Programmable by Design: Building Open Networking That Evolves at the Pace of AI

OCP Educational Webinar Series

November 10th, 2026

Community-driven hyperscale innovation for all

Programmable by Design:
Building Open Networking That Evolves at the Pace of AI



Cliff Grossner, Ph.D.

Chief Innovation Officer
OCP Foundation



Sameh Boujelbene

VP, AI Networks
Dell'Oro group



Shane Corban

Senior Director of
Technical Marketing
AMD

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking is Critical for AI Innovation
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking Critical for Open Ecosystems
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

Different Stages of the AI Journey

1	2	3	4
Generative AI	**Agentic AI**	**Multi-Agent AI**	**Physical AI**
CREATE	ACT	COLLABORATE	INTERACT
1× token intensity Prompt → response	**~4×** token intensity Tools + workflows	**~15×** token intensity Agents talk to agents	**Real time** token intensity Sense → reason → act

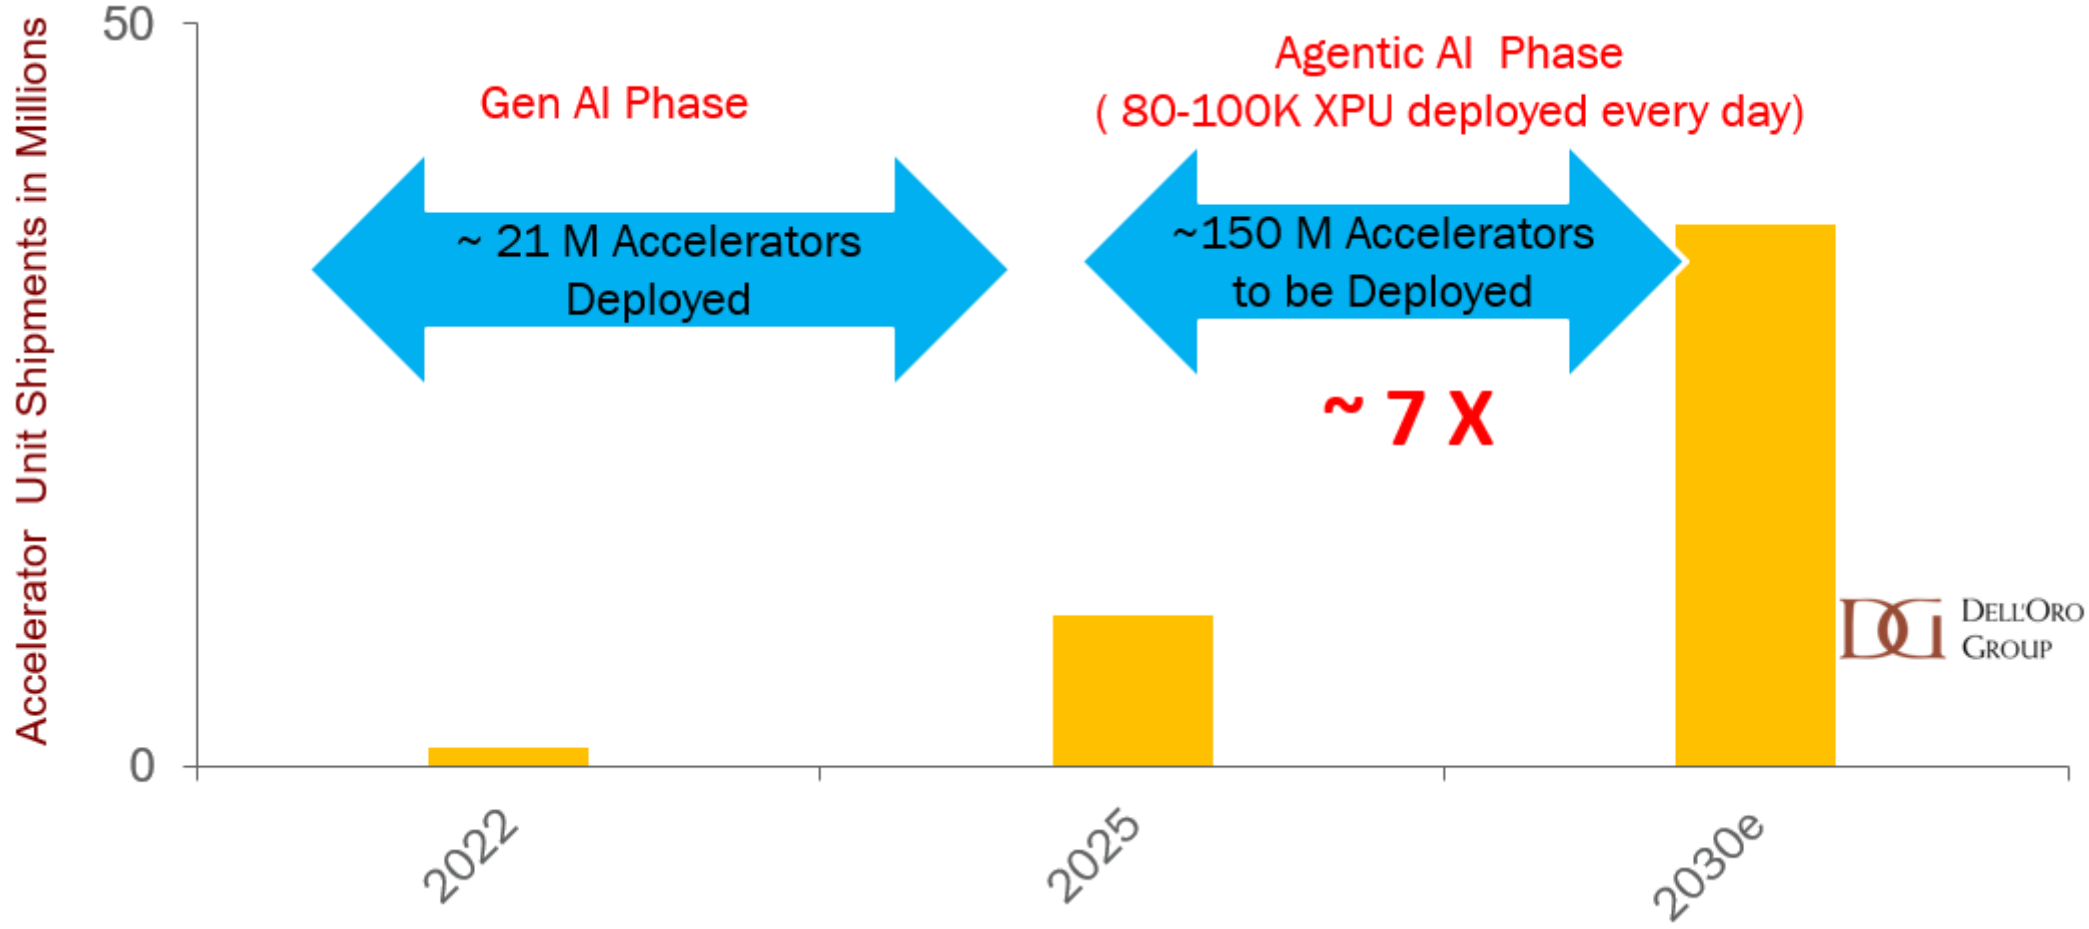
2022-2025 **2026+** **2028+**



- Various stages of AI journey
- Each stage drives new requirements
- Compute demand and network intensity increasing with each stage

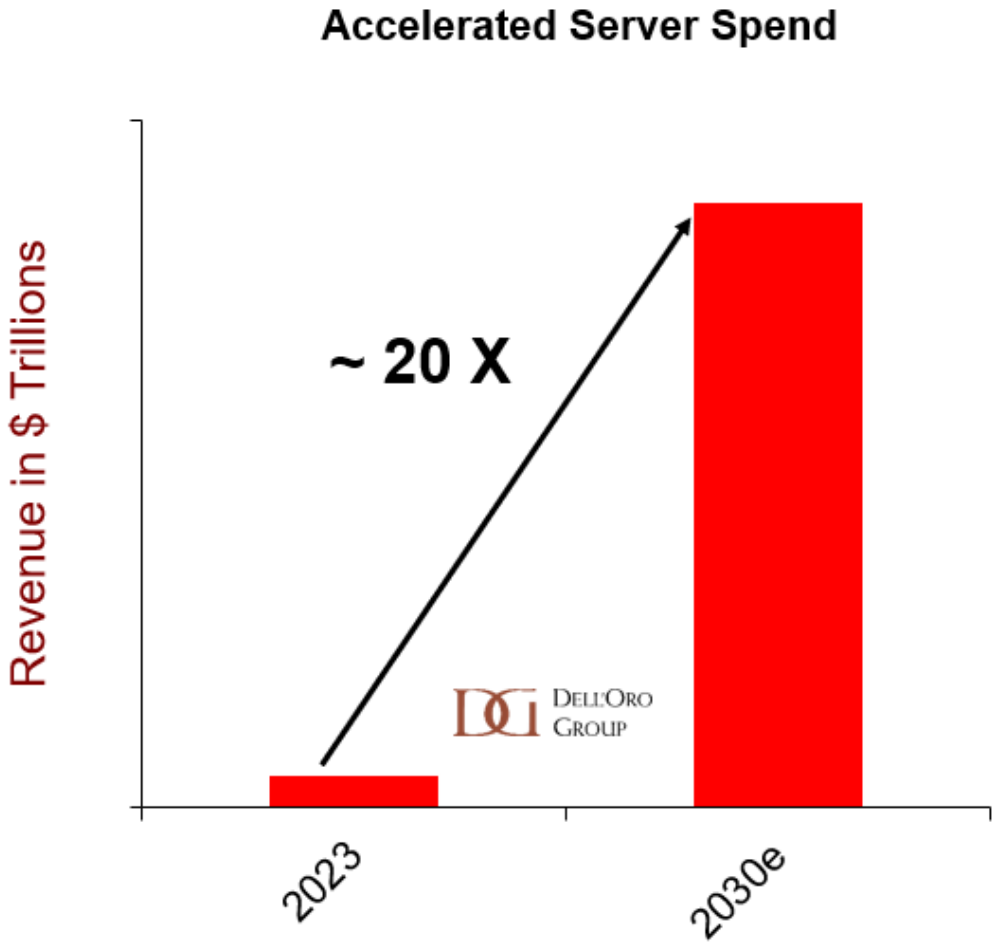
Agentic AI to Enable a 7X Expansion in Accelerator Deployment

The next five years will shift the market from tens of millions to 150 M XPUs

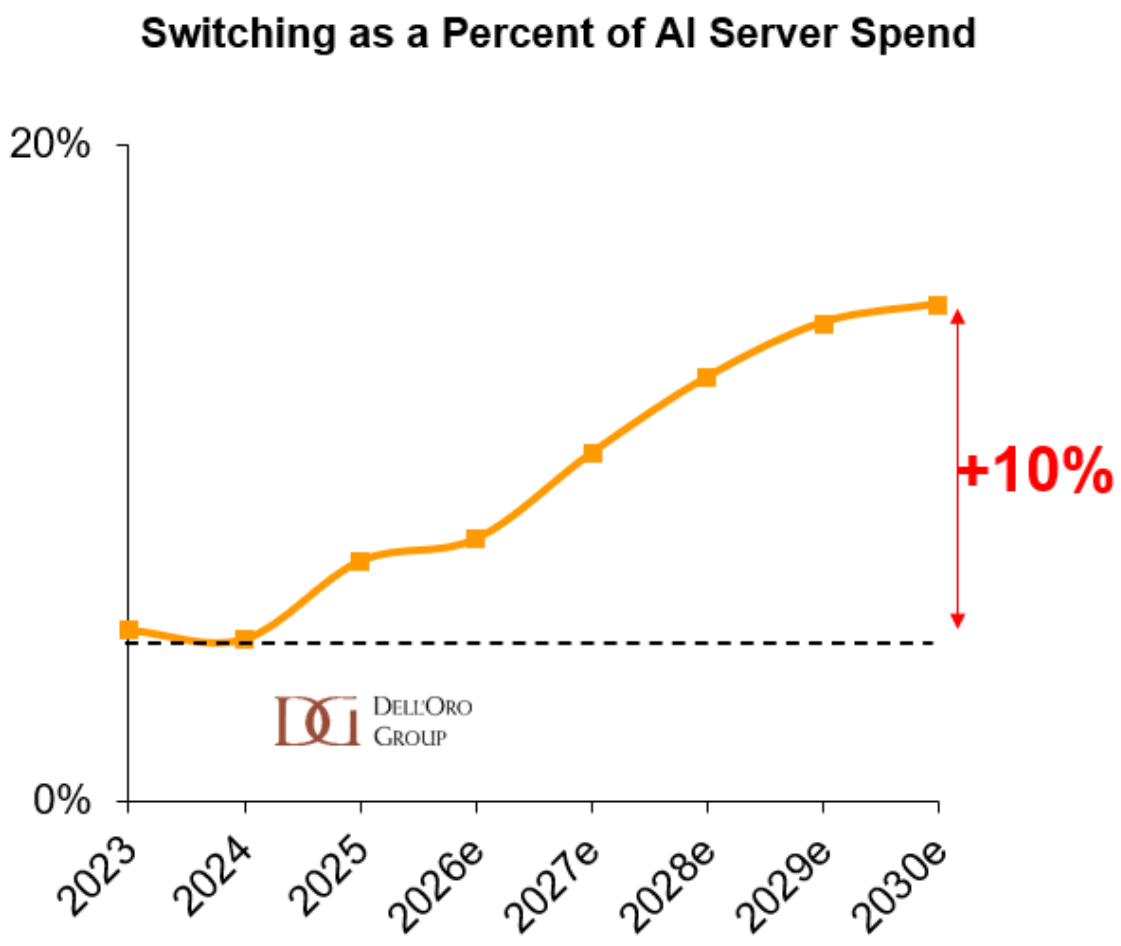


Source: Dell'Oro Data Center IT Semiconductors and Components Report

AI Server Spend is Rising Rapidly but Network Spend is Rising Even Faster



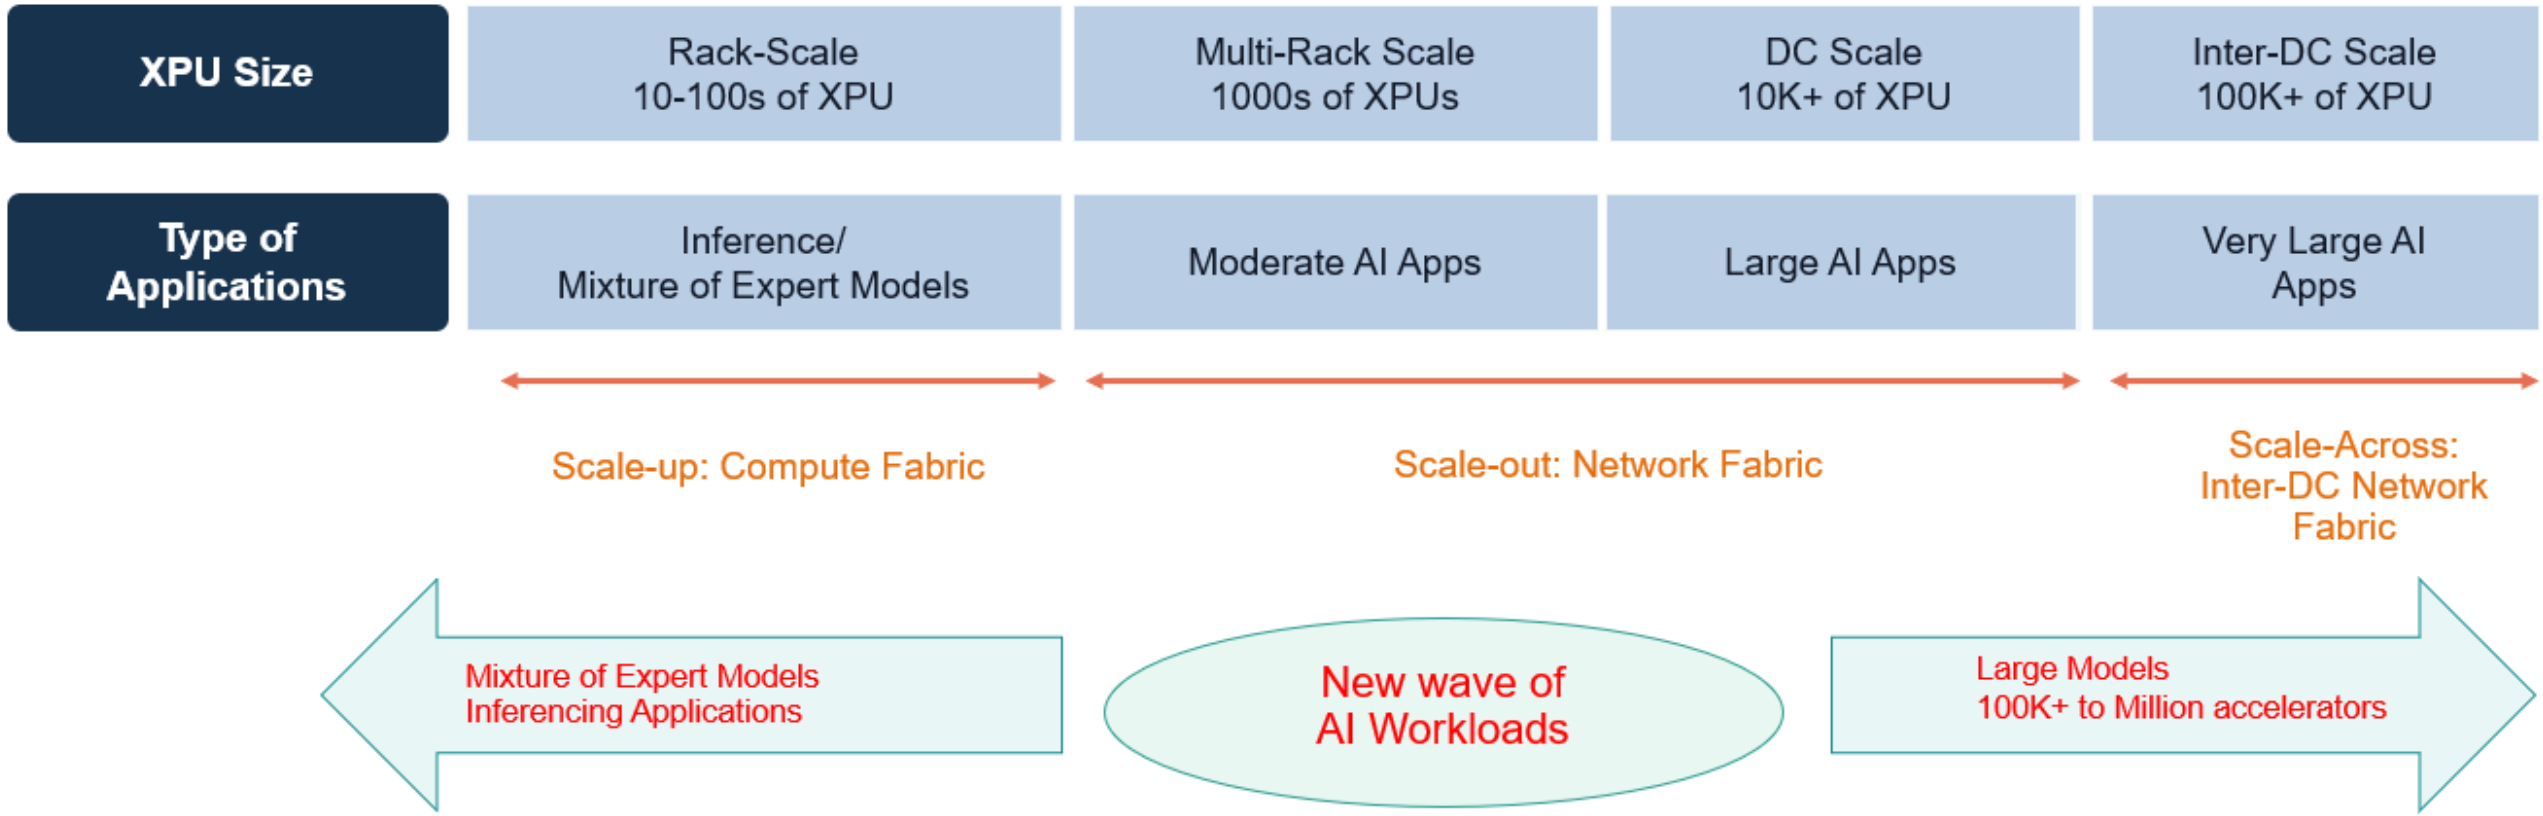
Source: Dell'Oro Data Center IT Capex Report



Source: Data Center Switch – AI Back-end Networks

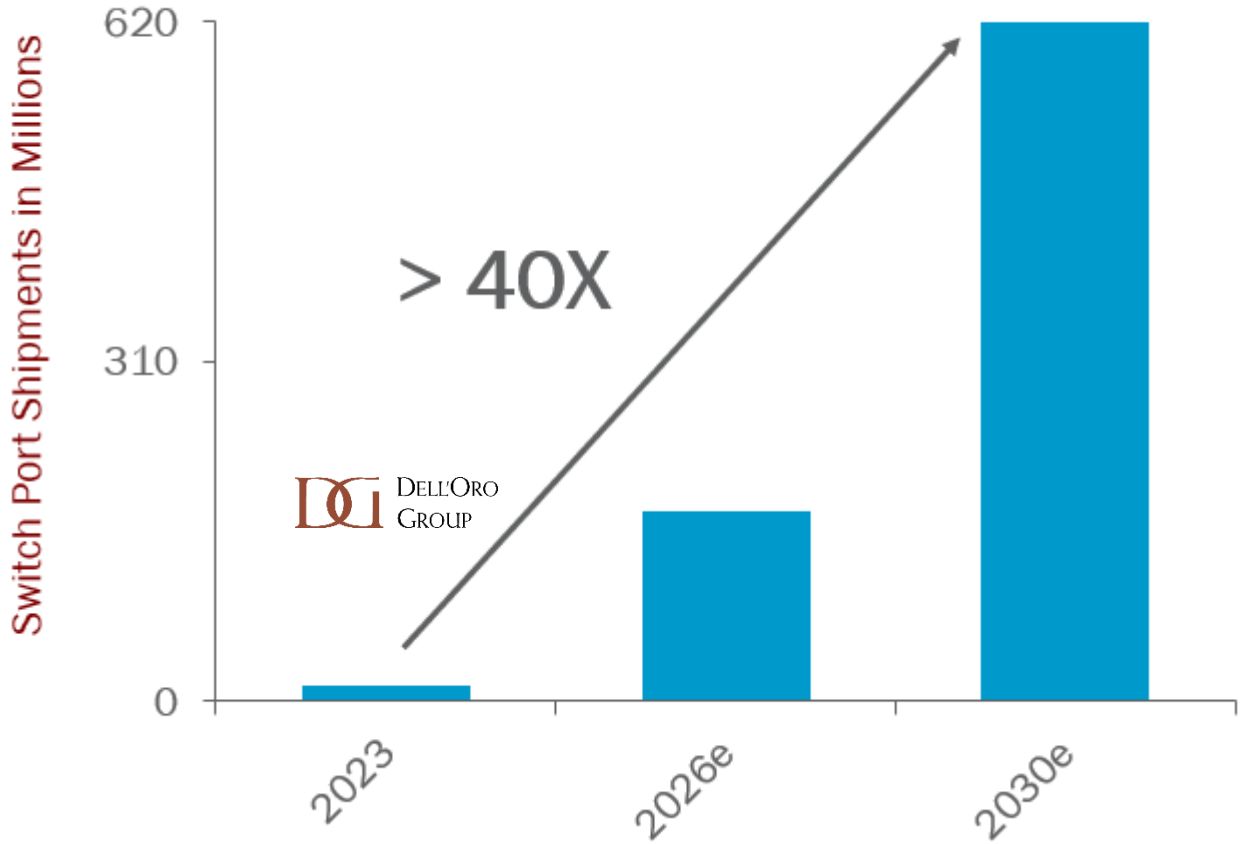
Dell'Oro Group

Back-end Networks Expanded Across Multiple Domains to Meet Evolving AI Workload Requirements



- Source: Dell'Oro Data Center Switch – AI Back-end Networks Report

Data Center Switch Port Shipments in AI Back-end Networks



- Exponential Increase in volume and scale in less than a decade
- Supply chain not prepared for such expansion in such a short amount of time
- Expect demand to continue to surpass supply over the next few years

Source: Data Center Switch – AI Back-end Networks

Poll Question 1:

Q: Which AI networking domains are you currently operating in?

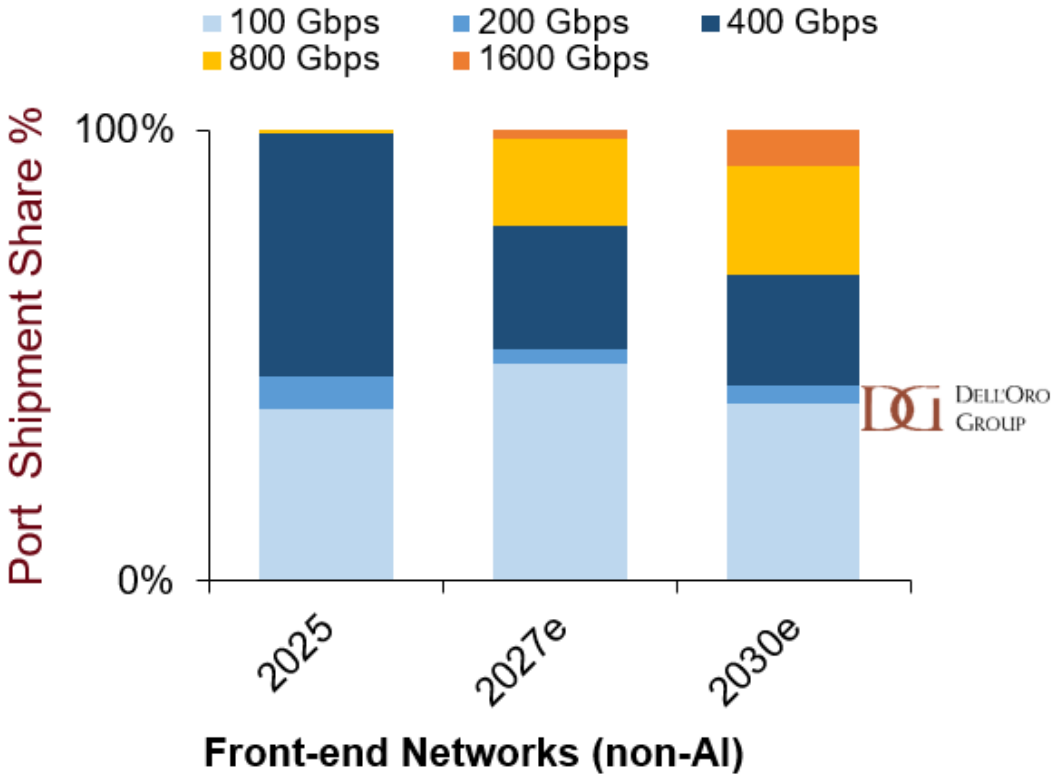
Responses
A. Front-End
B. Front-End, Scale-Up
C. Front-End, Scale-Up, Scale-Out
D. Front-End, Scale-up, Scale-out, Scale-Across

Audience Poll Results

- Some discussion facilitated by host and analyst of the results with optional comments from sponsors

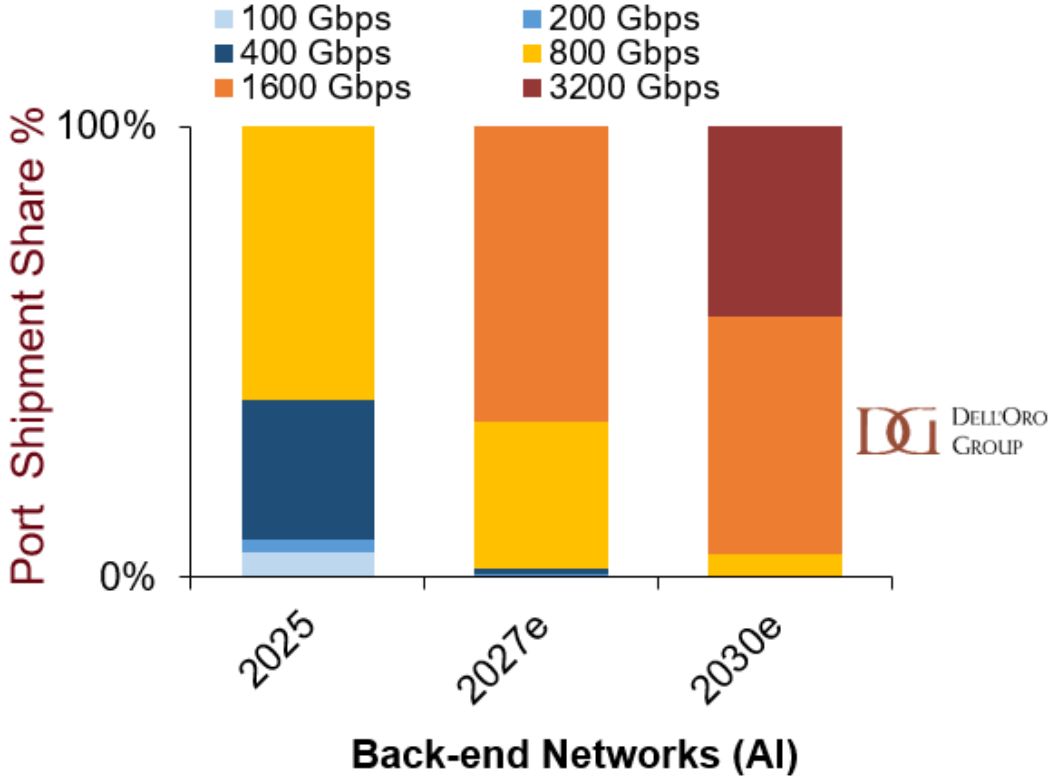
Back-end Speed Upgrade Cycles are Happening on Fast Forward

Back-end networks are two speed upgrade cycles ahead of the Front-end networks



Front-end Networks (non-AI)

- Compute-bound
- Low network utilization rate (average of 50%)
- Refresh cycle is 5 years or more



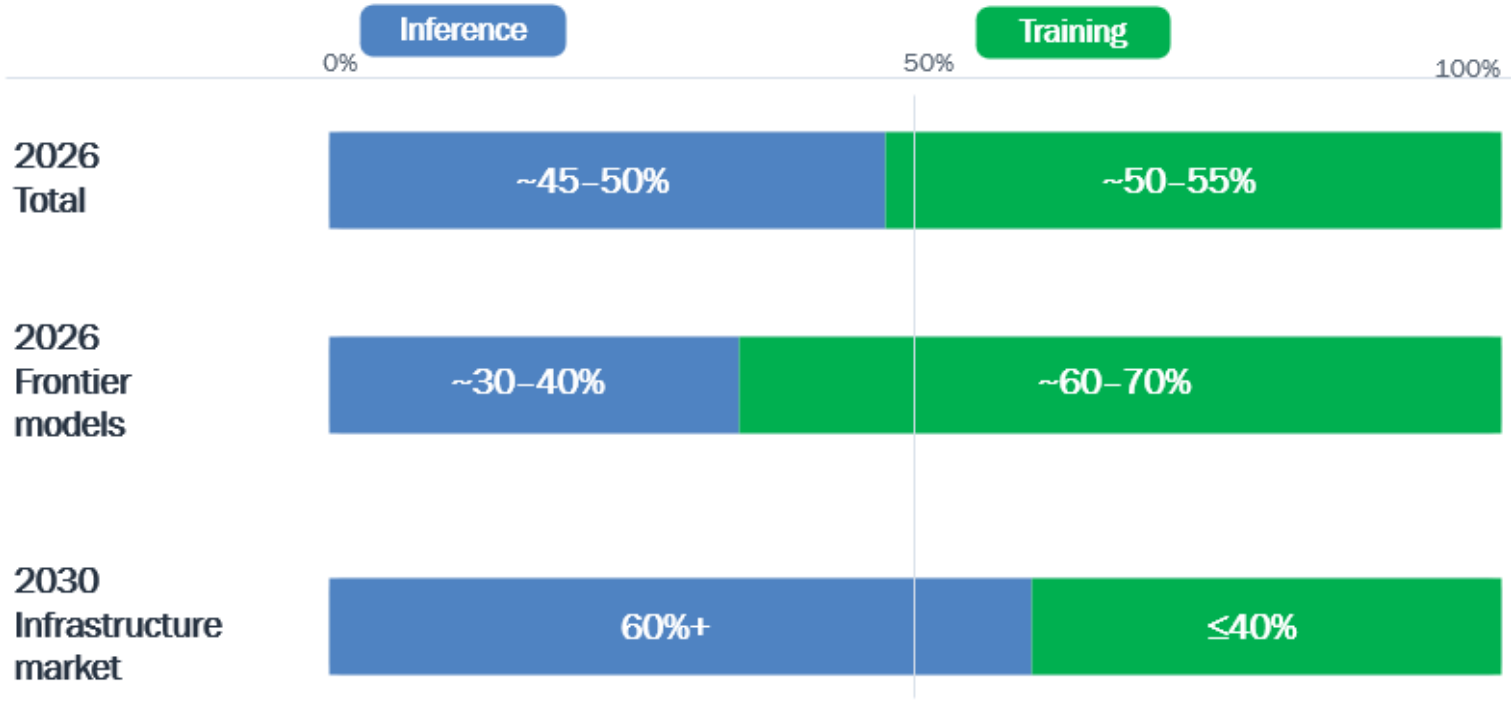
Back-end Networks (AI)

- Network bound
- Network needs to operate at nearly 100%
- Refresh cycle is 18 to 24 months

Dell'Oro Group

Inference Becomes the Dominant Driver of AI Infrastructure

AI back-end network architectures must optimize not only for training clusters but also for large-scale inference and agentic AI traffic patterns.



Key takeaways

Market remains roughly at parity between inference and training.

Frontier models remain more training-heavy than the broader market.

Inference becomes the dominant driver of infrastructure spending.

AI Back-End Networks: Key Issues & Requirements

AI infrastructure requires networks that are scalable, open, adaptable, and efficient.

Key Issues

What is changing in AI back-end networks

- Workload diversity across training, inference, agentic AI, storage, and collectives
- Rising cluster scale and ongoing supply constraints
- Scale-up, scale-out, and scale-across require different network behaviors
- Innovation cycles are moving faster than silicon refresh cycles
- Higher speeds increase power, cost, and technical complexity

Resulting Requirements

What the network must deliver

- Adaptable network architecture
- Open standards and ecosystem support
- Programmability for rapid feature evolution
- Supply diversity for resilience and faster innovation
- Better power and cost efficiency

Why Ethernet Is the Ultimate Choice for AI Back-End Networks

Open, multi-vendor Ethernet is becoming the long-term architecture for AI back-end networks

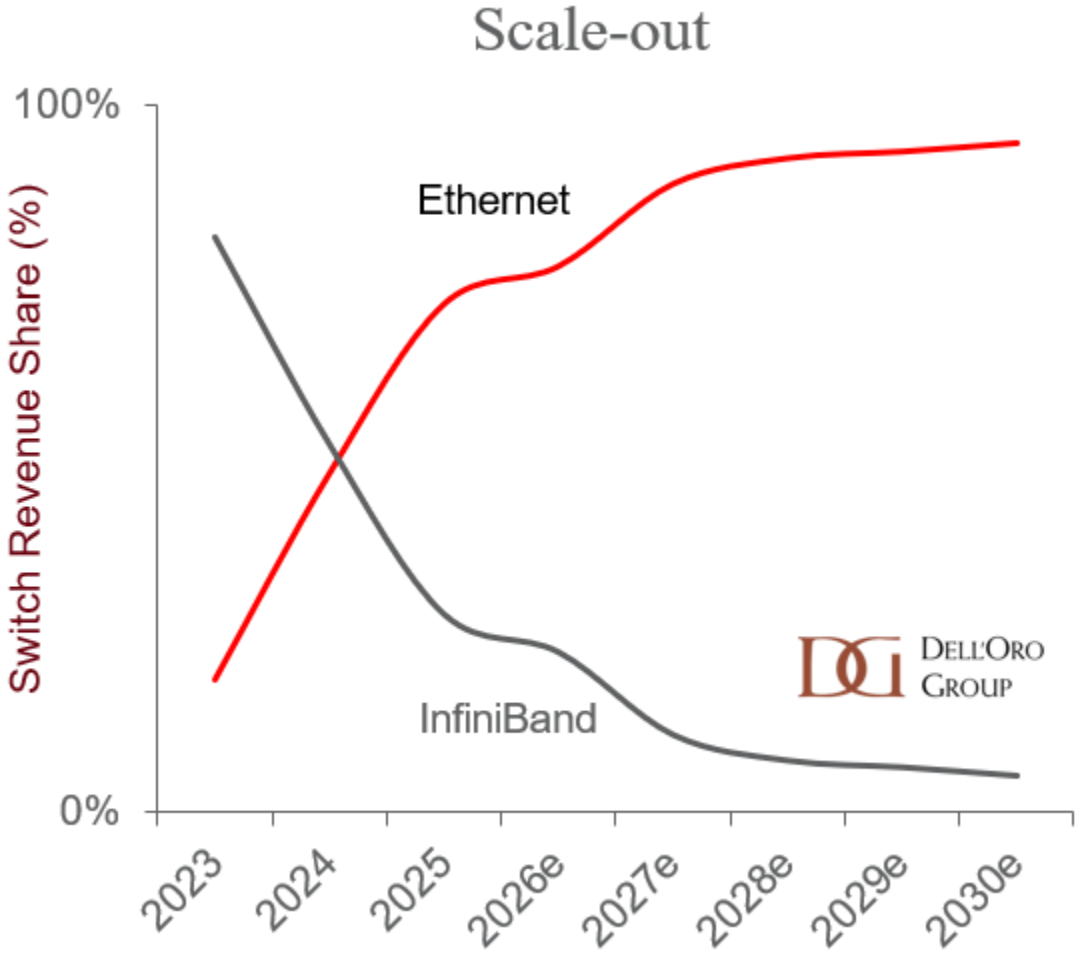
Why Proprietary Approaches Hit a Wall

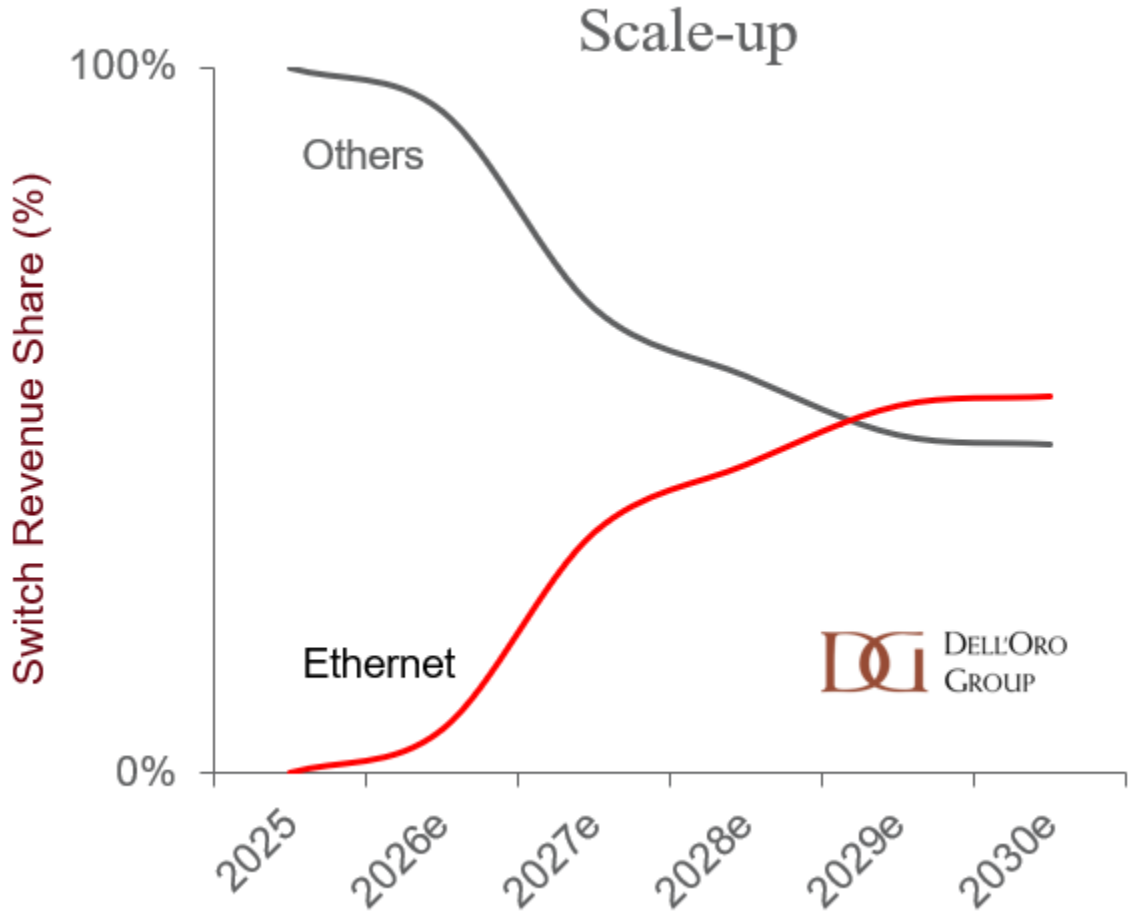
- Vendor lock-in raises cost and supply risk
- Closed ecosystems struggle to keep up with AI infrastructure needs
- Accelerated upgrade cycles require faster innovation
- Fragmented approaches make scaling harder across evolving AI workloads

Why Open Ethernet Wins

- Open standards enable supply diversity and reduce lock-in
- A broad ecosystem accelerates innovation and time to market
- Programmability supports fast-changing AI requirements
- Momentum is building around UEC and Ethernet for scale-up
- Common foundation for scale-out, scale-across, and emerging scale-up

Ethernet is Winning Everywhere:
Scale-up, Scale-out and Scale-across





Source: Dell'Oro Data Center Switch – AI Back-end Networks Report

Poll Question 2:

Q: Where do you plan to run Ethernet within your AI network domains?

Responses
A. Front-End Only
B. Front-End, Scale-Up
C. Front-End, Scale-Up, Scale-Out

Audience Poll Results

- Some discussion facilitated by host and analyst of the results with optional comments from sponsors

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking is Critical for AI Innovation
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

AI Demands a New Networking Foundation



AI scales at the speed of data movement

AMD

AI is Evolving Faster Than Networks Can Adapt

Bandwidth and Expertise Can't Keep Pace

Adoption outpaces deployment and operational readiness

Endpoints Must Handle More Complex Processing

Network, storage, security and transport demands are growing

Hardware Cycles Lag AI Innovation

New workloads and protocols emerge between refreshes

More bandwidth alone cannot keep pace with AI innovation

AMD

AI Workloads Expose Different Network Bottlenecks

Front-end

Service and context access delays AI responses



Scale-up

Communication delays and link failures stall tightly coupled GPUs



Scale-Out

Congestion and uneven traffic distribution slow distributed jobs



Scale Across

Distance and bandwidth constraints limit inter-site performance

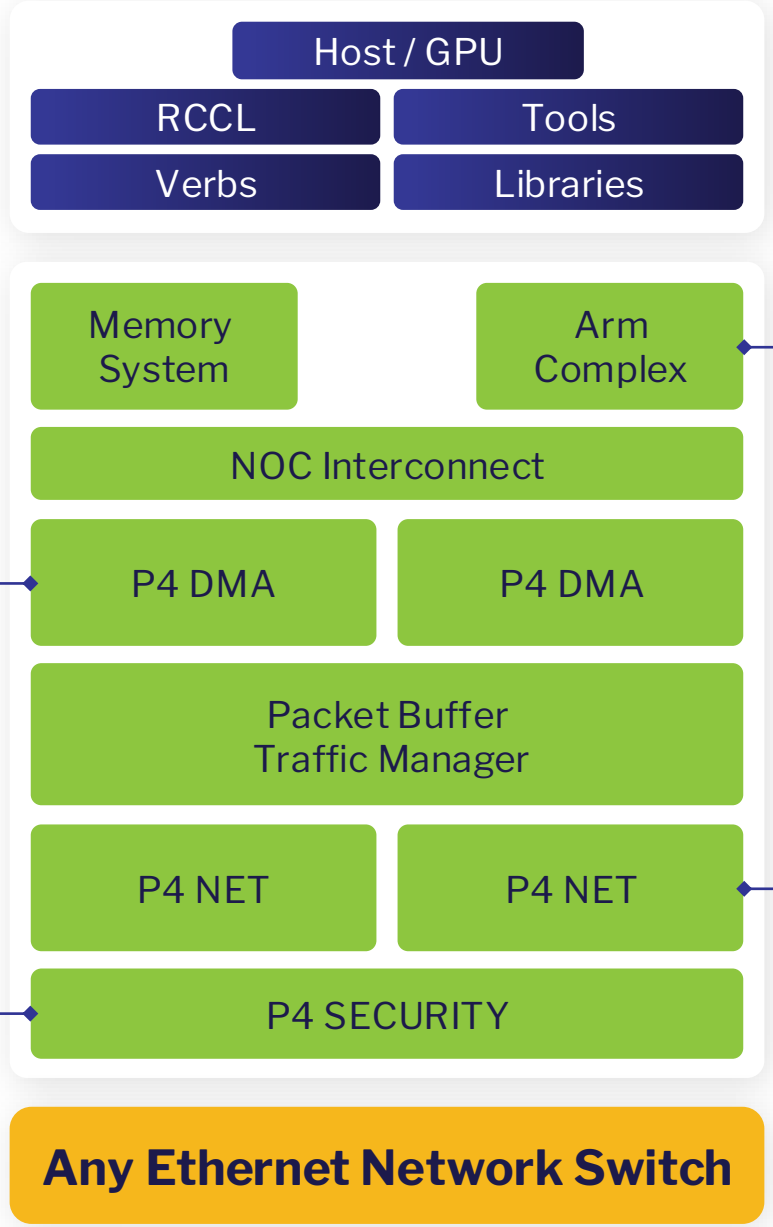


AMD

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking is Critical for AI Innovation
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

Modular P4 System Architecture



Storage & Data Path Acceleration:

- RoCEv2, UEC Ready RDMA & Custom Transports
- Address Translation Services (ATS)
- NVMe Virtualization, GDS, VirtIO-Net

Crypto & Security Functions:

- Stateful firewall*, Encryption, IPSec, PSP

Infrastructure Software & Services:

- Management Software
- Control Plane

Software Define Networks & Classic Offload:

- Overlay, TSO, Checksum, VLAN, VXLAN, Geneve
- SR-IOV Dynamic Switching
- Enhanced Broadcast/Multicast

AMD

The Expanding Role of DPUs in AI

DPUs preserve CPU capacity and accelerate secure data access as AI workloads grow.

DPU

Data Processing Unit

A programmable processor that takes infrastructure work off the host CPU for front-end and cloud networking.

Programmable services

Networking	Traffic management
Security	Isolation and encryption
Storage	Data access and movement

Add or adapt services in software

01 **Cloud infrastructure**

Offload shared services and isolate tenants.

02 **AI training**

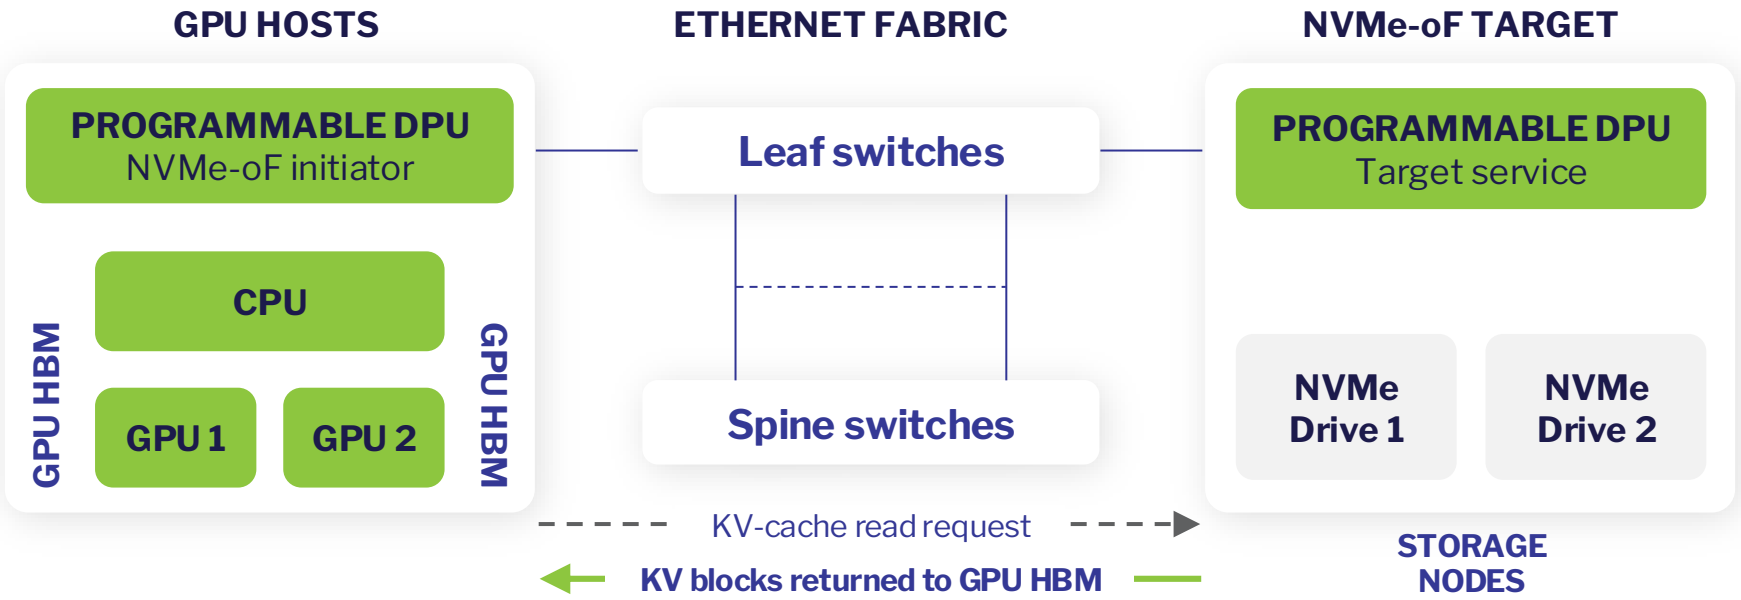
Accelerate data access and manage front-end traffic.

03 **Distributed Inference and Agentic AI**

Support more service interactions, persistent context storage, and overall acceleration of inference.

DPU Use Case for Agentic AI

- Free up GPU HBM capacity by extending KV cache to NVMe storage
- Offload KV-cache data movement through the programmable DPU data path
- Enable low-latency access to disaggregated NVMe storage over Ethernet



AMD

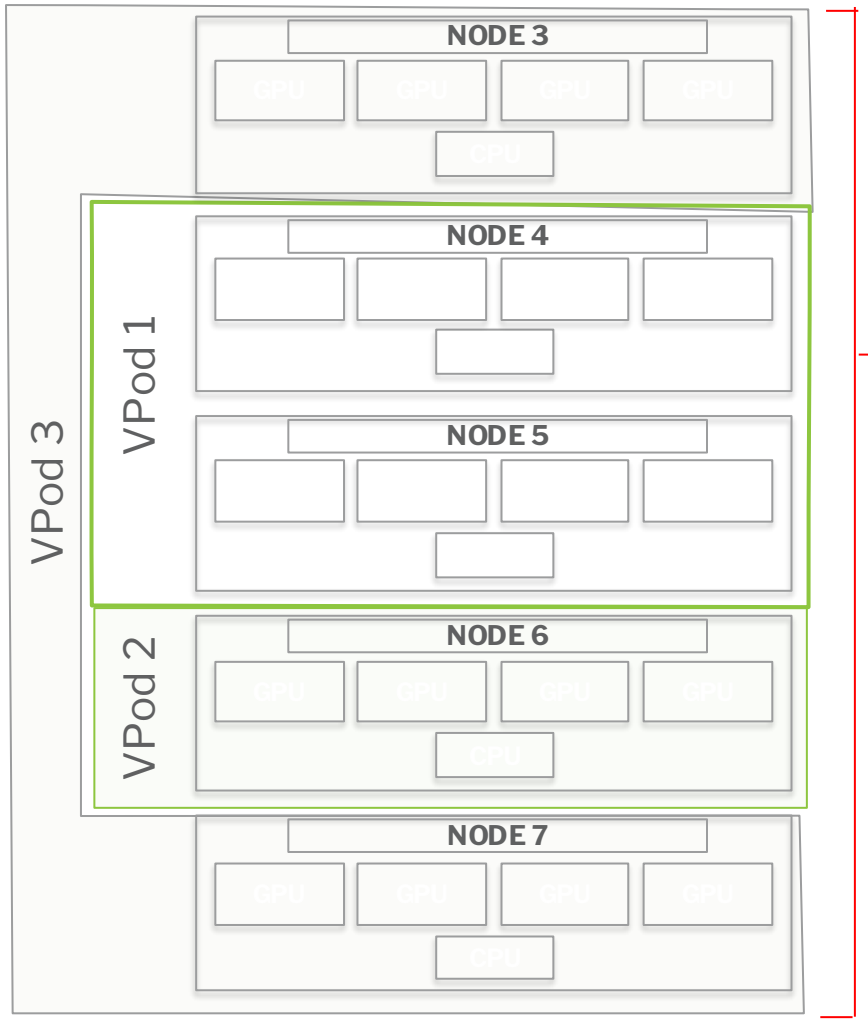
Scale Up: Virtual Pods for Flexible Partitioning

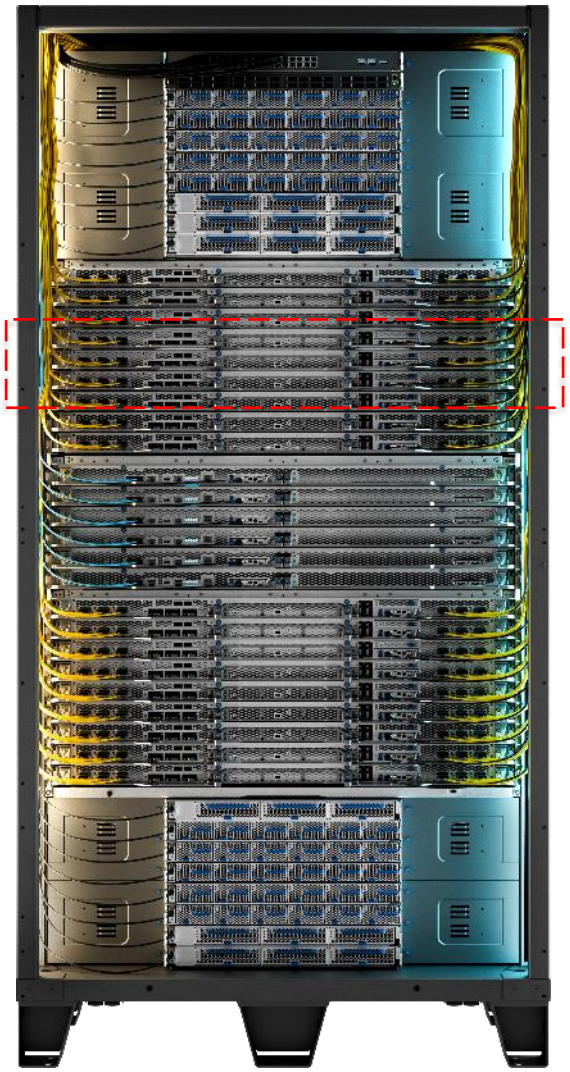
KEY TRAITS OF VIRTUAL PODS

FLEXIBLE: Combine subset of nodes to form a Virtual Pod

SECURE: VPods are isolated from each other

RESILIANCE: Node level failures are contained to a VPod





Scale-Up Requires Fast Communication and Resilience

Tightly Coupled GPU Work Fuels Tensor Parallelism

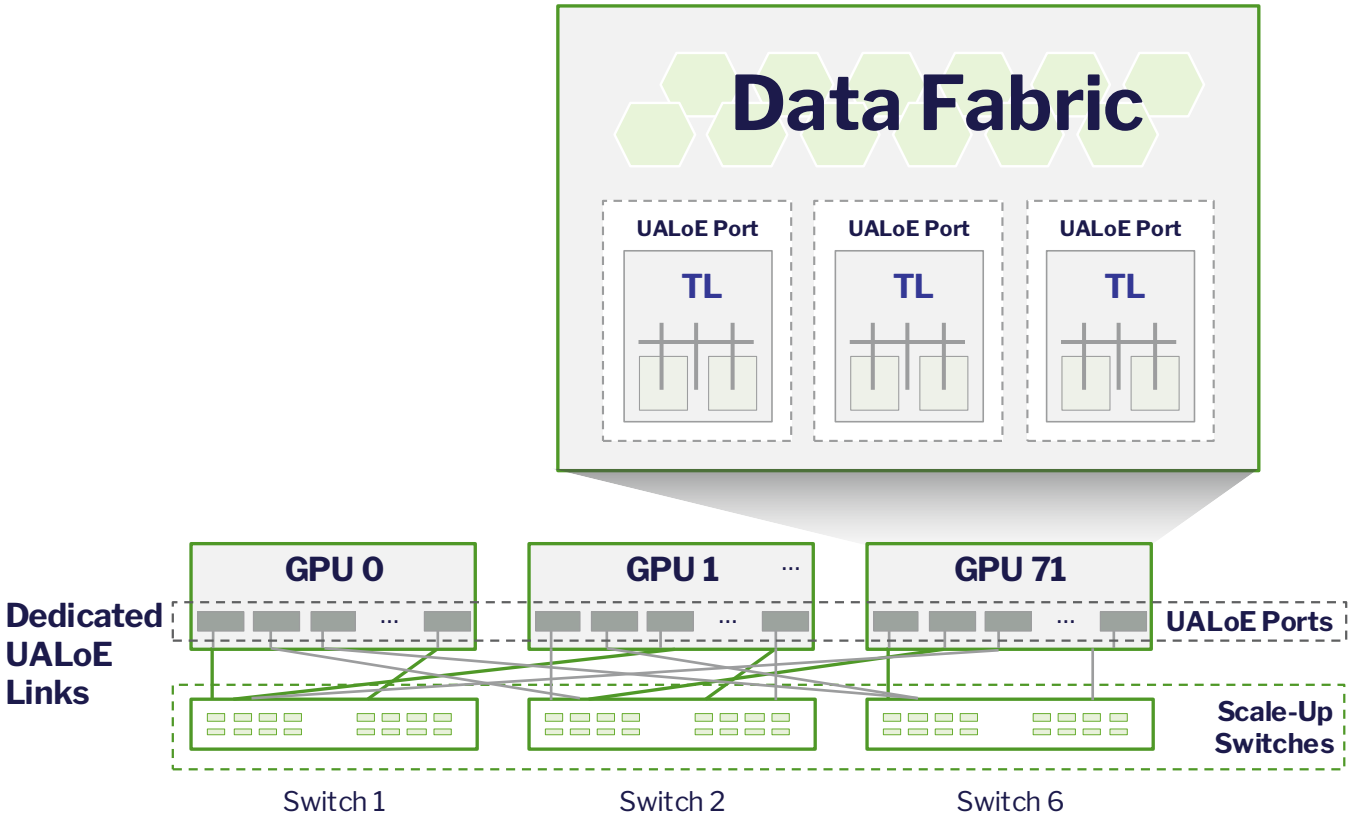
- Model partitions repeatedly exchange data
- Communication delays can stall participating GPUs
- GPUs directly read, write and run atomic operations

Redundant Communication Paths

- Multiple planes provide alternate paths
- Failure detection and transport recovery determine whether communication continues
- Multiple planes for single-hop paths between GPUs

Fault Isolation and Recovery

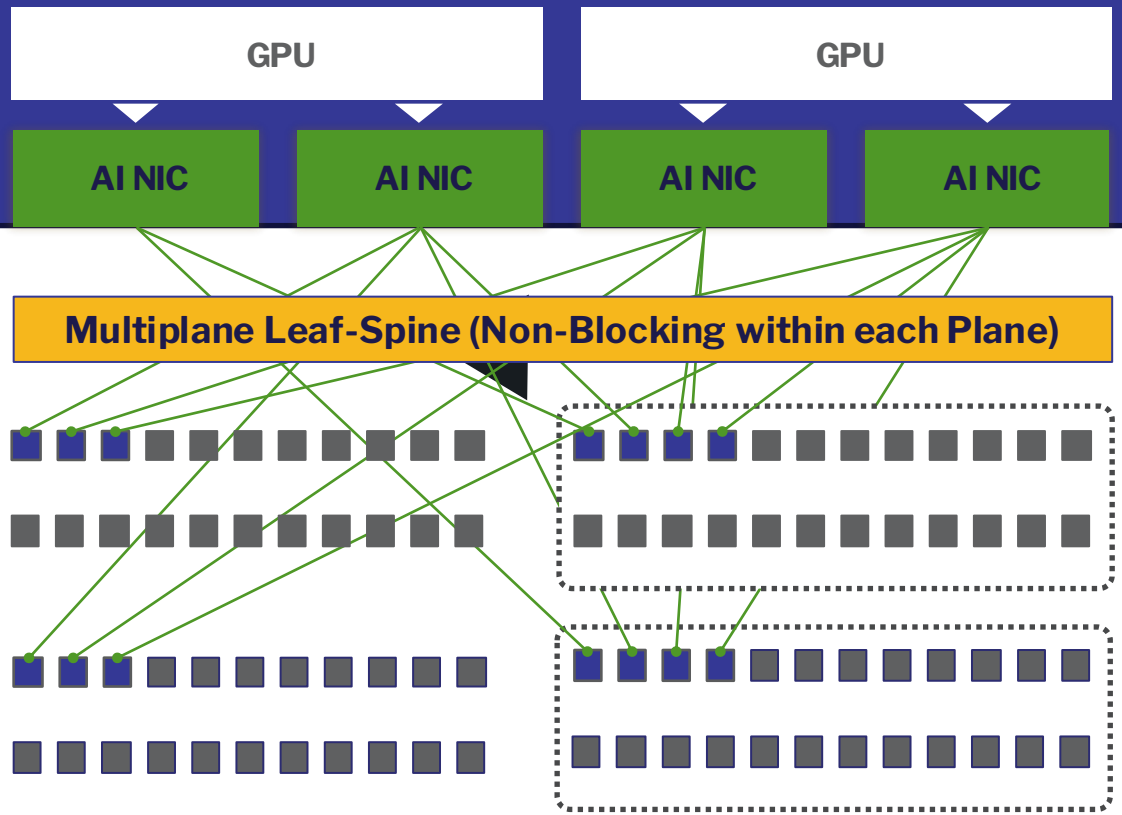
- Hardware link-layer replay corrects bit errors
- End-to-end retransmission recovers dropped packets
- Transfers retry on another plane if a link or switch fails



Multi-plane Architecture are Critical to Scaling Workloads

The Benefits of Multi-Plane

- Linear Scalability
- Resilience and Fault Isolation
- Proven at Scale
- Ready for growing workloads



Common Approach

- Programmability enables Multi-plane
- Reduces network complexity
- Improves fault isolation

AMD

MRC Network Load Balancing: NIC-based Source Routing

Customers Requirement:

- Deterministic control of from Operator
- Multi-plane

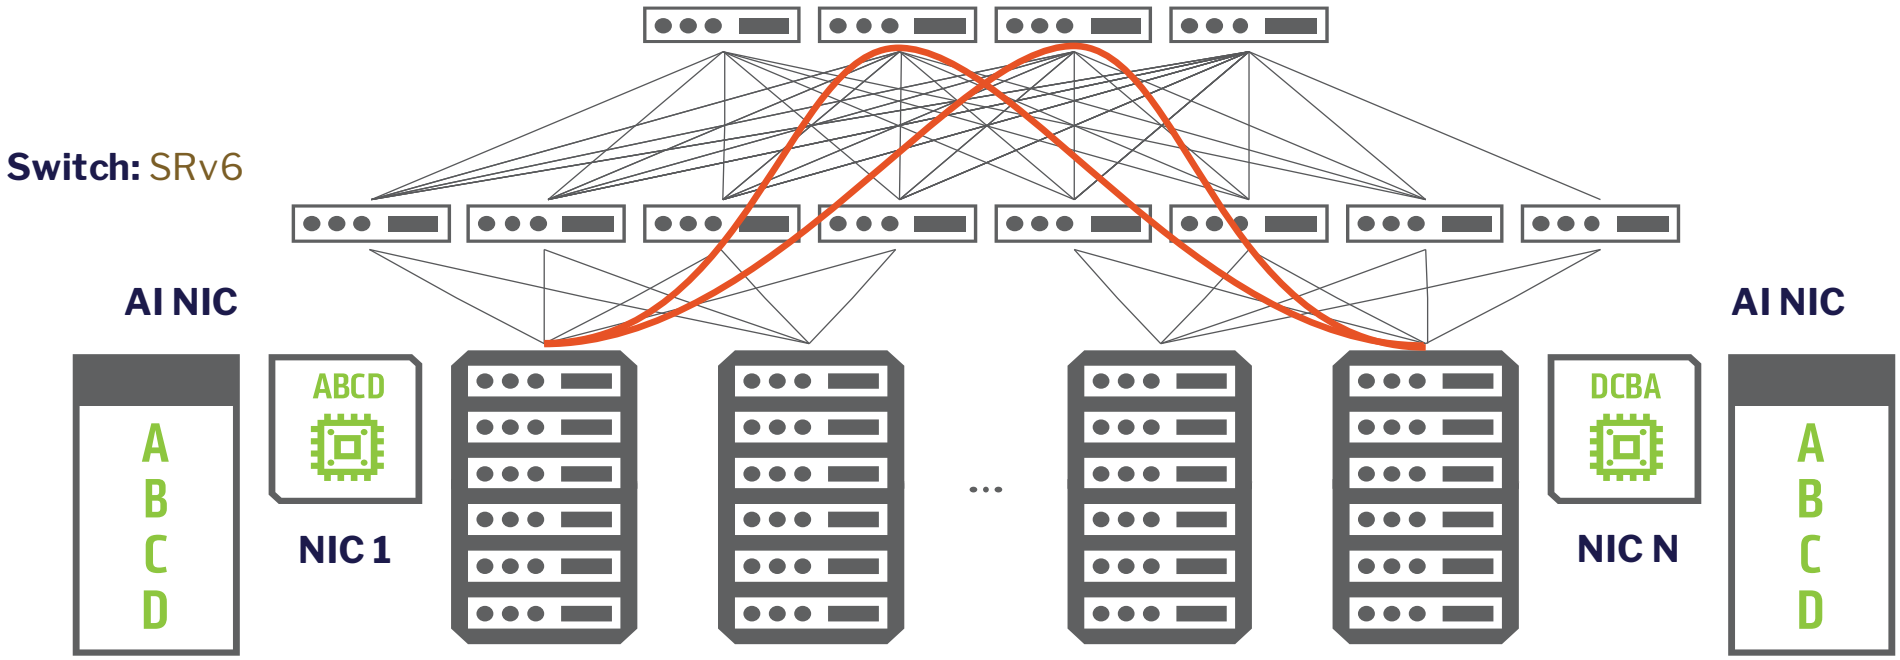
Source-Routed Network with IPv6 Segment-Routing (SRv6)

01. Control-plane Orchestration for EV-set mappings to SRv6 Micro-segment ID (uSID) list

02. NIC-based SRv6 path selection via per-packet uSID Carrier-Address assignment

03. Switches to do static / dynamic SRv6 uSID routing (no encap / decap)

04. Network link failures handled by fast NIC path failure detection and re-balancing of traffic



AMD

Poll Question 3:

Q: Which benefit of open, programmable networking matters most for your AI infrastructure?

Responses
A. Faster AI jobs and higher GPU utilization
B. Adding new capabilities on existing hardware
C. Freedom to choose vendors and transport protocols
D. Offloading networking, storage, and security tasks

Audience Poll Results

- Some discussion facilitated by host and analyst of the results with optional comments from sponsors

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking is Critical for AI Innovation
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

Programmable DPUs Accelerate Infrastructure Services at Scale

Customer deployments demonstrate networking performance and infrastructure efficiency

Microsoft

100×

Connections per second improvement*

DPU-based appliance and software for Accelerated Connections

IBM Cloud

30%

TCO savings*

Greater workload density and preserved CPU cores

ORACLE Cloud Infrastructure

2×

Network processing efficiency*

Efficient networking while preserving host CPU resources

Programmable DPU services	Virtual networking • Security policy enforcement • Storage acceleration
Relevance to agentic AI	Accelerate infrastructure services while preserving CPU resources

*Customer-reported results from separate deployments and workloads, not agentic AI benchmarks

Source: AMD, Industry Support for AMD Data Center Technologies • Oracle, Microsoft Azure and IBM Cloud statements

AMD

Comparison of Open Transport Protocols for Scale-Out

	Classic RoCEv2	Transport Protocols Built for AI	
		MRC	**MetaRoCE**
Bandwidth Utilization	No Multipathing	Multipath with Stateless Packet Spray	Multipath with Stateful Packet Spray
Multiplane Support	None	Yes	
Out of Order Delivery	None	In-Order Message Delivery	
Loss Recovery	Go-Back-N	SACK & Retransmit Packet Trimming/NACK	SACK & Retransmit
Congestion Control	PFC/DCQCN	Path Aware Congestion Control	
Network Failure Recovery	Slow Error/Fault Recovery	Fast Fault Recovery EV Probing	RTO-Based Path Shift

25%

Improved RCCL Performance
(UEC-ready RDMA vs. RoCEv2)*

Higher GPU Utilization

50%

Higher Cluster Uptime
(UEC-ready RDMA vs RoCEv2)*

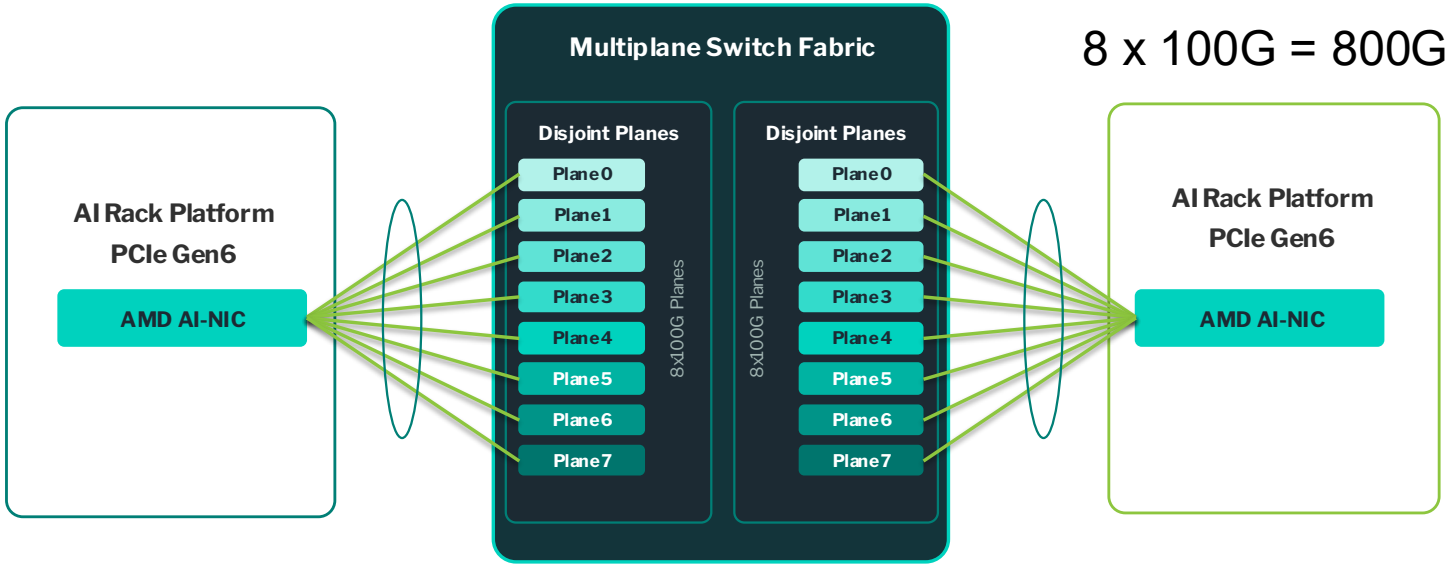
Improved Operational Efficiency

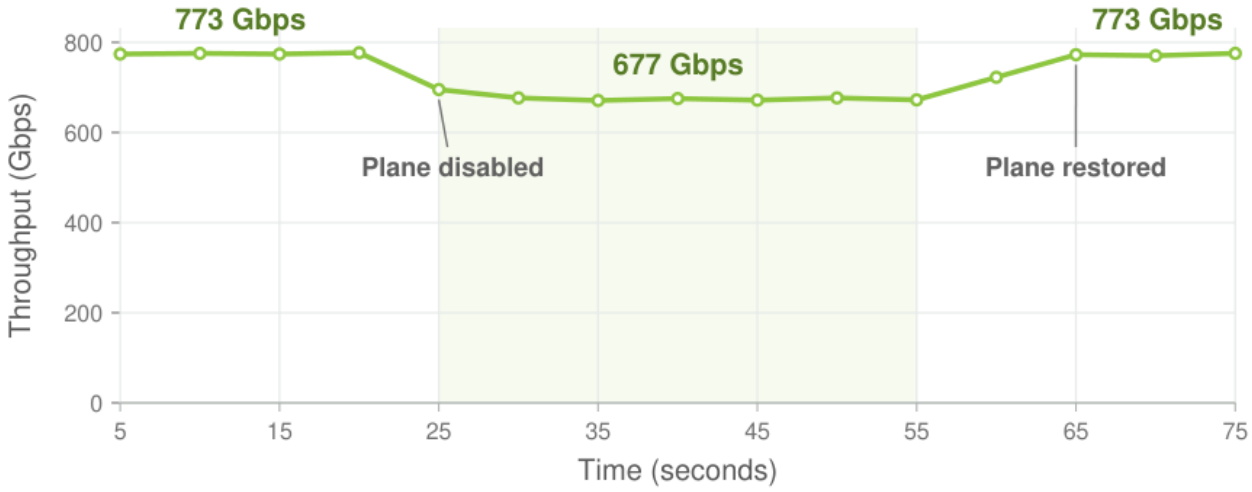
**End-notes: PEN-016, PEN-019

AMD

MetaRoCE: Programmable Scale-Out Multiplane Networking

- Solution Role
  - Programmable transport optimized for collective communication
  - Multipath networking improves utilization of available bandwidth
- Solution Benefits
  - Reduced communication bottlenecks
  - Computation continues through link failures
  - Graceful recovery





AMD

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking is Critical for AI Innovation
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

AMD Helios™ Rackscale Solution

Leadership Platform for Rack Scale AI

18	**6**	**18**
Compute Trays	Scale-Up Switch Trays	Front End DPUs
2.4 Tbps	**260 TB/s**	**43 TB/s**
Scale-Out BW Per GPU	Scale Up Bandwidth	Scale Out Bandwidth



AMD Fabric Manager

Rack Management for
Scale-Up Network

OPEN Compute Project

Ultra Ethernet Consortium

ULTRA ACCELERATOR LINK™

ESUN

AMD

AMD Helios™ Rackscale Solution

Benefits for Training, Distributed Inference and Agentic AI

Stronger Data Protection

Encrypt network and storage data with DPU offload.

Secure Multi-Tenancy

Isolate tenants and share GPU resources with vPods.

Resilient AI Operations

Recover network paths; contain faults within vPods.

Shorter AI Training Times

Accelerate GPU synchronization and collective communication.

Faster First Responses

KV-cache reuse can reduce time to first token (TTFT).

Greater Agentic Capacity

Free CPU cores for agents; expand shared context storage.



Up to 30%

More Token Per Dollar

Up to 34x

Higher Token Throughput

Performance varies by workload and configuration. TTFT benefit requires cache reuse; ROCm AIC support for Helios is planned.

AMD Helion Agentic AI Rack

Up to 32,768 “Zen 6”	**Up to 16 MI350P**	**Up to 102 Tbps Networking**
x86 CPU Cores	74 PetaFlops - FP4	Fully Programmable

ORv3 44OU Rack	**Standard Modular Compute Board**
Native OCP Spec Compliant	Native OCP Spec Compliant



AMD

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking is Critical for AI Innovation
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

Take Aways

- AI is Driving an Unprecedented Expansion in Networking Demand: Agentic AI, distributed inference, and larger GPU clusters are accelerating network bandwidth requirements, scale, and complexity.
- Networking Is Becoming a Critical Determinant of AI Performance: Network performance increasingly dictates GPU utilization, job completion times, infrastructure efficiency, and ultimately AI economics.
- Open Ethernet Is Emerging as the Foundation of AI Networking: Ethernet's expanding role across scale-up, scale-out, and scale-across enables greater interoperability, supplier diversity, and innovation.
- Programmability Helps Keep Pace with AI Innovation: AI workloads are evolving faster than hardware refresh cycles, making programmable NICs, DPUs, and network infrastructure increasingly critical.

Agenda Overview

- AI Networking: Market Overview
- New Demands of AI Networking
- Programmable Networking is Critical for AI Innovation
- Results from Scale Deployments
- Sponsor Products
- Conclusion
- Live Q&A

Programmable by Design:
Building Open Networking That Evolves at the Pace of AI



Cliff Grossner, Ph.D.

Chief Innovation Officer
OCP Foundation



Sameh Boujelbene

VP, AI Networks
Dell'Oro group



Shane Corban

Senior Director of
Technical Marketing
AMD

Thank you for attending

Programmable by Design: Building Open Networking That Evolves at the Pace of AI

OCP Educational Webinar Series

November 10th 2026

OPEN Compute Project®

Dell'Oro Group

AMD

Community-driven hyperscale innovation for all

End-Notes and Footnotes

- "PEN-016 - Testing conducted by AMD Performance Labs as of [28th April 2025] on the [AMD Pensando™ Pollara 400 AI NIC], on a production system comprising of: 2 Nodes of 8xMI300X AMD GPUs (16 GPUs): Broadcom Tomahawk-4 based leaf switch (64x400G) from MICAS network; CLOS Topology; AMD Pensando Pollara AI NIC – 16 NICs; CPU Model in each of the 2 nodes - Dual socket 5th gen Intel® Xeon® 8568 - 48 core CPU with PCIe® Gen-5 BIOS version 1.3.6 ; Mitigation - Off (default)System profile setting - Performance (default) SMT- enabled (default); Operating System Ubuntu 22.04.5 LTS, Kernel 5.15.0-139-generic. Following operation were measured: AllreduceAverage 25% for All-Reduce operations with 4QP and using UEC ready RDMA vs the RoCEv2 for multiple different message size samples (512MB, 1GB, 2GB, 4GB, 8GB, 16GB) . The results are based on the average at least 8 test runs."
- PEN-019: Testing conducted by AMD Performance Labs as of [15 September 2025] on the AMD Pensando Pollara AI NIC, on a test system comprising of SMC-300X server for GPU-GPU communication: 2x AMD Pensando Pollara AI NIC, 2P AMD EPYC 9454 48-Core -2P Processor, 8x AMD Instinct MI300X GPU, Ubuntu 22.04.5 LTS, kernel 5.15.0-139-generic, ROCm 6.4.1.0-83-69b59e5. Testing running Llama-3.1-8B, Model Configuration: SEQ_LEN=2048, TP=1, PP=1, CP=1,FP8=1, MBS=10, GBS = 5120. Iteration = 2, No. of paths/QP : 128. Results may vary based on factors including but not limited to system configuration and software settings.
- "PEN-020: Testing conducted by AMD Performance Labs as of [15 September 2025] on the AMD Pensando Pollara AI NIC running Llama 3.1-405B @ 64 global batch size (GBS) with 8K sequence length, on a test system comprising of 8 node SMC-300X server for GPU-GPU communication using 2x AMD Pensando Pollara AI NIC or 2x Nvidia CX-7, 2P AMD EPYC 9454 48-Core 2P - Processor, 8x AMD Instinct MI300X GPUs, Ubuntu 22.04.5 LTS, kernel 5.15.0-139-generic, ROCm 6.4.1.0-83-69b59e5 Following operation are part of gateway functionConfiguration : Num layers=4, Data Type=BF16, DCN - TP=1, PP=1, SP=1, DP=1, FSDP=-1, ICI - TP=1, PP=1, SP=1, DP=1, FSDP=8. AINIC container: jax-private:rocm6.4.0-jax0.5.0-py3.10.12-tedev2.1-20250801_training. Results may vary based on factors including but not limited to system configuration and software settings."
- Satinder Sethi, GM, IBM Cloud Infrastructure Services: "By collaborating with AMD and utilizing their programmable DPU, and turnkey SDN and security services software, IBM Cloud has been able to accelerate to market our award-winning secure high-performance "bare-metal servers for VPC" offering for VMware, Red Hat and SAP workloads, driving significant network throughput improvements for customers, and realizing 30% TCO savings from greater workload density and preserved CPU cores." -- published on amd.com: https://www.amd.com/en/solutions/data-center/data-center-ai-premiere/industry-support.htm
- Gerald De Grace, Lead, Microsoft Azure Accelerated Connections: "Our collaboration with AMD using the AMD Pensando DPU based smart appliance and full-stack AMD software in production for Accelerated Connections has resulted in 100X improvement in CPS performance over our existing solution and will extend to broader tiers within the Azure network in the future."
- Oracle Introduces New Cloud Networking Capabilities for Any Workload