Hitachi iQ with Hammerspace and Hitachi Virtual Storage Platform One

Reference Architecture Guide

MK-SL-326-00
April 2025

Feedback

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

Thank you!

Revision history

Changes	Date
Initial release	April 2025

Reference Architecture Guide

Introduction

As part of the Hitachi iQ portfolio of data solutions, Hitachi Vantara has integrated Hammerspace software with the Virtual Storage Platform One data platform, adding new data orchestration capabilities to Hitachi iQ's AI-ready infrastructure offerings, ensuring seamless and transparent access to distributed data for AI workloads. This document specifically defines a reference architecture that combines Hammerspace software with Hitachi Virtual Storage Platform One Block storage.

Overview

Hyperscale NAS is the shared storage architecture used by Hammerspace. It has two defining characteristics:

- High performance typical of specialized parallel file systems, delivered using 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, and auditing.

Hitachi Vantara has entered into a strategic resell agreement with Hammerspace. Hitachi Vantara has integrated Hammerspace software with the VSP One storage platform, expanding Hitachi iQ's capabilities to address different data management requirements for dataset creation, processing, governance and protection.

Traditional storage systems struggle with data sprawl, governance, security, and integration with cloud and GPU clusters. Hammerspace's data orchestration capabilities significantly benefit large enterprises by optimizing the data for AI workflows and pipelines. It addresses the challenges of managing vast amounts of unstructured data across siloed storage environments, which are intensified by AI-driven workloads requiring speed, scale, and flexibility.

Hitachi iQ is an industry-optimized AI solution suite designed to empower organizations investing in AI/ML, GenAI, and other demanding GPU-driven workloads. By leveraging cutting-edge infrastructure and AI solutions, Hitachi iQ enables businesses to automate processes, accelerate time-to-insights, and unlock innovation, ensuring they stay ahead in the competitive landscape. As part of the Hitachi iQ portfolio of data solutions, Hitachi Vantara has integrated Hammerspace software with the VSP One storage platform adding new data orchestration capabilities to Hitachi iQ's AI-ready infrastructure offerings. This combined VSP One and Hammerspace solution ensures seamless and transparent access to distributed data for GenAI workloads.

Experience the unparalleled capabilities of Hitachi's technologies combined with Hammerspace's data orchestration. With Hitachi's robust infrastructure and Hammerspace's true global file system, you can seamlessly span multiple sites, clouds, and storage systems under a single global namespace. Effortlessly utilize data to support your Hitachi iQ outcomes.

Solution components

The following components are used in this solution.

Anvil Metadata Nodes are the brains of the operation, controlling the organization of and access to the data being stored, and tracking all file system metadata.

DSX Data Service Nodes are the workhorses of the Hammerspace platform, providing scale-out engines that connect data and storage to users and applications, performing all I/O operations, replication, data movement, etc.

Virtual Storage Platform One (VSP One) storage systems are the devices that are responsible for providing storage for all data under Hammerspace Management. Any device that supports NFSv3 or can directly provide storage to the DSX Nodes can be used as a storage system.

Clients are the consumers of the storage services made available by the Hammerspace system. Typical NAS clients are GPU servers or clusters, individual workstations, or application servers or clusters.

Protocols

The Hammerspace components communicate with clients using protocols defined by pNFS Flex Files. Clients mount shares and conduct other metadata activity by communicating with metadata servers using NFSv4.2. After they have obtained a layout from the metadata server, clients read and write data directly to the storage servers using NFSv3. Metadata servers communicate to storage servers using NFSv3 to manage access to data.

Hammerspace and Hyperscale NAS

Hammerspace software creates a global data environment that can span multiple sites and multiple cloud environments. For Hyperscale NAS deployments, a redundant pair of Hammerspace Anvil servers act as the metadata servers. Storage servers may be Hammerspace DSX Data Services Nodes or plain Linux NFSv3 servers. DSX nodes are used when data orchestration capability is desired, or there is a need to export the file system using NFSv3, SMB, S3, or Kubernetes CSI in addition to NFSv4.2.

Hitachi VSP One Block

VSP One Block offers worry-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 workload.

Select between Triple-Level Cell (TLC) and Quad-Level Cell (QLC) media depending on performance needs and overall cost considerations. VSP One Block with TLC flash is ideal for scenarios requiring superior performance and greater endurance, while QLC is advantageous for workloads that prioritize high capacity and cost-effective storage where data is infrequently modified.

- COMPACT: All NVMe storage platforms with the ability to support up to 3.6 PBe in just 2U minimizing rack space and reducing 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 & 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.

Deployment architecture

Hitachi Vantara HA810 G3 (Intel) and HA 815 G3 (AMD) servers are used for the Anvil and DSX servers while storage is provided by the Hitachi Vantara VSP One Block 28 NVMe storage platform. The same basic design can be adjusted for use by other VSP One Block platforms to meet the required performance and cost profile.

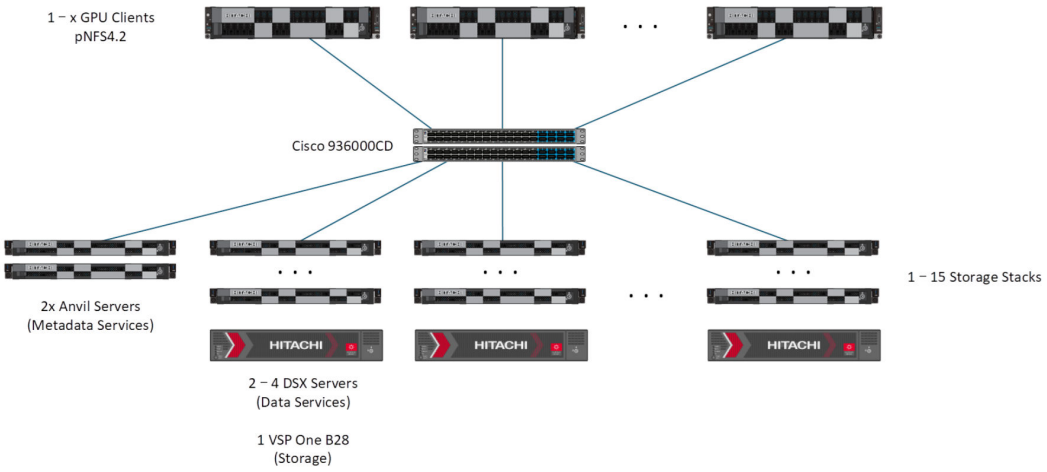

Server specifications

The following servers can be used in this solution.

AMD-based servers

Anvil	Server	Hitachi Advanced Server HA815 G3 1 RU, 2 AMD CPU, 8 SFF NVMe
	CPU	2 × AMD EPYC 9224, 24C, 2.25GHz, 200W (48 Cores Total)
	Memory	16 × 64 GB 2Rx4 PC5-4800B-R (1 TB Total)
	Boot Device	480 GB RI NVMe DS M.2 SSDs Mirrored
	Network 1	Broadcom 57414 10/25 GbE 2p SFP28
	Network 2	Mellanox CX6 100 GbE 2p QSFP56
	NVMe	2 × 6.4 TB NVMe SFF BC U.3 PM1735a (12 TB Useable)
DSX	Server	Hitachi Advanced Server HA815 G3 1 RU, 2 AMD CPU, 8 SFF NVMe
	CPU	2 × AMD EPYC 9334, 32C, 2.6GHz, 210W (64 Cores Total)
	Memory	16 × 32 GB 1Rx4 PC5-4800B-R (512 GB Total)
	Boot Device	480 GB RI NVMe DS M.2 SSDs Mirrored
	Network 1	Broadcom 57414 10/25 GbE 2p SFP28
	Network 2	Mellanox CX6 100 GbE 2p QSFP56
	Network 3	Intel E810 100 GbE 2p QSFP28

Intel-based servers

Anvil	Server	Hitachi Advanced Server HA810 G3 1 RU, 2 Intel CPU, 8 SFF NVMe
	CPU	2 × Intel Xeon-G 5418Y CPU, 24C, 2.0GH, 185W (48 Cores Total)
	Memory	16 × 64 GB 2Rx4 PC5-4800B-R (1 TB Total)
	Boot Device	480 GB RI NVMe DS M.2 SSDs Mirrored
	Network 1	Broadcom 57414 10/25 GbE 2p SFP28
	Network 2	Mellanox CX6 100 GbE 2p QSFP56
	NVMe	2 × 6.4 TB NVMe SFF BC U.3 PM1735a (12 TB Useable)
DSX	Server	Hitachi Advanced Server HA810 G3 1 RU, 2 Intel CPU, 8 SFF NVMe
	CPU	2 × Intel Xeon-G 6548Y+ CPU, 32C, 2.5GHz,250W (64 Cores Total)
	Memory	16 × 32 GB 1Rx4 PC5-4800B-R (512 GB Total)
	Boot Device	480 GB RI NVMe DS M.2 SSDs Mirrored
	Network 1	Broadcom 57414 10/25 GbE 2p SFP28
	Network 2	Mellanox CX6 100 GbE 2p QSFP56
	Network 3	Mellanox CX6 100 GbE 2p QSFP56

VSP One Block specification

The following table lists details of the VSP One Block specification.

	VSP One Block 26QLC	VPS One Block 28
Max Effective Capacity	1.6 PB/1.48 PiB	14.1 PB/12.8 PiB
Max Raw Capacity	720TB (24 × 30 TB)	4.32PB (72 × 60 TB)
Effective Capacity Guarantee Ratio	4:1	4:1
Maximum Number of Drives	24 NVMe	72 NVMe
Drive Options	30 TB NVMe QLC SSD	60 TB NVMe SSD 30 TB NVMe SSD 15 TB NVMe SSD 7.6 TB NVMe SSD 3.8 TB NVMe SSD 1.9 TB NVMe SSD
Expansion Drive Trays	-	2 (2U 24 SFF Drive Tray)

Deployment example

The following two deployment examples using VSP One Block 28 and VSP One Block 26QLC demonstrate ease of entry for the solution along with the ability to cost-effectively scale the solution based on performance and capacity needs.

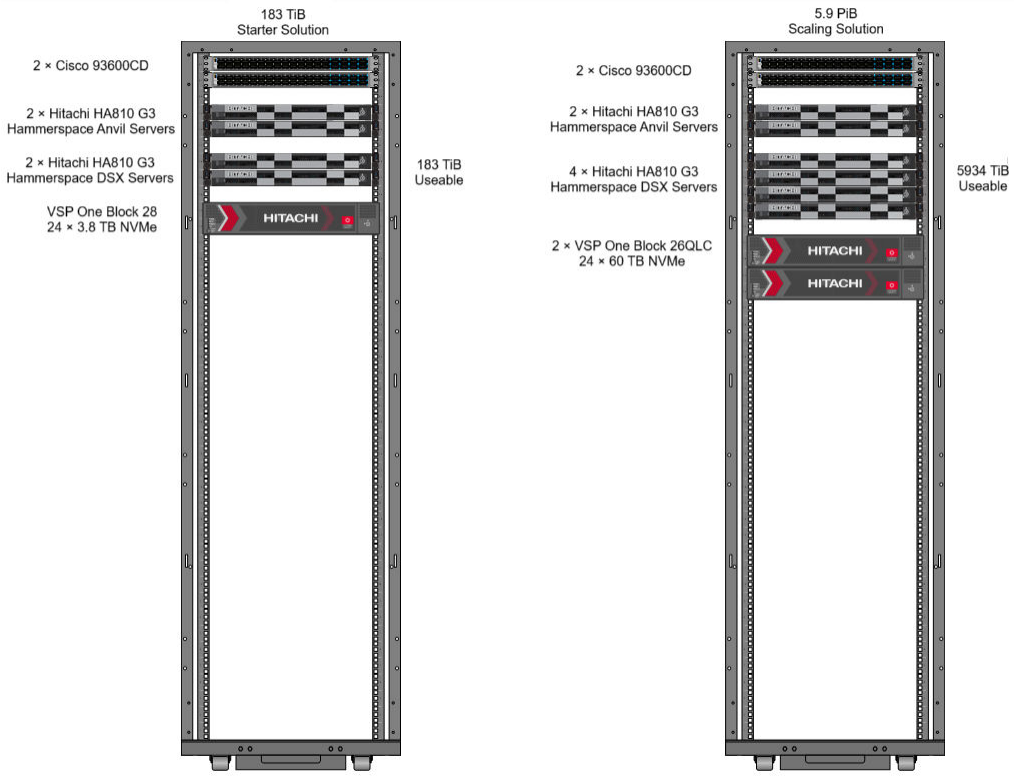

VSP One Block connectivity

To provide reliable and resilient storage to Hammerspace, storage is provided with Hitachi VSP One Block using direct connected 100 Gb TCP NVMe connections to DSX Nodes. VSP One Block 28 supports 8 × 100 Gb TCP NVMe connections which allow for up to 4 DSX Storage nodes to be connected without the need for consuming switch ports.

The following illustration provides a network connectivity example.

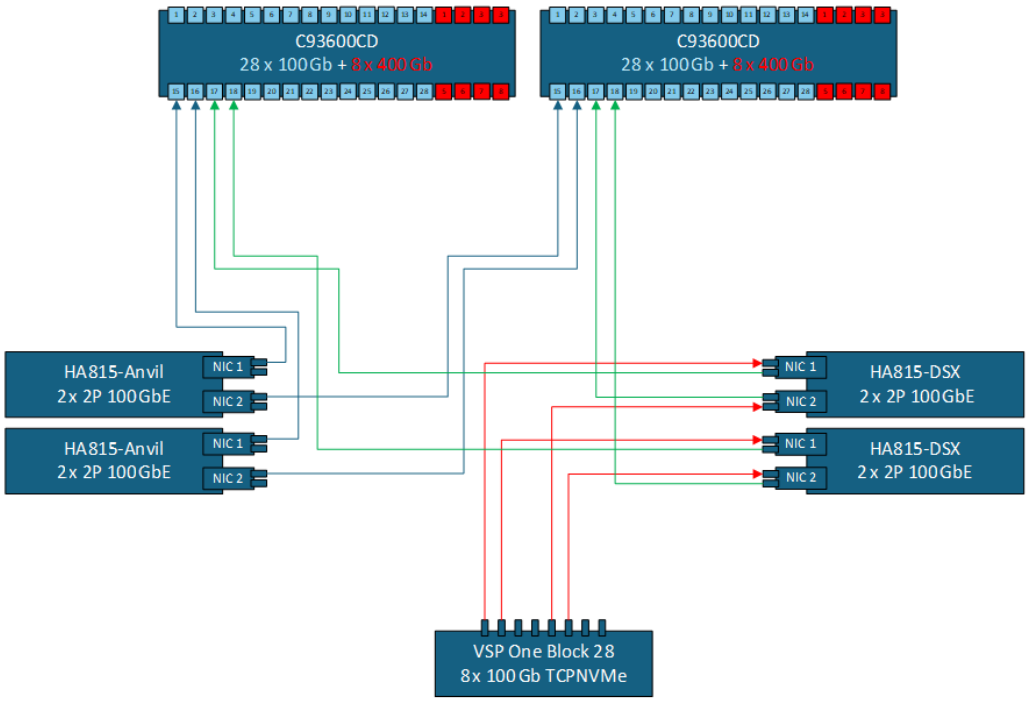

While the Cisco Nexus C936000CD was chosen for this example, any suitable switching environment can be used based on the network bandwidth and latency needs required by the application. Any networking technology that supports TCP/IP may be used. Where RDMA is desired, this includes InfiniBand (with IP over IB) and RDMA over Converged Ethernet v2 (RoCEv2).

Note: Hammerspace DSX servers require Mellanox ConnectX-4 or later NICs to use RDMA.

Capacity deployment and scaling

Capacity for Hammerspace with Hitachi storage can be provided from the VSP One Family of products: VSP One Block, VSP One File, and VSP One Object. For the design, VSP One Block 28 is shown, however any VSP One Block can use the same build patterns.

The Hammerspace with Hitachi VSP One Block solution can scale for both performance and capacity independently to meet the environment requirements. The solution design has been optimized with a maximum of 4 to 1 DSX Storage nodes to each VSP One Block storage system with no expansion drive trays. For additional capacity additional drive trays can be added, or if additional performance is needed, additional DSX nodes can be added to the namespace.

LUN sizing recommendations

Sizing LUNs for use in the Hammerspace solution is based on the VSP One Block controllers' capacity fully dedicated to the solution and only using the 24 internal drives of the storage system. If expansion drive trays are added, then LUN sizes will need to be updated to reflect the additional capacity added and be evenly distributed across the 24 provisioned LUNs.

LUN distribution across DSX nodes should be uniform to ensure proper capacity distribution. This ensures that Hammerspace is able to correctly distribute file operations evenly across all DSX storage nodes to ensure that no one node becomes a hotspot.

LUN Deployments	
LUNs Per Array	24
LUNs Per DSX Node	
2 Node DSX	12
3 Node DSX	8
4 Node DSX	6

		Capacity Per Node (TiB)		
Drive Size	**LUN Size (GiB)**	**2 Node**	**3 Node**	**4 Node**
1.9	3,372	39.52	26.34	19.76
3.8	6,745	79.04	52.70	39.52
7.6	13,490	158.09	105.39	79.04
15.3	26,870	314.88	209.92	157.44
30	53,780	630.23	420.16	315.12
60	107,600	1,260.94	840.63	630.47

Storage resiliency

To ensure maximum uptime for the solution, it is recommended to deploy all shares with an Objectives set including a minimum of availability-3-nines. This will ensure that two copies of the data are on different DSX volumes in the cluster. The solution uses VSP One Block for all storage volumes and with Hitachi Advanced Data Reduction (ADR). The second volume copy of the data is eliminated at the storage layer. This provides a solution that maximizes the potential performance of the DSX Nodes and provides enterprise resiliency in a space and cost-efficient manner.

An Objective rule of availability-3-nines will create 2 copies of the data on a 4 DSX node solution, allowing up to 2 nodes to fail while maintaining full uptime with degraded performance. If additional resilience is required, an Objective rule of availability-5-nines can be used instead allowing 3 copies of the data to be kept across 4 DSX nodes. While this provides additional resiliency with minimal impact on storage capacities because of ADR, availability-5-nines in the solution can impart a negative impact to write performance in some workloads.

The following illustration shows example objectives.

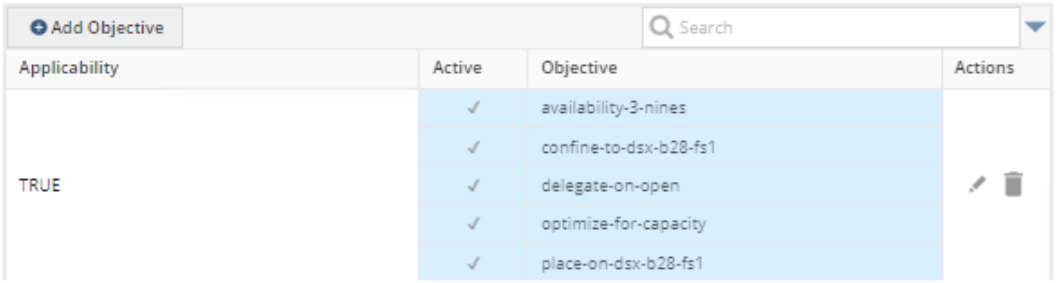

Volume and objective recommendations

To best leverage the storage flexibility that can be provided by Hitachi VSP One Block, create logical volume grouping objectives for each tier type that is added to the environment. For example, if a VSP One Block 28 with TLC is used alongside a VSP One Block 26QLC then it is recommended to create two rules that span across all volumes that are backed by LUNs for their respective controllers. This will allow easier management of the storage tier and ensure the applications are receiving the expected performance of the tier they are using.

With each tier type in a dedicated volume group Objectives can be used to ensure that data is always written and confined to the expected tier without needing to know what the individual storage controllers performance profiles are. Use of the confine-to and place-on rules ensures the data is always where it is expected to be and if the data needs to be moved to a different tier it simply becomes a matter of selecting the new volume-group for the share's Objective.

The following is an example of an Objective ensuring all data is written only to the VSP One Block 28-backed volumes.

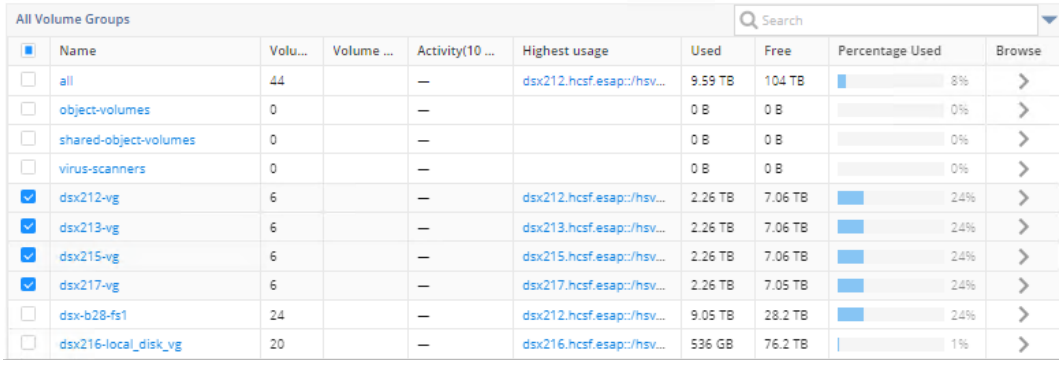

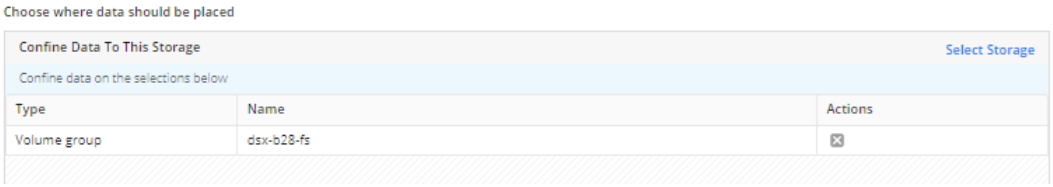

Performance notes

The Hammerspace with Hitachi VSP One Block solution is capable of scaling linearly based on the number of storage blocks that are added to the solution. The following performance numbers are examples of what each 2 and 4 DSX node with VSP One Block is able to achieve assuming no client or network bottlenecks. Each solution was tested based on the configuration recommendations outlined in this reference architecture.

B28 Patterns				
	Rnd Read (K/IO)	Rnd Write (K/IO)	Seq Read (GB/s)	Seq Write (GB/s)
2 × DSX	1080	360	50.24	7.58
4 × DSX	3036	429	169.9	8.68

B26QLC Patterns				
	Rnd Read (K/IO)	Rnd Write (K/IO)	Seq Read (GB/s)	Seq Write (GB/s)
2 × DSX	786	249	45.3	4.18
4 × DSX	1222	249	94.19	4.33

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact