THE PROBLEM

Your Network Wasn't Built for Distributed AI Workloads

AI training and inference expose latency, oversubscription, and East-West bottlenecks your network was never designed to handle.

Performance	Legacy networks introduce jitter, packet loss, and congestion that break GPU efficiency and model training timelines
Scale	AI data flows grow exponentially – traditional architectures can't scale East-West traffic without redesign
Security	AI expands machine-to-machine communication, creating blind spots traditional security models weren't built to see

THE SOLUTION

Purpose-Built AI Networking - From Fabric to Security

Architected for GPU clusters, high-throughput fabrics, and secure AI data pipelines – end-to-end.

Performance	Lossless, low-latency fabrics optimized for GPU-to-GPU communications (RoCEv2/AI transport aware design)
Scale	Leaf-spine and AI fabric architectures designed for horizontal GPU scaling without re-architecture
Security	Identity-aware Zero Trust segmentation for AI workloads, data pipelines, and East-West traffic

The Partner That Designs AI Infrastructure - Not Just Networks

We design AI-ready architectures from workload to fabric to security – before hardware decisions are made.

01

We Design From the AI Workload Back

We start with AI training and inference patterns – GPU communication, storage access, and latency sensitivity – then design the network to match.

02

End-to-End Network Expertise

From campus to AI data center to cloud, we design integrated architectures across:

- AI fabrics (Leaf-Spine, CLOS)
- Data pipelines
- WAN / Edge inference

One architecture. No silos.

03

Security Built Into the Architecture

We secure AI environments at the workload level – protecting models, data pipelines, and machine identities with Zero Trust and SASE frameworks.

Embedded. Not bolted on.