


Equinix Distributed AI™ Hub

Run AI where your data lives, connecting every cloud, model and AI provider through private, intelligent interconnection

AI is distributed, and that changes everything

Enterprise data no longer lives in one place, and neither does AI. The inference era has accelerated this shift: Training favors centralization, but inference follows the data, and agentic AI multiplies that complexity across every model, provider, and geography while regulations redraw where data can live and move.

Centralized infrastructure was not built for this. Getting from GPU cluster to production workload means navigating interconnection, latency, compliance, and ecosystem access, which are distribution and trust challenges that compute alone cannot solve.



Four problems blocking enterprise AI



Shadow AI

Every department adopts tools independently, creating duplicated effort and inconsistent quality.



Security gaps

Sensitive data pasted into public LLMs leads to breaches, IP loss and compliance violations.



Spiraling cost

With no centralized tracking, a \$50K pilot quietly becomes a \$500K production bill.



Network inadequacy

AI over the public internet means nondeterministic latency and degraded user experience.

The root cause

AI is being treated as a feature to bolt on, not as the platform infrastructure it requires. It demands the same platform-thinking enterprises already apply to cloud, security and data.

The solution: Equinix Distributed AI Hub

The Equinix Distributed AI Hub is the most comprehensive vendor-neutral framework for connecting every cloud, every model, and every AI provider. Built on the world's most vibrant neutral ecosystem and connected through private, intelligent interconnection, it places training where compute is abundant, sensitive processing where privacy and sovereignty demand, and inference where data is generated. It expands the proven Equinix Performance Hub model rather than replacing it, so enterprises extend infrastructure they already operate.

One neutral place Where the ecosystem converges

Connect private and public data, open and closed models, and a rapidly expanding partner ecosystem of neoclouds, cloud service providers, foundation model APIs, and AI SaaS on a single backbone, with instant access to a diverse set of AI tools and services.

Without lock-in Freedom by design

Hyperscalers optimize for their own platform, not neutrality. Equinix provides the neutral fabric that connects all of them, giving enterprises model choice, cost control, and data sovereignty across a fragmented AI landscape.

281

AI data centers

77

metros

36

countries

~10 ms


to most
internet users

99.999%

SLA with redundant
Equinix Fabric connections

One architecture, from private infrastructure to the AI ecosystem

The Equinix Distributed AI Hub unifies four layers: AI-ready data centers and private compute, intelligent Equinix Fabric® interconnection; a control plane of gateway, security and observability; and direct reach into the full AI ecosystem of clouds, neoclouds, foundation models and AI SaaS.



The freedom to run AI everywhere it matters

Connect data, models and a rich ecosystem

Intelligent connectivity that places every workload where it performs best.

Inference at the metro edge

Run AI where users and data live, with the proximity inference demands.

Privacy and sovereignty

Data residency controls and privacy protections your business requires.

The data center AI needs

High-density GPU infrastructure, liquid cooling and managed power capacity.

Deploy the model that fits

Option 1

Physical

Deploy DAI Hub core systems on physical hardware designed for your performance and security needs. Core items can include AI Guardrails, AI Gateways, and Observability. Directly connect to an Equinix Performance Hub, with your security controls, to your workloads and data with optimized interconnection to all public clouds, neoclouds, and the broader AI ecosystem.

Option 2

Hybrid physical and virtual

Deploy the same DAI Hub core systems globally and on-demand as agile VNFs using Equinix Network Edge and Fabric Cloud Router. Accelerate deployment, increase agility and minimize additional Capex all while maintaining the same benefits of immediate proximity to the AI Ecosystem through private, optimized connectivity.

Option 3

Cloud

Deploy the same services natively in Amazon Web Services (AWS), Microsoft Azure or Google Cloud Platform (GCP), connected through Equinix Fabric for private multicloud reach and one unified control plane.

Five capabilities in the control plane



Unified AI gateway

A single API across providers, plus MCP and agent (A2A) gateways with rate limiting, cost tracking, load balancing and failover.



AI security & guardrails

Threat detection, data loss prevention and policy enforcement with AI-aware inspection at every location.



AI observability

End-to-end tracing of every interaction, provider comparison, cost-per-request charge-back and team attribution.



AI optimization

Prompt and semantic caching plus semantic routing cut cost and latency while improving user experience.



Optimized interconnection

Private, low-latency Equinix Fabric delivering 50–70% lower cloud egress and 99.999% SLA with redundant Equinix Fabric connections.

Land and expand

Start with a single hub, add high availability, then scale to a multiregion footprint as workloads grow with centralized governance throughout.

Why Equinix, and why now?

Equinix vs. a private data center

- Built for one location; AI is multistage and multilocation by nature.
- No path to the clouds, model providers and GPU infrastructure AI depends on.
- Power and cooling become a full rebuild, not an upgrade

Equinix vs. other colocation

- Regional presence is not global infrastructure for inference at the edge.
- A connection without ecosystem access is just a cable, not an architecture.
- No privacy-preserving connectivity across cloud, edge and corporate environments.

Equinix vs. hyperscalers

- Optimized for their platform and lock-in, not neutrality.
- Centralized billing makes distributed AI data transfer costs unpredictable.
- Marketplaces favor first-party tools over best-of-breed integration.

Proven outcomes

6x

faster agentic rollouts (Zetaris)

22x

faster data transfer (AiHPC)

10x

faster performance (Merck)

50%

lower AI workload cost vs. public cloud

Sub-10 ms

latency for real-time inference

About Equinix

Equinix, Inc. (Nasdaq: EQIX) shortens the path to boundless connectivity anywhere in the world. Its digital infrastructure, data center footprint and interconnected ecosystems empower innovations that enhance our work, life and planet. Equinix connects economies, countries, organizations and communities, delivering seamless digital experiences and cutting-edge AI—quickly, efficiently and everywhere.

[Equinix.com](https://equinix.com)

Move AI from pilot to production where your data lives.

Run AI closer to users and data, connect across the neutral ecosystem and apply enterprise security guardrails by design.

[Learn more](#)

Equinix Distributed AI, Equinix Fabric and other Equinix marks are trademarks or registered trademarks of Equinix, Inc. in the United States and other countries. Third-party names referenced may be trademarks or registered trademarks of their respective owners.

[Equinix.com](https://equinix.com)

© 2026 Equinix, Inc. SB_Distributed-AI-Hub_EN | 6716 | v082426 | Q425

