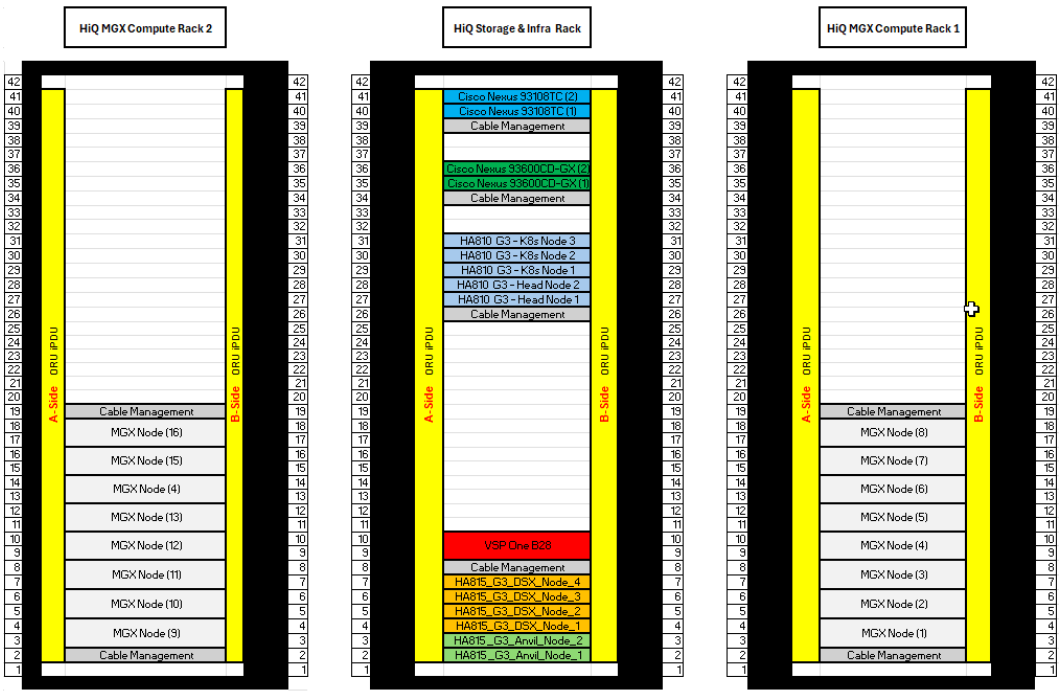

Physical architecture

The following is an abstract description of how MGX compute servers, Hammerspace, and VSP One Block storage are configured as part of Hitachi iQ MGX solution.

- HA810 G3 servers are used as management nodes.
  - 2 × NVIDIA Base Command Manager Head nodes.
  - 3 × K8s control plane nodes.
- MGX servers with GPUs are used as compute nodes to form K8S Cluster worker nodes.
- HA815 G3 servers are used as Hammerspace storage nodes.
  - 2 × Anvil nodes are used to configure highly available metadata server.
  - 2 to 4 DSX nodes are used to configure data nodes through which clients will access the data volumes and the global filesystem.
- 1 × VSP One Block 28 storage to allocate block volumes to DSX nodes.
- 2 × Cisco 93600CD switches are used for 100 GbE storage network that will be used for all IO communication between client and storage node.
- 2 × Cisco 93108TC-FX3 switches are used for 1/10 GbE in-band and OOB network.

The following illustration shows the architecture diagram:

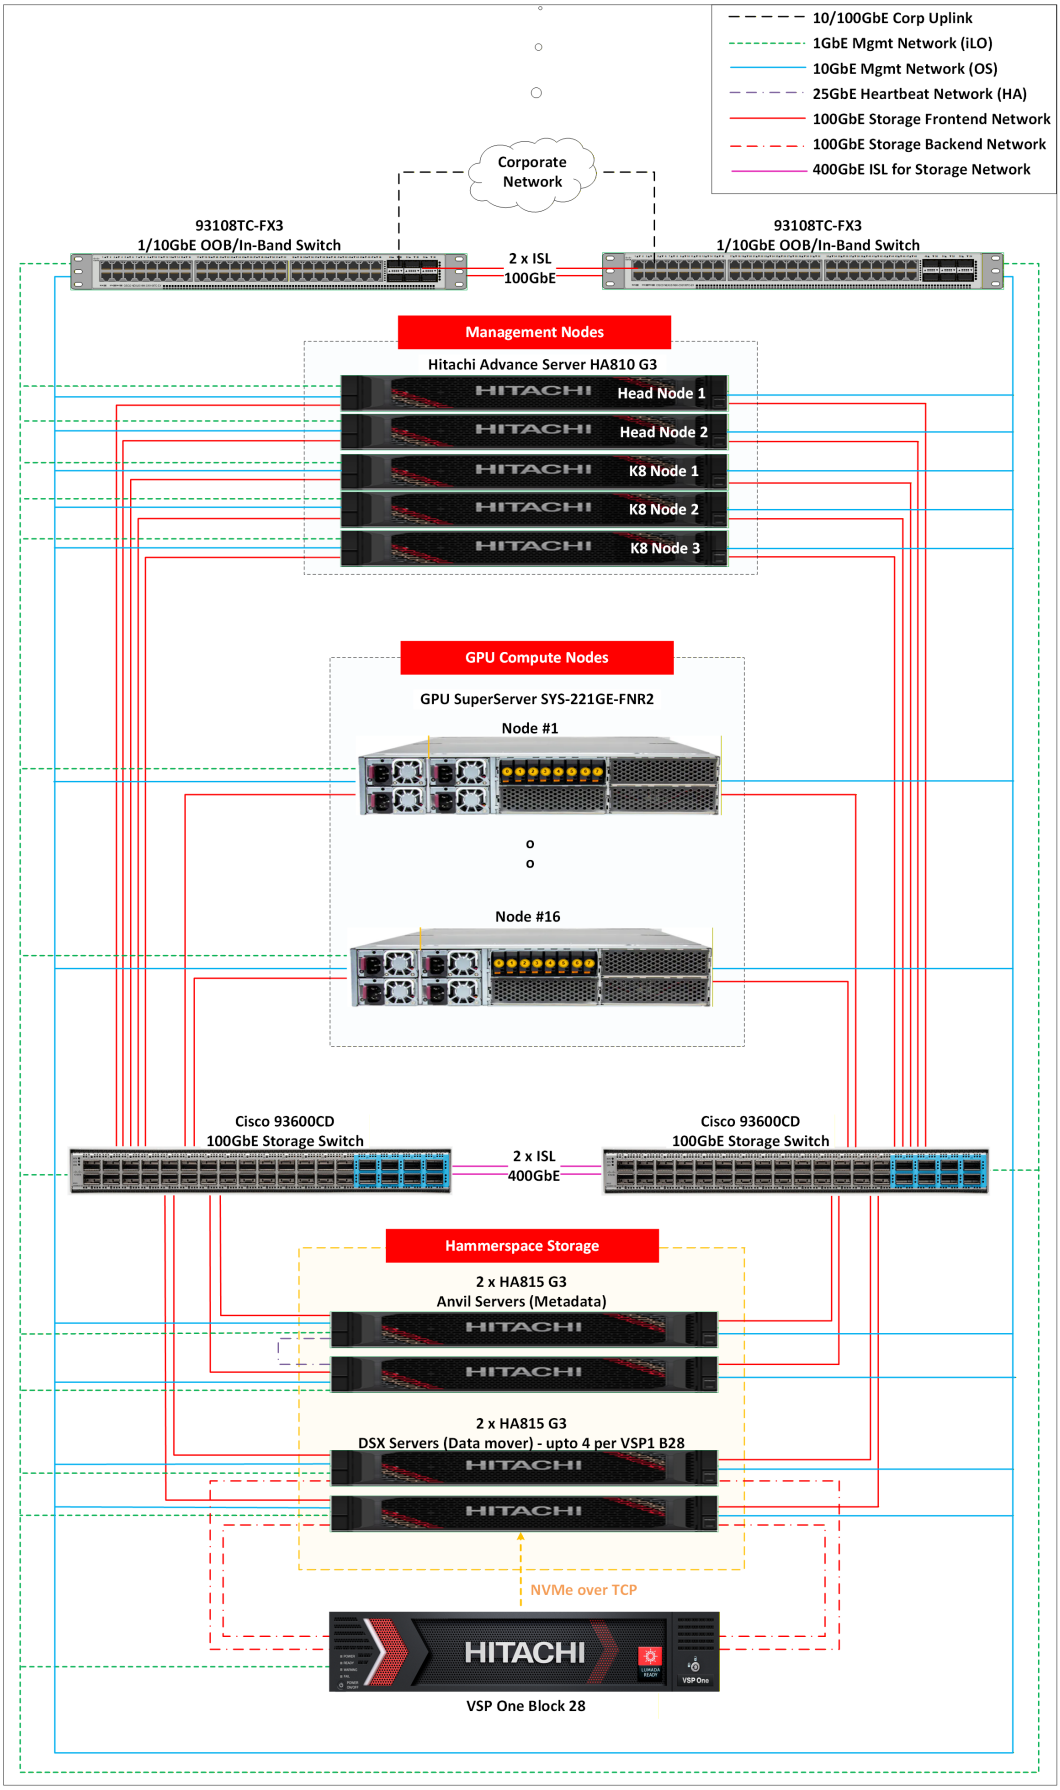
10/100GbE Corp Uplink
1GbE Mgmt Network (iLO)
10GbE Mgmt Network (OS)
25GbE Heartbeat Network (HA)
100GbE Storage Frontend Network
100GbE Storage Backend Network
400GbE ISL for Storage Network
Corporate Network
93108TC-FX3
1/10GbE OOB/In-Band Switch
2 x ISL
100GbE
93108TC-FX3
1/10GbE OOB/In-Band Switch
Management Nodes
Hitachi Advance Server HA810 G3
HITACHI Head Node 1
HITACHI Head Node 2
HITACHI K8 Node 1
HITACHI K8 Node 2
HITACHI K8 Node 3
GPU Compute Nodes
GPU SuperServer SYS-221GE-FNR2
Node #1
o
o
Node #16
Cisco 93600CD
100GbE Storage Switch
2 x ISL
400GbE
Cisco 93600CD
100GbE Storage Switch
Hammerspace Storage
2 x HA815 G3
Anvil Servers (Metadata)
HITACHI
HITACHI
2 x HA815 G3
DSX Servers (Data mover) - upto 4 per VSP1 B28
HITACHI
HITACHI
NVMe over TCP
HITACHI
VSP One Block 28

Note: In this architecture, Hammerspace DSX nodes have dedicated 100 GbE interfaces for clients and backend storage connectivity.

Network overview

On a high level four different networks are configured.

- Frontend Storage Network:
  - This is configured over a 100 GbE Ethernet switch.
  - A private L2 Ethernet network is set up for connecting clients with Anvil and DSX nodes.
- Backend Storage Network:
  - Dedicated direct connections are made between DSX nodes and Backend VSP B28 storage.
  - 2 × 100 GbE links are connected from each DSX nodes to backend VSP1 B28 storage
  - A maximum of 4 DSX nodes can be connected to VSP1 B28 storage.
- Management Network:
  - This is configured over a 1/10 GbE Ethernet switch.
  - The network is used for OOB (iLO) and in-band (OS) management.
  - We can configure a single VLAN or different VLAN for OOB and in-band management as needed.
- HA Network:
  - This is configured directly between 2 Anvil nodes using a 25 GbE network.
  - This is a private L2 heartbeat network for Anvil HA configuration.

The following are the available ports on the MGX GPU server that will be used for storage fabric and in-band management network.

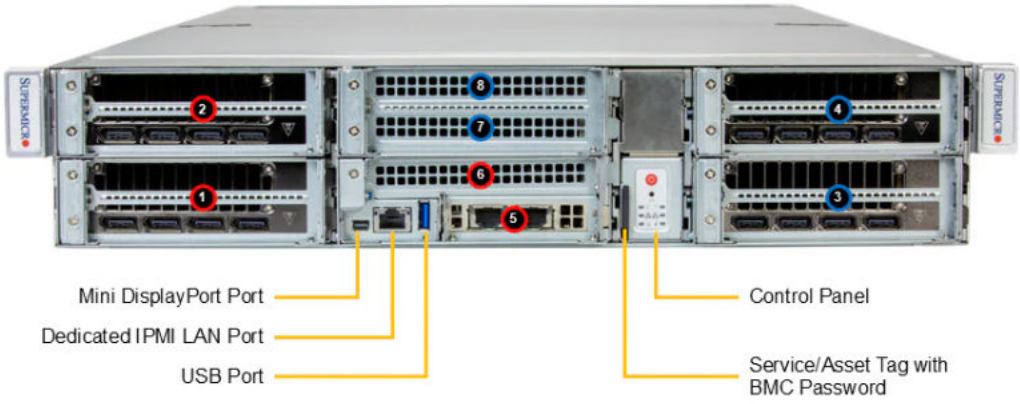

	Slot Description
1 2 3 4	PCIe 5.0 x16 FHFL Double-width GPU
5 6 7 8	PCIe 5.0 x16 FHFL Networking/DPU

CPU1 CPU2

- Slots 1 to 4 are reserved for GPU connections.
- Slot 7 is populated with dual port 200GbE CX-7 cards for connecting to Cisco 93600CD switches for storage connectivity.
- Slots 6 and 8 are kept free.
- For in-band management, a 2 × RJ45 10GbE interface provided by an Intel® X550 add-on NIC is used to connect to the Cisco 93108TC-FX3 switch as a bonded interface.
- For OOB management, a dedicated IPMI port will connect to the Cisco 93108TC-FX3 switch.

The following are the available ports on the Anvil/DSX servers that will be used for the storage fabric and in-band management network.

HA815 G3 DSX Storage Node I/O Slot Assignment			
	Slot Description	**Installed I/O Card**	**I/O Allocation**
1	PCIe 5.0 × 16 (FHFL) Slot 1	Intel E810-CQDA2 100 GbE 2P	Data Network Ethernet
2	PCIe 5.0 × 16 (FHFL) Slot 2	MLX MCX623106AS 100 GbE 2P	Direct VSP Connection
3	NS204i-u M.2 controller (Internal)	Boot Device	
4	OCP 3.0 Slot 22	Available for use	
5	OCP 3.0 Slot 21	BCM 57416 10 GbE 2P OCP3	In-Band Management
6	Dedicated iLO Management Port	iLO	Out-Of-Band Management

HA815 G3 Anvil Storage Node I/O Slot Assignment			
	Slot Description	**Installed I/O Card**	**I/O Allocation**
1	PCIe 5.0 × 16 (FHFL) Slot 1	Intel E810-CQDA2 100 GbE 2P	Data Network Ethernet
2	PCIe 5.0 × 16 (FHFL) Slot 2	Available for use	
3	NS204i-u M.2 controller (Internal)	Boot Device	
4	OCP 3.0 Slot 22	BCM 57414 10/25 GbE 2P OCP3	HA
5	OCP 3.0 Slot 21	BCM 57416 10 GbE 2P OCP3	In-Band Management
6	Dedicated iLO Management Port	iLO	Out-Of-Band Management

Note: Slots 1, 3, and 4 are NUMA0 (CPU 1)

Slots 2 and 5 are NUMA 1 (CPU 2)

Conclusion

Hitachi iQ offers comprehensive solutions with a scalable architecture and the flexibility to address the complexities of modern AI-driven workloads.

Hammerspace Software with Hitachi VSP One Block storage is capable of scaling linearly based on the number of storage “blocks” that are added to the solution allowing organizations to start with a smaller footprint and grow within the scalable architecture of Hitachi iQ.

This reference architecture serves as a foundational guide, providing the necessary tools and best practices to leverage the combined strengths of Hammerspace Software and Hitachi VSP One Block Storage, ultimately driving success in the era of data-driven decision-making.

Key benefits of this reference architecture include:

- Enhanced performance and scalability for demanding ML/AI workloads.
- Reliable and secure data management, ensuring data integrity and compliance.
- Simplified deployment and management, reducing operational complexities.
- Comprehensive support for advanced AI workflows, driving innovation and insights.

By adopting this solution, organizations can effectively harness the power of AI, transforming data into actionable insights and achieving strategic objectives with greater efficiency.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact