



# Cross-Connect or Crossed Wires?

Why AI Will Break Your Network Before It Breaks Even.

*The Rise of East-West Traffic and Why Your Network is the Weakest Link in Your AI Stack*

**PTC'26 -- HUB Lightning Talk**

Travis Ewert – Global Head of Platform Development

January 19, 2026



**6**  
Continents

**25+**  
Countries

**50+**  
Metros

**300+**  
Data Centers

**5,000+**  
Customers

**231,000**  
Cross Connects

**1,300+**  
Global Carriers (Carrier Neutral)

*Connectivity Ranking*

**#1**  
In Europe,  
LATAM & MENA

**#2**  
In NAM

**#4**  
In APAC on Cloudscene's Data  
Center Ecosystem Leaderboard

**2.8 GW**  
In-Place IT Capacity  
(1.5 GW renewable)

**750 MW**  
Under Construction

**>4 GW**  
Future development  
capacity

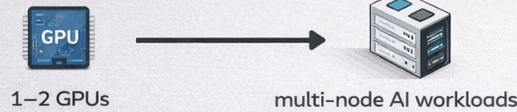
Note: As of October-2025



# GPU Rich... Network Poor

*Chips secured. Densified Space/Power/Cooling in place. Did you get the Network right?*

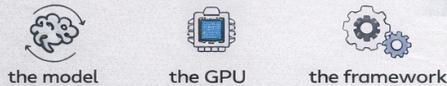
## The Reality Most Teams Discover Too Late



They hit problems that don't show up in proof-of-concepts:

- ✗ GPUs waiting on data
- ✗ Inference latency spikes
- ✗ Training jobs stalling unpredictably
- ✗ Costs rising without throughput gains

Most enterprise AI performance problems aren't caused by:



They're caused by how GPUs talk to each other and to storage.

## Where the Real Bottlenecks Live

Traditional data center networks were designed for:

- north-south traffic
- client/server apps
- bursty workloads

That mismatch is where performance dies.

AI is:

- continuous
- latency-sensitive
- bandwidth-hungry

## What Breaks First in the Real World

- Oversubscribed leaf-spine designs
- Insufficient RDMA / low-latency tuning
- Monitoring tools that don't see GPU stalls
- Networks not sized for AI-scale ops.

## The Takeaway



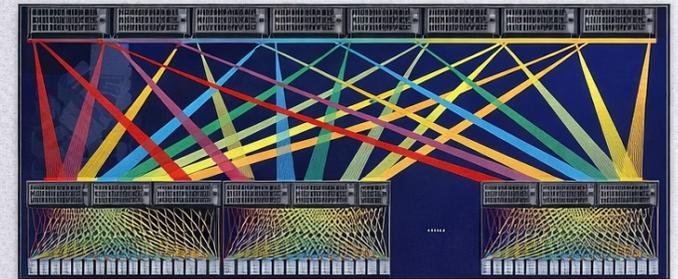
If your AI architecture plan says:  
**“We'll figure out networking later”**



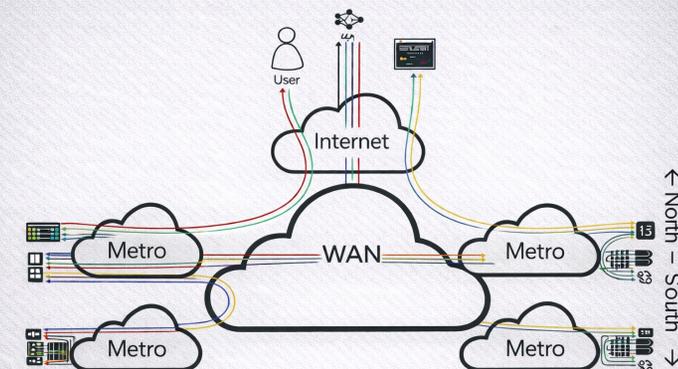
You're not building an AI platform.  
You're building an expensive waiting room for GPUs.

**AI success is determined between the servers, not inside them.**

If you're planning multi-GPU AI and networking isn't a first-class design concern yet — that's the conversation to have before hardware arrives.



Ex. 8K+ node GPU Cluster

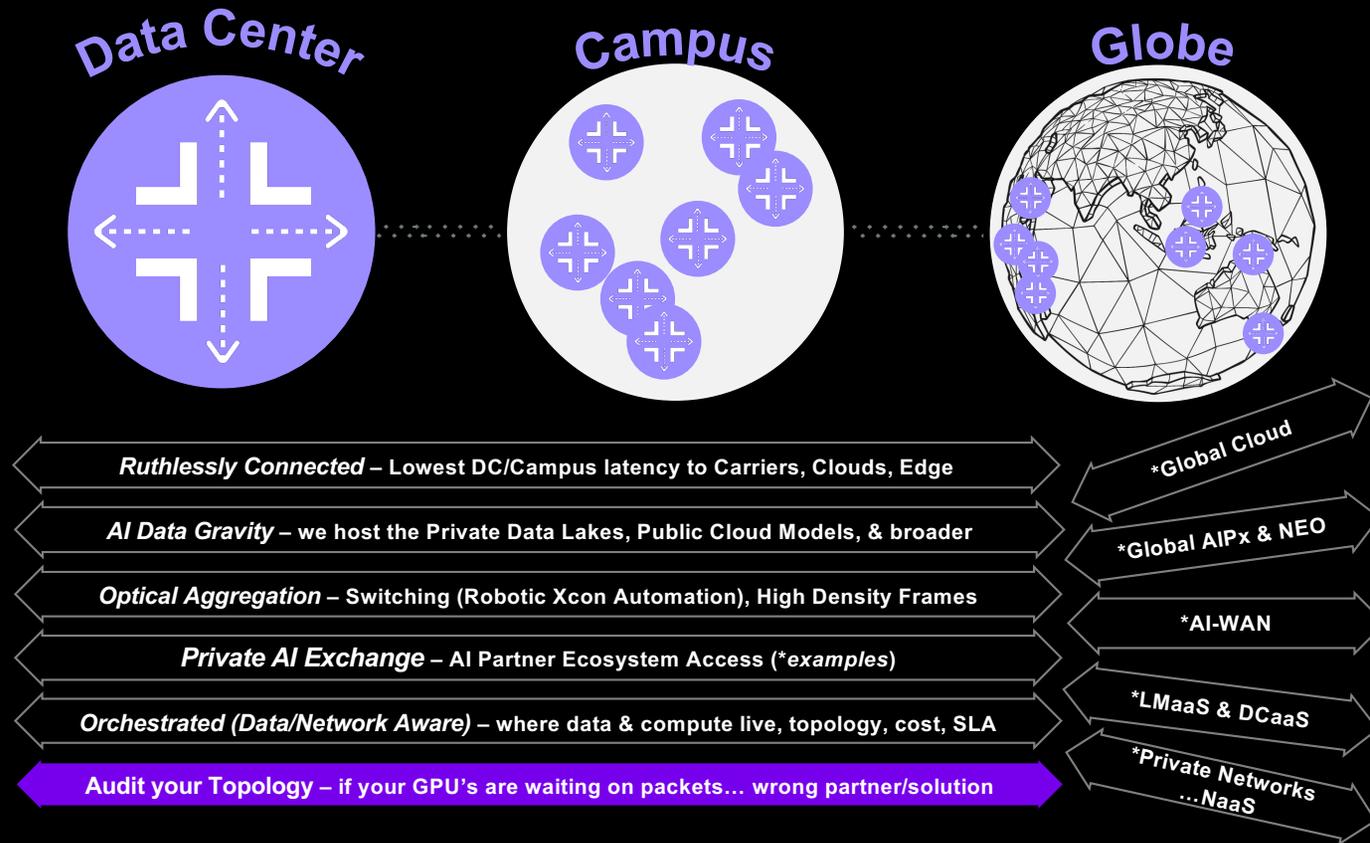


← West — East →

← North — South →

# AI Network Deployment Variants

Scale across all dimensions... while ensuring performance, security, sovereignty, orchestration



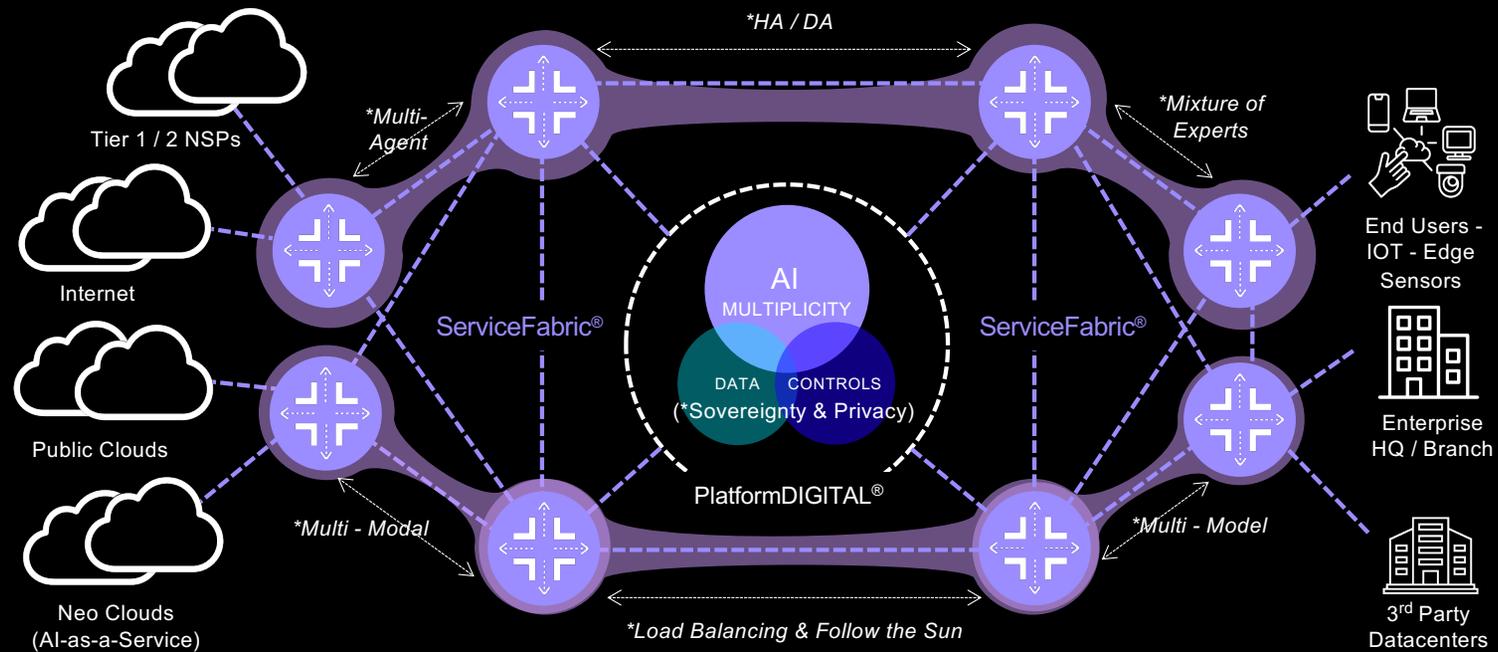
# Intelligent workflows require AI multiplicity

ServiceFabric® support of AI Inference and Distributed Compute

Workflows across points of presence on PlatformDIGITAL®

KEY

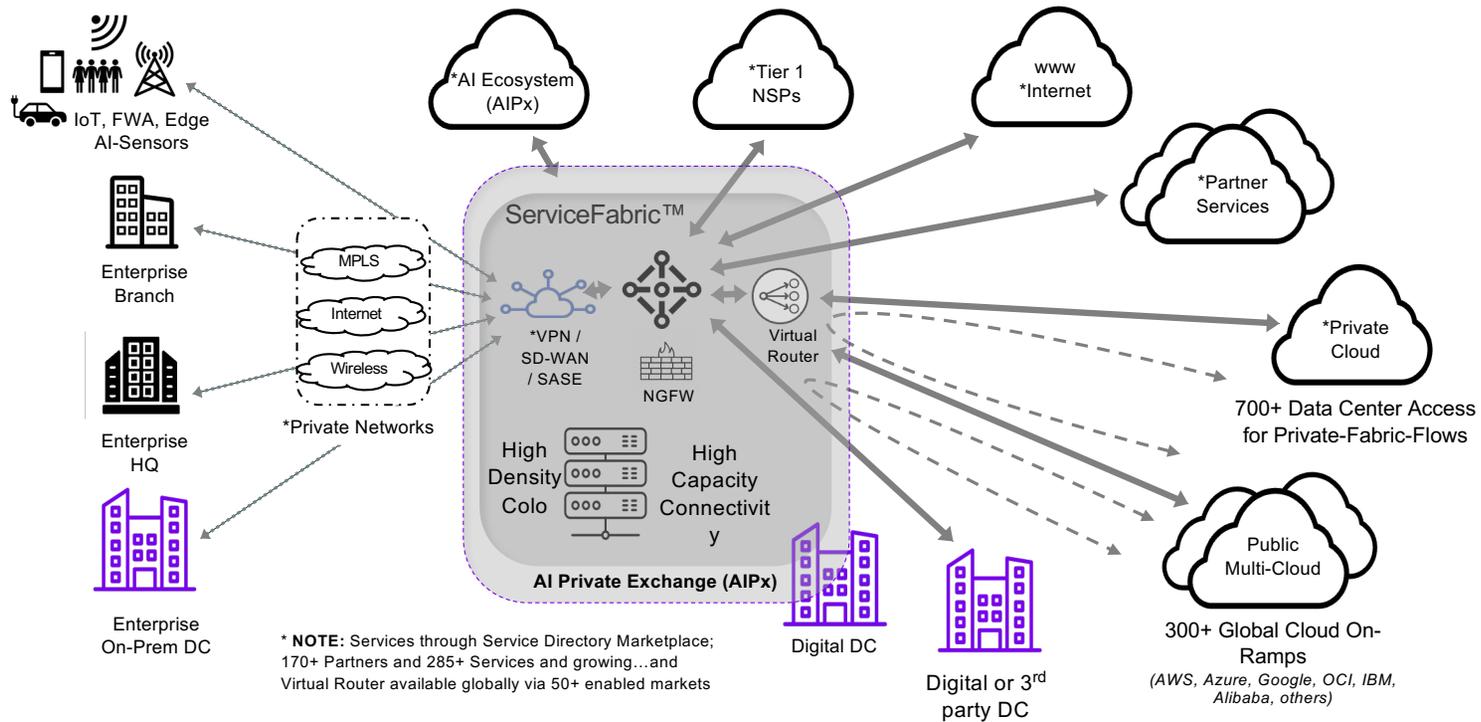
-  AI Inference Points of Presence
-  ServiceFabric® Interconnection
-  Site Selection Criteria



\*NOTE: Inference or Distributed-Compute use case solutioning examples – extended across any site or in any combination.

# ServiceFabric® Core Framework – Partner Ecosystem

Global Interconnect Platform | Multi-Cloud | Hybrid AI | Private AI Exchange |  
Open, Neutral, Fabric-of-Fabrics for On-Demand Services



**AI Private Exchange:** Industry First Announcement – spring of 2024. Patented technologies for Orchestrated SDN Services (Hybrid/Multi-Cloud, AI Onramps, on-demand AI Exchange, broader)