



Kamiwaza AI

accelerated by **intel**.

Helps Enterprises
Bring GenAI
Anywhere and
Everywhere

Kamiwaza's AI Orchestration Engine simplifies modern AI adoption by integrating our Stack, Inference Mesh, and Locality-Aware Data Engine. With enterprise data spread across Cloud, On-Prem, and Edge environments, Kamiwaza enables local processing, keeping data where it lives and only passing inference requests and results between locations. To ensure seamless integration, Kamiwaza exposes its API, SDK, and Model Context Protocol (MCP), allowing new and existing enterprise applications to connect directly with the Orchestration Engine. This enables AI workflows and processes across all locations and data sources while seamlessly operating within the enterprise security framework.

Key Features			
AI Anywhere and Everywhere	Inference Mesh	Local Data Engine	Hardware Agnostic

	Vertical Markets:	Use Case:	Country/Geo:
	<ul style="list-style-type: none">RetailFinancial ServicesGovernment	<ul style="list-style-type: none">CommercialEnterprise AI Software	<ul style="list-style-type: none">Asia Pacific & JapanEurope, Middle East, AfricaIndia, North & South Americas

“By integrating Gaudi into our orchestration engine, we ensure enterprises have access to cost-effective, high-performance AI across cloud, on-prem, and edge environments-without vendor lock-in.”

Luke Norris
Chief Executive Officer &
Co-Founder, Kamiwaza AI

- Intel Products and Technologies
- [Intel® Xeon® Scalable Processors](#)
 - [Intel® Gaudi® AI Accelerators](#)

- Learn more:
- [Kamiwaza AI Website](#)
 - [Kamiwaza AI Solution Page](#)