

# **Hitachi iQ with HGX H100 Powered with Hitachi Content Software for File and Hitachi Content Platform**

---

Reference Architecture Guide

© 2024 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

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

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

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

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

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

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

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html>.

## Feedback

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

## Revision history

Changes	Date
Updated content	May 2024
Initial release	April 2024

---

# Reference Architecture Guide

AI and Generative AI are revolutionizing industries by enhancing capabilities in art creation, fraud detection, and even software development. Organizations across various sectors recognize the transformative potential of AI, appreciating its significant value in shaping market dynamics and driving business success. From graphic design and code generation to crafting marketing slogans, AI applications are not only versatile but are proving to be integral in delivering immediate benefits to businesses. As these technologies continue to evolve, the possibilities for innovation and efficiency appear to be limitless.

Organizations are constantly looking for ways to automate their business, accelerate time-to-market, and develop new insights, products, or innovations to propel their business forward. Hitachi iQ allows organizations to realize these goals for their business through intelligent, performant, scalable, and flexible AI and GenAI solutions.

Whether 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.

## Industry specific outcomes

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.

The following are additional use cases powered by NVIDIA and Hitachi iQ:

- Health and Life Sciences (medical devices, genomics, drug discovery, smart hospitals)
- Automotive (parts design, autonomous vehicles)
- Telecommunications (network security, 5G network planning)
- Large Language Models (training, fine-tuning)

From entry level to enterprise capacity, Hitachi iQ is a versatile solution for AI needs. This reference architecture focuses on Hitachi iQ with HGX H100 powered by Hitachi Content Software for File.

This document is for the following audiences:

- Data scientists and data engineers
- AI developers
- Data analysts
- Security operations
- Storage administrators

- System administrators
- IT professionals

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

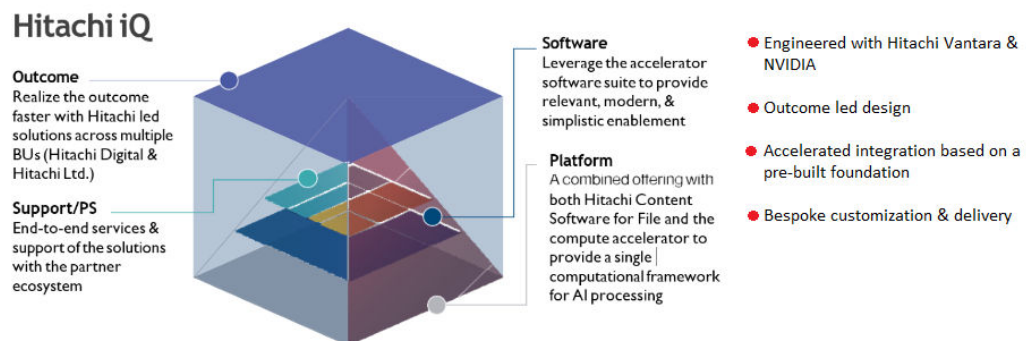
- GPU-based server products
- General storage concepts
- Common IT storage practices

## Solution overview

Unlike conventional AI offerings, Hitachi iQ transcends basic integration and storage capabilities by also layering industry-specific AI outcomes within the AI solution. This approach ensures outcomes are finely tuned to each organization's unique needs and objectives.

Hitachi iQ is the only accelerated solution on the market that provides unified access to data irrespective of where it resides while ensuring explainability, lineage, data accuracy, security, and traceability at any given point for mission critical solutions.

Optimize AI deployments with an end-to-end software-defined AI and data analytics software to streamline the development and deployment of production-grade AI applications, from pilot to production.



## Solution benefits

Hitachi iQ with HGX H100 has been specifically engineered to provide a simplified approach to supporting advanced AI workloads within the enterprise. With the Hitachi iQ portfolio, customers will benefit from a flexible infrastructure delivered and supported by the award-winning support and services organizations of Hitachi Vantara. In addition, customers will realize an improved time to value, better economics at scale, and the flexibility necessary to scale the solution to meet all of their AI demands.

## Accelerated time to value

Most customers struggle when implementing an AI initiative and, unfortunately, too many fail. This reference architecture creates the foundation for an enterprise's AI strategy and has been validated to meet even the highest demanding workloads. Customers can now avoid the trial-and-error approaches that are associated when implementing AI capabilities.

The data pipeline for Artificial Intelligence/Machine Learning (AI/ML) and Generative AI (GenAI) needs significant amounts of accelerated computing and the necessary network and storage platforms to support it. Hitachi iQ is designed to offer high-bandwidth, non-blocking network connectivity between the accelerated compute and the high-performance storage it relies on. The benefit is up to a 20× improvement in GPU utilization, improving the efficiency of the overall solution.

Using an end-to-end non-blocking NDR 400 Gb/s InfiniBand infrastructure, Hitachi iQ can ensure that even the most high-demanding workloads are satisfied within this platform. Additionally, as these workloads scale, this guide provides a foundation to increase the capabilities of the platform through established sizing guidelines and quick deployment methodologies. As the customer's demand increases, right-sizing the environment for various additional AI/ML/GenAI workloads becomes simplified.

The solution is delivered as an integrated system for fast onsite installation and setup, enabling quick access for data scientists, data engineers, application developers, and end users. All the components of the Hitachi iQ platform have been pre-validated and integrated through the global manufacturing process of Hitachi Vantara before shipment.

This solution empowers customers to run analytics workloads at accelerated speed, facilitating decision-making and insight compared to other market offerings.

## Flexible infrastructure

Hitachi iQ offers the flexibility of seamless execution of the workloads in on-premises or hybrid cloud environments. They can start at smaller scale, facilitating small scale workloads or experimentation with minimal initial investment, later scaling to meet the demands of a full production enterprise workload. Customers can independently scale accelerated compute and high-performance storage without any networking bottlenecks.

The carefully optimized CPU-to-GPU ratio enables concurrent execution of both data and AI pipelines. The extensive choice of hardware components such as processors, memory, interface storage – performance and capacity, and software components such as Kubernetes gives customers the flexibility to adopt this infrastructure for their specific requirements and preferences. The Hitachi iQ platform can start with as little as two HGX H100 servers scaling to thousands of nodes across the Hitachi iQ cluster following the practices established by NVIDIA with their BasePOD and SuperPOD architecture.

For improved economics, Hitachi Content Platform, Hitachi Vantara's award-winning object storage technology, can be included in the Hitachi iQ infrastructure, providing high-density scale supported by object storage APIs. This high-density option provides increased density of storage and inclusion of object storage capabilities while lowering the total cost of the solution.

Through the robustness of the Hitachi Vantara portfolio, improved flexibility while offering world class offerings with improved economics, ensures that Hitachi iQ customers are given the most flexibility possible to support their future demands.

**Simplified deployment**

When customers try to implement their own AI solution, they are quickly faced with the challenge of creating complicated systems and performing the necessary integrations on their own. Instead, with the simplification of Hitachi iQ, the entire platform is pre-built in the Hitachi Vantara distribution center, where it has been engineered for maximum performance and reliability. This ensures that when it arrives at the customer's site, it can readily be deployed without all of the complexity of cabling, configuration, interoperability validation, and troubleshooting saving a significant amount of time and improving the return-on-investment (ROI) of the solution.

**Cost effective performance**

Hitachi Vantara has engineered Hitachi iQ to use the latest in AI computing power, provided by NVIDIA and their Hopper GPU product line. Supporting both the H100 and the H200 GPU architectures, Hitachi iQ ensures that today's modern AI workloads (training, tuning, and inferencing) are able to be accelerated with the fastest GPU solution on the market.

Using the NVIDIA H200 Tensor Core GPU, this platform supercharges AI, Generative AI, and high-performance computing (HPC) workloads with the fastest performance possible. By being the first GPU equipped with HBM3e memory, the H200's larger and faster memory fuels the acceleration of Generative AI and large language models (LLMs) while advancing scientific computing for HPC workloads.

Ensuring the heart of the AI solution, the GPU, is effectively utilized throughout the AI pipeline is primarily the concern of the underlying storage system. Hitachi Content Software for File provides the fastest and most economic platform for AI workloads. Using the latest in Generation5 flash and network connectivity, Content Software for File can provide up to a 20× improvement in GPU utilization and performance.

The performance and density of Content Software File, combined with the cost effective, high capacity storage of Hitachi Content Platform (HCP), allows for the best scaling of high-performance and high-density storage. HCP provides an extension of the high-performance filesystem, allowing for customers to scale their AI storage solutions in the most cost-effective way possible.

**Consolidated enterprise support**

Investing in Hitachi iQ is more than just the engineering and validation or the high-performance nature of the storage. Every Hitachi iQ customer will also benefit from the award-winning, world-class support, provided by the Hitachi Vantara Global Support Center. This 24×7×365 support organization ensures that every customer is backed with the confidence of knowing that a team of highly talented support engineers are available, should the need arise.

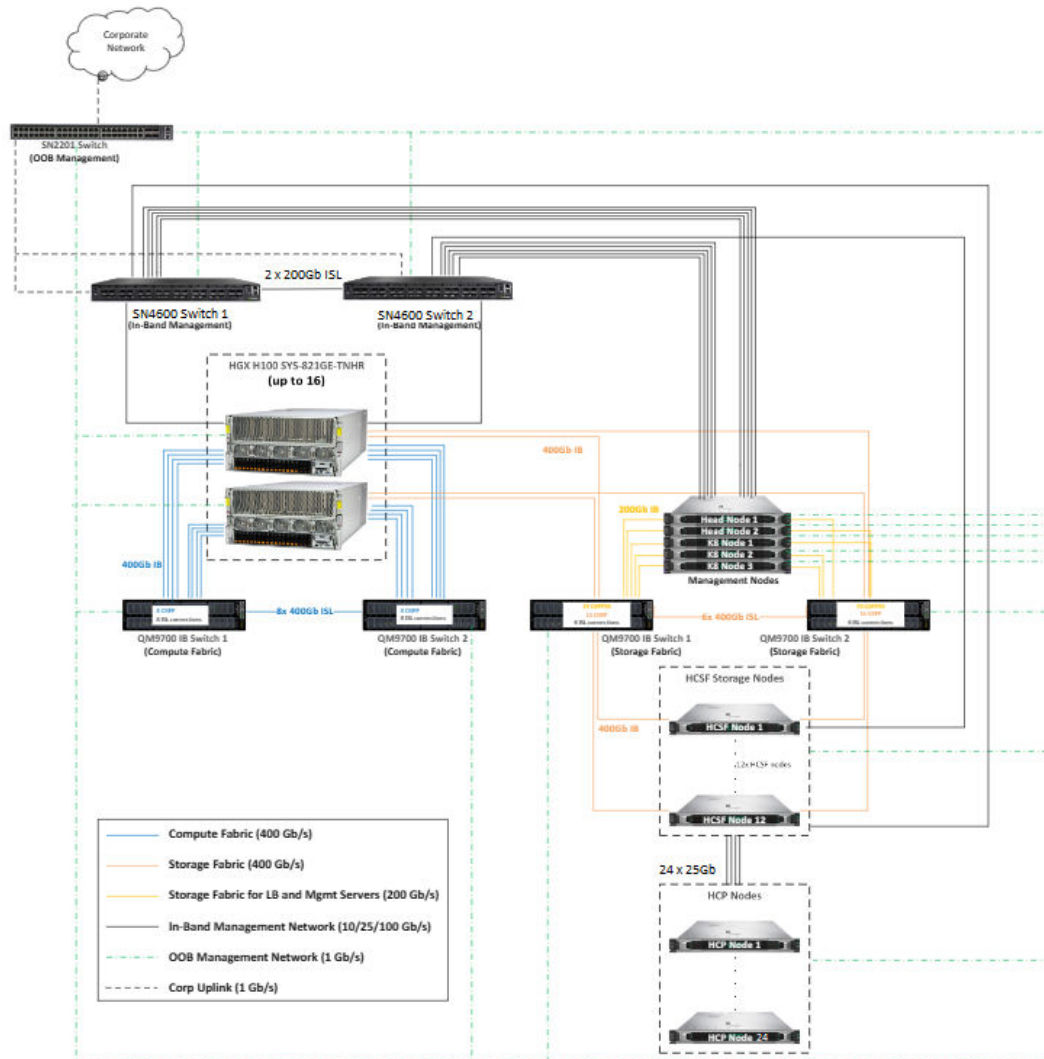
Equipped with the knowledge of not just how the infrastructure operates, but with an understanding of the ecosystem of applications and solutions around each Hitachi Vantara product, the Global Support Center is staffed to ensure each customer receives the highest quality support, at their most critical moments.

Complementing their support capabilities, Hitachi Vantara can also provide Hitachi iQ with a host of different managed services and as-a-service options, creating new ways of consuming the solution while also providing the expertise necessary for these solutions to be integrated and operated in the most effective manner possible.

## Solution architecture

Following this design, the Hitachi iQ with HGX H100 solution includes:

- Fully redundant hardware and network connectivity
- High-performance NDR 400 Gb/s InfiniBand and 100 Gb/s Ethernet connectivity between all components
- Hitachi Content Platform (HCP) for object storage capabilities, data tiering, and data protection



The network is comprised of multiple network fabrics, responsible for:


- Redundant compute interconnectivity via NDR 400 Gb/s
- Redundant storage interconnectivity via NDR400 Gb/s
- Redundant in-band connectivity via 100 Gb/s Ethernet
- Out of band (OOB) management connectivity via 1 Gb/s Ethernet

## Network overview

The following four networks must be configured for this solution:

- Compute Fabric (ibnet1): InfiniBand network for internode GPU communication, connecting all HGX servers via NDR400 Gb/s.
- Storage Fabric (ibnet2): InfiniBand network connecting HGX servers and Hitachi Content Software for File storage via NDR 400 Gb/s.
- In-Band Network (internalnet): Network used exclusively within the cluster for in-band management and communication to external resources (applications, developers, Internet).
- Out-of-Band Management Network (ipminet): Network for out-of-band management, connecting BMC/iLO ports.

The following illustration shows the available ports on the HGX server expansion tray that is used for the compute/storage fabric and in-band management network.



### Compute fabric

For ibnet1-1, ports associated with slots 1 to 4 are connected to switch QM9700-1. For ibnet1-2 the ports associated with slots 5 to 8 are connected to switch QM9700-2 IB.

A total of  $8 \times$  NDR400 links from each HGX system are connected to  $2 \times$  QM9700 IB switches providing a total aggregate compute bandwidth of 3.2 Tb/s for per HGX system.

As the cluster resources grow, additional QM9700 InfiniBand switches can be added to the deployment, configured in a non-blocking CLOS spine-leaf topology.

### Storage fabric

For ibnet2-1, ports associated with slots 10 are connected to switch QM9700-3. For inbet2-2, ports associated with slots 12 are connected to switch QM9700-4.

A total of  $2 \times$  NDR400 links from each HGX system are connected to  $2 \times$  QM9700 IB switches, providing a total aggregate storage bandwidth of 800 Gb/s per HGX system.

### In-band management network

A total of  $2 \times$  100 Gb bonded links from each HGX system are connected to  $2 \times$  SN4600 switches.

Ports 9/10 are used for connectivity, 1 in each SN4600 switch.

For management servers, a 10/25 GB bonded link is connected to the SN4600 switch based on available NICs on the servers.



### Out-of-band management network

The management network connects all the base management controller (BMC) ports, as well as other devices that should be physically isolated from system users to SN2201 switch using a 1 Gb link.



**Note:** An IPMI subnet route must be advertised to the in-band network so that BCM can control hosts using the IPMI network.

## Key solution components

The key solution components for this solution are listed in the following tables. The versions in the tables might be updated over time. Consult your Hitachi technical representative to verify the most current information.

### Hardware components

The following table lists the hardware components used in this solution.

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	HA810 G3	NVIDIA Management servers	OS: RHEL 9.2+ or RHEL 8.7+ Kubernetes: 1.27	5
		CPU: INT Xeon-G 5418Y CPU (185W, 24C, 2.0GHz)		
		RAM: 512 GB		
		Drive: 1 × 6.4 TB NVMe		
		NIC: 2 × IB HDR 200 Gb 2p QSFP56, 1 × 10/25 GbE 2p SFP28		
Hitachi Vantara	HGX H100	GPU SuperServer SYS-821GE-TNHR	OS: RHEL 9.2+ or RHEL 8.7+ BIOS:2.1 BMC Firmware:01.00.26	2 to 16 to scale in 2 server increments

Vendor	Hardware	Detail Description	Version	Quantity
		CPU: 2 × Intel® Xeon® Platinum 8480+ Processors 105M Cache, 2.00 GHz, Total Cores 56, Total Threads 112		
		GPU: 8 × H100 80 GB Tensor Cores		
		GPU memory: 640 GB		
		Memory: 32 × 64 GB DIMMs (2 TB)		
		Drive: M.2 Boot device - 2 × Micron 7450 PRO 1.9 TB NVMe PCIe 4.0 M.2 22x110mm 3D TLC Internal SSD - 8 × Kioxia CD6-R 3.84TB NVMe PCIe 4x4 2.5" 15mm SIE 1DWPD		
		NIC: 8 × NVIDIA 900-9X766-003N -SQ0 PCIe 1-port IB and Ethernet 400 GbE OSFP Gen5x16 CX7, RoHS		

Vendor	Hardware	Detail Description	Version	Quantity
		<p>1 × PCIe 2-port 100Gb/s InfiniBand or Ethernet, QSFP56, Gen 3.0/4.0 x16, CX-6 VPI</p> <p>2 × NVIDIA 900-9X766-003N -SQ0 PCIe 1-port IB and Ethernet 400 GbE OSFP Gen5x16 CX7, RoHS</p> <p>1 × Standard low-profile 2-port 10GbE RJ45 based on Intel X550-AT2</p>		
Hitachi Vantara	HCSF 31116	<p>Hitachi Content Software for File Storage Nodes</p> <p>CPU: Genoa 9534 DP/UP 64C/128T 2.45G 256M 280W SP5</p> <p>RAM: 768 GB</p> <p>Drive: 15 × 3.84TB NVMe</p> <p>NIC: 2 × CX7 IB 400 Gb OSFP, AIOM Dual Port 10G, AIOM Dual Port 25G</p>	OS: Rocky Linux 8.8	8
NVIDIA	NVIDIA QM9700	64 port 400 Gb/s InfiniBand switches	MLNX OS 3.11	4

Vendor	Hardware	Detail Description	Version	Quantity
NVIDIA	NVIDIA SN4600	2RU 64-port 200 GbE Ethernet switch for In-Band Management	Cumulus Linux v5.x	2
NVIDIA	NVIDIA SN2201	48 port 1 Gb/s Ethernet switch for OOB management	Cumulus Linux v5.x	1
Mellanox NIC	ConnectX-7	For HGX node adapter firmware	v28.39.2048-LTS	-
	ConnectX-6	For HGX node adapter firmware	v20.39.2048-LTS	-

### Software components

The following table lists software components used in this solution.

Software	Version	Function
HGX OS	OS: RHEL 8.7+ or RHEL 9.2+ Kubernetes: 1.27	HGX H100 OS
NVIDIA Driver and Tools	NVIDIA datacenter drivers NVIDIA Fabric Manager (FM) NVIDIA NVSwitch Configuration and Query Library (NSCQ) NVIDIA DCGM CUDA Toolkit OFED: v23.10-1.1.9.0-LTS	NVIDIA software to run CUDA applications
HGX Management Servers	RHEL 8.7+ or RHEL 9.2+ Kubernetes 1.27 Base Command Manager 10	SDDC infrastructure
Hitachi Content Software for File OS	Rocky Linux 8.8	Storage servers

Software	Version	Function
Hitachi Content Software for File software	v4.2.10.x	Storage software
QM9700 OS	v3.11	InfiniBand switches – Compute and Storage
MSN2201 OS	Cumulus Linux 5.x	Out-of-band Ethernet - management

## Engineering validation

This section summarizes the key lab verification tests performed on the Hitachi iQ solution outlined in this reference architecture.

The integration tests include the following:

- Validation that the performance achieved by the Hitachi iQ with HGX H100 solution is comparable to a like sized solution using Hitachi iQ with NVIDIA DGX H100 solution.
- NVIDIA Base Command Manager is compatible with the Hitachi HGX H100 server for provisioning and management.
- NVIDIA Base Command Manager is capable of monitoring the Hitachi HGX H100 cluster infrastructure.
- The solution does not introduce any Single Points of Failure (SPOF) within the network fabric or compute resources.
- The Hitachi iQ with HGX H100 solution is compatible with NVIDIA AI Enterprise Software, ensuring no interoperability issues in the solution software stack.

## Hitachi Content Platform (HCP)

Hitachi iQ also provides additional benefits by leveraging HCP for multiple business critical, data management, and data protection capabilities, while offering an economical way to tier data to object storage. HCP Object Storage also helps businesses meet governance and regulatory compliance matters. High-performance and all-flash capable object storage options are also available for high-performance workloads such as AI/ML, and analytics; data warehouses add additional capabilities, while keeping things resilient and reliable in the face of ransomware threats.

## Conclusion

Gain intelligent, scalable, and adaptable GenAI solutions and services that outpace conventional approaches with the Hitachi iQ platform. Unlike basic integration methods, Hitachi iQ elevates outcomes by integrating industry-specific capabilities into the AI solution stack, ensuring more precise and pertinent business results.

**Hitachi Vantara**

Corporate Headquarters  
2535 Augustine Drive  
Santa Clara, CA 95054 USA



[HitachiVantara.com/contact](https://HitachiVantara.com/contact)