



*Solution Profile*

# Supercharge Your AI Data Workflows with Hitachi and Hammerspace

## AI Without Complexity: Achieve Faster Insights, Make Smarter Decisions, and Achieve Cost-Effective Growth.

### Optimize Data for AI Workflows

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.

[A recent study by Hitachi Vantara and the Enterprise Strategy Group \(ESG\)](#) on AI buyer insights revealed that 97% of organizations with an ongoing GenAI project consider it a top 5 priority. Despite the potential benefits, Fortune 2000 companies face challenges in AI adoption. Some of the top concerns include data availability and quality, cost/technical debt, and integration challenges.

Hitachi and Hammerspace bridge these challenges and keep data environments running non-disruptively. Furthermore, customers gain 4:1 data reduction and 100% data availability guaranteed for data stored on the Virtual Storage Platform One (VSP One) Data Platform.

### Is Your AI Strategy Reaching its Full Potential?

In recent years, nearly every company has embarked on an AI initiative, yet many have struggled to achieve significant value or meaningful ROI due to persistent data challenges, such as:

- **Data Silos:** Restricts seamless access to datasets, limiting the effectiveness and accuracy of AI models.
- **Data Sprawl:** Complicates AI by making it difficult to manage, secure, and analyze dispersed datasets efficiently.
- **Slow Data Movement:** Delays the processing and analysis of data, reducing the efficiency and responsiveness of AI applications.
- **Legacy System Limitations:** Impedes AI by failing to support the speed, scale, and complexity required for modern AI workloads.



- **Disparate Environments:** Complicates AI by making it difficult to achieve cohesive data management and analysis, leading to inefficiencies and inconsistent results.

If these data challenges are left unresolved, AI initiatives can suffer from inefficiencies, inaccuracies, and limited scalability, ultimately hindering the potential to deliver valuable insights and drive innovation.

### Tackling the Data Challenge of AI with the Hitachi iQ Portfolio

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. The solution offers:

- **Global Data Access:** Call up data from any location, regardless of where it's stored.
- **Abstracted Storage Management:** Access data via a unified global data platform.

- **Data Orchestration:** Make your data visible and accessible in the right place at the right time while optimizing performance, cost, and compliance.
- **High Performance:** Ensure quick access, even for remote or distributed users, with data that can be placed proactively based on data orchestration policies.
- **Scalability:** Add storage resources to solutions as data requirements increase — without any disruption.
- **VSP One Platform Customer Guarantees:** Count on 100% data availability and 4:1 data reduction guaranteed (for data stored on VSP One Data Platform).

## Global Data, Local Speed: Revolutionizing AI Workflows.

One of the key benefits of this integrated solution is its ability to consolidate various storage systems into a unified, high-performance global data environment. For enterprises managing AI workflows, this means that data spread across on-premises systems, edge devices, and multiple cloud platforms can be accessed as if it were local. This eliminates the delays and inefficiencies associated with traditional data silos. The global namespace ensures that massive datasets such as petabytes of data that feed AI workloads are readily available to GPU clusters, regardless of the location of the data or compute resources. This seamless access reduces latency, keeps GPUs supplied with data, maximizes computational efficiency, and accelerates the time-to-insight.

## Seamless Data Management - No IT Overhaul Required.

Operationally, Hammerspace streamlines the management of AI data pipelines. IT teams are relieved from juggling multiple point solutions or manually migrating or copying data between silos. The system's metadata-driven orchestration automates tasks such as data ingestion, tiering, and protection, allowing resources to be redirected towards higher-value activities like model tuning or strategic planning.

## Keep GPU and Compute Resources Fully Utilized

Maximize resource utilization by fully leveraging the high performance required by AI pipelines. Unlike traditional HPC file systems, this solution ensures that your infrastructure resources are always operating at peak efficiency while retaining enterprise-grade features. Integration with NVIDIA® GPUDirect® technology improves AI pipelines for RAG data sources. This is transformative for enterprises running generative AI, deep learning, or real-time analytics, where delays can lead to lost productivity or innovation.

## Optimize Workloads for AI

AI projects require infrastructure that unifies, orchestrates, and optimizes data access across distributed environments, delivering seamless scalability, high performance, and compliance to meet the demands of diverse analytics, AI, and machine learning workflows.

### LLM Customization

- LLM customization requires specific computational resources and high-performance data pipelines for fine-tuning.
- Bottlenecks in data access or movement can dramatically slow down development cycles.

*The solution provides the benefits of HPC file systems with standards based NFS to accelerate AI data pipelines and GPU computing.*

### Fine-tuning & RAG

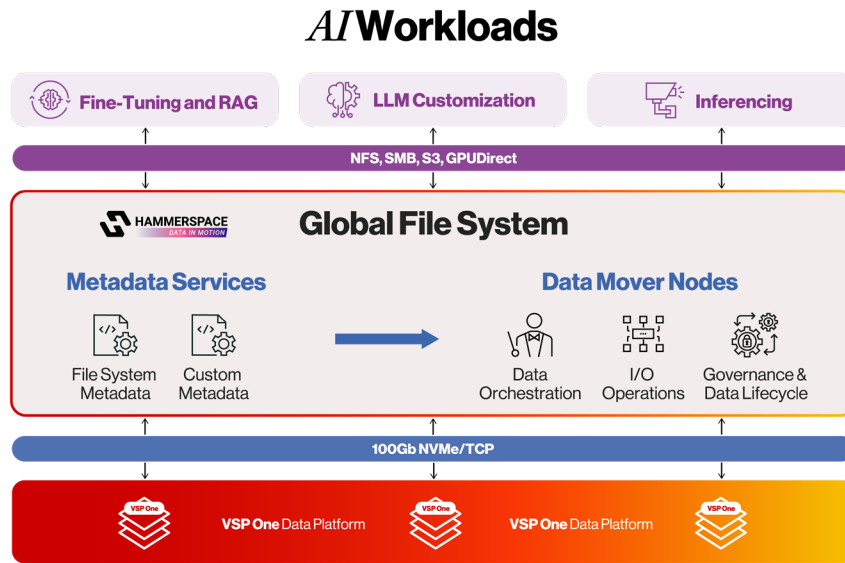
- Fine-tuning models require synchronized access to both the retriever (search/indexing) and the generator (LLM).
- Data pipeline inefficiencies can lead to mismatches, latency, and poor performance.
- Retrieval augmented generation (RAG) allows the optimization of AI models by connecting them with external knowledge bases to deliver more relevant responses at a higher quality.

*The solution provides high speed data access and automates data placement in the appropriate storage tier, cloud region, or cloud service non-disruptively with automated Data Orchestration.*

### Inferencing

- AI inferencing requires immediate access to relevant data, especially in latency-sensitive applications like chatbots, recommendation systems, and autonomous systems.
- Delays in retrieving data can degrade user experience and model effectiveness.

*With high-speed access to global data, this solution provides quick and seamless access to the data necessary for inferencing, crucial to ensuring real-time applications like chatbots and recommendation systems provide the most relevant and up to date information.*



## Solution Overview

### AI Workloads

- **Fine Tuning & RAG** - Adapt models to specific tasks and leveraging external data sources for more accurate and contextually relevant outputs.
- **LLM Customization** - Tailor large language models to specific tasks or domains.
- **Inferencing** - Enable real-time decisions, generate predictions, and drive intelligent actions.

### Hammerspace Global File System

- **High Performance Data Access** - Using industry standard protocols
- **Metadata Services** - Manages all the metadata within the system, and controls all the I/O and other actions performed by the Data Mover Nodes.
- **Data Mover Nodes** - Provides scale-out engines that connect data and storage to users and applications, and performs all I/O operations, replication, data movement, and so on.

### VSP One Data Platform

- VSP One is a next-generation, all-in-one solution that powers infrastructure, data and apps without constraints. It's a scale out hybrid cloud data platform backed by 100% data availability and 4:1 data reduction guaranteed.

## Hitachi Vantara and Hammerspace: Your Competitive Edge for the AI Era

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. Plus, if the data is running on VSP One, you are guaranteed 100% data availability and 4:1 data reduction, making it the only GenAI data platform that can offer this level of assurance.

**Unlock the full potential of your AI initiatives with Hitachi Vantara today.**



**Corporate Headquarters**  
 2535 Augustine Drive  
 Santa Clara, CA 95054 USA  
[hitachivantara.com](http://hitachivantara.com) | [community.hitachivantara.com](http://community.hitachivantara.com)

**Contact Information**  
 USA: 1-800-446-0744  
 Global: 1-858-547-4526  
[hitachivantara.com/contact](http://hitachivantara.com/contact)

© Hitachi Vantara LLC 2025. All Rights Reserved. HITACHI and Pentaho are trademarks or registered trademarks of Hitachi, Ltd. Hammerspace and the Hammerspace logo are trademarks of Hammerspace, Inc. All other trademarks, service marks and company names are properties of their respective owners.

HV-BTD-SP-VSP-One-With-Hammerspace-25Feb25-A

