Empowering AI Workloads in UEC

Unleashing Potential from the Physical Layer

OIF 448Gbps Signaling for AI Workshop, April 15-16, 2025

Presented By: Mark Nowell, Fellow, Cisco

“Ultra Ethernet?” - Agenda

- Who is Ultra Ethernet Consortium (UEC)?
- AI Workloads and Challenges
- Physical Layer Considerations

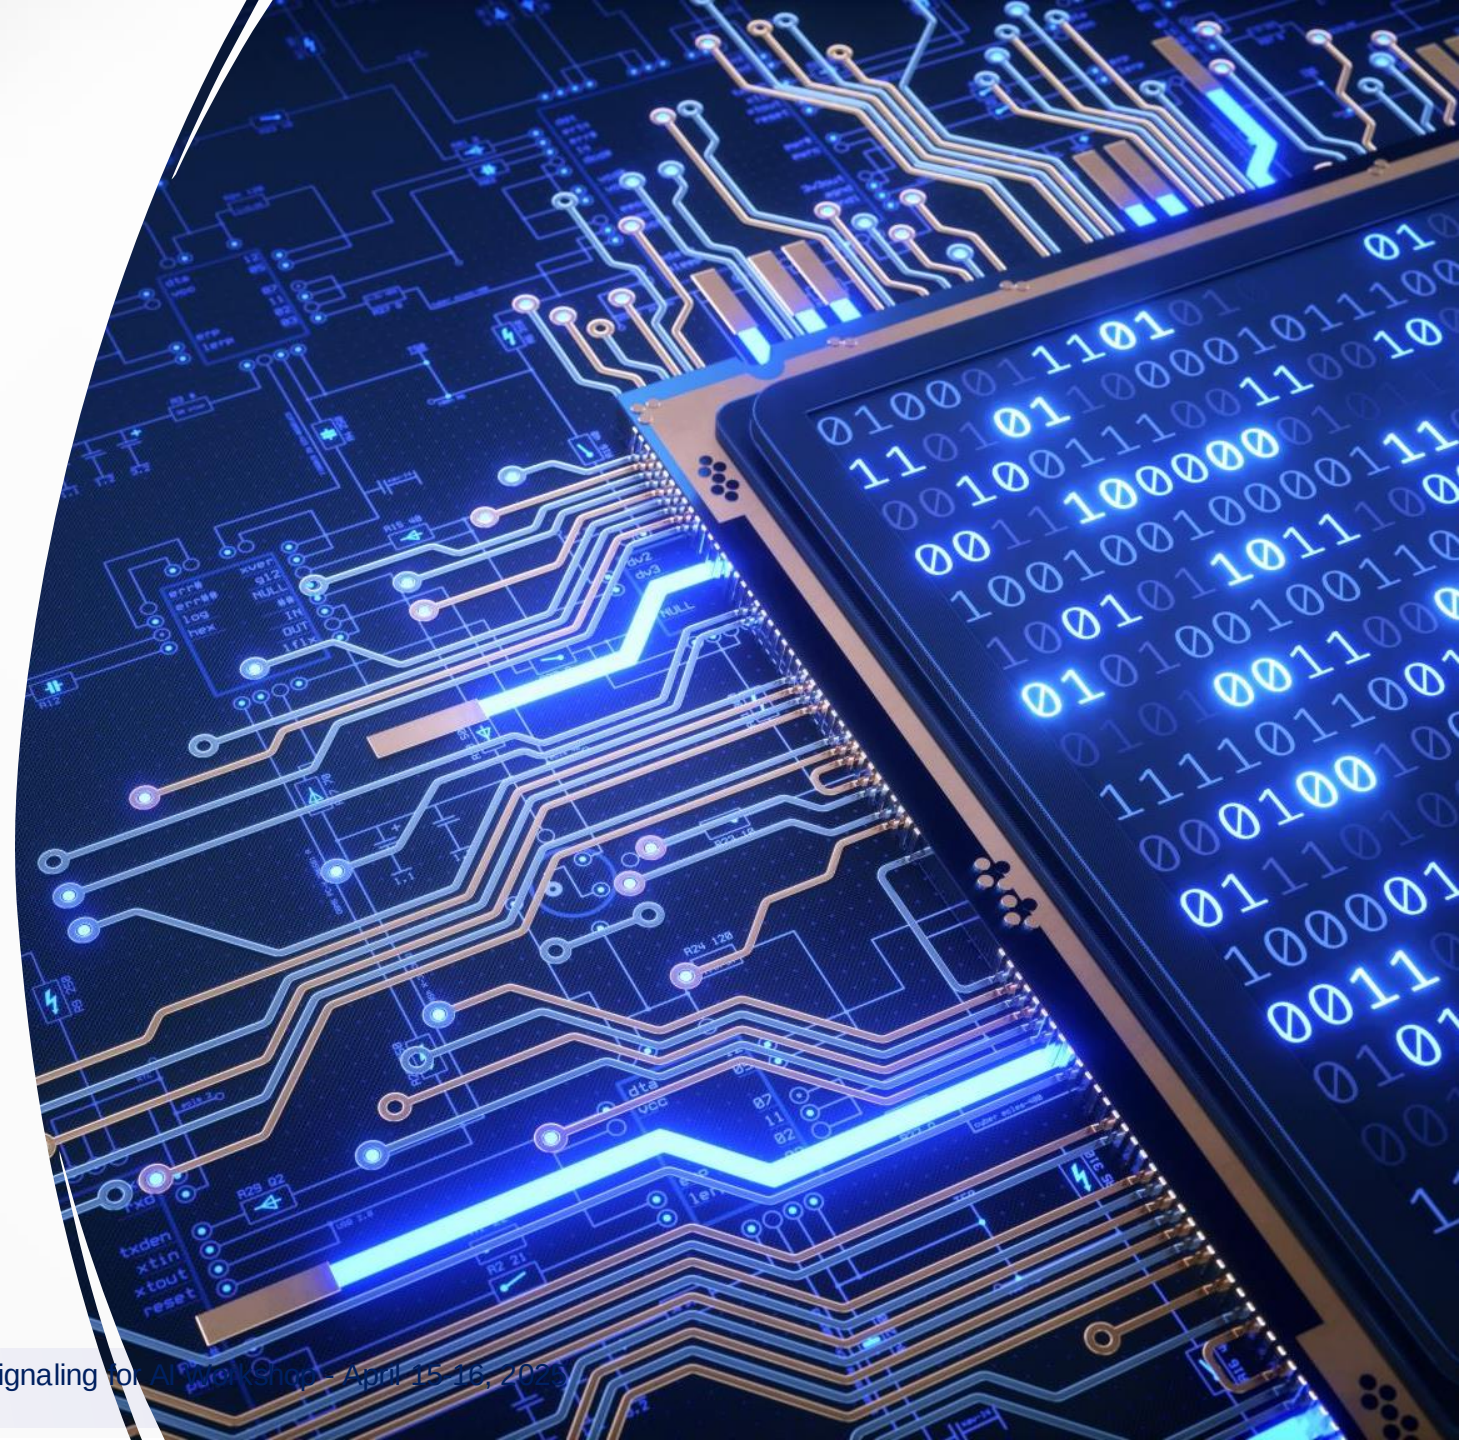

Introducing: The Promise of Ultra Ethernet

https://*ultra*ethernet.org/

THE NEW ERA NEEDS A NEW NETWORK

Ultra Ethernet

As **performant** as a supercomputing interconnect

As **ubiquitous** and **cost-effective** as Ethernet

As **scalable** as a cloud data center

What is UEC and Its Holistic Approach to AI & HPC

- Full Standards Development Organization
- (One of the?) Fastest growing projects in Linux Foundation
- 120+ Companies
- 1500+ individual active contributor volunteers

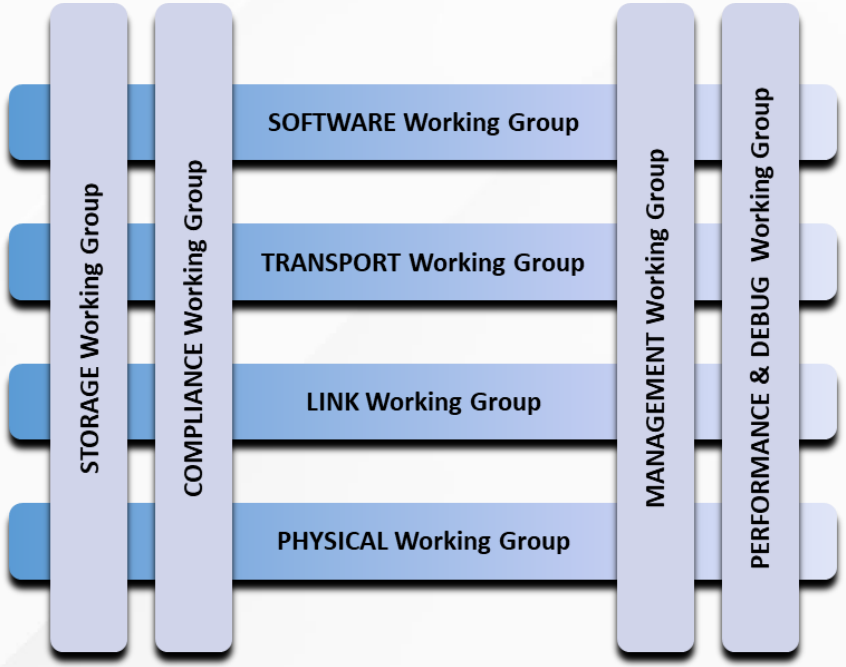

UEC Stack

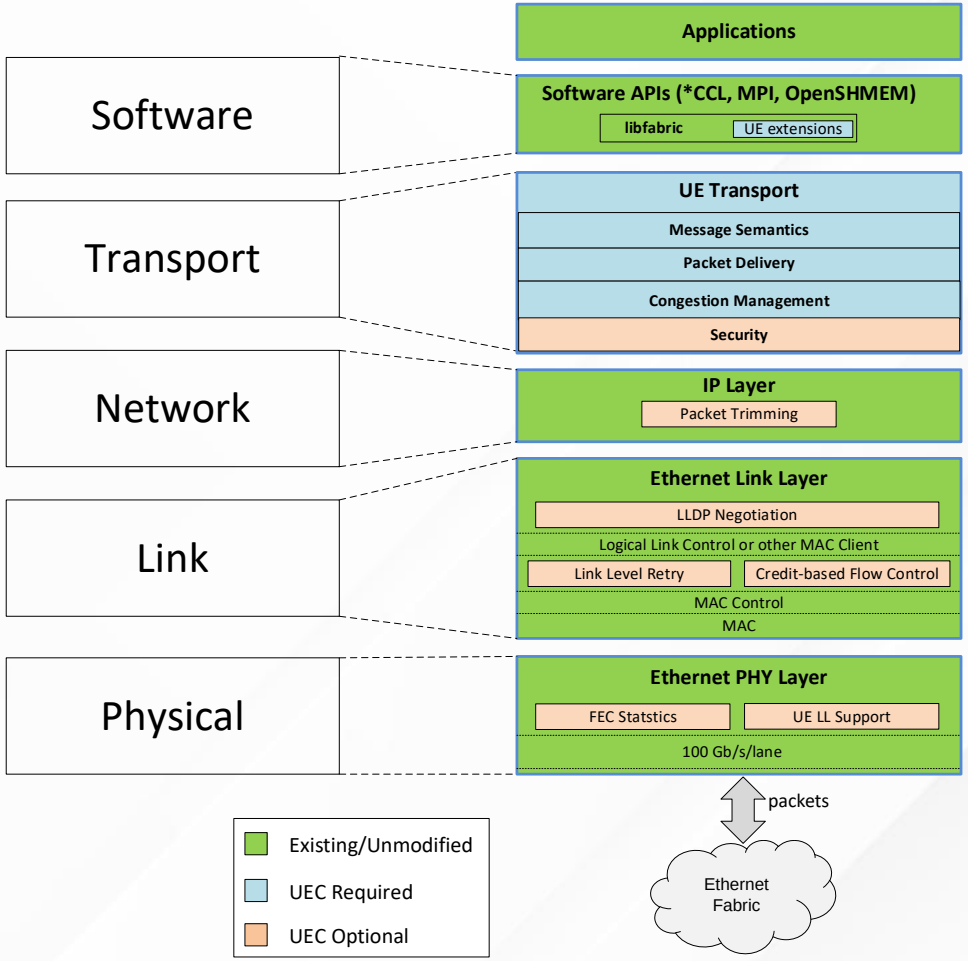

Which Network?

General Purpose vs. Scale-Up versus Scale-Out (UEC) Networks

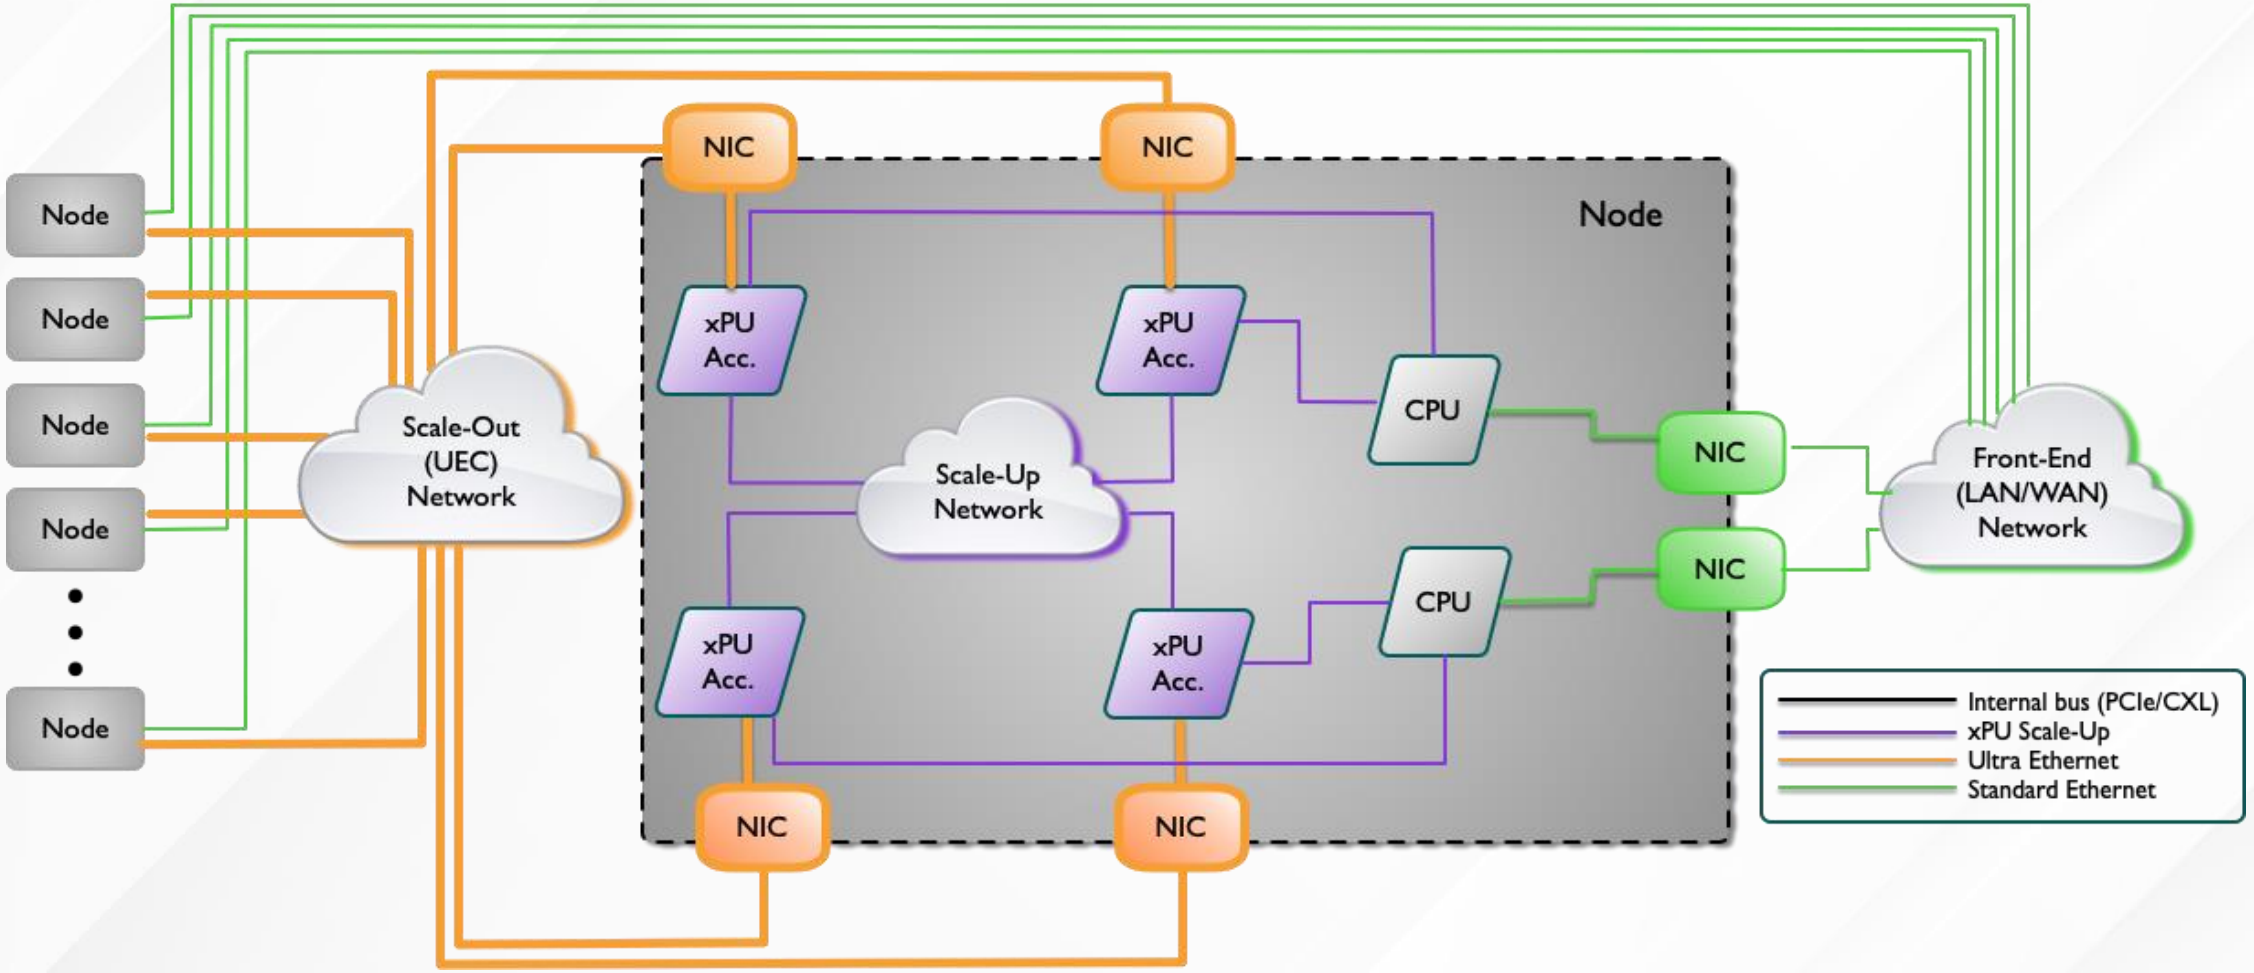

Why Ethernet?

- Aggressive bandwidth roadmap
- Latency improvements
- Topology flexibility
- Enhanced reliability and QoS
- Scalability
- Energy Efficiency
- Advanced Security Features
- AI-driven network management
- Cost and mature ecosystem support
- Broad interoperability
- Day 0, Day 1, and Day 2 management
- **Diverse, open ecosystem**

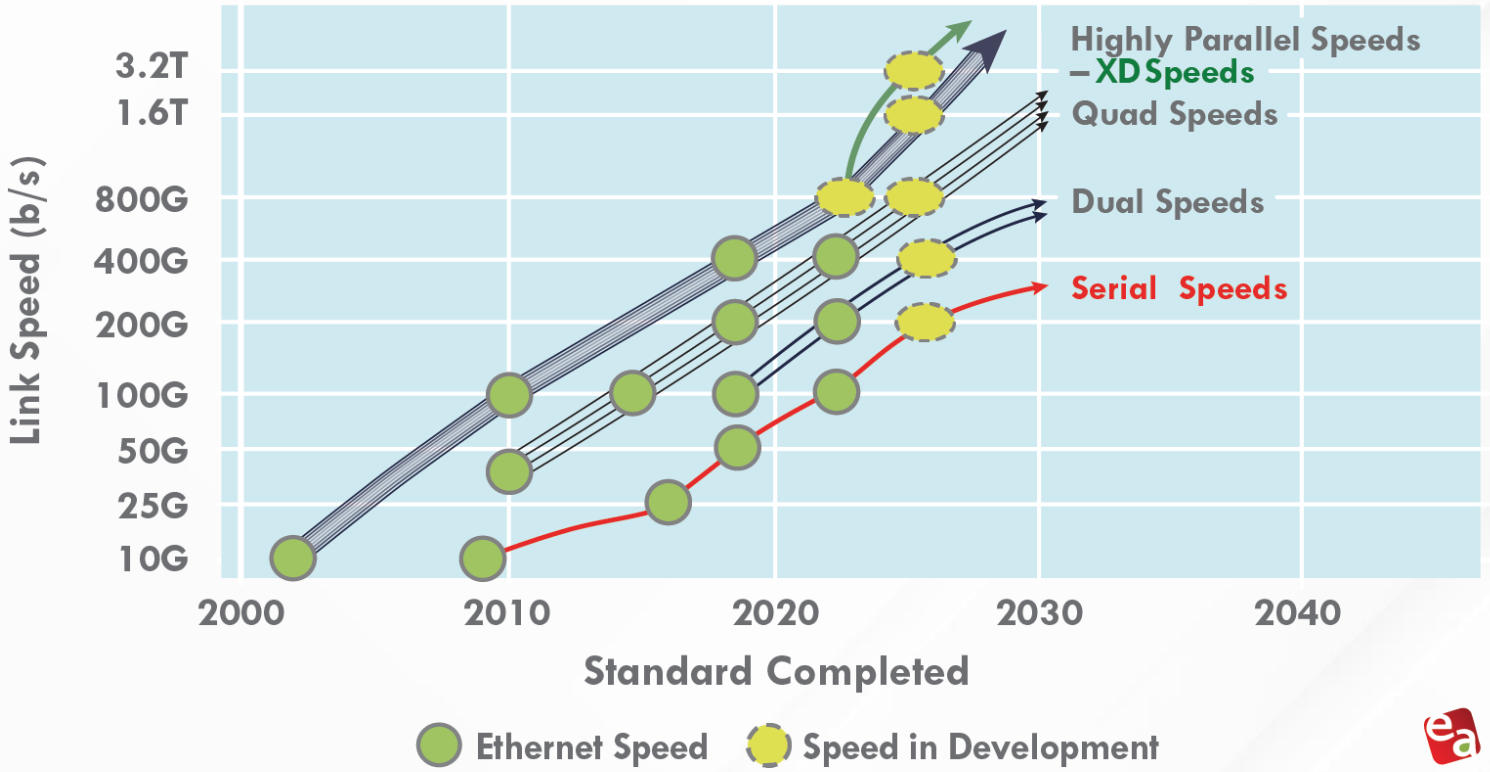

Source: Ethernet Alliance (https://ethernetalliance.org/technology/ethernet-roadmap/)

AI for Networking, or Networking for AI?

- Many articles/blogs have talked about how AI can change the networking infrastructure
  - … but what network infrastructure do you need to have enough AI to change the networking infrastructure?
  - Is it more than just superfast speeds and feeds?
  - Massive data sets, parallel processing requirements
  - Where does the data need to be, and when?

The AI Monster

- AI workloads need
  - Ever-increasing Memory Bandwidth
  - Ever-increasing Memory Capacity
  - (Near) Instantaneous Data Access (Exabytes)
- Intermittent data surges
- ”Straggler” data (tail latency) significantly impacts completion time
- Extended operation duration (hours, days)

Scale Out Considerations

- Amplification
  - Topology Amplification
  - Packet error, block error, or signaling error perpetuates through the upper layers of all nodes
  - Power Amplification
  - Error Amplification
    - ”Without fault recovery through memory reconstruction, more time would be spent restarting the training run in 100,000 GPU clusters due to optics failures, than it would be advancing the model forward.”
- Reach and Distance
  - Current spec assumes 150m for intra-Data Center and 1.5km for inter-Data Center
  - Longer distances increase latency times due to retry efforts
  - Tradeoff: Distance adds latency, FEC adds latency but FEC improves link performance (reducing retries)
  - Estimation of where optics is likely to have adoption:
    - Scale-out (node-to-node) → definitely
    - Scale-up (within a rack) → possible. Needs to add value compared to copper (higher speeds? Better error rates and consistency?

*Patel, Dylan, and Nishball, Daniel. (2024). “100,000 H100 Clusters: Power, Network Topology, Ethernet vs InfiniBand, Reliability, Failures, Checkpointing.” https://www.semianalysis.com/p/100000-h100-clusters-power-network

Summary/Conclusion

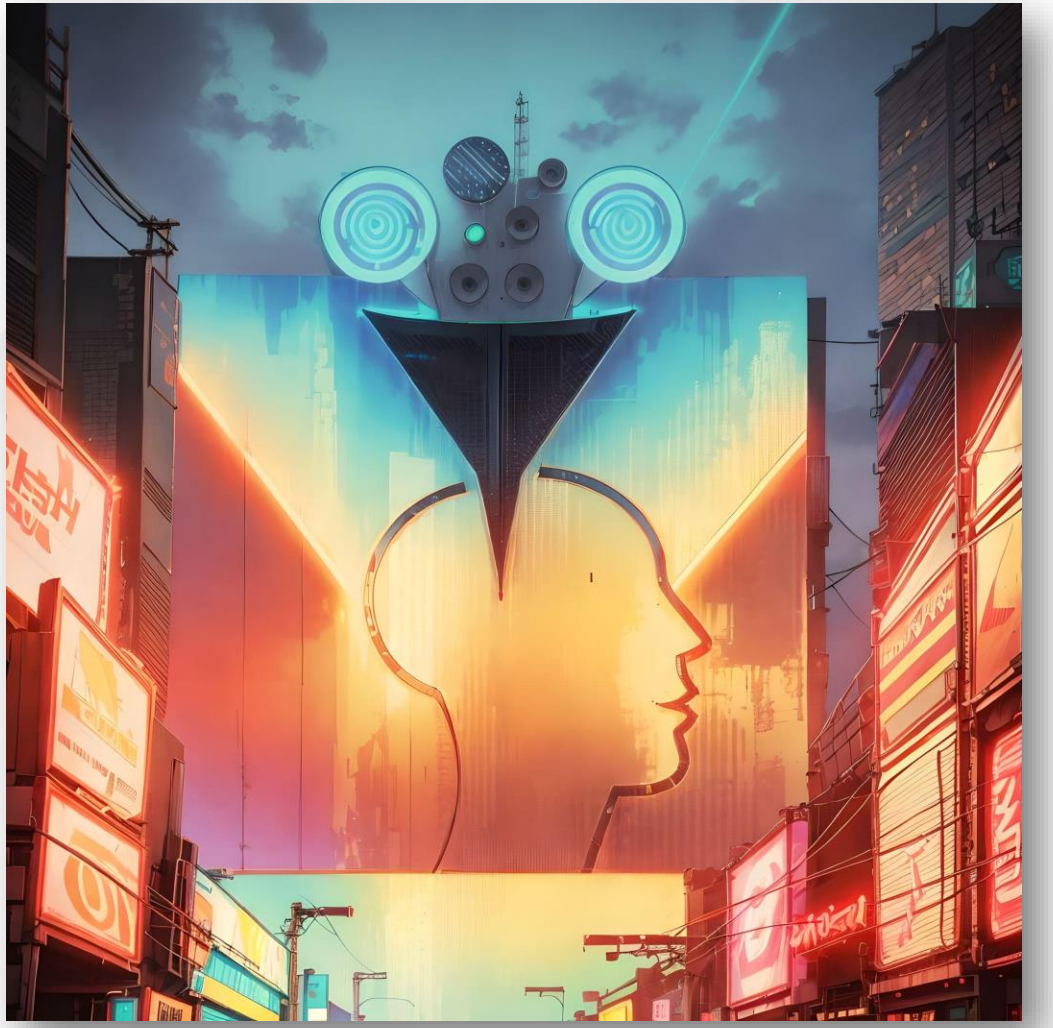

- If UEC is to succeed, there needs to be a paradigm shift
- Parallelism is the answer, but brings with it more questions
  - Different communication needs have complex interactions and dependencies
- Increase scale reduces margin/tolerance for error
  - Adjusting topologies and network architectures *increase* inter-reliance on physical interconnects
  - Better physical layer reliability reduces significant pain points at the software layer
- The importance of collaboration and innovation cannot be understated

Ultra Ethernet Consortium

LEARN MORE AT

www.ultraethernet.org


Ultra Ethernet
Consortium