

AI orchestration: Turn data chaos into competitive advantage

What's the offer?

Traditional AI platforms assume enterprises have centralized data with unlimited compute power. But for many enterprises, the reality is different: decades of legacy systems, petabytes of distributed data, and strict compliance requirements.

Kamiwaza enables enterprises to securely orchestrate AI across distributed data without costly migrations, expensive hardware, or compromising security. The Kamiwaza AI Orchestration Engine allows enterprises to securely run any AI model on any infrastructure, wherever the data resides. This flexible approach accelerates decision-making, automates complex workflows, and helps keep sensitive data secure.

Who will buy it?

Kamiwaza's software stack facilitates advanced analytics, agentic AI, workflow automation, document processing, and many other use cases for the following industries:

- Retail
- Manufacturing
- Healthcare and life sciences
- Government and defense
- Financial services (banking, asset management, insurance)
- Critical infrastructure (energy, utilities, transportation)
- Media

Business value

- Pre-integrated software stack that is ready to deploy with customized AI agents
- Broad-based support, including dedicated architects who provide solution design and optimization, guidance, and outcome support
- Measurable business impact
- Silicon-agnostic
- No data centralization needed
- Compatible with any large language model (LLM)

A Fortune 500 consumer packaged goods company

**reduced document
processing time by 87%**


using AI-driven solutions from Kamiwaza.

Kamiwaza helped a public sector-focused organization achieve a

**75% reduction in total
cost of ownership**

by replacing legacy robotic process automation (RPA) with AI-driven orchestration.¹

¹ Source: kamiwaza.ai/document-processing



Customer pain points and solution benefits

Pain point	Solution benefit
Disconnected data sources that make data centralization slow and costly	Distributed Data Engines (DDEs) run where enterprise data resides and bring the right intelligence to it, unlocking insights across the entire data estate
Operational complexity hinders AI scalability	An Inference Mesh coordinates the execution of AI workloads across cloud, on-premises, and edge to distribute inference where it runs best
Security and compliance regulations for regulated industries prevent data from being moved or centralized	Relationship-based access control (ReBAC) manages permissions through dynamic, graph-like relationships between entities, accounting for delegation and change over time
Implementing new solutions can bring vendor integration challenges	Solution-in-a-box: Fully validated, integrated (hardware and software) solutions built at an HPE factory. Right-sized configurations to fit any deployment
Unpredictable pricing models and the need to manage time-consuming, complex procurement processes	Ready-to-quote same day with competitive prices

Solution components

- Kamiwaza AI Orchestration Engine, Distributed Data Engine, and Inference Mesh bring data, context, and action together inside the enterprise’s security perimeter
- Foundation models, domain-specific and fine-tuned models, and task-specific agents make sense of enterprise data and take informed action
- HPE Private Cloud AI and/or HPE ProLiant Compute DL380a Gen12 servers provide an enterprise-grade hardware foundation
- Intel® Xeon® 6 processors deliver the scalability and performance that AI demands

What are my next steps?

Reach out to hello@kamiwaza.ai to engage with you on your next opportunity.

More resources

- [Kamiwaza website](#)
- [Kamiwaza/HPE partnership page](#)
- [HPE ProLiant Compute servers](#)
- [HPE Private Cloud AI](#)
- [Discover Press Studio 2025 - Kamiwaza AI video](#)



HEWLETT PACKARD ENTERPRISE
hpe.com

