Link Oregon

SPONSORED LIGHTNING TALK

IP Routing and Switching, Engineered for the AI Era

CIENA/PIER GROUP

Speaker:
Ken Abel, Solutions Architect

THE NEW REALITY

What AI Is Asking of Your Network

AI applications don't behave like traditional business software. They place fundamentally different demands on the network infrastructure beneath them.

Massive Data Movement

Training and inference move large volumes of data continuously — often between compute clusters, storage, and end users simultaneously.

Low Latency at Scale

Real-time AI inference — copilots, assistants, analytics — is sensitive to network delay. Milliseconds matter when users expect instant responses.

Unpredictable Traffic Patterns

AI workloads generate traffic bursts that are difficult to predict, requiring networks that can adapt dynamically — not just handle a fixed peak.

THE GAP

Where Most Networks Fall Short

Bandwidth & Capacity

Most enterprise and education networks were designed for human-scale traffic — email, video calls, file sharing. AI workloads can dwarf these assumptions.

Ask yourself:

Can your core and distribution layers absorb sudden 10x spikes without degrading other services?

Segmentation & Security

AI systems often access sensitive institutional data. Flat or loosely segmented networks create risk when AI tools need broad data access.

Ask yourself:

Are your network zones aligned with where AI tools will reach — and where they shouldn't?

Latency & Consistency

Inconsistent latency — jitter — is often invisible in traditional apps but highly disruptive to AI inference and real-time processing pipelines.

Ask yourself:

Do you have visibility into latency consistency across your network, not just average throughput?

Resilience & Redundancy

AI-powered workflows are increasingly mission-critical. A network outage doesn't just slow down email — it can halt operations that staff now depend on.

Ask yourself:

Is your redundancy designed for the uptime expectations that come with AI-dependent operations?

The question isn't whether AI will change your network requirements — it's whether your network is ready when it does.

BEFORE YOU PLAN

What Your Organization Needs to Consider

Understanding the technical gap is only the first step. The path forward requires honest assessment across operations, governance, and investment.

Operational Readiness

Do your network and IT teams have the skills to design, deploy, and operate AI-optimized infrastructure? Gaps in expertise can be as limiting as gaps in hardware.

Network Planning Horizon

AI adoption timelines move faster than traditional network refresh cycles. Planning must account for capacity needs 3–5 years out, not just current demand.

Governance & Policy

AI tools accessing institutional data require clear network access policies. Who governs which systems reach which data — and how is that enforced at the network level?

Investment Prioritization

Not everything needs to be upgraded at once. Prioritize investments where AI workloads are heaviest today — core switching, inter-site connectivity, and data center access.

THE RESPONSE

How the Industry Is Evolving to Meet the Demand

LAN / Campus

25G / 100GbE Access

Switches now deliver 25G and 100G to the server and compute edge — replacing 1G/10G links that can't sustain AI workload throughput.

EVPN / VXLAN Fabric

EVPN with VXLAN overlays replaces legacy VLAN spanning tree with a scalable, segmented fabric — essential for isolating AI workloads and sensitive data.

PFC / DCB — Lossless Ethernet

Priority Flow Control (802.1Qbb) and Data Center Bridging eliminate packet loss in AI training fabrics where dropped packets cause costly retransmits across GPU clusters.

WAN / Metro / Core

Segment Routing (SR-MPLS / SRv6)

SR encodes the full forwarding path in the packet header, enabling precise traffic engineering — routing latency-sensitive AI flows on optimized paths end to end.

400G / 800G Optical

Coherent optical platforms now deliver 400G and 800G per wavelength — providing the raw capacity needed to move AI datasets and model traffic between data centers and campuses.

SRv6 + Network Slicing

SRv6 combined with network slicing allows carriers to carve dedicated logical paths for AI traffic — guaranteeing bandwidth and latency independent of other network users.

The technology exists. The question is whether your network roadmap is aligned with where these standards are heading.