

Reimagining Networking: My Vision for a Secure Data Exchange with an Agentic AI Overlay

A New Era for Networking

Last week at the AI Networking Summit in Dallas, I shared a vision for the future of networking that challenges us to rethink what our industry is fundamentally about. We're no longer just moving packets between devices; we're orchestrating a secure data exchange that powers the AI era. Our current enterprise compute infrastructure—rigid, siloed, and reliant on outdated protocols like BGP—is buckling under the weight of petabyte-scale networks, exponential AI growth, and unpredictable traffic patterns. To meet these demands, we must build a flexible, autonomous system that prioritizes data flows with agility and resilience. At the heart of this transformation is an Agentic AI Overlay, a framework I believe is critical to our future, powered by four key components: the Modal Context Protocol (MCP), Agent-to-Agent (A2A) protocols, an Agent Directory, and a stateless AI Router. Yes, there will be more components as we learn more, but this is a good start. Let me explain why this matters and how we can make it a reality.

The Secure Data Exchange Paradigm

A secure data exchange redefines networking as a system that facilitates and prioritizes data movement across applications, storage, and users—much like a stock exchange, but with data as the currency. Today's networks face unprecedented complexity, with north-south, east-west, and now unpredictable traffic patterns as data gravity has expanded from clouds and data centers to now the edge and AI data centers, driven by AI workloads that generate erratic latency and processing demands. The intersection of AI and cloud workloads generates “unique patterns, processing, and latency requirements” that invert the expectations of how networks should function. Unlike traditional routing protocols that focus on source and destination addresses, a data exchange must optimize flows based on a wide range of metrics: security posture, cost, latency, congestion, and real-time patterns. This requires a shift from rigid infrastructure to one that's autonomous yet controlled, capable of making decisions in milliseconds to enhance user experience and mitigate risks like the NotPetya virus, which cost FedEx and Merck half a billion dollars in minutes each.

The Agentic AI Overlay: A Dual-Service Framework

The Agentic AI Overlay is the backbone of this vision, delivering two critical services. First, it offers a new operational model by moving AI closer to the data—system notifications, logs, alarms, and alerts. Instead of funneling this data into centralized “haystacks” like Splunk,

Data Dog, Snowflake, S3 buckets etc where engineers hunt for needles among thousands of alerts, distributed AI agents process this data locally. These agents, deployed across owned or leased facilities, communicate directly with NOC and SOC systems, slashing response and resolution times. Second, these agents guide data movement under the direction of a stateless AI Router, ensuring optimal flow across the infrastructure. This dual role—operational intelligence and data orchestration—makes the overlay a game-changer.

Key Components of the Agentic AI Overlay

Let's break down the components that make this possible:

- **Modal Context Protocol (MCP):** MCP integrates applications, agents, tools, and AI-driven enrollment processes into a unified ecosystem, facilitating communication across the application dependency map. It ensures cohesive operation for agents and AI applications, whether in your data center or a cloud provider's environment.
- **Agent-to-Agent (A2A) Protocols:** These protocols enable secure, trusted communication between distributed agents, allowing them to share insights and coordinate in real time. This eliminates the inefficiencies of centralized data lakes and supports rapid incident response.
- **Agent Directory:** This acts as a registry, cataloging all agents and their capabilities, so the system can orchestrate resources efficiently.
- **Stateless AI Router:** This revolutionizes path selection by moving beyond BGP's sluggish 90-second convergence times with an updated millisecond reconfiguration rate. It uses AI to make routing decisions based on diverse metrics—security, cost, latency, and more—ensuring data flows optimally in dynamic environments.

Why This Matters: The AI Revolution Demands a New Network

By 2030, AI data centers could be as large as 5,000,000 to 10,000,000 GPUs powered by 10-100 Gigawatts to a Terawatt of power; we can remove regulations on power generation. With that kind of scale in just 5 years, AI data centers will be a black hole in which very few IT applications will escape. Applications will be both hosted in these massive compute centers or have access to them where applications gain AI functionality via MCP. This is a huge and massive tsunami heading our way. We had lift and shift in the cloud era; in the AI era, we will have AI rewriting all of our system of record applications and fundamentally transforming business. We need a new network for this era. Without a secure data exchange powered by an Agentic AI Overlay, our organizations risk falling behind. John Chambers said at the Dallas Summit that 70% of Fortune 500 companies will not be on that list if they don't embrace AI. These components—MCP, A2A protocols, Agent Directory, and stateless AI Router—provide the foundation for a network that's autonomous, controllable through policy engines, and aligned with business logic, compliance, and regulation.

Building the Future Together

The future of AI-driven networking is already unfolding. At the AI Networking Summit in Dallas, enterprises like eBay, Target, and Walmart showcased SONiC-based, flat, high-capacity data exchanges with AI agents positioned near the data. To advance this vision, the ONUG Collaborative is launching a working group to define the requirements for an Agentic AI Overlay. One key output of this working group will be an Agentic AI Overlay demonstration at the AI Networking Summit in NYC this fall. I invite all vendors to participate in this collaborative effort. This is a critical opportunity to shape the future of AI networking and prepare for the AI-driven era, benefiting our entire industry. Time is of the essence. To learn more, join us on Monday June 23rd at noon ET for a kick-off and information session zoom meeting. Click [here](#) to register. Let's seize this moment to redefine networking and build a secure data exchange for the Enterprise AI era.