Evolving UALink and UEC as the Gold Standard for Accelerator Connectivity in AI

Pankaj Goel
Associate Director, Siemens EDA
Questa One Avery VIP

Agenda

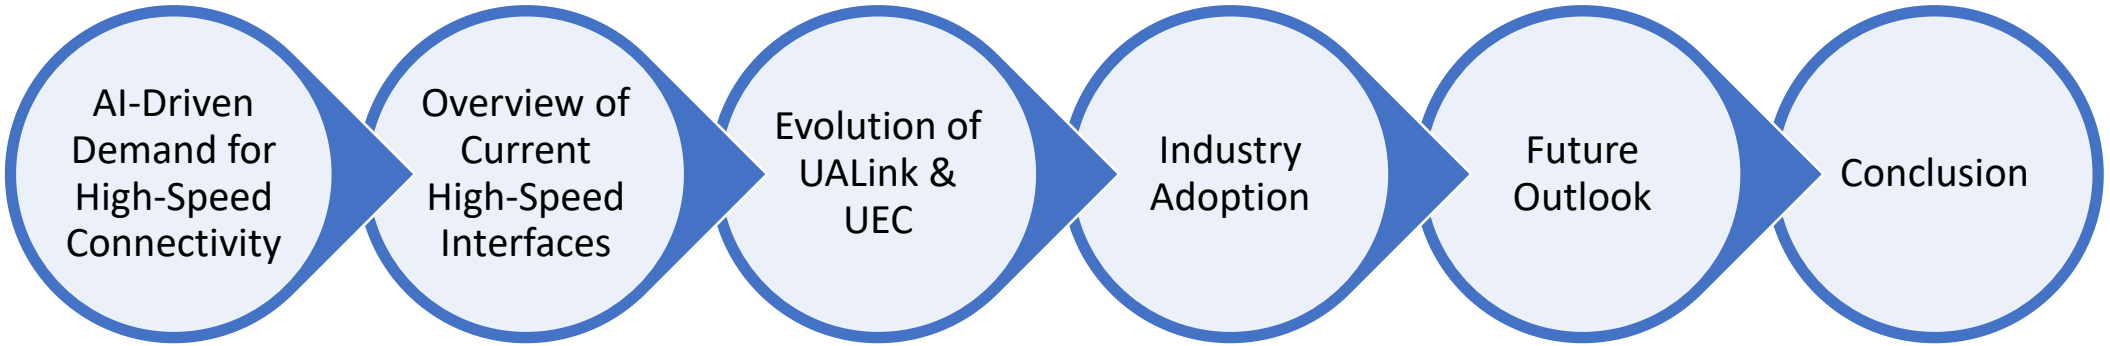

AI-Driven Demand for High-Speed Connectivity

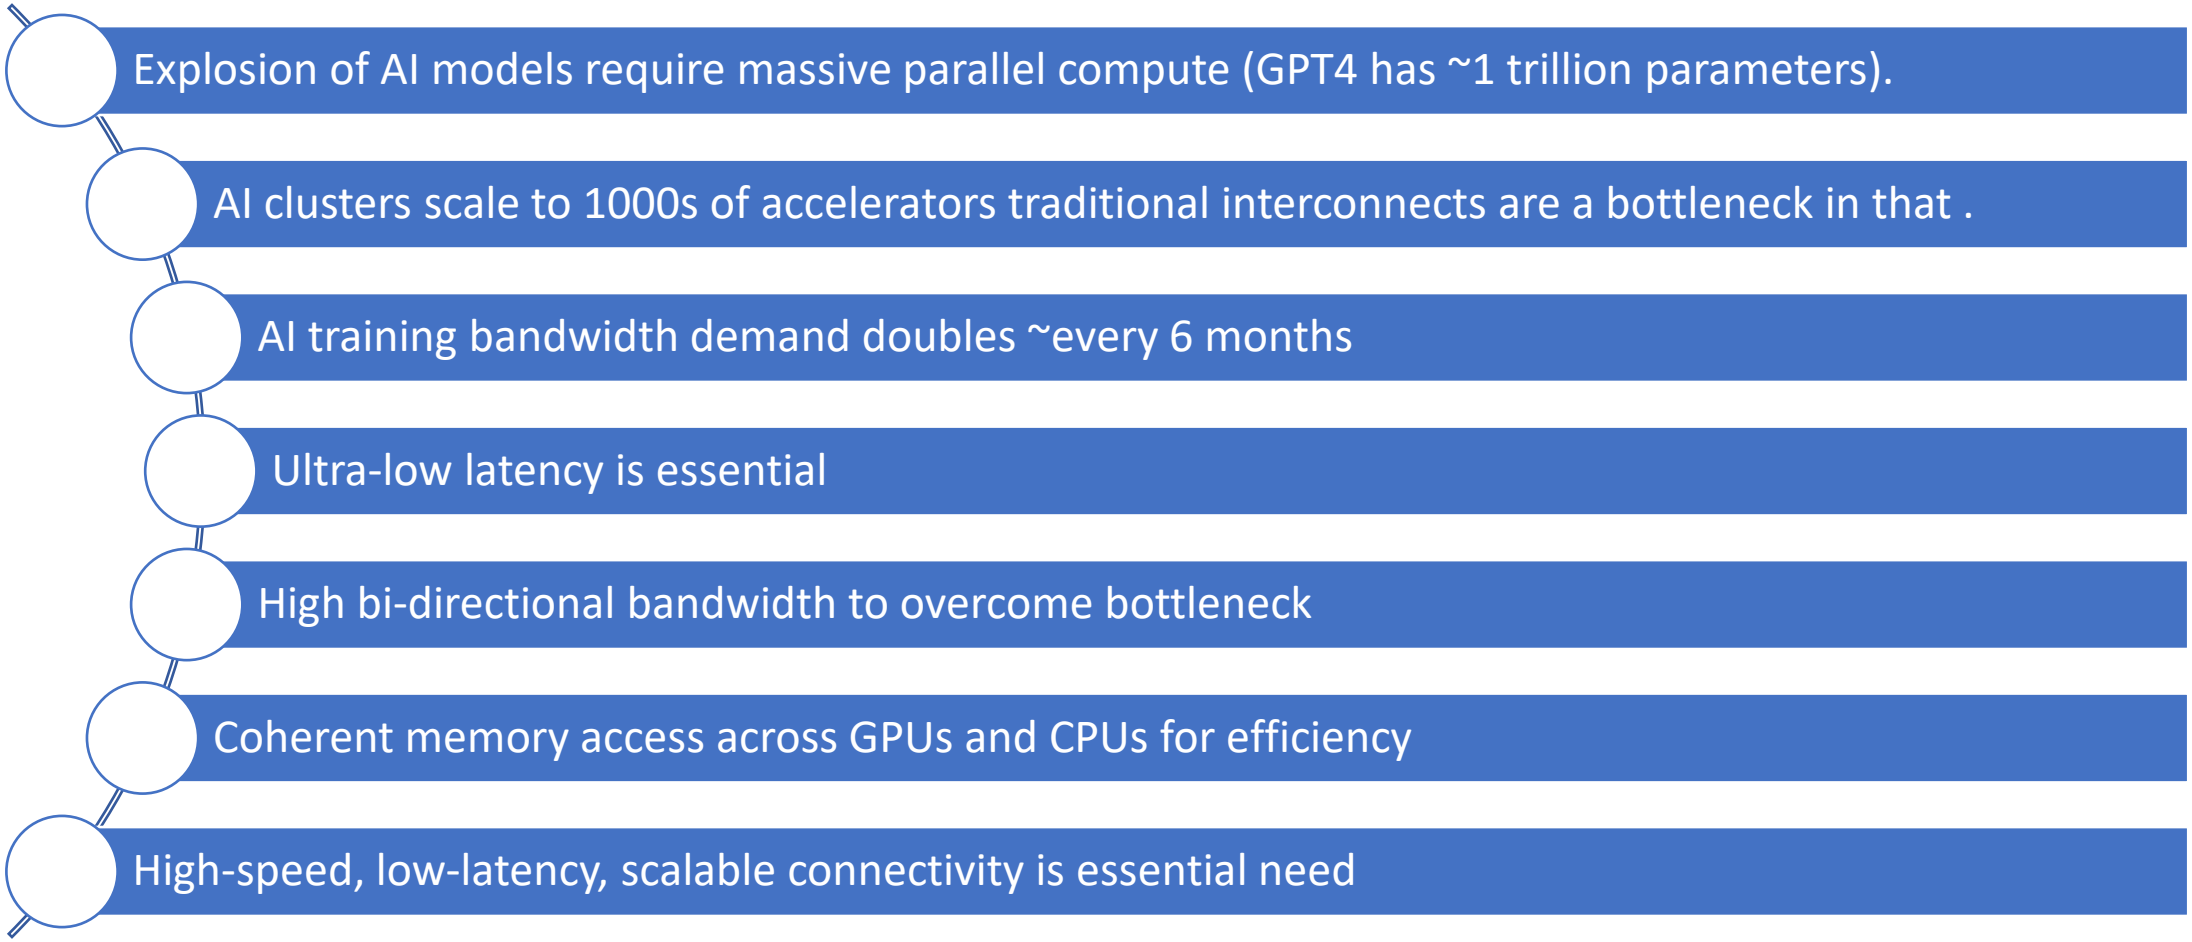

Worldwide Data Centre Computing Forecast, Accelerated and Non-Accelerated

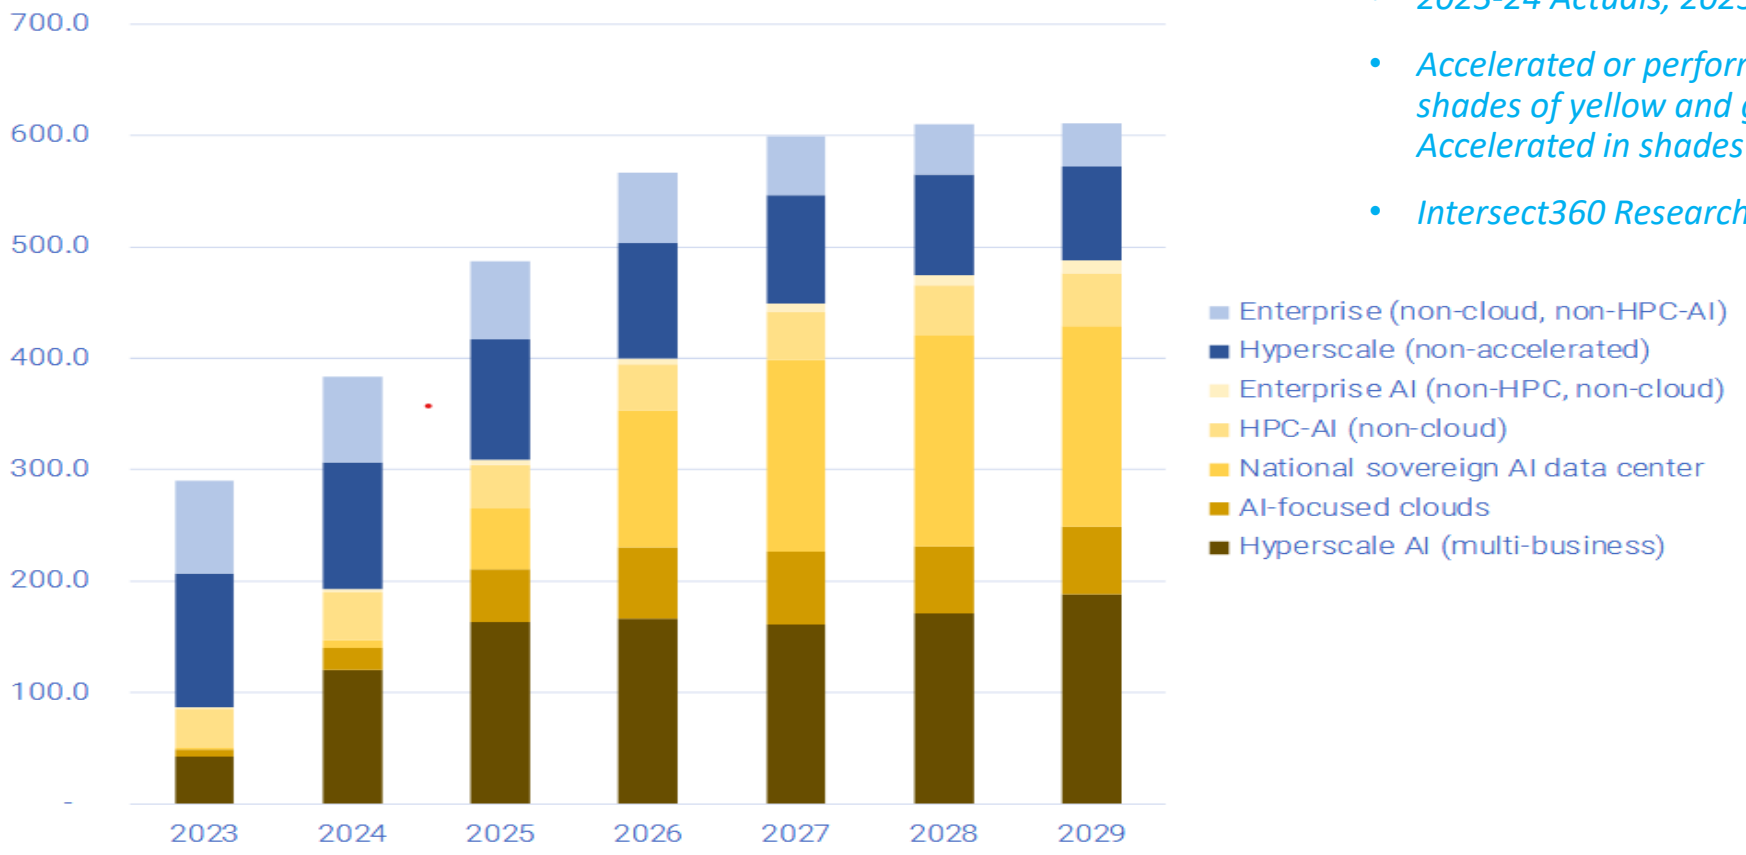

- *2023-24 Actuals, 2025-29 Forecast (in $B).*
- *Accelerated or performance-optimized in shades of yellow and gold; non-Accelerated in shades of blue.*
- *Intersect360 Research, 2025*

Popular High Speed connectivity Today

Technology	Latency	Bandwidth	Scope	Topology	Vendor Lock-in	Use Case
PCI Express (PCIe 5/6)	150–300 ns	~64 GB/s (Gen5 x16) / 128 GB/s (Gen6 x16)	CPU ↔ Device	Point-to-point	Open Standard	General compute, GPUs, storage
NVLink (v4/v5)	~50–80 ns	~200 Gbps per link (v5), ~900 GB/s total	GPU ↔ GPU	Mesh/Full mesh	Vendor Locked	AI training clusters
Infinity Fabric	~100 ns	64–128 GB/s	CPU ↔ Device ↔ Memory	Ring or Mesh	Vendor Locked	CPU-GPU-DRAM sharing
Scale-Up Ethernet (SUE)	~500 ns	400–800 Gbps	Pod-scale	Tree/Leaf-Spine	Vendor Locked	AI accelerator networks
✴ **UALink**	<100 ns (pin-to-pin), ~250 ns end-to-end	200 Gbps/lane (x4 = 800 Gbps per port)	GPU ↔ GPU, NPU ↔ NPU	Fully connected mesh or hybrid ring	Open Consortium	Scale-up AI pods, shared memory fabrics
✴ **UEC (Ultra Ethernet Consortium)**	~500 ns–1 µs	Multi-Tbps	Pod ↔ Pod, Cluster ↔ Cluster	Leaf-spine/Ethernet	Open Consortium	Scale-out, datacenter-wide AI workloads

PCIe

Widely adopted and cost-effective, but lacks peer-to-peer accelerator communication, making it inefficient for large-scale AI training.

NVLink Infinity Fabrics

Highly advanced, enabling ultra-high-bandwidth, low-latency accelerator communication, but proprietary and limits open architecture scalability

UALink

AI-native, open-standard interconnect for high-speed communication, but ecosystem support is still emerging

Unified Vision: UALink & UEC

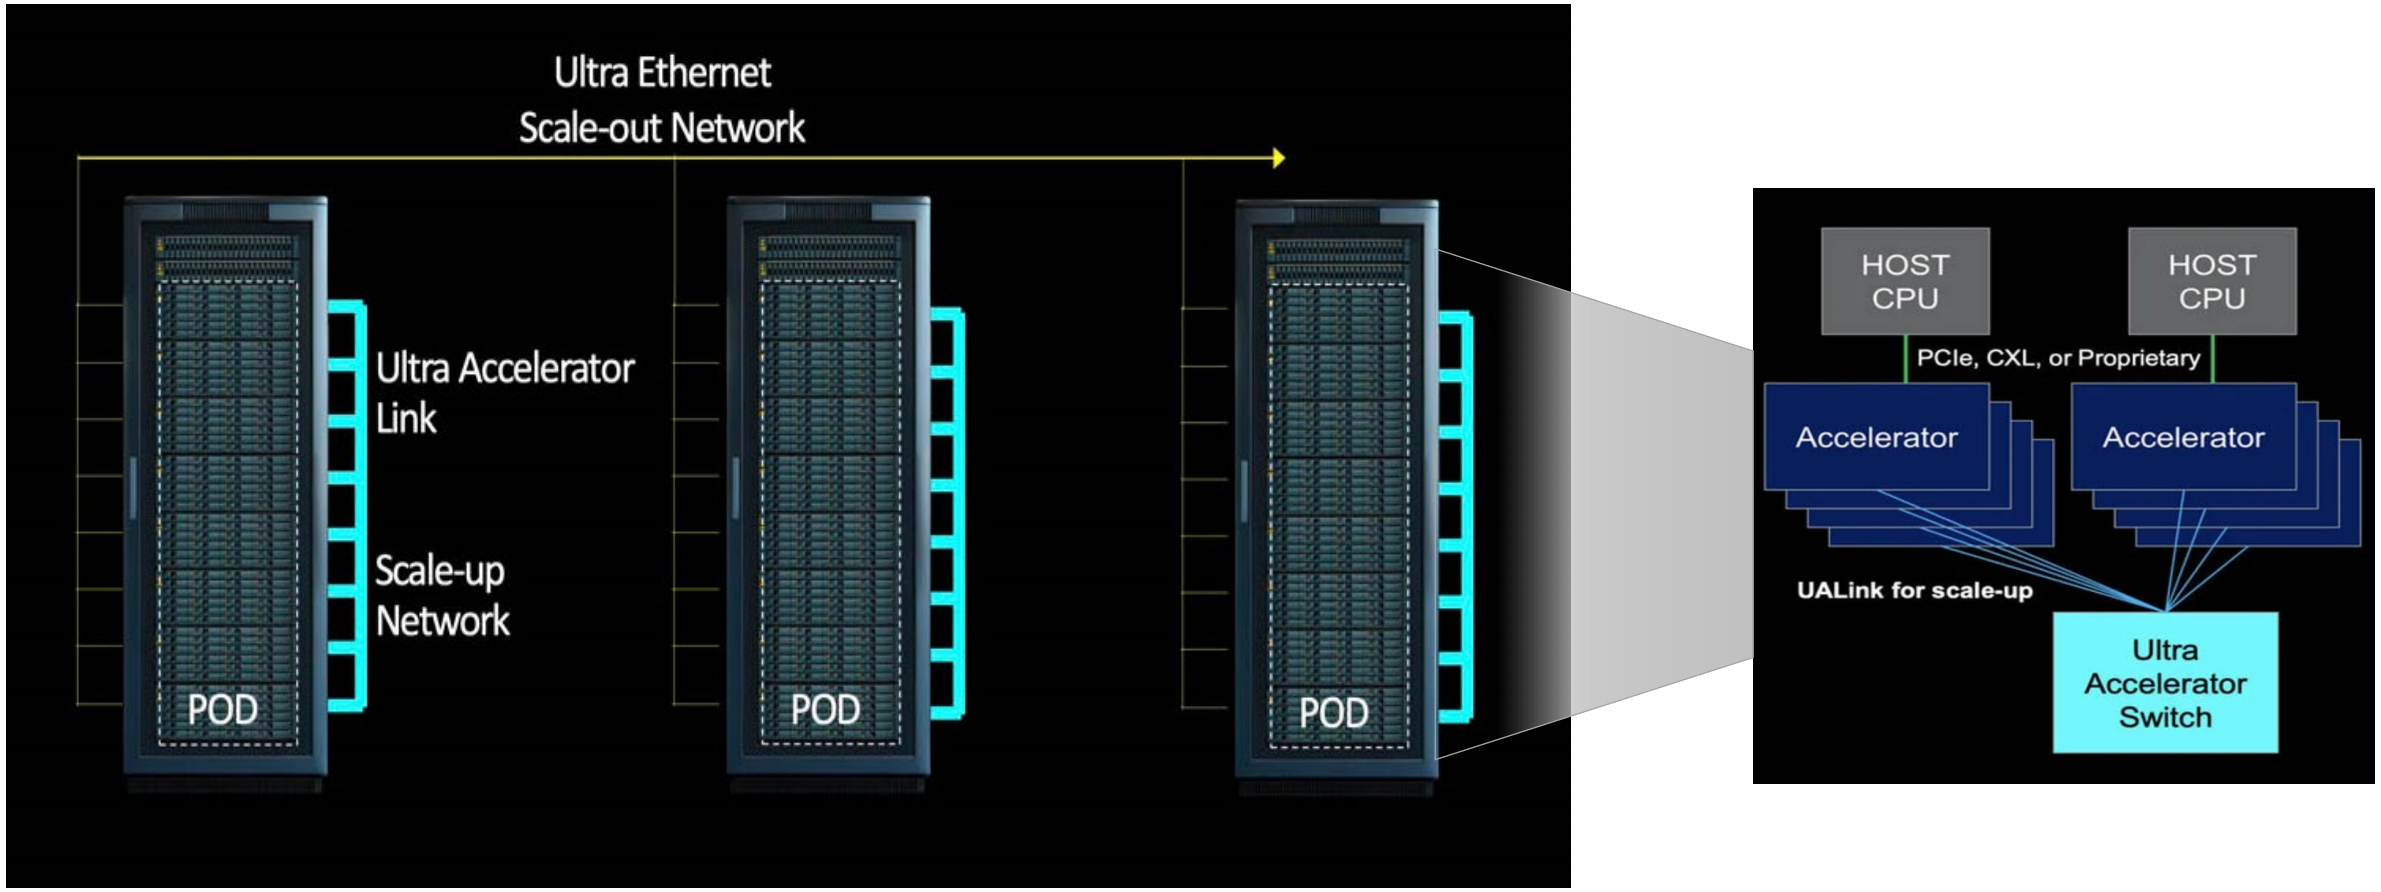

UALink

- The UALink interconnect is for Accelerator-to-Accelerator communication
  - The initial focus will be sharing DDR & HBM memory among accelerators
- Direct load, store, and atomic operations between accelerators (i.e. GPUs)
  - Low latency, high bandwidth fabric for 100's of accelerators in a pod
  - Simple load/store/atomics semantics
- Supports data rates of 128Gbps and 200Gbps per lane

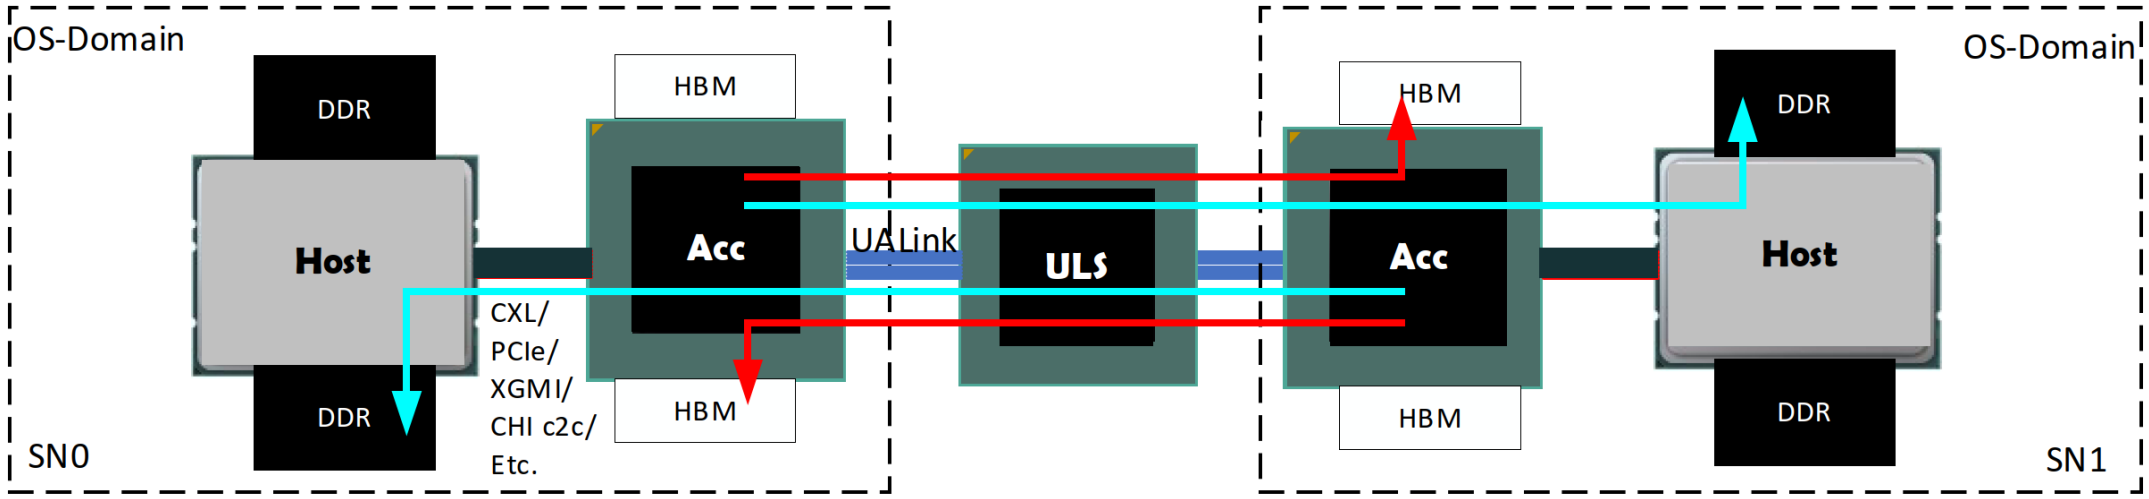

UALink stack

Ethernet

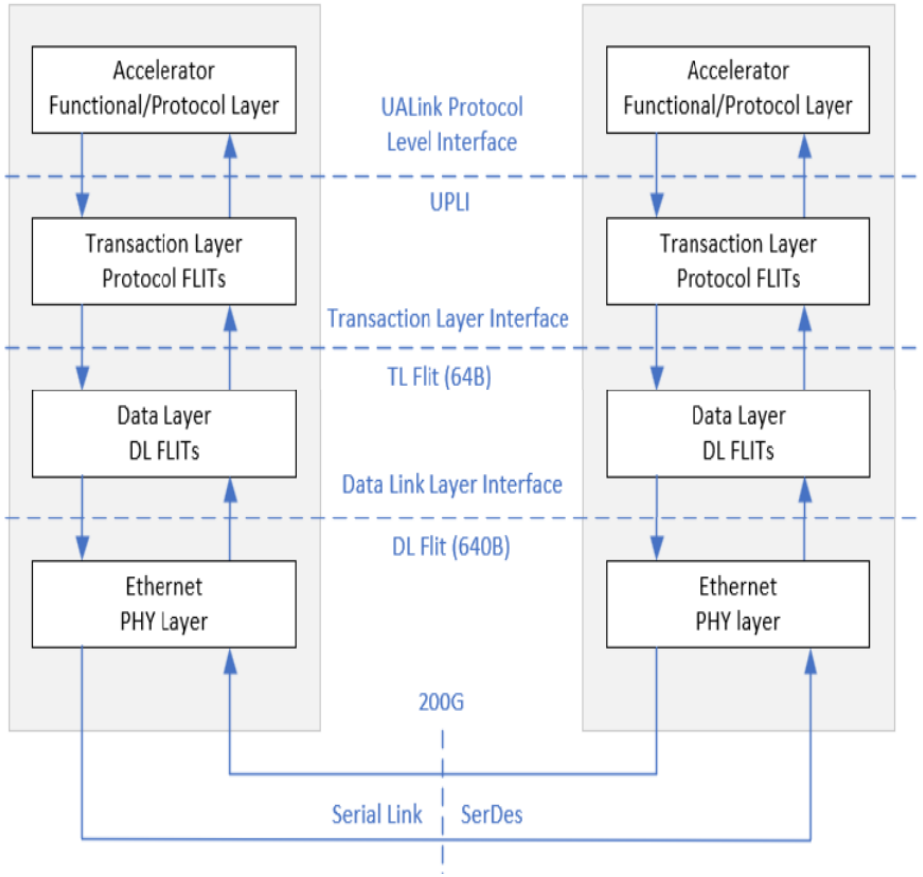

PCIe

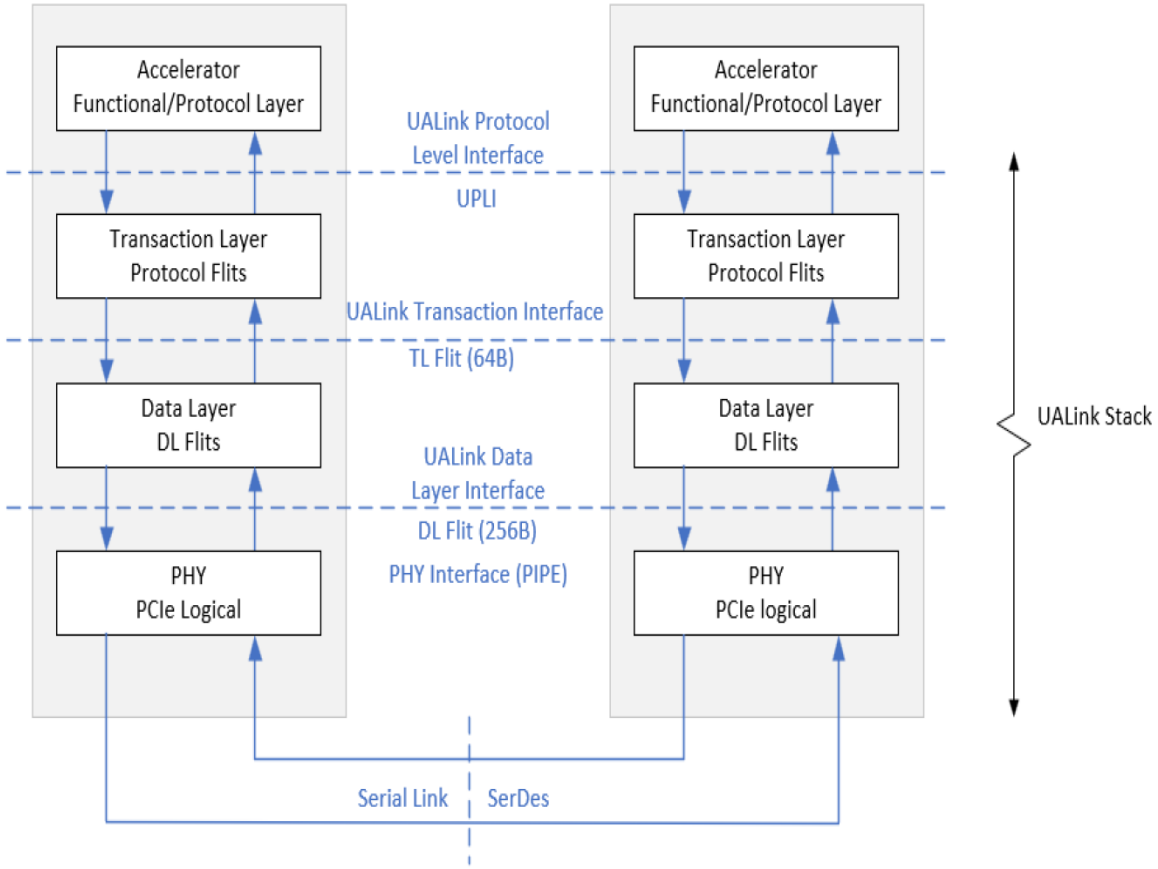

UALink: Architectural Edge

Memory-semantic Fabric

- Direct load/store/atomic between GPUs; pod can appear as a single large-memory accelerator .

Low latency & high bandwidth

- 200 Gbps per lane; 4-lane “Station” = 800 Gbps full duplex
- Switch latency: < 300 ns at full scale

Protocol stack

- Layered: physical (Ethernet PHY), data link (FEC, flow-control, replay), transaction (64 B flits), protocol (load/store)

Security

- Encryption/Authentication protects against hardware attacks

Scalability

- Supports up to 1,024 accelerators in a single pod
- Bandwidth efficiency: 88→95% via compression, efficient flit packing, FEC optimizations
- Power-saving: optimized protocol stack and PHY reduce die area and interconnect power by ~40%

UALink Roadmap

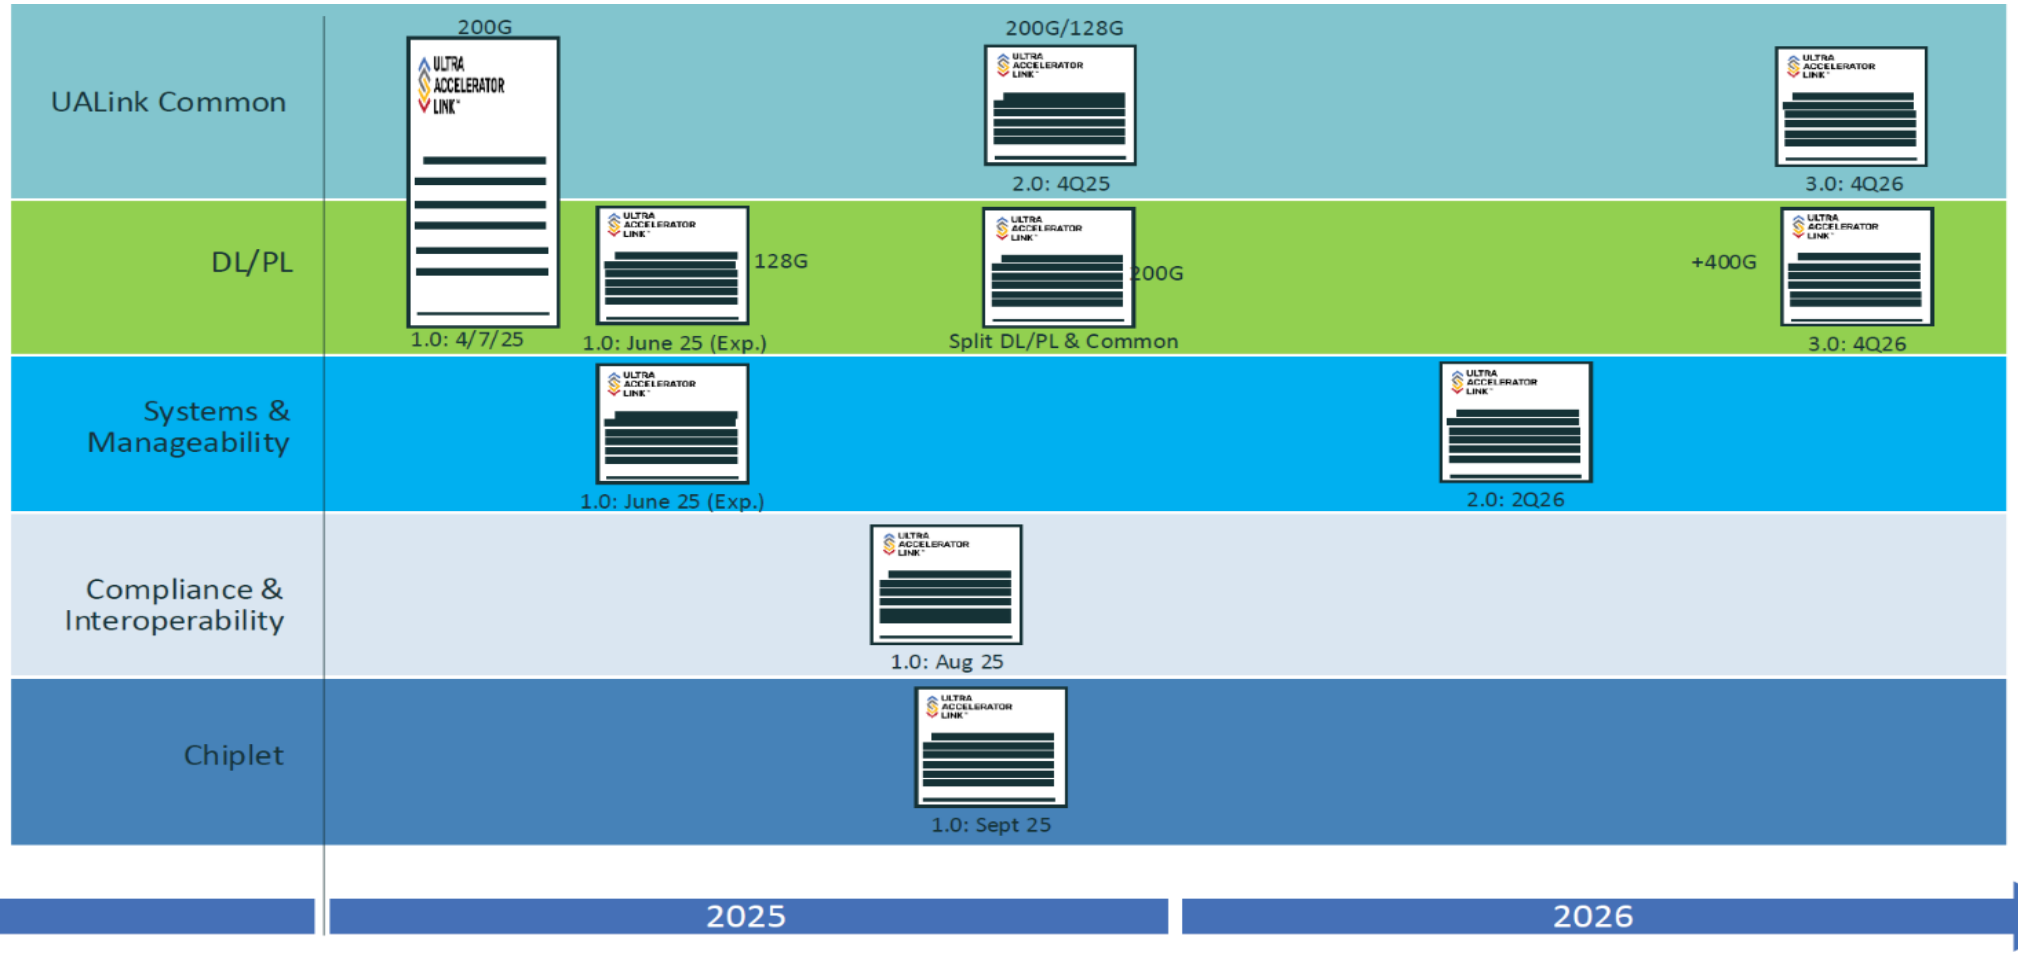

UALink Verification Use Case

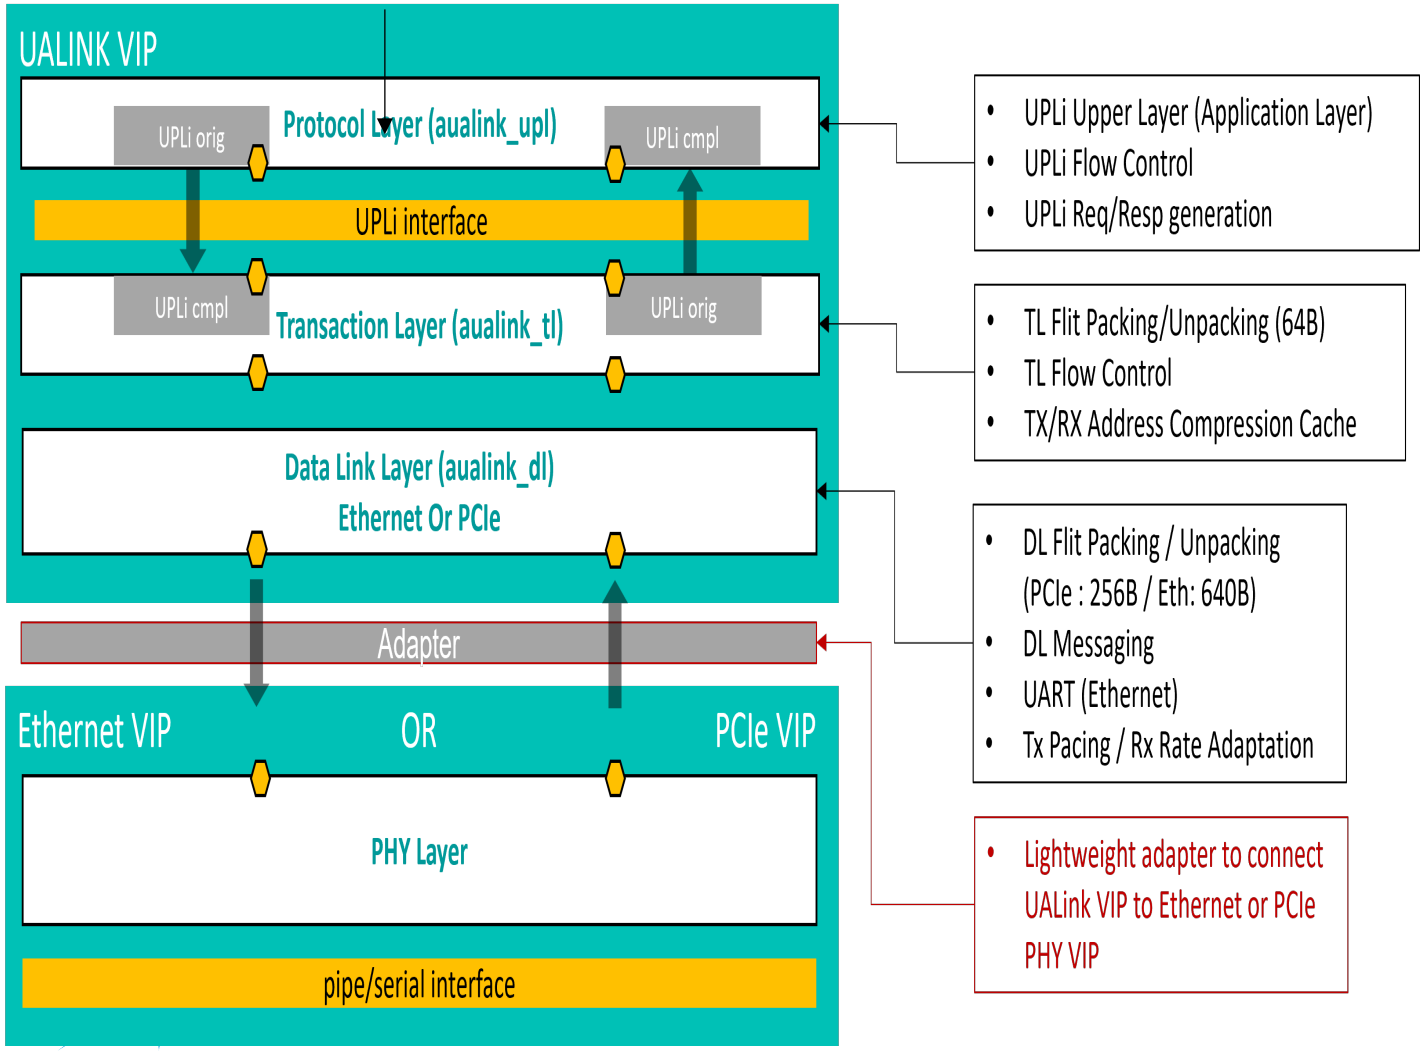

- VIP Vendor (Avery VIP) provides modular components
- Any component can be replaced by VIP or DUT
- Dedicated & clean connection
- In built Tests and Checks for compliance
- Coverage for spec completeness

Ultra Ethernet (UE): Standardization Ecosystem

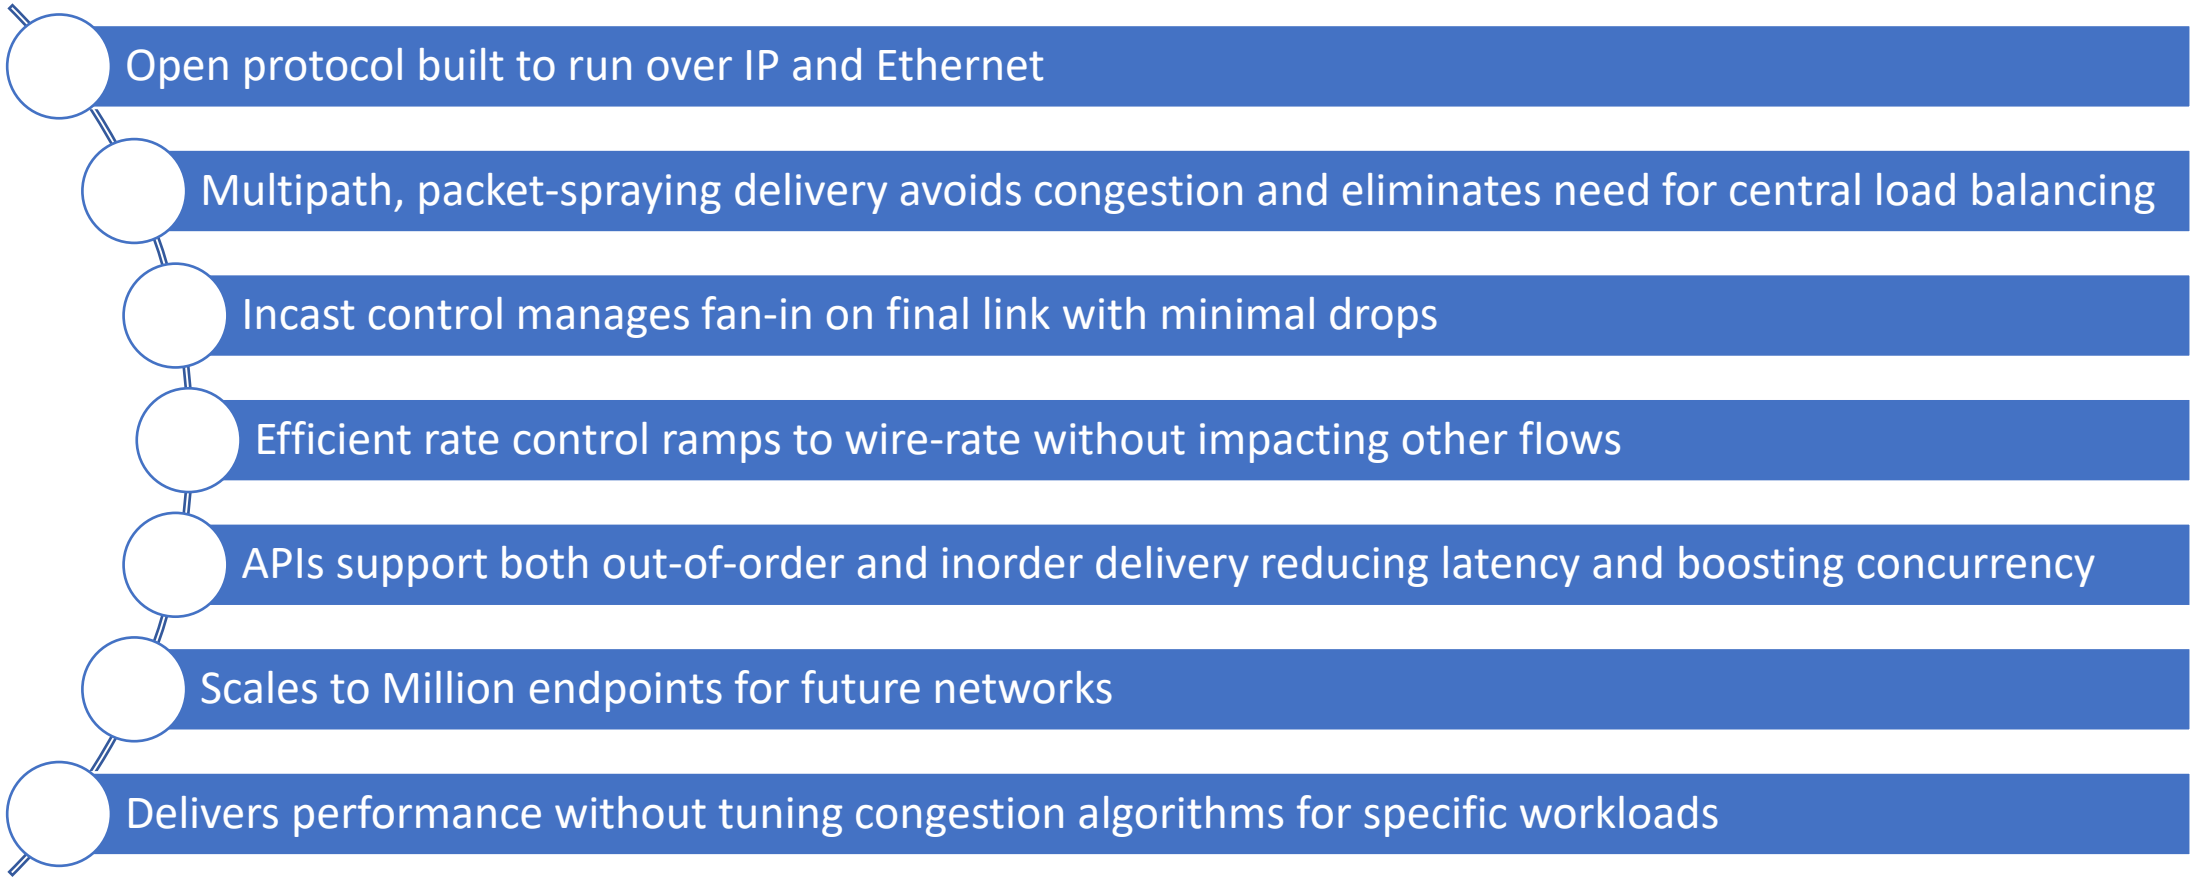

Ultra Ethernet Transport

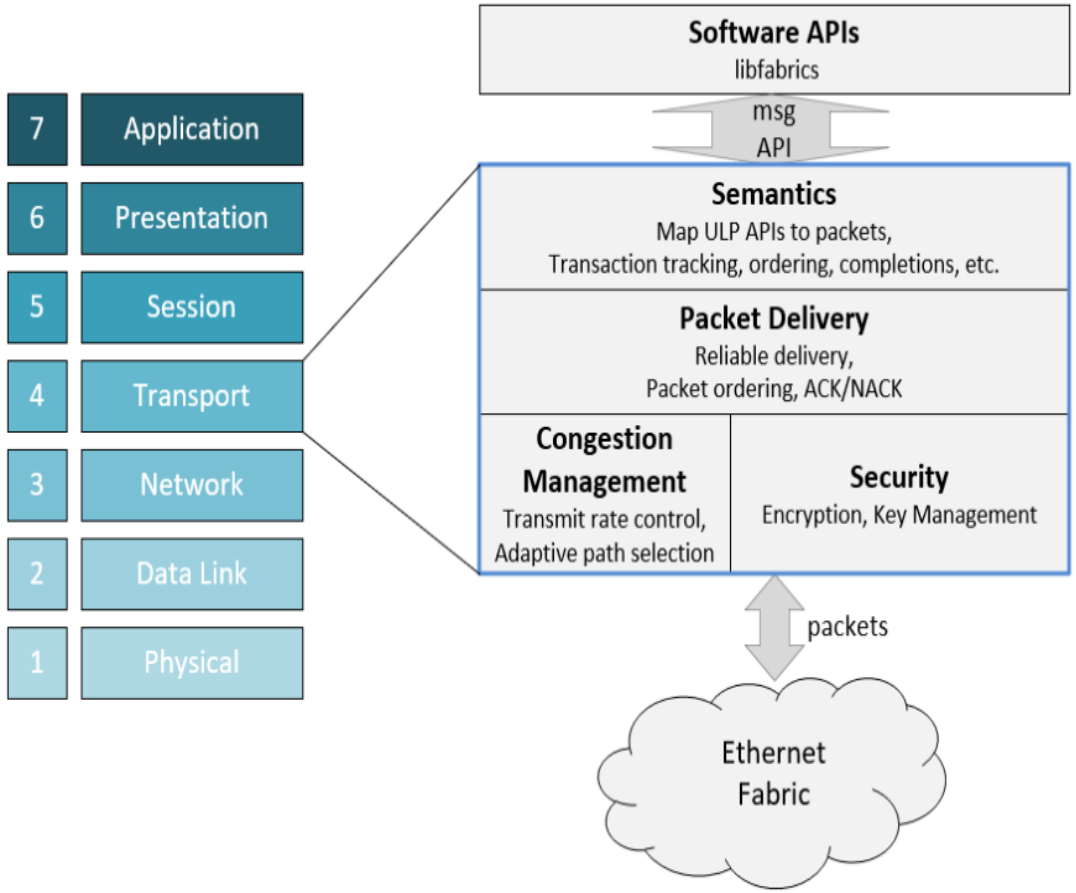

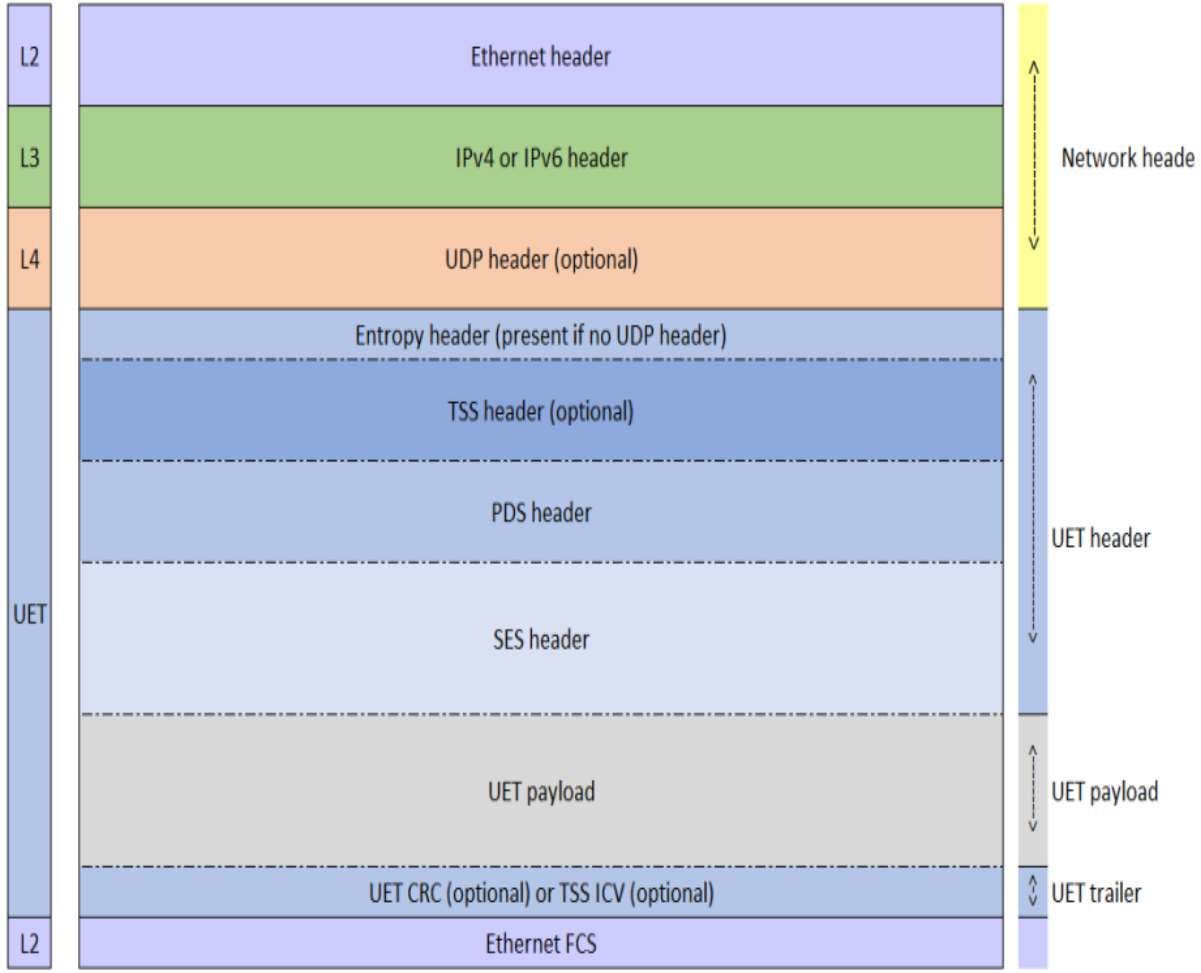

The Ethernet Advantage

Many large accelerators clusters for AI training already run on Ethernet-based IP networks

- Broad, multi-vendor ecosystem of switches, NICs, cables, optics, tools, and software
- Scalable IP routing for rack, building, and datacentre-wide networks
- Wide range of tools for testing, deployment, and operations
- Lower costs via competition and economies of scale
- Rapid, regular advancements through IEEE Ethernet standards across physical and optical layers

Ethernet vs Ultra Ethernet

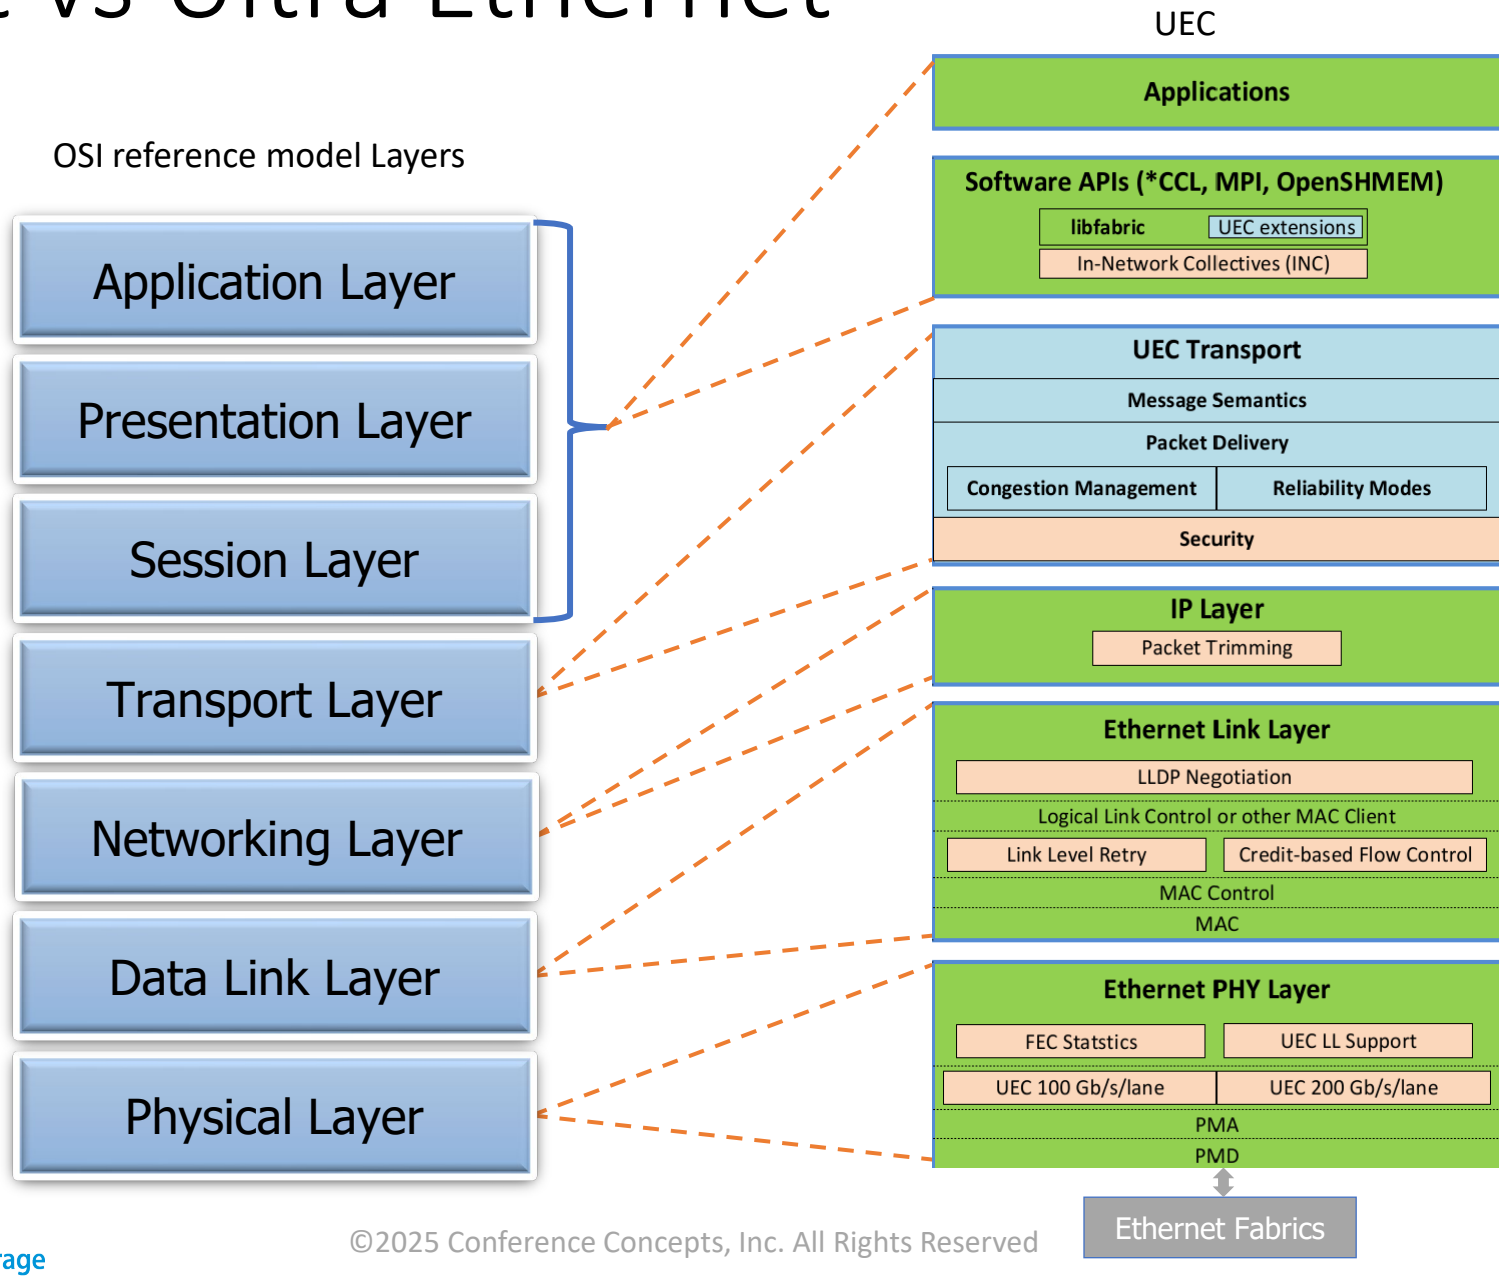

Ultra Ethernet Verification Use model

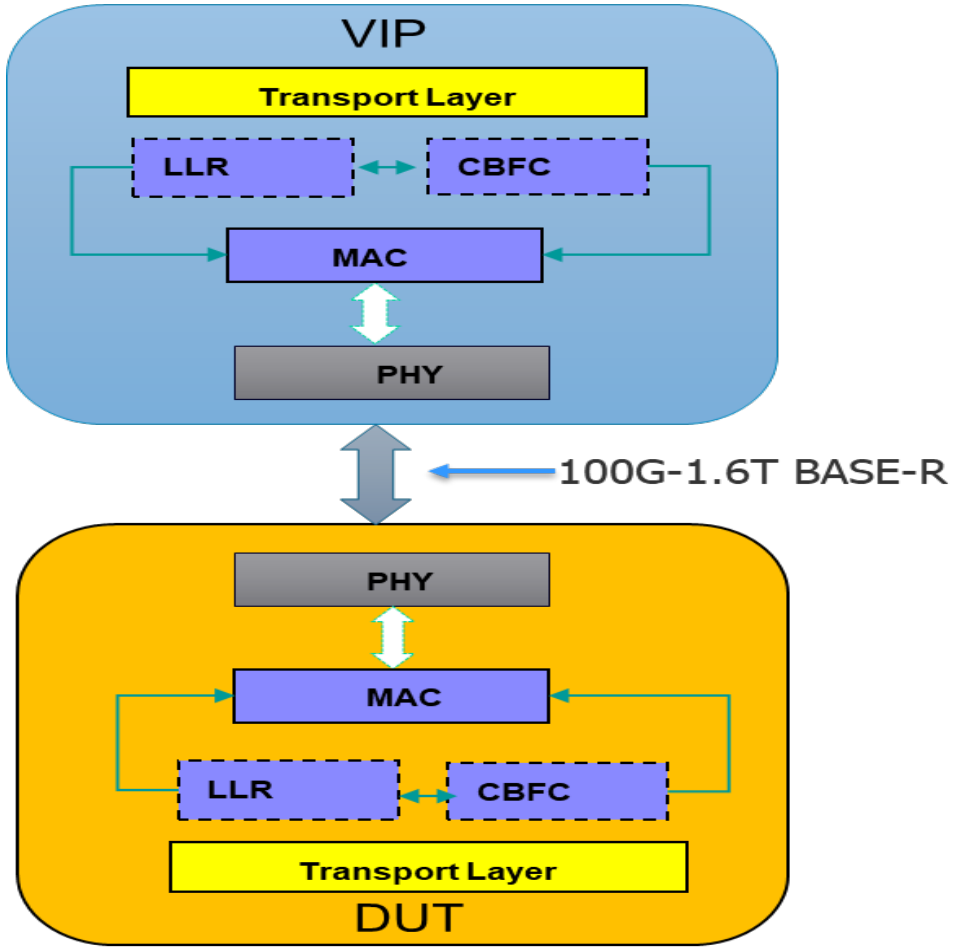

- VIP Vendor (Avery VIP) provides modular components
- Components can be VIP or DUT
- Configurable port speeds from 100G to 1.6T
- Dedicated & clean connection
- In built Tests and Checks for compliance
- Coverage for spec completeness

UEC Specification Rollout

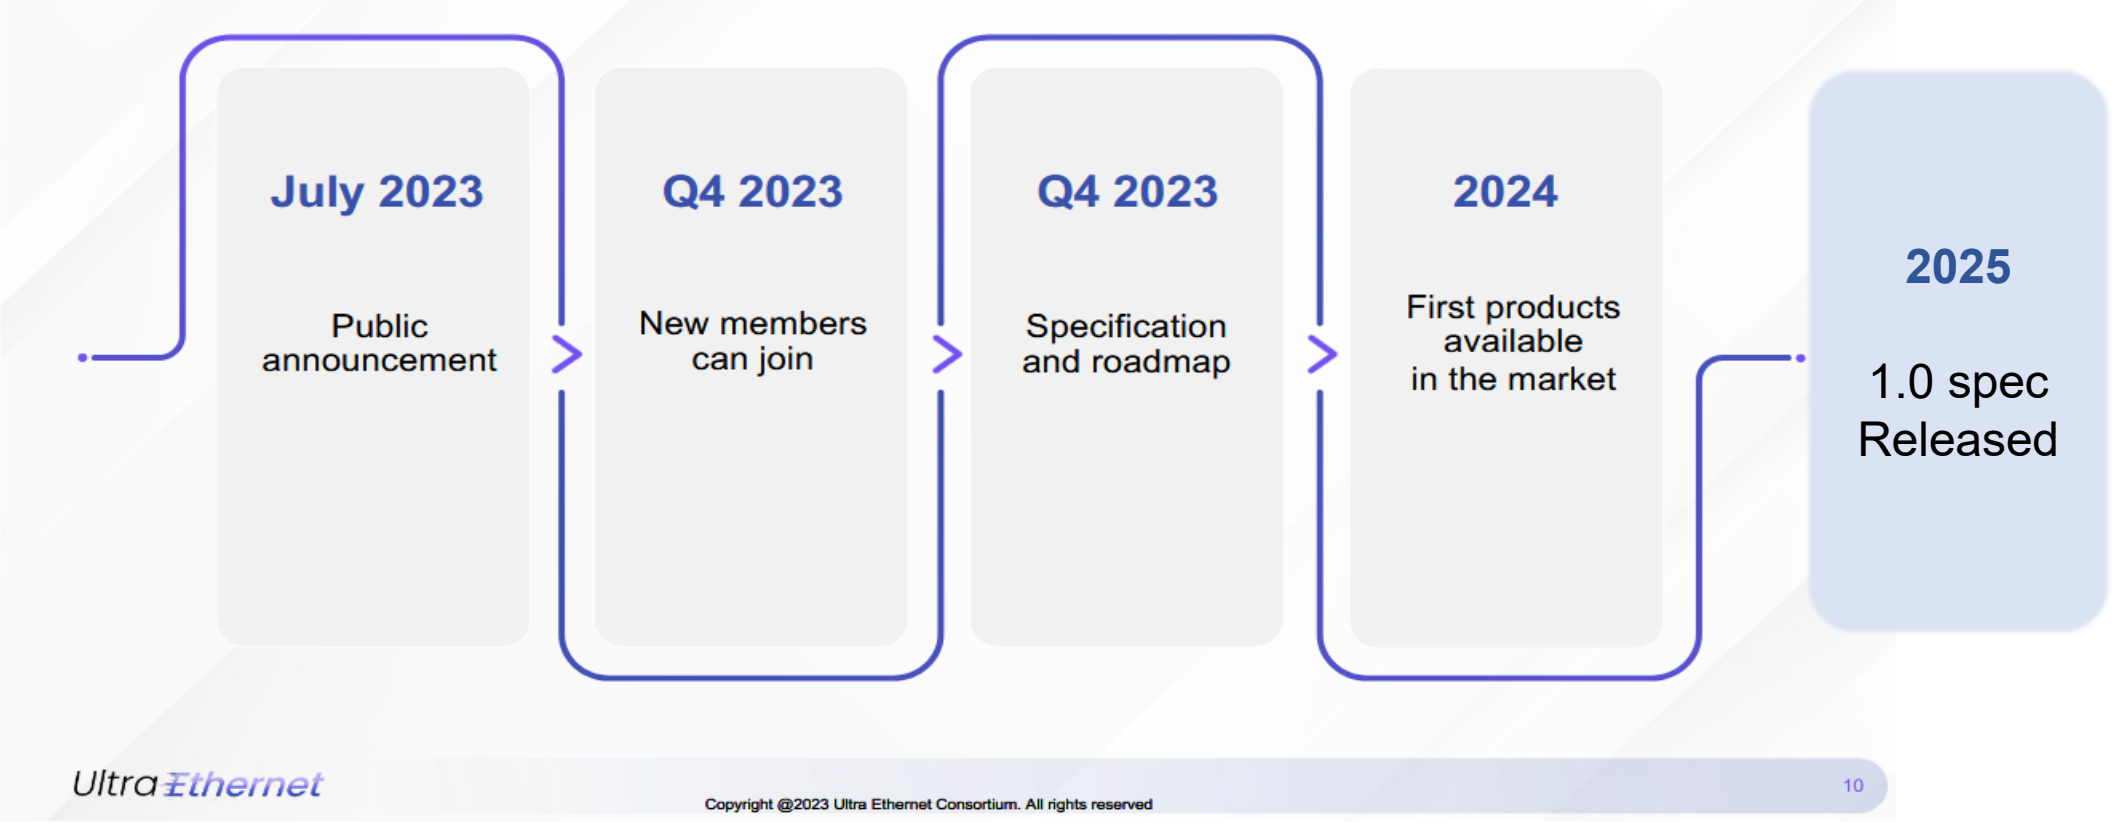

Industry Adoption Trends

Strong Industry Backing
– Founding members: AMD, Intel, Broadcom, HPE, Cisco

VIP Support
– Example: Siemens EDA Avery VIP announced UALink VIP

IP Availability
– Vendors offer controller + PHY IP

Switch Silicon Development
– Companies like Broadcom, Marvell working on UALink-compliant switches
– Potential for custom PHY + MAC

Software Ecosystem
– Early tools and drivers available
– Reference implementations on GitHub

Future Outlook

Both UE and UALink are going to play crucial role in AI/ML infrastructure

Expected significant adoption of UEC by data centers by 2027

Ultra Ethernet support for Scaleup network (ULN) is also expected in future

The Future is Ultra

- Ultra Link and Ultra Ethernet for Ultra scale of Data

Conclusion

- AI workloads are driving the need for scalable, low-latency connectivity
- UALink and UEC emerge as purpose-built solutions for AI and HPC fabrics
- Strong industry backing and ecosystem activity accelerating adoption
- Roadmaps show clear alignment with future performance targets (800G, 1.6T+)
- Open, interoperable standards ensure long-term scalability and vendor flexibility

References

- UALink Consortium
- UEC consortium
- Tom’s Hardware: UALink vs NVLink context & incorporation
- Cisco Nexus AI Networking
- AMD Arch
- NvLink