




AMAX AI Factory Networking for On-Prem Enterprise AI

AMAX's RackScale Solutions with NVIDIA Quantum-X800 InfiniBand and Spectrum-X Ethernet for high-performance training, inference, HPC, and AI cloud deployments

AMAX Validated RackScale AI Factory Platform

AMAX helps enterprises turn high-performance NVIDIA networking into deployable, production-ready AI factory infrastructure. From RackScale 32 with NVIDIA HGX B300 systems to RackScale 72 with the NVIDIA GB300 NVL72 platform, AMAX combines GPU-dense compute, NVIDIA Quantum-X800 InfiniBand or NVIDIA Spectrum-X Ethernet, power, cooling, rack-level design, validation, and software readiness for on-prem AI training, inference, HPC, simulation, and AI cloud workloads.

This delivery model helps customers reduce deployment risk, improve cluster performance, and scale from single-rack deployments to multi-rack enterprise AI environments with a clear path from design to production handoff.



AMAX AI Factory

How AMAX Supports On-Prem AI Factory Deployment

Design	Validate	Deploy	Enable	Scale
Rack-level architecture aligned to workload, power, cooling, networking, and software requirements.	Cluster testing to reduce deployment risk before production handoff.	Factory-built infrastructure prepared for enterprise data center rollout.	NVIDIA AI Enterprise support, OS and container runtime setup, cluster software, and workload readiness.	Network fabric choices that support single-rack, multi-rack, and AI cloud growth.

Why Networking Matters in AI Factories

As GPU clusters scale, network fabric performance has a direct impact on GPU utilization, job completion time, and application responsiveness. AMAX helps customers select the right NVIDIA networking architecture for their deployment model, whether the priority is maximum scale-out performance with InfiniBand or AI-optimized Ethernet for enterprise data center environments.

NVIDIA Quantum-X800 InfiniBand Performance-first fabric for AI training and HPC 	WORKLOADS: Foundational model training, HPC, simulation, advanced AI research KEY VALUE: 800 Gb/s connectivity, high-throughput data movement, SHARPV4 in-network computing DEPLOYMENT FIT: Performance-first clusters and multi-rack AI factories Best For: Enterprise and research teams building high-performance AI factories where scale-out bandwidth, low latency, and GPU utilization are critical.
NVIDIA Spectrum-X Ethernet AI-optimized Ethernet for enterprise AI clouds 	WORKLOADS: Enterprise AI, generative AI, inference, training clusters, AI cloud KEY VALUE: Standards-based Ethernet with NVIDIA AI networking optimization and higher performance than traditional Ethernet for generative AI workloads DEPLOYMENT FIT: Enterprise data centers and Ethernet-based AI cloud fabrics Best For: Organizations standardizing on Ethernet while building scalable AI platforms for training, inference, and generative AI workloads.

Ready to build your on-prem AI Factory? Contact AMAX to configure the right RackScale Architecture, NVIDIA networking fabric and software environment for your workload, data center, and growth plans. www.amax.com/contact