

Solution Profile

Conversational Agentic AI with VSP One Block, Hitachi iQ and iQ Studio



Deliver Performance

Accelerated compute for real-time inference, intent detection and agentic workflow execution

Ensure Reliability

Always-on availability for 24/7/365 operations

Retain Control

Sovereign AI for governance and control over models, data and AI outputs

Enterprises can look to secure, reliable conversational AI at scale to improve customer engagement and services availability.

Customer engagement volumes are rising rapidly across chat, voice and digital channels. Meanwhile expectations for personalized, always-available services continue to increase. And, at the same time, service and operations leaders are looking to improve outcomes with constrained resources.

Advances in large language models (LLMs) and agentic AI are expanding what organizations can automate and optimize. The real opportunity is not simply to launch another chatbot, but to build an enterprise-grade conversational platform that connects securely to business data, integrates with core workflows, and operates with the governance and control required at scale.

Meet Growing Demands for Speed, Responsiveness and Operational Scale

Hitachi iQ is built on a proven, enterprise-ready architecture designed to support conversational and agentic AI at scale. It combines GPU-powered compute for real-time AI processing, high-performance storage for fast access to operational data, resilient networking for large volumes of concurrent interactions, and redundant power for continuous availability. Together, these capabilities create a secure, reliable foundation for AI-driven engagement and automation across mission-critical environments.

The result is a platform that can improve responsiveness, streamline operations and deliver measurable business value in industries where performance, uptime and scale are essential.

Hitachi iQ takes a differentiated approach designed specifically for enterprise AI inference workloads. While many vendors focus on object- or file-based architectures, this solution uses block storage to support faster, more precise retrieval of trusted information. The platform is designed to scale seamlessly across compute, networking and storage, enabling organizations to meet growing demands for speed, responsiveness and operational scale.

Solution Components: Hardware

- Hitachi Virtual Storage Platform One Block (VSP One Block) is a high-performance, GPU-optimized block storage appliance that delivers ultra-low latency and high IOPS for mission-critical AI and ML workloads. It enables rapid retrieval of data and vector indexes for multimodal search, allowing the solution to access operational knowledge, such as maintenance logs or sensor readings, with minimal delay.
- Hitachi iQ is the Hitachi end-to-end enterprise AI and ML platform, combining high-performance infrastructure with integrated AI software to accelerate enterprise AI deployments.

- Accelerated compute is delivered through GPU-enabled nodes and high-performance servers with NVIDIA® GPU acceleration, providing the computational power required for real-time AI inference and complex multimedia processing.
- The high-speed network is a unified, high-bandwidth Cisco® fabric that connects all solution components, including compute nodes, storage systems and services, with low latency and massive throughput.

Solution Components: Software

- Red Hat Enterprise Linux® is a stable, secure enterprise Linux operating system that provides the foundation for the solution's servers and software components. It delivers enterprise-grade reliability, enhanced security and performance optimizations for AI workloads.
- Red Hat OpenShift® is a Kubernetes®-based container platform used to deploy and manage the solution's microservices and AI components at scale. It provides the orchestration and management capabilities needed to support enterprise-grade containerized workloads.

Use Case: The Art of the Possible

Enterprise conversational AI solutions go beyond just “smart chat bots”. They can engage humans and systems alike in multi-turn conversations in pursuit of solving a problem or to achieve a specific set of objectives. These solutions can be used in a variety of industries. Some examples of enterprise conversational AI include:

Finance

- Customer service and banking operations teams can use conversational AI to access account activity, transaction status, loan application progress and policy details through natural language, reducing the need to navigate multiple fragmented systems.
- Fraud and risk teams can use conversational AI to retrieve transaction histories, customer profiles, sanctions screening results and prior case notes, helping accelerate suspicious activity investigations.
- For claims and exceptions handling, agentic AI can identify anomalous transactions or disputed payments, gather supporting records, determine the appropriate workflow and coordinate next-step actions within a single conversational process.
- Bankers and advisors can use conversational AI to access portfolio summaries, product eligibility, client interaction histories and compliance-approved recommendations in real time during customer engagements.
- Regulatory and compliance teams can use conversational interfaces to generate audit-ready reports, retrieve policy evidence, explain decision trails and maintain documentation for know your customer, anti-money laundering and internal control reviews.

Healthcare

- Clinical operations teams can use conversational AI to access patient histories, lab results, medication records, discharge instructions and care pathways through natural language, reducing the need to switch between multiple clinical systems.
- Nurses, case managers and physicians can use conversational AI to retrieve referral status, prior authorizations, appointment availability and post-acute care plans, helping reduce delays in patient transitions.
- For clinical documentation support, agentic AI can capture encounter context, retrieve relevant patient history, draft structured summaries and route documentation for review within a single conversational workflow.
- Patient access and contact center teams can use conversational systems to manage appointment changes, benefits questions, preprocedural instructions and follow-up requests, reducing administrative burden while improving patient responsiveness.
- Quality, compliance and operational readiness teams can use conversational interfaces to surface protocol guidance, infection-control procedures, incident documentation and audit trails needed for accreditation, safety reviews and regulatory reporting.

Manufacturing

- Shop floor teams can use conversational AI to access production schedules, machine status and quality metrics through voice, allowing operators to stay focused on their work without leaving the workstation.
- Maintenance engineers can use conversational AI to retrieve sensor trends, fault-code histories and OEM repair procedures as soon as anomalies are detected, helping reduce response times and improve equipment uptime.
- In supply chain disruption scenarios, agentic AI can identify critical component shortages, surface alternative suppliers from procurement data and initiate approval workflows within a single conversational process.
- For shift handovers, conversational AI can generate summaries of production performance, yield deviations and outstanding issues, reducing reliance on lengthy manual documentation.

Rail Operations

- Control room teams can use conversational AI to query real-time train positioning, delay propagation and track conditions through natural language, reducing the need to navigate multiple fragmented systems.
- Field maintenance engineers can use voice-enabled conversational AI to retrieve asset fault histories, maintenance procedures and safety certifications hands-free while working trackside.

- In incident response scenarios, agentic AI can automatically retrieve maintenance records, identify the appropriate crew and coordinate dispatch within a single conversational workflow when signal failures or other operational issues occur.
- Conversational AI can help manage large-scale service disruptions by handling high volumes of passenger rebooking requests during network events without overwhelming contact center operations.
- Workforce and compliance teams can use conversational interfaces to manage rostering questions, confirm safety briefings and maintain regulatory audit trails more efficiently.

Energy

- Grid control room teams can use conversational AI to query real-time load distribution, fault locations and outage status through natural language, reducing the need to switch between disparate supervisory control and data activation (SCADA) screens.
- Field technicians at substations and wind farms can use voice-enabled conversational AI to retrieve isolation procedures, equipment specifications and safety permits hands-free, while working in the field.
- In outage management scenarios, agentic AI can detect faults, retrieve asset history, identify the nearest qualified crew and coordinate restoration workflows within a single automated process.
- Regulatory and emissions compliance teams can use conversational interfaces to log incidents, generate emissions reports and maintain audit-ready documentation in real time.

Support Conversational AI at Scale

Today's enterprises are challenged to address rapidly increasing customer engagement volumes and demand for always-available services while delivering better a customer experience within the constraints of company resources.

Look to the Hitachi iQ solution for enterprise conversational AI to create a secure, reliable foundation for AI-driven engagement and automation across mission-critical environments. The platform scales seamlessly across compute, networking and storage, enabling your organization to meet growing demands for speed, responsiveness and operational scale — delivering measurable business value.

[Learn More](#)

About Hitachi Vantara

Hitachi Vantara is transforming the way data fuels innovation. A wholly owned subsidiary of Hitachi, Ltd., we're the data foundation the world's leading innovators rely on. Through data storage, infrastructure systems, cloud management and digital expertise, we build the foundation for sustainable business growth.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
hitachivantara.com | community.hitachivantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
hitachivantara.com/contact

© Hitachi Vantara LLC 2026. All Rights Reserved. All other trademarks, service marks and company names are properties of their respective owners.

NVIDIA, CUDA and TensorRT are registered trademarks and NVIDIA AI Enterprise and Triton Inference Server are trademarks of NVIDIA Corporation in the United States and other countries. Cisco is a registered trademark of Cisco, Inc. Red Hat Enterprise Linux, Red Hat and OpenShift are registered trademarks of Red Hat, Inc. Kubernetes is a registered trademark of the Linux Foundation.

HV-BTD-SP-VSP-One-Block-for-Enterprise-Conversational-AI-29June26-A